# Time-Space Tradeoffs for Finding Multi-Collisions in Merkle-Damgård Hash Functions

**Akshima** ✉ 🏠
NYU Shanghai, China

## Abstract

We analyze the multi-collision resistance of Merkle-Damgård hash function construction in the auxiliary input random oracle model. Finding multi-collisions or $m$-way collisions, for some parameter $m$, in a hash function consists of $m$ distinct input that have the same output under the hash function. This is a natural generalization of the collision finding problem in hash functions, which is basically finding 2-way collisions. Hardness of finding collisions, or collision resistance, is an important security assumption in cryptography. While the time-space trade-offs for collision resistance of hash functions has received considerable attention, this is the first work that studies time-space trade-offs for the multi-collision resistance property of hash functions based on the popular and widely used Merkle-Damgård (MD) constructions.

In this work, we study how the advantage of finding $m$-way collisions depends on the parameter $m$. We believe understanding whether multi-collision resistance is a strictly easier property than collision resistance is a fundamental problem and our work facilitates this for adversaries with auxiliary information against MD based hash functions. Furthermore, in this work we study how the advantage varies with the bound on length of the $m$ colliding inputs. Prior works [1, 19, 3] have shown that finding "longer" collisions with auxiliary input in MD based hash functions becomes easier. More precisely, the advantage of finding collisions linearly depends on the bound on the length of colliding inputs. In this work, we show similar dependence for $m$-way collision finding, for any $m \geq 2$.

We show a simple attack for finding 1-block $m$-way collisions which achieves an advantage of $\tilde{\Omega}(S/mN)$. For $2 \leq B < \log m$, we give the best known attack for finding $B$-blocks $m$-way collision which achieves an advantage of $\tilde{\Omega}(ST/m^{1/(B-1)}N)$ when $m^{1/(B-1)}$-way collisions exist on every salt. For $B > \log m$, our attack achieves an advantage of $\tilde{\Omega}(STB/N)$ which is optimal when $SB \geq T$ and $ST^2 \leq N$. The main results of this work is showing that our attacks are optimal for $B = 1$ and $B = 2$. This implies that in the auxiliary-input random oracle model, the advantage decreases by a multiplicative factor of $m$ for finding 1-block and 2-block $m$-way collisions (compared to collision finding) in Merkle-Damgård based hash functions.

**2012 ACM Subject Classification** Security and privacy → Cryptography; Security and privacy → Information-theoretic techniques; Mathematics of computing → Probability and statistics

**Keywords and phrases** Collision, hash functions, multi-collisions, Merkle-Damgård, pre-computation, auxiliary input

**Digital Object Identifier** 10.4230/LIPIcs.ITC.2024.9

**Funding** Most of this work was done while Akshima was a PhD student at University of Chicago and was supported in part by NSF Grant No. 1928767. Akshima is supported by Ministry of Education of P.R.C., through an Overseas Postdoc Special Program (2023), National Natural Science Foundation of China Grant No. 6210226 and NYTP Grant No. 2012120. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding institutions or the governments.

**Acknowledgements** We thank the anonymous reviewers for their constructive comments on earlier drafts of the work. We thank David Cash and Siyao Guo for useful discussions.

© Akshima;
licensed under Creative Commons License CC-BY 4.0
5th Conference on Information-Theoretic Cryptography (ITC 2024).
Editor: Divesh Aggarwal; Article No. 9; pp. 9:1–9:22
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

# 1 Introduction

In this work we study multi-collision finding in Merkle-Damgård (MD) hash functions with auxiliary input generated during a pre-computation. Several recent works [11, 12, 1, 19, 3] have rigorously studied collision finding in MD hash functions with pre-computation but currently nothing is known about the upper and lower bounds for $m$-way collision finding with pre-computation, where an $m$-way collision consists of $m$ distinct messages, for some parameter $m$, whose hash values are identical.

### Auxiliary-Input Random Oracle Model

We will study the problem in the auxiliary-input(AI) random oracle (RO) model in order to obtain the time-space trade-offs. We note that AI-RO model is strong enough to capture pre-computing as well as non-uniform adversaries. We informally describe the model next.

An adversary $\mathcal{A}$ in this model can be thought of as a pair of algorithms $(\mathcal{A}_1, \mathcal{A}_2)$ each of which corresponds to a stage of $\mathcal{A}$:

1. Offline/pre-computation stage: $\mathcal{A}_1$ gets computationally unbounded access to the oracle (i.e., there is no bound on the number of oracle queries and computation made by $\mathcal{A}_1$) and outputs a fixed-size advice/information about the oracle. Let's denote this pre-computed information by $\sigma$.
2. $\mathcal{A}_2$ takes $\sigma$ outputted by $\mathcal{A}_1$ as input along with challenge, makes a bounded number of queries to the oracle.

In this work, we will use $S$ to denote the length of $\sigma$ in bits output by $\mathcal{A}_1$ and $T$ to denote the bound on the number of queries made by $\mathcal{A}_2$. We will refer to such an adversary $\mathcal{A}$ as $(S, T)$-AI adversary.

Specifically for a $m$-way collision finding $(S, T)$-AI adversary, the input challenge to $\mathcal{A}_2$ consists of a salt for the hash function and $\mathcal{A}_2$ needs to output $m$ colliding messages on the salt to win. There could be an additional restriction that the collisions found are at most $B$ blocks long.

### Why Study Multi-collisions

Multi-collision resistance is a natural generalization of a fundamental and widely used security property in cryptography, namely collision resistance. Multi-collision resistance property is seen as a relaxation of the collision resistance and it is a meaningful exercise to understand whether former is strictly an easier property than the standard collision resistance. In a recent work [29], Rothblum and Vasudevan studied this and gave a non-black box and non-constructive transformation from a multi-collision resistant hash function to collision resistant hash function. In this work, we understand how the complexity of the $m$-way collision finding in MD hash functions varies with $m$ for any $(S, T)$-AI adversary. If the complexity grows linearly or exponentially in $m$, then it could be better to reduce the security of protocols to $m$-way collisions for larger $m$'s. Reducing the security to $m$- way collision resistance for larger and larger $m$ could allow applications to reduce the length of the output of the hash function and hence make them more efficient.

For a random hash function (i.e., a hash function modelled as a random oracle) it was proven that any adversary making $T$ queries can achieve the advantage of the order $(T/m)^m/N^{m-1}$ for finding $m$-way collision. Note that the advantage decreases exponentially in the size of $m$. Works dating as early as 1990s had identified that finding $m$-way collisions gets harder as $m$ becomes larger and inspired by this observation, the notion of Multi-collision resistance in hash functions had been used in several existing constructions of applications such as digital signatures, commitment scheme and more.

The multi-collision resistance primitive has been used in signature schemes [8], identification schemes [20], micropayment scheme [28] and commitment scheme [26]. Bitansky et al. in [7], formally introduced multi-collision resistance as a cryptographic primitive, used it to construct secure keyless hash functions and study its applications. Berman et al., too, in [5] used the notion of multi-collision resistance to find secure keyless hash functions.

Liu and Zhandry in [27] analyzed multi-collisions in the quantum setting. As Grover's algorithm speeds up the classical birthday attack for finding collisions from $O(2^{n/2})$ to $O(2^{n/3})$, Liu and Zhandry study the speedup on multi-collisions in random hash function with quantum query. They conclude that it could be better to base the security of constructions on the multi-collision resistance in the quantum setting as they prove that $\Omega(2^{n/2\cdot(1-1/(2^m-1))})$ queries are required on average to find $m$-way collision for any constant $m$.

While it can be provably shown for random functions that finding $m$-way collisions becomes exponentially unlikely in $m$ (the probability of finding $m$-way collisions is $\theta(T^m/N^{m-1})$), same cannot be said to hold true for iterated hash function constructions such as Merkle-Damgård. In fact Joux in [25], showed an attack that requires just $\log m \cdot \sqrt{N}$ queries in expectation to find $m$-way collision in Merkle-Damgård based hashing algorithms. Note that this means that the number of queries required to find $m$-way collisions in Merkle-Damgård based hashing algorithms increases by just a multiplicative factor of $\log m$ with $m$ (unlike for random functions).

Here we note that Joux's result does not take into account two things: 1) What if there is a restriction on the length of permitted collisions? 2) What if adversaries allowed to perform pre-computation or learn some auxiliary information about the hash function? To the best of our knowledge, our work is the first to investigate on both of these fronts simultaneously.

From our investigation of the problem with pre-computation, we learn that finding $m$-way collisions in MD based hash functions do not get harder in the Auxiliary input random oracle (AI-RO) model as far as there is no restriction on the length of $m$-way collisions found. But what about when there is a restriction on the length of the $m$-way collisions found with pre-computation? The biggest question this work investigates is about the hardness of finding "short" $m$-way collisions. In fact, for some restricted parameter ranges we manage to show that any pre-computing adversary's advantage provably takes a linear hit as $m$ grows larger.

It is worth-noting that $m$-way collision finding is somewhat related to the fundamental $m$-distinctness problem. The difference is that our problem studies collisions on the same salt whereas $m$-distinctness applies even when there is an $m$-way collision with different salts. Nevertheless both the problems are looking for $m$-way collisions. In fact, solving $m$-distinctness is a trivial problem in the AI-RO model as long as the advice $\sigma$ is $\Omega(m \log N)$ bits long. This is because $\mathcal{A}_1$ can simply find the $m$-way collision on the hash function and store it as advice for $\mathcal{A}_2$ to win certainly. Thus, it is important that $\mathcal{A}_2$ gets a random salt as challenge and the outputted $m$ messages should collide with this salt under the hash function for $\mathcal{A}$ to win.

## 1.1 Our Contributions

### 1.1.1 Our Results

First, we summarize our results for $m$-way collision-finding in table 1. As stated before, let $S$ denote the size of pre-computed information in bits, $T$ be the number of the queries made to oracle implementing a random function $h$ in $[N] \times [M] \to [N]$ and $B$ be the number of blocks accepted in the multi-collisions.

**Table 1** Asymptotic security bounds on the security of finding $B$-block $m$-way collisions in Merkle-Damgård Hash Functions constructed from a random function $h : [N] \times [M] \mapsto [N]$ against $(S, T)$-algorithms. The attacks assume $M$ is "sufficiently" larger than $N$ for the required collisions to exist in $h$. For simplicity, logarithmic terms and constant factors are omitted.

| $B$ | Best attacks | Security bounds |
|---|---|---|
| $B = 1$ | $\frac{S}{mN} + \frac{(T/m)^m}{N^{m-1}}$ [Thm 10] | $\frac{S}{mN} + \frac{(2T/m)^m}{N^{m-1}}$ [Thm 10] |
| $B = 2$ | $\frac{ST}{mN} + \frac{(T/m)^m}{N^{m-1}}$ [Thm 8] | $\frac{ST}{mN} + \frac{(eT/m)^m}{N^{m-1}}$ [Thm 11] |
| $B < \log m$ | $\frac{ST}{m^{1/(B-1)}N} + \frac{(T/m)^m}{N^{m-1}}$ [Thm 8] | $\frac{STB}{N} \cdot \max\{1, \frac{ST^2}{N}\} + \frac{T^2}{N}$ [3] |
| $[\log m, 2\log m)$ | $\frac{ST}{N} + \left(\frac{T^2}{N \log^2 m}\right)^{\log m}$ [Thm 8] | $\frac{STB}{N} \cdot \max\{1, \frac{ST^2}{N}\} + \frac{T^2}{N}$ [3] |
| $[2\log m, T]$ | $\frac{STB}{N} + \left(\frac{T^2}{N \log^2 m}\right)^{\log m}$ [Thm 8] | $\frac{STB}{N} \cdot \max\{1, \frac{ST^2}{N}\} + \frac{T^2}{N}$ [3] |
| Unbounded | $\frac{ST^2}{N}$ [Thm 7] | $\frac{ST^2}{N}$ [11] |

We elaborate our results next.

1. For unbounded $B$ (or $B \geq T$), a reduction to the collision-finding problem upper bounds the advantage to $\tilde{O}(ST^2/N)$ using a result from [12]. Our interesting finding is a matching attack based on the attacks of Joux in [25] and Coretti et al. in [12]. Finding an attack that matches this trivial security bound up to poly-log factors is surprising because this attack shows that finding $m$-way collision in the AI model requires about the same effort for any $m \geq 2$. To put this into more perspective, this attack shows that an AI adversary can find an $\sqrt{N}$-way collision in about the same effort as required for finding a 2-way collision.
2. For any other $B$, the best we manage to do is again bound the security via reduction to the collision-finding problem as for unbounded $B$. The bounded-length version of our attack from (1) matches the bound when $B \geq \log m, SB \geq T$ and $ST^2 \leq N$.
   For $B < \log m$, the best attack we can find has a gap of a factor of $m^{1/(B-1)}$. To elaborate, we give attack that achieves:
   - advantage $\tilde{\Omega}\left(\frac{ST}{m^{1/(B-1)}N}\right)$ for any $B > 1$ when $M > (m^{1/(B-1)} - 1) \cdot N$ and $m = O(N^c)$ for some constant $c$.
     We restrict to $M > (m^{1/(B-1)} - 1) \cdot N$, so that an $m^{1/(B-1)}$-way collision is guaranteed to exist on every salt in $[N]$ under $h$ by the pigeonhole principle. However, $M > (m^{1/(B-1)} - 1) \cdot N$ is not a necessary condition for such collisions to exist. It is the case that our attack succeeds with advantage $\tilde{\Omega}\left(\frac{ST}{m^{1/(B-1)}N}\right)$ when $m^{1/(B-1)}$-way collisions exist on every salt in $[N]$.
   - advantage $\tilde{\Omega}\left(\frac{STB}{N}\right)$ when $B \geq (d+1)\log m$ and $M^d > N$ for some $d \geq 1$.
     Let's first understand this for $d = 1$. It means we restrict to $M > N$ but again this is not a necessary condition. Having 2-way collisions on every salt is sufficient for our attack to achieve $\tilde{\Omega}\left(\frac{STB}{N}\right)$ advantage. To get rid of the $M > N$ requirement, we can set $d$ to be the smallest value such that $M^d > N$ and replace $h$ with $h^d$. (Hence, $B \geq (d+1)\log m$.) This is so that the offline adversary will be able to find 2-way collisions on every salt under $h^d$ by the pigeonhole principle. Note that the smallest $d$ which satisfies this requirement will be at most $\log N + 1$ for any $M \geq 2$.

3. For $B = 1$, we show that advantage for any $(S, T)$-AI adversary is bounded by $\tilde{O}\left(\frac{S}{mN} + \frac{T^m}{N^{m-1}}\right)$. Note that this bound shows a loss of factor of $m$ in advantage when $S/mN$ is the dominating term. The trivial attack stated in section 5 matches this bound. The proof for this result is via generalizing a technique used in [16], namely "global" compression.
4. Our result for $B = 2$ is the most important contribution of this work in terms of techniques used. We show that no $(S, T)$-AI adversary can achieve better advantage than $\tilde{O}\left(\frac{ST}{mN}\right)$. This bound again shows a loss of factor of $m$ in advantage and matches our attack from (2).
   Several prior works have used the relation between the advantage of solving a problem in the AI model and another model, namely the multi-instance model (MI), so that they analyze the MI-security and obtain bounds for the AI-security. However, we reduce the problem to finding $m$-way collision in a variant of the "parallel" multi-instance model (MI). It is the way we relate the security in the two models which is new and we believe could be of independent interest.

Our results suggest that finding $m$-way collisions doesn't get harder (as $m$ increases) in the AI-RO model if the collisions are allowed to be sufficiently long or $m$ is really "small" (for instance $O(\log^c N)$ for some constant $c$ because we ignore poly-log factors in our results). From the practice standpoint, this could mean that if the constructors of applications want to gain in complexity by reducing their security to larger collision finding, it is imperative to make sure the relative size of the parameters $N, m$ and $B$ in their application actually makes that plausible.

## Discussion

As in several prior works [24, 1, 10, 19, 3], we also use that the security in the auxiliary input RO model is closely related to the security in another model that is easier to analyze, namely the Multi-instance (MI) model. How we relate these two models in this work is different from the prior works. This allows our analysis of the Multi-instance model to be significantly different and easier compared to the prior works.

We informally describe the "Parallel" MI model next and follow up with a high level idea of our technique. We note that there is a "Sequential" variation of the MI model that has been extensively used in several prior works but this work only uses the Parallel variation. So whenever we refer to MI model in this paper, we mean the Parallel MI model.

Informally, an adversary in the MI model gets multiple instances of challenges and gets to make bounded number of queries to the oracle for each challenge instance. The adversary gets no advice and it is required to succeed on every challenge instance in order to succeed. Specifically, our parameters in the MI model will be $S, T$ where $S$ will denote the number of challenge instances given to the adversary (note the parameter $S$ here corresponds to the size the advice from AI model for optimal reduction) and $T$ will denote the number of queries the adversary can make for each instance. Note that the adversary will get to make a total of $ST$ queries.

Our high level approach would be to identify all the types of 2-block $m$-way collisions. For 2-block $m$-way collision finding, there are several types of collisions which an adversary could find in order to win. We depict the types in fig. 4. (Note that an adversary could find a collision that is a mix of these types but there will exist an $(c \cdot m)$-way collision of one type where $c$ is a constant.) For a straight-forward AI to MI security reduction used in prior works, we need to analyze adversaries that could find any one of these five types of $m$-way collisions for each of the $S$ challenge salts in MI model and win. However, analyzing such adversaries seems very hard.

We will relate the AI-security to the MI-security for each of these type of collisions. This allows us to analyze adversary restricted to finding a certain type of collision for all the $S$ challenge instances in the MI model. This is the idea that makes the analysis in the MI model possible. We would like to note here that we are able to do per collision type of relation of AI-security to MI-security because our analysis in the MI model is insensitive to the order in which the blocks of colliding messages are found. For full details refer to section 6.

## 1.2 Related Works

We would like to first note that time-space lower bounds have been studied for several cryptographic primitives including collision resistance. Some other primitives are function inversion, discrete log, one-way functions, pseudo-random generators. Refer [11, 15, 9, 14, 13, 22, 21] for further details.

Here we will focus on the prior works for collision resistance in the auxiliary input model. As far as we know, all the prior works have focused on time-space trade-offs for 2-way collisions. Ours is the first work that studies the more general primitive, $m$-way collisions in hash functions.

Before discussing the prior works, we recall our notations. The adversary in the AI model works in two stages. The output of the first (offline) stage adversary is bounded length advice $\sigma$. We will denote this length by $S$-bit, i.e., $|\sigma| = S$. The adversary in the second (online) stage makes bounded number of queries to the oracle, denoted by $T$. The adversary could be required to output collisions that have bounded length. We will denote this bound on the length by $B$.

Dodis et al. in [16] were the first to study 2-way collision-resistance in AI-RO model. Dodis et al. studied it for random functions and proved the security bound of $\tilde{\theta}(S/N + T^2/N)$. They presented an attack and proved the matching security bound via "global" compression. To elaborate further, the authors showed that any pre-computing adversary that finds collisions on "too many" salts can be used to compress the random function. Since random functions cannot be compressed, this can be used to bound the size of the set of winning salts for any adversary. Our proof of security bound for 1-block $m$-way collisions is based on this technique of Dodis et al..

Coretti et al. in [12] were the first to study (2-way) collision-resistance for Merkle-Damgård based hash function in the AI-RO model. Coretti et al. proved a bound of $O(ST^2/N)$ on finding collisions when there is no restriction on the length of the collisions. Coretti et al. reduced the security of collision finding in the AI-RO model to another model, namely Bit-fixing random oracle (BF-RO) model which was inspired by the work of Unruh [30]. The matching attack presented in [12] was inspired by the Hellman's attack presented in the seminal work [23] and found collisions that could be order of $T$ blocks long.

The follow-up work of Akshima et al. [1] realized that when accepted collisions were restricted to $B$-blocks in length, the best attack they could find achieved an advantage of $\Omega(STB/N)$ and they conjectured it to be the optimal attack. However, Akshima et al. could prove the optimality of their attack only for a restricted class of pre-computing adversaries and not any pre-computing adversary. For any pre-computing adversary, they could show their attack to be optimal only when $B = 2$. The proof in [1] reduced to Sequential multi-instance game and proved the required bound via compression. In addition, Akshima et al. presented an attack that showed it was impossible to achieve the desired bound (bound that would give optimal bound in AI-RO model) in the parallel multi-instance game, effectively proving a gap between the two versions of the multi-instance game for collision-finding.

The follow-up work of Ghoshal and Komargodski [19] proved the conjecture in [1] for any constant $B$. Ghoshal et al. used similar techniques to those in [1] together with their observation that it is "unlikely" to find $\geq \log N$-way collision in a random function to obtain better results.

Inspired by the idea of [10] to analyze the sequential multi-instance game stage-wise instead of simultaneously analyzing all the $S$ stages, Akshima et al. in [3] proved the conjecture of [1] for any $B \in [2, T)$ when $ST^2 \leq N$. They proved a bound of $STB/N \cdot \max\{1, ST^2/N\}$. Our bounds for "long" $m$-way collisions is based on this result. Our attack for $B \geq \log m$ blocks long $m$-way collision, as in [3], matches the bound (up to poly-log factors) when $ST^2 \leq N$.

Freitag et al. in [17] studied 2-way collision-finding for Sponge based hash functions in the AI random permutation model. They presented an attack for $B = 1$ which uses inverse queries of permutation to beat the trivial attack for some parameter ranges. They also prove loose security bounds for $B = 1$ and $B = 2$ block collision finding. Their bound for $B = 1$ was improved in the follow-up work [2]. In [2], Akshima et al. also showed that relation between AI security and multi-instance security cannot be used to improve the security bounds any further for Sponge based constructions.

A recent work [18] by Freitag et al. uses Merkle trees to get a provably improved hash function construction that gives optimal collision resistance against adversaries with pre-computation.

### 1.3 Open Problems

1. An obvious open question is proving a tighter security bound on $B$-block $m$-way collision finding for $B < \log m$ or alternatively finding a better attack. We conjecture that the attack we have presented in this work for that parameter range is optimal.
2. We are aware of few works that have studied multi-collisions in sponge constructions without pre-processing. In the famous work [6], Bertoni et al. presented the sponge construction and briefly talked about realizing multi-collisions in sponges. Another work [4] from 2013 by AlAhmad, Alshaikhli and Nandi formally studied Joux's attack for sponges.

   As far as we are aware there is no work on security bounds of multi-collision resistance with pre-processing for Sponge constructions. As Sponge constructions use a permutation instead of a function, our lower bounds do not extend trivially to Sponge (however, our attacks do) and it is an interesting problem to consider for future work.

## 2 Preliminaries

### 2.1 Notation

For non-negative integers $N, k$, $[N] := \{1, \ldots, N\}$ and $\binom{[N]}{k}$ denotes the set of all $k$-sized subsets of $[N]$. For non-negative integers $a, b$ such that $a < b$, $(a, b)$ is the set of values between $a$ and $b$ excluding $a$ and $b$ themselves, i.e., $(a, b) := \{a+1, a+2, \ldots, b-1\}$. For a set $X$, $x \leftarrow_\$ X$ denotes $x$ is a uniform random variable on $X$. $X^+$ denotes a list of one or more elements from $X$.

$O$, $\Omega$ and $\Theta$ are the standard asymptotic notations. The asymptotic notations with $\tilde{}$, e.g. $\tilde{O}$, ignore the poly-logarithmic terms.

## 2.2 Merkle Damgård Hash Function

A cryptographic hash function is a function that maps inputs of varied length to a fixed length output. Merkle Damgård (MD) is one of the popular classical construction for hash functions. It uses fixed length compression functions. MD5, SHA1, SHA2 and many more standard hashing algorithms are based on this construction.

In this work, we study an abstraction of the MD construction that takes a salt and an arbitrary length message as input, breaks the message into blocks of fixed length and iteratively applies the compression function, denoted $h$ and modelled as a random oracle. For the remaining of the paper, unless specified otherwise, we will always think of $h$ as a random function from $[N] \times [M] \to [N]$.

For any salt $a \in [N]$ and message $m \in [M]$, we define $MD_h(a, m) = h(a, m)$. For message $m \in [M]^+$ such that $m$ can be written as $m = m_1 || \dots || m_\ell$ where $m_1, \dots, m_\ell \in [M]$, we define $MD_h$ on $a$ and $m$ as follows:

$$MD_h(a, m) := h(MD_h(a, m_1 || \dots || m_{\ell-1}), m_\ell).$$

We refer to $m_1, \dots, m_\ell$ as blocks.

## 2.3 Definitions

Next, we formally define the $m$-way multi-collision resistance game for MD hash functions in the auxiliary input (AI) RO model.

▶ **Definition 1** ((S, T, m)-AI adversary). *A pair of algorithms $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is an $(S, T, m)$-AI adversary against $m$-way collision resistance in $MD_h$ if*

- *$\mathcal{A}_1^h$ gets unbounded access to $h$ (i.e., gets to make unbounded number of oracle queries to $h$) and outputs $S$ bits of advice $\sigma$;*
- *$\mathcal{A}_2^h$ takes $\sigma$ and a salt $a \in [N]$ as input, issues $T$ queries to $h$ and outputs $\mathsf{msg}_1, \dots, \mathsf{msg}_m$*

*where $h$ is a function in $[N] \times [M] \to [N]$.*

The $m$-way multi-collision resistance security game, namely m-AICR, is defined next.

▶ **Definition 2.** *For a function $h$ in $[N] \times [M] \to [N]$ and salt $a \in [N]$, the game* m-AICR *is defined in fig. 1.*

```
Game $\text{m-AICR}_{h,a}(\mathcal{A})$
    $\sigma \leftarrow \mathcal{A}_1^h$
    $\mathsf{msg}_1, \dots, \mathsf{msg}_m \leftarrow \mathcal{A}_2^h(\sigma, a)$
    If $\mathsf{msg}_i \neq \mathsf{msg}_j$ and $\mathrm{MD}_h(a, \mathsf{msg}_i) = \mathrm{MD}_h(a, \mathsf{msg}_j)$  $\forall i \neq j \in [m]$
        Then Return 1
    Else Return 0
```

**Figure 1** Security game $\text{m-AICR}_{h,a}(\mathcal{A})$.

*For an $(S, T, m)$-AI adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, the advantage of $\mathcal{A}$, denoted as $\mathsf{Adv}^{\text{m-AICR}}(\mathcal{A})$, is defined as the probability of $\text{m-AICR}_{h,a}(\mathcal{A})$ returning 1 when $h$ is a random function in $[N] \times [M] \to [N]$ and $a$ is a random salt in $[N]$ drawn independently. We define the $(S, T, m)$-AI multi-collision resistance, denoted by $\mathsf{Adv}^{\text{m-AICR}}(S, T, m)$, as the maximum advantage taken over all $(S, T, m)$-AI adversaries in the game* m-AICR.

```
Game $\mathsf{m\text{-}AICR}^{B}_{h,a}(\mathcal{A})$
    $\sigma \leftarrow \mathcal{A}_1^h$
    $\mathsf{msg}_1, \ldots, \mathsf{msg}_m \leftarrow \mathcal{A}_2^h(\sigma, a)$
    If any of $\mathsf{msg}_1, \ldots, \mathsf{msg}_m$ have more than $B$ blocks
        Then Return 0
    If $\mathsf{msg}_i \neq \mathsf{msg}_j$ and $\mathrm{MD}_h(a, \mathsf{msg}_i) = \mathrm{MD}_h(a, \mathsf{msg}_j)$ $\forall i \neq j \in [m]$
        Then Return 1
    Else Return 0
```

**Figure 2** Security game $\mathsf{m\text{-}AICR}^{B}_{h,a}(\mathcal{A})$.

▶ **Definition 3.** *For a function $h$ in $[N] \times [M] \to [N]$ and salt $a \in [N]$, the game $\mathsf{m\text{-}AICR}^B$ is defined in fig. 2.*

*For an $(S, T, m)$-AI adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, the advantage of $\mathcal{A}$ in the game $\mathsf{m\text{-}AICR}^B$, denoted as $\mathsf{Adv}^{\text{m-AICR}}_{\mathsf{B}}(\mathcal{A})$, is defined as the probability of $\mathsf{m\text{-}AICR}^{B}_{h,a}(\mathcal{A})$ returning 1 when $h$ is a random function in $[N] \times [M] \to [N]$ and $a$ is a random salt in $[N]$ drawn independently. We define the $(S, T, m)$-AI $B$-length multi-collision resistance, denoted by $\mathsf{Adv}^{\text{m-AICR}}_{\mathsf{B}}(S, T, m)$, as the maximum advantage taken over all $(S, T, m)$-AI adversaries in the game $\mathsf{m\text{-}AICR}^B$.*

## 2.4 Relevant Results

### Compression arguments

▶ **Lemma 4** ([15], restated in [16])**.** *Say $\mathsf{Enc}$ is an encoding function in $\{0,1\}^{\ell_1} \to \{0,1\}^{\ell_2}$ and $\mathsf{Dec}$ is a decoding function in $\{0,1\}^{\ell_2} \to \{0,1\}^{\ell_1}$ such that for any $\ell_1$-bit $x$, $\mathsf{Dec}(\mathsf{Enc}(x)) = x$ with probability at least $\epsilon$, then $\ell_2 \geq \ell_1 - \log(1/\epsilon)$ holds true.*

### 2-way collision results

▶ **Lemma 5** (from [12])**.** *For any $S, T, N$*

$$\mathsf{Adv}^{\text{2-AICR}}(S, T, 2) = \tilde{O}\left(\frac{ST^2}{N}\right).$$

▶ **Lemma 6** (from [3])**.** *For any $S, T, N$ and $2 < B < T$*

$$\mathsf{Adv}^{\text{2-AICR}}_{B}(S, T, 2) = \tilde{O}\left(\frac{STB}{N} \cdot \max\left\{1, \frac{ST^2}{N}\right\}\right).$$

# 3 Unbounded Length Multi-Collisions

▶ **Theorem 7.** *For any $S, T, m, N$ and $M$ such that $m = O(N^c)$ for some constant $c$,*

$$\mathsf{Adv}^{\text{m-AICR}}(S, T, m) = \tilde{\Theta}\left(\frac{ST^2}{N}\right).$$

**Proof.** Note that the security bound is unsurprising and follows from lemma 5 via trivial reduction. We present the reduction in the full version of the paper for completeness.

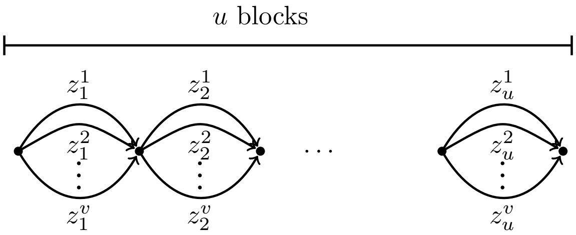

**Figure 3** $u$-length chain of $v$-way collisions.

## 3.1 Matching Attack

Next we present the matching attack achieving the bound in the theorem. This multi-collision finding attack is based on the attack of Joux in [25] and the pre-computing attacks given in [23, 12, 1]. We assume $M > N$ and thus 2-way collisions exist for every salt in $[N]$ under $h$ (that is because even if $M = N + 1$ there should exist a 2-way collision on every salt by pigeonhole principle). Note that it is possible to get rid of this assumption by replacing $h$ with $h^c$ for some $c$ such that $M^c > N$. Then our $(S, T, m)$- AI adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is as follows:

1. In *Offline stage*, $\mathcal{A}_1$ picks $s = S/(3 \log m \cdot \log M)$ random salts, denoted $\{a_1, \ldots, a_s\}$ and iteratively computes $a_j^i = h(a_j^{i-1}, 0)$ for $i \in [T/2]$ and $j \in [s]$ to obtain, say $\{a'_1, \ldots, a'_s\}$ set of salts. Then for each of the salt $a'_i$ in this set, $\mathcal{A}_1$ finds a $(\log m)$-length chain of 2-way collisions (refer fig. 3 for a pictorial depiction of $u$-length chain of $v$-way collisions) and store these along with $a'_i$ as advice.
2. In *Online stage*, $\mathcal{A}_2$ takes the advice output by $\mathcal{A}_1$ and random salt $a$ and iteratively queries $h(*, 0)$ up to $T$ times. If output of any of these queries is some $a' \in \{a'_1, \ldots, a'_s\}$, then $\mathcal{A}_2$ succeeds in outputting an $m$-way collision.

Let $\mathcal{G}$ be the subset of salts that are output of some $h(*, 0)$ query made in step 1 (Offline stage). We can show that $\mathbb{E}[|\mathcal{G}|] = \tilde{\Omega}(ST)$ in exactly the same manner as in lemma 14 of [12].

Then,

$$\mathsf{Adv}^{\text{m-AICR}}(\mathcal{A}) = \tilde{\Omega}\left(\frac{ST^2}{N}\right).$$

using that the adversary wins if output of any of its first $T/2$ queries in the online stage is in $\mathcal{G}$ and $\mathbb{E}[|\mathcal{G}|] = \tilde{\Omega}(ST)$. ◀

# 4 $B$-Block Multi-Collisions

▶ **Theorem 8.** *For any $S, T, m, B, N$ and $M$such that $m = O(N^c)$ for some constant $c$ and $B > 1$,*

$$\mathsf{Adv}_{\mathsf{B}}^{\text{m-AICR}}(S, T, m) = \tilde{\Omega}\left(\frac{ST}{m^{1/(B-1)}N}\right).$$

*when $M > (m^{1/(B-1)} - 1) \cdot N$ and*

$$\mathsf{Adv}_{\mathsf{B}}^{\text{m-AICR}}(S, T, m) = \tilde{\Omega}\left(\frac{STB}{N}\right)$$

*when $B \geq (d + 1) \log m$ and $M^d > N$ for the some $d \geq 1$.*

**Proof.** We give an attack that achieves the bound in the theorem. In the attack and its analysis we assume $d = 1$ for simplicity. However, it is straightforward to extend it for other $d$.

We also assume $m = O(N^c)$ for some constant $c$. Then our proposed $(S, T, m)$-AI adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is a variation of our attack proposed in the proof of theorem 7. Refer to Appendix A for more details. ◀

▶ **Theorem 9.** *For any $S, T, m, N$ and $M$*

$$\mathsf{Adv}_{\mathsf{B}}^{\text{m-AICR}}(S, T, m) = \tilde{O}\left(\max\left\{1, \frac{ST^2}{N}\right\} \cdot \frac{STB}{N} + \frac{T^2}{N}\right).$$

**Proof.** The proof of the theorem follows from the reduction of 2-way collision finding to $m$-way collision finding and lemma 6. ◀

Note that there is a gap between the best known attack and the bounds proved in Thm. 9 for several parameter ranges. In the following sections, we prove optimal security bounds for some of these parameter ranges, specifically $B = 1$ and $B = 2$ block $m$-way collisions.

## 5 1-Block Multi-Collisions

▶ **Theorem 10.** *For any $S, T, m, N, M$ and $B = 1$,*

$$\mathsf{Adv}_{\mathsf{B}}^{\text{m-AICR}}(S, T, m) = \tilde{\Theta}\left(\frac{S}{mN} + \frac{(T/m)^m}{N^{m-1}}\right).$$

We refer the readers to Appendix B for the proof.

## 6 2-Block Multi-Collisions

▶ **Theorem 11.** *For any $S, T, m, N, M$ and $B = 2$ such that $m \leq T$,*

$$\mathsf{Adv}_{\mathsf{B}}^{\text{m-AICR}}(S, T, m) = \tilde{O}\left(\frac{ST}{mN} + \frac{(eT/m)^m}{N^{m-1}}\right).$$

**Proof.** Let's fix an $(S, T, m)$-AI adversary $\mathcal{A}$. Next, we identify all the types of 2-block $m$-way collisions. See fig. 4 for a pictorial depiction of the types. Note that collision type 5 in fig. 4e can be thought of as a generalization of collision types in fig. 4cand fig. 4d but we treat them separately for simplicity.

Moreover it is possible that an adversary finds an $m$-way collision which is a mix of two types of collisions given in fig. 4. For some $\ell$ in $(0, m)$, consider an $m$-way collision on salt $a$ that comprises of $\mathsf{msg}_1, \ldots, \mathsf{msg}_\ell, (\mathsf{msg}^0_{\ell+1}, \mathsf{msg}^1_{\ell+1}), \ldots,$ $(\mathsf{msg}^0_m, \mathsf{msg}^1_m)$ where for every $i \in [\ell]$, $h(a, \mathsf{msg}_i) = a'$ and for every $i \in [m] \setminus [\ell]$, $h(h(a, \mathsf{msg}^0_i), \mathsf{msg}^1_i) = a'$. Such an $m$-way collision can be thought of as $\ell$-way collision of type 1 and $(m - \ell)$-way collision of type 3. Note that any $m$-way collision which is a mix of two types of collisions contains an $m'$-way pure collision (i.e., collision of a unique type) such that $m' \geq m/2$. Thus, we can reduce bounding the security of such a (mixed) $m$-way collision-finding adversary to bounding security of an $(m/2)$-way (pure) collision finding adversary (adding a multiplicative factor of 2 to our bound for pure collision-finding).

▷ Claim 12. Any $m$-way collision contains an $(m/2)$-way collision of exactly one of the types given in fig. 4.

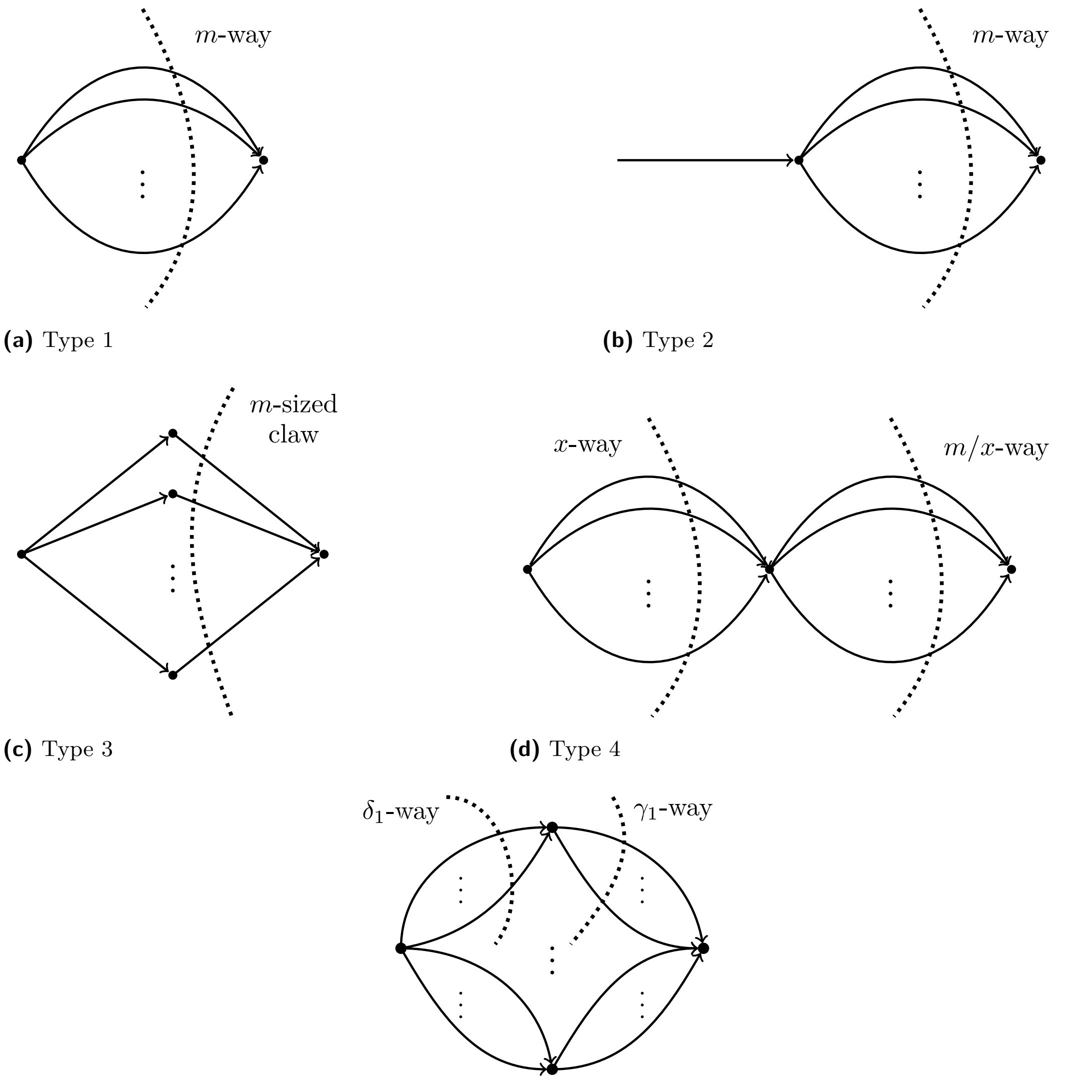

**(a)** Type 1

**(b)** Type 2

**(c)** Type 3

**(d)** Type 4

**(e)** Type 5

**Figure 4** Depiction of types of 2-block $m$-way collisions using function graph of $h$. The vertices denote the salt (in $[N]$) and the arrows correspond to a value in $[M]$.

*Proof.* For an $m$-way 2-block collision on an arbitrary salt $a$, the colliding messages can be denoted by $(\mathsf{msg}_i^0, \mathsf{msg}_i^1)_{i=1}^m$ such that $\mathsf{msg}_i^0 \in [M]$ and $\mathsf{msg}_i^1 \in [M] \cup \{\perp\}$ where $\mathsf{msg}_i^1 = \perp$ denotes that the colliding message is just 1-block, $\mathsf{msg}_i^0$. Let's denote the output salt on which the messages collide by $a'$, i.e., $h(h(a, \mathsf{msg}_i^0), \mathsf{msg}_i^1) = a'$ for all $i \in [m]$.

One possibility is that for all $i \in [m]$, $\mathsf{msg}_i^1 = \perp$. Then $\mathsf{msg}_1^0, \ldots, \mathsf{msg}_m^0$ will have to be distinct and satisfy $h(a, \mathsf{msg}_i^0) = a'$ for every $i$. This is type 1 collision shown in fig. 4a.

Other possibility is that $\mathsf{msg}_i^1 \neq \perp$ for all $i \in [m]$. Then exactly one of the following will hold:

- $\mathsf{msg}_1^0 = \cdots = \mathsf{msg}_m^0$, then $\mathsf{msg}_1^1, \ldots, \mathsf{msg}_m^1$ will have to be distinct and the messages will form type 2 collision shown in fig. 4b.
- $\mathsf{msg}_1^0, \ldots, \mathsf{msg}_m^0$ are distinct and $h(a, \mathsf{msg}_1^0), \ldots, h(a, \mathsf{msg}_m^0)$ are distinct, then the messages will form type 3 collision shown in fig. 4c.

- $\mathsf{msg}_1^0, \ldots, \mathsf{msg}_m^0$ are neither all equal nor all distinct but $h(a, \mathsf{msg}_1^0) = \cdots = h(a, \mathsf{msg}_m^0)$, then the colliding messages will form collision of types 4 shown in fig. 4d.
- $\mathsf{msg}_1^0, \ldots, \mathsf{msg}_m^0$ are not all equal (may or may not be all distinct) and $h(a, \mathsf{msg})$ for $m \in \{\mathsf{msg}_1^0, \ldots, \mathsf{msg}_m^0\}$ are neither all equal nor all distinct, then the colliding messages will form collision of types 5 shown in fig. 4e.

Finally, consider the possibility that $\mathsf{msg}_i^1 \neq \perp$ for some (not all) $i \in [m]$. Then the messages with $\mathsf{msg}_i^1 = \perp$ will be forming type 1 $m'$-way collision (where $m'$ is the number of messages with $\mathsf{msg}_i^1 = \perp$) and the remaining messages will satisfy one of the cases listed above for $(m - m')$-way collision. Then there will exist $m/2$ colliding messages that satisfy one of the types of collisions depicted in fig. 4. $\lhd$

Let $E_t$ be the event that $\mathcal{A}$ wins game $\mathsf{m\text{-}AICR}^2$ by outputting $(m/2)$-way collision of type $t$ in fig. 4. Then

$$\mathsf{Adv}_2^{\text{m-AICR}}(\mathcal{A}) \leq \sum_{t=1}^{5} \Pr\left[\mathsf{m\text{-}AICR}^2_{h,a}(\mathcal{A}) = 1 \wedge E_t\right]$$

Next, we formally define multi-collision resistance game for MD hash functions in the Multi-instance model. We first define the adversary in the model.

▶ **Definition 13** ((S, T, m)-MI adversary). *An algorithm $\mathcal{A}$ is an $(S, T, m)$-MI adversary against $m$-way collision resistance in* $\mathsf{MD}_h$ *if*

- $\mathcal{A}^h$ *receives $S$ salts from $[N]$ as input*
- $\mathcal{A}^h$ *gets to make $ST$ queries to $h$ and outputs* $(\mathsf{msg}_1^i, \ldots, \mathsf{msg}_m^i)_{i=1}^S$

*where $h$ is a function in $[N] \times [M] \to [N]$.*

▶ **Definition 14.** *For a function $h$ in $[N] \times [M] \to [N]$, fixed function $S$, and any $t \in [5]$, the game* $\mathsf{m\text{-}MICR}^{S,t}$ *is defined in fig. 5.*

```
Game $\mathsf{m\text{-}MICR}_h^{S,t}(\mathcal{A})$
    Sample $\{a_1, \ldots, a_S\} \leftarrow_{\$} \left[\binom{N}{S}\right]$
    $(\mathsf{msg}_1^i, \ldots, \mathsf{msg}_m^i)_{i=1}^S \leftarrow \mathcal{A}^h(\{a_1, \ldots, a_S\})$
    If for any $i \in [S]$ and any of $\mathsf{msg}_1^i, \ldots, \mathsf{msg}_m^i$ have more than 2 blocks
        Then Return 0
    If $\forall i \in [S]$: $(\mathsf{msg}_1^i, \ldots, \mathsf{msg}_m^i)$ is $m$-way collision of type $t$ on $a_i$
        Then Return 1
    Else Return 0
```

**Figure 5** Security game $\mathsf{m\text{-}MICR}_h^{S,t}(\mathcal{A})$.

*For an $(S, T, m)$-MI adversary $\mathcal{A}$, the advantage of $\mathcal{A}$ in the game* $\mathsf{m\text{-}MICR}^{S,t}$*, denoted as* $\mathsf{Adv}^{(\mathrm{m,t})\text{-MICR}}(\mathcal{A})$*, is defined as the probability of* $\mathsf{m\text{-}MICR}_h^{S,t}(\mathcal{A})$ *returning 1 when $h$ is a random function in $[N] \times [M] \to [N]$. We define the $(S, T, m)$-MI 2-block multi-collision resistance, denoted by* $\mathsf{Adv}^{(\mathrm{m,t})\text{-MICR}}(S, T, m)$*, as the maximum advantage taken over all $(S, T, m)$-MI adversaries.*

Next, we define a modified version of the $\mathsf{m\text{-}MICR}$ game, which we refer to as $\mathsf{m\text{-}mod\text{-}MICR}$ game. It is different from $\mathsf{m\text{-}MICR}$ game only in the way the challenge salts are sampled.

▶ **Definition 15.** *For a function $h$ in $[N] \times [M] \to [N]$, fixed function $S$, and any $t \in [5]$, the game* $\mathsf{m\text{-}mod\text{-}MICR}^{S,t}$ *is defined in fig. 6.*

*The advantages in the game* $\mathsf{m\text{-}mod\text{-}MICR}$ *is defined similarly to that in the game* $\mathsf{m\text{-}MICR}$.

```
Game $\textsf{m-mod-MICR}_h^{S,t}(\mathcal{A})$
    Sample $a_i \leftarrow_\$ [N]$ for $i \in [S]$
    $(\textsf{msg}_1^i, \ldots, \textsf{msg}_m^i)_{i=1}^S \leftarrow \mathcal{A}^h(\{a_1, \ldots, a_S\})$
    If for any $i \in [S]$ and any of $\textsf{msg}_1^i, \ldots, \textsf{msg}_m^i$ have more than 2 blocks
        Then Return 0
    If $\forall i \in [S]$: $(\textsf{msg}_1^i, \ldots, \textsf{msg}_m^i)$ is $m$-way collision of type $t$ on $a_i$
        Then Return 1
    Else Return 0
```

**Figure 6** Security game $\textsf{m-mod-MICR}_h^{S,t}(\mathcal{A})$.

Next, we present a lemma that establishes a relation between the advantage in the games $\textsf{m-MICR}$ and $\textsf{m-mod-MICR}$ for a MI adversary.

▶ **Lemma 16.** *For any $S, T, u, m, N$, any $t \in [5]$ and $\delta$ such that $\delta \geq S/N$, if $\textsf{Adv}^{(\mathrm{m,t})\text{-}\mathrm{MICR}}(u, T) \leq \delta^u$ for every $u \leq S$ then $\textsf{Adv}^{(\mathrm{m,t})\text{-}\mathrm{mod}\text{-}\mathrm{MICR}}(S, T) \leq (2\delta)^S$.*

**Proof.**

$$\begin{aligned}
&\textsf{Adv}^{(\mathrm{m,t})\text{-}\mathrm{mod}\text{-}\mathrm{MICR}}(S, T) \\
&\leq \sum_{u=1}^{S} \Pr[u \text{ distinct salts among } S \text{ random salts}] * \textsf{Adv}^{(\mathrm{m,t})\text{-}\mathrm{MICR}}(u, T) \\
&\leq \sum_{u=1}^{S} \binom{S}{u} \cdot \left(\frac{u}{N}\right)^{S-u} \cdot \delta^u \\
&\leq \left(\frac{S}{N} + \delta\right)^S \leq (2\delta)^S
\end{aligned}$$

where the third inequality follows from the binomial theorem and the last inequality uses that $\delta \geq S/N$. ◀

Now we present the theorem that allows bounding the security in the auxiliary input model by analyzing the multi-instance game. This theorem is a variation of Theorem 4.1 in [10] and Theorem 3 in [3].

▶ **Theorem 17.** *For any $S, T, m$, $t \in [5]$ and $\delta \in [0, 1]$, if $\textsf{Adv}^{(\mathrm{m}/2,\mathrm{t})\text{-}\mathrm{MICR}}(S, T, m/2) \leq \delta^S$, then $\Pr[\textsf{m-AICR}_{h,a}^2(\mathcal{A}) = 1 \wedge E_t] \leq 4\delta$.*

**Proof.** It is given that $\textsf{Adv}^{(\mathrm{m}/2,\mathrm{t})\text{-}\mathrm{MICR}}(S, T, m/2) \leq \delta^S$. Using lemma 16 we know that $\textsf{Adv}^{(\mathrm{m}/2,\mathrm{t})\text{-}\mathrm{mod}\text{-}\mathrm{MICR}}(S, T, m/2) \leq (2\delta)^S$ when $ST/mN \geq S/N$ (in other words $m \leq T$).

Say for some $t \in [5]$ there exists an $(S, T, m)$-AI adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that $\Pr[\textsf{m-AICR}_{h,a}^2(\mathcal{A}) = 1 \wedge E_t] > 4\delta$. Then we can give an $(S, T, m)$-MI adversary $\mathcal{A}'$ that uses $\mathcal{A}$ and achieves $\textsf{Adv}^{(\mathrm{m}/2,\mathrm{t})\text{-}\mathrm{mod}\text{-}\mathrm{MICR}}(\mathcal{A}') > (2\delta)^S$ which is a contradiction. We describe $\mathcal{A}'$ next.

1. $\mathcal{A}'$ randomly chooses $S$-bit advice, $\sigma$.
2. For each element $a$ in the input set of salts $\{a_1, \ldots, a_S\}$, $\mathcal{A}'$ runs $\mathcal{A}_2(\sigma, a)$ to obtain $m$ messages that collide under $MD_h$ and outputs the $m/2$ messages that form $(m/2)$-way collision of type $t$.

Let $\delta_h := \Pr_a[\textsf{m-AICR}^2_{h,a}(\mathcal{A}) = 1 \wedge E_t]$ for a fixed choice of $h$. Then we know $\mathbb{E}_h[\delta_h] = \Pr_{h,a}[\textsf{m-AICR}^2_{h,a}(\mathcal{A})] > 4\delta$. Then

$$\begin{aligned}\textsf{Adv}^{(\mathrm{m/2,t})\text{-MICR}}(S,T,m/2) &= \Pr_{h,a_1,\ldots,a_S}[(\textsf{m/2})\text{-}\textsf{MICR}^{S,t}_h(\mathcal{A}) = 1] \\ &= \mathbb{E}_h\left[\Pr_{a_1,\ldots,a_S}[(\textsf{m/2})\text{-}\textsf{MICR}^{S,t}_h(\mathcal{A}) = 1]\right] \\ &= \mathbb{E}_h\left[\Pr_{a_1,\ldots,a_S}[(\textsf{m/2})\text{-}\textsf{MICR}^{S,t}_h(\mathcal{A}) = 1 | \sigma = \mathcal{A}^h_1] \cdot \Pr[\sigma = \mathcal{A}^h_1]\right] \\ &= \mathbb{E}_h\left[\delta^S_h \cdot \frac{1}{2^S}\right] \geq \frac{\mathbb{E}_h[\delta_h]^S}{2^S} > (2\delta)^S,\end{aligned}$$

where the second to last inequality is by Jensen's inequality. ◀

▶ **Lemma 18.** *For any $S, T, N, m$ and $t \in [5]$ such that $m = \omega(\log^2 N)$ and $m \leq ST$,*

$$\textsf{Adv}^{(\mathrm{m,t})\text{-MICR}}(S,T,m) = \left(\tilde{O}\left(\frac{ST}{mN} + \frac{(eT/m)^m}{N^{m-1}}\right)\right)^S.$$

It is worth noting that in the lemma we make an additional assumption that $m = \omega(\log^2 N)$ which does not contradict with the statement of Thm 11. And that is because for $m = O(\log^c N)$, the theorem holds trivially from the reduction of 2-way collision finding to $m$-way collision finding and the result of [1].

Now proof of theorem 11 follows from theorem 17 and lemma 18 as follows:

$$\begin{aligned}&\forall t \in [5] : \Pr[\textsf{m-AICR}^2_{h,a}(\mathcal{A}) = 1 \wedge E_t] \leq 4 \cdot \left(\textsf{Adv}^{(\mathrm{m,t})\text{-MICR}}(S,T,m)\right)^{1/S} \\ \implies &\textsf{Adv}^{\text{m-AICR}}_2(\mathcal{A}) \leq \sum_{t=1}^{5} \Pr[\textsf{m-AICR}^2_{h,a}(\mathcal{A}) = 1 \wedge E_t] \\ &\leq 4 \cdot \sum_{t=1}^{5} (\textsf{Adv}^{(\mathrm{m,t})\text{-MICR}}(S,T,m))^{1/S} \\ &= \tilde{O}\left(\frac{ST}{mN} + \frac{(eT/m)^m}{N^{m-1}}\right)\end{aligned}$$

To complete the proof of theorem 11, it only remains to prove lemma 18.

**Proof.** Let's fix a $(S,T,m)$-MI adversary $\mathcal{A}$ and let $h$ be a random oracle. We denote the input of random $S$-sized set of salts by $\{a_1, \ldots, a_S\}$. Let $\mathbf{X}^t_i$ be the indicator variable that $\mathcal{A}$ wins on salt $a_i$, i.e., $\mathcal{A}$ finds $m$-way collision of type $t$ on salt $a_i$ for $t \in [5]$.

To prove the lemma , we need to show for each $t \in [5]$, that $\textsf{Adv}^{(\mathrm{m,t})\text{-MICR}}(S,T,m) = \left(\tilde{O}\left(\frac{ST}{mN} + \frac{(eT/m)^m}{N^{m-1}}\right)\right)^S$.

Type 1 collisions are intuitively the easiest to analyze. Consider the adversary gets $S$ distinct salts, say $\{a_1, \ldots, a_S\}$, as input and in order to win on $a_i$ for any $i \in [S]$, $\mathcal{A}$ needs to make $m$ queries of the form $(a_i, *)$ such that all their outputs are equal under $h$. If $\mathcal{A}$ wins on "too many" $S$-sized subsets of $[N]$, we give an encoder that could use this $\mathcal{A}$ to compress $h$ but that is impossible. Hence, there exists no such $\mathcal{A}$.

The encoding algorithm will be as follows:

- Store the $Sm$ queries among the $ST$ queries that are involved in collisions for the $S$ challenge salts in an unordered set, say $W$. This would require $\log\binom{ST}{Sm}$ bits.
- For all $i \in [S]$, delete the output of all the queries of the form $(a_i, *)$ except the first occurring query in the unordered set $W$ from the table of $h$. This saves $S(m-1)\log N$ bits.

Say $\mathsf{Adv}^{(\mathrm{m,1})\text{-}\mathrm{MICR}}(\mathcal{A}) = \epsilon$. Then

$$\log \epsilon \leq \log \binom{ST}{Sm} - S(m-1)\log N$$
$$\implies \epsilon \leq \frac{\binom{ST}{Sm}}{N^{S(m-1)}} \leq \left[\frac{(eT/m)^m}{N^{m-1}}\right]^S.$$

Analysis for collisions of type 2-5 are more evolved. We refer the readers to Appendix C for a careful analysis of type 2 collisions. Analysis for remaining collision types are based on similar ideas. However, due to lack of space we omit the complete analysis here. We refer the readers to the full version of the paper. ◀

◀

## References

1 Akshima, David Cash, Andrew Drucker, and Hoeteck Wee. Time-Space Tradeoffs and Short Collisions in Merkle-Damgård Hash Functions. In *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 157–186. Springer, 2020.

2 Akshima, Xiaoqi Duan, Siyao Guo, and Qipeng Liu. On Time-Space Lower Bounds for Finding Short Collisions in Sponge Hash Functions. In *Theory of Cryptography - 21st International Conference, TCC 2023, Taipei, Taiwan, November 29 - December 2, 2023, Proceedings, Part III*, volume 14371 of *Lecture Notes in Computer Science*, pages 237–270. Springer, 2023. doi:10.1007/978-3-031-48621-0_9.

3 Akshima, Siyao Guo, and Qipeng Liu. Time-Space Lower Bounds for Finding Collisions in Merkle-Damgård Hash Functions. In *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 192–221. Springer, 2022. doi:10.1007/978-3-031-15982-4_7.

4 Mohammad A AlAhmad, Imad Fakhri Alshaikhli, and Mridul Nandi. Joux Multicollisions Attack in Sponge Construction. In *Proceedings of the 6th International Conference on Security of Information and Networks*, pages 292–296, 2013.

5 Itay Berman, Akshay Degwekar, Ron D Rothblum, and Prashant Nalini Vasudevan. Multi-Collision Resistant Hash Functions and their Applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 133–161. Springer, 2018.

6 Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge Functions. In *ECRYPT hash workshop*, volume 2007, 2007.

7 Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-Collision Resistance: a Paradigm for Keyless Hash Functions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 671–684, 2018.

8 Ernest Brickell, David Pointcheval, Serge Vaudenay, and Moti Yung. Design Validations for Discrete Logarithm Based Signature Schemes. In *International Workshop on Public Key Cryptography*, pages 276–292. Springer, 2000.

9 Dror Chawin, Iftach Haitner, and Noam Mazor. Lower Bounds on the Time/Memory Tradeoff of Function Inversion. In *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part III*, pages 305–334, 2020.

10 Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight Quantum Time-Space Tradeoffs for Function Inversion. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 673–684. IEEE, 2020.

11 Sandro Coretti, Yevgeniy Dodis, and Siyao Guo. Non-Uniform Bounds in the Random-Permutation, Ideal-Cipher, and Generic-Group Models. In *Annual International Cryptology Conference*, pages 693–721. Springer, 2018.

12 Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John Steinberger. Random Oracles and Non-Uniformity. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 227–258. Springer, 2018.

13 Henry Corrigan-Gibbs and Dmitry Kogan. The Discrete-Logarithm Problem with Preprocessing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 415–447. Springer, 2018.

14 Henry Corrigan-Gibbs and Dmitry Kogan. The Function-Inversion Problem: Barriers and Opportunities. In *TCC*, 2019. Also, Crypto ePrint 2019/1046.

15 Anindya De, Luca Trevisan, and Madhur Tulsiani. Time Space Tradeoffs for Attacks against One-Way Functions and PRGs. In *Annual Cryptology Conference*, pages 649–665. Springer, 2010.

16 Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. Fixing Cracks in the Concrete: Random Oracles with Auxiliary Input, Revisited. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 473–495. Springer, 2017.

17 Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. Time-Space Tradeoffs for Sponge Hashing: Attacks and Limitations for Short Collisions. In *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 131–160. Springer, 2022. `doi:10.1007/978-3-031-15982-4_5`.

18 Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. Optimal Security for Keyed Hash Functions: Avoiding Time-Space Tradeoffs for Finding Collisions. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part IV*, volume 14007 of *Lecture Notes in Computer Science*, pages 440–469. Springer, 2023. `doi:10.1007/978-3-031-30634-1_15`.

19 Ashrujit Ghoshal and Ilan Komargodski. On Time-Space Tradeoffs for Bounded-Length Collisions in Merkle-Damgård Hashing. In *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 161–191. Springer, 2022. `doi:10.1007/978-3-031-15982-4_6`.

20 Marc Girault and Jacques Stern. On the Length of Cryptographic Hash-Values Used in Identification Schemes. In *Annual International Cryptology Conference*, pages 202–215. Springer, 1994.

21 Alexander Golovnev, Siyao Guo, Spencer Peters, and Noah Stephens-Davidowitz. Revisiting Time-Space Tradeoffs for Function Inversion. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part II*, volume 14082 of *Lecture Notes in Computer Science*, pages 453–481. Springer, 2023. `doi:10.1007/978-3-031-38545-2_15`.

22 Nick Gravin, Siyao Guo, Tsz Chiu Kwok, and Pinyan Lu. Concentration Bounds for Almost k-wise Independence with Applications to Non-Uniform Security. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2404–2423. SIAM, 2021.

23 M. Hellman. A Cryptanalytic Time-memory Trade-off. *IEEE Trans. Inf. Theor.*, 26(4):401–406, July 1980.

24 Russell Impagliazzo and Valentine Kabanets. Constructive Proofs of Concentration Bounds. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, pages 617–631, 2010.

25 Antoine Joux. Multicollisions in Iterated Hash Functions. Application to Cascaded Constructions. In *Annual International Cryptology Conference*, pages 306–316. Springer, 2004.

26 Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision Resistant Hashing for Paranoids: Dealing with Multiple Collisions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 162–194. Springer, 2018.

27 Qipeng Liu and Mark Zhandry. On Finding Quantum Multi-Collisions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 189–218. Springer, 2019.

28 Ronald L Rivest and Adi Shamir. PayWord and MicroMint: Two Simple Micropayment Schemes. In *International workshop on security protocols*, pages 69–87. Springer, 1996.

29 Ron D. Rothblum and Prashant Nalini Vasudevan. Collision-Resistance from Multi-Collision-Resistance. In *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 503–529. Springer, 2022. doi:10.1007/978-3-031-15982-4_17.

30 Dominique Unruh. Random Oracles and Auxiliary Input. In *Annual International Cryptology Conference*, pages 205–223. Springer, 2007.

## A Bounded Length Multi-Collision Attack

The attack is as follows:

1. In the *Offline* stage,
   a. $\mathcal{A}_1$ randomly picks $s$ salts, denoted $\{a_1^0, \ldots, a_s^0\}$. Let's define $x := \max\{0, \lfloor (B - \log_2 m)/2 \rfloor\}$ and $y := \min\{\log_2 m, B-1\}$. Then for every $i \in [s]$ and $j \in [x]$, $\mathcal{A}_1$ first iteratively computes $a_i^j = h(a_i^{j-1}, 0)$ to obtain the set $\{a_1^x, \ldots, a_s^x\}$.
   b. Next, $\mathcal{A}_1$ finds a $y$-length chain of $m^{1/y}$-way collision (recall fig. 3) on $a_i^x$ for every $i \in [s]$.
   c. Finally, $\mathcal{A}_1$ stores the tuple of $a_i^x$ and the corresponding $y$-length chain for every $i \in [s]$ as advice.
2. In the *Online* stage, $\mathcal{A}_2$ is given the advice from $\mathcal{A}_1$ and a random salt, denoted $a$, as challenge. $\mathcal{A}_2$ runs the following steps for $k \in [\lfloor T/2x \rfloor]$ when $x$ is non-zero and $k \in [T]$ when $x = 0$:
   - Set $b = h(a, k)$ and $\mathsf{counter} = 1$
   - While $a, b \notin \{a_1^x, \ldots, a_s^x\}$ and $\mathsf{counter} < 2x$:
     - query and set $b = h(b, 0)$
     - increment $\mathsf{counter}$ by 1.
   - Say $b = a_\ell^x$ for some $\ell$, then $\mathcal{A}_2$ can learn the $y$-length chain of $m^{1/y}$-way collision on $a_\ell^x$ from the advice and output $m$ colliding messages consisting of $k$ concatenated with $(\mathsf{counter} - 1)$ 0's followed by $m$ different combinations in the chain.

The analysis of this attack can be found in the full version of the paper.

## B Proof of Theorem 10

We first present the attack for finding 1-block $m$-way collisions.

Note that the following trivial adversary, say $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, achieves $\tilde{\Omega}\left(\frac{S}{mN} + \frac{(T/m)^m}{N^{m-1}}\right)$ advantage.

1. In the *Offline* stage, $\mathcal{A}_1$ stores (1-block) $m$-way collisions for $\tilde{\Omega}(S/m)$ salts as part of the advice.
2. In the *Online* stage, if the input random salt, say $a$, is one of the salts for which the advice contains a $m$-way collision, adversary returns the corresponding $m$-way collision. Else the adversary tries to find $m$-way collision using it's $T$ queries to the oracle $h$.

Next we prove the security bound from the theorem. We prove it via "global" compression which is based on a result in [16].

Fix an $(S, T, m)$-AI adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ having advantage $\epsilon$. Let $\sigma$ denote the $S$-bit advice output by $\mathcal{A}_1$ and let $\mathcal{G}$ be the set of salts in $[N]$ on which $\mathcal{A}_2$ succeeds in finding $m$-way collision when given $\sigma$ and makes $T$ queries to $h$.

We want to show that the set $\mathcal{G}$ cannot be large for any $(S, T, m)$-AI adversary. In other words, it is impossible for any $(S, T, m)$-AI adversary to succeed in finding $m$-way collisions on "too" many salts, irrespective of what information is contained in the $S$-bit advice $\sigma$. The idea is to give an encoding algorithm that uses such an adversary that succeeds on a "lot" of salts to compress $h$, which should be impossible as $h$ is a random function.

Note that since $\mathcal{A}_2$ finds $m$-way collisions on each of the salts in $\mathcal{G}$, for every $a \in \mathcal{G}$ there should exist $m$ queries of the type $(a, *) \in [N] \times [M]$ such that outputs of all the $m$ queries are same under $h$. Our encoding algorithm uses this repetition in outputs to compress as follows:

1. Store the advice $\sigma$ (requires $S$-bits)
2. Store the size of the set $\mathcal{G}$ (in other words, the number of elements in $\mathcal{G}$), denoted $|\mathcal{G}|$ (requires $\log N$ bits)
3. Store the set $\mathcal{G}$ (requires $\log \binom{N}{|\mathcal{G}|}$ bits)
4. Store the unordered set of $|\mathcal{G}| \cdot m$ queries corresponding to the $m$-way collisions for each salt in $\mathcal{G}$. Let's denote this set by $\mathcal{X}$. Note that $\mathcal{X}$ is a subset of the $|\mathcal{G}| \cdot T$ queries made by $\mathcal{A}_2$ using the assumption that $\mathcal{A}_2$ has to query it's outputs. Thus, storing $\mathcal{X}$ requires $\log \binom{|\mathcal{G}|T}{|\mathcal{G}|m}$ bits.
5. Delete the outputs of $\mathcal{G}|(m-1)$ queries in $\mathcal{X}$ from table of $h$. (Note that the table for $h$ is stored as follows: table contains output of $h$ on the queries made by $\mathcal{A}_2$ on the set $\mathcal{G}$ followed by the output of $h$ on the remaining entries of queries in $[N] \times [M]$ in lexicographic order.) This saves $|\mathcal{G}|(m-1) \cdot \log N$ bits.

Via the compression argument of De et al. [15] (stated as lemma 4 in this paper), the following holds:

$$S + \log N + \log \binom{N}{|\mathcal{G}|} + \log \binom{|\mathcal{G}|T}{|\mathcal{G}|m} + NM \log N - |\mathcal{G}|(m-1) \cdot \log N \geq NM \log N$$

$$\implies S + \log N + |\mathcal{G}| \log \left(\frac{N}{|\mathcal{G}|}\right) + |\mathcal{G}| m \log \left(\frac{|\mathcal{G}|T}{|\mathcal{G}|m}\right) \geq |\mathcal{G}|(m-1) \log N$$

$$\implies S + \log N \geq |\mathcal{G}| \log \left(\frac{N^{m-1}|\mathcal{G}|}{N(T/m)^m}\right)$$

We take expectation on both sides of the above equation and use convexity of the function $x \log x$ along with $\mathbb{E}[|\mathcal{G}|] = \epsilon N$ as follows:

$$\begin{aligned}
\mathbb{E}[S + \log N] &\geq \mathbb{E}\left[|\mathcal{G}| \log\left(\frac{N^{m-1}|\mathcal{G}|}{N(T/m)^m}\right)\right] \\
\implies S + \log N &\geq \mathbb{E}[|\mathcal{G}|] \log\left(\frac{N^{m-1}\mathbb{E}[|\mathcal{G}|]}{N(T/m)^m}\right) \\
&\geq \epsilon N \cdot \log\left(\frac{N^{m-1} \cdot \epsilon N}{N(T/m)^m}\right) \geq \epsilon N \cdot \log\left(\frac{N^{m-1} \cdot \epsilon}{(T/m)^m}\right)
\end{aligned}$$

Consider the following two cases:

1. If

$$\frac{N^{m-1}\epsilon}{(T/m)^m} \leq 2^m$$

This simply implies

$$\epsilon \leq \frac{(2T/m)^m}{N^{m-1}}$$

2. Otherwise

$$\frac{N^{m-1}\epsilon}{(T/m)^m} > 2^m$$

which implies

$$S + \log N \geq \epsilon N \log 2^m \implies \epsilon \leq \frac{S + \log N}{mN}$$

This completes the proof for the security bound.

## C Bounding MI-security for Type 2 collisions

The analysis would be trivial and same as for $t = 1$ if $m$-way collision for each salt used distinct queries. However, queries can be reused among collisions for different salts. Refer to fig. 4b for a visual depiction.

Say $x$ salts use the same $m$-way collision structure. Refer fig. 7 for a visual depiction. Say the adversary finds $y$ such structures where $x = \delta_i$ for the $i^{\text{th}}$ structure.

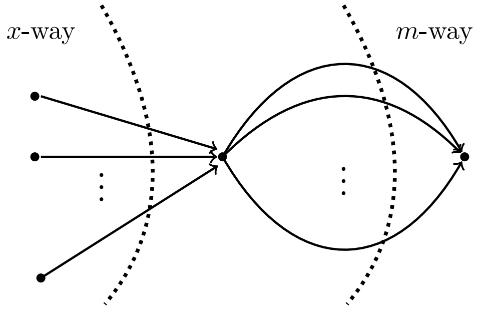

**Figure 7** Depiction of $x$ salts “sharing” a $m$-way 1-block collision.

Then it holds that

$$\delta_1 + \cdots + \delta_y \geq S$$

The encoding algorithm will be as follows:

1. Store $y$ which requires $\log S$ bits.
2. Store the subset of $ST$ queries forming $m$-way collisions in an unordered fashion. From the table of $h$, delete the entries corresponding to the queries in the set except the first query. There will be $y$ such $m$-sized sets.
3. Store the number of $i \in [y]$ such that $\delta_i \geq m$ which requires $\log S$-bits.
4. For $i \in [y]$: if $\delta_i > m$ do the following
   - Store $\delta_i$
   - Store the $\delta_i$-sized set of queries that have the same output. From the table of $h$, delete the entries corresponding to the queries in the set except the first query.

Say $\epsilon$ is the advantage. For decoding purposes, the encoder must store these sets of collisions but also store the values of $y, \{\delta_i | i \in [y], \delta_i > m\}$ and $|\{\delta_i | i \in [y], \delta_i > m\}|$. It holds by the compression argument that

$$\begin{aligned}\log \epsilon \leq \log S + y \log \binom{ST}{m} - y \cdot (m-1) \log N + \log S + \\ \sum_{i \in [y]: \delta_i \geq m} \left[ \log S + \log \binom{ST}{\delta_i} - (\delta_i - 1) \log N \right] \\ = \log S + \sum_{i \in [y]: \delta_i < m} \left[ \log \binom{ST}{m} - (m-1) \log N \right] + \log S + \\ \sum_{i \in [y]: \delta_i \geq m} \left[ \log \binom{ST}{m} - (m-1) \log N + \log S + \log \binom{ST}{\delta_i} - (\delta_i - 1) \log N \right]\end{aligned}$$

In the analysis, we consider the following cases:

1. $x \geq m$
   It would suffice to take into account the probability of $x$ of the $ST$ queries having the same output, which is:
   $$\frac{\binom{ST}{x}}{N^{x-1}} \leq \left(\frac{eST}{xN}\right)^x \cdot N \leq \left(\frac{2eST}{xN}\right)^x \cdot \frac{1}{2^x} \cdot N \leq \left(\frac{2eST}{mN}\right)^x$$
   where the last inequality uses that $\frac{1}{2^x} \leq \frac{1}{2^m} \leq N$ as $x \geq m \geq \log N$.
2. $x < m$
   Then the probability of an $m$-way collision is:
   $$\leq \frac{\binom{ST}{m}}{N^{m-1}} \leq \left(\frac{eST}{mN}\right)^m \cdot N \leq \left(\frac{2eST}{mN}\right)^m \cdot \frac{1}{2^m} \cdot N \leq \left(\frac{2eST}{mN}\right)^x$$
   where the last inequality holds for $m \geq \log N$ and $2eST/mN \leq 1$.

Hence, the probability of finding the structure where $x$ salts share a $m$-way 1-block collision such that we get type 2 collision on the $x$ salts in the MI setting is at most $(2eST/mN)^x$ irrespective of whether $x < m$ or $x \geq m$. This shows that in either case, we can get the desired bound exponentially small in $x$.

Then, we know for each $i$ such that $\delta_i < m$,

$$\log \binom{ST}{m} - (m-1) \log N \leq \log \left(\frac{2eST}{mN}\right)^{\delta_i}$$

and for each $i$ such that $\delta_i \geq m$

$$\begin{aligned}
&\log \binom{ST}{m} - (m-1)\log N + \log S + \log \binom{ST}{\delta_i} - (\delta_i - 1)\log N \\
&\leq \log S + \log \binom{ST}{\delta_i} - (\delta_i - 1)\log N \leq \log \left[ S \cdot \left( \frac{2eST}{mN} \right)^{\delta_i} \right]
\end{aligned}$$

using $2eST \leq mN$. Then,

$$\begin{aligned}
\implies \epsilon &\leq S^2 \cdot \left[ \prod_{i \in [y] : \delta_i < m} \left( \frac{2eST}{mN} \right)^{\delta_i} \right] \cdot \left[ \prod_{i \in [y] : \delta_i \geq m} S \cdot \left( \frac{2eST}{mN} \right)^{\delta_i} \right] \\
&\leq S^2 \cdot \left[ \prod_{i \in [y] : \delta_i < m} \left( \frac{2eST}{mN} \right)^{\delta_i} \right] \cdot \left[ \prod_{i \in [y] : \delta_i \geq m} S \cdot \frac{1}{2^m} \cdot \left( \frac{4eST}{mN} \right)^{\delta_i} \right] \\
&\leq 2^S \cdot \prod_{i \in [y]} \left( \frac{4eST}{mN} \right)^{\delta_i} = 2^S \cdot \left( \frac{4eST}{mN} \right)^{\delta_1 + \cdots + \delta_y} \leq 2^S \cdot \left( \frac{4eST}{mN} \right)^{S} \\
&= \left( \frac{8eST}{mN} \right)^{S}
\end{aligned}$$

where the last inequality uses that $\delta_1 + \delta_2 + \delta_y >= S$ as these structures should give collisions on $S$ distinct salts.