

Optimal Security for Keyed Hash Functions: Avoiding Time-Space Tradeoffs for Finding Collisions

Cody Freitag¹[0000–0002–6307–204X], Ashrujit Ghoshal²[0000–0003–2436–0230],
and Ilan Komargodski³[0000–0002–1647–2112]

¹ Cornell Tech

`cfreitag@cs.cornell.edu`

² Paul G. Allen School of Computer Science & Engineering
University of Washington, Seattle, Washington, USA

`ashrujit@cs.washington.edu`

³ School of Computer Science and Engineering, Hebrew University of Jerusalem
and NTT Research 91904 Jerusalem, Israel

`ilank@cs.huji.ac.il`

Abstract. Cryptographic hash functions map data of arbitrary size to a fixed size digest, and are one of the most commonly used cryptographic objects. As it is infeasible to design an individual hash function for every input size, variable-input length hash functions are built by designing and bootstrapping a single fixed-input length function that looks sufficiently random. To prevent trivial preprocessing attacks, applications often require not just a single hash function but rather a family of keyed hash functions.

The most well-known methods for designing variable-input length hash function families from a fixed idealized function are the Merkle-Damgård and Sponge designs. The former underlies the SHA-1 and SHA-2 constructions and the latter underlies SHA-3. Unfortunately, recent works (Coretti et al. EUROCRYPT 2018, Coretti et al. CRYPTO 2018) show non-trivial time-space tradeoff attacks for finding collisions for both. Thus, this forces a parameter blowup (i.e., efficiency loss) for reaching a certain desired level of security. We ask whether it is possible to build families of keyed hash functions which are *provably* resistant to any non-trivial time-space tradeoff attacks for finding collisions, without incurring significant efficiency costs.

We present several new constructions of keyed hash functions that are provably resistant to any non-trivial time-space tradeoff attacks for finding collisions. Our constructions provide various tradeoffs between their efficiency and the range of parameters where they achieve optimal security for collision resistance. Our main technical contribution is proving optimal security bounds for converting a hash function with a fixed-sized input to a keyed hash function with (potentially larger) fixed-size input. We then use this keyed function as the underlying primitive inside the standard Merkle-Damgård and Merkle tree constructions. We strongly believe that this paradigm of using a keyed inner hash function in these constructions is the right one, for which non-uniform security has not been analyzed prior to this work.

1 Introduction

A cryptographic hash function is a (deterministic) algorithm that takes arbitrary length input data and outputs a fixed length digest. It is one of the most fundamental tools in modern applications of cryptography, underlying numerous widely used applications. For example, it facilitates the hash-and-sign paradigm, proofs-of-work for blockchains, and more. While it is empirically believed that concrete cryptographic hash functions satisfy various useful security properties, formalizing this seems to be currently out of reach. Thus, in the context of provable security, cryptographic hash functions are usually modeled as random oracles, i.e., completely random functions [6]. This allows us to analyze specific properties and argue about the concrete security of systems that use them. In this work, we focus on the property of a hash function being *collision resistant*, i.e., the idea that although collisions exist in abundance in a compressing function, it should be computationally hard to find them.

The task of finding collisions in a given compressing function is only interesting if the adversary is uniform. That is, the adversary is “fixed” before the hash function. Indeed, otherwise, a non-uniform attacker can simply have collisions hardwired. However, the uniform model of security does not capture many real-world adversaries, and therefore it is common to model adversaries as non-uniform in theoretical cryptography. Specifically, non-uniform security captures adversaries that have been designed to attack specific instances, adversaries that have gone through an expensive preprocessing stage, or even protect against (currently unknown) future attacks. Non-uniform security is also necessary for composition within larger systems [18]. For all of these reasons, it is widely believed by the theoretical community that modeling attackers as non-uniform is the right thing to do, despite potentially being overly conservative and including unrealistic attackers.

Dealing with non-uniform attackers in the context of hashing and collision finding makes it necessary to consider a family of keyed hash functions, rather than a single hash function. Collision finding is then defined via the following two-stage game. First, a (keyed) family H of hash functions is fixed, and the attacker can depend arbitrarily on H . Second, a random key key is sampled, and the adversary needs to find a collision in H relative to key . Intuitively, in order to attack the hash function (e.g., find a collision), a non-uniform attacker must either (a) have some hard-coded information about key , or (b) can essentially be treated as uniform.

For applications, we typically want each member of H to operate on unbounded input lengths. That is, $H: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ should be viewed as a two-input function, operating on (key, m) , where $\text{key} \in \{0, 1\}^\kappa$ is the key and $m \in \{0, 1\}^*$ is an arbitrary length input. Since it is practically infeasible to design a different hash function for every input length, what happens is that a single basic compressing function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ for some $a > n$ is designed, and then it is iterated in some way to get a hash function that compresses arbitrarily. For instance, the well-known Merkle-Damgård design [24, 13] iterates such a basic compressing function in order to get a variable-input-length hash

function that can operate on arbitrary sized data up to some maximum length (e.g., 2^{64} bits).

AI-ROM. Since we consider non-uniform security in the random oracle model, we model attackers using the *auxiliary-input random oracle model* (AI-ROM), formally defined by Unruh [27] although implicitly used earlier, for example, by Hellman [20], Yao [28], and Fiat and Naor [15]. In this model, we assume a hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ with $a > n$ modeled as a completely random one, i.e., a random oracle [6]. The AI-ROM models preprocessing adversaries as two-stage algorithms $(\mathcal{A}_1, \mathcal{A}_2)$ parameterized by S (for “space”) and T (for “time”). We refer to such an attacker as an (S, T) -attacker. The first part \mathcal{A}_1 (i.e., the offline phase) has unbounded access to h , and its goal is to compute an S -bit “advice” σ for \mathcal{A}_2 . The second part \mathcal{A}_2 (i.e., the online phase) gets the advice σ , can make at most T queries to h , and attempts to accomplish some task involving h . In our case, \mathcal{A}_2 gets a random key $\text{key} \leftarrow \{0, 1\}^\kappa$ as a challenge and its goal is to come up with a collision in $H(\text{key}, \cdot)$. Aside from the restrictions that $|\sigma| \leq S$ and that \mathcal{A}_2 can make at most T queries to h , both \mathcal{A}_1 and \mathcal{A}_2 are allowed to be computationally unbounded.

Building a keyed hash from a single hash function. Observe that for every keyed hash construction, there is an (S, T) -attacker that finds a collision relative to a random key with probability⁴ $\Omega(S/2^\kappa + T^2/2^n)$ via the following attack. First, the preprocessing adversary outputs $\Omega(S)$ collisions with respect to arbitrary distinct keys. The online adversary receives a random key. If key is in the remembered list from the preprocessing phase, it outputs the corresponding collision. Otherwise, it performs a T -query birthday-style attack. The adversary wins if either the challenge key appears in one of its preprocessed collisions (giving the $S/2^\kappa$ term) or if the birthday attack succeeds (giving the $T^2/2^n$ term). We refer to this attack as the *naive attack*, and say that a construction is *optimally secure* if there is provably no better attack. This brings us to the main question we consider in this work.

Can we build a keyed hash function (i.e., $H: \{0, 1\}^\kappa \times \{0, 1\}^ \rightarrow \{0, 1\}^n$) from non-keyed one (i.e., $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$) with optimal non-uniform security?*

If we could design an $h_a: \{0, 1\}^a \rightarrow \{0, 1\}^n$ for every $a \in \mathbb{N}$, then the above task is easy. We can simply parse the input to the appropriate h_a into two parts, one for the key and the other for the input to $H: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$. That is, define $H(\text{key}, m) = h_{\kappa+|m|}(\text{key} \| m)$, where $\|$ stands for string concatenation and $|\cdot|$ stands for bit length. For this construction, Dodis, Guo and Katz [14] showed that the best attack achieves advantage $O(S/2^\kappa + T^2/2^n)$, matching the advantage of the naive attack.

Unfortunately, it is infeasible to design a different hash function for every input length as discussed above. The design of a new h is a delicate and lengthy

⁴ To simplify notation throughout the introduction, we suppress **poly** factors in n in the asymptotic $O(\cdot)$ and $\Omega(\cdot)$ notation.

process that could take many years to test and standardize. Having a single hash function is therefore more robust security-wise. Thus, the standard procedure is to design a hash function h with fixed input size and then iterate it in some way to get a hash function that supports arbitrary input lengths.

It may seem that standard domain extension techniques for hash functions (like Merkle-Damgård, Sponge, or Merkle trees) provide a solution for this problem. Indeed, their goal is to take a hash function on a small domain and turn it into a hash function with arbitrary-size domain. But, as we point out next, the standard constructions suffer from a significant security loss. A priori, it is not even clear that this security loss is avoidable.

The security of existing constructions. First, consider (a keyed variant of) the Merkle-Damgård (MD) construction [24,13], perhaps the most widely popular design for getting a hash function on long inputs from one on fixed input sizes. This design is not only extremely fundamental in cryptographic theory, but it also underlies popular hash functions used in practice, most notably MD5, SHA-1, and SHA-2. The MD: $\{0,1\}^\kappa \times \{0,1\}^* \rightarrow \{0,1\}^n$ construction iterates the basic hash function $h: \{0,1\}^a \rightarrow \{0,1\}^n$ by feeding in input blocks of size $s = a - \max\{\kappa, n\}$ one by one. It first pads the message appropriately such that it is a multiple of s bits. For key $\text{key} \in \{0,1\}^\kappa$ and input $m \in \{0,1\}^s$, define $\text{MD}(\text{key}, m) = h(\text{key} \| m)$. Then, for a longer input $m \in (\{0,1\}^s)^\ell$, viewed as ℓ blocks m_1, \dots, m_ℓ each from $\{0,1\}^s$, recursively define $\text{MD}(\text{key}, (m_1, \dots, m_\ell)) = h(\text{MD}(\text{key}, (m_1, \dots, m_{\ell-1})), m_\ell)$. We note that in the standard MD construction (studied, for example, in [11,1,17,2]), key is only explicitly included once when processing the first message block.

Collision resistance of MD in the AI-ROM was first studied by Coretti, Dodis, Guo, and Steinberger [11] and more recently by [1,17,2]. It is known that there is an attack, loosely based on the idea of rainbow tables [20,25], which succeeds in finding a collision with probability $\Omega(S/2^\kappa + ST^2/2^n)$. In typical settings of parameters, the $ST^2/2^n$ term dominates the above expression and in this case it is evident that MD suffers from a significant security loss.

Concretely, in the SHA-1 construction, $a = 678$ and $\kappa = n = 160$. If we model the underlying primitive $h: \{0,1\}^{678} \rightarrow \{0,1\}^{160}$ as a perfectly random function, an (S, T) -attacker with $S = 2^{53}$ and $T = 2^{50}$ will find a collision with probability $\approx 2^{-7}$ (essentially completely breaking the scheme).⁵ On the other hand, the best one could hope is a construction with maximal advantage $O(S/2^\kappa + T^2/2^n) \approx 2^{-60}$ (obtained by the naive attack).

Another construction we mention is the Sponge [8,7] construction, an alternative to the Merkle-Damgård design that underlies the modern SHA-3 hashing standard. As opposed to MD, the Sponge construction relies on a random *permutation* $\Pi: \{0,1\}^n \rightarrow \{0,1\}^n$. Sponge iterates Π by feeding in blocks of size $r < n$ from the input one at a time in a certain way. It results with a keyed hash function $\text{Sp}: \{0,1\}^\kappa \times \{0,1\}^* \rightarrow \{0,1\}^r$ with $\kappa + r = n$. Coretti et al. [11] (see

⁵ These parameters roughly correspond to an attacker with ≈ 1000 terabytes of memory that uses optimized hardware that can compute 3 billion hashes per second for a long weekend.

also [16]) showed that there is a collision finding (S, T) -attack with advantage $\Omega(ST^2/2^\kappa + T^2/2^r)$ against **Sp** relative to a random key. Again, we see that there is a non-trivial security loss in this construction.

It is important to note that for every choice of S and T the above attacks on Merkle-Damgård and Sponge beat the naive attack. In particular, there is no non-trivial choice of parameters where **MD** or **Sp** achieve the optimal security bound.

Lastly, we mention two other popular (variable-input-length) hash function designs: Merkle trees [23] and the BLAKE family [3,4]. The former (Merkle trees) is a popular design that has important features like local opening and can be easily parallelized. Although it is extremely popular both in theory and in practice, we are not aware of a keyed variant that has been studied in the non-uniform setting. The latter (BLAKE) is a runner-up in NIST’s competition to create a new hashing standard (where Sponge ended up as the winner). This design is based on the MD design, but they allow the inner hash function h to be keyed at every invocation. We are not aware of a formal study of its security in the non-uniform setting. Looking ahead, two of our main contributions are a proposal and analysis of the non-uniform security of Merkle tree and the MD/BLAKE design, where the inner hash function h is keyed in every invocation. Concretely, we believe that this is the right notion to consider moving forward, in terms of non-uniform security.

A different perspective. Above, we considered the scenario where $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ is given, and we want to build an $H: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ which is as secure as possible for every (S, T) -attacker. A different perspective, slightly more target oriented, is to first fix a desired security level (say 2^{-50}) and the power of adversaries (say $S = T = 2^{60}$) and then understand which h is needed in order to get the desired H . If we use **MD** (for concreteness) for H , we will need $n \geq 230$, but if we had an optimally secure construction of H , we would need only $n \geq 170$. The latter could potentially be easier to design and argue about.

1.1 Our Results

We provide several constructions of keyed hash functions from non-keyed ones that do not suffer from any security loss (i.e., the naive attack that has advantage $\Theta(S/2^\kappa + T^2/2^n)$ is provably optimal). Our constructions provide various tradeoffs between their efficiency and the range of parameters (S and T) where they achieve optimal security.

Merkle-Damgård and Merkle trees with a keyed inner hash. All of our constructions can be viewed within a framework that builds on the Merkle-Damgård and Merkle tree constructions.

We start by discussing the MD-based approach. We consider an iterative hashing design where a compression phase is performed in every step using an “inner hash” function. The input for the compression phase is the current state and the next input. At the end, the compression phase outputs the next state.

Of course, the inner hash function in the compression phase can use h as a subroutine. Abstractly, the compression phase for the MD-based construction is

$$y := \text{compress}(\text{key}, y, m),$$

where importantly the **compress** function takes **key** as input. See Figure 1 for an illustration. With this notation, the compression function of the standard MD function (at least as studied in numerous recent works including [11, 1, 17, 2]) is simply $\text{compress}(\text{key}, y, m) = h(y, m)$, and for the first step, y is initialized to **key**. Notably **key** is not included in every compression phases.




Fig. 1: Our framework for building keyed hash functions based on the Merkle-Damgård construction with a keyed inner compression function.

We next consider a parallelizable hashing design that generalizes the Merkle tree hash function. Here, each input is fed into a “leaf” of the Merkle tree, along with the **key** value. The compression function is then used to recursively combine outputs in previous levels until a final output is generated. Crucially, we always include **key** in the compression function. See Figure 2 for an illustration. This framework provides an alternative to the generalized MD approach described above. It requires at most a factor of two more calls to **compress**, but it is extremely parallelizable. Further, it provides a local opening property, where someone can prove that an individual message block m_i was included in the hash, without providing the full message.

Our constructions are obtained by different implementations of **compress**, namely viewing **compress** as an inner keyed hash function used in the MD and Merkle tree designs. Quantitatively, the MD and Merkle tree approaches give similar results, so we focus our attention on instantiating **compress** in the case of the MD-based framework. However, all of our main results extend to the setting of the generalized Merkle tree framework, which we provide in the full version.

For simplicity of presentation of our results, we slightly simplify notation and assume that κ (the key length) is equal to n (the output size of the hash function).⁶ In our formal theorem statements in the technical sections, κ and n are treated independently when relevant.

⁶ We note that there are constructions that use $\kappa \neq n$ by design (e.g., BLAKE hash [3, 4] uses $\kappa = n/2$).




Fig. 2: Our framework for building keyed hash functions based on the Merkle tree construction.

Efficiency: We measure efficiency of a given construction by the number of calls to h needed to evaluate H at a single point. For example, in the standard MD construction with an underlying hash that maps $\{0, 1\}^a$ to $\{0, 1\}^n$, to hash a b -bit input, the query complexity is $b/(a - n)$ (ignoring rounding⁷). Indeed, every application of h takes as input the previous output (n bits) and so it can process $a - n$ bits from the input each time.

Assuming a large inner hash. Our first result shows that optimal security loss is achievable. That is, we show that there is a way to take a random oracle that operates on a fixed input length and get a keyed hash H that operates on arbitrary-length inputs with the following security guarantee: for any S, T , any (S, T) -attacker has minimal possible advantage in finding a collision in H relative to a random key. In words, the new construction is a variant of MD where we also feed key as input in every block. We refer to this construction as the MD construction with a keyed inner hash, in contrast to the standard MD construction where key is only fed in the first block. At a high level, feeding the key into every invocation of h allows us to reduce the probability of finding a long collision in H to that of finding a collision in h , which achieves optimal security $O(S/2^n + T^2/2^n)$ [14]. Refer to Figure 3 for an illustration of how the construction works.⁸

⁷ To be more precise, MD requires $\lceil (b + n + 1)/(a - n) \rceil$ calls to h after padding the input with its length followed by a 1 and a sequence of 0s to fill the remaining current block. However, for ease of presentation, we ignore rounding in the introduction. In the formal theorem statements, we give exact efficiency bounds.

⁸ Essentially the same construction appears in Goldwasser-Bellare's lecture notes [19, §8.5] where it is shown that this construction is collision resistant in the uniform setting. Our result shows that this holds in the non-uniform (AI-ROM) setting as well.

Theorem 1 (Informal; see Theorem 8). Assume $h: \{0,1\}^a \rightarrow \{0,1\}^n$ is modeled as a random oracle with $a > 2n$. Then, there is an $H_1: \{0,1\}^n \times \{0,1\}^* \rightarrow \{0,1\}^n$ such that:

1. For any $S, T \in \mathbb{N}$ and any (S, T) -attacker, their advantage in finding a collision in $H_1(\text{key}, \cdot)$ relative to a random $\text{key} \leftarrow_s \{0,1\}^n$ is $O(S/2^n + T^2/2^n)$.
2. One evaluation of H_1 on a given key and a b -bit message requires $b/(a-2n)$ queries to h .




Fig. 3: The construction H_1 underlying Theorem 1 given a hash function $h: \{0,1\}^a \rightarrow \{0,1\}^n$ for $a > 2n$.

The above result is optimal in terms of security and is almost as efficient as standard MD if $a \geq 2n + \Omega(n)$. For example, if $a = 3n$, processing a b -bit input of H_1 requires querying h as many as $b/(a-2n) = b/n$ times. In the standard MD construction, only $b/(a-n) = b/(2n)$ queries are required, so our construction is less efficient than MD by a small constant factor at most 2 when $a = 3n$.

However, H_1 is significantly less efficient than MD if a is roughly $2n$, i.e. h compresses by a factor of 2. For example, if $a = 2n + 1$, then processing a b -bit input of H_1 requires invoking h as many as b times. However, MD requires only $b/(n+1)$ queries. This is a significant difference. We emphasize that having an efficient construction even when $a \approx 2n$ is not only a technicality but is rather important: concretely, assuming that the basic compressing function shrinks by a factor 2 is extremely common, both in theory and in practice. Thus, our next results are focused on closing this gap.

Instantiating the keyed inner hash with standard MD. To this end, we start by considering a construction H_2 that works for *any* $a > n$ and only incurs a factor of 2 overhead in terms of efficiency relative to MD. While this may seem too good to be true, we pay in terms of the assumptions we need to make to claim optimal security for collision resistance. Namely, the scheme has “optimal security,” meaning any (S, T) -attacker can find a collision with probability at most $O(S/2^n + T^2/2^n)$, only whenever $S \leq T$ and $ST^2 \leq 2^n$.

This main idea behind the construction H_2 is to instantiate the compress function in the MD-based framework of Figure 1 with a standard MD hash

function. We use `key` as the key for MD, and we treat $y_{i-1} \| m_i$ as the message. If we use a message block size $|m_i| = n$, this results in only a factor of two overhead relative to MD (essentially, half of the invocations of h incorporate bits of the message m_i , and half of the invocations incorporate bits of the previous output y_{i-1}). This construction is depicted in Figure 4 and gives the following result.

Theorem 2 (Informal; see Theorem 10). *Assume $h: \{0,1\}^a \rightarrow \{0,1\}^n$ is modeled as a random oracle with $a > n$. Then, there is an $H_2: \{0,1\}^\kappa \times \{0,1\}^* \rightarrow \{0,1\}^n$ such that:*

1. *For any $S, T \in \mathbb{N}$ such that $S \leq T$, $ST^2 \leq 2^n$, and any (S, T) -attacker, their advantage in finding a collision in $H_2(\text{key}, \cdot)$ relative to a random $\text{key} \leftarrow \{0,1\}^n$ is $O(S/2^n + T^2/2^n)$.*
2. *One evaluation of H_2 on a given key and a b -bit input requires $2 \cdot b/(a - n)$ queries to h .*




Fig. 4: The construction H_2 underlying Theorem 2 given a hash function $h: \{0,1\}^a \rightarrow \{0,1\}^n$ for $a > n$. The gray dotted boxes represent the `compress` function, instantiated with the Merkle-Damgård construction, that uses $\text{key} \in \{0,1\}^n$ as the key and $y_{i-1} \| m_i \in \{0,1\}^{2n}$ as the message.

We note that the assumption that $ST^2 \leq 2^n$ in the construction above comes from the fact that best currently known time-space tradeoffs for the collision resistance of standard MD (culminating in [17, 2] following the works of [11, 1, 9]) require this assumption to get optimal bounds when analyzing the ℓ -block MD construction when $\ell \in \omega(1)$. In the special case where we only use a 2-block variant of MD as the underlying `compress` function, [1] give tight bounds that do not require that $ST^2 \leq 2^n$ (and furthermore, [17] gave tight bounds for all constants ℓ). This motivates our next construction, H_3 , which uses the 2-block MD construction for `compress`, but requires that the input size to h satisfies $a > 3n/2$. See Figure 5 for an illustration, and the corresponding result is given in the following theorem.

Theorem 3 (Informal; see Corollary 1). *Assume $h: \{0,1\}^a \rightarrow \{0,1\}^n$ is modeled as a random oracle with $a > 3n/2$. Then, there is an $H_3: \{0,1\}^\kappa \times \{0,1\}^* \rightarrow \{0,1\}^n$ such that:*

1. For any $S, T \in \mathbb{N}$ such that $S \leq T$ and any (S, T) -attacker, their advantage in finding a collision in $H_3(\text{key}, \cdot)$ relative to a random $\text{key} \leftarrow_{\$} \{0, 1\}^n$ is $O(S/2^n + T^2/2^n)$.
2. One evaluation of H_3 on a given key and a b -bit input requires $2 \cdot b / (2a - 3n)$ queries to h .




Fig. 5: The construction H_3 underlying Theorem 3 given a hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ for $a > 3n/2$. The gray dotted boxes represent the **compress** function, instantiated with a two-block Merkle-Damgård construction, that uses $\text{key} \in \{0, 1\}^n$ as the key and $y_{i-1} || m_i \in \{0, 1\}^{2a-2n}$ as the message. The $2a - 2n$ bit message is split evenly into the first and second call to h , indicated in the figure by a diamond.

Instantiating the keyed inner hash with a 2-level Merkle tree. For our final construction, we seek to build a hash function H_4 that is both efficient and optimally secure whenever $a \approx 2n$ (h is only compressing by a factor of 2), without assuming that $S \leq T$. In particular, $S \gg T$ makes sense in many practical scenarios: the pre-processing attacker may have much more than time T to generate its advice string of size S , the online attacker may have easy random access to a structured advice string, or the online time T may be small for applications that enforce a timeout with fixed-time communication session (see [5] as an example of an attack on a TLS session that requires relatively heavy computation in an offline phase). Lastly, we mention that the bounds we obtain on H_2 and H_3 are tight—there is a non-trivial attack whenever $S > T$ that scales with advantage $\Omega(ST\ell/2^n)$ for standard ℓ -block MD [11, 1].

For H_4 , we instantiate the **compress** function from the framework of Figure 1 using a keyed variant of a Merkle tree construction (this has never been formally defined or analyzed to the best of our knowledge). In our variant, we feed key into all leaves of the Merkle tree. Concretely, in our construction, we use a 2-level Merkle tree and feed key to both of them, corresponding to two distinct invocations of h , and we split the “message” $y_{i-1} || m_i$ into the remaining input bits for the leaves. The second level of the Merkle tree combines the two outputs from the first level, to produce a n bit output for **compress**. The full construction of H_4 is illustrated in Figure 6. We conjecture that this Merkle tree-based

construction is optimally secure for collision resistance (see Remark 1 for more details), but analyzing its security turns out to be highly non-trivial. In particular, our current analysis is only optimally secure when $ST^2 \leq 2^n$, as stated in the following theorem.

Theorem 4. Assume $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ is modeled as a random oracle with $a \geq 2n$. Then, there is an $H_4: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ such that:

1. For any $S, T \in \mathbb{N}$ such that $ST^2 \leq 2^n$ and any (S, T) -attacker, their advantage in finding a collision in $H_4(\text{key}, \cdot)$ relative to a random $\text{key} \leftarrow \{0, 1\}^\kappa$ is $O(S/2^n + T^2/2^n)$.
2. One evaluation of H_4 on a given key and a b -bit input requires $3 \cdot b / (2a - 3n)$ queries to h .




Fig. 6: The construction H_4 underlying Theorem 4 given a hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ for $a \geq 2n$. The gray dotted boxes represent the compress function, instantiated with a two-level Merkle tree, that uses $\text{key} \in \{0, 1\}^n$ in each leaf and $y_{i-1} || m_i \in \{0, 1\}^{2a-2n}$ as the message. The $2a - 2n$ bit message is split evenly between the two leaves of the Merkle tree, indicated by a diamond in the figure. We require $a \geq 2n$ so that both outputs from the leaves can be fed into the next layer of the Merkle tree.

We summarize our results in Table 1.

Remark 1 (A conjecture on the security of keyed Merkle trees). The main building block in our construction of Theorem 4 is a Merkle tree where the key value key is included only at the leaves. Concretely, we include a key key in each of the leaves of a Merkle tree and then fill in the rest of the leaves with bits of some message $m \in \{0, 1\}^*$, and then run the Merkle tree construction (with an unkeyed hash) as normal to get an n -bit output. In this work, we analyze the simplest case where the Merkle tree has depth 2 with only two leaves. However, this naturally generalizes to any number of ℓ leaves resulting in a tree of depth $O(\log \ell)$.

	Advantage	Efficiency	Input Size	Assumptions
MD	$\Theta(ST^2/2^n)$	$b/(a-n)$	$a > n$	None
H_1 (Thm. 1)	$\Theta(S/2^n + T^2/2^n)$	$b/(a-2n)$	$a > 2n$	None
H_2 (Thm. 2)	$\Theta(S/2^n + T^2/2^n)$	$2 \cdot b/(a-n)$	$a > n$	$S \leq T, ST^2 \leq 2^n$
H_3 (Thm. 3)	$\Theta(S/2^n + T^2/2^n)$	$2 \cdot b/(2a-3n)$	$a > 1.5n$	$S \leq T$
H_4 (Thm. 4)	$\Theta(S/2^n + T^2/2^n)$	$3 \cdot b/(2a-3n)$	$a \geq 2n$	$ST^2 \leq 2^n$

Table 1: A summary of our results as well as the standard MD construction for reference. The advantage of a construction is given in terms of the probability of (S, T) -attackers to find collisions relative to a random key. The efficiency is measured in terms of the number of calls to h that maps a -bit inputs to n when processing a b -bit input. The assumptions column specifies conditions on various parameters.

We conjecture that this approach, where only the leaves are keyed, is as secure as the Merkle tree approach of Figure 2 where the inner hash function is keyed at every invocation, including interior nodes. This latter approach requires a larger, more complicated, inner hash function, so we would like to avoid this if at all possible.

First, for the simple case of a depth two tree with keyed leaves, we conjecture that the bound we show in this work is not tight (see Theorem 12 for the exact bound we show). Namely, we believe that we should not need to assume $ST^2 \leq 2^n$ (or make any assumptions on S, T) in order to get optimal security. Second, we believe that this intuition should extend to the arbitrary depth Merkle trees that are keyed at the leaves, and we conjecture that it should also achieve optimal (S, T) security without any assumptions on S, T . However, even getting a bound in this case that is optimal in the setting $ST^2 \leq 2^n$ we believe would be very interesting. Additionally, handling deeper Merkle trees could potentially allow for constructions that do not necessarily have even a two-to-one structure, meaning that we could build a keyed hash function based on Merkle trees without assuming the hash function h has input size $a \geq 2n$.

1.2 Related Work

The motivation for this work comes from a recent line of results on the non-uniform security loss of various hashing mechanisms.

For Merkle-Damgård’s construction [24, 13], this was first studied by Coretti et al. [11] who showed how to find collisions with probability $\Omega(S/2^\kappa + ST^2/2^n)$. The idea is reminiscent of the rainbow tables attack due to Oechslin [25] (in turn building on Hellman [20]). The collisions they get are rather long (of length

proportional to T). Akshima et al. [1] generalized the attack to get an ℓ -block collision with probability $\Omega(S/2^\kappa + ST\ell/2^n)$ and showed that this attack is optimal for $\ell = 2$. Ghoshal and Komargodski [17] showed that this attack is optimal for all constant values of ℓ and Akshima, Guo, and Liu [2] almost proved the tightness of the bound for all ℓ s by showing that the best possible attack has advantage $O(ST\ell/2^n \cdot (1 + ST^2/2^n) + T^2/2^n)$. For a single-block Merkle-Damgård (i.e., just a compressing random oracle), Dodis, Guo, and Katz [14] showed that including a random key (optimally) defeats preprocessing attacks.

For Sponge [8,7], Coretti, Dodis, and Guo [10] stated a related attack with advantage $\Omega(ST^2/2^\kappa + T^2/2^r)$ (with r being a “rate” parameter of the scheme). Again, this attack resulted in very long collisions. The attack was formalized and extended to ℓ -block collisions with advantage $\Omega(ST\ell/2^\kappa + T^2/2^r)$ by Freitag, Ghoshal and Komargodski [16]. Freitag et al. also proved several upper bounds on the advantage of any attacker, but their bounds are not known to be tight.

Indifferentiability. Our work focuses on collision resistance, but there are other security properties of interest (such as inversion, second preimage resistance, pseudo-randomness, and unpredictability). In the uniform security setting there is a well-known framework called *indifferentiability* (due to Maurer, Renner, and Holenstein [22]) that is used to show that a (wide) class of security goals are simultaneously met. This allows to modularly transition to a (simpler) hybrid world where a complicated hash function construction is replaced with a monolithic random oracle (see, for example, [12]). Such transitions are known to work for all single-stage games but not for multi-stage games [26]. Our non-uniform security model is fundamentally a two-stage model and therefore the indifferentiability framework (as is) does not apply. It is an interesting open problem to find an analogue in the non-uniform setting.

2 Technical Overview

In this section, we give a high level overview of our main techniques. Recall, our goal is to construction a variable-input length, keyed, hash function $H^h: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ from an idealized, fixed-input length hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$.

Non-uniform security in the AI-ROM. We consider non-uniform (S, T) -attackers $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the auxiliary-input random oracle model (AI-ROM) [27] with the following structure. First, h is randomly sampled from the space of all a -bit to n -bit functions. Then, in the preprocessing phase, \mathcal{A}_1 has unbounded access to h and outputs an advice string σ such that $|\sigma| \leq S$. The online phase \mathcal{A}_2 receives auxiliary input σ and a random key $\text{key} \leftarrow \{0, 1\}^\kappa$ as input, and then has to find two distinct messages msg, msg' such that $H^h(\text{key}, \text{msg}) = H^h(\text{key}, \text{msg}')$ while making at most T queries to h . Our goal is to give constructions H^h such that no (S, T) -attacker as above can find a collision with better than $O(S/2^\kappa + T^2/2^n)$ probability. This is “optimal” in the sense that this matches a naive attack against a purely random H : the preprocessing attacker stores

collisions for $\Omega(S)$ keys, and the online attacker either gets “lucky” and receives one of those keys as input or performs a standard birthday-style attack.

Merkle-Damgård framework with a keyed inner hash. We consider a general framework based on the Merkle-Damgård (MD) transformation where we instantiate the inner hash function with a keyed one. Let $s \in \mathbb{N}$ be the desired message block size. Then, given a function $g: \{0,1\}^{\kappa+n+s} \rightarrow \{0,1\}^n$, we can build a function H^g where any attack on the collision resistance of H^g implies an attack on g . The idea behind H^g is as follows. We first break our message up into blocks m_1, \dots, m_ℓ of size $a - \kappa - n$. We initialize the value $y_0 = 0^n$, and for $i = 1, 2, \dots, \ell$ we compute $y_i = g(\text{key} \| y_{i-1} \| m_i)$. Finally, we output y_ℓ . It is known (e.g. see Section 8.5 of [19]) that if you can find a collision in the MD construction H^g for a keyed g in the uniform setting, then this implies you can find a collision in g . Indeed, this reasoning extends to the non-uniform setting with (S, T) -attackers in the AI-ROM. Hence, this shows that H^g is as secure as g . So, our new goal is to construct such a g with “optimal security” given an idealized hash function $h: \{0,1\}^a \rightarrow \{0,1\}^n$.

Our first observation is that if $a \geq \kappa + n + s$, then we can simply use $g = h$. Furthermore, h has optimal security $O(S/2^\kappa + T^2/2^n)$ (first formalized by [14] in the AI-ROM), so we are done! So our next goal is to try to use an h from minimal assumptions. Namely, can we get a keyed hash function with arbitrary length input from any $h: \{0,1\}^a \rightarrow \{0,1\}^n$ where a is much smaller, i.e. even $a = n + 1$? This will allow us to focus on building as simple a primitive as possible which we can bootstrap to a full variable-input length hash function with optimal security.

Next, we note that for any $a > \max(\kappa, n)$, we can always do the standard Merkle-Damgård transformation using h to construct g , where key is not fed into every invocation of h . For standard MD, it makes sense to set $\kappa = n$ since we use the key key as the initialization vector. Recall, the MD^h construction sets $y_0 = \text{key}$, computes $y_i = h(y_{i-1} \| m_i)$ for $i = 1, \dots, \ell$, and outputs y_ℓ . This approach has a major downside in that instantiating MD without inserting key into each invocation of h suffers non-trivial time-space tradeoffs. In general, there is an attack on the general MD^h construction with advantage $\Omega(ST^2/2^n)$. This is strictly worse than the optimal bound of $O(S/2^n + T^2/2^n)$ for any setting of parameters with $S, T \gg 1$. However, this attack finds very large—roughly length T —collisions. In our setting, we only care about using MD^h to get a function g with inputs of size $2n + s$. Thus, we leverage a recent line of work (see [11, 1, 9, 17, 2]) that shows that if you only use MD^h on ℓ -block messages, then the best known attack has advantage at most $O(ST\ell/2^n + T^2/2^n)$ (and further this is provably tight for constant ℓ [1, 17] and, when $ST^2 \leq 2^n$, is provably tight for all ℓ [2]).

So, if we instantiate g in our framework with a fixed-length MD^h construction where we set κ and s to be equal to the output length n , we get a construction for H^h where any (S, T) -attacker has advantage at most $O(ST/2^n + T^2/2^n)$ (up to $\text{poly}(n)$ factors). This is only “optimal,” however, under the (strong) assumption that $S \leq T$.

Our main technical contribution is instantiating g in the framework above using a new keyed Merkle tree approach, which does not require the assumption that $S \leq T$.

2.1 Keyed Merkle Tree Analysis

For the rest of this technical overview, we focus on our analysis of the keyed Merkle tree construction.

Construction. We start by defining the 2-level construction. We want a keyed function $g: \{0,1\}^\kappa \times \{0,1\}^{n+s} \rightarrow \{0,1\}^n$ from a hash function $h: \{0,1\}^a \rightarrow \{0,1\}^n$ where $a \geq 2n$. To do so, we split the $(n+s)$ -bit input into two parts, call them m_L, m_R , of size at most $a - \kappa$ (hence we require here that $a - \kappa \geq (n+s)/2$, or $a \geq \kappa + n/2 + s/2$). We concatenate each part with the key key and compute $y_L \leftarrow h(\text{key} \| m_L)$ and $y_R \leftarrow h(\text{key} \| m_R)$. We then concatenate y_L with y_R and feed the resulting string into h to get the output $z \leftarrow h(y_L \| y_R)$ (we require here that $a \geq 2n$). For technical purposes, we “domain separate” each call to h , so $y_L = h_1(\text{key} \| m_L)$, $y_R = h_2(\text{key} \| m_R)$, and $z = h_3(y_L \| y_R)$.

Analysis. We want to bound the probability that any (S, T) -attacker can find a collision in this keyed Merkle tree construction of g . To simplify the analysis here, we consider the case where κ and s are equal to n . Hence, each call to h takes two n -bit inputs and has one n -bit output. We also assume from the start that $ST^2 \leq 2^n$.

We start with the following observation. The probability that any (S, T) -attacker finds a collision at the leaves, corresponding to the calls to h_1, h_2 , is at most $O(S/2^n + T^2/2^n)$. Because these calls include the key key , if such an event happened, we could reduce to finding a collision in h directly. So, the challenge is to reason about the advantage of an (S, T) -attacker finding a collision at the second level of the Merkle tree, which is only implicitly related to the key key in the Merkle tree construction. We have to somehow characterize all possible ways that an (S, T) -attack can encode information about h in its advice string from the preprocessing phase.

One of the first tools one often turns to in such analysis is to use the *pre-sampling* technique from [27, 11]. We note that if we were to use the presampling technique, we would obtain a term of the form $ST/2^n$ in our bound, which is optimally secure only in the range $S \leq T$. Our main technical contribution is getting an optimally secure protocol for the range $S \gg T$, which therefore requires techniques other than presampling.

The AGL [2] framework: reducing to multi-instance games. To make our lives significantly easier, we use the multi-instance framework of [2], previously used in somewhat different forms in [21, 1, 9, 17]. At a very high level, this framework gives a way to reason about (S, T) -attackers using an *average-case* advice string rather than a worst-case one. In more detail, they show how to bound the advantage that any (S, T) -attacker \mathcal{A} finds a collision in g^h by the advantage of a (uniform) attacker \mathcal{B} in the following game. First, a random function h is

sampled. Then, \mathcal{B} has to win the following game for all $i = 1, \dots, S$ sequentially, where it is allowed to maintain arbitrary state (that it generates) between each successive game. In each game i , the attacker \mathcal{B} receives a random key key_i , its state from the previous games, and has to come up with a pair of messages $\text{msg}_i, \text{msg}'_i$ such that $g^h(\text{key}_i, \text{msg}_i) = g^h(\text{key}_i, \text{msg}'_i)$ using at most T queries to h . [9,2] show that if the advantage of \mathcal{B} is at most δ^S , then the advantage of the (S, T) -attacker \mathcal{A} is at most 2δ . The magic of this framework is that we can analyze the advantage of \mathcal{B} in each game i only given its state from the previous games, instead of having to reason about arbitrary advice strings as in the case of (S, T) -attackers. Namely, we can lazily sample h on any point that \mathcal{B} has not queried, in a way that is independent of \mathcal{B} 's current state.

Note that it suffices to show that the advantage of \mathcal{B} “in game i ” is at most δ *given* it has won all previous games. Let W_i be the event that \mathcal{B} wins game i and $W_{<i}$ be the event that \mathcal{B} wins all games before i . This follows since $\Pr[W_1 \wedge \dots \wedge W_S] = \prod_{i=1}^S \Pr[W_i | W_{<i}] \leq \delta^S$ if $\Pr[W_i | W_{<i}] \leq \delta$ for all $i \in [S]$. Hence, our goal is to show that $\Pr[W_i | W_{<i}] \leq O(S/2^n + T^2/2^n)$, up to $\text{poly}(n)$ factors.

Knowledge gaining event: bounding “hitting” queries. Now, to even further simplify the analysis of $\Pr[W_i | W_{<i}]$, we define a key “knowledge gaining event” (based on the techniques of [2]) representing the kind of information that \mathcal{B} may have encoded into its state based on the queries it made to h before game i has started. At a high level, this is an event that we show happens with very small probability (technically at most $2^{-2^{i \cdot n}}$) for an average-case advice string for \mathcal{B} at the start of game i . Then, assuming this event does not occur, we can more easily characterize the strategies of \mathcal{B} .

To define this event, we introduce some notation to characterize \mathcal{B} 's queries. We refer to all $(i-1) \cdot T$ queries \mathcal{B} makes before the start of game i as “offline” queries, and we refer to the T queries made in game i as “online” queries. An offline query is said to be “hitting” if its output is equal to an output or either of the two inputs to some prior query, i.e. it “hits” a previous query. We are now ready to state our key knowledge gaining event.

- We say that E_{hit}^i holds if there are more than $i \cdot \text{poly}(n)$ hitting queries among the $(i-1) \cdot T$ offline queries.

Briefly, we justify why E_{hit}^i holds with very small probability. The output for each query is uniformly sampled, and there are at most $3 \cdot (i-1) \cdot T$ values to hit across inputs/outputs in the previous $(i-1) \cdot T$ offline queries. So the probability each offline query is a hitting query is at most $3iT/2^n$, meaning we expect at most $3i^2T^2/2^n$ hitting queries accounting for all $(i-1) \cdot T$ offline queries. Furthermore, we show using a Chernoff bound that there will not be more than $i \cdot \text{poly}(n) \cdot \max(1, iT^2/2^n) = O(i)$ hitting queries (assuming $ST^2 \leq 2^n$) with high probability (recall that we ignore $\text{poly}(n)$ terms).

A case analysis based on collision queries. Now, assuming there are at most $i \cdot \text{poly}(n)$ hitting queries, we are ready to show that $\Pr[W_i | W_{<i}] \leq O(S/2^n +$

$T^2/2^n$). To do so, we look at the following “collision” queries corresponding to the valid collision $\text{msg}_i = (m_L, m_R) \neq \text{msg}'_i = (m'_L, m'_R)$ that \mathcal{B} outputs in game i (we assume that \mathcal{B} makes all of these queries at some point during or before game i).

- Q_1, Q_2, Q_3 are the queries $y_L \leftarrow h_1(\text{key}_i \| m_L)$, $y_R \leftarrow h_2(\text{key}_i \| m_R)$, $z \leftarrow h_3(y_L \| y_R)$, respectively.
- Q'_1, Q'_2, Q'_3 are the queries $y'_L \leftarrow h_1(\text{key}_i \| m'_L)$, $y'_R \leftarrow h_2(\text{key}_i \| m'_R)$, $z \leftarrow h_3(y'_L \| y'_R)$, respectively.

Recall that we assumed the collision occurs among Q_3, Q'_3 (not at the leaves), so it must be the case that $(y_L, y_R) \neq (y'_L, y'_R)$ and queries Q_3, Q'_3 are distinct.

If all of these collision queries are online (were first made during game i), then clearly $\Pr[W_i | W_{<i}] \leq O(T^2/2^n)$. Specifically, as Q_3 and Q'_3 are online (and distinct by assumption) and form a collision, this follows by a birthday bound on at most T online—and hence lazily sampled—queries that \mathcal{B} makes during game i . The challenge comes when analyzing the cases where \mathcal{B} may have made some of these queries *before* game i , so it could have encoded information about these queries in its state. To do so, we consider the following remaining cases, which cover all possible strategies that \mathcal{B} may employ.

As we already considered when both Q_3, Q'_3 are online queries, it must be the case that one of Q_3, Q'_3 must be an offline query. Assume without loss of generality that Q_3 is offline. Then either **(A)** both Q_1, Q_2 are online, **(B)** exactly one of Q_1, Q_2 are online, or Q_1, Q_2 are both offline. The latter case implies that Q_1, Q_2, Q_3 are all offline. Then either is the case that **(C)** Q'_3 is online, or Q'_3 is also offline. If Q'_3 is also offline, then either we reduce to case (A) or (B) above by symmetry, or it holds that **(D)** all $Q_1, Q_2, Q_3, Q'_1, Q'_2, Q'_3$ are offline. So, it suffices to show in cases (A-D) that $\Pr[W_i | W_{<i}] \leq \delta \leq O(S/2^n + T^2/2^n)$, at least assuming $ST^2 \leq 2^n$. We proceed to give the main ideas behind each of these cases.

(A) Q_3 is offline but Q_1, Q_2 are online.

There are at most $(i-1) \cdot T$ options for the offline query Q_3 . Both online queries Q_1 and Q_2 have to hit such a query, which happens with at most $(i-1) \cdot T \cdot (T/2^n)^2 = O(iT^3/2^{2n}) \leq O(T/2^n)$ when $ST^2 \leq 2^n$ since $i \leq S$.

(B) Q_1, Q_3 are offline but Q_2 is online (symmetrically for Q_1 online and Q_2 offline).

In this case, we claim we can “associate” the key key_i in game i to the query Q_3 since both Q_1 and Q_3 are offline queries. If we can associate key_i to at most k possible Q_3 queries, then this implies the probability the output of some online query hits an input of such an associated query is at most $k \cdot T/2^n$. But how many Q_3 queries can we associate to a key key_i ?

In the worst case, key_i may be associated to $(i-1) \cdot T$ many Q_3 queries, but this implies a suboptimal bound of $O(iT^2/2^n)$. But this cannot be true for too many values of key simultaneously. In particular, if a Q_3 query is associated with more than k possible values of key , this means there are k hitting queries, so each Q_3 query can be associated with at most $O(i)$ keys.

This implies there are at most $O(i^2T)$ pairs of associated key values with potential Q_3 queries, meaning a random key_i value will be associated with at most $O(i^2T/2^n)$ potential Q_3 values on average. Plugging this average-case bound into k above, this implies a bound of $O(i^2T^2/2^{2n}) \leq O(S/2^n)$ assuming $ST^2 \leq 2^n$ given $i \leq S$.

(C) Q_1, Q_2, Q_3 are offline but Q'_3 is online.

In this case, Q'_3 is a distinct query from Q_3 by assumption, but again we can “associate” Q_3 with key_i as above. Then, since Q'_3 must share an output with Q_3 , the same argument as above gives a bound of $O(S/2^n)$ in this case.

(D) All collision queries $Q_1, Q_2, Q_3, Q'_1, Q'_2, Q'_3$ are offline.

In this case, we show that every full collision structure with respect to some key among the offline queries leads to a hitting query. Furthermore, two collision structures cannot share the same hitting query. So if there are at most $O(i)$ hitting queries, the probability $Q_1, Q_2, Q_3, Q'_1, Q'_2, Q'_3$ are all offline for a random key_i is at most $O(i/2^n) \leq O(S/2^n)$.

Thus, we showed that $\Pr[W_i | W_{<i}] \leq O(S/2^n + T^2/2^n)$ no matter when $Q_1, Q_2, Q_3, Q'_1, Q'_2, Q'_3$ were queried before or during game i . Further, recall the last case where there is a collision at one of the leaves (corresponding to Q_1, Q'_1 or Q_2, Q'_2), which can happen with at most $O(S/2^n + T^2/2^n)$ probability since such a collision directly involves key_i . Thus, in all possible cases, we have shown $\Pr[W_i | W_{<i}] \leq \delta \leq O(S/2^n + T^2/2^n)$, at least assuming $ST^2 \leq 2^n$. Finally, by the framework of [9, 2], this implies the same bound (up to a multiplicative factor of 2) on the advantage of finding a collision for any (S, T) -attacker.

3 Preliminaries

We let $\mathbb{N} = \{1, 2, 3, \dots\}$ denote the natural numbers. The set of all functions with domain D and range R is denoted by $\text{Fcs}(D, R)$. We let $*$ denote a wildcard element. For example $(*, z) \in L$ is true if there is an ordered pair in L where z is the second element (the type of the wildcard element shall be clear from the context). For a random variable X we use $\mathbb{E}[X]$ to denote its expected value. We use $x \leftarrow_s D$ to denote sampling x uniformly sampling from the elements of D . All logarithms in this paper are for base 2 unless otherwise specified.

For a bit-string s , we use $|s|$ to denote the number of bits in s . For two strings s_1, s_2 , we use $s_1 \| s_2$ to denote the concatenation of two strings. We use standard regular expression notation where s^* denotes 0 or more copies of s , s^+ denotes one or more copies of s , and s^k denotes k copies of s . Similarly, for a set S , we use S^* , S^+ , and S^k to represent 0 or more, 1 or more, of k elements takes from a set S . In particular, we use $\{0, 1\}^*$ to represent any arbitrary string of bits. We use the notation $\{0, 1\}^{\leq k}$ to represent a string of length at most k .

Chernoff bound. We state a Chernoff bound which we use in the technical part of the paper.

Proposition 1. *Let $n \in \mathbb{N}$. Let X_1, X_2, \dots, X_n be independent 0-1 random variables. Let $X = \sum_{i=1}^n X_i$. Let μ' be such that $\mathbb{E}[X] \leq \mu'$. Then we have that*

$$\Pr[X \geq (1 + \delta)\mu'] \leq e^{-\frac{\delta\mu'}{3}}.$$

Notice that this version is somewhat non-standard as it even works when we know only an upper bound on the expectation (usually, in standard formulations of Chernoff bound, we need to know the expectation exactly). We have a proof of this Chernoff bound in the full version.

Auxiliary-input Random Oracle Model (AI-ROM). The auxiliary-input random oracle model, introduced by Unruh [27], captures the power of non-uniform adversaries against random oracles. An attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in this model is formalized as a two stage adversary. In its first stage, which is referred to as the preprocessing phase, \mathcal{A}_1 has unbounded access to the random oracle h , and outputs any arbitrary S -bit advice string or auxiliary input σ . In the second stage, referred to as the online phase, gets σ as input, \mathcal{A}_2 can make at most T queries to its oracle h . Its aim is to accomplish some task involving h , e.g. find a collision in a construction based on h . We refer to such an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ as an (S, T) -attacker.

Collision resistance of g^h in AI-ROM. We next formalize the keyed-collision resistance of an iterated hash function construction g relative to a hash function $h : \{0, 1\}^a \rightarrow \{0, 1\}^n$ in the AI-ROM. The construction g has a parameter κ associated with it, where κ is the bit length of key used in g . It first samples a random function $h : \{0, 1\}^a \rightarrow \{0, 1\}^n$. The adversary \mathcal{A}_1 gets unbounded access to h , and it outputs an advice string σ . At this time, \mathcal{A}_2 is given the auxiliary input σ , a randomly sampled key from $\{0, 1\}^\kappa$, as well as oracle access to h , and it needs to find $\text{msg} \neq \text{msg}'$ such that $g^h(\text{key}, \text{msg}) = g^h(\text{key}, \text{msg}')$. This game, denoted $\mathsf{G}_{g^h}^{\text{ai-cr}}$ is formally defined in Fig. 7.

Definition 1 (AI-CR Advantage). *The advantage of an adversary \mathcal{A} against the collision resistance of g^h in the AI-ROM is*

$$\text{Adv}_{g^h}^{\text{ai-cr}}(\mathcal{A}) = \Pr \left[\mathsf{G}_{g^h}^{\text{ai-cr}}(\mathcal{A}) = \text{true} \right].$$

For parameters $S, T \in \mathbb{N}$, we overload notation and denote

$$\text{Adv}_{g^h}^{\text{ai-cr}}(S, T) = \max_{\mathcal{A}} \left\{ \text{Adv}_{g^h}^{\text{ai-cr}}(\mathcal{A}) \right\},$$

where the maximum is over all (S, T) -attackers.

Throughout the paper, for any (S, T) -attacker \mathcal{A} that outputs messages msg, msg' that causes $\mathsf{G}_{g^h}^{\text{ai-cr}}(\mathcal{A})$ to output true on a key key , we assume that \mathcal{A} has fully queried $g^h(\text{key}, \text{msg})$ and $g^h(\text{key}, \text{msg}')$. This is true without loss of generality (up to constant factors in the advantage) as if there exists any (S, T) -attacker \mathcal{A} that does not, you can construct an $(S, T + 2\ell)$ -attacker \mathcal{B} that does,

<p>Game $G_{g^h}^{\text{ai-cr}}(\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2))$</p> <ol style="list-style-type: none"> 1. $h \leftarrow \text{Fcs}(\{0, 1\}^a, \{0, 1\}^n)$ 2. $\sigma \leftarrow \mathcal{A}_1(h)$ 3. $\text{key} \leftarrow \{0, 1\}^\kappa$ 4. $(\text{msg}, \text{msg}') \leftarrow \mathcal{A}_2^h(\sigma, \text{key})$ 5. Return true if: <ol style="list-style-type: none"> (a) $\text{msg} \neq \text{msg}'$, and (b) $g^h(\text{key}, \text{msg}) = g^h(\text{key}, \text{msg}')$ 6. Else, return false

Fig. 7: The collision resistance game $G_{g^h}^{\text{ai-cr}}$ in AI-ROM for a function g^h based on a random oracle $h : \{0, 1\}^a \rightarrow \{0, 1\}^n$. The construction g^h has a parameter κ associated with it, where κ is the bit length of **key** used in g .

where g^h requires at most ℓ invocations of h to compute either $g^h(\text{key}, \text{msg})$ or $g^h(\text{key}, \text{msg}')$. As $\ell \leq T$, the resulting attacker will have comparable advantage up to constant factors in T . We note that this is a standard assumption in existing related works in the AI-ROM.

On padding. The variable-input length hash functions we consider of this work all act on messages which have been parsed into many fixed size blocks of some specified size s . We therefore need a padding function that takes arbitrary length inputs and converts them to a sequence of fixed-size blocks. We need to ensure that this padding function maintains certain properties like injectivity in order to guarantee that if an adversary finds a collision on the padded versions of messages, then it implies a collision with respect to the underlying messages as well. For the purpose of this paper, we define the following padding function, which is a slightly simplified version of the padding function used by the SHA family of hash functions (see [19, Section 8.5] for more discussion on MD-compliant padding functions). The function **pad** we use takes in a message $\text{msg} \in \{0, 1\}^*$, an integer $s \in \mathbb{N}$ representing the size of each block, and an integer n that stipulates that $|\text{msg}| \leq 2^n$. The construction is formally defined as follows.

pad(msg, s, n):

1. Let $k = s - ((|\text{msg}| + n) \bmod s + 1)$.
2. Interpret $|\text{msg}| \in [2^n]$ as an n -bit string.
3. Output $(m_1, \dots, m_\ell) \in (\{0, 1\}^s)^\ell$ where $m_1 \parallel \dots \parallel m_\ell = \text{msg} \parallel |\text{msg}| \parallel 1 \parallel 0^k$.

We formalize the guarantees we use for this padding function in the following theorem.

Theorem 5 (Padding). *Let $s, n \in \mathbb{N}$. The function **pad**(msg, s, n) on messages $\text{msg} \in \{0, 1\}^{2^n}$ satisfies the following properties:*

1. $|\text{pad}(\text{msg}, s, n)| \in (\{0, 1\}^s)^\ell$ for $\ell = \lceil (|\text{msg}| + n + 1)/s \rceil$.

Construction $\text{KMD}^h(\text{key}, \text{msg})$:

1. $(m_1, m_2, \dots, m_\ell) \leftarrow \text{pad}(\text{msg}, s, n)$ where $\ell = \lceil (|\text{msg}| + n + 1)/s \rceil$ given by Theorem 5.
2. Initialize $y_0 = 0^n$, and compute $y_i \leftarrow h(\text{key}, y_{i-1} \| m_i)$ for $i = 1, \dots, \ell$.
3. Output y_ℓ .

Fig. 8: The keyed Merkle-Damgård construction $\text{KMD}^h: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ given any underlying function $h: \{0, 1\}^\kappa \times \{0, 1\}^{n+s} \rightarrow \{0, 1\}^n$, where κ is the key length, n is the output length, and s is the message block size.

2. *There is a unique decoding procedure that outputs msg given $\text{pad}(\text{msg}, s, n)$, and outputs \perp on invalid padded messages.*
3. *If $\text{pad}(\text{msg}, s, n) = \text{pad}(\text{msg}', s, n)$, then $\text{msg} = \text{msg}'$.*
4. *If $|\text{msg}| < |\text{msg}'|$, then $\text{pad}(\text{msg}, s, n)$ is not a suffix of $\text{pad}(\text{msg}', s, n)$.*

We defer the proof of this theorem to the full version.

4 Merkle-Damgård Framework with a Keyed Inner Hash

In this section, we lay out the general framework for our main results, based on the Merkle-Damgård transform using a keyed inner hash. We note that this framework has been explicitly considered in the uniform setting in Section 8.5 of the lecture notes of Goldwasser and Bellare [19].⁹ We extend this framework to the preprocessing setting, modeled by the AI-ROM of Unruh [27], noting that the high level ideas are similar.

For a key length κ , output length n , and message block size s , we assume an underlying primitive $h: \{0, 1\}^\kappa \times \{0, 1\}^{n+s} \rightarrow \{0, 1\}^n$. In other words, viewing the primitive h as a function from $\{0, 1\}^a$ to $\{0, 1\}^n$, this implies that $a = \kappa + n + s$.

Given such a primitive h , we define the following keyed Merkle-Damgård hash function $\text{KMD}^h: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$. On input key $\text{key} \in \{0, 1\}^\kappa$ and message $\text{msg} \in \{0, 1\}^*$ of length at most 2^n , the function KMD^h first pads the message to split it into $\ell = \lceil (|\text{msg}| + n + 1)/s \rceil$ message blocks of size s as in Theorem 5. It then essentially computes the Merkle-Damgård hash function using the underlying hash function h , except that the key key is inserted into every invocation of h . This is formalized in Figure 8.

Since the key key is included in every call to the underlying primitive h , it follows that any (S, T) -attacker that finds a collision in KMD^h with respect to a key key also finds a collision in h with respect to key . This is formalized via a reduction, which gives the following theorem. Again, we note that the following

⁹ The existence of this variant of the Merkle-Damgård transform has gone completely unnoticed in recent works studying non-uniform security of this transformation [11, 1, 17, 2].

theorem very closely follows the reduction given in [19, Section 8.5], but we give the full details in the AI-ROM for completeness.

Theorem 6. *Let $\kappa, n, s \in \mathbb{N}$. Let $h: \{0, 1\}^\kappa \times \{0, 1\}^{n+s} \rightarrow \{0, 1\}^n$ be any function, and let $\text{KMD}^h: \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^n$. Then, for every $S, T \in \mathbb{N}$, it holds that*

$$\text{Adv}_{\text{KMD}^h}^{\text{ai-cr}}(S, T) \leq \text{Adv}_h^{\text{ai-cr}}(S, T).$$

The proof of this theorem is a straightforward reduction, and we defer it to the full version.

Next, we recall that if h is a keyed function modeled as a random oracle (in the AI-ROM of Unruh [27]), Dodis, Guo, and Katz [14] give the following bound on the success probability that any (S, T) -attacker can find a collision in h with respect to a random key.

Theorem 7 ([14]). *Let $h: \{0, 1\}^\kappa \times \{0, 1\}^b \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM. Then, for any $S, T \in \mathbb{N}$,*

$$\text{Adv}_h^{\text{ai-cr}}(S, T) \leq \frac{2S + 2\kappa}{2^\kappa} + \frac{50T^2}{2^n}.$$

Combining Theorems 6 and 7, we get the following result.

Theorem 8. *Let $a, \kappa, n \in \mathbb{N}$ be such that $a > \kappa + n$. Let $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM. Then, there is an $H^h: \{0, 1\}^\kappa \times \{0, 1\}^{<2^n} \rightarrow \{0, 1\}^n$ such that:*

1. *For any $S, T \in \mathbb{N}$,*

$$\text{Adv}_{H^h}^{\text{ai-cr}}(S, T) \leq \frac{2S + 2\kappa}{2^\kappa} + \frac{50T^2}{2^n}.$$

2. *One evaluation of H^h on messages of length b requires $\lceil (b+n+1)/s \rceil$ queries to h , where $s = a - \kappa - n$.*

Proof. As $a > \kappa + n$, we define $s = a - \kappa - n > 0$. Then, we view h as a function from $\{0, 1\}^\kappa \times \{0, 1\}^{n+s}$ to $\{0, 1\}^n$ and use it in the construction KMD^h of Figure 8 to get the hash function H^h required by the theorem.

The bound on the advantage immediately follows as a corollary of Theorems 6 and 7. As for efficiency, we note that padding a message $\text{msg} \in \{0, 1\}^b$ via Theorem 5 results in a message consisting of ℓ blocks each of length s , where $\ell = \lceil (b+n+1)/s \rceil$. Each block of the message requires a single invocation of h , so evaluating KMD^h on msg requires $\ell = \lceil (b+n+1)/s \rceil$ queries to h as required. ■

Construction $\text{FMD}_b^h(\text{key}, \text{msg})$:

1. Parse msg into ℓ blocks $(m_1, m_2, \dots, m_\ell)$ each of size s where $\ell = \lceil b/s \rceil$ and m_ℓ is padded with 0s if needed.
2. Initialize $y_0 = \text{key}$, and compute $y_i \leftarrow h(y_{i-1}, m_i)$ for $i = 1, \dots, \ell$.
3. Output y_ℓ .

Fig. 9: The standard Merkle-Damgård construction with fixed input message length b $\text{FMD}_b^h: \{0, 1\}^n \times \{0, 1\}^b \rightarrow \{0, 1\}^n$ given any underlying function $h: \{0, 1\}^n \times \{0, 1\}^s \rightarrow \{0, 1\}^n$, where n is the key and output length, and s is the message block size.

5 Instantiating the Inner Hash: Standard MD

We next consider instantiating the Merkle-Damgård framework of Section 4 whenever the underlying hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ has input length a such that $n < a < \kappa + n + s$, where κ is the key length, n is the output length, and s is the desired message block size. Specifically, our goal is to use such a primitive $h: \{0, 1\}^a \times \{0, 1\}^n$ to build a larger (fixed-length) hash function $g^h: \{0, 1\}^\kappa \times \{0, 1\}^{n+s}$ that *can* be plugged into the construction KMD^{g^h} .

The first approach we consider is by simply building g from h using the standard Merkle-Damgård construction where h is not keyed in every invocation. We emphasize that in the standard version of MD, a random initialization vector/key is still included in the first invocation of h . However, it is not included in the subsequent invocations of h , allowing h to take in smaller inputs overall.

Given an underlying hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ where $a > \max(\kappa, n)$, we define the standard Merkle-Damgård hash function with fixed input message length b $\text{FMD}_b^h: \{0, 1\}^\kappa \times \{0, 1\}^b \rightarrow \{0, 1\}^n$ as follows. For sake of simplicity and due to the nature of the MD construction, we will assume that the key length κ is equal to the output length n . Let $s = a - n$ be the message block size we will include in each invocation of the underlying h . On input a key $\text{key} \in \{0, 1\}^n$ and a message $\text{msg} \in \{0, 1\}^b$, the function FMD_b^h splits the message msg into $\ell = \lceil b/s \rceil$ message blocks of size s (adding 0s to the last block if needed). It initializes $y_0 = \text{key}$, and for $i = 1, \dots, \ell$, computes y_i as the hash of y_{i-1} concatenated with m_i using h . The output of FMD_b^h is then y_ℓ . This is formalized in Figure 9.

If we instantiate $g^h := \text{FMD}_b^h$ only on input messages of size at most $b = n + s$ as required to instantiate it inside KMD^{g^h} , then we only need to worry about ℓ -block collisions for $\ell = \lceil (n + s)/(a - n) \rceil$. Akshima, Guo, and Liu [2] currently show the best known upper bound on the advantage of finding general ℓ -block collisions in MD^h , given by the following theorem.

Theorem 9 ([2]). *Let $b, n, s \in \mathbb{N}$ such that $b > s$, and set $\ell = \lceil b/s \rceil$. Let $h: \{0, 1\}^n \times \{0, 1\}^s \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM.*

Then, for any $S, T \in \mathbb{N}$,

$$\text{Adv}_{\text{FMD}_b^h}^{\text{ai-cr}}(S, T) \leq \begin{cases} \frac{200n \cdot (ST + T^2)}{2^n} & \text{if } \ell = 2, \text{ and} \\ \frac{34n \cdot ST\ell}{2^n} \cdot \max\left(1, \frac{ST^2}{2^n}\right) + \frac{2 \cdot T^2}{2^n} & \text{if } \ell > 2. \end{cases}$$

We note that Ghoshal and Komargodski [17] give a bound of $O(ST/2^n + T^2/2^n)$ whenever ℓ is a constant, which doesn't require the assumption that $ST^2 \leq 2^n$. However, their bound does not extend to super constant ℓ .

Combined with Theorem 6, we get the following result.

Theorem 10. *Let $a, n, s \in \mathbb{N}$ be such that $a > n$. Let $\ell = \lceil (n + s)/(a - n) \rceil$. Let $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM. Then, there is an $H^h: \{0, 1\}^\kappa \times \{0, 1\}^{<2^n} \rightarrow \{0, 1\}^n$ such that:*

1. *For any $S, T \in \mathbb{N}$,*

$$\text{Adv}_{H^h}^{\text{ai-cr}}(S, T) \leq \begin{cases} \frac{200n \cdot (ST + T^2)}{2^n} & \text{if } \ell = 2, \text{ and} \\ \frac{34n \cdot ST\ell}{2^n} \cdot \max\left(1, \frac{ST^2}{2^n}\right) + \frac{2 \cdot T^2}{2^n} & \text{if } \ell > 2. \end{cases}$$

2. *One evaluation of H^h on messages of length b requires $\ell \cdot \lceil (b + n + 1)/s \rceil$ queries to h .*

Proof. As $a > n$, we parse $h: \{0, 1\}^n \times \{0, 1\}^{a-n} \rightarrow \{0, 1\}^n$ and use it to construct FMD_{n+s}^h as defined in Figure 9. We then use FMD_{n+s}^h as the primitive underlying our MD-based hash function of Figure 8. So, we set

$$H^h := \text{KMD}^{\text{FMD}_{n+s}^h}.$$

The bound on the advantage follows as a corollary to Theorems 6 and 9, where the message length required for the FMD_{n+s}^h construction is only $n + s$. Furthermore, this implies that FMD_{n+s}^h requires $\ell = \lceil (n + s)/(a - n) \rceil$ invocations of h per invocation of FMD_{n+s}^h , and $\text{KMD}^{\text{FMD}_{n+s}^h}$ requires $\lceil (b + n + 1)/s \rceil$ invocations of FMD_{n+s}^h , giving the resulting efficiency bound. ■

We emphasize that because our reduction in Theorem 6 is generic, any improvement on the bound of [2, 17] for finding an ℓ -block collision in MD will immediately imply an improved bound in Theorem 10.

We next state a corollary of Theorem 10 where we restrict to using $\ell = 2$ invocations of h in the underlying FMD_{n+s}^h construction. For this setting, the bound of [2] above is optimal (up to $\text{poly}(n)$ factors), and we observe that the resulting bound matches our desired bound of $O(S/2^n + T^2/2^n)$ whenever $S \leq T$.

Corollary 1. *Let $a, n, s \in \mathbb{N}$ be such that $a \geq 3n/2 + s/2$. Let $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM. Then, there is an $H^h: \{0, 1\}^\kappa \times \{0, 1\}^{<2^n} \rightarrow \{0, 1\}^n$ such that:*

Game $G_{g,S}^{\text{mi-cr}}(\mathcal{B})$

1. $h \leftarrow \$ \text{Fcs}(\{0,1\}^a, \{0,1\}^n)$
2. $\text{key}_1, \dots, \text{key}_S \leftarrow \$ \{0,1\}^\kappa$
3. $\text{st} \leftarrow \perp$
4. For $i = 1$ to S
 - (a) $(\text{msg}, \text{msg}', \text{st}) \leftarrow \$ \mathcal{B}^h(\text{st}, \text{key}_i)$
 - (b) Output false if:
 - i. $\text{msg} = \text{msg}'$, or,
 - ii. $g^h(\text{key}_i, \text{msg}) \neq g^h(\text{key}_i, \text{msg}')$
5. Otherwise, output true

Fig. 10: The multi-instance game $G_{g,S}^{\text{mi-cr}}$, where \mathcal{B} is a uniform adversary with oracle access to the function h .

1. For any $S, T \in \mathbb{N}$,

$$\text{Adv}_{H^h}^{\text{ai-cr}}(S, T) \leq \frac{200n \cdot (ST + T^2)}{2^n}$$

2. One evaluation of H^h on messages of length b requires $2 \cdot \lceil (b + n + 1)/s \rceil$ queries to h .

Proof. Restricting to $\ell = 2$ in Theorem 10, we require that $\ell = \lceil (n + s)/(a - n) \rceil \leq 2$. This holds as long as $a \geq 3n/2 + s/2$, as required. ■

6 Instantiating the Inner Hash: Two-Level Merkle Tree

In this section, we instantiate the Merkle-Damgård framework of Section 4, that uses a keyed inner hash, in the setting where the underlying hash function $h: \{0,1\}^a \rightarrow \{0,1\}^n$ satisfies $a \geq \max(2n + 2, \kappa + \lceil n/2 \rceil + 3)$.

The compression function in this instantiation of the MD-based framework, is a Merkle tree with two leaves, where we additionally input the key into each leaf. We describe next the framework introduced in [2] that we use to analyze the collision-resistance of this construction in AI-ROM.

6.1 The AGL [2] Framework

In this section, we briefly introduce the framework given by [2] which is useful in analyzing non-uniform security. An earlier version of this framework was introduced by [9,1] inspired by techniques used in proving constructive Chernoff bounds in [21] and later refined by [1,17,2] to upper bound $\text{Adv}_{g^h}^{\text{ai-cr}}(S, T)$. This framework involves upper bounding the advantage of an (S, T) -attacker using the advantage of a uniform adversary for a multi-instance game that has to find collisions for S randomly chosen values of key.

We define the “multi-instance” game $G_{g^h, S}^{\text{mi-cr}}(\mathcal{B})$ in Fig. 10. We refer an adversary playing $G_{g^h, S}^{\text{mi-cr}}$ and making at most T queries for each key as a (S, T) -MI adversary. For any (S, T) -MI adversary, define

$$\text{Adv}_{g^h}^{\text{mi-cr}}(\mathcal{B}) = \Pr \left[G_{g^h, S}^{\text{mi-cr}}(\mathcal{B}) = \text{true} \right].$$

Further,

$$\text{Adv}_{g^h}^{\text{mi-cr}}(S, T) = \max_{\mathcal{B}} \text{Adv}_{g^h}^{\text{mi-cr}}(\mathcal{B}),$$

where the maximum is taken over all (S, T) -MI adversaries. The following key lemma relates $\text{Adv}_{g^h}^{\text{ai-cr}}(S, T)$ to $\text{Adv}_{g^h}^{\text{mi-cr}}(S, T)$, which is proven in [2].

Lemma 1 ([2]). *Fix $S, T \in \mathbb{N}$ and $0 \leq \delta \leq 1$, if $\text{Adv}_{g^h}^{\text{mi-cr}}(S, T) \leq \delta^S$, then $\text{Adv}_{g^h}^{\text{ai-cr}}(S, T) \leq 2\delta$.*

Offline and Online queries. Since the adversary in the multi-instance game is stateful, we can assume without loss of generality that it does not repeat queries since they can simply remember the answers. Additionally, [2] formalized the notion of “offline” and “online” queries during a particular instance of the game. When running the adversary on key_i , the queries that were made while the adversary was run on $\text{key}_1, \dots, \text{key}_{i-1}$ are collectively known as the “offline” queries, and the queries made while running on key_i are “online queries”.

6.2 Two-Level Merkle Tree

In this section, we present our construction of a keyed Merkle tree and analyze its collision resistance in the AI-ROM using the framework in the previous section. Specifically, given an underlying hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ where $a \geq \max(2n + 2, \kappa + \lceil n/2 \rceil + 3)$, we define a keyed, 2-level Merkle tree 2MT_b^h for message length $b \leq 2a - 2\kappa - 4$.

Before we define 2MT_b^h , we introduce notation that allows use to “domain-separate” h into three separate functions. Given a fixed hash function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$, we define three domain-separated functions $h_1, h_2, h_3: \{0, 1\}^{a-2} \rightarrow \{0, 1\}^n$, where $h_i(x)$ outputs $h(\hat{i} \| x)$ where $\hat{i} \in \{0, 1\}^2$ is the 2-bit binary representation of i . Moreover, we refer to a query $h(\hat{i} \| *)$ as a query to the function h_i (which is also clearly a query to h).

To construct $2\text{MT}_b^h: \{0, 1\}^\kappa \times \{0, 1\}^b \rightarrow \{0, 1\}^n$, we use h_1 and h_2 above to process the two leaves of the depth-2 Merkle tree, where we include $\text{key} \in \{0, 1\}^\kappa$ in each leaf. We then feed those outputs as input to h_3 to get the output of 2MT_b^h . This construction is formalized in Figure 11.

Our main result of this section is the following theorem, which bounds the probability that any (S, T) -attacker finds a collision in 2MT_b^h .

Theorem 11. *Let $a, \kappa, n \in \mathbb{N}$ be such that $a \geq \max(\kappa + \lceil n/2 \rceil + 3, 2n + 2)$. Let $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM. Then, for $s = 2a - 2\kappa - n - 4$, the construction $2\text{MT}_{n+s}^h: \{0, 1\}^\kappa \times \{0, 1\}^{n+s} \rightarrow \{0, 1\}^n$ of Figure 11 satisfies the following.*

$2\text{MT}_b^h(\text{key}, \text{msg})$

1. Parse $\text{msg} \in \{0, 1\}^{2a-2\kappa-4}$ as $\text{msg}_L \parallel \text{msg}_R$ where $\text{msg}_L, \text{msg}_R \in \{0, 1\}^{a-\kappa-2}$.
2. Compute $y_L \leftarrow h_1(\text{key} \parallel \text{msg}_L)$.
3. Compute $y_R \leftarrow h_2(\text{key} \parallel \text{msg}_R)$.
4. Output $z \leftarrow h_3(y_L \parallel y_R)$.

Fig. 11: The two-level, keyed Merkle tree construction $2\text{MT}_b^h: \{0, 1\}^\kappa \times \{0, 1\}^b \rightarrow \{0, 1\}^n$ with fixed input message length b given any underlying function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$, where $a \geq \max(2n + 2, \kappa + 3)$ and $b \leq 2a - 2\kappa - 4$. h_1, h_2, h_3 are all domain-separated using the first two bits of h to encode 1, 2, 3, respectively.

– For any $S, T \in \mathbb{N}$,

$$\text{Adv}_{g^h}^{\text{ai-cr}}(S, T) \leq \left(\frac{S}{2^\kappa} \cdot (14n + 42n\gamma + 42n\gamma^2) + \frac{T^2}{2^n} \cdot \left(2 + \frac{2\gamma}{T} + \frac{2}{T^2} \right) \right),$$

where $\gamma = ST^2/2^n$.

We prove this theorem using the framework described in Section 6.1. So, by Lemma 1, it suffices to prove the following lemma which bounds the advantage of an (S, T) -MI adversary.

Lemma 2. *Let $S, T \in \mathbb{N}$. Then*

$$\text{Adv}_{g^h}^{\text{mi-cr}}(S, T) \leq \left(\frac{S}{2^\kappa} \cdot (7n + 21n \cdot \gamma + 21n \cdot \gamma^2) + \frac{T^2}{2^n} \cdot \left(1 + \frac{\gamma}{T} + \frac{1}{T^2} \right) \right)^S,$$

where $\gamma = ST^2/2^n$.

Proof. Following the techniques of [2], we reduce the task of bounding $\text{Adv}_{g^h}^{\text{mi-cr}}(S, T)$ to that of bounding any T -query adversaries advantage of succeeding in iteration i given that it has succeeded in all previous iterations. Fix any (S, T) -MI attacker \mathcal{A} . Let W_i be the indicator random variable that \mathcal{A} wins on key_i in $G_{g^h, S}^{\text{mi-cr}}$. Define the random variable $W_{<i} := W_1 \wedge \dots \wedge W_{i-1}$. We have that

$$\text{Adv}_{g^h}^{\text{mi-cr}}(\mathcal{A}) = \Pr[W_1 \wedge W_2 \wedge \dots \wedge W_S] = \prod_{i=1}^S \Pr[W_i | W_{<i}].$$

We prove in the full version that for every \mathcal{A} and each $i \in [S]$, $\Pr[W_{<i+1}] \leq (\delta_S)^i$ where

$$\delta_S = \frac{T^2}{2^n} + 7n \cdot \frac{S}{2^\kappa} + 21n \cdot \frac{S^2 T^2}{2^{n+\kappa}} + \frac{ST^3}{2^{2n}} + 21n \cdot \frac{S^3 T^4}{2^{2n+\kappa}} + \frac{1}{2^n}.$$

It follows that for any (S, T) -MI attacker \mathcal{A} ,

$$\text{Adv}_{g^h}^{\text{mi-cr}}(\mathcal{A}) = \Pr[W_{<S+1}] = \prod_{i=1}^S \Pr[W_i | W_{<i}] \leq (\delta_S)^S.$$

As this holds for any such \mathcal{A} , it follows that $\text{Adv}_{g^h}^{\text{mi-cr}}(S, T) \leq (\delta_S)^S$, as required by the lemma statement. This completes the proof of Lemma 2. ■

6.3 Variable-Input Length Hash from Two-Level Merkle Trees

We combine the construction of Section 6.2 with the framework of Section 4 to get a variable-input length hash function. This construction is optimally secure as long as $ST^2 \leq 2^n$ and requires an underlying function $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ where $a \geq \max(\kappa + \lceil n/2 \rceil + 3, 2n + 2)$. This results in the following theorem.

We note if we modify values of κ, n by additive constant factors, we can get the same result as below with only $O(1)$ multiplicative loss in security. In this sense, we can achieve the theorem below assuming a function $h: \{0, 1\}^{2n} \rightarrow \{0, 1\}^n$, i.e. $a = 2n$ so only compressing by a factor exactly two.

Theorem 12. *Let $a, \kappa, n \in \mathbb{N}$ be such that $a \geq \max(\kappa + \lceil n/2 \rceil + 3, 2n + 2)$. Let $h: \{0, 1\}^a \rightarrow \{0, 1\}^n$ be modeled as a random oracle in the AI-ROM. Then, there is an $H^h: \{0, 1\}^\kappa \times \{0, 1\}^{<2^n} \rightarrow \{0, 1\}^n$ such that:*

1. *For any $S, T \in \mathbb{N}$,*

$$\text{Adv}_{H^h}^{\text{ai-cr}}(S, T) \leq \left(\frac{S}{2^\kappa} \cdot (14n + 42n\gamma + 42n\gamma^2) + \frac{T^2}{2^n} \cdot \left(2 + \frac{2\gamma}{T} + \frac{2}{T^2} \right) \right),$$

where $\gamma = ST^2/2^n$.

2. *One evaluation of H^h on a message b bits long requires $3 \cdot \lceil (b + n + 1)/s \rceil$ queries to h where $s = 2a - 2\kappa - n - 4$.*

Proof. We define $H := \text{KMD}^{2\text{MT}_{n+s}^h}$ where 2MT_{n+s}^h is defined in Figure 11. From Theorem 6, we have that $\text{Adv}_{H^h}^{\text{ai-cr}}(S, T)$ is upper bounded by $\text{Adv}_{2\text{MT}_{n+s}^h}^{\text{ai-cr}}(S, T)$. Therefore, the bound on the advantage of any (S, T) -attacker on H^h follows from Theorem 11.

We have that $2\text{MT}_{n+s}^h \{0, 1\}^\kappa \times \{0, 1\}^{n+s} \rightarrow \{0, 1\}^n$, for $s = 2a - 2\kappa - n - 4$. A b bit message, after padding, will result in $\lceil (b + n + 1)/s \rceil$ message blocks that are fed into 2MT_{n+s}^h . For each call of 2MT_{n+s}^h , we need 3 calls to h , which implies the bound on the efficiency of H^h . ■

Acknowledgements

Ashrujit Ghoshal's work was partially supported by NSF grants CNS-2026774, CNS-2154174, a JP Morgan Faculty Award, a CISCO Faculty Award, and a gift

from Microsoft. Part of Ashrujit Ghoshal’s work was done during an internship at NTT Research. Cody Freitag is supported in part by the National Science Foundation Graduate Research Fellowship under Grant No. DGE–2139899 and DARPA Award HR00110C0086. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the Defense Advanced Research Projects Agency (DARPA). Ilan Komargodski is the incumbent of the Harry & Abe Sherman Senior Lectureship at the School of Computer Science and Engineering at the Hebrew University, supported in part by an Alon Young Faculty Fellowship, by a grant from the Israel Science Foundation (ISF Grant No. 1774/20), and by a grant from the US-Israel Binational Science Foundation and the US National Science Foundation (BSF-NSF Grant No. 2020643).

References

1. Akshima, Cash, D., Drucker, A., Wee, H.: Time-space tradeoffs and short collisions in merkle-damgård hash functions. In: *Advances in Cryptology - CRYPTO*. pp. 157–186 (2020). https://doi.org/10.1007/978-3-030-56784-2_6
2. Akshima, Guo, S., Liu, Q.: Time-space lower bounds for finding collisions in Merkle-Damgård hash functions. *CRYPTO* (2022). https://doi.org/10.1007/978-3-031-15982-4_7
3. Aumasson, J.P., Henzen, L., Meier, W., Phan, R.C.W.: SHA-3 proposal BLAKE. Submission to NIST **92** (2008)
4. Aumasson, J., Meier, W., Phan, R.C., Henzen, L.: The Hash Function BLAKE. *Information Security and Cryptography* (2014). <https://doi.org/10.1007/978-3-662-44757-4>
5. Aviram, N., Schinzel, S., Somorovsky, J., Heninger, N., Dankel, M., Steube, J., Valenta, L., Adrian, D., Halderman, J.A., Dukhovni, V., Käsper, E., Cohnen, S., Engels, S., Paar, C., Shavitt, Y.: DROWN: breaking TLS using sslv2. In: *USENIX*. pp. 689–706 (2016). <https://doi.org/10.5555/3241094.3241148>
6. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: *CCS*. pp. 62–73 (1993). <https://doi.org/10.1145/168588.168596>
7. Bertoni, G., Daemen, J., Peeters, M., Assche, G.V.: On the indistinguishability of the sponge construction. In: *Advances in Cryptology - EUROCRYPT*. pp. 181–197 (2008). https://doi.org/10.1007/978-3-540-78967-3_11
8. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sponge functions. In: *ECRYPT hash workshop*. vol. 2007. Citeseer (2007)
9. Chung, K., Guo, S., Liu, Q., Qian, L.: Tight quantum time-space tradeoffs for function inversion. In: *FOCS*. pp. 673–684 (2020). <https://doi.org/10.1109/FOCS46700.2020.00068>
10. Coretti, S., Dodis, Y., Guo, S.: Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In: *Advances in Cryptology - CRYPTO*. pp. 693–721 (2018). https://doi.org/10.1007/978-3-319-96884-1_23
11. Coretti, S., Dodis, Y., Guo, S., Steinberger, J.P.: Random oracles and non-uniformity. In: *Advances in Cryptology - EUROCRYPT*. pp. 227–258 (2018). https://doi.org/10.1007/978-3-319-78381-9_9

12. Coron, J., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-damgård revisited: How to construct a hash function. In: CRYPTO. pp. 430–448 (2005). https://doi.org/10.1007/11535218_26
13. Damgård, I.: A design principle for hash functions. In: Advances in Cryptology - CRYPTO. pp. 416–427 (1989). https://doi.org/10.1007/0-387-34805-0_39
14. Dodis, Y., Guo, S., Katz, J.: Fixing cracks in the concrete: Random oracles with auxiliary input, revisited. In: Advances in Cryptology - EUROCRYPT. pp. 473–495 (2017). https://doi.org/10.1007/978-3-319-56614-6_16
15. Fiat, A., Naor, M.: Rigorous time/space trade-offs for inverting functions. SIAM J. Comput. **29**(3), 790–803 (1999). <https://doi.org/10.1137/S0097539795280512>
16. Freitag, C., Ghoshal, A., Komargodski, I.: Time-space tradeoffs for sponge hashing: Attacks and limitations for short collisions. In: Advances in Cryptology - CRYPTO (2022). https://doi.org/10.1007/978-3-031-15982-4_5
17. Ghoshal, A., Komargodski, I.: On time-space tradeoffs for bounded-length collisions in Merkle-Damgård hashing. In: Advances in Cryptology - CRYPTO (2022). https://doi.org/10.1007/978-3-031-15982-4_6
18. Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. SIAM J. Comput. **25**(1), 169–192 (1996). <https://doi.org/10.1137/S0097539791220688>
19. Goldwasser, S., Bellare, M.: Lecture notes on cryptography (2008), <https://cseweb.ucsd.edu/~mihir/papers/gb.pdf>
20. Hellman, M.E.: A cryptanalytic time-memory trade-off. IEEE Trans. Inf. Theory **26**(4), 401–406 (1980). <https://doi.org/10.1109/TIT.1980.1056220>
21. Impagliazzo, R., Kabanets, V.: Constructive proofs of concentration bounds. In: APPROX, pp. 617–631 (2010). https://doi.org/10.1007/978-3-642-15369-3_46
22. Maurer, U.M., Renner, R., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In: TCC (2004). https://doi.org/10.1007/978-3-540-24638-1_2
23. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Advances in Cryptology - CRYPTO. pp. 369–378 (1987). https://doi.org/10.1007/3-540-48184-2_32
24. Merkle, R.C.: A certified digital signature. In: Advances in Cryptology - CRYPTO. pp. 218–238 (1989). https://doi.org/10.1007/0-387-34805-0_21
25. Oechslin, P.: Making a faster cryptanalytic time-memory trade-off. In: Advances in Cryptology - CRYPTO. pp. 617–630 (2003). https://doi.org/10.1007/978-3-540-45146-4_36
26. Ristenpart, T., Shacham, H., Shrimpton, T.: Careful with composition: Limitations of the indifferentiability framework. In: EUROCRYPT. pp. 487–506 (2011). https://doi.org/10.1007/978-3-642-20465-4_27
27. Unruh, D.: Random oracles and auxiliary input. In: Advances in Cryptology - CRYPTO. pp. 205–223 (2007). https://doi.org/10.1007/978-3-540-74143-5_12
28. Yao, A.C.: Coherent functions and program checkers (extended abstract). In: STOC. pp. 84–94 (1990). <https://doi.org/10.1145/100216.100226>