

Better Than Advertised: Improved Collision-Resistance Guarantees for MD-Based Hash Functions

Mihir Bellare*
University of California San Diego
La Jolla, California
mihir@eng.ucsd.edu

Joseph Jaeger†
University of California San Diego
La Jolla, California
jsjaeger@eng.ucsd.edu

Julia Len
University of California San Diego
La Jolla, California
jlen@ucsd.edu

ABSTRACT

The MD transform that underlies the MD and SHA families iterates a compression function h to get a hash function H . The question we ask is, what property X of h guarantees collision resistance (CR) of H ? The classical answer is that X itself be CR. We show that weaker conditions X , in particular forms of what we call constrained-CR, suffice. This reduces demands on compression functions, to the benefit of security, and also, forensically, explains why collision-finding attacks on compression functions have not, historically, lead to immediate breaks of the corresponding hash functions. We obtain our results via a definitional framework called RS security, and a parameterized treatment of MD, that also serve to unify prior work and variants of the transform.

1 INTRODUCTION

The so-called MD transform [15, 24] iterates a compression function h to get a hash function H . The question we ask is, what property X of h guarantees collision resistance (CR) of H ? The classical answer is that X itself be CR [15, 24]. We show that weaker conditions X , in particular forms of what we call constrained-CR, suffice.

The benefit is that if we ask less of compression functions (as we can now do), they are less likely to disappoint. Put another way, our result lowers the bar for the compression function designer, and raises it for the compression function attacker. It also explains an historical cryptanalytic phenomenon, namely that collision-finding attacks on compression functions [16, 30] have not immediately led to breaks of the corresponding hash functions. (Our explanation is that the attacks on the compression functions did not break constrained collision resistance.) In this (second) light, our work formalizes existing cryptanalytic intuition.

We obtain our results via a broader treatment that also serves to unify prior work and different variants of the transform, and to formalize folklore. It involves (1) a definitional framework called RS security that allows us to formulate both classical and new

security goals in a unified way, and (2) a modular treatment of MD that parameterizes it via a splitting function and a compression function.

The MD transform was used in the MD series of hash functions (MD4 [27] and MD5 [26]) and now underlies the most widely used hash functions in practice, namely the SHA series (SHA-1, SHA-256, SHA-512) [25]. An improved understanding of its security, as we provide, is thus of both historical and current interest.

MD FRAMEWORK. We formulate MD in a general, parameterized way, as a transform taking (1) a compression function $h : \{0, 1\}^k \times (\{0, 1\}^\mu \times \{0, 1\}^\sigma) \rightarrow \{0, 1\}^\sigma$, (2) a splitting function $\text{Split} : D \rightarrow (\{0, 1\}^\mu)^*$ and (3) a set $S \subseteq \{0, 1\}^\sigma$ of starting points (also called initial vectors) to return a hash function $H = \text{MD}[h, \text{Split}, S] : D \rightarrow \{0, 1\}^\sigma$. The compression function takes a key k and an input $x = (m, c)$ consisting of a message block m and chaining variable c , and returns output $c' = h_k(m, c)$. The domain D is intended to be large, usually the set of all strings of length up to some big maximum length. The key (k, s) for H consists of a random key k for h and a random starting point s from S . The splitting function breaks the input M to $H_{(k,s)}$ into a sequence $\mathbf{m} = \mathbf{m}[1] \dots \mathbf{m}[n]$ of μ -bit blocks. To compute $H_{(k,s)}(M)$, set $c[1] \leftarrow s$, iterate the compression function via

For $i = 1, \dots, n$ do $c[i + 1] \leftarrow h_k(\mathbf{m}[i], c[i])$,

and return $c[n + 1]$ as the value of $H_{(k,s)}(M)$.

CHARACTERIZING CR PRESERVATION. We start by revisiting the classical question of showing that H is CR assuming h is CR ($X = \text{CR}$). Several works have noted that suffix-freeness of Split is sufficient for this purpose [1, 5, 17, 18]. (Some of these attribute the result to [15, 24], but neither paper appears to actually contain such a claim.) For completeness, we give, in our setting, a formal claim (suffix-freeness of Split plus CR of h implies CR of H , Theorem 5.3) together with the (easy) proof. We then complement this with a novel result: we show that the sufficient condition of suffix-freeness on Split is also *necessary*. We do this by showing that given *any* Split that is not suffix-free, we can construct a compression function h and set S such that (1) h is CR but (2) $H = \text{MD}[h, \text{Split}, S]$ is *not* CR. This fully characterizes MD for the (classical) case where the assumption X made on h is CR.

UNIFYING VARIANTS. Papers, textbooks and standards present variants of the MD transform that differ in details. We can capture them as special cases, corresponding to different choices of splitting function Split and set S of starting points. Together with our above-mentioned characterization, this unifies prior work.

To elaborate, a basic version of MD, from Merkle [23], MOV [22] and KL [20], corresponds to the splitting function that pads the

*Supported in part by NSF grants CNS-1526801 and CNS-1717640, ERC Project ERC FP7/615074 and a gift from Microsoft corporation.

†Supported in part by grants of first author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '17, October 30–November 3, 2017, Dallas, TX, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00

<https://doi.org/10.1145/3133956.3134087>

message M to a multiple of the block length μ and appends a block encoding the length of M . Stinson’s [31] version corresponds to the last block encoding the amount of padding rather than the message length. Damgård’s version [15] starts each block of the padded message with a 1 bit except the first, which it starts with a 0 bit, and also appends a block encoding the amount of padding. The SHA functions [25] use yet another variant where the (padded) message may spill into the last block so that the latter does not encode just the message length (cf. Fig. 3). In papers and textbooks the starting point s is usually 0^σ [15, 20, 22, 23, 31], but, in the SHA series, s differs across hash functions. (For example, for SHA-256 it is the first 32 bits of the square roots of each of the first 8 primes.) All these MD variants can be captured in our framework as making particular choices of suffix-free splitting functions Split and (singleton) spaces S .

CR IS NOT NECESSARY. We would like to show CR of $H = \text{MD}[h, \text{Split}, S]$ under an assumption X on the compression function h that is weaker than CR. We first ask, is this even possible? Or, is CR of h necessary for CR of H ? We show in Section 6 that CR of h is *not* necessary. Given a suffix-free Split, we build a compression function h and set S such that (1) h is *not* CR, yet (2) $H = \text{MD}[h, \text{Split}, S]$ is CR. This opens the door to proving CR of H under relaxed assumptions on h .

RS SECURITY. But what would these assumptions be? Towards finding and formulating them, we step back to give a framework to define security goals for h . Security is parameterized by a relation R and a set S . The game gives the adversary $\mathcal{A}$ a random key k for h and a random point s in S . It returns an object denoted *out*, and wins if R , given k, s, out , returns true. Its RS-advantage is the probability that it wins. Classical $X=\text{CR}$ is captured by viewing *out* as a pair of strings that R checks are a collision under h_k , with s not being involved, formally $S = \{\varepsilon\}$. A form of pre-image resistance that we will use is captured by having R check that *out* gives a pre-image of $s \in S = \{0, 1\}^\sigma$ under h_k . We can also capture constrained forms of collision resistance (ccr), so called because extra requirements are made on the collision, thereby constraining it. In particular, we define R_{ccrS} security. Here winning requires that *out* contains, not only a collision $(m_1, c_1), (m_2, c_2)$ for h_k , but also, for both $j = 1, 2$, if $c_j \neq s$, a further pre-image of it under h_k . Providing the auxiliary information in addition to a collision makes the adversary’s job harder, so $X=R_{\text{ccrS}}$ security is a weaker assumption on h than CR. We can define other relaxations of CR as well.

CR FROM CCR. Theorem 6.4 relaxes the CR assumption, made on the compression function h in Theorem 5.3, to R_{ccrS} security. That is, we show that if Split is suffix free and h is R_{ccrS} secure, then hash function $H = \text{MD}[h, \text{Split}, S]$ is CR secure. The first consequence of this is that the bar is lowered for the compression function designer (their design only needs to provide R_{ccrS} security, which is easier than providing CR) and raised for the cryptanalyst (their attack needs to violate R_{ccrS} security, which is harder than violating CR). We now discuss another consequence, namely to (possibly) better understand some cryptanalytic history.

Already in 1996, Dobbertin had found collisions for the compression function md5 of MD5 [16]. This did not, however, yield collisions on MD5 itself. This, to us, was an indication that MD

was “better than advertised:” it was (possibly) able to promote a non-CR compression function to a CR hash function. Our work is an attempt to capture this intuition formally. Now, it is true that in this particular case the hope was not realized, meaning MD5 failed to be CR, as shown by direct attack [33, 34]. What that tells us is that the compression function md5 is even weaker than we thought: not only is it not CR, it is not even R_{ccrS} . In fact, starting from known MD5 collisions, our reduction will construct collisions, and accompanying auxiliary information, to violate R_{ccrS} security of md5. The story repeats with SHA-1, where collisions found for the compression function sha-1 [30] did not immediately yield collisions for SHA-1, but the latter have now been found [29]. Again, it means sha-1 is not even R_{ccrS} . This, in our view, improves our understanding of compression function security.

SPEEDING UP MD. Suppose the message $M = 0^{2\mu-1}$ to be hashed is a bit short of twice the block length μ . A typical suffix-free encoding (for example that of SHA-256) will pad M and append an encoding of the length $|M|$, to result in a 3-block string $\mathbf{m}$, over which the compression function is iterated. The compression function is thus called 3 times. One might hope for better, just 2 calls. More generally, the savings from dropping one compression function call are significant since messages in practice are often short. This leads us to ask why not use a minimal splitting scheme, like just 10^* pad the message, which in our example results in $\mathbf{m} = 0^{2\mu-1}1$ being 2 blocks long, so that MD will use only two calls to the compression function. But this splitting function is not suffix-free, and did we not show that the suffix-freeness assumption on Split is necessary for CR of H ? Yes, but that was when the assumption on h is CR. Hence it is of course also true when the assumption is R_{ccrS} , since that is implied by CR, but in Section 7 we show that mere injectivity of Split (in particular 10^* padding of the message) does guarantee CR of $H = \text{MD}[h, \text{Split}, S]$ under *alternative* assumptions on h , specifically that it is both ccr and pre-image resistant. The assumption seems quite plausible compared to CR so the performance gain and simplicity of the splitting could make this version of MD attractive.

Our result generalizes, formalizes and strengthens folklore understanding suggested by the following quote from [3]: “It was already known that the plain Merkle-Damgård iteration (so without length strengthening) preserves collision resistance provided it is hard to find a pre-image for the initial vector ... the latter is implied by everywhere pre-image resistance ...” Specifically, our notion of pre-image resistance is weaker than everywhere pre-image resistance: we assume only CCR, rather than CR, of the compression function, we allow starting points (initial vectors) chosen from a distribution, our result applies with any injective Split rather than just plain Merkle-Damgård (a particular choice of Split), and we give formal definitions, result statements and proofs.

REDUCTION COMPLEXITY. As indicated above, many prior works have claimed or proved that CR of h implies CR of H , either for particular choices of Split or assuming the latter is suffix free. It is interesting that, with the exception of a work on formal verification [5], not only papers [1, 5, 15, 17, 18, 24], but also textbooks [20, 22, 31], fail to explicitly specify the reduction underlying the proof. This takes attention away from, and makes it difficult to address, the important question of the (computational) complexity (efficiency) of the reduction. Whether in showing CR or CCR of h implies CR of

H, we in contrast are interested in the precise complexity of the reduction. We accordingly give explicit, pseudocode reductions. In the main sections, we give the reductions that emanate naturally from the proof. Then, in Section 9, we revisit the question of complexity to give alternative reductions that are more memory-efficient [4].

DISCUSSION AND RELATED WORK. MD-based hash functions are also used for HMAC [8]. If we contemplate changes in splitting functions, we want to ensure HMAC security is preserved. However, current analyses of HMAC security [6, 19] show that suffix-free, and even injective, splitting functions suffice.

Our focus is on MD as a way to achieve collision resistance. Other works have looked at it for other ends. Use of MD with prefix-free (as opposed to suffix-free) encodings has been shown in [7, 9] to preserve PRF security. Its ability to provide indistinguishability from a random oracle is studied in [18]. More broadly, MD is one of many possible domain extension methods, and some works [2, 10] consider methods that preserve multiple properties.

2 NOTATION AND CONVENTIONS

If $\mathbf{m}$ is a vector then $|\mathbf{m}|$ denotes its length, $\mathbf{m}[i]$ denotes its i -th coordinate and $\mathbf{m}[i..j]$ denotes the vector consisting of coordinates i through j of $\mathbf{m}$. For example if $\mathbf{m} = (010, 11, 10, 111)$ then $|\mathbf{m}| = 4$, $\mathbf{m}[2] = 11$, $\mathbf{m}[2..4] = (11, 10, 111)$. By ε we denote the empty vector, which has length 0. If D is a set, we say that $\mathbf{m}$ is a vector over D if all its components belong to D , and we let D^* denote the set of all finite-length vectors over D . If $\mathbf{m}, \mathbf{y}$ are vectors, their concatenation, denoted $\mathbf{m}||\mathbf{y}$, is the vector $(\mathbf{m}[1], \dots, \mathbf{m}[|\mathbf{m}|], \mathbf{y}[1], \dots, \mathbf{y}[|\mathbf{y}|])$. For example $(01, 11, 1)|| (10, 000) = (01, 11, 1, 10, 000)$.

A string y is identified with a vector over $\{0, 1\}$, so that $|y|$ denotes its length, $y[i]$ denotes its i -th bit and $y[i..j]$ denotes bits i through j of y . For example if $y = 0100$ then $|y| = 4$, $y[2] = 1$ and $y[2..4] = 100$. In this case, ε denotes the empty string, $\{0, 1\}^*$ is the set of all binary strings, $x||y$ denotes the concatenation of strings x, y . For example $010||11 = 01011$. By $\bar{y}$ we denote the bitwise complement of string y . (For example if $y = 010$ then $\bar{y} = 101$.)

By $\mathbb{N} = \{0, 1, 2, \dots\}$ we denote the set of all non-negative integers. For $p \in \mathbb{N}$ with $p \geq 2$, we let $\mathbb{Z}_p = \{0, 1, \dots, p-1\}$ denote the set of integers modulo p . If $x, n \in \mathbb{N}$ satisfy $0 \leq x < 2^n$ then $\langle x \rangle_n$ denotes the encoding of x as a binary string of length (exactly) n . For example $\langle 7 \rangle_4 = 0111$.

If X is a finite non-empty set, we let $x \leftarrow_s X$ denote picking an element of X uniformly at random and assigning it to x . Algorithms may be randomized unless otherwise indicated. Running time and memory usage are worst case. If A is an algorithm, we let $y \leftarrow A(x_1, \dots; r)$ denote running A with random coins r on inputs $x_1, \dots$ and assigning the output to y . We let $y \leftarrow_s A(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We let $[A(x_1, \dots)]$ denote the set of all possible outputs of A when invoked with inputs $x_1, \dots$.

We use the code based game playing framework of [11]. (See Fig. 1 for an example.) By $\Pr[G]$ we denote the probability of the event that the execution of game G results in the game returning true. We adopt the convention that the running time of an adversary refers to the worst-case execution time of the game with the adversary. We adopt the analogous convention for the memory

usage. This means that usually in reductions, adversary time and memory complexity can be roughly maintained.

3 RS SECURITY FRAMEWORK

FUNCTION FAMILIES. A function family $F : F.\text{Keys} \times F.\text{Inp} \rightarrow F.\text{Out}$ is a 2-argument function taking a key fk in the keyspace $F.\text{Keys}$ and an input x in the input space $F.\text{Inp}$ to return an output $F(fk, x)$ in the output space $F.\text{Out}$. For $fk \in F.\text{Keys}$ we let $F_{fk} : F.\text{Inp} \rightarrow F.\text{Out}$ be defined by $F_{fk}(x) = F(fk, x)$ for all $x \in F.\text{Inp}$.

RS SECURITY. Our definition of security for a function family F is parameterized by a relation $R : \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\text{true}, \text{false}\}$ and a set $S \subseteq \{0, 1\}^*$. Different choices of the pair (R, S) allow us to recover classical definitions including collision resistance, and to specify extensions and variants including constrained collision resistance. The formalism considers game $G_F^{\text{RS}}(\mathcal{A})$ of Fig. 1 associated to R, S, F and adversary $\mathcal{A}$. The latter is given the key fk and a challenge point s drawn randomly from S , and returns some output denoted out . It wins (the game returns true) if relation R returns true on inputs fk, s, out . The advantage of $\mathcal{A}$ relative to R, S , also called its RS advantage, is defined as $\text{Adv}_F^{\text{RS}}(\mathcal{A}) = \Pr[G_F^{\text{RS}}(\mathcal{A})]$, the probability that the game returns true.

COLLISION RESISTANCE. Recall that a *collision* for a function f is a pair of distinct points x_1, x_2 in the domain of f such that $f(x_1) = f(x_2)$. Classical collision resistance of a function family F asks that it be hard for an adversary $\mathcal{A}$, given fk , to find a collision x_1, x_2 for the function F_{fk} . In our framework, this is $R_{\text{cr}}S_\varepsilon$ security, where $S_\varepsilon = \{\varepsilon\}$ consists of just the empty string and $R_{\text{cr}}(fk, s, out)$ parses out as a pair, $(x_0, x_1) \leftarrow out$, and returns true iff $F(fk, x_1) = F(fk, x_2)$ and $x_1 \neq x_2$, meaning x_1, x_2 is a collision for F_{fk} . We recover familiar notation for collision resistance by letting game $G_F^{\text{cr}}(\mathcal{A}) = G_F^{R_{\text{cr}}S_\varepsilon}(\mathcal{A})$ and $\text{Adv}_F^{\text{cr}}(\mathcal{A}) = \text{Adv}_F^{R_{\text{cr}}S_\varepsilon}(\mathcal{A})$.

PRE-IMAGE RESISTANCE. This is a form of one-wayness where the adversary, given fk and challenge s , tries to recover a pre-image of s under F_{fk} . Generalizing [28], our formalization is parameterized by the set S from which s is drawn, and is obtained via our RS framework, as follows. Let $R_{\text{pre}}(fk, s, out)$ return true iff $F_{fk}(out) = s$, meaning $out \in F.\text{Inp}$ is a pre-image of s under F_{fk} . Then $R_{\text{pre}}S$ security captures pre-image resistance for challenges drawn from S . We further discuss this notion, and its relation to other types of pre-image resistance, in Section 8.

RESTRICTED COLLISION RESISTANCE. Restricted collision resistance makes the adversary's job harder by asking that the collision x_1, x_2 satisfy some additional condition that will be specified by R . We will describe the particular restriction we are interested in later in Section 6.

4 THE MD TRANSFORM

COMPRESSION FUNCTIONS. Let h be a family of functions with domain $h.\text{Inp} = h.\text{Bl} \times h.\text{Out}$, meaning $h : h.\text{Keys} \times (h.\text{Bl} \times h.\text{Out}) \rightarrow h.\text{Out}$. A point in the domain is a pair (m, c) where c , called the chaining variable, is in the range of h , and m , called a message block, is in the space $h.\text{Bl}$ of message blocks. Such an h is called a compression function. For example the compression function $h = \text{sha256}$ of

Game $G_F^{RS}(\mathcal{A})$
$fk \leftarrow \$ F.Keys ; s \leftarrow \$ S ; out \leftarrow \$ \mathcal{A}(fk, s)$
Return $R(fk, s, out)$

R	out	$R(fk, s, out)$ returns true iff	
R_{cr}	(x_1, x_2)	$x_1 \neq x_2$ and $F_{fk}(x_1) = F_{fk}(x_2)$	Collision resistance
R_{pre}	x	$F_{fk}(x) = s$	Pre-image resistance
R_{ccr}	$((m_1, c_1), (m_2, c_2), (a_1, a_2))$	$R_{cr}(fk, \varepsilon, (m_1, c_1), (m_2, c_2)) \wedge (c_1 \in \{s, F_{fk}(a_1)\}) \wedge (c_2 \in \{s, F_{fk}(a_2)\})$	Constrained CR

Figure 1: Top: Game for defining R-security of function family F. Bottom: Some relations we will use. For R_{cr} we have $s = \varepsilon$. For R_{ccr} , function family $F : F.Keys \times (F.Bl \times F.Out) \rightarrow F.Out$ is a compression function.

$H_{(k,s)}(M)$
$\mathbf{m} \leftarrow \text{Split}(M) ; c \leftarrow s ; n \leftarrow \mathbf{m} $
For $i = 1, \dots, n$ do $c \leftarrow h_k(\mathbf{m}[i], c)$
Return c

Figure 2: Function family $H = \text{MD}[h, \text{Split}, S]$ obtained by applying the MD transform to compression function h , splitting function Split and space S of initial vectors.

SHA256 has $\text{sha256.Bl} = \{0, 1\}^{512}$ and $\text{sha256.Out} = \{0, 1\}^{256}$. Its key space $\text{sha256.Keys} = \{k\}$ is a singleton, where k consists of 64 32-bit strings which are the first 32 bits of the fractional parts of the cube roots of the first 64 primes [25].

SPLITTING FUNCTIONS. Let $\text{Split} : \text{Split.Inp} \rightarrow \text{Split.Bl}^*$ be a function that takes a message $M \in \text{Split.Inp}$ and returns a vector $\mathbf{m} = \text{Split}(M)$ over a set Split.Bl . We require that this function is injective, and there is an inverse $\text{Split}^{-1} : \text{Split.Bl}^* \rightarrow \text{Split.Inp} \cup \{\perp\}$ such that $\text{Split}^{-1}(\mathbf{m}) = M$ if $\mathbf{m} = \text{Split}(M)$ and $\perp$ otherwise. We call Split a splitting function. The domain Split.Inp is expected to be a large set, usually all strings of length up to some very high maximum. In usage, $\text{Split.Bl} = h.Bl$ will be the set of message blocks for a compression function.

THE MD TRANSFORM. Let $h : h.Keys \times (h.Bl \times h.Out) \rightarrow h.Out$ be a compression function. Let $\text{Split} : \text{Split.Inp} \rightarrow h.Bl^*$ be a splitting function whose range Split.Bl^* , as the notation indicates, is $h.Bl^*$. Let $S \subseteq h.Out$ be a set of *starting points*, also called *initial vectors*. The MD transform $\text{MD}[h, \text{Split}, S]$ associates to them the family of functions H that is defined as follows. Let $H.Inp = \text{Split.Inp}$ be the set of messages that are possible inputs to the splitting function. Let $H.Out = h.Out$. Let $H.Keys = h.Keys \times S$, so that a key for H is a pair (k, s) consisting of a key k for the compression function and a particular starting point (initial vector) $s \in S$. Then H is specified in Fig. 2.

SPLITTING IN SHA. Our rendition of MD generalizes prior ones, both from the literature [1, 14, 23] and from standards [25], all of which can be seen as particular choices of Split and S . We illustrate by recovering SHA256 as $\text{MD}[\text{sha256}, \text{SplitSha}_{(\mu, \epsilon)}, \{s\}]$ for choices of the components that we now specify. The compression function $\text{sha256} : \{k\} \times (\{0, 1\}^{512} \times \{0, 1\}^{256}) \rightarrow \{0, 1\}^{256}$ is of course the compression function of SHA256 as per [25], with k the 64 · 32 bit

key discussed above. The starting point s , as specified in [25], is a 256-bit string, viewed as 8 32-bit blocks which are the first 32 bits of the square roots of the first 8 primes. We define $\text{SplitSha}_{(\mu, \epsilon)}$ as the general splitting function for the SHA function families: SHA1, SHA256, SHA512. It is parameterized by μ , the block length, and ϵ , the length of the encoding of the message length. These values are shown for each SHA function in Fig. 4. Specifically for SHA256, $\mu = 512$ and $\epsilon = 64$. To define $\text{SplitSha}_{(\mu, \epsilon)}$, first define function $\text{pad}_{(\mu, \epsilon)}$ to take as input an integer L , with $0 \leq L < 2^\epsilon$, and return $1\|0^\ell\|\langle L \rangle_\epsilon$, where $\langle L \rangle_\epsilon$ is an ϵ -bit encoding of L , and $\ell \geq 0$ is the smallest integer such that $L + \epsilon + 1 + \ell$ is a multiple of μ . Let SplitSha.Inp be the set of all strings of length at most 2^ϵ , and let $\text{SplitSha.Bl} = \{0, 1\}^\mu$. The function $\text{SplitSha} : \text{SplitSha.Inp} \rightarrow \text{SplitSha.Bl}^*$, on input M , lets $L = |M|$ be the length of M , and lets $X = M\|\text{pad}_{(\mu, \epsilon)}(L)$. Note that the length of string X is a multiple of μ . Let $n \leftarrow |X|/\mu$, and let $\mathbf{m}[i] = X[1 + \mu(i - 1)..\mu i]$ be the μ bit-block consisting of bits $1 + \mu(i - 1)$ through μi of X , for $1 \leq i \leq n$. Then $\text{SplitSha}(M)$ returns $\mathbf{m}$, which is a vector over $\{0, 1\}^\mu$.

5 CR PRESERVATION OF MD

Here we recall the classical problem of showing collision resistance of the hash function $H = \text{MD}[h, \text{Split}, S]$ assuming *only* collision resistance of the compression function h . As noted in the introduction, several works have noted that suffix-freeness of Split is sufficient for this purpose [1, 5, 17, 18]. For completeness, we will provide a formal claim together with the (easy) proof in our setting.

We then expand on this knowledge to establish a novel result that the property of being suffix-free is *precisely* the property required for this proof; it is a necessary condition in addition to being sufficient. To demonstrate this, we construct, for any splitting function which is not suffix-free, a compression function that is collision resistant in isolation, but for which the result of applying the MD transform is not collision resistant.

SUFFIX-FREENESS. Let $x, y \in D^*$ be vectors over a set D . We say that x is a suffix of y , written $y \supseteq x$, if there exists a vector $z \in D^*$ such that $y = z\|x$. (The notation $y \supseteq x$ is intended to visualize x being the right-hand side of y .) For example, $(10, 11)$ is a suffix of $(00, 11, 10, 11)$, namely $(00, 11, 10, 11) \supseteq (10, 11)$, by letting $z = (00, 11)$. However, $(01, 11)$ is not a suffix of $(00, 11, 10, 11)$, namely $(00, 11, 10, 11) \not\supseteq (01, 11)$. We say that splitting function $\text{Split} : \text{Split.Inp} \rightarrow \text{Split.Bl}^*$ is *suffix-free* if for any two distinct

$\text{SplitSha}_{(\mu, \epsilon)}(M)$ $L \leftarrow M $ $X \leftarrow M \parallel \text{pad}_{(\mu, \epsilon)}(L); n \leftarrow X /\mu$ For $1 \leq i \leq n$ do $\mathbf{m}[i] \leftarrow X[1 + \mu(i-1) \dots \mu i]$ Return $\mathbf{m}$
$\text{pad}_{(\mu, \epsilon)}(L)$ $\ell \leftarrow (\mu - \epsilon - 1 - L) \bmod \mu$ Return $1 \parallel 0^\ell \parallel \langle L \rangle_\epsilon$

Figure 3: SplitSha and pad, the splitting function and padding function, respectively, of the SHA function families. They are parameterized by μ , the block length, and ϵ , the length of the encoding of the message length.

Function	μ	σ	ϵ
SHA1	512	160	64
SHA256	512	256	64
SHA512	1024	512	128

Figure 4: Choices of parameters across different hash functions.

messages $M_1, M_2 \in \text{Split.Inp}$ we have $\text{Split}(M_1) \not\supseteq \text{Split}(M_2)$, that is, $\text{Split}(M_2)$ is not a suffix of $\text{Split}(M_1)$.

SUFFIX-FREENESS OF SplitSha. We discussed above how SHA256 is underlain by a particular splitting function that we defined and called $\text{SplitSha}_{(\mu, \epsilon)}$. Here we show that this function is suffix-free to provide an example of a suffix-free scheme.

PROPOSITION 5.1. *The function $\text{SplitSha}_{(\mu, \epsilon)}$ is suffix-free.*

PROOF. (of Proposition 5.1) Let $M_1, M_2 \in \text{SplitSha}_{(\mu, \epsilon)}. \text{Inp}$ be distinct. Consider when $|M_1| \neq |M_2|$. Then the last blocks of vectors $\text{SplitSha}_{(\mu, \epsilon)}(M_1)$ and $\text{SplitSha}_{(\mu, \epsilon)}(M_2)$ are, respectively, $\text{pad}_{(\mu, \epsilon)}(|M_1|)$ and $\text{pad}_{(\mu, \epsilon)}(|M_2|)$. But then $\langle |M_1| \rangle_\epsilon \neq \langle |M_2| \rangle_\epsilon$, which implies $\text{pad}_{(\mu, \epsilon)}(|M_1|) \neq \text{pad}_{(\mu, \epsilon)}(|M_2|)$ and so neither vector can be a suffix of the other.

We now consider the case when $|M_1| = |M_2|$. This will imply that $|\text{SplitSha}_{(\mu, \epsilon)}(M_1)| = |\text{SplitSha}_{(\mu, \epsilon)}(M_2)|$. Then in order for $\text{SplitSha}_{(\mu, \epsilon)}(M_1) \supseteq \text{SplitSha}_{(\mu, \epsilon)}(M_2)$ or the opposite to hold it must be that $\text{SplitSha}_{(\mu, \epsilon)}(M_1) = \text{SplitSha}_{(\mu, \epsilon)}(M_2)$. Notice that $\text{SplitSha}_{(\mu, \epsilon)}(M)$ prepends the message M to its output. Since $M_1 \neq M_2$, we have that $\text{SplitSha}_{(\mu, \epsilon)}(M_1) \neq \text{SplitSha}_{(\mu, \epsilon)}(M_2)$ and so $\text{SplitSha}_{(\mu, \epsilon)}(M_1) \not\supseteq \text{SplitSha}_{(\mu, \epsilon)}(M_2)$ as required. $\square$

THE STRUCTURE OF MD COLLISIONS. We now proceed to a simple lemma about the structure of collisions in the MD transform. This will be used for our proof that the MD transform preserves collision resistance when the splitting function is suffix-free. This lemma argues the correctness of an algorithm B_{cr} to formalize the observation that if M_1, M_2 form a collision for the MD transform with a suffix-free splitting function, then by examining the computation

of the hash function on these inputs we can easily find a collision for the underlying compression function.

LEMMA 5.2. *Let h be a compression function, let Split be a splitting function with $\text{Split.Bl} = h.\text{Bl}$ and let $S \subseteq h.\text{Out}$ be a set of possible starting points. Let $H = \text{MD}[h, \text{Split}, S]$ be the hash function associated to these components via the MD transform of Fig. 2. Let $k \in h.\text{Keys}, s \in S$. Suppose $M_1, M_2 \in \text{Split.Inp}$ are a pair of distinct messages satisfying (1) $\text{Split}(M_1) \not\supseteq \text{Split}(M_2)$ and $\text{Split}(M_2) \not\supseteq \text{Split}(M_1)$ and (2) M_1, M_2 are a collision for $H_{(k, s)}$. Then, on inputs $(k, s), M_1, M_2$, algorithm B_{cr} of Fig. 5 returns $(x_1, c_1), (x_2, c_2)$ that form a collision for h_k .*

PROOF. (of Lemma 5.2) From algorithm B_{cr} , let $\mathbf{m}_1 = \text{Split}(M_1)$, $\mathbf{m}_2 = \text{Split}(M_2)$, $n_1 = |\mathbf{m}_1|$, and $n_2 = |\mathbf{m}_2|$. First B_{cr} computes the vectors of chaining variables, c_1 and c_2 as shown in the pseudocode. Assume (without loss of generality) that $n_1 \geq n_2$, i.e. that $|\mathbf{m}_1| \geq |\mathbf{m}_2|$. Since M_1, M_2 are a collision for $H_{(k, s)}$, we have $H_{(k, s)}(M_1) = H_{(k, s)}(M_2)$. Because $\mathbf{m}_1 \not\supseteq \mathbf{m}_2$ and $\mathbf{m}_2 \not\supseteq \mathbf{m}_1$, there must exist $i \in \{0, \dots, n_2 - 1\}$ such that $(\mathbf{m}_1[n_1 - i], c_1[n_1 - i]) \neq (\mathbf{m}_2[n_2 - i], c_2[n_2 - i])$. Let j represent the minimal such value. Then it will hold that $c_1[n_1 - j + 1] = c_2[n_2 - j + 1]$. Thus, the pair $(\mathbf{m}_1[n_1 - j], c_1[n_1 - j]), (\mathbf{m}_2[n_2 - j], c_2[n_2 - j])$, return by B_{cr} will form a collision for h_k . $\square$

SUFFIX-FREENESS PRESERVES CR. Finally, we show a reduction from the collision resistance of the hash function $H = \text{MD}[h, \text{Split}, S]$ to the collision resistance of the compression function h when using a suffix-free splitting function Split.

THEOREM 5.3. *Let h be a compression function, let Split be a suffix-free splitting function with $\text{Split.Bl} = h.\text{Bl}$ and let $S \subseteq h.\text{Out}$ be a set of possible starting points. Let $H = \text{MD}[h, \text{Split}, S]$ be the hash function associated to these components via the MD transform of Fig. 2. Given an adversary $\mathcal{A}_H$, let $\mathcal{A}_h$ be the adversary of Fig. 5 using B_{cr} .*

Then

$$\text{Adv}_H^{\text{cr}}(\mathcal{A}_H) \leq \text{Adv}_h^{\text{cr}}(\mathcal{A}_h). \quad (1)$$

The time complexity of $\mathcal{A}_h$ is the sum of the time complexities of $\mathcal{A}_H$ and B_{cr} . The memory complexity of $\mathcal{A}_h$ is the maximum of the memory complexity of $\mathcal{A}_H$ and the memory complexity of B_{cr} .

In Section 9, we revisit this and other reductions to give alternative reductions that are more memory-efficient [4].

PROOF. (of Theorem 5.3) It is clear that the time and memory complexity of adversary $\mathcal{A}_h$ are as stated in the theorem.

Let $k \in h.\text{Keys}, s \in S$ be the values sampled when $\mathcal{A}_h$ is executed and $M_1, M_2 \in \text{Split.Inp}$ be the values returned by $\mathcal{A}_h$. Suppose they form a collision for $H_{(k, s)}$.

Then we have $\text{Split}(M_1) \not\supseteq \text{Split}(M_2)$ and $\text{Split}(M_2) \not\supseteq \text{Split}(M_1)$ because Split is suffix-free, so they fulfill the conditions of Lemma 5.2 and is B_{cr} guaranteed to return a collision for h_k . As an immediate result Equation (1) holds, completing the proof. $\square$

NECESSITY. We can now complete the picture for splitting functions when assuming the compression function is collision resistant by showing that the suffix-free restriction is *precisely* the correct restriction on the splitting function. In particular, we will establish

Adversary $\mathcal{A}_h(k, \varepsilon)$ $s \leftarrow S$; $(M_1, M_2) \leftarrow \mathcal{A}_H((k, s), \varepsilon)$ Return $B_{cr}((k, s), M_1, M_2)$
Algorithm $B_{cr}((k, s), M_1, M_2)$ $\mathbf{m}_1 \leftarrow \text{Split}(M_1)$; $\mathbf{m}_2 \leftarrow \text{Split}(M_2)$; $n_1 \leftarrow \mathbf{m}_1$; $n_2 \leftarrow \mathbf{m}_2$ $c_1[1] \leftarrow s$; $c_2[1] \leftarrow s$ For $i = 1, \dots, n_1$ do $c_1[i+1] \leftarrow h_k((\mathbf{m}_1[i], c_1[i]))$ For $i = 1, \dots, n_2$ do $c_2[i+1] \leftarrow h_k((\mathbf{m}_2[i], c_2[i]))$ $n \leftarrow \min(n_1, n_2)$ For $i = 0, \dots, n-1$ do $(x_1, c_1) \leftarrow (\mathbf{m}_1[n_1 - i], c_1[n_1 - i])$ $(x_2, c_2) \leftarrow (\mathbf{m}_2[n_2 - i], c_2[n_2 - i])$ If $(x_1, c_1) \neq (x_2, c_2)$ then return $((x_1, c_1), (x_2, c_2))$ Return $\perp$

Figure 5: Adversary $\mathcal{A}_h$ for the proof of Theorem 5.3 and algorithm B_{cr} for Lemma 5.2.

that Split being suffix-free is a necessary condition for proving that $\text{MD}[h, \text{Split}, S]$ is secure under the assumption that h is collision resistant, in addition to a sufficient one.

Given an arbitrary splitting function Split and a pair of inputs M_1, M_2 such that $\text{Split}(M_1) \supseteq \text{Split}(M_2)$, we construct a compression function h which is collision resistant (from another function which we assume to be collision resistant), but for which the pair M_1, M_2 is a collision for $\text{MD}[h, \text{Split}, S]$ (with high probability over the choice of $s \in S$ in the case that $S = h.\text{Out}$).

For simplicity, we will first consider the simpler case when S consists of a single, fixed value s on which our choice of compression function can depend. For this case, we can directly construct the compression function so that when chained with the starting value s on the blocks contained uniquely in $\text{Split}(M_1)$ but not those in $\text{Split}(M_2)$, it "loops" back to s . We then extend this technique to cover the case when S is some larger set from which s is sampled randomly.

It will be convenient to describe our results in terms of the MD transform applied to messages that have already been split into blocks. For any set Bl we let $l : \text{Bl}^* \rightarrow \text{Bl}^*$ be the splitting function which simply outputs its input unchanged. For some compression function h and set S , let $H^l = \text{MD}[h, l, S]$. We will informally say that h loops on $(s, \mathbf{u})$ if $H^l_{(k,s)}(\mathbf{u}) = s$ for all $k \in h.\text{Keys}$. The following lemma observes that if h loops on $(s, \mathbf{u})$ and $\text{Split}(M_1) = \mathbf{u} \parallel \text{Split}(M_2)$, then the pair M_1, M_2 is a collision for $\text{MD}[h, \text{Split}, S]$.

LEMMA 5.4. *Let $h : h.\text{Keys} \times (h.\text{Bl} \times h.\text{Out}) \rightarrow h.\text{Out}$ be a compression function, $\text{Split} : \text{Split.Inp} \rightarrow h.\text{Bl}^*$ be a splitting function, and $S \subseteq h.\text{Out}$ be a set of starting points. Let $H = \text{MD}[h, \text{Split}, S]$ be the hash function associated to these components via the MD transform of Fig. 2. Let l be the splitting function described above and $H^l = \text{MD}[h, l, S]$ be the corresponding hash function obtained via the MD transform. Suppose $M_1, M_2 \in \text{Split.Inp}$ are a pair of distinct messages satisfying $\text{Split}(M_1) \supseteq \text{Split}(M_2)$. Let $\mathbf{u}$ be the vector for which $\text{Split}(M_1) = \mathbf{u} \parallel \text{Split}(M_2)$. For any choice of $(k, s) \in H.\text{Keys}$, if $H^l_{(k,s)}(\mathbf{u}) = s$ then M_1, M_2 is a collision for $H_{(k,s)}$.*

PROOF. (of Lemma 5.4) First note that for any vectors $\mathbf{m}, \mathbf{y} \in \text{Bl}^*$, $H^l_{(k,s)}(\mathbf{m} \parallel \mathbf{y}) = H^l_{(k,s')}(\mathbf{y})$ where $s' = H^l_{(k,s)}(\mathbf{m})$. This is a simple observation from the code of the MD transform shown in Fig. 2. The chaining variable $c[|\mathbf{m}| + 1]$ obtained during the computation of $H^l_{(k,s)}(\mathbf{m} \parallel \mathbf{y})$ would be the output of $H^l_{(k,s')}(\mathbf{m})$. The rest of the computation then exactly mirrors H^l applied to $\mathbf{y}$ with $c[|\mathbf{m}| + 1]$ serving the role of the starting point.

Using this observation, the proof is straightforward. We can rewrite H on input M_1 as follows

$$\begin{aligned} H_{(k,s)}(M_1) &= H^l_{(k,s)}(\mathbf{u} \parallel \text{Split}(M_2)) \\ &= H^l_{(k,s')}(\text{Split}(M_2)) \\ &= H_{(k,s')}(M_2) \end{aligned}$$

where $s' = H^l_{(k,s)}(\mathbf{u})$. From our assumption, this equals s . Thus $H_{(k,s)}(M_1) = H_{(k,s)}(M_2)$, as desired. $\square$

First we will handle the case when the s used for the MD transform is an a priori fixed value.

PROPOSITION 5.5. *Let Split be a splitting function and $M_1, M_2 \in \text{Split.Inp}$ satisfy $\text{Split}(M_1) \supseteq \text{Split}(M_2)$. Let $\mathbf{u}$ be the vector for which it holds that $\text{Split}(M_1) = \mathbf{u} \parallel \text{Split}(M_2)$. Let $b \in \mathbb{N}$ and $a = b + |\mathbf{u}|$. Let f be a family of functions with $f.\text{Inp} = \text{Split.Bl} \times \mathbb{Z}_a$ and $f.\text{Out} = \mathbb{Z}_b$.*

Then we can build a compression function g^u (shown in Fig. 6, with $g.\text{Inp} = f.\text{Inp}$ and $g.\text{Out} = \mathbb{Z}_a$) such that for all adversaries $\mathcal{A}$, $\text{Adv}_{g^u}^{\text{cr}}(\mathcal{A}) \leq \text{Adv}_f^{\text{cr}}(\mathcal{A})$. Furthermore, letting $G = \text{MD}[g^u, \text{Split}, \{0\}]$, we can build an efficient adversary $\mathcal{B}$ (shown in Fig. 6) such that

$$\text{Adv}_G^{\text{cr}}(\mathcal{B}) = 1.$$

The compression function g^u above is specifically defined in a contrived way so that it loops on $(0, \mathbf{u})$ and thus M_1, M_2 is a collision for the MD transform.

In the above we fixed the starting point s to 0 and the set of chaining variables $g^u.\text{Out}$ to $\mathbb{Z}_a$. This is without loss of generality because the lemma can easily be extended to any reasonable choice of $g^u.\text{Out}$ and fixed $s \in g^u.\text{Out}$ by choosing an efficiently computable and invertible mapping $e(\cdot) : g^u.\text{Out} \rightarrow \mathbb{Z}_{|g^u.\text{Out}|}$ which satisfies $e(s) = 0$.

PROOF. (of Proposition 5.5) We will first show that any collision for g^u is also a collision for f by proving that if $g_k^u((m, c)) = g_k^u((m', c'))$ it either holds that $(m, c) = (m', c')$ or that $f_k((m, c)) = f_k((m', c'))$. As such, suppose $g_k^u((m, c)) = g_k^u((m', c'))$.

Note that the first return statement of g^u always outputs a value less than $|\mathbf{u}|$ while the second always outputs a value greater than $|\mathbf{u}|$. We can consider these two cases separately.

Let us first suppose that $g_k^u((m, c)) < |\mathbf{u}|$. This then means that $c + 1 = c' + 1 \pmod{|\mathbf{u}|}$, so c and c' must be the same (because the condition of the if statement guarantees that both are less than $|\mathbf{u}|$). The if statement inside g^u must evaluate to true for both pairs, so we have $m = \mathbf{u}[c + 1] = \mathbf{u}[c' + 1] = m'$ and so $(m, c) = (m', c')$.

Now consider the other case, that $g_k^u((m, c)) \geq |\mathbf{u}|$. Then this must mean that $f_k((m, c)) + |\mathbf{u}|$ and $f_k((m', c')) + |\mathbf{u}|$ are the same and so $f_k((m, c)) = f_k((m', c'))$.

Because any collision for g^u is also a collision for f , for any adversary $\mathcal{A}$ it must hold that $\text{Adv}_{g^u}^{\text{cr}}(\mathcal{A}) \leq \text{Adv}_f^{\text{cr}}(\mathcal{A})$.

$g_k^u((m, c))$
If $(c < u)$ and $(m = u[c + 1])$
Return $c + 1 \bmod u $
Return $f_k((m, c)) + u $
Adversary $\mathcal{B}(k, s)$
Return (M_1, M_2)
$h_k^u((m, c))$
$(b_c, q_c, r_c) \leftarrow c$
If $(b_c = 0)$ and $(m = u[r_c + 1])$
Return $(0, q_c, (r_c + 1 \bmod u))$
$(q, r) \leftarrow f_k((m, c))$
Return $(1, q, r)$

Figure 6: Compression functions and adversary used for Proposition 5.5 and Theorem 5.6.

To prove that $\mathcal{B}$ has advantage 1, we will make use of Lemma 5.4 by showing that $H_{(k,0)}^1(u) = 0$ where H^1 is defined as in the lemma. Let c be the vector of values that would be obtained in the computation of $H_{(k,0)}^1(u)$; that is, let $c[1] = 0$ and $c[i + 1] = g_k^u((u[i], c[i]))$ for $i = 1, \dots, |u|$.

Following the code of g^u we can then see that its if statement will always evaluate to true in this computation and so $c[i + 1] = c[i] + 1 \pmod{|u|}$ holds for all i . Consequently, $c[i] = i$ for $i = 1, \dots, |u|$ and then $c[|u| + 1] = 0$. The latter is the value returned by H^1 so $H_{(k,0)}^1(u) = 0$ and the pair M_1, M_2 is a collision for H . It follows that the advantage of $\mathcal{B}$ is exactly 1. $\square$

The lemma above might seem somewhat contrived, because we allowed our compression function g to depend on the starting point used for the MD transform. This makes it, in some senses, a weak result and one might naturally wonder this dependency was necessary for the result. It is not. At the cost of some added complexity and lost success probability for $\mathcal{B}$ we can extend this to the case when s is randomly chosen from some set instead of fixed.

THEOREM 5.6. *Let $a \in \mathbb{N}$. Let Split be a splitting function. Suppose $\text{Split}(M_1) \supseteq \text{Split}(M_2)$ and in particular $\text{Split}(M_1) = u \parallel \text{Split}(M_1)$. Let f be a family of functions with $f.\text{Inp} = \text{Split}.Bl \times (\mathbb{Z}_2 \times \mathbb{Z}_a \times \mathbb{Z}_{|u|})$ and $f.\text{Out} = \mathbb{Z}_a \times \mathbb{Z}_{|u|}$.*

Now let $h^u : \text{Split}.Bl \times (\mathbb{Z}_2 \times \mathbb{Z}_a \times \mathbb{Z}_{|u|}) \rightarrow \mathbb{Z}_2 \times \mathbb{Z}_a \times \mathbb{Z}_{|u|}$ be the compressions function shown in Fig. 6. For all adversary $\mathcal{A}$, it holds that $\text{Adv}_{h^u}^{\text{cr}}(\mathcal{A}) \leq \text{Adv}_f^{\text{cr}}(\mathcal{A})$. Furthermore, letting $S = h^u.\text{Out}$ and $H = \text{MD}[h^u, \text{Split}, S]$, we can build an efficient adversary $\mathcal{B}$ (shown in Fig. 6) satisfying,

$$\text{Adv}_H^{\text{cr}}(\mathcal{B}) \geq 1/(2|u|).$$

The compression function h^u above is specifically designed so that it loops on s, u for any s of the form $(0, q, 0)$, giving the desired collision between M_1 and M_2 with the specified probability over the random choice of s .

Again, this theorem can be extended to cover any reasonable choice of $h^u.\text{Out}$. One can first map $h^u.\text{Out}$ to $\mathbb{Z}_{|h^u.\text{Out}|}$ as discussed earlier. Then from $c \in \mathbb{Z}_{|h^u.\text{Out}|}$ one can obtain the tuple (b_c, q_c, r_c) via $b_c \leftarrow c \bmod 2$, $y \leftarrow \lfloor c/2 \rfloor$, $r_c \leftarrow y \bmod |u|$, and $q_c \leftarrow \lfloor y/2 \rfloor$.

There are technical details to be considered regarding the fact that 2 and $|u|$ may not be divisors of $|h^u.\text{Out}|$ and that S may not be “nicely” distributed in $\mathbb{Z}_{|h^u.\text{Out}|}$, but for any reasonable choice of $h^u.\text{Out}$ and S this should not be an issue.

For our theorems we have assumed that we were given a pair M_1, M_2 such that $\text{Split}(M_1) \supseteq \text{Split}(M_2)$. It is not difficult to come up with (contrived) splitting functions which are not suffix-free, but for which we believe it is computationally difficult to find such a pair. We chose our formalization that M_1 and M_2 are a priori known because, for specific, prior splitting functions, either they were suffix-free, or it was trivially easy to find M_1 and M_2 violating suffix-freeness. An alternative way to address this would be to make suffix-freeness a computational condition, and then say that, given an adversary returning M_1 and M_2 violating suffix-freeness with high probability, we build our compression function and adversary. (Of course, one might then ask about finding the adversary, analogous to keyless collision resistance, but the philosophical position would at least seem on par with prior ones.)

PROOF. (of Theorem 5.6) The basic structure of this proof closely follows that of the proof for Proposition 5.5. Throughout this proof for a string c we will let b_c, q_c, r_c denote the corresponding values used by h^u on input (m, c) for some c .

To start, we will show that any collision for h^u is also a collision for f by proving that if $h_k^u((m, c)) = h_k^u((m', c'))$ it either holds that $(m, c) = (m', c')$ or that $f_k((m, c)) = f_k((m', c'))$. As such, suppose $h_k^u((m, c)) = h_k^u((m', c'))$.

Note that the first return statement of h^u always outputs a tuple whose first element is 1 while the second always outputs a tuple whose first element is 0. We will consider these two cases separately. Let $y = h_k^u((m, c))$ and $y' = h_k^u((m', c'))$.

Let us first suppose that $y[1] = 0$. This means that $q_c = q_{c'}$ and $r_c + 1 = r_{c'} + 1 \pmod{|u|}$. The if statement in h^u must have evaluated to true on both inputs so we have $b_c = 1 = b_{c'}$ and $m = u[r_c + 1] = u[r_{c'} + 1] = m'$. Putting this all together, we have shown that $(m, c) = (m', c')$.

Now consider the other case when $y[1] = 1$. Then we have that $f_k((m, c)) = f_k((m', c'))$.

Because any collision for h^u is also a collision for f , for any adversary $\mathcal{A}$ it must hold that $\text{Adv}_h^{\text{cr}}(\mathcal{A}) \leq \text{Adv}_f^{\text{cr}}(\mathcal{A})$.

To prove our statement about the advantage of $\mathcal{B}$ we will make use of Lemma 5.4 and bound the probability that $H_{(k,s)}^1(u) = s$ over the random choice of s (where H^1 is defined as in the lemma).

Suppose s is of the form $(0, q, 0)$ for some $q \in \mathbb{Z}_a$. Let c be the vector of values that would have been obtained in the computation of $H_{(k,s)}^1(u)$; that is, let $c[1] = s$ and $c[i + 1] = h_k^u((u[i], c[i]))$ for $i = 1, \dots, |u|$.

Following the code of h^u we can see that the if statement will always return true in this computation, and so $c[i] = (0, q, i)$ for $i = 1, \dots, |u|$ and then $c[|u| + 1] = (0, q, 0) = s$. The latter is the value returned by H^1 so $H_{(k,s)}^1(u) = s$ as desired and the pair M_1, M_2 is a collision for H .

Then the advantage of $\mathcal{B}$ is bounded by the probability that s is of the form $(0, q, 0)$ which is exactly $1/(2|u|)$. $\square$

```

 $h_k^s((m, c))$ 
If  $c[1] = s[1]$  then return  $s[1] || h'_k((m, c[2...|c|]))$ 
Return  $\bar{s}$ 

```

Figure 7: h_k^s for Proposition 6.1, Proposition 6.2, and Proposition 6.3

6 WEAKENING ASSUMPTIONS ON H

In this section we improve on the classic result that the collision resistance of h guarantees that H will be collision resistant. In particular, we will explore the possibility of weakening the assumption made of h and provide a natural, less stringent variation of collision resistance from which we are able to assure the collision resistance of h obtained via the MD transform.

USING A NON-CR h . We have shown that the collision resistance of the compression function h implies the collision resistance of the hash function H obtained by the MD transform. However, the collision resistance of H may not always rely on h being collision resistant. We will show by construction that H can be collision resistant even when h is not.

Let $b, c \in \mathbb{N}$. Given a compression function $h' : h'.\text{Keys} \times (\{0, 1\}^b \times \{0, 1\}^c) \rightarrow \{0, 1\}^c$ and some $s \in \{0, 1\}^{c+1}$, we construct the compression function $h^s : h'.\text{Keys} \times (\{0, 1\}^b \times \{0, 1\}^{c+1}) \rightarrow \{0, 1\}^{c+1}$ shown in Fig. 7. Let Split be a suffix-free splitting function with $\text{Split}.Bl = h.Bl$ and define the set of starting points by $S = \{s\}$. Let $H = \text{MD}[h^s, \text{Split}, S]$ be the hash function associated to these components via the MD transform of Fig. 2. We will think of h' as being a good collision resistant compression function. Then we will show that while h^s is a poor collision resistant compression function, H nonetheless remains a good collision resistant hash function.

The idea motivating our construction of h^s should be clear. Creating a collision for h^s is trivial by making the if statement evaluate to false. However, when h^s is used inside of the MD transform with s as a starting point, this case will never occur.

PROPOSITION 6.1. *Let $M_1 = (0^b, \bar{s})$, $M_2 = (1^b, \bar{s})$, and $\mathcal{B}$ be the adversary shown in Fig. 6. Then $\text{Adv}_{h^s}^{\text{cr}}(\mathcal{B}) = 1$.*

Put simply, the above tells us that h^s is not collision resistant because $\mathcal{B}$ is clearly efficient.

PROOF. (of Proposition 6.1) When we compute $h_k^s(M_1)$, we see that $\bar{s}[1] \neq s[1]$, so $\bar{s}$ is returned. Similarly, $\bar{s}$ is returned when we compute $h_k^s(M_2)$. Notice that $M_1 \neq M_2$ yet $h_k^s(M_1) = h_k^s(M_2)$. Thus, h^s is not collision resistant. $\square$

The following proposition is a useful stepping stone for showing that H is collision resistant if h' is.

PROPOSITION 6.2. *Let $k \in h'.\text{Keys}$. Then for each iteration of h_k^s in the computation of H_k , h_k^s never returns $\bar{s}$.*

PROOF. (of Proposition 6.2) Fix $M \in \text{Split}.Inp$ and let $\mathbf{m}, \mathbf{c}$ be the vectors computed by $H_k(M)$. Suppose, for a contradiction, that for some i from 1 to $|\mathbf{m}|$, $h_k^s(\mathbf{m}[i], \mathbf{c}[i]) = \bar{s}$ and let $d = \bar{s}[1]$. Note then the first bit of $\mathbf{c}[i]$ must be d because the if statement in h_k^s must

```

Adversary  $\mathcal{A}_{h'}(k, \epsilon)$ 
 $s \leftarrow S : (M_1, M_2) \leftarrow \mathcal{A}_H((k, s), \epsilon)$ 
 $((m_1, c_1), (m_2, c_2)) \leftarrow B_{\text{cr}}((k, s), M_1, M_2)$ 
Return  $((m_1, c_1[2...|c_1|]), (m_2, c_2[2...|c_2|]))$ 

```

Figure 8: Adversary $\mathcal{A}_{h'}$ for the proof of Proposition 6.3

have evaluated to false. Essentially the same reasoning implies the first bit of $\mathbf{c}[i - 1]$ is d .

We can continue this argument for each i back to 1. However, this contradicts the fact that $\mathbf{c}[1] = s$, so h_k^s never returns $\bar{s}$. $\square$

PROPOSITION 6.3. *Given an adversary $\mathcal{A}_H$, let $\mathcal{A}_{h'}$ be the adversary of Fig. 8. Then*

$$\text{Adv}_H^{\text{cr}}(\mathcal{A}_H) \leq \text{Adv}_{h'}^{\text{cr}}(\mathcal{A}_{h'}). \quad (2)$$

The time complexity of $\mathcal{A}_{h'}$ is the sum of the time complexities of $\mathcal{A}_H$ and B_{cr} . The memory complexity of $\mathcal{A}_{h'}$ is the maximum of the memory complexity of $\mathcal{A}_H$ and the memory complexity of B_{cr} .

Notice that Equation (2) tells us that if h' is collision resistant, then H is as well. Let $\mathcal{A}_H$ be a practical adversary against H . Then $\mathcal{A}_{h'}$ is also practical because its efficiency is about that of $\mathcal{A}_H$. This means that if h' is collision resistant, $\text{Adv}_H^{\text{cr}}(\mathcal{A}_H)$ is low. Equation (2) tells us that $\text{Adv}_H^{\text{cr}}(\mathcal{A}_H)$ will be at most $\text{Adv}_{h'}^{\text{cr}}(\mathcal{A}_{h'})$, which means H is also collision resistant.

PROOF. (of Proposition 6.3) The facts about the time and memory of $\mathcal{A}_{h'}$ are clear from its pseudocode.

Now we claim that if the message pair M_1, M_2 returned by $\mathcal{A}_H$ is a collision for $H_{(k,s)}$ then $\mathcal{A}_{h'}$ will return a collision for h'_k . Adversary $\mathcal{A}_{h'}$ takes input $k \in h'.\text{Keys}$. It then runs $\mathcal{A}_H$ on input ϵ given key (k, s) to get a pair of messages (M_1, M_2) in $\text{Split}.Inp$. Then it runs B_{cr} to obtain a pair of inputs to h , which we will refer to as (m_1, c_1) and (m_2, c_2) . It then returns these (after removing the first bits of c_1 and c_2).

Suppose M_1, M_2 is a collision for $H_{(k,s)}$. Since Split is suffix-free, $\text{Split}(M_1) \not\supseteq \text{Split}(M_2)$ and $\text{Split}(M_2) \not\supseteq \text{Split}(M_1)$. Then by Lemma 5.2, we know that B_{cr} will have returned a collision for h_k^s .

From Proposition 6.2 we also know that $h_k^s((m_1, c_1)) \neq \bar{s}$ and $h_k^s((m_2, c_2)) \neq \bar{s}$. Then it must be the case that they cause the if statement in h^s to evaluate to true and so $h_k^s((m_1, c_1[2...|c_1|])) = h'_k((m_1, c_1[2...|c_1|]))$. Furthermore, $(m_1, c_1) \neq (m_2, c_2)$ and $c_1[1] = c_2[1] = s[1]$, so $(m_1, c_1[2...|c_1|]) \neq (m_2, c_2[2...|c_2|])$ and thus they form a collision for h'_k .

Therefore, adversary $\mathcal{A}_{h'}$ finds a collision in h'_k whenever $\mathcal{A}_H$ does for $H_{(k,s)}$. This justifies Equation (2), completing the proof. $\square$

DEFINING A NEW CONSTRAINT FOR CR. The previous example established that traditional definitions of collision resistance with the MD transform do not fully capture the security behind the construction. Although the compression function h used to construct the hash function H was not collision resistant, we were still able to prove the collision resistance of H .

An obvious question at this point is whether there is a natural, weaker assumption we could place on h from which we can still

prove H is collision resistant. We answer this in the affirmative with a new security definition in the RS security framework. For this, we now define a new relation which is strictly harder for the adversary to satisfy than R_{cr} , making it a weaker assumption on h . Despite this, we can recover our result that the MD transform is *fully* collision resistant under the assumption that h is $R_{cr}S$ secure for any suffix-free splitting function. We call our new security definition constrained collision resistance, or R_{ccr} , and provide the pseudocode for the relation below. We previously defined R_{ccr} in Fig. 1.

Relation $R_{ccr}(k, s, out)$

$(x_1, x_2, a_1, a_2) \leftarrow out ; (m_1, c_1) \leftarrow x_1 ; (m_2, c_2) \leftarrow x_2$
 $coll \leftarrow R_{cr}(k, \varepsilon, ((m_1, c_1), (m_2, c_2)))$
 $valid \leftarrow ((c_1 \in \{s, h_k(a_1)\}) \text{ and } (c_2 \in \{s, h_k(a_2)\}))$
 Return ($coll$ and $valid$)

This relation makes the adversary's job harder than for collision resistance by putting further restrictions on the collisions it is allowed to submit. In particular, it requires that for both chaining variables in the collision submitted by the adversary, this chaining variable must be s or the adversary must know a pre-image for it.

Now we proceed to proving that the MD transform gives a collision resistant hash function if the splitting function is suffix-free and the compression function is constrained-collision resistant. This result helps provide some theoretical understanding to the observation that collisions in the compression functions underlying MD-style hash functions tend not to immediately result in the entire hash function being broken.

THEOREM 6.4. *Let h be a compression function, let $Split$ be a suffix-free splitting function with $Split.BI = h.BI$ and let $S \subseteq h.Out$ be a set of possible starting points. Let $H = MD[h, Split, S]$ be the hash function associated to these components via the MD transform of Fig. 2. Given an adversary $\mathcal{A}_H$, let $\mathcal{A}_h$ be the adversary of Fig. 9 using algorithm B_{ccr} . Then*

$$Adv_H^{cr}(\mathcal{A}_H) \leq Adv_h^{R_{ccr}S}(\mathcal{A}_h). \quad (3)$$

The time complexity of $\mathcal{A}_h$ is the sum of the time complexities of $\mathcal{A}_H$ and B_{ccr} . The memory complexity of $\mathcal{A}_h$ is the maximum of the memory complexity of $\mathcal{A}_H$ and the memory complexity of B_{ccr} .

The algorithm B_{ccr} mentioned above (and defined in Fig. 9) is an extension of B_{cr} to also return the values a_1, a_2 expected by R_{ccr} . We discuss it in more detail in the proof.

Equation (1) tells us that if h is constrained-collision resistant, then H is collision resistant. Let $\mathcal{A}_H$ be a practical adversary against H . Then $\mathcal{A}_h$ is also practical because its efficiency is about that of $\mathcal{A}_H$. If h is constrained collision resistant, then $Adv_h^{R_{ccr}S}(\mathcal{A}_h)$ will be low. Equation (3) tells us that $Adv_H^{cr}(\mathcal{A}_H)$ will be at most $Adv_h^{R_{ccr}S}(\mathcal{A}_h)$, which means H is collision resistant.

PROOF. (of Theorem 6.4) The claimed bounds on the complexity of $\mathcal{A}_h$ are clear from its pseudocode.

Adversary $\mathcal{A}_h$ takes as input a random $(k, s) \in h.Keys \times S$. It runs $\mathcal{A}_H$ on input ε and key (k, s) to get a pair of messages (M_1, M_2) in $Split.Inp$. Note this exactly matches the input distribution $\mathcal{A}_H$

expects to be given. Adversary $\mathcal{A}_h$ can then run the algorithm B_{ccr} shown in Fig. 9 with inputs $((k, s), M_1, M_2)$ for it to extract a collision and appropriate information about the pre-images of this collision, if required.

Assume that M_1, M_2 is a collision for $H_{(k, s)}$. Since $Split$ is suffix-free, $Split(M_1) \not\subseteq Split(M_2)$ and $Split(M_2) \not\subseteq Split(M_1)$.

We may think of B_{ccr} as a similar algorithm to B_{cr} , with the added task of finding pre-images for the chaining variables in its colliding messages. Indeed, B_{ccr} creates the vectors of chaining variables, c_1 and c_2 , and searches for a collision in the same way as B_{cr} , returning this message pair at which it found a collision. Thus, Lemma 5.2 guarantees that the pair $(m_1, c_1), (m_2, c_2)$ forms a collision for h_k . We must verify that a_1, a_2 returned by B_{ccr} additionally satisfies $c_1 \in \{s, h_k(a_1)\}$ and $c_2 \in \{s, h_k(a_2)\}$. Let m_1, c_1, n_1 and m_2, c_2, n_2 be the values calculated by B_{ccr} when run by $\mathcal{A}_{HH}$.

First suppose that B_{ccr} halts in the middle of the execution of its for loop. Then it is clear for the manner they were created that a_1 will be a pre-image for c_1 and a_2 will be a pre-image for c_2 .

Now suppose that B_{ccr} does not halt until after the for loop is complete. We will separately analyze the case that $n_1 = n_2$ and the case that $n_1 \neq n_2$. In the former case the chaining variables c_1 and c_2 specifying the collision are $c_1[1]$ and $c_2[1]$, respectively. Since these are both equal to s , $\mathcal{A}_h$ does not need to provide a pre-image for them and we are done. In the latter case the above reasoning tells us that $c_b = s$ (because c_b corresponds to the shorter vector) and so a pre-image is not required for it. This will, presumably, not hold for c_{3-b} (which corresponds to the longer vector), so a pre-image is required for it. As with our earlier analysis we can see that a_{3-b} is a pre-image for c_{3-b} under h_k . Then for compactness $\mathcal{A}_h$ arbitrarily returns this pre-image for both messages and we are again done.

Thus, on any input (k, s) , adversary $\mathcal{A}_h$ finds a constrained collision in h_k when $\mathcal{A}_H$ finds a collision in $H_{(k, s)}$. This justifies Equation (3). $\square$

A CCR h . With the introduction of $R_{ccr}S$ security, one might ask whether this assumption is necessary for any h to produce an MD transform that is collision resistant. It is, in fact, not necessary, although it is sufficient. Indeed, the compression function h_k^s given in Fig. 7 is itself not $R_{ccr}S$ secure, yet we have shown that it results in a collision resistant MD transform. We prove this result below. This shows that an assumption on h even weaker than CCR could suffice, and the benefit of our framework is that one could easily do so. However, one has to make some value judgment about the tradeoff between the assumptions and the result. In the extreme, the assumption on h could just be that the MD transform on it is $R_{cr}S$ secure, which is not a useful result. The advantage of $R_{ccr}S$ is that it is appropriately balanced: it is meaningfully weaker than $R_{cr}S$, yet the implication that the MD transform is $R_{cr}S$ secure is still non-trivial.

PROPOSITION 6.5. *Let $M_1 = a_1 = (0^b, \bar{s})$, $M_2 = a_2 = (1^b, \bar{s})$, and $\mathcal{B}_{h^s}$ be the adversary shown in Fig. 10. Then $Adv_{h^s}^{R_{ccr}S}(\mathcal{B}_{h^s}) = 1$.*

Put simply, the above tells us that h^s is not constrained collision resistant because $\mathcal{B}_{h^s}$ is clearly efficient.

Algorithm $B_{\text{ccr}}((k, s), M_1, M_2)$ $\mathbf{m}_1 \leftarrow \text{Split}(M_1); \mathbf{m}_2 \leftarrow \text{Split}(M_2); n_1 \leftarrow \mathbf{m}_1 ; n_2 \leftarrow \mathbf{m}_2$ $c_1[1] \leftarrow s; c_2[1] \leftarrow s$ For $i = 1, \dots, n_1$ do $c_1[i+1] \leftarrow h_k((\mathbf{m}_1[i], c_1[i]))$ For $i = 1, \dots, n_2$ do $c_2[i+1] \leftarrow h_k((\mathbf{m}_2[i], c_2[i]))$ $b \leftarrow \text{argmin}_d(n_d)$ For $i = 0, \dots, n_b - 2$ do $(m_1, c_1) \leftarrow (\mathbf{m}_1[n_1 - i], c_1[n_1 - i])$ $(m_2, c_2) \leftarrow (\mathbf{m}_2[n_2 - i], c_2[n_2 - i])$ $a_1 \leftarrow (\mathbf{m}_1[n_1 - i - 1], c_1[n_1 - i - 1])$ $a_2 \leftarrow (\mathbf{m}_2[n_2 - i - 1], c_2[n_2 - i - 1])$ If $(m_1, c_1) \neq (m_2, c_2)$ then return $((m_1, c_1), (m_2, c_2), a_1, a_2)$ If $n_1 = n_2$ then $(m_1, c_1) \leftarrow (\mathbf{m}_1[1], c_1[1]); (m_2, c_2) \leftarrow (\mathbf{m}_2[1], c_2[1])$ $a_1 \leftarrow 1; a_2 \leftarrow 2$ Return $((m_1, c_1), (m_2, c_2), a_1, a_2)$ $(m_1, c_1) \leftarrow (\mathbf{m}_1[n_1 - n_b + 1], c_1[n_1 - n_b + 1])$ $(m_2, c_2) \leftarrow (\mathbf{m}_2[n_2 - n_b + 1], c_2[n_2 - n_b + 1])$ $a_{3-b} \leftarrow (\mathbf{m}_{3-b}[n_{3-b} - n_b], c_{3-b}[n_{3-b} - n_b])$ $a_b \leftarrow a_{3-b}$ Return $((m_1, c_1), (m_2, c_2), a_1, a_2)$
Adversary $\mathcal{A}_h(k, s)$ $(M_1, M_2) \leftarrow \mathcal{A}_H((k, s), \varepsilon)$ Return $B_{\text{ccr}}((k, s), M_1, M_2)$

Figure 9: Adversary $\mathcal{A}_h$ for the proof of Theorem 6.4

Adversary $\mathcal{B}_{h^s}(k, s)$ Return (M_1, M_2, a_1, a_2)
$h_k^\dagger((m, c))$ If $(m, c) \in \{(0^b, 1\ 0^c), (1^b, 1^2\ 0^{c-1})\}$ Return 1^{c+1} Return $0\ h'_k((m, c))$
Adversary $\mathcal{B}_{h''}(k, \varepsilon)$ $((m_1, c_1), (m_2, c_2), a_1, a_2) \leftarrow \mathcal{A}_{h^\dagger}(k, \varepsilon)$ Return $((m_1, c_1), (m_2, c_2))$

Figure 10: $\mathcal{B}_{h^s}$ for Proposition 6.5, $h_k^\dagger$ for Proposition 6.6, and $\mathcal{B}_{h''}$ for Proposition 6.6

PROOF. (of Proposition 6.5) As shown in the proof of Proposition 6.1, M_1 and M_2 form a collision for h^s . Notice that for $M_1 = (0^b, \bar{s})$, a preimage for $\bar{s}$ is simply M_1 itself. Similarly M_2 is a preimage for $\bar{s}$. We thus let $a_1 = M_1$ and $a_2 = M_2$. Therefore, h^s is not constrained collision resistant. $\square$

A natural question that might arise is whether R_{ccrS} security is actually strictly weaker than R_{crS} . With the revelation that h_k^s in Fig. 7 is not R_{crS} secure, is there any compression function that is R_{ccrS} secure yet is not R_{crS} secure? We claim that such a compression function does exist and give an example in Fig. 10.

We again let $b, c \in \mathbb{N}$. Given a good collision resistant function $h'' : h''.\text{Keys} \times (\{0, 1\}^b \times \{0, 1\}^{c+1}) \rightarrow \{0, 1\}^c$, we construct the compression function $h^\dagger : h''.\text{Keys} \times (\{0, 1\}^b \times \{0, 1\}^{c+1}) \rightarrow \{0, 1\}^{c+1}$ shown in Fig. 10.

It is clear that $h^\dagger$ is not collision resistant, since for the distinct inputs $(0^b, 1\|0^c)$ and $(1^b, 1^2\|0^{c-1})$ it returns 1^{c+1} . Despite this, we now show that $h^\dagger$ is instead constrained collision resistant.

PROPOSITION 6.6. *Given an adversary $\mathcal{A}_{h^\dagger}$, let $\mathcal{B}_{h''}$ be the adversary of Fig. 10. Then*

$$\text{Adv}_{h^\dagger}^{\text{RccrS}}(\mathcal{A}_{h^\dagger}) \leq \text{Adv}_{h''}^{\text{cr}}(\mathcal{B}_{h''}) \quad (4)$$

and both the time and memory complexity of $\mathcal{B}_{h''}$ are about that of $\mathcal{A}_{h^\dagger}$.

Notice that Equation (4) tells us that if h'' is collision resistant, then $h^\dagger$ is constrained collision resistant. Let $\mathcal{A}_{h^\dagger}$ be a practical adversary against $h^\dagger$. Then $\mathcal{B}_{h''}$ is also practical because its efficiency is about that of $\mathcal{A}_{h^\dagger}$. This means that if h'' is collision resistant, $\text{Adv}_{h''}^{\text{cr}}(\mathcal{B}_{h''})$ is low. Equation (4) tells us that $\text{Adv}_{h^\dagger}^{\text{RccrS}}(\mathcal{A}_{h^\dagger})$ will be at most $\text{Adv}_{h''}^{\text{cr}}(\mathcal{B}_{h''})$, which means $h^\dagger$ is constrained collision resistant.

PROOF. (of Proposition 6.6) The running time claim is analogous to our prior ones.

We claim that if the tuple (M_1, M_2, a_1, a_2) returned by $\mathcal{A}_{h^\dagger}$ is a constrained collision for $h^\dagger$ then $\mathcal{B}_{h''}$ will return a collision for h'' . Adversary $\mathcal{B}_{h''}$ takes input $k \in h''.\text{Keys}$. It then runs $\mathcal{A}_{h^\dagger}$ on input the given key k and ε to get the tuple $((m_1, c_1), (m_2, c_2), a_1, a_2)$, where $(m_1, c_1), (m_2, c_2) \in \{0, 1\}^b \times \{0, 1\}^{c+1}$ and $a_1, a_2 \in \{0, 1\}^{c+1}$. Since $\mathcal{A}_{h^\dagger}$ returns a constrained collision, it must be true that $(m_1, c_1) \neq (m_2, c_2)$, $h_k^\dagger((m_1, c_1)) = h_k^\dagger((m_2, c_2))$, $c_1 \in \{s, h_k^\dagger(a_1)\}$, and $c_2 \in \{s, h_k^\dagger(a_2)\}$.

Since $h_k^\dagger$ will only output strings of all 1s or strings that start with 0, it can never output $1\|0^c$ or $1^2\|0^{c-1}$. Thus, neither string has a pre-image, so $(m_1, c_1), (m_2, c_2) \notin \{(0^b, 1\|0^c), (1^b, 1^2\|0^{c-1})\}$. On input (m_1, c_1) , the if statement in $h_k^\dagger$ will be false, so $h_k^\dagger$ will return $0\|h'_k((m_1, c_1))$. Similarly, on input (m_2, c_2) , $h_k^\dagger$ will return $0\|h'_k((m_2, c_2))$. Since $h_k^\dagger((m_1, c_1)) = h_k^\dagger((m_2, c_2))$, we can conclude that $h'_k((m_1, c_1)) = h'_k((m_2, c_2))$, therefore forming a collision for h'' . Adversary $\mathcal{B}_{h''}$ returns this message pair, so it finds a collision in h'' whenever $\mathcal{A}_{h^\dagger}$ finds a constrained collision for $h^\dagger$. This justifies Equation (4), completing the proof. $\square$

7 A MINIMAL TRANSFORM

Having to use a suffix-free splitting function necessarily adds some computational overhead to the computation of the MD hash function over what would have been necessary if we were able to use a minimal splitting function. For instance, one such splitting function could be padding M with a single one bit and then with as many zeros as necessary to be of a block size. If the message M is particularly short, this padding scheme may only require one invocation of h when a suffix-free padding function would likely have increased the length of M enough to require a second such invocation.

As such, in some use cases it would be beneficial to use such a minimal splitting function in the transform. We saw earlier that we cannot hope to use a splitting function which is not suffix-free assuming only that the underlying compression function is collision resistant, so it may seem that this efficiency gain would be countered by a loss in provable security. In this section we show this is not the case, establishing that if the compression function is constrained collision resistant and it is difficult to find a pre-image for a randomly chosen s from S , then it suffices for the splitting function to be injective. This result serves to generalize, formalize, and strengthen the informal folklore claim of Andreeva and Stam [3] that the plain MD transform preserves collision resistance assuming it is hard to find a pre-image of the initial vector.

It is important to emphasize here that our proof assumes only the constrained collision resistance of the compression function and, thus, this use of MD may similarly enjoy collision resistance even after a collision is found in the underlying compression function. We leave it to others to decide when this gain in efficiency is worth the security tradeoff required in assuming an additional security property of the hash function.

It will be convenient to first prove a lemma we will use in our proof. The lemma establishes that any collision in the MD transform must necessarily give either a collision in the underlying compression function or a pre-image of s in that compression function. This builds on Lemma 5.2 to classify MD collisions by additionally considering what happens when the splitting function is not necessarily suffix-free. It again is not a computational statement; it is a fact about the structure of collisions for the MD transform.

LEMMA 7.1. *Let h be a compression function, Split be a splitting function with $\text{Split.Bl} = h.\text{Bl}^*$, and $S \subseteq h.\text{Out}$ be a set of possible starting points. Let $H = \text{MD}[h, \text{Split}, S]$. Let $k \in h.\text{Keys}$, $s \in S$ and suppose $M_1, M_2 \in \text{Split.Inp}$ form a collision for $H_{(k,s)}$. Then at least one of the following two conditions holds:*

- (1) *On inputs $(k, s), M_1, M_2$, the algorithm B_{ccr} shown in Fig. 9 returns $((m_1, c_1), (m_2, c_2), a_1, a_2)$ such that $(m_1, c_1), (m_2, c_2)$ form a collision for h_{Hk} and both $(c_1 \in \{s, h_{Hk}(a_1)\})$ and $(c_2 \in \{s, h_{Hk}\})$ hold.*
- (2) *On input $(k, s), M_1, M_2$, algorithm B_{pre} of Fig. 11 returns a pre-image for s under h_{Hk} .*

PROOF. (of Lemma 7.1) Let $\mathbf{m}_1 = \text{Split}(M_1)$, $\mathbf{m}_2 = \text{Split}(M_2)$, $n_1 = |\mathbf{m}_1|$, and $n_2 = |\mathbf{m}_2|$, as defined in algorithms B_{ccr} and B_{pre} .

As detailed in the proof of Theorem 6.4 and from Lemma 5.2, if $\mathbf{m}_1 \not\sqsupseteq \mathbf{m}_2$ and $\mathbf{m}_2 \not\sqsupseteq \mathbf{m}_1$, then the first condition holds. So suppose without loss of generality that $\mathbf{m}_1 \supseteq \mathbf{m}_2$. Note it then must hold that $n_1 \geq n_2$.

First suppose there exists an $i \in \{0, \dots, n_2 - 1\}$ such that $(\mathbf{m}_1[n_1 - i], c_1[n_1 - i]) \neq (\mathbf{m}_2[n_2 - i], c_2[n_2 - i])$. Let j be the smallest such value. It will then hold that $c_1[n_1 - (j - 1)] = c_2[n_2 - (j - 1)]$, so $(\mathbf{m}_1[n_1 - j], c_1[n_1 - j])$ and $(\mathbf{m}_2[n_2 - j], c_2[n_2 - j])$ form a collision for h_k . During the execution of B_{ccr} it will then return this collision when $i = j$ in the for loop (or after the for loop in the case that $j = n_2 - 1$). The same reasoning as in the proof of Theorem 6.4 tells us that the condition $(c_1 \in \{s, h_{Hk}(a_1)\})$ and $(c_2 \in \{s, h_{Hk}\})$ will hold of the values returned by B_{ccr} .

Algorithm $B_{\text{pre}}((k, s), M_1, M_2)$

```

 $\mathbf{m}_1 \leftarrow \text{Split}(M_1); n_1 \leftarrow |\mathbf{m}_1|$ 
 $\mathbf{m}_2 \leftarrow \text{Split}(M_2); n_2 \leftarrow |\mathbf{m}_2|$ 
 $b \leftarrow \text{argmax}_d(n_d)$ 
For  $i = 1, \dots, n_b - n_{3-b}$  do
   $c_b[i + 1] \leftarrow h_k((\mathbf{m}_b[i], c_b[i]))$ 
Return  $(\mathbf{m}_b[n_b - n_{3-b}], c_b[n_b - n_{3-b}])$ 

```

Adversary $\mathcal{B}_h(k, s)$

```

 $(M_1, M_2) \leftarrow \mathcal{A}_H(k, s)$ 
Return  $B_{\text{pre}}(k, s)$ 

```

Figure 11: Adversary $\mathcal{B}_h$ for Theorem 7.2.

Now suppose $(\mathbf{m}_1[n_1 - i], c_1[n_1 - i]) \neq (\mathbf{m}_2[n_2 - i], c_2[n_2 - i])$ for all $i \in \{0, \dots, n_2 - 1\}$. This implies, in particular, that $c_1[n_1 - (n_2 - 1)] = c_2[n_2 - (n_2 - 1)] = s$. Note it must then hold that $n_1 > n_2$ (because otherwise we would have $\mathbf{m}_1 = \mathbf{m}_2$ contradicting the fact that $M_1 \neq M_2$). Hence, $h_k((\mathbf{m}_1[n_1 - n_2], c_1[n_1 - n_2])) = s$. Noting then that $(\mathbf{m}_1[n_1 - n_2], c_1[n_1 - n_2])$ will be returned by B_{pre} , we have that in all cases at least one of the two conditions of the lemma hold as desired. $\square$

Having established the above lemma we can move on to the main result of this section, that Split being merely injective (and not necessarily suffix-free) suffices to prove that the MD transform gives collision resistance if it is hard to find a pre-image for a starting point randomly chosen from S . Mirroring Theorem 6.4, we will in fact show the result assuming only the weaker notion of constrained collision resistance for the compression function h .

THEOREM 7.2. *Let h be a compression function, let Split be an injective splitting function with $\text{Split.Bl} = h.\text{Bl}^*$ and let $S \subseteq h.\text{Out}$ be a set of possible starting points. Let $H = \text{MD}[h, \text{Split}, S]$. Given an adversary $\mathcal{A}_H$, let $\mathcal{A}_h$ be the adversary of Fig. 9 and $\mathcal{B}_h$ be the adversary of Fig. 11. Then*

$$\text{Adv}_H^{\text{cr}}(\mathcal{A}_H) \leq \text{Adv}_h^{\text{RccrS}}(\mathcal{A}_h) + \text{Adv}_h^{\text{RpreS}}(\mathcal{B}_h). \quad (5)$$

The time complexity of $\mathcal{A}_h$ is about that of $\mathcal{A}_H$ plus that of B_{ccr} . The memory complexity of $\mathcal{A}_h$ is the maximum of that of $\mathcal{A}_H$ and that of B_{ccr} . The time complexity of $\mathcal{B}_h$ is about that of $\mathcal{A}_H$ plus that of B_{pre} . The memory complexity of $\mathcal{A}_H$ is the maximum of that of $\mathcal{A}_h$ and that of B_{pre} .

Stating that Split is injective is redundant (splitting functions are required to be injective), but we state this explicitly above to emphasize that injectivity is the *only* property we assume of Split.

The theorem proceeds fairly easily from Lemma 7.1 because the adversaries simply run B_{ccr} and B_{pre} .

PROOF. (of Theorem 7.2) Consider the view of adversary $\mathcal{A}_H$ when run by either $\mathcal{A}_h$, $\mathcal{B}_h$, or in $\mathbf{C}_F^{\text{cr}}(\mathcal{A}_H)$. In each, it consists of a key k and starting point s , both of which were chosen uniformly at random from their respective sets. From Lemma 7.1, we know if $\mathcal{A}_H$ successfully finds a collision in H for (k, s) it must be the case that one of $\mathcal{A}_h$ or $\mathcal{B}_h$ will be successful in their respective games (because they simply run B_{ccr} and B_{pre} , respectively). Then

we have the following inequality, that establishes the result,

$$\begin{aligned}\text{Adv}_H^{\text{cr}}(\mathcal{A}_H) &= \Pr[G_F^{\text{cr}}(\mathcal{A}_H)] \\ &\leq \Pr[G_h^{\text{RccrS}}(\mathcal{A}_h)] + \Pr[G_h^{\text{RpreS}}(\mathcal{B}_h)] \\ &= \text{Adv}_h^{\text{RccrS}}(\mathcal{A}_h) + \text{Adv}_h^{\text{RpreS}}(\mathcal{B}_h).\end{aligned}$$

The claims on the time and memory complexities of the adversaries are clear. $\square$

8 ASSUMPTION OF PRE-IMAGE RESISTANCE

We showed that the MD transform can be simplified by additionally assuming that the compression function h satisfies a pre-image resistance property. To help understand this result we restate some results for the literature in our language, as well as prove some new results to understand how strong of an assumption this is.

The notion of R_{preS} , defined in Fig. 1, with $S = h.\text{Out}$ has been considered under several different names. In [13], Brown refers to it as both one-wayness and pre-image resistance and uses it as one of several assumptions on a hash function to prove a digital signature scheme is secure. BRS [12] refer to it as inversion resistance and analyze the security of hash functions in an idealized model. In [21], Laccetti and Schmid refer to it as pre-image resistance and analyze the success probability of a brute force attack. In [32], Stinson refers to it as the pre-image problem and bounds the advantage of an adversary in the random oracle model as well as giving reductions between it and some other security notions. Andreeva and Stam [3] refer to it as range-oriented pre-image finding and relate its security to various other notions of pre-image resistance. Finally, DRS [18] refer to it as a variant of pre-image resistance and cite [17] for the definition, though the latter actually considers a weaker notion.

The notion of R_{preS_0} (where $S_0 = \{0\}$ is the singleton set containing only some fixed zero string) is considered by both [13] and [32], where it is referred to, respectively, as zero-finder-resistance and the zero pre-image problem.

The everywhere pre-image resistant (ePre) security notion of [28] defines the advantage of an adversary as $\max_{S \in \mathcal{S}} \text{Adv}_F^{\text{RpreS}}(\mathcal{A})$ where $\mathcal{S}$ is defined as the set of all sets $S \subseteq \{0, 1\}^*$ with $|S| = 1$. They consider the relationship between this and six other security notions. In their framework, this is equivalent to a definition in which s is chosen from an arbitrary adversarially chosen distribution. This is the case when only reasoning about worst case adversaries, but is not equivalent when one worries about explicit reductions, as we do.

Most of these works [12, 13, 17, 18, 21, 32] were interested only in keyless hash functions while [28] was interested only in keyed hash functions. We intentionally model both situations.

NEGATIVE RESULTS. Pre-image resistance and collision resistance are not, in general, implied by each other. To see that pre-image resistance does not imply collision resistance, consider the hash function g which on any input x returns its randomly chosen key k as output. This is trivially pre-image resistant but not collision resistant as formalized by the following proposition.

PROPOSITION 8.1. *Let Out be a set. Let g be the hash function shown in Fig. 12 where $g.\text{Keys} = g.\text{Out} = \text{Out}$ and $g.\text{Inp}$ is arbitrary. Let $S \subseteq \text{Out}$. Then for any $\mathcal{A}$ it holds that $\text{Adv}_g^{\text{RpreS}}(\mathcal{A}) \leq 1/|\text{Out}|$.*

$g_k(x)$	$h_k((m, c))$	$f_{(k,p)}(x)$
Return k	If $((m, c[1]) = (0, 0))$	$y \leftarrow f'_k(x)$
	Return c	Return $y \oplus p$
adversary $\mathcal{B}_g(k, s)$	$r \leftarrow h'_k((m, c))$	adversary $\mathcal{B}_{f'}(k, s')$
Return (x_1, x_2)	Return $1 r$	$s \leftarrow s; p \leftarrow s \oplus s'$
	adversary $\mathcal{B}_h(k, s)$	$x \leftarrow \mathcal{A}((k, p), s)$
	Return $(0, s)$	Return x

Figure 12: Compression functions and adversaries used for Proposition 8.1, Proposition 8.2, and Proposition 8.4.

Additionally, if x_1, x_2 are two distinct elements of $g.\text{Inp}$, then the adversary $\mathcal{B}_g$ shown in Fig. 12 satisfies $\text{Adv}_g^{\text{cr}}(\mathcal{B}_g) = 1$.

PROOF. (of Proposition 8.1) The first part of this claim holds because the only way that a pre-image of the chosen $s \in S$ will even *exist* is if k happens to equal it. This happens with an exact probability of $1/|\text{Out}|$. The second part of this claim holds because two distinct elements of $g.\text{Inp}$ always form a collision for g_k . $\square$

It is also the case that pre-image resistance is not implied by collision-resistance. This is already implicit in our work from the combination of Theorem 7.2 and Theorem 5.6 for $S = h.\text{Out}$ or Proposition 5.5 for $S = \{0\}$. As noted there, those techniques could be used to show the same result for any “reasonable” choice of S . For concreteness, we extract out the core idea of these theorems.

First consider $S = \{s\}$ for some fixed string s and a hash function h' whose output space is all bit strings of the same length as s . Then define h by $h_k(x) = h'_k(x) \oplus h'_k(0) \oplus s$, where 0 is some fixed string in the input space of h . Then any collision in h is a collision in h' and 0 is trivially a pre-image of s . So h is not S pre-image resistant, but it is collision resistant (assuming h' was).

The following proposition establishes the result for a particular, larger $S = h.\text{Out}$ and can easily be extended to any reasonable choice of a “large” S .

PROPOSITION 8.2. *Let $a \in \mathbb{N}$. Let h' be a family of functions with $h'.\text{Inp} = \{0, 1\} \times \{0, 1\}^a$ and $h'.\text{Out} = \{0, 1\}^{a-1}$. Then let h be the compression function shown in Fig. 12 with $h.\text{Inp} = h'.\text{Inp}$, $h.\text{Out} = \{0, 1\}^a$, and $h.\text{Keys} = h'.\text{Keys}$. Then for all adversaries $\mathcal{A}$, $\text{Adv}_h^{\text{cr}}(\mathcal{A}) \leq \text{Adv}_{h'}^{\text{cr}}(\mathcal{A})$. Furthermore, letting $S = h.\text{Out}$, we can build an efficient adversary $\mathcal{B}_h$ (shown in Fig. 12) satisfying,*

$$\text{Adv}_h^{\text{RpreS}}(\mathcal{B}_h) \geq 1/2.$$

PROOF. (of Proposition 8.2) To see that the first claim of the above statement holds, note that no collisions in h are possible unless the if statement evaluates to false for both inputs of the collision. In this case, a collision for h is immediately a collision for h' . To see that the second claim holds, observe that for half the choices of $s \in S$, specifically those whose first bit are 0, the compression function h_k returns s when given as input $(0, s)$. $\square$

POSITIVE RESULTS. We just showed that pre-image resistance is not, in general, implied by collision resistance. This implies that the pre-image resistance assumed for Theorem 7.2 is necessarily

Game $G_f^{\text{unif}}(\mathcal{A})$	Adversary $\mathcal{A}_1(k, s)$	Adversary $\mathcal{A}_2(k, \epsilon)$
$b \leftarrow \{0, 1\}$ $k \leftarrow f.\text{Keys}$ $x \leftarrow f.\text{Inp}$ $s_0 \leftarrow f_k(x)$ $s_1 \leftarrow f.\text{Out}$ $b' \leftarrow \mathcal{A}(k, s_b)$ Return $(b = b')$	$x \leftarrow \mathcal{A}(k, s)$ If $(f_k(x) = s)$ Return 1 Return 0	$x_1 \leftarrow f.\text{Inp}$ $s \leftarrow f_k(x_1)$ $x_2 \leftarrow \mathcal{A}(k, s)$ Return (x_1, x_2)

Figure 13: Game defining uniformity of compression function F and adversaries used in proof of Theorem 8.3.

Game $[G_0, G_1]$	Game G_2
$k \leftarrow f.\text{Keys}; \text{bad} \leftarrow \text{false}$ $x_1 \leftarrow f.\text{Inp}$ $s \leftarrow f_k(x_1); \boxed{s \leftarrow f.\text{Out}}$ $x_2 \leftarrow \mathcal{A}(k, s)$ If $(x_1 = x_2)$ then $\text{bad} \leftarrow \text{true}$ Return $(f_k(x_2) = s)$	$k \leftarrow f.\text{Keys}; \text{bad} \leftarrow \text{false}$ $x_1 \leftarrow f.\text{Inp}$ $s \leftarrow f_k(x_1)$ $x_2 \leftarrow \mathcal{A}(k, s)$ If $(x_1 = x_2)$ then $\text{bad} \leftarrow \text{true}$ Return false Return $(f_k(x_2) = s)$

Figure 14: Games used in proof of Theorem 8.3. Boxed code is only executed in the comparably boxed game.

a separate assumption than the assumption of collision resistance. There is, however, an assumption we can make on the structure of the compression function for which pre-image resistance will not be a separate assumption. In particular, collision resistance will imply pre-image resistance when the image of a random point is indistinguishable from a random range point.

The ideas of this proof are very similar to the method used by [12] to show that R_{preS} is essentially equivalent to the typical notion of one-way function security for their particular idealized hash functions. Stinson [32] likewise showed that pre-image resistance was implied by collision resistance in certain cases. Specifically, he showed it for the case that the pre-image resistance adversary was equally likely to succeed for all inputs or the case that every range point had a “large” number of pre-images. Our result can be considered a generalization of this latter result. It strengthens these prior results by showing it for potentially keyed hash functions and by giving an explicit reduction to the uniformity of the hash function. Thus, the result will hold even if it is only computationally (but not information theoretically) difficult to distinguish between the output of the hash function and uniform output.

We will define the uniformity of a hash function by the game $G_f^{\text{unif}}(\mathcal{A})$ shown in Fig. 13. This game measures an adversary’s ability to distinguish between the pairs (k, y) and $(k, f_k(x))$ when y is picked at random from $f.\text{Out}$ and x is picked randomly from $f.\text{Inp}$. The advantage of an adversary is defined by $\text{Adv}_f^{\text{unif}}(\mathcal{A}_1) = 2 \Pr[G_f^{\text{unif}}(\mathcal{A})] - 1$.

It is important to note that this requires the output of f to look uniformly random *even given* k . We now present our theorem which tells us that if f is sufficiently uniform then collision resistance will imply pre-image resistance.

THEOREM 8.3. *Let f be a family of functions, $S = f.\text{Out}$, and $\mathcal{A}$ be an adversary. Then we can build adversaries $\mathcal{A}_1$ and $\mathcal{A}_2$ (shown in Fig. 13) such that,*

$$\text{Adv}_f^{\text{RpreS}}(\mathcal{A}) \leq \text{Adv}_f^{\text{unif}}(\mathcal{A}_1) + \text{Adv}_f^{\text{cr}}(\mathcal{A}_2) + |f.\text{Out}|/|f.\text{Inp}|.$$

Both $\mathcal{A}_1$ and $\mathcal{A}_2$ have approximately the same time and memory complexities as $\mathcal{A}$.

Recall from above that the particular case we are interested in is when f is a compression function and thus $f.\text{Inp} = f.\text{Bl} \times f.\text{Out}$. Therefore, the ratio will equal $1/|f.\text{Bl}|$ and the desired implication will hold as long as $f.\text{Bl}$ is large (which is typically the case in practice). To prove the result we first use the uniformity of f to switch to a game in which $\mathcal{A}$ is trying to find a pre-image for a point $f_k(x)$ where x is chosen at random instead of a random s . Then we consider the standard collision resistance adversary which chooses a random x_1 , asks $\mathcal{A}$ to produce a pre-image of $f_k(x_1)$, then returns that together with x_1 as its collision. Analyzing the success of this adversary requires bounding the probability that the pre-image produced by $\mathcal{A}$ is itself x_1 because (x_1, x_1) is not a valid collision.

PROOF. (of Theorem 8.3) Consider the sequence of games G_0 , G_1 , and G_2 shown in Fig. 14. The boxed code is only included in G_0 , meaning that s will be chosen uniformly at random. Note that G_2 only differs from G_1 after the bad flag is set to true. Game G_0 is identical to $G_f^{\text{RpreS}}(\mathcal{A})$ with f and $\mathcal{A}$ hardcoded. In both, $\mathcal{A}$ is given a random key k and string s chosen at random from $f.\text{Out}$. It wins if it correctly returns a pre-image of s under f_k . Thus we have

$$\begin{aligned} \text{Adv}_f^{\text{RpreS}}(\mathcal{A}) &= \Pr[G_f^{\text{RpreS}}(\mathcal{A})] \\ &= \Pr[G_0] \\ &= (\Pr[G_0] - \Pr[G_1]) + (\Pr[G_1] - \Pr[G_2]) + \Pr[G_2]. \end{aligned}$$

To bound the first difference, consider the view of $\mathcal{A}$ when run by $\mathcal{A}_1$ (during the execution of $G_f^{\text{unif}}(\mathcal{A}_1)$). Let b_{unif} denote the bit chosen by the game $G_f^{\text{unif}}(\mathcal{A}_1)$ and b_1 denote the bit output by $\mathcal{A}_1$. When $b_{\text{unif}} = 1$, the view of $\mathcal{A}$ is k and s chosen uniformly at random. Then $\mathcal{A}_1$ returns $b_1 = 1$ if $\mathcal{A}$ returns a pre-image for s . Similarly when $b_{\text{unif}} = 0$, $\mathcal{A}$ is given k and $s = f_k(x)$ for a uniformly random x . In this case, $\mathcal{A}_1$ once again returns $b_1 = 1$ if $\mathcal{A}$ returns a pre-image for s . Thus we have

$$\begin{aligned} \Pr[G_0] - \Pr[G_1] &= \Pr[b_1 = 1 | b_{\text{unif}} = 1] - \Pr[b_1 = 1 | b_{\text{unif}} = 0] \\ &= \text{Adv}_f^{\text{unif}}(\mathcal{A}_1). \end{aligned}$$

That the latter equality holds by a standard conditioning argument.

Now games G_1 and G_2 are identical until bad, so the fundamental lemma of game playing [11] says that $\Pr[G_1] - \Pr[G_2] \leq \Pr[G_1 \text{ sets bad}]$. In particular, the bad flag is set when $\mathcal{A}$ happens to choose x_1 exactly as its pre-image for s . Let k be fixed and $f_k^{-1}(s)$ be the set of all pre-images of s under f_k . The probability that $\mathcal{A}$ is given a particular s is then exactly $|f_k^{-1}(s)|/|f.\text{Inp}|$ and the probability that bad will be set given that $\mathcal{A}$ is given s is at most $1/|f_k^{-1}(s)|$.

Summing over all possible choices of s we then have that bad is set with probability at most $|f.\text{Out}|/|f.\text{Inp}|$ for our chosen k . This was independent of the key k , so the probability averaged over all choices of k will be bounded by the same probability. Hence, $\Pr[G_1 \text{ sets bad}] \leq |f.\text{Out}|/|f.\text{Inp}|$.

Finally, note that the probability $\mathcal{A}$ succeeds in G_2 is identical to the probability that adversary $\mathcal{A}_2$ succeeds in $G_f^{\text{cr}}(\mathcal{A}_2)$. Each succeeds if $\mathcal{A}$, given $f_k(x_1)$, produces an x_2 which is a collision with x_1 for f_k . Hence $\Pr[G_2] = \text{Adv}_f^{\text{cr}}(\mathcal{A}_2)$.

Combining the given equation gives the stated bound. The stated complexities are apparent from the code of $\mathcal{A}_1$ and $\mathcal{A}_2$. $\square$

From the above, believing that f is uniform suffices to remove the necessity that it is pre-image resistant from Theorem 7.2, but only in the case that the set of starting points is all of $f.\text{Out}$. However, for hash functions based on the MD transform which are used in practice, it is clearly the case that the initial point s was not chosen uniformly at random from $f.\text{Out}$. In the following theorem we show how to build a compression function which is S pre-image resistant for an arbitrary S from one which is pre-image resistant for $S' = f.\text{Out}$. Put simply, we include a random pad in the key of our new compression function which is XORd with the output of f .

PROPOSITION 8.4. *Let f' be a family of functions with $f'.\text{Out} = \{0, 1\}^a$ for some $a \in \mathbb{N}$. and f be the compression function shown in Fig. 12 with $f.\text{Keys} = f'.\text{Keys} \times f'.\text{Out}$. Let $S' = f.\text{Out}$ and $S \subseteq \{0, 1\}^a$ be a set. Then given an adversary $\mathcal{A}$, we can build an efficient adversary $\mathcal{B}_f$ (shown in Fig. 12) satisfying,*

$$\text{Adv}_f^{\text{Rpre}S}(\mathcal{A}) \leq \text{Adv}_{f'}^{\text{Rpre}S'}(\mathcal{B}_f).$$

The time and memory complexity of $\mathcal{B}_f$ are essentially that of $\mathcal{A}$.

The use of this new compression function in the MD transform could alternatively be viewed as a variation of the transform in which an additional key is added that gets XORd with all intermediate chaining variables. In terms of collision resistance security, however, this does not gain anything that would not have been equivalently obtained by simply choosing the starting point uniformly at random.

PROOF. (of Proposition 8.4) The view of $\mathcal{A}$ when run by $\mathcal{B}_f$ consists of a random key $k \in f'.\text{Keys}$, a random pad $p \in f'.\text{Out}$, and a random element s from S . This is identical to the view it expects when run in $G_f^{\text{Rpre}S}(\mathcal{A})$, so it will have the same probability of returning an x such that $s = f_{(k,p)}(x) = f'_k(x) \oplus p$. Then from the perspective of $\mathcal{B}_f$, such an x satisfies $f'_k(x) = s \oplus p = s \oplus (s \oplus s') = s'$. This gives us the stated advantage bound because $\mathcal{B}_f$ wins whenever $\mathcal{A}$ would.

The stated complexities are apparent from the code of $\mathcal{A}_f$. $\square$

9 REDUCTION COMPLEXITY

We will now briefly revisit some of our reductions to provide alternative reductions that are more memory-efficient [4].

Consider, for example, Theorem 5.3. The primary technical component underlying that theorem is the collision-finding algorithm B_{cr} which, given as input a collision for the hash function H , finds a collision for the underlying compression function h . This collision-finding algorithm naturally emanates from various proofs of the

MD transform's collision resistance which do not explicitly give an algorithm. [1, 5, 15, 17, 18, 20, 22, 24, 31]. We observe that B_{cr} is a less memory efficient algorithm than the algorithm specified by BBBGKSZ in [5].

Recall the algorithm B_{cr} shown in Fig. 5. It first precomputes the entire vectors $\mathbf{m}_1$, $\mathbf{m}_2$, $\mathbf{c}_1$, and $\mathbf{c}_2$, then processes them *backwards* to find the collision described in the proof of the lemma. Its memory complexity is thus the memory required to store the entire precomputed vectors.

However, this reduction can be done in a more memory efficient manner. We present such an algorithm, called B_{mem} , in Fig. 5. It scans for the collision in the opposite direction, as compared to B_{cr} . By doing so, it avoids the need to precompute the vector $\mathbf{c}$ and instead only computes the individual blocks of $\mathbf{c}$ that it needs in a streaming fashion. Furthermore, for most choices of splitting function considered in practice (e.g. SplitSha), the individual blocks of Split(M) can be computed independently in a memory-efficient manner which would allow B_{mem} to use only a *constant* amount of memory overhead.

THEOREM 9.1. *Let h be a compression function, let Split be a suffix-free splitting function with Split.Bl = h.Bl and let $S \subseteq h.\text{Out}$ be a set of possible starting points. Let $H = \text{MD}[h, \text{Split}, S]$ be the hash function associated to these components via the MD transform of Fig. 2. Given an adversary $\mathcal{A}_H$, let $\mathcal{A}_h$ be the adversary of Fig. 15 using B_{mem} from the same figure.*

Then

$$\text{Adv}_H^{\text{cr}}(\mathcal{A}_H) \leq \text{Adv}_h^{\text{cr}}(\mathcal{A}_h). \quad (6)$$

The time complexity of $\mathcal{A}_h$ is the sum of the time complexities of $\mathcal{A}_H$ and B_{cr} . The memory complexity of $\mathcal{A}_h$ is the maximum of the memory complexity of $\mathcal{A}_H$ and the memory complexity of B_{mem} .

We make no value statement on how likely the improved memory usage of B_{mem} is to matter in practice. Our point is simply that it is more memory efficient and that these differences are hidden by proofs in which the explicit reduction algorithms are not provided.

The proof of the theorem mirrors that of Theorem 5.3, only requiring arguing that B_{mem} also correctly returns a collision whenever $\mathcal{A}_H$ does.

PROOF. (of Theorem 9.1) It is clear that the time and memory complexity of adversary $\mathcal{A}_h$ are as stated in the theorem.

Let $k \in h.\text{Keys}$, $s \in S$ be the values sampled when $\mathcal{A}_h$ is executed and $M_1, M_2 \in \text{Split}.\text{Inp}$ be the values returned by $\mathcal{A}_H$. We have $\text{Split}(M_1) \not\equiv \text{Split}(M_2)$ and $\text{Split}(M_2) \not\equiv \text{Split}(M_1)$, so if they form a collision for $H_{(k,s)}$, then they fulfill the conditions of Lemma 5.2 so B_{cr} would be guaranteed to return a collision for h_k . It is clear from examining the code that if B_{cr} finds a collision on any input, then B_{mem} will find a collision on the same input (though they might output different collisions). As an immediate result Equation (6) holds, completing the proof. $\square$

In Fig. 16 we present an analogous memory efficient algorithm $B_{\text{mem}2}$ and the corresponding $\mathcal{A}_h$ which would obtain the same sorts of memory savings for Theorem 6.4 and Theorem 7.2. For notational convenience, the presented pseudocode of $B_{\text{mem}2}$ uses vectors $\mathbf{c}_1$ and $\mathbf{c}_2$, but we note that they can easily be computed only as needed using constant memory. Furthermore, we use the

Algorithm $B_{\text{mem}}((k, s), M_1, M_2)$ $\mathbf{m}_1 \leftarrow \text{Split}(M_1); \mathbf{m}_2 \leftarrow \text{Split}(M_2); n_1 \leftarrow \mathbf{m}_1 ; n_2 \leftarrow \mathbf{m}_2$ $n \leftarrow \min(n_1, n_2)$ $c_1 \leftarrow s; c_2 \leftarrow s$ If $(n_1 > n_2)$ then For $i = 1, \dots, n_1 - n_2$ do $c_1 \leftarrow h_k((\mathbf{m}_1[i], c_1))$ If $(n_2 > n_1)$ then For $i = 1, \dots, n_2 - n_1$ do $c_2 \leftarrow h_k((\mathbf{m}_2[i], c_2))$ For $i = 1, \dots, n$ do $c'_1 \leftarrow h_k((\mathbf{m}_1[n_1 - n + i], c_1))$ $c'_2 \leftarrow h_k((\mathbf{m}_2[n_2 - n + i], c_2))$ If $(c'_1 = c'_2)$ and $(\mathbf{m}_1[n_1 - n + i], c_1) \neq (\mathbf{m}_2[n_2 - n + i], c_2)$ then Return $((\mathbf{m}_1[n_1 - n + i], c_1), (\mathbf{m}_2[n_2 - n + i], c_2))$ $c_1 \leftarrow c'_1; c_2 \leftarrow c'_2$ Return $\perp$	
Adversary $\mathcal{A}_h(k, \epsilon)$ $s \leftarrow \\$S; (M_1, M_2) \leftarrow \mathcal{A}_H((k, s), \epsilon)$ Return $B_{\text{mem}}((k, s), M_1, M_2)$	

Figure 15: Memory efficient algorithm B_{mem} and adversary $\mathcal{A}_h$ used for Theorem 9.1.

Algorithm $B_{\text{mem2}}((k, s), M_1, M_2)$ $\mathbf{m}_1 \leftarrow \text{Split}(M_1); \mathbf{m}_2 \leftarrow \text{Split}(M_2); n_1 \leftarrow \mathbf{m}_1 ; n_2 \leftarrow \mathbf{m}_2$ $c_1[1] \leftarrow s; c_2[1] \leftarrow s; n \leftarrow \min(n_1, n_2)$ If $(n_1 > n_2)$ then For $i = 1, \dots, n_1 - n_2$ do $c_1[i + 1] \leftarrow h_k((\mathbf{m}_1[i], c_1[i]))$ If $(n_2 > n_1)$ then For $i = 1, \dots, n_2 - n_1$ do $c_2[i + 1] \leftarrow h_k((\mathbf{m}_2[i], c_2[i]))$ For $i = 1, \dots, n$ do $\mathbf{m}_1 \leftarrow \mathbf{m}_1[n_1 - n + i]; c_1 \leftarrow c_1[n_1 - n + i]$ $\mathbf{m}_2 \leftarrow \mathbf{m}_2[n_2 - n + i]; c_2 \leftarrow c_2[n_2 - n + i]$ $c'_1 \leftarrow h_k((\mathbf{m}_1, c_1))$ $c'_2 \leftarrow h_k((\mathbf{m}_2, c_2))$ If $(c'_1 = c'_2)$ and $(\mathbf{m}_1, c_1) \neq (\mathbf{m}_2, c_2)$ then $a_1 \leftarrow (\mathbf{m}_1[n_1 - n + i - 1], c_1[n_1 - n + i - 1])$ $a_2 \leftarrow (\mathbf{m}_2[n_2 - n + i - 1], c_2[n_2 - n + i - 1])$ Return $((\mathbf{m}_1, c_1), (\mathbf{m}_2, c_2), a_1, a_2)$ $c_1[n_1 - n + i + 1] \leftarrow c'_1$ $c_2[n_2 - n + i + 1] \leftarrow c'_2$ Return $\perp$	
Adversary $\mathcal{A}_h(k, s)$ $(M_1, M_2) \leftarrow \mathcal{A}_H((k, s), \epsilon)$ Return $B_{\text{mem2}}((k, s), M_1, M_2)$	

Figure 16: Memory efficient algorithm B_{mem2} and adversary $\mathcal{A}_h$ to improve Theorem 6.4 and Theorem 7.2.

convention that out of bounds accesses to an array are accesses to its first element (this simplifies notation for the case that the initialization vector is part of the collision).

We note that ACFK [4] make the claim that collision-resistance is not a memory sensitive problem, saying for example “ t -collision-resistance is not memory sensitive for $t = 2$.” This seems to imply

that there is no reason to worry about memory tightness in our setting because we are doing a reduction to collision resistance. However, this statement is somewhat deceptive. When one unpacks this statement, the actual claim they are making is that the *best known generic attack* does not require much memory. This tells us nothing about whether there may exist better, not yet known, generic attacks or whether there exist better non-generic attacks against specific hash functions for which memory is a dominating factor.

We also observe that, as a community, there is much work to be done in this setting to determine how the memory usage of an adversary “should” be measured to best capture the reality. For example, in their work ACFK observe that many reductions in the random oracle model are highly inefficient in terms of memory complexity. They then show that, in some cases, a PRF can be used to make reductions more tight. However, the value of these points depends heavily on the fact that they adopt a convention of the memory used by the underlying game not counting towards the memory complexity of the adversary. When using our convention that the memory complexity of the adversary includes the memory used by the game in which it is executed (this way of measuring of memory complexity is referred to as **LocalMem** in their work), the value of this observation disappears. In this setting, the straightforward security reductions typically being done in the literature would *already* be memory-tight. By giving our reduction algorithms explicitly, we aim to make it easy for their memory complexity to be analyzed using whichever convention one desires.

10 CONCLUSION

This paper revisited the MD transform to unify prior work and variants, improve security guarantees and formalize folklore results. We introduced the RS security framework for hash functions with which we simultaneously capture several standard notions of security for hash functions and introduce our new notion of constrained collision resistance. Our new security notion allows us to understand ways in which an MD hash function can satisfy collision resistance despite collisions being known for its underlying compression function. In more detail, we have considered a parameterized MD transform that constructs a hash function $H = \text{MD}[h, \text{Split}, S]$ from a compression function h , splitting function Split , and set S of starting points. We have then comprehensively investigated what assumptions on h and Split guarantee collision resistance (CR) of H . We have shown that MD is better than advertised in the sense that conditions on h weaker than CR, formalized in our RS framework as constrained collision resistance (R_{ccrS}), suffice for H to be CR. This strengthens guarantees on hash functions and partially explains why, historically, attacks on compression functions have not immediately translated to attacks on the hash functions. The consequences are the usual benefits of weakening assumptions, namely that weaker compression functions are easier to design, harder to break and more likely to last. Furthermore, we have also shown how to speed up hashing by using very simple Split functions.

REFERENCES

- [1] Elena Andreeva, Bart Mennink, and Bart Preneel. 2011. Security Reductions of the Second Round SHA-3 Candidates. In *ISC 2010 (LNCS)*, Mike Burmester, Gene

- Tsudik, Spyros S. Magliveras, and Ivana Ilic (Eds.), Vol. 6531. Springer, Heidelberg, 39–53.
- [2] Elena Andreeva, Gregory Neven, Bart Preneel, and Thomas Shrimpton. 2007. Seven-Property-Preserving Iterated Hashing: ROX. In *ASIACRYPT 2007 (LNCS)*, Kaoru Kurosawa (Ed.), Vol. 4833. Springer, Heidelberg, 130–146.
 - [3] Elena Andreeva and Martijn Stam. 2011. The Symbiosis between Collision and Preimage Resistance. In *13th IMA International Conference on Cryptography and Coding (LNCS)*, Liqun Chen (Ed.), Vol. 7089. Springer, Heidelberg, 152–171.
 - [4] Benedikt Auerbach, David Cash, Manuel Fersich, and Eike Kiltz. 2017. Memory-Tight Reductions. In *CRYPTO 2017, Part 1 (LNCS)*, Jonathan Katz and Hovav Shacham (Eds.), Vol. 10401. Springer, Heidelberg, 101–132. https://doi.org/10.1007/978-3-319-63688-7_4
 - [5] Michael Backes, Gilles Barthe, Matthias Berg, Benjamin Grégoire, César Kunz, Malte Skoruppa, and Santiago Zanella Béguelin. 2012. Verified security of merkle-damgård. In *Computer Security Foundations Symposium (CSF), 2012 IEEE 25th. IEEE*, 354–368.
 - [6] Mihir Bellare. 2006. New Proofs for NMAC and HMAC: Security without Collision-Resistance. In *CRYPTO 2006 (LNCS)*, Cynthia Dwork (Ed.), Vol. 4117. Springer, Heidelberg, 602–619.
 - [7] Mihir Bellare, Daniel J. Bernstein, and Stefano Tessaro. 2016. Hash-Function Based PRFs: AMAC and Its Multi-User Security. In *EUROCRYPT 2016, Part 1 (LNCS)*, Marc Fischlin and Jean-Sébastien Coron (Eds.), Vol. 9665. Springer, Heidelberg, 566–595. https://doi.org/10.1007/978-3-662-49890-3_22
 - [8] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. 1996. Keying Hash Functions for Message Authentication. In *CRYPTO'96 (LNCS)*, Neal Koblitz (Ed.), Vol. 1109. Springer, Heidelberg, 1–15.
 - [9] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. 1996. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th FOCS. IEEE Computer Society Press*, 514–523.
 - [10] Mihir Bellare and Thomas Ristenpart. 2006. Multi-Property-Preserving Hash Domain Extension and the EMD Transform. In *ASIACRYPT 2006 (LNCS)*, Xuejia Lai and Kefei Chen (Eds.), Vol. 4284. Springer, Heidelberg, 299–314.
 - [11] Mihir Bellare and Phillip Rogaway. 2006. The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In *EUROCRYPT 2006 (LNCS)*, Serge Vaudenay (Ed.), Vol. 4004. Springer, Heidelberg, 409–426.
 - [12] John Black, Phillip Rogaway, and Thomas Shrimpton. 2002. Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV. In *CRYPTO 2002 (LNCS)*, Moti Yung (Ed.), Vol. 2442. Springer, Heidelberg, 320–335.
 - [13] Daniel R. L. Brown. 2002. Generic Groups, Collision Resistance, and ECDSA. Contributions to IEEE P1363a. (Feb. 2002). Updated version for “The Exact Security of ECDSA.” Available from <http://grouper.ieee.org/groups/1363/>.
 - [14] Ivan Damgård. 1988. Collision Free Hash Functions and Public Key Signature Schemes. In *EUROCRYPT'87 (LNCS)*, David Chaum and Wyn L. Price (Eds.), Vol. 304. Springer, Heidelberg, 203–216.
 - [15] Ivan Damgård. 1990. A Design Principle for Hash Functions. In *CRYPTO'89 (LNCS)*, Gilles Brassard (Ed.), Vol. 435. Springer, Heidelberg, 416–427.
 - [16] Hans Dobbertin. 1996. Cryptanalysis of MD5 Compress. (1996).
 - [17] Yevgeniy Dodis and Prashant Puniya. 2008. Getting the Best Out of Existing Hash Functions; or What if We Are Stuck with SHA?. In *ACNS 08 (LNCS)*, Steven M. Bellovin, Rosario Gennaro, Angelos D. Keromytis, and Moti Yung (Eds.), Vol. 5037. Springer, Heidelberg, 156–173.
 - [18] Yevgeniy Dodis, Thomas Ristenpart, and Thomas Shrimpton. 2009. Salvaging Merkle-Damgård for Practical Applications. In *EUROCRYPT 2009 (LNCS)*, Antoine Joux (Ed.), Vol. 5479. Springer, Heidelberg, 371–388.
 - [19] Peter Gaži, Krzysztof Pietrzak, and Michal Rybár. 2014. The Exact PRF-Security of NMAC and HMAC. In *CRYPTO 2014, Part 1 (LNCS)*, Juan A. Garay and Rosario Gennaro (Eds.), Vol. 8616. Springer, Heidelberg, 113–130. https://doi.org/10.1007/978-3-662-44371-2_7
 - [20] Jonathan Katz and Yehuda Lindell. 2014. *Introduction to modern cryptography*. CRC press.
 - [21] G. Laccetti and G. Schmid. 2004. On a Probabilistic Approach to the Security Analysis of Cryptographic Hash Functions. Cryptology ePrint Archive, Report 2004/324. (2004). <http://eprint.iacr.org/2004/324>.
 - [22] Alfred J Menezes, Paul C Van Oorschot, and Scott A Vanstone. 1996. *Handbook of applied cryptography*. CRC press.
 - [23] Ralph C. Merkle. 1990. A Fast Software One-Way Hash Function. *Journal of Cryptology* 3, 1 (1990), 43–58.
 - [24] Ralph C. Merkle. 1990. One Way Hash Functions and DES. In *CRYPTO'89 (LNCS)*, Gilles Brassard (Ed.), Vol. 435. Springer, Heidelberg, 428–446.
 - [25] NIST. August 2015. FIPS 180-4, Secure Hash Standard. (August 2015).
 - [26] Ronald Rivest. 2004. The MD5 message-digest algorithm, 1992. *RFC1321, Internet Engineering Task Force* (2004).
 - [27] Ronald L. Rivest. 1991. The MD4 Message Digest Algorithm. In *CRYPTO'90 (LNCS)*, Alfred J. Menezes and Scott A. Vanstone (Eds.), Vol. 537. Springer, Heidelberg, 303–311.
 - [28] Phillip Rogaway and Thomas Shrimpton. 2004. Cryptographic Hash-Function Basics: Definitions, Implications, and Separations for Preimage Resistance, Second-Preimage Resistance, and Collision Resistance. In *FSE 2004 (LNCS)*, Bimal K. Roy and Willi Meier (Eds.), Vol. 3017. Springer, Heidelberg, 371–388.
 - [29] Marc Stevens, Elie Bursztein, Pierre Karpman, Ange Albertini, and Yarik Markov. 2017. The First Collision for Full SHA-1. In *CRYPTO 2017, Part 1 (LNCS)*, Jonathan Katz and Hovav Shacham (Eds.), Vol. 10401. Springer, Heidelberg, 570–596. https://doi.org/10.1007/978-3-319-63688-7_19
 - [30] Marc Stevens, Pierre Karpman, and Thomas Peyrin. 2016. Freestart Collision for Full SHA-1. In *EUROCRYPT 2016, Part 1 (LNCS)*, Marc Fischlin and Jean-Sébastien Coron (Eds.), Vol. 9665. Springer, Heidelberg, 459–483. https://doi.org/10.1007/978-3-662-49890-3_18
 - [31] Douglas R Stinson. 2005. *Cryptography: theory and practice*. CRC press.
 - [32] Douglas R Stinson. 2006. Some observations on the theory of cryptographic hash functions. *Designs, Codes and Cryptography* 38, 2 (2006), 259–277.
 - [33] Xiaoyun Wang, Dengguo Feng, Xuejia Lai, and Hongbo Yu. 2004. Collisions for Hash Functions MD4, MD5, HAVAL-128 and RIPEMD. Cryptology ePrint Archive, Report 2004/199. (2004). <http://eprint.iacr.org/2004/199>.
 - [34] Xiaoyun Wang and Hongbo Yu. 2005. How to Break MD5 and Other Hash Functions. In *EUROCRYPT 2005 (LNCS)*, Ronald Cramer (Ed.), Vol. 3494. Springer, Heidelberg, 19–35.