On Time-Space Tradeoffs for Bounded-Length Collisions in Merkle-Damgård Hashing

Ashrujit Ghoshal[1] and Ilan Komargodski[2]

[1] Paul G. Allen School of Computer Science & Engineering
University of Washington, Seattle, Washington, USA
`ashrujit@cs.washington.edu`
[2]Hebrew University of Jerusalem and NTT Research
`ilank@cs.huji.ac.il`

Abstract

We study the power of preprocessing adversaries in finding bounded-length collisions in the widely used Merkle-Damgård (MD) hashing in the random oracle model. Specifically, we consider adversaries with arbitrary S-bit advice about the random oracle and can make at most T queries to it. Our goal is to characterize the advantage of such adversaries in finding a B-block collision in an MD hash function constructed using the random oracle with range size N as the compression function (given a random salt).

The answer to this question is completely understood for very large values of B (essentially $\Omega(T)$) as well as for $B = 1, 2$. For $B \approx T$, Coretti et al. (EUROCRYPT '18) gave matching upper and lower bounds of $\tilde{\Theta}(ST^2/N)$. Akshima et al. (CRYPTO '20) observed that the attack of Coretti et al. could be adapted to work for any value of $B > 1$, giving an attack with advantage $\tilde{\Omega}(STB/N + T^2/N)$. Unfortunately, they could only prove that this attack is optimal for $B = 2$. Their proof involves a compression argument with exhaustive case analysis and, as they claim, a naive attempt to generalize their bound to larger values of B (even for $B = 3$) would lead to an explosion in the number of cases needed to be analyzed, making it unmanageable. With the lack of a more general upper bound, they formulated the *STB conjecture*, stating that the best-possible advantage is $\tilde{O}(STB/N + T^2/N)$ for any $B > 1$.

In this work, we confirm the STB conjecture in many new parameter settings. For instance, in one result, we show that the conjecture holds for all constant values of B, significantly extending the result of Akshima et al. Further, using combinatorial properties of graphs, we are able to confirm the conjecture even for super constant values of B, as long as some restriction is made on S. For instance, we confirm the conjecture for all $B \leqslant T^{1/4}$ as long as $S \leqslant T^{1/8}$. Technically, we develop structural characterizations for bounded-length collisions in MD hashing that allow us to give a compression argument in which the number of cases needed to be handled does not explode.

Contents

1 Introduction

Starting from the seminal work of Hellman [Hel80], there have been significant efforts to understand the power of preprocessing attacks in various applications and constructions (e.g., [Yao90, FN99, Oec03, BBS06, Unr07, DTT10, DGK17, AAC⁺17, CDGS18, CDG18, ACDW20, CK19, CHM20]). Preprocessing attacks, i.e., ones that utilize a bounded amount of auxiliary information, capture the standard modeling of attackers as *non-uniform*, allowing them to obtain some arbitrary (but bounded length) "advice" before attacking the system. In this work, we continue the recent line of works studying the power of preprocessing adversaries in the context of finding collisions in the widely used Merkle-Damgård (MD) design.

Collision resistance of salted MD. The Merkle-Damgård hash function construction [Mer82, Mer87, Mer89, Dam87] is a popular design for building an arbitrary-size-input compression function from a fixed-size-input compression function. This design is not only extremely fundamental in cryptographic theory, but it also underlies popular hash functions used in practice, most notably MD5, SHA-1, and SHA-2.

The MD construction is defined relative to a compressing function $h: [N] \times [M] \to [N]$, modeled as a random oracle, as follows. First, for $a \in [N]$ and $\alpha \in [M]$, let $\mathsf{MD}_h(a, \alpha) = h(a, \alpha)$. Then, define recursively

$$\mathsf{MD}_h(a, (\alpha_1, \ldots, \alpha_B)) = h(\mathsf{MD}_h(a, (\alpha_1, \ldots, \alpha_{B-1})), \alpha_B)$$

for $a \in [N]$ and $\alpha_1, \ldots, \alpha_B \in [M]$. The a is referred to as *salt* (sometimes also called IV) and each of the following B elements are referred to as *blocks*.

Due to the ubiquitous influence of this hashing paradigm, both in theory and practice, characterizing the complexity of finding collisions in MD_h (on a random salt) is a fundamental problem. The well-known birthday attack gives a T-query attacker with $\Theta(T^2/N)$ advantage. However, this attack is very generic: it neither takes advantage of the structure of MD_h nor does it utilize the fact that the attacker may have access to some limited amount of "advice" about h due to a long preprocessing phase. But, there is a good reason that security against non-uniform attackers has become the standard notion of security in the cryptographic literature: it captures the natural idea that an adversary may have been designed to attack specific instances, guaranteeing security against an expensive preprocessing stage, or even unknown future attacks. On the whole, it is widely believed by the theoretical community that non-uniformity is the right cryptographic modeling of attackers, despite being overly conservative and including potentially unrealistic attackers. Therefore, understanding the complexity of finding collisions in MD, allowing preprocessing, is a fundamental problem.

The auxiliary-input random oracle model. The concrete hash functions h used in real-life do not have solid theoretical foundations from the perspective of provable security. Therefore, when analyzing the security of the MD construction, the function h is typically modeled as a completely random one, i.e., a random oracle. We follow the standard approach and model preprocessing adversaries using the influential extension of the random oracle model termed *auxiliary-input random oracle model* (AI-ROM). This model was (implicitly) used, for example, in the classical works of Yao [Yao90] and Fiat and Naor [FN99], and formally defined in the influential work of Unruh [Unr07] which was recently revisited by Dodis, Guo, and Katz [DGK17] and Coretti et al. [CDGS18].

The AI-ROM models preprocessing adversaries as two-stage algorithms $(\mathcal{A}_1, \mathcal{A}_2)$ parametrized by S (for "space") and T (for "time"). The first part $\mathcal{A}_1$ has unbounded access to the random oracle h, and its goal is to compute an S-bit "advice" σ for $\mathcal{A}_2$. The second part $\mathcal{A}_2$ gets the advice σ, can make at most T queries to the random oracle, and attempts to accomplish some task involving h. In our case, $\mathcal{A}_2$ gets a random salt a as a challenge and its goal is to come up with a collision in $\mathsf{MD}_h(a, \cdot)$. Both $\mathcal{A}_1$ and $\mathcal{A}_2$ are unbounded in running time.

Known results. Collision resistance of salted MD hash functions in the AI-ROM was first studied by Coretti, Dodis, Guo, and Steinberger [CDGS18]. Among other results, they showed an attack, loosely based on the idea of rainbow tables [Hel80, Oec03], with advantage $\tilde{\Omega}(ST^2/N)$.[1] [2] They further showed that

---

[1] Throughout the paper, the ˜ notation suppresses poly-logarithmic terms in N.

[2] By "advantage" we mean the probability of finding a collision.

this attack is optimal. Namely, no attack can have an advantage better than $\tilde{O}(ST^2/N)$. (Notice that this attack beats the naive birthday attack mentioned above for typical values of S.) In a more recent work of Akshima, Cash, Drucker, and Wee [ACDW20] observed that the attack of Coretti et al. [CDGS18] results in very long collisions, of the order of T blocks, which may limit their practical usefulness. While formally, a length T collision does violate collision resistance, it is hard to imagine a natural application where it is useful. Indeed, for reasonable values of T, say $T = 2^{60}$, it is unlikely that such a collision, which is several petabytes long, could damage any widely-used application.

Akshima et al. [ACDW20] therefore raise the very natural question of what is the complexity of finding *short* collisions.

What is the complexity, as a function of S and T (the allowed space and query bounds, respectively), of finding a B-block collision in salted MD?

Although this question is very natural and clean, as mentioned, a complete answer is known only in the extreme cases, either when $B = 2$ or when $B = \tilde{\Omega}(T)$. Indeed, when B is very close to T, the result above of Coretti et al. [CDGS18] implies that the advantage is $\tilde{\Theta}(ST^2/N)$. The case of $B = 2$ was resolved by Akshima et al. [ACDW20] who showed that the advantage is $\tilde{\Theta}(ST/N + T^2/N)$. Even for $B = 3$, a complete answer is not known: The analysis of Akshima et al. [ACDW20] consists of an elaborate case analysis tailored to the $B = 2$ case, and they claim that even for $B = 3$ the proof of their lower bound "... *would be too long and complex to write down*".[3]

The STB conjecture. In terms of upper bounds, Akshima et al. [ACDW20] showed that a variant of Coretti et al.'s (rainbow tables inspired) attack could be generalized to get a B-block collision with advantage $\tilde{\Omega}(STB/N)$. This attack generalized the attack of Coretti et al. [CDGS18] which gives an $O(T)$-block collision with probability $\tilde{\Omega}(ST^2/N)$. With the lack of better bounds on the best possible attack for a wide range of B's (anywhere between $B = 3$ and $B \ll T$), Akshima et al. [ACDW20] put forward the "*STB conjecture*" which posits that the optimal attack for finding length B collisions has advantage $\tilde{\Theta}(STB/N + T^2/N)$ (i.e., the better between their attack and the generic birthday attack).

We believe that our current understanding of the exact security that MD-style constructions could ideally achieve is insufficient. Therefore, given how widespread MD-based hash functions are, progress towards resolving the conjecture is highly important.

1.1 Our Results

Our main result confirms the STB conjecture in many new parameter settings. Specifically, we prove two new upper bounds on the advantage of the best attack for finding short collisions in salted MD hash functions in the AI-ROM. The first bound confirms the STB conjecture for all constant values of B. The second result confirms the STB conjecture even for super constant values of B but only for moderately large values of S compared to T.

STB conjecture is true for all constant B. We show that for any $B \in O(1)$, the advantage of any S-space T-query attacker in finding a length B collision is bounded by $\tilde{O}(ST/N + T^2/N)$, matching the known attack up to poly-logarithmic factors.

Theorem 1.1 (Informal; See Theorem 5.1)**.** *For every constant B, the STB conjecture is true.*

This theorem is obtained as a special case of a more general bound on the advantage of any S-space T-query attacker in finding a length B collision of the form

$$\tilde{O}\left(\frac{STB^2(\log^2 S)^{B-2}}{N} + \frac{T^2}{N}\right).$$

Note that this bound is meaningful when B is a constant (or slightly bigger) but becomes vacuous when say, $B = \log N$.

---

[3]For $B = 1$ a tight bound of $\Theta(S/N + T^2/N)$ is known [DGK17].

STB conjecture is true for all $SB \ll T$. We show that as long as $S, B \ll T$, the conjecture is true again. Specifically, we show that whenever $S^4 B^2 \in \tilde{O}(T)$, the maximal advantage of any S-space T-query attacker in finding a length B collision is obtained by the birthday attack, up to poly-logarithmic factors. For example, when $SB \leqslant T^{1/4}$, the maximal advantage is $O(T^2/N)$, and therefore the STB conjecture holds.

Theorem 1.2 (Informal; See Theorem 6.1). *For every $S^4 B^2 \in \tilde{O}(T)$, the STB conjecture is true. For instance, the conjecture holds if either*

- $B \in \mathsf{poly} \log N$ *and* $S \in \tilde{O}(T^{1/4})$, *or*

- $B \in \tilde{O}(T^{1/4})$ *and* $S \in \tilde{O}(T^{1/8})$.

This theorem is obtained as a special case of a more general bound on the advantage of any S-space T-query attacker in finding a length B collision of the form

$$\tilde{O}\left(\frac{S^4 T B^2}{N} + \frac{T^2}{N}\right).$$

A concrete comparison between the results. The two bounds are generally incomparable. While the bound from Theorem 1.1 is asymptotically tight whenever B is constant (independent of S, T), it becomes vacuous for say $B = \log N$. On the other hand, the bound from Theorem 1.2 is meaningful for all $B \in o(N^{1/2})$, as long as $S^4 \cdot B^2 \ll N$. For instance, assume that $S = N^{1/16}$ and $B \in \Theta(N^\epsilon)$ (for $0 < \epsilon < 1/8$). In this setting, the bound from Theorem 1.1 is trivial. On the other hand, the bound from Theorem 1.2 gives that any successful attack must satisfy $T \in \tilde{\Omega}(N^{1/2})$ which is strictly better than what the generic $\tilde{O}(ST^2/N)$ bound gives (it only gives $T \in \tilde{\Omega}(N^{15/32})$).

Technical highlight. The main technical component in both of our bounds is a compression argument that uses a "too-good-to-be-true" attacker to non-trivially compress a uniformly random sequence of bits, thereby getting a contradiction. The setup is somewhat similar to the one of Akshima et al. [ACDW20] (although slightly more modular), but our compression argument deviates from theirs significantly. Their argument inherently relied on the fact that there are at most two blocks in the collision, therefore greatly simplifying the possible structures to consider. In contrast, we consider arbitrary length collisions, and thus we have to deal with all possible structures of collisions. Our proof identifies and analyzes a general structure for MD collisions and unveils a natural combinatorial problem that influences the resulting upper bound on the advantage of preprocessing adversaries. Specifically, it turns out that the "dominant extra" terms in both of our bounds ($(\log^2 S)^{B-2}$ in the first bound and S^3 in the second) emerge due to the need to encode a *reverse path* in a general (fan-out 1, but possibly large fan-in) directed graph, where the graph is the one induced by the queries that the adversary makes to the random oracle. Any improvement on this encoding would immediately imply a better upper bound, a fact that we hope will lead to better bounds in the future.

1.2 Discussion

As mentioned, the MD paradigm underlies numerous hash function constructions that are central building blocks in many applications. There are several popular variants of the MD paradigm implemented in practice. In this work, we follow previous works and focus on the cleanest variant for concreteness. One prominent variant withstands length extension attacks by padding the input message with its length. (In fact, this is the version suggested by Merkle and Damgård.) We remark that our results directly apply to this padded variant. Specifically, the "STB" attack finds a collision with the same number of blocks, so it readily extends to this padded variant. Our bounds on the best possible attacks also extend to this setting since the argument did not use any specific property on the collision blocks. It is interesting to study other practically used variants and understand if similar results can be obtained. To this end, we hope that the techniques we develop in this work will be helpful.

From a theoretical perspective, no single function can be collision-resistant (in the plain model), as a non-uniform attacker can trivially hardwire a collision. This is why collision resistance is considered with

respect to a family of hash functions indexed by a key called *salt*. The salt is chosen after the attacker is fixed (and so is the non-uniform advice about the family of functions). Still, in practice, a single hash function is typically defined by fixing an IV, making it insecure against non-uniform attackers. This contrasts with how we define the collision resistance game, where the IV is chosen randomly after the preprocessing phase. Thus, it may seem that the expensive preprocessing needed for attacks in our model does not represent real-life scenarios. However, often the hash function used in a particular application (relying on collision-resistance) is *salted* by prepending a random salt value to the input. One such well-known application is *password hashing* [ST79]. Such salting essentially corresponds to the random-IV setting considered here, and, therefore, the attack becomes relevant again.

The primary motivation for our work is to make progress towards the STB conjecture, which we view as a fundamental problem. To this end, we focused on asymptotic bounds as a function of S, T, and B. We hope that the concrete bounds could be improved in future works, affecting design choices of real-life hash functions.

2 Our Techniques

In this section, we provide a high-level overview of our techniques. Both of our results follow a similar high-level rationale, and thus throughout this overview, we mainly focus on the techniques for proving the STB conjecture whenever $B \in O(1)$ (Theorem 1.1). Towards the end, we describe the additional ideas needed to obtain the result for $SB \ll T$ (Theorem 1.2).

Before explaining the ideas, let us describe the challenge more precisely. We are given a compressing function $h : [N] \times [M] \to [N]$, modeled as a random oracle, and we want to upper bound the probability of a non-uniform attacker in finding a collision in an MD_h instance with a random salt. We model non-uniform attackers by thinking of them as two-stage adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. The offline part $\mathcal{A}_1$ is unbounded in running time, and its only restriction is that it can output only S bits. This output is the non-uniform advice given to the online part $\mathcal{A}_2$ which is then allowed to make up to T queries after which it must output a B-block collision for MD_h and terminate. We assume unbounded running time for both parts $\mathcal{A}_1$ and $\mathcal{A}_2$ and only restrict the output-size for $\mathcal{A}_1$ and the number of queries for $\mathcal{A}_2$. We refer to such a two-stage adversary $(\mathcal{A}_1, \mathcal{A}_2)$ as a (S, T)-adversary. In the context of length-B collisions in MD_h, the game is as follows:

- $\mathcal{A}_1$ has unbounded access to h and it outputs $\sigma \in \{0, 1\}^S$

- $\mathcal{A}_2$ gets σ as input along with a random salt $a \in [N]$,

- $\mathcal{A}_2$ outputs α, α'.

- $\mathcal{A}$ wins if $\alpha \neq \alpha'$, α, α' consist of $\leqslant B$ blocks, and $\mathsf{MD}_h(a, \alpha) = \mathsf{MD}_h(a, \alpha')$.

There are essentially two main generic approaches known in the literature for proving bounds of this sort. The first is the so-called *pre-sampling* technique, originally due to Unruh [Unr07], and the second is a compression argument. The first technique reduces the problem from considering general hash functions and adversaries $(\mathcal{A}_1, \mathcal{A}_2)$ as above to a simpler model (and associated attacker) called the *bit-fixing* model. The advantage of the latter model is that it is typically easier to analyze and results in clean proofs. The second technique is based on a simple information-theoretic idea that *random bits cannot be compressed.*[4] Thus, an attacker that succeeds in finding collisions is used to compress some random information that is used in the game, and thereby contradiction is reached. This technique, while being extremely influential in many fields and problems in computer science (e.g., "Algorithmic Lovász Local Lemma" [MT10], lower bounds on cryptographic constructions [GT00, DTT10, GT20], analyzing hardness of problems in the non-uniform setting [DGK17, CK18] and time-space tradeoffs for quantum algorithms [CGLQ20]), often results in technical and complex proofs.

It would have been convenient if any non-trivial bound on our problem could be obtained using the bit-fixing technique. Unfortunately, Akshima et al. [ACDW20] observed that finding short collisions is relatively

---

[4]Specifically, it is impossible to save w bits of information about a random string, except with probability 2^{-w}.

easy in the bit-fixing model. Hence, the only remaining potentially helpful technique is based on compression. Indeed, Akshima et al. [ACDW20], as their primary technical contribution, managed to carry out such an argument for the particular case of $B = 2$, and already then their proof is highly non-trivial and consists of a tedious case analysis. We distill some of the main ideas underlying their general framework approach next—this will be useful for us, as well.

The framework. We reduce the task of handling arbitrary (S, T)-adversaries to the problem of handling (S, T)-adversaries, where the preprocessing part $\mathcal{A}_1$ is degenerate and outputs a fixed string σ, independent of h. Specifically, we define a game, parameterized by $u \in \mathbb{N}_{>0}$, where $\mathcal{A}_2$ has an arbitrary size S string σ hard-coded, and its goal is to find a collision relative to a given salt. $\mathcal{A}_2$ wins the game if it succeeds in finding a collision when executed with *every one* of u uniformly random salts (there is no $\mathcal{A}_1$ in this game). The reduction shows that if $\mathcal{A}_2$ has advantage ϵ in the modified game, then the best advantage of an (S, T)-adversary in the original game is (roughly) $O(\epsilon^{1/u})$ for $u \approx S$. This reduction, formalized in Lemma 4.1, is adapted from Akshima et al. [ACDW20] and it uses the beautiful "constructive Chernoff bound" of Impagliazzo and Kabanets [IK10].[5]

The advantage of considering the new game is that there is no $\mathcal{A}_1$, so it is easier to handle. But, to obtain a meaningful result for the original (S, T) game, say an upper bound of ϵ, we need to prove a somewhat stronger upper bound for the new game, that is, roughly ϵ^u. This means, in other words, that we need to show how to compress about $\log(1/\epsilon)$ bits *per each one of the u salts*. (Actually, keep in mind that it suffices to achieve this on average!) For our target ϵ, we therefore have the following goal.

Main challenge: For every one of the u salts, we need to "save/compress" (on average) roughly the following number of bits:

$$\log \left(\min \left\{ \frac{N}{uT(\log u)^{2(B-2)}B^2}, \frac{N}{T^2} \right\} \right).$$

By impossibility of non-trivial compression, this would imply that $\mathcal{A}_2$ must succeed with probability at most

$$\epsilon \leqslant O \left(\left(\frac{uT(\log u)^{2(B-2)}B^2}{N} + \frac{T^2}{N} \right)^u \right),$$

which would give our result when plugged into the framework.

The compression argument. The random string that we shall compress consists of the set of salts denoted U, as well as the function h. We give encoding and decoding algorithms that use $\mathcal{A}_2$ first to encode the pair (U, h) and then use the result to fully decode them whenever $\mathcal{A}_2$ wins the game. If $\mathcal{A}_2$ wins with good enough probability, the output of the encoding procedure will be non-trivially short with good probability, which is a contradiction.

Remark 2.1. *Akshima et al. [ACDW20] used the same approach for $B = 2$, but their proof does not seem to scale for larger values of B. Specifically, their proof involves an exhaustive case analysis. It seems like a naive attempt to generalize their bound to larger values of B would proliferate the number of cases needed to be analyzed, making it unmanageable. One of our key conceptual insights is a structural characterization of collisions in MD_h that prevents this explosion in the number of cases needed to be handled. While our analysis applies to any B, the number of total cases we consider is roughly the same as Akshima et al.*

A first attempt and a glimpse at the challenge. Let us make a strong (and typically false) assumption: *the adversary $\mathcal{A}_2$ never repeats any query to h across all the u runs*.[6] Since we can assume (w.l.o.g) that if $\mathcal{A}_2$ outputs (α, α') when run on salt a, it has queried h at all values needed to compute $\mathsf{MD}_h(a, \alpha)$ and

---

[5]The use of this reduction is the main (and perhaps only) point of similarity between our proof and [ACDW20]'s.

[6]While we can assume without loss of generality that $\mathcal{A}_2$ does not repeat queries within a single execution (since it is not memory-bounded), it is not very reasonable to assume that it will never repeat queries across different executions on different salts.
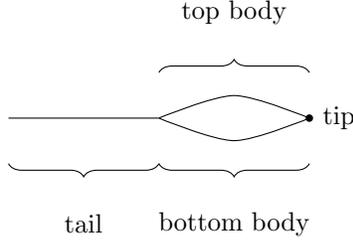
Figure 1: A mouse structure. For ease of visual representation we do not draw the nodes and edges of the graph, instead represent it as a continuous structure.

$\mathsf{MD}_h(a, \alpha')$, we are guaranteed that for each of the salts $u \in U$ there are at least two *distinct* queries which have the same answer, i.e., a collision. The indices of these queries reside in $[T]$ as this is the query complexity of $\mathcal{A}_2$ when executed on the particular salt a. Thus, we could avoid encoding the answer of the second query and instead encode these two indices in T and remove the answer of the second query from evaluations of h. This saves us $\log(N/T^2)$ bits for every salt, giving us even more savings than what we are aiming for. Such compression, in turn, would imply that the birthday attack is optimal, no matter what B is (which makes sense given our assumption but is clearly false for general attackers).

A naive way to get rid of the assumption (that queries never repeat across different u runs) would be to encode the index of the other query among all uT queries made (instead of T). But, this would eventually result in another multiplicative S term in our bound. Namely, we would only be able to save $\log(N/(ST^2))$ bits per salt which is too little for us (as it leads to a trivial upper bound). This motivates us to look more closely at how MD collisions are formed, what kind of queries could be involved, and how we could leverage the fact that collisions are short to get more efficient encoding.

The mouse structure and query types. We consider the graph implicitly formed through the queries made by $\mathcal{A}_2$ when running on salts $a_1, \ldots, a_u \in U$. The nodes of the graph are the possible salts and there is a directed edge from salt a to a' with label α if $h(a, \alpha) = a'$ and this query was made by $\mathcal{A}_2$.

Suppose that $\mathcal{A}_2$, when run on salt a_i, outputs (α, α'). Since $\mathcal{A}_2$ wins on every salt in U, its output on salt a_i, denoted α, α' must satisfy: (1) $\mathsf{MD}_h(a_i, \alpha) = \mathsf{MD}_h(a_i, \alpha')$, (2) $\alpha \neq \alpha'$, and (3) α, α' are at most B blocks long. Without loss of generality, assume that that the last blocks of α and α', respectively, are distinct. Considering the *core* sub-graph of the query graph that is induced by queries made by $\mathcal{A}_2$ that are *required* to evaluate $\mathsf{MD}_h(a, \alpha)$ and $\mathsf{MD}_h(a, \alpha')$, we obtain a structure that we call a **mouse structure**. Important parts of the mouse structure are the tail, top and bottom body, and the tip, as depicted in Fig. 1. Our entire approach is based on studying these mouse structures to come up with encoding strategies.

Given the concept of a mouse structure and our discussion from above about the adversary, possibly repeating queries motivates us to classify each query into one of three types. The first is called NEW and refers to queries made for the first time. The rest of the queries are called *repeated*, and they are further classified into two types, depending on whether they previously appeared in some mouse structure or not. Specifically, a repeated query is called REPEATEDMOUSE if the same query was already made by $\mathcal{A}_2$ when executed on a previous salt, and otherwise, a repeated query is called REPEATEDNONMOUSE.

Intuitively, we want to save bits when the answer of a NEW query was the input salt of a repeated query. A REPEATEDMOUSE query facilitates such savings since we can encode the answer to the query by storing its index along with the index of the previous query and the corresponding index within the mouse structure—a total of $\approx \log(uTB)$ bits instead of $\log N$—which eventually turns into an STB/N term in the upper bound, as conjectured. The problem is with encoding REPEATEDNONMOUSE queries—there seems to be no trivial way to write the index of the previously-made query with less than $\log(uT^2)$ bits, which is too much (since it will eventually turn into a ST^2/N term in the upper bound).

Some "easy" mouse structures. We observe that some cases of mouse structures readily give us a way to have efficient encoding and save sufficiently many bits. We offer three examples to convey intuition on

8
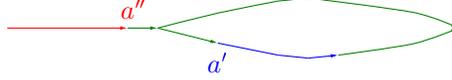
Figure 2: An example of a "hard" mouse structure. The NEW queries are drawn in red, REPEATEDMOUSE queries are drawn in blue and the REPEATEDNONMOUSE queries are drawn in green.

how our analysis is done. Throughout, let us assume that every mouse structure contains at least one NEW query—otherwise, we will ignore this mouse structure altogether.[7]

As a first example, if two NEW queries form the tip of the mouse structure of salt a_j, we can simply encode the index of these queries within the queries made while handling this salt—this uses $2 \log T$ bits instead of $\log N$ bits which is sufficient. As another example, if a self-loop forms the body of the mouse structure and the self loop query is NEW or REPEATEDNONMOUSE, we can simply encode the index of the self loop query in the list of queries used to handle this salt and avoid encoding the answer of the query—this uses $\log(uT)$ bits instead of $\log N$ which is again sufficient. As the last example, suppose the answer to a NEW query is a salt that appeared in some earlier mouse structure. In this case, we can avoid encoding the answer of the NEW query and instead encode the index of the query and which of the salts in the previous mouse structures is the answer—this uses $\log(2uB)$ bits (because there are at most $2B$ salts in mouse structures and at most u mouse structures) instead of $\log N$ which is also sufficient.

Some "hard" mouse structures. The aforementioned easy cases give rise to a relatively easy encoding that results in an optimal STB/N term. Next, we focus on the more complex cases, which cause our bound to have the extra $\approx B \cdot (\log S)^B$ factor.

Assume that there is a mouse structure where there are two salts a' and a'' such that a' is the input salt to a REPEATEDMOUSE query, a'' is the answer to a NEW query, and the path from a'' to a' in the mouse structure consists of only REPEATEDNONMOUSE queries (as shown in Figure 2). By definition, the distance between a' and a'' is at most B.[8] Potential savings could be achieved by not encoding a'', the answer of the NEW query, as it can essentially be extracted from already-observed queries. We can easily encode the appropriate index of the query in the mouse structure and the salt a' using roughly $\log(uB)$ bits, but how can we encode the information about the path back from a' to a''?

The non-triviality is that for each node on this path, there might be many possible ways to reach it among all the different queries that have been already made, i.e., if each node on this path has fan-in m, namely an m-multi-collision, then the natural encoding of the path back would cost at most $\log(m^B)$ bits, specifying which back edge to take for every node. Here, m could be very large, e.g., as large as S or even larger, making the whole result meaningless for most reasonable parameters settings. (We also need to encode the length of the path, which would lead to the additional multiplicative B factor, but we ignore it in the discussion here.)

We get around this problem of m potentially being very large by observing that one of the following two cases holds:

- *Many small multi-collisions:* either the fan-in of every node along the path (i.e., the number of previously-made queries whose result is a node on this path) is smaller than $\log u$, or

- *One large multi-collision:* there is (at least one) node on the path where the fan-in is at least $\log u$.

In the earlier case, we encode the path back as mentioned above, where we write the index of each back edge. This costs us $B \cdot \log \log u$ bits, which eventually translates to an extra $(\log S)^B$ term,[9] as we have in

---

[7]In the technical section, we refer to salts a_j in U that were not the input salt of a query when running $\mathcal{A}_2$ on a_i for $i < j$ as *fresh*. It follows that mouse structures for fresh salts will always have a NEW query. For salts that are not fresh, it is relatively straightforward to achieve some compression by avoiding storing these salts in the encoding of U. For now, the reader can imagine that all salts are fresh for simplicity.

[8]By being slightly more careful, we can show that the distance is $B - 2$ but we ignore this fact for the overview.

[9]Remember that the actual term is $(\log S)^{B-2}$ and that is why the proof of Akshima et al. [ACDW20], in which it was assumed that $B = 2$, did not have an extra term that depends on S.

our main theorem. The question is, what do we do when the second case occurs.

The idea is to leverage the fact that there is a large multi-collision and obtain our savings from a completely different place. Specifically, we remember the index of all the queries involved in the multi-collision and the common answer. A calculation reveals that if the collision consists of $\approx \log u$ edges, we can already save enough. One subtlety is that the same multi-collision might repeat throughout many different mouse structures, and we need to make sure not to double count the savings from a single multi-collision more than what we get. The reason why we do not double count is that a single large-enough (i.e., with $\log u$ edges) multi-collision saves us enough bits for about $\log u$ different structures, and at the same time, such a multi-collision can appear in at most $\log u$ mouse structures.[10]

Let us finally remark that while the above description conveys the main idea underlying our compression strategy, it is somewhat simplified and glossed over many technicalities and the complete specification of all possible cases that our full proof covers. We refer to Section 5 for full details.

Proving the STB conjecture for $S \cdot B \ll T$. The proof of our second upper bound follows the same overall structure. The only difference is, naturally, in the way we encode reverse paths in mouse structures that have nodes corresponding to REPEATEDNONMOUSE queries between the output a'' of a NEW query and a node a' corresponding to a REPEATEDMOUSE query.

In the earlier proof, when we needed to locate the salt a'' from salt a', we encoded the number of edges in between and all the edges on the path. Here, we prove a purely graph-theoretic lemma saying that if for salt a' there are more than $\approx u^2$ salts a'' such that there are $d > 0$ edges on the shortest path back from a' to a'' in the query graph, then there must be $t \geqslant 1$ multi-collisions in the graph, such that the total number of edges involved in the t multi-collisions is at least $\approx u^2$. While the proof of this lemma is a simple inductive argument, it turns out to be extremely helpful for us. Specifically, if there are $t \geqslant 1$ different multi-collisions such that a total of at least $\approx u^2$ different queries are involved in the t multi-collisions, *we can save enough by only encoding these multi-collisions and nothing else.* To prove this fact, we consider the minimum savings we can get from encoding these t multi-collisions. We show using some elementary calculus that if there are $\approx u^2$ queries involved in t different multi-collisions, the minimum saving is more than the total amount of savings we need.

Equipped with this fact, we split our analysis into the two following scenarios.

1. The first is where for each case where we need to encode the location of a'' from a' at a distance d, there are at most $\approx u^2$ salts – here we simply encode the index of the "right" a'' using $\approx \log u^2$ bits.

2. The other scenario is when for at least one case, there are more than $\approx u^2$ salts. Here we can save enough by encoding the t multi-collisions involving at least $\approx \log u^2$ that the graph-theoretic lemma guarantees us. We get enough savings by only encoding these t multi-collisions.

The u^2 term from the first scenario above turns into an additional factor of the form S^2 in the final bound. Due to additional technicalities that we glossed over during this overview, we suffer another multiplicative factor S in our bound, which amounts to having an S^4 term. Full details appear in Section 6.

3   Preliminaries

Let $\mathbb{N} = \{0, 1, 2, \ldots\}$ denote the set of all natural numbers and let $\mathbb{N}_{>0} = \mathbb{N}\backslash\{0\}$. For $N \in \mathbb{N}_{>0}$, let $[N] = \{1, 2, \ldots, N\}$ and for $k \in \mathbb{N}$ such that $k \leqslant N$, let $\binom{[N]}{k}$ denote the set of k-sized subsets of $[N]$. For a set X, let $|X|$ be its size and X^+ denote one or more elements of X. We denote $\mathsf{Fcs}(D, R)$ the set of all functions mapping elements in D to the elements of R. We let $x \leftarrow_\$ \mathcal{D}$ denote sampling x according to the distribution $\mathcal{D}$. We let $*$ denote a wildcard element. For example $(*, z) \in L$ is true if there is an ordered pair in L where z is the second element (the type of the wildcard element shall be clear from the context). If D is a set, we overload notation and let $x \leftarrow_\$ D$ denote uniformly sampling from the elements of D. For a bit-string s we use $|s|$ to denote the number of bits in s.

---

[10]We mention that multi-collisions in hash functions have been studied on their own right (e.g., [Jou04, BDRV18, BKP18, KNY18, Din20]), but our context is totally different.

Game $\mathsf{G}^{\mathsf{ai\text{-}cr}}_{N,M,B}(\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2))$	Subroutine $\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2))$
1. $h \leftarrow\!\!\$\ \mathsf{Fcs}([N] \times [M], [N])$	1. $\sigma \leftarrow\!\!\$\ \mathcal{A}_1(h)$
2. $a \leftarrow\!\!\$\ [N]$	2. $(\alpha, \alpha') \leftarrow\!\!\$\ \mathcal{A}_2^h(\sigma, a)$
3. Return $\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A})$	3. Return true if:
	$\quad$ (a) $\alpha \neq \alpha'$,
	$\quad$ (b) α, α' consist of $\leqslant B$ blocks from $[M]$, and
	$\quad$ (c) $\mathsf{MD}_h(a, \alpha) = \mathsf{MD}_h(a, \alpha')$
	4. Else, return false

Figure 3: The bounded-length collision resistance game of salted MD hash in the AI-ROM, denoted $\mathsf{G}^{\mathsf{ai\text{-}cr}}_{N,M,B}$.

When referring to directed graphs in this paper, we mean directed multi-graphs, i.e., these directed graphs might have parallel edges. All logarithms in this paper are for base 2 unless otherwise stated.

Merkle-Damgård (MD) hashing. For $N, M \in \mathbb{N}_{>0}$, let $h : [N] \times [M] \to [N]$ be a compression function. We recursively define Merkle-Damgård (MD) hashing $\mathsf{MD}_h : [N] \times [M]^+ \to [N]$ as $\mathsf{MD}_h(a, \alpha) = h(a, \alpha)$ for $a \in [N], \alpha \in [M]$ and

$$\mathsf{MD}_h(a, (\alpha_1, \ldots, \alpha_B)) = h(\mathsf{MD}_h(a, (\alpha_1, \ldots, \alpha_{B-1})), \alpha_B)$$

for $a \in [N]$ and $\alpha_1, \ldots, \alpha_B \in [M]$. The elements of $[M]$ shall be referred to as *blocks* and a, the first input to MD_h is referred to as the *salt*. This is the same abstraction of plain MD hashing as the one used in [ACDW20].

Auxiliary-input Random Oracle Model (AI-ROM). We use the Auxiliary-Input Random Oracle Model (AI-ROM) introduced by Unruh [Unr07] to study non-uniform adversaries in the Random Oracle Model. This model is parameterized by two non-negative integers S and T and an adversary $\mathcal{A}$ is divided into two stages $(\mathcal{A}_1, \mathcal{A}_2)$. Adversary $\mathcal{A}_1$, referred to as the preprocessing phase of $\mathcal{A}$, has unbounded access to the random oracle h and outputs an S-bit auxiliary input σ. Adversary $\mathcal{A}_2$, referred to as the online phase, gets σ as input and can make T queries to h, attempting to accomplish some goal involving the function h. Formally, we say that $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is an (S, T)-AI adversary if $\mathcal{A}_1$ outputs S bits and $\mathcal{A}_2$ issues T queries to its oracle. We next formalize the salted-collision resistance of MD hash functions in AI-ROM.

Salted short collision resistance of MD in AI-ROM. We formalize the hardness of bounded-length collision resistance of salted MD hash functions in the AI-ROM. The game is parametrized by N, M, and B. The game first samples a function h uniformly at random from $\mathsf{Fcs}([N] \times [M], [N])$ and a salt a uniformly at random from $[N]$. Then, $\mathcal{A}_1$ is given unbounded access to h, and it outputs σ. At this time, $\mathcal{A}_2$ is given the auxiliary input σ, a salt a, as well as oracle access to h, and it needs to find $\alpha \neq \alpha'$ such that (1) $|\alpha|, |\alpha'| \leqslant B \cdot M$, and (2) $\mathsf{MD}_h(a, \alpha) = \mathsf{MD}_h(a, \alpha')$. This game, denoted $\mathsf{G}^{\mathsf{ai\text{-}cr}}_{N,M,B}$, is explicitly written in Fig. 3. In Fig. 3, we write the adversary's execution in its own subroutine only for syntactical purposes (as we shall use it later in our proof).

Definition 3.1 (AI-CR Advantage)**.** For parameters $N, M, B \in \mathbb{N}$, the advantage of an adversary $\mathcal{A}$ against the bounded-length collision resistance of salted MD in the AI-ROM is

$$\mathsf{Adv}^{\mathsf{ai\text{-}cr}}_{\mathsf{MD},N,M,B}(\mathcal{A}) = \Pr\left[\mathsf{G}^{\mathsf{ai\text{-}cr}}_{N,M,B}(\mathcal{A}) = \mathsf{true}\right]$$

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

$$\mathsf{Adv}^{\mathsf{ai\text{-}cr}}_{\mathsf{MD},N,M,B}(S, T) = \max_{\mathcal{A}} \left\{\mathsf{Adv}^{\mathsf{ai\text{-}cr}}_{\mathsf{MD},N,M,B}(\mathcal{A})\right\},$$

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

The compression lemma. Our proof uses the well-known technique of finding an "impossible compression". The main idea, formalized in the following proposition, is that it is impossible to compress a random element in set $\mathcal{X}$ to a string shorter than $\log|\mathcal{X}|$ bits long, even relative to a random string.

Proposition 3.2 (E.g., [DTT10])**.** Let Encode be a randomized map from $\mathcal{X}$ to $\mathcal{Y}$ and let Decode be a randomized map from $\mathcal{Y}$ to $\mathcal{X}$ such that

$$\Pr_{x \leftarrow\$ \mathcal{X}} \left[\mathsf{Decode}(\mathsf{Encode}(x)) = x\right] \geqslant \epsilon.$$

Then, $\log|\mathcal{Y}| \geqslant \log|\mathcal{X}| - \log(1/\epsilon)$.

Constructive Chernoff bound. We use the following "constructive" Chernoff bound taken from Akshima et al. [ACDW20] which in turn is a variant of a similar lemma of Impagliazzo and Kabanets [IK10]. We refer to [ACDW20, Section 9] for a proof of this proposition.

Proposition 3.3 ([ACDW20, Theorem 1])**.** Let $N \in \mathbb{N}_{>0}$, $u \in \mathbb{N}$ such that $u \leqslant N$, $0 < \delta < 1$ and let $X_1, \ldots, X_N$ be $0/1$ random variables. Let $X = X_1 + \cdots + X_N$ and let U be an independent random subset of $[N]$ of size u.

$$\text{If} \quad \Pr\left[\bigwedge_{i \in U} X_i\right] \leqslant \delta^u, \quad \text{then} \quad \Pr\left[X \geqslant \max\{6\delta N, u\}\right] \leqslant 2^{-u}.$$

4 The Framework: Reducing the Problem to a Multi-instance Collision Finder

Our task here is to upper-bound the advantage of an adversary in finding a short collision in a salted MD, according to the game $\mathsf{G}^{\mathsf{ai-cr}}_{N,M,B}$ described in Figure 3. First, without loss of generality, in what follows, we assume that the adversary is deterministic. This follows since we can transform any probabilistic attacker into a deterministic one by hard-wiring the best randomness (see Adleman [Adl78]).

We reduce the task of bounding the advantage of an attacker in finding a short collision in a salted MD, according to the game $\mathsf{G}^{\mathsf{ai-cr}}_{N,M,B}$, to a "multi-instance" game where the adversary does not have a preprocessing phase but instead only has a non-uniform auxiliary input, chosen *before* the random oracle h. The latter game is easier to analyze. Although the statement and reduction below were implicit in the work of Akshima et al. [ACDW20], we make it formal and hopefully useful for future works.

We define the following "multi-instance" game $\mathsf{G}^{\mathsf{mi-cr}}_{N,M,B,u}(\sigma, \mathcal{A}_2)$, where the preprocessing part of the adversary $\mathcal{A}_1$ is degenerate and outputs the fixed string σ. More precisely, the game has the following steps:

1. $h \leftarrow\$ \mathsf{Fcs}([N] \times [M], [N])$

2. $U \leftarrow\$ \binom{[N]}{u}$

3. Define $\mathcal{A}_1$ to be the algorithm that always outputs the string σ.

4. Return true if $\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)) = \mathsf{true}$ for every $a \in U$. Otherwise, return false.

For a string σ and an adversary $\mathcal{A}_2$, define

$$\mathsf{Adv}^{\mathsf{mi-cr}}_{\mathsf{MD},N,M,B,u}(\sigma, \mathcal{A}_2) = \Pr\left[\mathsf{G}^{\mathsf{mi-cr}}_{N,M,B,u}(\sigma, \mathcal{A}_2)\right].$$

Lemma 4.1. *Fix $N, M, B, S, T, u \in \mathbb{N}_{>0}$. Then,*

$$\mathsf{Adv}^{\mathsf{ai-cr}}_{\mathsf{MD},N,M,B}(S,T) \leqslant 6 \cdot \left(\max_{\sigma, \mathcal{A}_2}\left\{\mathsf{Adv}^{\mathsf{mi-cr}}_{\mathsf{MD},N,M,B,u}(\sigma, \mathcal{A}_2)\right\}\right)^{\frac{1}{u}} + 2^{S-u},$$

where the maximum is taken over all $\sigma \in \{0,1\}^S$ and T-query algorithms $\mathcal{A}_2$.

Proof. Define

$$\gamma := \left(\max_{\sigma, \mathcal{A}_2} \left\{ \mathsf{Adv}_{\mathsf{MD}, N, M, B, u}^{\mathsf{mi\text{-}cr}}(\sigma, \mathcal{A}_2) \right\} \right)^{\frac{1}{u}}.$$

We need to upper bound $\Pr_{h,a} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \right]$ for every (S, T)-AI adversary $\mathcal{A}$. We say that h is *easy* for adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ if $\mathcal{A}$ wins the $\mathsf{AI\text{-}CR}_{h,a}$ game on a random $a \leftarrow\!\!{\scriptstyle\$}\, [N]$ with probability 6γ. That is,

$$h \text{ is easy for } \mathcal{A} \iff \Pr_{a \leftarrow\!\!{\scriptstyle\$}\, [N]} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \right] \geqslant 6\gamma.$$

By the law of total probability,

$$\Pr_{h,a} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \right] = \Pr_{h,a} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \mid h \text{ is easy for } \mathcal{A} \right] \cdot \Pr_h[h \text{ is easy for } \mathcal{A}] +$$

$$\Pr_{h,a} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \mid h \text{ is not easy for } \mathcal{A} \right] \cdot \Pr_h[h \text{ is not easy for } \mathcal{A}]$$

$$\leqslant \Pr_h[h \text{ is easy for } \mathcal{A}] + \Pr_{h,a} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \mid h \text{ is not easy for } \mathcal{A} \right]$$

$$\leqslant \Pr_h[h \text{ is easy for } \mathcal{A}] + 6\gamma.$$

The lemma readily follows from the following claim where we prove an upper bound on $\Pr_h[h \text{ is easy for } \mathcal{A}]$.

Claim 4.2. *For any (S, T)-AI adversary $\mathcal{A}$, it holds that $\Pr_h[h \text{ is easy for } \mathcal{A}] \leqslant 2^{S-u}$.*

Proof. Fix an adversary as in the statement $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. Define $\mathcal{A}_\sigma$ the canonical preprocessing adversary that outputs the string $\sigma \in \{0, 1\}^S$, ignoring its input. Denote $\mathcal{A}_\sigma$ the (S, T)-AI adversary $(\mathcal{A}_\sigma, \mathcal{A}_2)$. By definition and by a union bound, it holds that

$$\Pr_h[h \text{ is easy for } \mathcal{A}] \leqslant \Pr_h \left[\exists \sigma \in \{0, 1\}^S : h \text{ is easy for } \mathcal{A}_\sigma \right]$$

$$\leqslant 2^S \cdot \max_{\sigma \in \{0,1\}^S} \Pr_h \left[h \text{ is easy for } \mathcal{A}_\sigma \right]. \tag{1}$$

For a fixed $a \in [N]$, $\sigma \in \{0, 1\}^S$, and $h \in \mathsf{Fcs}([N] \times [M], [N])$, define the indicator $X_{h,a}^\sigma = 1$ if and only if $\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}_\sigma) = \mathsf{true}$. By definition and with this notation, for every $\sigma \in \{0, 1\}^S$, it holds that

$$\Pr_h \left[h \text{ is easy for } \mathcal{A}_\sigma \right] = \Pr_h \left[\Pr_{a \leftarrow\!\!{\scriptstyle\$}\, [N]} \left[\mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}_\sigma) = \mathsf{true} \right] \geqslant 6\gamma \right]$$

$$= \Pr_h \left[\Pr_{a \leftarrow\!\!{\scriptstyle\$}\, [N]} \left[X_{h,a}^\sigma \right] \geqslant 6\gamma \right]$$

$$= \Pr_h \left[\sum_{a \in [N]} X_{h,a}^\sigma \geqslant 6N\gamma \right]. \tag{2}$$

By definition, $\mathsf{Adv}_{\mathsf{MD}, N, M, B, u}^{\mathsf{mi\text{-}cr}}(\sigma, \mathcal{A}_2)$ is the advantage of the T-query attacker $\mathcal{A}_2$ given auxiliary input σ in finding collisions in salted MD relative to h and on u uniformly random salts denoted U. Written differently, for every $\sigma \in \{0, 1\}^S$ and T query adversary $\mathcal{A}_2$, it holds that

$$\Pr_{h,U} \left[\bigwedge_{a \in U} X_{h,a}^\sigma \right] = \mathsf{Adv}_{\mathsf{MD}, N, M, B, u}^{\mathsf{mi\text{-}cr}}(\sigma, \mathcal{A}_2)$$

$$\leqslant \max_{\sigma^*, \mathcal{A}_2^*} \left\{ \mathsf{Adv}_{\mathsf{MD}, N, M, B, u}^{\mathsf{mi\text{-}cr}}(\sigma^*, \mathcal{A}_2^*) \right\} = \gamma^u,$$

13

where the last equality follows from the definition of γ. Applying Proposition 3.3, we get that for every $\sigma \in \{0,1\}^S$, it holds that

$$\Pr_h \left[\sum_{a \in [N]} X_{h,a}^\sigma \geqslant 6N\gamma \right] \leqslant 2^{-u}.$$

Plugging this back into (2), we get that $\Pr_h\left[h \text{ is easy for } \mathcal{A}_\sigma\right] \leqslant 2^{-u}$. In turn, plugging the latter back into (1), we get that $\Pr_{h,a}\left[\text{AI-CR}_{h,a}(\mathcal{A}) = \text{true}\right] \leqslant 2^{S-u}$, as needed. ∎

∎

5 Proving the STB Conjecture for $B \in O(1)$

This section proves an upper bound on the advantage of any auxiliary-input adversary in the bounded-length collision resistance game of salted MD hash in the AI-ROM. The main theorem is stated next.

Theorem 5.1. *Let $C = 2^{16} \cdot 6 \cdot e^2$. For any $N, M, B, S, T \in \mathbb{N}_{>0}$ and fixing $\hat{S} := S + \log N$, it holds that*

$$\text{Adv}_{\text{MD},N,M,B}^{\text{ai-cr}}(S,T) \leqslant C \cdot \max\left\{ \left(\frac{\hat{S}TB^2 \left(\frac{3e \log \hat{S}}{\log \log \hat{S}} \right)^{2(B-2)}}{N} \right), \left(\frac{T^2}{N} \right) \right\} + \frac{1}{N} \ .$$

Theorem 5.1 follows as a direct corollary of Lemma 4.1 together with the following lemma and setting $u = S + \log N$.

Lemma 5.2 (Hardness for a multi-instance collision finder)**.** *Fix N, M, B, S, T, $u \in \mathbb{N}_{>0}$ and $\sigma \in \{0,1\}^S$. Then, for any $\mathcal{A}_2$ that makes at most T queries to its oracle, it holds that*

$$\text{Adv}_{\text{MD},N,M,B,u}^{\text{mi-cr}}(\sigma,\mathcal{A}_2) \leqslant$$
$$\left(2^{16}e^2 \cdot \max\left\{ \left(\frac{uTB^2(3e \log u/\log \log u)^{2(B-2)}}{N} \right), \left(\frac{T^2}{N} \right) \right\} \right)^u .$$

The rest of this section is devoted to the proof of Lemma 5.2. Unlike the proof of Lemma 4.1, the proof of this lemma is novel and differs completely from that of Akshima et al. [ACDW20]. The key conceptual insight is a structural characterization of collisions in MD_h that prevents the explosion in the number of cases that [ACDW20] faced during the case analysis.

We are interested in bounding the advantage of the best strategy, i.e., a pair $(\sigma, \mathcal{A}_2)$ where $\sigma \in \{0,1\}^S$ is a fixed string and $\mathcal{A}_2$ is a T-query algorithm, of finding bounded-length collisions in a salted MD with respect to the game $\text{G}_{N,M,B,u}^{\text{mi-cr}}(\sigma, \mathcal{A}_2)$. Recall that in this game, $\mathcal{A}_2$ needs to find proper collisions for u randomly chosen salts, denoted U. The main idea in the proof is to use any such adversary $(\sigma, \mathcal{A}_2)$ to represent the function h *as well as* the set of random salts U with as few bits as possible. If the adversary is "too good to be true," we will get an impossible representation, contradicting Proposition 3.2.

Non-trivial range. If either

$$\frac{T^2}{N} > 1 \quad \text{or} \quad \frac{uTB^2(3e \log u/\log \log u)^{2(B-2)}}{N} > 1,$$

then Lemma 5.2 is trivially true. Hence, from now on we assume that both of the above left hand side terms are upper bounded by 1.

Setup. Denote

$$\zeta^* := \left(2^{16}e^2 \cdot \max\left\{ \left(\frac{uTB^2(3e \log u/\log \log u)^{2(B-2)}}{N} \right), \left(\frac{T^2}{N} \right) \right\} \right)^u .$$

14

Assume the existence of an adversary $\mathcal{A} = (\sigma, \mathcal{A}_2)$, where $\sigma \in \{0,1\}^S$ is a string and $\mathcal{A}_2$ is a T-query adversary, that contradict the inequality stated in the lemma. That is, there is $\zeta > \zeta^*$ such that

$$\mathsf{Adv}^{\mathsf{mi\text{-}cr}}_{\mathsf{MD},N,M,B,u}(\mathcal{A}) := \zeta > \zeta^*. \tag{3}$$

Define $\mathcal{G}$ to be the set of functions-sets of salts pairs for which the attacker succeeds in winning the game for every salt in the set relative to the function, That is,

$$\mathcal{G} = \left\{ (U,h) \,\middle|\, \begin{array}{c} U \in \binom{[N]}{u}, \\ h \in \mathsf{Fcs}([N] \times [M], [N]), \end{array} \forall a \in U : \mathsf{AI\text{-}CR}_{h,a}(\mathcal{A}) = \mathsf{true} \right\}.$$

Recall that ζ is defined to be the advantage of $\mathcal{A}$ in the game $\mathsf{G}^{\mathsf{mi\text{-}cr}}_{N,M,B,u}(\mathcal{A})$ in which h and U are chosen uniformly, and then $\mathcal{A}$ needs to find a collision with respect to every one of the u salts in U. Therefore,

$$|\mathcal{G}| = \zeta \cdot \binom{N}{u} \cdot N^{MN}.$$

In what follows we define an encoding and a decoding procedure such that the encoding procedure gets as input U, h such that $U \in \binom{[N]}{u}$ and $h \in \mathsf{Fcs}([N] \times [M], [N])$, and it outputs an L bit string, where $L = \log\left(\zeta^* \cdot \binom{N}{u} \cdot N^{MN}\right)$. The decoding procedure takes as input the string L and outputs U^*, h^*. It will hold that $U^* = U$ and $h^* = h$ with probability ζ.[11] Using Proposition 3.2, this would give us that

$$\log \zeta \leqslant L - \log\left(\binom{N}{u} \cdot N^{MN} \right) \implies \zeta \leqslant \zeta^*$$

which is a contradiction to the assumption (see (3)).

Notation and Definitions. Fix $(U,h) \in \mathcal{G}$. Let $U = \{a_1, \ldots, a_u\}$ where the a_i's are ordered lexicographically. Let $\mathsf{Qrs}(a) \in ([N] \times [M])^T$ be the list of queries that $\mathcal{A}_2$ makes to h when executed with input (σ, a). Namely, for $a \in [N]$,

$$\mathsf{Qrs}(a) = \left\{ (a', \alpha') \in [N] \times [M] \mid \mathcal{A}_2(\sigma, a) \text{ queries } h \text{ on } (a', \alpha') \right\} .$$

Note that $\mathsf{Qrs}(a)$ is indeed a set as we can assume (without loss of generality) that $\mathcal{A}_2$ never repeats queries in a single execution (since $\mathcal{A}_2$ can just store all of its past queries).

We say that $a' \in \mathsf{Slts}(a)$ if there is some $\alpha' \in [M]$ such that (a', α') is an entry in $\mathsf{Qrs}(a)$. Namely, for $a, a' \in [N]$,

$$a' \in \mathsf{Slts}(a) \iff \exists \alpha' \in [M] \text{ s.t. } (a', \alpha') \in \mathsf{Qrs}(a).$$

We define the set of *fresh* salts in U. A salt a_i for $i \in [u]$ is called fresh if it was never used as the salt in any query performed by $\mathcal{A}_2$ while being executed on salts a_j for $j \leqslant i - 1$ which are fresh. The first salt a_1 is always fresh. A salt a_i for $i \geqslant 2$ is fresh if for any fresh a_j for $j \leqslant i - 1$, $a_i \notin \mathsf{Slts}(a_j)$. Namely, denoting the set of fresh salts by U_{fresh}, we have the following inductive (on $i \in [u]$) definition:

$$a_i \in U_{\mathsf{fresh}} \iff \forall j \leqslant i - 1, a_j \in U_{\mathsf{fresh}} : a_i \notin \mathsf{Slts}(a_j).$$

Looking ahead, we define U_{fresh} like this because we run $\mathcal{A}_2$ on the salts in U_{fresh} in lexicographical order, and this definition ensures that each salt that $\mathcal{A}_2$ is executed on was not queried by it previously. Denote

$$F := |U_{\mathsf{fresh}}| \quad \text{and} \quad U_{\mathsf{fresh}} = \{a'_1, \ldots, a'_F\} \text{ (ordered lexicographically)}.$$

---

[11]Essentially, we will show that for all $(U,h) \in |\mathcal{G}|$, if the encoding procedure produces output L, then the decoding procedure on input L outputs U^*, h^* such that $U^* = U$ and $h^* = h$.

Denote

$$\forall i \in [F] : \mathsf{Q}_i := \mathsf{Qrs}(a'_i) \quad \text{and} \quad \mathsf{Q}_{\mathsf{fresh}} := \mathsf{Q}_1 \parallel \ldots \parallel \mathsf{Q}_F,$$

where $\parallel$ is the concatenation operator. Let $\mathsf{Q}_{\mathsf{fresh}}[r]$ be the rth query in the list $\mathsf{Q}_{\mathsf{fresh}}$. Note that $r \in [F \cdot T]$. For every $a \in U \backslash U_{\mathsf{fresh}}$, let t_a be the minimum value such that $\mathsf{Q}_{\mathsf{fresh}}[t_a]$ is a query with salt a. Define the set of *prediction* queries as

$$\mathsf{P} := \{ t_a \mid a \in U \backslash U_{\mathsf{fresh}} \}.$$

The encoding algorithm will output $U_{\mathsf{fresh}}, \mathsf{P}$, which suffices to recover the set U by running $\mathcal{A}_2$.

We let $\tilde{h}$ be the list of $h(a, \alpha)$ values when executed on distinct queries in $\mathsf{Q}_{\mathsf{fresh}}$, in the same order as they appear in $\mathsf{Q}_{\mathsf{fresh}}$, followed by the evaluation of h on the following values in lexicographical order of the inputs.

$$\{ (a, \alpha) : a \in [N], \alpha \in [M] \} \backslash \mathsf{Q}_{\mathsf{fresh}}.$$

Therefore, $\tilde{h}$ is initialized to contain the evaluation of h at all points in its domain. Looking ahead, in the encoding procedure, we will remove elements from $\tilde{h}$ as needed to compress h.

Function and Query graphs. A notion that will be useful is that of a "function graph".

Definition 5.3 (Function graph)**.** For a function $h : [N] \times [M] \to [N]$, consider the following directed graph: it has N nodes labelled with elements of $[N]$ and each node has exactly M outgoing edges, each labelled with elements of $[M]$. There is an edge from node a_i to a_j labelled α if and only if $h(a_i, \alpha) = a_j$.

We define the notion of query graph for an adversary as follows.

Definition 5.4 (Query graph)**.** Execution of an adversary $\mathcal{A}_2$ on salts $a'_1, \ldots, a'_F$ defines a query graph as follows. Initially the graph is empty. Whenever $\mathcal{A}_2$ queries (a, α) to h, we add a node with label a if not already present and add an edge $(a, h(a, \alpha))$ with label α if not already present.

Fact 5.5. *The query graph is always a sub-graph of the function graph of h.*

Structure of collisions: The mouse structure. Since adversary $\mathcal{A}_2$ succeeds on all of the salts in U, it holds that for every $j \in [F]$, the output of the adversary is (α_j, α'_j) such that $\alpha_j \neq \alpha'_j$, $\mathsf{MD}_h(a_j, \alpha_j) = \mathsf{MD}_h(a_j, \alpha'_j)$ and both α_j, α'_j are at most B blocks long. We can assume without loss of generality that the last blocks of α_j and α'_j are distinct (because otherwise, we can trim α_j, α'_j to obtain a shorter collision). The evaluations of h in order to compute $\mathsf{MD}_h(a'_j, \alpha_j)$ and $\mathsf{MD}_h(a'_j, \alpha'_j)$ induce a structure that we call a *mouse structure* as shown in Fig. 4. More explicitly, suppose the output of the $\mathcal{A}_2$ is $(\alpha_j = (\alpha_{j,1}, \ldots, \alpha_{j,B_1}), \alpha'_j = (\alpha'_{j,1}, \ldots, \alpha'_{j,B_2}))$ for $B_1, B_2 \leqslant B$ such that $\alpha_{j,i} = \alpha'_{j,i}$ for all $1 \leqslant i \leqslant k$ where $k \geqslant 0$. Define $(x_1, \ldots, x_k, y_1, \ldots, y_{k'}, z_1, \ldots, z_{k''})$ where $k' = B_1 - k, k'' = B_2 - k$ as follows.

$$x_1 = h(a'_j, \alpha_{j,1}) , \; x_i = h(x_{i-1}, \alpha_{j,i}) \text{ for } 1 < i \leqslant k$$

$$y_1 = \begin{cases} h(a'_j, \alpha_{j,1}) & \text{if } k = 0 \\ h(x_k, \alpha_{j,k+1}) & \text{otherwise} \end{cases} , \; y_i = h(y_{i-1}, \alpha_{j,i+k}) \text{ for } 1 < i \leqslant B_1 - k$$

$$z_1 = \begin{cases} h(a'_j, \alpha'_{j,1}) & \text{if } k = 0 \\ h(x_k, \alpha'_{j,k+1}) & \text{otherwise} \end{cases} , \; z_i = h(z_{i-1}, \alpha'_{j,i+k}) \text{ for } 1 < i \leqslant B_2 - k$$

Then $(x_1, \ldots, x_k, y_1, \ldots, y_{k'}, z_1, \ldots, z_{k''})$ form a mouse structure as specified in Figure 4. Without loss of generality, we can assume that the mouse structure is present in the query graph of $\mathcal{A}_2$ before it outputs the answer for salt a'_j. We refer to this structure as the mouse structure for salt a'_j.

We define $\mathsf{MouseQrs}_j$ to be the set of queries in the mouse structure. Similarly, we define $\mathsf{MouseSlts}_j$ as the set of salts that comprise the mouse structure. By definition, $1 \leqslant |\mathsf{MouseQrs}_j|, |\mathsf{MouseSlts}_j| \leqslant 2B$.

16
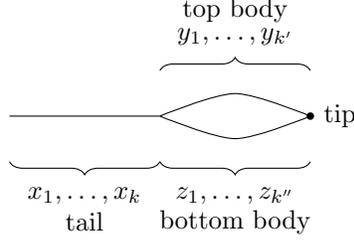
Figure 4: A mouse structure. For ease of visual representation we do not draw the nodes and edges of the graph, instead represent it as a continuous structure.

Classifying queries. We classify every one of the queries in $\mathsf{Q}_{\mathsf{fresh}}$ into one of 3 types by scanning through them in order. Recall that $\mathsf{Q}_{\mathsf{fresh}}$ consists of F blocks, each consisting of T queries. Each block contains a mouse structure as in Figure 4. The first type of query is called NEW. A NEW query did not appear in any previous block. Non-NEW queries are called *repeated* and they are classified further into one of 2 types: REPEATEDMOUSE, and REPEATEDNONMOUSE. A query (a, α) would be a REPEATEDMOUSE query if it was made as part of a mouse structure during some earlier salt in U_{fresh}. Lastly, a REPEATEDNONMOUSE query is one that was made before but is not part of any mouse structure. That is,

1. NEW: A query with index $r \in [F \cdot T]$ is NEW if there does not exist $r' < r$ such that $\mathsf{Q}_{\mathsf{fresh}}[r'] = \mathsf{Q}_{\mathsf{fresh}}[r]$.

2. A non-NEW queries is called *repeated*. We classify the latter into two subcategories:

   (a) REPEATEDMOUSE: A query in Q_j with index $s \in [T]$ such that $h(\mathsf{Q}_j[s]) = a$ is REPEATEDMOUSE if it is not NEW and $\mathsf{Q}_j[s] \in \mathsf{MouseQrs}_i$ for some $i < j$.

   (b) REPEATEDNONMOUSE: A query in Q_j with index $s \in [T]$ such that $h(\mathsf{Q}_j[s]) = a$ is REPEATEDNONMOUSE if it is not NEW, $\mathsf{Q}_j[s] \notin \mathsf{MouseQrs}_i$ for all $i < j$.

Note that this classification covers all queries made during execution. The following is a simple observation.

Claim 5.6. *Every mouse structure has at least one NEW query.*

Proof. For every $j \in [F]$, the queries in Q_j with salt a'_j are necessarily NEW because we defined U_{fresh} to contain a'_j's that were not queried earlier by $\mathcal{A}_2$ when run on a'_i for $i < j$, and also assumed that $\mathcal{A}_2$ does not repeat queries during a single execution. ∎

5.1 The Compression Argument

As mentioned, our goal is to compress (U, h), and we will achieve this by using our collision finding adversary $\mathcal{A}_2$. The encoding procedure shall output the set U_{fresh}, the set P, the list $\tilde{h}$ with some entries removed and additional lists and sets. We will be describing the details of these lists and sets below and which entries we remove from $\tilde{h}$.

It is instructive to think about how decoding would work to understand the intuition behind our encoding strategy. The decoding procedure will run $\mathcal{A}_2$ on the salts in U_{fresh} and answer its queries using the entries in $\tilde{h}$ and the other lists and sets it gets as input. We need to encode enough information in these lists and sets so that the decoding procedure can correctly answer the queries whose answers will be removed from $\tilde{h}$. The salts in $U \backslash U_{\mathsf{fresh}}$ will be recovered using P.

Our main goal is to show that we are compressing when we remove entries of $\tilde{h}$ and instead use additional lists and set. Our ways to compress will depend on the induced mouse structure in each Q_j for $j \in [F]$. To this end, we first classify the mouse structures into six broad cases. We classify the jth mouse structure for each $j \in [F]$ into the first of the following six cases it satisfies, e.g., if a mouse structure satisfies both cases 2 and 3, we categorize it into 2.

1. There is a NEW query (a, α) such that $h(a, \alpha) = a$.

17

2. There are two distinct NEW queries (a_1, α_1), (a_2, α_2) such that $h(a_1, \alpha_1) = h(a_2, \alpha_2)$.

3. There is a NEW query (a, α) such that $h(a, \alpha) = a'$ and $a' \in \mathsf{MouseSlts}_i$ for some $i < j$.

4. There is a REPEATEDNONMOUSE query (a, α) such that $h(a, \alpha) = a$.

5. There is at least one salt a such that $a \in \mathsf{MouseSlts}_i$ for some $i < j$ and there is a path of at most $B - 2$ edges in the mouse structure from a back to a' where a' is an answer to a NEW query.

6. There are no REPEATEDMOUSE queries in the mouse structure.

Note that these cases cover all mouse structures.

Claim 5.7. *Every mouse structure can be categorized into one of the cases 1 to 6.*

Proof. We will show that if a mouse structure does not satisfy case 6, it has to satisfy case 5, which suffices to prove our claim. Since the mouse structure is not in 6, it has a REPEATEDMOUSE query. Let the input salt of this query be a. Moreover, since the first query of the mouse structure has to be new, let the answer salt of this query be a'. Since the longest path in the mouse structure is of length B, it follows that there are at most $B - 2$ edges in the mouse structure between a and a'. Hence, case 5 is satisfied. ∎

Compression budget. Recall that we need to prove that the size of the output of the encoding procedure is

$$L = \log\left(\left(2^{16}e^2 \cdot \max\left\{\left(\frac{uTB^2 m_0^{2(B-2)}}{N}\right), \left(\frac{T^2}{N}\right)\right\}\right)^u \cdot \binom{N}{u} \cdot N^{MN}\right)$$

bits, where $m_0 = 3e \log u / \log \log u$. In other words, we need to show that the encoding procedure saves at least

$$u \cdot \log\left(\min\left\{\frac{N}{4T^2}, \frac{N}{4uTB^2 m_0^{2(B-2)}}\right\}\right) - 2u \log e - 14u \tag{4}$$

bits overall.

Required savings in $\tilde{h}$. As mentioned earlier, the output of the encoding algorithm will consist of $U_{\mathsf{fresh}}, \mathsf{P}, \tilde{h}$, and some additional sets and lists. The lists U_{fresh} and P will suffice to recover the set U. The list $\tilde{h}$ and the additional sets and lists are used to recover h.

Denoting $|U_{\mathsf{fresh}}| = F$ and $|U| = u$, we can describe P using $\binom{FT}{u-F}$ bits. Therefore, U, which is trivially described using $\log\binom{N}{u}$ bits, can be encoded using

$$\log\left(\binom{FT}{u-F}\binom{N}{F}\right).$$

Therefore, the saving in bits in the description of U is at least

$$\log\binom{N}{u} - \log\left(\binom{FT}{u-F}\binom{N}{F}\right) \geq \log\left(\frac{\left(\frac{N}{u}\right)^u}{\left(\frac{eFT}{u-F}\right)^{u-F}\left(\frac{eN}{F}\right)^F}\right)$$

$$= \log\left(\frac{N^{u-F}(u-F)^{u-F}F^F}{e^u(FT)^{u-F}u^u}\right) = (u-F)\log\left(\frac{N}{FT}\right) + \log\left(\frac{(u-F)^{u-F}F^F}{e^u u^u}\right)$$

$$= (u-F)\log\left(\frac{N}{FT}\right) - \log\left(e^u\left(\frac{u}{F}\right)^F\left(\frac{u}{u-F}\right)^{u-F}\right)$$

$$\geq (u-F)\log\left(\frac{N}{uT}\right) - u\log 4e . \tag{5}$$

where the first inequality uses the basic bounds for binomial coefficients $(n/r)^r \leqslant \binom{n}{r} \leqslant (en/r)^r$, and the last inequality follows since $x \leqslant 2^x$ for every $x \geqslant 0$, and $u \geqslant F$.

By subtracting (5) (how much we save in U) from (4) (how much we need to save in total), it suffices to show that we save at least

$$u \cdot \log \left(\min \left\{ \frac{N}{4T^2}, \frac{N}{4uTB^2 m_0^{2(B-2)}} \right\} \right) - 2u \log e - 14u - (u - F) \log \left(\frac{N}{uT} \right) + u \log 4e$$

bits while encoding h. Since $\log(N/uT) \geqslant \log \left(\min \left\{ N/4T^2, N/4uTB^2 m_0^{2(B-2)} \right\} \right)$, this is at most

$$F \cdot \log \left(\min \left\{ \frac{N}{4T^2}, \frac{N}{4uTB^2 m_0^{2(B-2)}} \right\} \right) - u \log e - 12u. \tag{6}$$

bits.

To show that the compression indeed achieves the savings from (6), we will show that for every salt in U_{fresh}, we can save at least the following number of bits, except for a few cases.

$$\log \left(\min \left\{ N/4T^2, N/4uTB^2 m_0^{2(B-2)} \right\} \right) .$$

In the rare cases where we cannot save as much, we will incur a small penalty. We will show that the cumulative penalty we incur is at most $7u + u \log e$ bits. Additionally, we will label each of the salts in F with a few bits that describe its "type" (according to the cases described above – case 5 will have 3 subcategories, and case 6 will have 10 subcategories), and for this 5 bits will suffice. This will cost, in total, another $5u$ bits, and therefore the total size of the encoding will indeed be bounded by the term from (6).

We now describe the details of how we handle each case. Assuming that the mouse structure for salt a'_j satisfies a particular case, we describe the encoding procedure, calculate the amount of compression we get, and then explain how decoding would work. In Section 5.2 we handle Cases 1 to 4, in Section 5.3, we handle case 5, and lastly in Section 5.4, we handle case 6. We note that here we describe the details such that they are locally verifiable, i.e., we describe the encoding, show that we save enough and describe the decoding for each case locally. We made this choice due to the complexity of the encoding and decoding procedures- we felt it was much easier to parse in the locally verifiable form. For sake of completeness, we do provide the full pseudocode of encoding and decoding in Appendix A.

5.2 Handling Cases 1 to 4

In each of the four cases below, we will be saving more bits than we need, i.e., more bits than

$$\log \left(\min \left\{ N/4T^2, N/4uTB^2 m_0^{2(B-2)} \right\} \right) .$$

Case 1. The jth mouse structure contains a NEW query (a, α) such that $h(a, \alpha) = a$, as depicted in Fig. 5a. The encoding procedure stores the index of the query (a, α) in Q_j in a list L_1 and removes the entry $h(a, \alpha)$ from $\tilde{h}$.

In decoding, if the current (jth) salt is categorized as case 1, then it removes the front index in the list L_1, and denote the index by i. It answers the ith h query, denoted (a, α), with a and sets $h(a, \alpha) = a$.

Since the index of the query (a, α) in Q_j is in $[T]$, and we remove one element of $\tilde{h}$, we save $\log(N/T)$ which is more than what we need to save.

Case 2. The jth mouse structure contains two distinct NEW queries (a_1, α_1), (a_2, α_2) such that $h(a_1, \alpha_1) = h(a_2, \alpha_2)$, and (a_1, α_1) is queried before (a_2, α_2). This is depicted in Fig. 5b. The encoding procedure stores the pair indices of the queries (a_1, α_1), (a_2, α_2) in Q_j in a list L_2 and removes the entry $h(a_2, \alpha_2)$ from $\tilde{h}$.

In decoding, if the current (jth) salt is categorized as case 2, then it removes the front element (i_1, i_2) in the list L_2. Suppose that the i_1th h query while running on salt a'_j is on (a_1, α_1). The decoding procedure

19

(a) Case 1   (b) Case 2

(c) Case 3   (d) Case 4
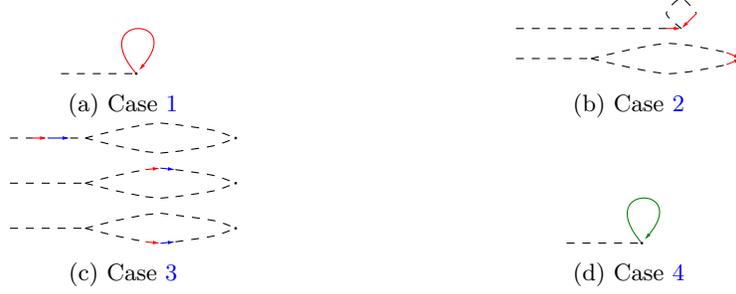
Figure 5: Cases 1 to 4. The NEW queries are drawn in red, the REPEATEDMOUSE queries are drawn in blue, and the REPEATEDNONMOUSE queries are drawn in green. The black dashed lines indicate zero or more queries of any type.
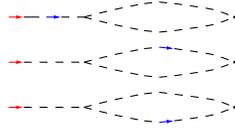
Figure 6: Mouse structure for case 5

gets the answer to this query from $\tilde{h}$. It answers the i_2th h query on (a_2, α_2) with $h(a_1, \alpha_1)$ and sets $h(a_2, \alpha_2) = h(a_1, \alpha_1)$.

Since the pair of indices of the queries $(a_1, \alpha_1), (a_2, \alpha_2)$ in Q_j are in $[T]$, and we remove one element of $\tilde{h}$, we save $\log(N/T^2)$ bits which is more than what we need to save.

Case 3. The jth mouse structure contains a NEW query (a, α) such that $h(a, \alpha) = a'$ and $a' \in \mathsf{MouseSlts}_i$ for some $i < j$. This is depicted in Fig. 5c. The encoding procedure stores the tuple consisting of i, the index of the query (a, α) in Q_j, and the lexicographical order of a' in $\mathsf{MouseSlts}_i$ in a list L_3 and removes the entry $h(a, \alpha)$ from $\tilde{h}$.

In decoding, if the current (jth) salt is categorized as case 3, it removes the front element (i_1, i_2, i_3) in the list L_3. Suppose the i_2th h query while running on salt a'_j is on (a, α). It answers the query with a' such that a' is the salt in $\mathsf{MouseSlts}_{i_1}$ whose lexicographical order is i_3. It sets $h(a, \alpha) = a'$.

Since $i \in [F]$, $F \leqslant u$, the index of (a, α) in Q_j is in $[T]$, the lexicographical index of a' in $\mathsf{MouseSlts}_i$ is in $[2B]$, and we remove one element of $\tilde{h}$, we save $\log(N/(2uTB))$ bits which is more than what we need to save.

Case 4. The jth mouse structure contains a REPEATEDNONMOUSE query (a, α) such that $h(a, \alpha) = a$. This is depicted in Fig. 5d. The encoding procedure stores the smallest index of the query (a, α) in $\mathsf{Q}_{\mathsf{fresh}}$ in a set S and removes the entry $h(a, \alpha)$ from $\tilde{h}$. Note that we never add an index corresponding to the same query multiple times to S. Indeed, if an index associated with (a, α) already appears in S, the next query on (a, α) will be a REPEATEDMOUSE query and not a REPEATEDNONMOUSE one, meaning that it cannot be added to S again.

In decoding, if the current (jth) salt is categorized as case 4, it checks for every h query on (a, α) whether S contains the index of the query in $\mathsf{Q}_{\mathsf{fresh}}$. If so it answers with a and sets $h(a, \alpha) = a$.

Since the smallest index of the query (a, α) in $\mathsf{Q}_{\mathsf{fresh}}$ is in $[FT]$, and we remove one element of $\tilde{h}$, we save at least $\log(N/(uT))$ which is more than necessary.

5.3   Handling Case 5

In this section, we describe our compression strategy in case the jth mouse structure for salt a'_j is categorized as case 5. That is, there is at least one salt d such that $d \in \mathsf{MouseSlts}_i$ for some $i < j$ and there are at most

20

$B - 2$ edges in the mouse structure between d and s where s is an answer to a NEW query. If there are several possible candidate pairs, we choose one where the number of edges is the smallest between the source and the destination salt.

Intuition. We refer to the salt s in the answer to the NEW query as the *source* salt, and the salt that appears in some earlier mouse structure as the *destination*. We are guaranteed that the path in the mouse structure from the source salt to the destination salt consists of at most $B - 2$ edges that are REPEATEDNONMOUSE queries. (There is at least one intermediate edge between the source and the destination because otherwise, the answer of a new query will be the input of a REPEATEDMOUSE query, in which case this scenario would have been classified into case 3. Additionally, all the intermediate edges must be REPEATEDNONMOUSE since we consider the shortest possible path from such source to such destination.)

Let the NEW query, whose answer is s, be (a, α). Suppose the path from the s to the d in the mouse structure consists of REPEATEDNONMOUSE queries $(a_1, \alpha_1), \ldots, (a_p, \alpha_p)$ where $a_1 = s$, $p \leqslant B - 2$. The main idea is to avoid encoding $s = h(a, \alpha)$ and recover it by encoding the lexicographical order of d in $\mathsf{MouseSlts}_i$ and encoding the path required to backtrack from d to s in the query graph at the time (a, α) is queried, i.e., encoding which of the past queries were $(a_p, \alpha_p), \ldots, (a_1, \alpha_1)$ (in this order). The problem is that, in general, the path back from d to s might be too expensive to encode. This depends on the number of "other" edges incident on the nodes on the real path back from t to s. If all nodes on the path have very few edges incident on them, say less than m_0 of them, we encode each back edge using $\log m_0$ bits per node, which requires with at most $\log(B(m_0)^{B-2})$ bits for the whole path back (the term B comes from encoding the length of the path). But, if some node has many adjacent edges, namely a *multi-collision* with more than m_0 edges, we will need to take advantage of this fact to obtain our savings (and not encode the path back from d to s).

Definition 5.8 (Large multi-collision). We say that queries $q_1, \cdots, q_m$ form an m-way multi-collision if all the q_i's are distinct and all the $h(q_i)$ are equal. We say that the multi-collision is *large* if $m \geqslant m_0$, where $m_0 := 3e \log u / \log \log u$.

If any of the queries in $(a_p, \alpha_p), \ldots, (a_1, \alpha_1)$ are involved in a large multi-collision of REPEATEDNONMOUSE queries in the query graph so far, we say we have "encountered a large multi-collision". In what follows, we explain how to obtain the required compression if a large multi-collision was *not* encountered.

Encoding when no large multi-collision. Suppose that the NEW query whose answer is the source salt s is (a, α), the destination salt is d and $d \in \mathsf{MouseSlts}_i$ for some $i < j$. The path back from d to s contains only nodes that have at most m_0 adjacent edges in the corresponding query graph since we have not encountered a large multi-collision. The encoding procedures constructs a tuple consisting of:

- the index i,

- the index of query (a, α) in Q_j,

- the lexicographical index of d in $\mathsf{MouseSlts}_i$, and

- the path back from d to s in the query graph when (a, α) is queried.

It stores the tuple in a list L_5. Finally, it removes the entry $h(a, \alpha)$ from $\tilde{h}$.

In decoding, if the current (jth) salt is categorized as case 5 *without a large multi-collision*, it detects the query (a, α) from its index in Q_j, then finds the salt d using the index i and the lexicographical order of d in $\mathsf{MouseSlts}_i$, and finally finds s using the path back from d to s. It answers the query with s.

Since $i \in [F]$, the index of (a, α) in Q_j is in $[T]$, the lexicographical index of a' in $\mathsf{MouseSlts}_i$ is in $[2B]$, the path back from d to s can be encoded in $\log(B(m_0)^{B-2})$ bits, and we remove one element of $\tilde{h}$, we save at least $\log(N/(2uTB^2(m_0)^{B-2}))$ bits which is more than necessary.

Encoding with large multi-collision. Suppose that the NEW query whose answer is the source salt s is (a, α), the destination salt is d and $d \in \mathsf{MouseSlts}_i$ for some $i < j$. Suppose further that the path back from d to s contains at least one node that has $m \geqslant m_0$ adjacent edges in the corresponding query graph (i.e., an

m-multi-collision) such that these m edges are REPEATEDNONMOUSE queries when running $\mathcal{A}_2$ on a_j. First, observe that the multi-collision does not involve a self-loop. Indeed, if any node in the mouse structure has a REPEATEDNONMOUSE query whose answer is itself, the mouse structure would be classified into case 4, and therefore we will never reach this case.

At this point, we argue that we can record the multi-collision by encoding the indices of all queries associated with the multi-collision and the center, and remove their answers from $\tilde{h}$. To this end, we store $\log N + \log \binom{FT}{m}$ bits and remove $m \log N$ bits. We have that the saving is at least

$$
\begin{aligned}
m \log N - \log N - \log \binom{FT}{m} &\geqslant \log\left(N^{m-1} \cdot \left(\frac{m}{eFT}\right)^m\right) \\
&= \log\left(\left(\frac{N}{FT}\right)^{m-2} \cdot \frac{N}{T^2} \cdot \frac{m^m}{e^m F^2}\right) \\
&\geqslant \log\left(\left(\frac{N}{uT}\right)^{m-2} \cdot \frac{N}{T^2} \cdot \frac{m^m}{e^m u^2}\right) \\
&\geqslant \log\left(\left(\frac{N}{uT}\right)^{m-2} \cdot \frac{N}{T^2}\right)
\end{aligned}
\tag{7}
$$

bits, where the first inequality follows by using the binomial inequality $\binom{FT}{m} \leqslant (eFT/m)^m$, the second inequality follows since $F \leqslant u$, and the last inequality follows since for $m \geqslant m_0 = 3e \log u / \log \log u$ it holds that $m^m \geqslant e^m u^2$. Therefore, we have

Claim 5.9. *The number of bits saved is at least*

$$
(m-1) \cdot \log\left(\min\left\{\frac{N}{T^2}, \frac{N}{uTB^2 m_0^{2(B-2)}}\right\}\right).
$$

Proof. The claim follows since the minimum between the two terms is always upper bound by $N/(uT)$ and upper bounded by N/T^2. ∎

Thus, every m-multi-collision we record allows us to save the number of bits corresponding to $m-1$ mouse structures. It is left to argue that we do not over-count, namely, that we do not count the removal of the same element from $\tilde{h}$ twice. Indeed, the same multi-collision may be encountered in several mouse structures. To this end, observe that if a salt in a mouse structure is the center of a large multi-collision which we had recorded earlier, we will be in one of the following two cases:

1. There is a query in this mouse structure whose answer is the center of the multi-collision, and this query was in an earlier mouse structure.

2. There is a query in this mouse structure whose answer is the center of the multi-collision, and this query was not in any earlier mouse structure. (Note that by the structure of collisions, i.e., a mouse structure, there could be either one such query or two.)

Case 1 need not be handled. The reason is that if the condition in it holds, then the multi-collision is, in fact, outside of the (shortest) path from the source to the destination. Therefore the scenario will either be classified as a mouse structure *without* a large multi-collision or a different large multi-collision will be encountered for this mouse structure.

In case 2, first note that when we encounter a multi-collision, we add the relevant queries to the multi-collision if they were not already recorded and only then remove the corresponding entry from $\tilde{h}$. Therefore, we never remove anything twice. The last point we need to argue is that the we get enough savings. Above we showed that we have sufficient saving for $m-1$ mouse structures. Recall that we defined that we encounter a large multi-collision if at least m_0 REPEATEDNONMOUSE queries are involved in a multi-collision on some

path we care about. It follows from this that a multi-collision of size m will be relevant in at most $m - m_0$ mouse structures- beyond that, it will no longer be a multi-collision because there will be less than m_0 REPEATEDNONMOUSE queries among the queries of the multi-collision. Since $m_0 \geqslant 1$, savings for $m - 1$ mouse structures is sufficient for us.

Overall, as claimed earlier, case 5 has 3 different further categorizations– no multi-collisions and the two cases for multi-collisions.

5.4 Handling Case 6

Handling this case, where there are no REPEATEDMOUSE edges, is somewhat more complicated compared to the previous ones as we need to further divide the compression strategy into the four following sub-cases.

(a) The body of the mouse structure has at least one NEW query, and there is a NEW query at the tip of the mouse structure

(b) The body of the mouse structure has at least one NEW query, but there are no NEW queries at the tip of the mouse structure

(c) The body of the mouse structure has no NEW queries, and the body of the mouse structure is a directed cycle

(d) The body of the mouse structure has no NEW queries, but the body of the mouse structure is not a directed cycle.

The cases (a) and (b) are easier to deal with, and the strategy we will use for them mirrors the strategy used for case 5. The cases (c) and (d) are more challenging to deal with. We present the strategies for each of them individually *assuming that we do not encounter any large multi-collisions while encoding paths*. The case for large multi-collisions is handled the same as it was handled in Section 5.3.

Handling sub-case (a). In this case, the mouse structure looks like one of the structures shown in Fig. 7a. We make the following claim.

Claim 5.10. *A mouse structure that has been categorized in case 6 (a) must contain three distinct queries q_1, q_2, q_3 such that*

- *q_1, q_2 are NEW queries.*

- *q_3 is a REPEATEDNONMOUSE query.*

- *q_2, q_3 form the tip of the mouse structure.*

- *The path back from q_3 to q_1 has only REPEATEDNONMOUSE queries.*

The proof of this claim follows readily from the definition of the case, and it is easiest to see from Fig. 7a. Nevertheless, below we provide detailed proof of this claim and remark that we shall omit the detailed proof in similar claims made below.

Proof. [of Claim 5.10] Let the NEW query at the tip of the mouse structure be q_2 and let the REPEATEDNONMOUSE query at the tip of the mouse structure be q_3. (Indeed, we are guaranteed that the other query at the tip of the mouse structure is a REPEATEDNONMOUSE query because otherwise this mouse structure would have been categorized into either case 2 or case 5). Let q_1 be the first NEW query when backtracking from q_3 along with the mouse structure — we are guaranteed such a NEW query because on backtracking, eventually, we will reach the salt a'_j and by the definition of fresh salts, any query with a'_j as the input salt is NEW. We are guaranteed that q_1 is different from q_2 because q_1 is not at the tip of the mouse structure. Given these definitions of q_1, q_2, q_3 the claim follows. ∎

Given this claim, the encoding procedure defines a tuple consisting of the following:

(a) Mouse structures with at least one NEW query in their body and a NEW query at the tip

(b) Mouse structures with at least one NEW query in their body and no NEW query at the tip

(c) Mouse structure with no NEW query in the body, but a NEW query at the tip

(d) Mouse structure with no NEW query in the body, and no NEW query at the tip
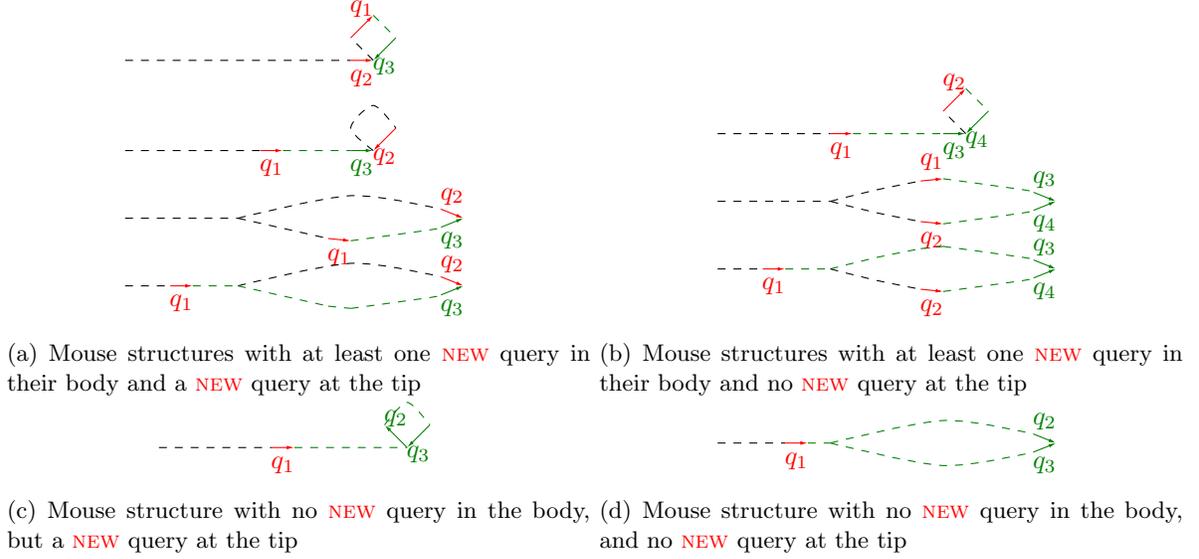
Figure 7: Mouse structures categorized in case 6. The NEW queries are drawn in red, and the REPEATED-NONMOUSE queries are drawn in green. The black dashed lines indicate zero or more NEW or REPEATED-NONMOUSE queries.

- the indices in Q_j of the q_1, q_2,

- the smallest index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$, and

- the path back from input salt of q_3 to the answer of q_1 in the query graph when q_1 was queried

This tuple will be stored in a list. The answers corresponding to queries q_1, q_2 is removed from $\tilde{h}$.

The decoding procedure figures out the answer to query q_2 from the answer of q_3, and the answer to q_1 using the path back from the input salt of q_3 to the answer of q_1 in the query graph.

The indices of q_1, q_2 in Q_j can be encoded in $2 \log T$ bits, the index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$ can be encoded in $\log(FT) \leqslant \log(uT)$ bits, and the path requires at most $\log(Bm_0^{B-2})$ bits (the calculation is the same as was done in Section 5.3 for case 5). Therefore, we save a total of $\log N^2/(uT^3 Bm_0^{B-2}) = \log N/T^2 + \log N/(uTBm_0^{B-2})$ bits which is sufficient.

Handling sub-case (b). In this case the mouse structure looks like one of the structures shown in Fig. 7b. We claim that a mouse structure that is categorized in case 6 (b) must contain distinct queries q_1, q_2, q_3, q_4 such that

- q_1, q_2 are NEW queries

- q_3, q_4 are REPEATEDNONMOUSE queries such that q_3 was made for the first time before q_4

- q_3, q_4 form the tip of the mouse structure

- The paths back from q_3 to q_1 and q_4 to q_2 have only REPEATEDNONMOUSE queries

The proof of this claim can be verified from Fig 7b. Since the details are very similar to the proof of Claim 5.10, we omit them. The encoding procedure, in this case, will define a tuple consisting of the following:

- the indices in Q_j of the q_1, q_2

- the smallest indices of q_3, q_4 in $\mathsf{Q}_{\mathsf{fresh}}$

24

- the path back from the input salt of q_3 to the answer of q_1 in the query graph when q_1 was queried.

- the path back from the input salt of q_4 to the answer of q_2 in the query graph when q_2 was queried.

This tuple is stored in a list. The answer of the queries q_1, q_2, q_4 are removed from $\tilde{h}$.

The decoding procedure figures out the answer to the query q_4 from the answer of query q_3 which came before it and the answer of queries q_2 and q_1 using the paths back from q_4 to q_2 and q_3 back to q_1 in the query graph, respectively.

Therefore, by a calculation similar to the one above, we save at least the following number of bits which is sufficient.
$$\log N^3/(u^2 T^4 B m_0^{2(B-2)}) = \log N/T^2 + 2\log N/(uTBm_0^{B-2})\,.$$

Handling sub-case (c). In this case the mouse structure looks like one of the structures shown in Fig. 7c. We claim that a mouse structure that has been categorized in case 6 (c) must have three queries q_1, q_2, q_3 in it such that

- q_1 is a NEW query

- q_2, q_3 are REPEATEDNONMOUSE queries

- q_3 is one of the queries at the tip of the mouse structure

- The answer of q_3 is the input salt of q_2

- The body of the mouse structure is a single directed cycle starting with q_2 and ending with q_3

- The queries in the mouse structure between q_1 and the answer of q_3 are REPEATEDNONMOUSE queries.

The proof of this claim can be verified from Fig 7c. Since the details are similar to the proof of Claim 5.10, we omit it.

Remark 5.11 (Naive encoding fails). *Note that a similar encoding strategy to the one we used in sub-cases (a) and (b), i.e., encoding the indices of q_2, q_3 in $\mathsf{Q}_{\text{fresh}}$, the index of q_1 in Q_j and the paths required to backtrack, will not suffice. Indeed, it will save no more than $\log N^2/(F^2 T^3) = \log N/T^2 + \log N/(F^2 T)$ bits which might not suffice since $N/(F^2 T)$ could be less than 1, in general.*

The *key idea* for compression is that we will achieve enough cumulative compression over all the mouse structures that satisfy this sub-case by storing some of the components of the encoding in an unordered set. We provide the details below.

Scenario 1. First, let us additionally assume that the body of the mouse structure is formed in sequence, i.e., the query q_3 was made for the first time after all the other queries in the body of the mouse had been made for the first time. In this case, we encode the index of queries q_2 and q_1 in a list, but the index of the query q_3 and the path back from q_3 to q_2 is encoded in an unordered set. In more detail, the encoding procedure stores the following:

- A tuple consisting of the minimum index of the query q_2 in $\mathsf{Q}_{\text{fresh}}$, index of query q_1 in Q_j, and the path back to the answer salt of q_1 from answer salt of q_3 in the query graph when q_1 was queried. This tuple is stored in a list ordered by j.

- A tuple consisting of the minimum index of the query q_3 in $\mathsf{Q}_{\text{fresh}}$ along with the path back from the input salt of q_3 to the input salt of q_2 in the query graph when q_3 was queried. This is stored in an unordered set.

The answers of queries q_1 and q_3 are removed from $\tilde{h}$.

The decoding procedure answers q_3 as follows. It detects which query is q_3 and the path back from the input salt of q_3 to the input salt of q_2 from the unordered set, locates the right q_2, and answers with the input salt of q_2. It answers q_1 by using the path from the answer salt of q_3 back to the answer of q_1.

If there are y mouse structures that are captured by this sub-case, then the total number of bits saved over the y mouse structures is at least

$$y \log N^2 - y \log FT - y \log T - \log \binom{FT}{y} - 2y \log(Bm_0^{B-2})$$

$$\geqslant y \left(\log \frac{N}{T^2} + \log \frac{N}{uTB^2 m_0^{2(B-2)}} \right) - u - y \log e,$$

where the last inequality follows since $F \leqslant u$ and

$$\binom{FT}{y} \leqslant \left(\frac{eFT}{y} \right)^y \leqslant e^y T^y \left(\frac{F}{y} \right)^y \leqslant e^y T^y (2^{F/y})^y \leqslant e^y T^y 2^F \leqslant e^y T^y 2^u.$$

Ignoring the missing $u + y \log e$ bits in the encoding size, we indeed save as much as we need to save. The missing $u + y \log e$ bits are consumed by the additional terms we introduced in the compression budget (see (6)): u is consumed by the penalty budget of $14u$ and $y \log e$ by the penalty budget of $u \log e$.

Scenario 2. Here, the mouse structure is not formed in sequence. Namely, there is a query q_4 in the body such that it was the last query to be made in the body. The encoding procedure stores the following:

- A tuple consisting of the minimum index of the query q_3 in $\mathsf{Q}_{\mathsf{fresh}}$, index of query q_1 in Q_j, and the path back to the answer salt of q_1 from the answer salt of q_3 in the query graph when q_1 was queried. This tuple is stored in a list ordered by j.

- A tuple consisting of the minimum index of the query q_4 in $\mathsf{Q}_{\mathsf{fresh}}$ along with the path back to the answer salt of q_4 from the input salt of q_4 in the query graph when q_4 was queried. This tuple is stored in an unordered set.

The answers corresponding of queries q_1, q_4 are removed from $\tilde{h}$.

The decoding procedure answers q_4 by using the path corresponding to q_4 and finding the answer salt of q_4 in the query graph. It answers q_1 using the path back from answer of q_3.

The compression analysis is the same as in scenario 1, so we avoid repetition. We remark that here we incur a similar penalty of $y \log e + u$ bits, as well.

Handling sub-case (d). In this case the mouse structure looks the structure shown in Fig. 7d. We claim that a mouse structure that has been categorized in case 6 (d) must have three queries q_1, q_2, q_3 in it such that

- q_1 is a NEW query

- query q_2 was made for the first time before the query q_3 was made for the first time

- q_2, q_3 are REPEATEDNONMOUSE queries that form the tip of the mouse structure

- The paths from q_1 to q_2 and q_1 to q_3 in the mouse structure have only REPEATEDNONMOUSE queries.

The proof of this claim can be verified from Fig 7d. The details are similar to the proof of Claim 5.10 and are therefore omitted.

The main idea for the compression is similar to that in sub-case (c) i.e., we achieve enough cumulative compression over all the mouse structures that satisfy this sub-case by storing some of the components of the encoding in an unordered set.

Definition 5.12 (Frontal tail queries). Queries in the mouse structure that occur after the query q_1 but before the body are the *frontal tail queries*.

Scenario 1. First, let us additionally assume that the last query to be made for the first time among the queries in the body of the mouse structure and the frontal tail queries is q_3. In this scenario, the encoding procedure stores the following.

- A tuple consisting of the minimum index of the query q_2 in $\mathsf{Q}_{\mathsf{fresh}}$, index of query q_1 in Q_j, the path back from the input salt of q_2 to the salt that is the answer of q_1 when in the query graph when q_3 is queried. This tuple is stored in a list ordered by j.

- A tuple consisting of the minimum index of the query q_3 in $\mathsf{Q}_{\mathsf{fresh}}$ and the path back from the input salt of q_3 to the salt that is the answer of q_1 when in the query graph when q_3 is queried. This tuple is stored in an unordered set.

The answers of queries q_1, q_3 are removed from $\tilde{h}$.

The decoding procedure answers q_3 by using the path corresponding to q_3 to locate the answer of q_1, and then for every candidate q_2 it checks whether the corresponding path back leads to the answer of q_1. Note that, only for the correct q_2 this would be true. Otherwise, if it were true for some other q_2 as well, then the mouse structure corresponding to the other q_2 would have been categorized in case 5. Once the correct q_2 has been found, it answers with the answer salt of q_2. It has already located the answer of q_1, in this case.

If there are y mouse structures that are captured by this sub-case, then the total number of bits saved over the y mouse structures is at least

$$y \log N^2 - y \log FT - y \log T - \log \binom{FT}{y} - y \log(Bm_0^{B-2})$$

$$\geqslant y \left(\log \frac{N}{T^2} + \log \frac{N}{uTBm_0^{B-2}} \right) - u - y \log e,$$

where the last inequality follows since $F \leqslant u$ and $\binom{FT}{y} \leqslant e^y T^y 2^u$, as in sub-case (c). Note that in this case, as in sub-case (c), we incur another (global) $u + y \log e$ penalty which we can tolerate by the total compression budget; ignoring this penalty, we compress sufficiently.

Scenario 2. Suppose there is a query q_4 in the body that is not at the tip of the mouse structure, and it was the last query to be made for the first time among the body queries and the frontal tail queries. Assume without loss of generality that q_2 is the query in the same body curve as q_4. The encoding procedure stores the following

- A tuple consisting of the minimum index of the queries q_2, q_3 in $\mathsf{Q}_{\mathsf{fresh}}$, index of query q_1 in Q_j, the path back to the answer salt of q_4 from the input salt of query q_2 in the query graph when q_4 is queried, and the path back from the input salt of q_3 to the answer salt of query q_1 in the query graph when q_4 is queried. This tuple is stored in a list ordered by j.

- A tuple consisting of the minimum index of the query q_4 in $\mathsf{Q}_{\mathsf{fresh}}$, and the path back from the input salt of q_4 to the answer salt of query q_1 in the query graph when q_4 is queried. This tuple is stored in an unordered set.

The answers of queries q_1, q_3, q_4 are removed from $\tilde{h}$.

The decoding procedure answers q_3 using the answer of the corresponding q_2. It answers q_4 by first locating the right q_3 by using the paths back from input salt of q_4 to the answer salt of q_1 and from input salt q_3 to the answer salt of q_1. This automatically locates the right q_2– it locates the answer of q_4 using the path back from the input salt of q_2 to the answer salt of q_4. It has already found the answer of q_1 in this case.

A similar calculation as in the previous scenario shows that this encoding procedure saves at least this many bits:

$$y \left(\log \frac{N}{T^2} + 2 \log \frac{N}{uTB^2 m_0^{2(B-2)}} \right) - y \log e - u$$

which is again within our budget (incurring with a global $u + y \log e$ penalty).

In this remaining scenario, there is a query q_4 in the frontal tail that is made for the first time after all the body queries and the other frontal tail queries. We split it again into two cases, depending on the order in which the queries were made.

Scenario 3. We will first consider the case when one of the queries q_2, q_3 at the tip of the mouse structure is the last to be made for the first time among the body queries. Recall the first body node of the mouse structure was the node from which the two body curves of the mouse separated (see Fig. 4). The encoding procedure here will store the following.

- A tuple consisting of the minimum index of the queries q_2 in $\mathsf{Q}_{\mathsf{fresh}}$, index of query q_1 in Q_j, the minimum index of the query q_4 in $\mathsf{Q}_{\mathsf{fresh}}$, the path back to the answer salt of q_4 from the first body node of the mouse structure in the query graph when q_4 is queried, the path back from the input salt of q_2 to the first body node in the query graph when q_3 was queried for the first time, and the path back from the input salt of q_4 to the answer salt of query q_1 in the query graph when q_1 is queried. This tuple is stored in a list ordered by j.

- A tuple consisting of the minimum index of the query q_3 in $\mathsf{Q}_{\mathsf{fresh}}$, the path back from the input salt of q_3 to the first body node of the mouse structure in the query graph when q_3 is queried for the first time. This tuple is stored in an unordered set.

It removes the answers of queries q_3, q_4, q_1 from $\tilde{h}$. We can do similar calculations like the previous scenario and show here that we save at least $y(\log N/T^2 + 2 \log N/uTB^2 m_0^{2(B-2)}) - y \log e - u$ bits.

The argument that decoding will succeed is similar to the earlier scenarios, and we omit it here.

Scenario 4. And finally, for the other scenario when a different query q_5 (without loss of generality in the same body curve as q_2) which is not at the tip of the mouse structure is the last query to be made for the first time among the body queries, the encoding procedure will store the following.

- A tuple consisting of the minimum index of the queries q_2, q_3 in $\mathsf{Q}_{\mathsf{fresh}}$, index of query q_1 in Q_j, the minimum index of the query q_4 in $\mathsf{Q}_{\mathsf{fresh}}$, the path back to the answer salt of q_5 from the input salt of q_2 when q_5 was queried for the first time, the path back from the input salt of q_3 to the first body node when q_5 was queried for the first time, the path back to the answer salt of q_4 from the first body node of the mouse structure when q_5 was queried for the first time, and the path back from the input salt of q_4 to the answer salt of query q_1 in the query graph when q_4 is queried for the first time. This tuple is stored in a list ordered by j.

- A tuple consisting of the minimum index of the query q_5 in $\mathsf{Q}_{\mathsf{fresh}}$, the path back from the input salt of q_5 to the first body node when q_5 was queried for the first time. This is stored in an unordered set.

It removes the answers of queries q_3, q_4, q_5, q_1 from $\tilde{h}$. We can do similar calculations like before and show here that we save at least the following number of bits:

$$y(\log N/T^2 + 3 \log N/uTB^2 m_0^{2(B-2)}) - y \log e - u \ .$$

.

The argument that decoding will succeed is similar to the earlier scenarios, and we omit it here.

Remark 5.13 (Total penalties incurred throughout sub-cases (c) and (d).)**.** *There are two scenarios we considered in sub-case (c) and in each case, the penalty incurred was $u + y \log e$ bits. For sub-case (d), we considered four scenarios, and in each scenario the penalty incurred was $u + y \log e$ bits. So the total penalty we incur is at most $6u + u \log e$ bits. Overall, case 6 has $8 \ (= 1 + 1 + 2 + 4)$ different further categorizations if there are no multi-collisions and 2 categorizations concerning multi-collisions, leading to 10 sub-categories as claimed earlier.*

6 Proving the STB Conjecture for $SB \ll T$

This section proves another upper bound on the advantage of any auxiliary-input adversary in the bounded-length collision resistance game of salted MD hash in the AI-ROM. The main theorem is stated next.

Theorem 6.1. *Let $C = 2^9 \cdot 6 \cdot e^4$. For any $N, M, B, S, T \in \mathbb{N}_{>0}$ and fixing $\hat{S} := S + \log N$, it holds that*

$$\mathsf{Adv}^{\mathsf{ai\text{-}cr}}_{\mathsf{MD},N,M,B}(S,T) \leqslant C \cdot \max\left\{ \left(\frac{T^2}{N}\right), \left(\frac{\hat{S}^4 T B^2}{N}\right) \right\} + \frac{1}{N} \ .$$

The proof of this theorem mostly mirrors that of Theorem 5.1, except in the way that few of the cases are handled in the compression argument. Technically, we derive the following Lemma 6.2 (an analogue of Lemma 5.2), and combine it with Lemma 4.1 to get the claimed bound in Theorem 6.1.

Lemma 6.2. *Fix $N, M, B, S, T, u \in \mathbb{N}_{>0}$, $\sigma \in \{0,1\}^S$. Then, for any $\mathcal{A}_2$ that makes at most T queries to its oracle, it holds that*

$$\mathsf{Adv}^{\mathsf{mi\text{-}cr}}_{\mathsf{MD},N,M,B,u}(\sigma, \mathcal{A}_2) \leqslant \left(2^9 e^4 \max\left\{ \left(\frac{u^4 T B^2}{N}\right), \left(\frac{T^2}{N}\right) \right\} \right)^u .$$

The proof of this lemma is in the same spirit as the proof of Lemma 5.2 with several key differences. The main difference is how we encode the path back from a given destination node to the associated source node. In Section 5, we do this in a somewhat straightforward manner by encoding the length of the path and then the index of every edge to take, where the index might be large if there is a large multi-collision associated with that node. Large-enough multi-collisions were handled separately, so we had a bound ($m_0 \approx \log u / \log \log u$) for the range of the index of each back-edge. In this section, we encode the source node by just writing its lexicographic index among all possible sources within a given distance in the query graph. Of course, there might be too many possible sources at a given distance, making it too expensive to encode. But, and this is our main technical observation in this section, if there are too many possible sources, then there must be many large multi-collision, and therefore we can save enough bits by taking advantage of it.

We proceed by setting up the graph-theoretic definitions and lemmas. In the lemma below, we show that if in the query graph there is a node v such that there are at least p nodes which have a shortest path of length d to the node v, then there must be at least $p - 1$ edges involved in a multi-collision in the induced sub-graph.

Definition 6.3 (*d*-neighborhood of a vertex)**.** *Let $G = (V, E)$ be a directed graph. We say that a node $v_1 \in V$ is in the d-neighborhood of $v_2 \in V$ if the shortest directed path from v_1 to v_2 in G consists of (exactly) d edges.*

Definition 6.4 (multi-collision of edges in a graph)**.** *For an edge $e = (a, b)$, we refer to b as the target of the edge. We say that edges $e_1, \cdots, e_m$ form a m-multi-collision if all of them share a common target. We refer to the common target as the center of the multi-collision. We refer to m as the size of the multi-collision.*

We prove the following lemma.

Lemma 6.5. *Let $G = (V, E)$ be a directed graph. Let $d \in \mathbb{N}_{>0}$ and $p \in \mathbb{N}_{>0}$ such that $p \geqslant 2$. Suppose that there is a node $v \in V$ such that there are p distinct nodes in its d-neighborhood. Then, there are $t \geqslant 1$ distinct multi-collisions in G, each of size β_i for $i = 1, \ldots, t$, such that each $\beta_i \geqslant 2$ and $\sum_{i=1}^{t} (\beta_i - 1) \geqslant p - 1$.*

Proof. The proof is via induction on p. The base case is for $p = 2$. If there are at least 2 nodes in the d-neighborhood of v, it is easy to see that the graph has to have a 2-multicollision. Here $t = 1$ and $\beta_1 = 2$ and therefore, the statement holds.

For the induction hypothesis, assume that the statement holds for $p = k$, and we will show that it also holds for $p = k + 1$. Since $p = k + 1$, the graph G has a node v with $k + 1$ nodes in its d-neighborhood. Let this set of $k + 1$ nodes be S. Consider the subgraph G' of G with only the edges on the shortest directed path from the nodes in S to the node v (choose one shortest path arbitrarily if there are multiple paths of the same length). By definition, in G', the node v still has $k + 1$ nodes in its d-neighborhood. Consider a node $v' \in S$ such that it has in-degree 0 in G' (note that such a v' must exist as there are no cycles in G'). Further, consider the set of edges P' that are on the path from v' to v but not on the path from the other edges in S to v. In other words, denoting $E_{v* \to v}$ the set of edges on the path from v^* to v, we define

$$P' = E_{v' \to v} \setminus \bigcup_{v'' \in S \setminus \{v'\}} E_{v'' \to v}.$$

Definition 6.6 (Frontal edge of path)**.** For a directed path in a graph, we define the edge whose target node has out-degree zero to be the frontal edge of the path.

Define a subgraph H' of G' by removing the edges in P' from G'. Observe that P' will consist of a path and the frontal edge of the path will necessarily be incident on the target of an edge in H' – this is true because the first possibility is that this path is incident directly on v. In this case, there has to be at least one other path incident of v since $p > 2$ and we did not remove edges on the shortest path from other nodes in S to v. Hence, it is incident on the target of some edge in H'. The other possibility is that the frontal edge of this path incident on the end-point of some other edge in H' because we know this is part of a path from v' to v. The end-point has to be the target of some edge because otherwise, the length of the path from v' to v would be more than d.

It follows from the definition of P' that v has p distinct nodes in its d-neighborhood in H. Using the induction hypothesis H has t multi-collisions each of size β_i such that $\beta_i \geqslant 2$ and $\sum_{i=1}^{t} (\beta_i - 1) \geqslant k - 1$. When we add back the edges in G', there are two possible cases. In the first case, one of the edges we add back is incident to one of the m multi-collisions, in which case G' has t multi-collisions each of size β_i' such that $\beta_i' \geqslant 2$ and $\sum_{i=1}^{t} (\beta_i' - 1) \geqslant k - 1 + 1 = k$, because one of the β_i' will be larger than β_i by 1. In the other case, one of the edges we add back is incident on the target of an edge in G which was not involved in a multi-collision. But this in now a 2-multi-collision in G'. So G' has $t + 1$ multi-collisions each of size β_i' such that $\beta_i' \geqslant 2$ and $\sum_{i=1}^{t} (\beta_i' - 1) \geqslant k - 1 + 1 = k$. Since G' is a subgraph of G, it follows that for some t, G has t multi-collisions each of size β_i' such that $\beta_i' \geqslant 2$ and $\sum_{i=1}^{t} \beta_i' \geqslant k$. ∎

A direct corollary is that in the query graph, if for any salt s there are at least $p + 1$ salts in its d-neighborhood, then there must be p queries involved in multi-collisions on the paths from the nodes in its d-neighborhood to s. We obtain non-trivial compression for large enough p by encoding those multi-collisions.

Non-trivial encoding for multiple multi-collisions. We consider t multi-collisions of sizes $\beta_1, \ldots, \beta_t$. We encode these t multi-collisions by encoding the t centers of the multi-collision, and for each center, we encode the index of the queries in $\mathsf{Q_{fresh}}$ that form the multi-collision in a set. Recall that the indices of the queries in $\mathsf{Q_{fresh}}$ are in $[FT]$. The total number bits we need to encode is

$$\log \left(\binom{N}{t} \binom{FT}{\beta_1} \binom{FT - \beta_1}{\beta_2} \cdots \binom{FT - \sum_{i=1}^{t-1} \beta_i}{\beta_t} \right) \leqslant \log \left(\frac{N^t (uT)^{\sum_{i=1}^{t} \beta_i}}{t! \beta_1! \beta_2! \cdots \beta_t!} \right)$$

$$= \log \left(\frac{N^t (u^4 T)^{\beta - 2t} T^{2t} u^{8t - 3\beta}}{t! \beta_1! \beta_2! \cdots \beta_t!} \right), \qquad (8)$$

where the inequality follows by using $\binom{n}{r} \leqslant \frac{n^r}{r!}$ and $F \leqslant u$, and the equality follows by letting $\beta = \sum_{i=1}^{t} \beta_i$ and rearranging.

30

Claim 6.7. *If $\beta \geqslant e^3 u^2/2$, then Eq. (8) $\leqslant \log\left(N^t(u^4 T)^{\beta-2t}T^{2t}\right)$.*

Proof. If $8t < 3\beta$, then the claim is immediate. Otherwise, we use the fact that $\beta_1!\beta_2!\cdots\beta_t! \geqslant (\beta/(te))^{\beta}$ (see Claim 6.8 for a proof) and the fact that $t! \geqslant (t/e)^t$ in order to get

$$(8) \leqslant \log\left(\frac{N^t(u^4 T)^{\beta-2t}T^{2t}u^{8t-3\beta}t^{\beta-t}e^{\beta+t}}{\beta^{\beta}}\right).$$

Since $3\beta/8 \leqslant t \leqslant \beta/2$ (with the upper bound since each $\beta_i \geqslant 2$ and the lower bound by our assumption), we can show that $\beta^{\beta} \geqslant u^{8t-3\beta}t^{\beta-t}e^{\beta+t}$; See Claim 6.9 for a proof. Then, by plugging this bound into the above inequality we get the claim. ∎

From this claim it follows that when $\beta \geqslant e^3 u^2/2$, by storing the multi-collisions as above, the amount of bits saved is at least

$$\log N^{\beta} - \log\left(N^t(u^4 T)^{\beta-2t}T^{2t}\right) = t\log\left(\frac{N}{T^2}\right) + (\beta-2t)\log\left(\frac{N}{u^4 T}\right)$$
$$\geqslant (\beta-t)\log\left(\min\left\{\frac{N}{T^2}, \frac{N}{u^4 T}\right\}\right) \geqslant u\log\left(\min\left\{\frac{N}{T^2}, \frac{N}{u^4 T}\right\}\right),$$

where the second inequality follows since $t \leqslant \beta/2$, and $\beta/2 \geqslant e^3 u^2/4 \geqslant u^2 \geqslant u$.

For the bound that we want to get in this section, the amount of bits we need to save *per mouse structure* is $\log\left(\min\left\{\frac{N}{T^2}, \frac{N}{u^4 T B^2}\right\}\right)$ (see explanation below), and so *we indeed save enough from encoding one such set of multi-collisions.*

The compression argument. We encode a source node from a destination node by encoding the distance and which of the nodes at the given distance from the destination node is the source node. If the number of candidate nodes at the specified distance is larger than $e^3 u^2/2$, by Lemma 6.5, we are guaranteed that there exist t multi-collisions of size $\beta_1, \beta_2, \ldots, \beta_t$ such that $\beta \geqslant \sum_{i=1}^{t}(\beta_i-1) \geqslant e^3 u^2/2$. This implies, by Claim 6.7, that we can already save enough by only encoding this set of multi-collisions. Using this argument we prove Lemma 6.2 which in turn implies Theorem 6.1, as explained above. We sketch the proof of Lemma 6.2 below. (We note that in Appendix B we state the changes in the pseudocode of encoding and decoding algorithms, compared to the algorithms used in the proof of Theorem 5.1.)

Proof. [Sketch of Lemma 6.2] We continue using the setup and notation we defined in the proof of Lemma 5.2. After encoding U into $U_{\mathsf{fresh}}, \mathsf{P}$ and accounting for the list that classifies the mouse structure into the various cases, here the budget of bits per mouse structure will be

$$\log\left(\min\left\{\frac{N}{e^3 4 T^2}, \frac{N}{4e^3 u^4 T B^2}\right\}\right).$$

The first difference here is how we encode source salts from destination salts, an action that we need to do to handle Cases 5 and 6. Except this, there are no differences in the argument, namely, handling Cases 1 to 4 remains the same.

Encoding source salts. We let d be the minimum number of edges between the source and destination nodes. If there are at most $e^3 u^2/2$ possible salts such that their shortest path to the source node in the query graph so far has d edges, we encode the distance and which of the nodes is the actual destination node. This would take $\log(e^3 u^2 B/2)$ bits. By repeating the same proof with this term, we obtain that the term $B(m_0)^{B-2}$ is replaced by $e^3 u^2 B/2$. In Case 5 and sub-cases (a) and (b) of Case 6, this strategy will save enough directly.

If there are more than $e^3 u^2/2$ possible salts such that their shortest path to the destination node in the query graph so far has d edges, we encode the set of multi-collisions that we are guaranteed by Lemma 6.5 and save at least $u\log\left(\min\left\{\frac{N}{e^3 T^2}, \frac{N}{e^3 u^4 T}\right\}\right)$ bits, which is more than what we need across all mouse structures.

In this case, the entire encoding just removes the answers of queries from $\tilde{h}$ involved in these multi-collisions and just outputs the rest of $\tilde{h}$ and these multi-collisions.

Handling sub-cases (c) and (d) of Case 6. For sub-cases (c) and (d) of Case 6, if we use the new strategy to encode the destination nodes, we might not save enough. We explain why this is the case. Recall that in sub-case (c) of Case 6, per mouse structure, our savings was

$$\log \frac{N}{T^2} + \log \frac{N}{uTB^2 m_0^{2(B-2)}}$$

bits. Now, when we replace Bm_0^{B-2} with $e^3 u^2 B/2$, the per mouse structure savings is

$$\log \frac{N}{T^2} + \log \frac{N}{e^3 u^5 TB^2}$$

bits. For our current budget, this is not enough. To accommodate this, we could increase our budget per mouse structure by replacing the u^4 term with a u^5 term.[12] But, by slightly modifying the encoding/decoding procedures, we can use the allocated budget.

The idea here is that because we have a larger budget in terms of u, we no longer need to use the unordered set trick. In more detail, we were encoding indices of some queries in unordered sets and paths related to that query to figure out the order of these queries. Here we can encode all queries in their right order and no longer encode the paths required to recover the order of these queries. We explain this through the example of how we handle Scenario 1 of sub-case (c).

The encoding procedure now stores the following:

- Minimum index of the query q_2, q_3 in $\mathsf{Q}_{\mathsf{fresh}}$, index of query q_1 in Q_j. These are stored in lists.

- The number of edges d on the shortest path from the answer salt of query q_3 in the query graph when q_1 was queried and the lexicographic index i of the answer salt of q_1 among the salts whose shortest path to the answer salt of query q_3 is d. This is stored in a list ordered by j.

The answers of queries q_1 and q_3 are removed from $\tilde{h}$.

The decoding procedure answers q_3 with the input salt of q_2. On query q_1, it locates the salts in the query graph such that they have d edges on the shortest path answers with the salt with lexicographic index i among these salts.

Here the total number of bits saved per mouse structure is at least

$$\log N^2 - 2\log FT - \log T - \log B - \log(e^3 u^2/2)$$
$$\geq \log \frac{N}{T^2} + \log \frac{N}{e^3 F^2 u^2 TB}$$
$$\geq \log \frac{N}{T^2} + \log \frac{N}{e^3 u^4 TB},$$

where the last inequality follows since $F \leq u$. Note that this is within our budget. We can similarly show that we can save enough bits (as required by our budget) for other scenarios of sub-cases (c) and (d). We avoid repetition and omit the full details. Otherwise, there are no other differences in this proof compared to Lemma 5.2. ∎

Finally, we prove the two claims we used in calculations earlier.

Claim 6.8. *Let $t \in \mathbb{N}_{>0}$ and $a_1, a_2, \ldots, a_t \in \mathbb{N}_{>0}$. Let $a = a_1 + a_2 + \ldots + a_t$. Then*

$$a_1! a_2! \cdots a_t! \geq \left(\frac{a}{te}\right)^a$$

---

[12] This would result in an S^5 term in the total upper bound on the advantage of the attacker. Our improved compression strategy gives an S^4 term.

Proof. From Sterling's approximation, we have that $a_i! \geqslant \left(\frac{a_i}{e}\right)^{a_i}$. So it suffices to prove that

$$\prod_{i=1}^{t} a_i^{a_i} \geqslant \left(\frac{a}{t}\right)^a .$$

Since $a = a_1 + a_2 + \ldots + a_t$, this is equivalent to proving

$$\prod_{i=1}^{t} \left(\frac{a}{t a_i}\right)^{a_i} \leqslant 1 .$$

We have that

$$\frac{a}{t a_i} = \left(\frac{1}{t} + \sum_{j \in [t] \setminus \{i\}} \frac{a_j}{t a_i}\right) = 1 + \left(\sum_{j \in [t] \setminus \{i\}} \frac{a_j}{t a_i} - \frac{t-1}{t}\right) .$$

Using $1 + x \leqslant e^x$ we get that

$$\prod_{i=1}^{t} \left(\frac{a}{t a_i}\right)^{a_i} \leqslant \prod_{i=1}^{t} e^{a_i \sum_{j \in [t] \setminus \{i\}} \frac{a_j}{t a_i} - \frac{t-1}{t}} .$$

Simplifying we have

$$\prod_{i=1}^{t} \left(\frac{a}{t a_i}\right)^{a_i} \leqslant e^{\sum_{j=1}^{t} a_j \left(\frac{t-1}{t} - \frac{t-1}{t}\right)} = e^0 = 1 .$$

Hence, $\prod_{i=1}^{t} \left(\frac{a}{t a_i}\right)^{a_i} \leqslant 1$, as needed. ∎

Claim 6.9. *Let $\beta, u, t \in \mathbb{N}_{>0}$ such that $3\beta/8 \leqslant t \leqslant \beta/2$ and $\beta \geqslant e^3 u^2/2$. Then*

$$\beta^\beta \geqslant u^{8t-3\beta} t^{\beta-t} e^{\beta+t} .$$

Proof. Let $g(t) = u^{8t-3\beta} t^{\beta-t} e^{\beta+t} / \beta^\beta$. Let $f(t) = \ln g(t)$ where $\ln$ is the natural logarithm. Since $\ln$ is an increasing function, $g(t)$ is maximized where $f(t)$ is maximized and vice versa. We show that $g(t)$ is maximized at $t = \beta/2$ (see why below). For this, it suffices to show that $f(t)$ is maximized at $t = \beta/2$ for $3\beta/8 \leqslant t \leqslant \beta/2$ and $\beta \geqslant e^3 u^2/2$.
We have

$$f(t) = (8t - 3\beta) \ln u + (\beta - t) \ln t + (\beta + t) - \beta \ln \beta .$$

Differentiating with respect to t, we have

$$f'(t) = 8 \ln u + \frac{\beta - t}{t} - \ln t + 1 = \frac{\beta}{t} + \ln \frac{u^8}{t} .$$

Since $\beta \geqslant e^3 u^2/2$, and $3\beta/8 \leqslant t \leqslant \beta/2$, $f'(t)$ is positive throughout in the domain we are interested in. So, the maximum occurs for $t = \beta/2$. At $t = \beta/2$,

$$g(t) = \frac{u^\beta \beta^{\frac{\beta}{2}} e^{\frac{3\beta}{2}}}{2^{\frac{\beta}{2}} \beta^\beta} = \left(\frac{e^3 u^2}{2\beta}\right)^{\frac{\beta}{2}} .$$

Since $\beta \geqslant e^3 u^2/2$, this is at most 1. Since at $t = \beta/2$ the value of $g(t)$ is maximized for the range we are interested in, the claim follows. ∎

Acknowledgements

References

[AAC+17]   Hamza Abusalah, Joël Alwen, Bram Cohen, Danylo Khilko, Krzysztof Pietrzak, and Leonid Reyzin. Beyond hellman's time-memory trade-offs with applications to proofs of space. In *Advances in Cryptology - ASIACRYPT*, pages 357–379, 2017.

[ACDW20]   Akshima, David Cash, Andrew Drucker, and Hoeteck Wee. Time-space tradeoffs and short collisions in merkle-damgård hash functions. In *Advances in Cryptology - CRYPTO*, pages 157–186, 2020.

[Adl78]   Leonard Adleman. Two theorems on random polynomial time. In *Symposium on Foundations of Computer Science, SFCS*, pages 75–83, 1978.

[BBS06]   Elad Barkan, Eli Biham, and Adi Shamir. Rigorous bounds on cryptanalytic time/memory tradeoffs. In *Advances in Cryptology - CRYPTO*, pages 1–21, 2006.

[BDRV18]   Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In *Advances in Cryptology - EURO-CRYPT*, pages 133–161, 2018.

[BKP18]   Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In *STOC*, pages 671–684, 2018.

[CDG18]   Sandro Coretti, Yevgeniy Dodis, and Siyao Guo. Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In *Advances in Cryptology - CRYPTO*, pages 693–721, 2018.

[CDGS18]   Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In *Advances in Cryptology - EUROCRYPT*, pages 227–258, 2018.

[CGLQ20]   Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *FOCS*, pages 673–684, 2020.

[CHM20]   Dror Chawin, Iftach Haitner, and Noam Mazor. Lower bounds on the time/memory tradeoff of function inversion. In *Theory of Cryptography - TCC*, pages 305–334, 2020.

[CK18]   Henry Corrigan-Gibbs and Dmitry Kogan. The discrete-logarithm problem with preprocessing. In *Advances in Cryptology - EUROCRYPT*, pages 415–447, 2018.

[CK19]   Henry Corrigan-Gibbs and Dmitry Kogan. The function-inversion problem: Barriers and opportunities. In *Theory of Cryptography - TCC*, pages 393–421, 2019.

[Dam87]   Ivan Damgård. Collision free hash functions and public key signature schemes. In *Advances in Cryptology - EUROCRYPT*, pages 203–216, 1987.

[DGK17]   Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. Fixing cracks in the concrete: Random oracles with auxiliary input, revisited. In *Advances in Cryptology - EUROCRYPT*, pages 473–495, 2017.

[Din20]   Itai Dinur. Tight time-space lower bounds for finding multiple collision pairs and their applications. In *Advances in Cryptology - EUROCRYPT*, pages 405–434, 2020.

[DTT10]   Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and prgs. In *Advances in Cryptology - CRYPTO*, pages 649–665, 2010.

[FN99]   Amos Fiat and Moni Naor. Rigorous time/space trade-offs for inverting functions. *SIAM J. Comput.*, 29(3):790–803, 1999.

[GT00]   Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *FOCS*, pages 305–313, 2000.

[GT20]   Ashrujit Ghoshal and Stefano Tessaro. On the memory-tightness of hashed elgamal. In *Advances in Cryptology - EUROCRYPT*, pages 33–62, 2020.

[Hel80]   Martin E. Hellman. A cryptanalytic time-memory trade-off. *IEEE Trans. Inf. Theory*, 26(4):401–406, 1980.

[IK10]   Russell Impagliazzo and Valentine Kabanets. Constructive proofs of concentration bounds. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 617–631. Springer, 2010.

[Jou04]   Antoine Joux. Multicollisions in iterated hash functions. application to cascaded constructions. In *Advances in Cryptology - CRYPTO*, pages 306–316, 2004.

[KNY18]   Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In *Advances in Cryptology - EUROCRYPT*, pages 162–194, 2018.

[Mer82]   Ralph C. Merkle. *Secrecy, Authentication and Public Key Systems*. PhD thesis, UMI Research Press, Ann Arbor, Michigan, 1982.

[Mer87]   Ralph C. Merkle. A digital signature based on a conventional encryption function. In *Advances in Cryptology - CRYPTO*, pages 369–378, 1987.

[Mer89]   Ralph C. Merkle. A certified digital signature. In *Advances in Cryptology - CRYPTO*, pages 218–238, 1989.

[MT10]   Robin A. Moser and Gábor Tardos. A constructive proof of the general lovász local lemma. *J. ACM*, 57(2):11:1–11:15, 2010.

[Oec03]   Philippe Oechslin. Making a faster cryptanalytic time-memory trade-off. In *Advances in Cryptology - CRYPTO*, pages 617–630, 2003.

[ST79]   Robert H. Morris Sr. and Ken Thompson. Password security - A case history. *Commun. ACM*, 22(11):594–597, 1979.

[Unr07]   Dominique Unruh. Random oracles and auxiliary input. In *Advances in Cryptology - CRYPTO*, pages 205–223, 2007.

[Yao90]   Andrew Chi-Chih Yao. Coherent functions and program checkers (extended abstract). In *STOC*, pages 84–94, 1990.

A   Description of our Encoding and Decoding Procedures of Theorem 5.1

This section describes the encoding and decoding procedure used in the proof of Theorem 5.1. The encoding and decoding procedures for Theorem 6.1 require very small changes, and so the modifications are given in Appendix B.

A.1 The Encoding Procedure

This section gives the decoding procedure. Recall that the sets $\mathsf{Q_{fresh}}, \mathsf{Slts}, \mathsf{Qrs}, \mathsf{Q}_j, \mathsf{MouseSlts}_j$ are defined as mentioned in the setup in Section 5. We do not explicitly repeat these definitions to avoid cluttering.

The encoding algorithm $\mathsf{Encode}(\sigma, U, h)$:

1. Parse U as $\{a_1, \ldots, a_u\}$ where $a_1, \ldots, a_u$ are in lexicographical order.

2. Set $j \leftarrow 0$.

3. Initialize the set P to be empty, the lists Cases and $\tilde{h}$ to be empty, and the set of nodes in the query graph to $[N]$ and the set of edges to be empty.

4. For $i = 1, \ldots, u$, do:

   (a) If $a_i \in \mathsf{Slts}(a)$ for some $a \in U_{\mathsf{fresh}}$, let $r \in [jT]$ be the smallest index such that $\mathsf{Q_{fresh}}[r] = (a_j, *)$. Add r to P, and continue to next iteration.

   (b) Otherwise, increment j, set $a'_j \leftarrow a_i$. Run $\mathcal{A}_2$ on (σ, a'_j), answering its queries to h. For every query made by $\mathcal{A}_2$ to h on (a, α) with answer a', add an edge (a, a') with label α to the query graph. Add $h(a, \alpha)$ to the list $\tilde{h}$ if it was not already added earlier. Classify the query as NEW, REPEATEDMOUSE, SEMIFAMILIARREPEATED, REPEATEDNONMOUSE as defined in Section 5.

   (c) Determine the category of the mouse structure using the classification in Section 5.1. According to the categorization do the following.

   - If the mouse structure gets categorized in case 1, then the mouse structure contains a NEW query (a, α) such that $h(a, \alpha) = a$. Let the index of the query in Q_j be n. Add n to the list L_1, set $\mathsf{Cases}[j] \leftarrow 1$ and remove the entry $h(a, \alpha)$ from $\tilde{h}$.

   - Else, if the mouse structure gets categorized in case 2, the mouse structure contains two distinct NEW queries $(a_1, \alpha_1), (a_2, \alpha_2)$ such that $h(a_1, \alpha_1) = h(a_2, \alpha_2)$, and (a_1, α_1) is queried before (a_2, α_2). Let the indices of the two queries be n_1, n_2, respectively, in Q_j. Add (n_1, n_2) to the list L_2, set $\mathsf{Cases}[j] \leftarrow 2$ and remove the entry $h(a_2, \alpha_2)$ from $\tilde{h}$. Continue to next iteration.

   - Else, if the mouse structure gets categorized in case 3, then it contains a NEW query (a, α) such that $h(a, \alpha) = a'$ and $a' \in \mathsf{MouseSlts}_k$ for some $k < j$. Let the index of the NEW query in Q_j be n and the lexicographic order of the salt a' in $\mathsf{MouseSlts}_k$ be o. Add (k, o, n) to a list L_3, set $\mathsf{Cases}[j] \leftarrow 3$ and remove the entry $h(a, \alpha)$ from $\tilde{h}$. Continue to next iteration.

   - Else, if the mouse structure gets categorized in case 4, then it contains a REPEATEDNON-MOUSE query (a, α) such that $h(a, \alpha) = a$. Let n the smallest index of the query (a, α) in $\mathsf{Q_{fresh}}$. Add n to the set S_4, set $\mathsf{Cases}[j] \leftarrow 3$ and remove the entry $h(a, \alpha)$ from $\tilde{h}$. Continue to next iteration.

   - Else, if the mouse structure is categorized in case 5, there is at least one salt d such that $d \in \mathsf{MouseSlts}_i$ for some $i < j$ and there are at most $B - 2$ edges in the mouse structure between d and s where s is an answer to a NEW query. If there are several possible candidate pairs, choose one where the number of edges is the smallest between the source and the destination salt. Let the path back from d to s in the query graph when (a, α) is queried be P. If we have "encountered a large multi-collision" (as defined in Section 5.3) in P, then set $\mathsf{Cases}[j] \leftarrow 5m$ and goto MULCOL below. Else, suppose that the NEW query that results with the source salt s is (a, α), the destination salt is d and $d \in \mathsf{MouseSlts}_i$ for some $i < j$. Construct the tuple consisting of the following.[13]

---

[13] We list the elements of the tuple in a table along with the sizes of each element to enable a reader cross check the size calculations we did in section 5.

36

Element	Size in bits
i	$\log u$
index of query (a, α) in Q_j	$\log T$
lexicographical index of d in $\mathsf{MouseSlts}_i$	$\log B$
path P	$\log B m_0^{(B-2)}$

Store the tuple in a list L_5, set $\mathsf{Cases}[j] \leftarrow 5a$ it removes the entry $h(a, \alpha)$ from $\tilde{h}$. Continue to next iteration.

- Else, if the mouse structure gets categorized in case 6, sub-categorize into sub-cases (a) to (d) as specified in Section 5.4.

  - If the mouse structure is categorized in (a), define queries q_1, q_2, q_3 as defined under the heading "Handling sub-case (a)" in Section 6. Let the path back from input salt of q_3 to the answer of q_1 in the query graph when q_1 was queried be path P. If there was a large multi-collision encountered on path P, set $\mathsf{Cases}[j] = 6m$ and goto line MULCOL. Otherwise define a tuple consisting of the following.

Element	Size in bits
the indices of q_1, q_2 in Q_j	$2 \log T$
the smallest index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P	$\log B m_0^{(B-2)}$

    Add the tuple to the list L_{6a}. Remove the answers corresponding to queries q_1, q_2 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6a$.

  - If the mouse structure is categorized in (b), define queries q_1, q_2, q_3, q_4 as defined under the heading "Handling sub-case (b)" in Section 6. Let the path back from the input salt of q_3 to the answer of q_1 in the query graph when q_1 was queried be P_1. Let the path back from the input salt of q_4 to the answer of q_2 in the query graph when q_2 was queried be P_2. If there was a large multi-collision encountered on either P_1 or P_2, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise define the following tuple.

Element	Size in bits
the indices of q_1, q_2 in Q_j	$2 \log T$
the smallest indices of q_3, q_4 in $\mathsf{Q}_{\mathsf{fresh}}$	$2 \log FT$
paths $\mathsf{P}_1, \mathsf{P}_2$	$\log B^2 m_0^{2(B-2)}$

    Add this tuple to the list L_{6b}. Remove the answers corresponding to queries q_1, q_2, q_4 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6b$.

  - If the mouse structure is categorized in (c), define queries q_1, q_2, q_3 as defined under the heading "Handling sub-case (c)" in Section 6. Further, we categorize to scenarios 1 or 2, as defined there.

    i. If it is in scenario 1 let the path back to the answer salt of q_1 from answer salt of q_3 in the query graph when q_1 was queried be P_1 and let the path back from the input salt of q_3 to the input salt of q_2 in the query graph when q_3 was queried be P_2. If there is a large multi-collision encountered either P_1 or P_2, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise define the following tuple.

Element	Size in bits
the index of q_1 in Q_j	$\log T$
the smallest index of q_2 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_1	$\log B m_0^{(B-2)}$

    Add the tuple to the list L_{6c1}. Define the following tuple.

Element	Size in bits
the smallest index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_2	$\log B m_0^{(B-2)}$

Add the tuple to the set S_{6c1}.

Remove the answers corresponding to queries q_1, q_3 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6c1$.

ii. Otherwise if it is in scenario 2, define q_4 as mentioned in scenario 2 under heading "Handling sub-case (c)" in Section 6. Let the path back to the answer salt of q_1 from the answer salt of q_3 in the query graph when q_1 was queried be P_1. Let the path back to the answer salt of q_4 from the input salt of q_4 in the query graph when q_4 was queried be P_2. If there is a large multi-collision encountered on either P_1 or P_2, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise, define the following tuple.

Element	Size in bits
the index of q_1 in Q_j	$\log T$
the smallest index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
paths P_1	$\log Bm_0^{(B-2)}$

Add the tuple to the list L_{6c2}. Define the following tuple.

Element	Size in bits
the smallest index of q_4 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_2	$\log Bm_0^{(B-2)}$

Add this tuple to the list L_{6c2}. Remove the answers corresponding to queries q_1, q_4 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6c2$.

– If the mouse structure is categorized in (d), define queries q_1, q_2, q_3 as defined under the heading "Handling sub-case (d)" in Section 6. Further, categorize into scenarios 1, 2, 3 or 4, as defined there.

i. If it is in scenario 1 let the path back from the input salt of q_2 to the salt that is the answer of q_1 when in the query graph when q_3 is queried be P_1 and let the path back from the input salt of q_3 to the salt that is the answer of q_1 when in the query graph when q_3 is queried be P_2 If there is a large multi-collision encountered either P_1 or P_2, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise define the following tuple.

Element	Size in bits
the index of q_1 in Q_j	$\log T$
the smallest index of q_2 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_1	$\log Bm_0^{(B-2)}$

Add the tuple to the list L_{6d1}. Define the following tuple.

Element	Size in bits
the smallest index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_2	$\log Bm_0^{(B-2)}$

Add the tuple to the set S_{6d1}.

Remove the answers corresponding to queries q_1, q_3 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6d1$.

ii. If it is in scenario 2, define q_4 as mentioned in scenario 2 under heading "Handling sub-case (d)" in Section 6. Let the path back to the answer salt of q_4 from the input salt of query q_2 in the query graph when q_4 is queried be P_1. Let the path back from the input salt of q_3 to the answer salt of query q_1 in the query graph when q_4 is queried be P_2. Let the path back from the input salt of q_4 to the answer salt of query q_1 in the query graph when q_4 is queried be P_3. If there is a large multi-collision encountered either $\mathsf{P}_1, \mathsf{P}_2$ or P_3, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise define the following tuple.

Element	Size in bits
the index of q_1 in Q_j	$\log T$
the smallest indices of q_2, q_3 in $\mathsf{Q}_{\mathsf{fresh}}$	$2 \log FT$
paths $\mathsf{P}_1, \mathsf{P}_2$	$\log B^2 m_0^{2(B-2)}$

Add the tuple to the list L_{6d2}. Define the following tuple.

Element	Size in bits
the smallest index of q_4 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_3	$\log Bm_0^{(B-2)}$

Add the tuple to the set S_{6d2}.

Remove the answers corresponding to queries q_1, q_3, q_4 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6d2$.

iii. If it is in scenario 3, let the path back to the answer salt of q_4 from the first body node of the mouse structure in the query graph when q_4 is queried be P_1. Let the path back from the input salt of q_2 to the first body node in the query graph when q_3 was queried for the first time be P_2. Let the path back from the input salt of q_4 to the answer salt of query q_1 in the query graph when q_1 is queried be P_3. Let the path back from the input salt of q_3 to the first body node of the mouse structure in the query graph when q_3 is queried for the first time be P_4. If there is a large multi-collision encountered either $\mathsf{P}_1, \mathsf{P}_2, \mathsf{P}_3$ or P_4, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise define the following tuple.

Element	Size in bits
the index of q_1 in Q_j	$\log T$
the smallest indices of q_2, q_4 in $\mathsf{Q}_{\mathsf{fresh}}$	$2\log FT$
paths $\mathsf{P}_1, \mathsf{P}_2, \mathsf{P}_3$	$3\log Bm_0^{(B-2)}$

Add the tuple to the list L_{6d3}. Define the following tuple.

Element	Size in bits
the smallest index of q_3 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
path P_4	$\log Bm_0^{(B-2)}$

Add the tuple to the set S_{6d3}. Remove the answers corresponding to queries q_1, q_3, q_4 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6d3$.

iv. If it is in scenario 4, define q_5 as mentioned in scenario 4 under heading "Handling sub-case (d)" in Section 6. Let the path back to the answer salt of q_5 from the input salt of q_2 in the query graph when q_5 was queried for the first time be P_1. Let the path back from the input salt of q_3 to the first body node in the query graph when q_5 was queried for the first time be P_2. Let the path back to the answer salt of q_4 from the first body node of the mouse structure in the query graph when q_5 was queried for the first time be P_3. Let the path back from the input salt of q_4 to the answer salt of query q_1 in the query graph when q_4 is queried for the first time be P_4. Let the path back to the input salt of q_5 to the first body node in the query graph when q_5 was queried for the first time be P_5. If there is a large multi-collision encountered either $\mathsf{P}_1, \mathsf{P}_2, \mathsf{P}_3, \mathsf{P}_4$, or P_5, set $\mathsf{Cases}[j] = 6m$ and goto MULCOL. Otherwise define the following tuple.

Element	Size in bits
the index of q_1 in Q_j	$\log T$
the smallest indices of q_2, q_3, q_4 in $\mathsf{Q}_{\mathsf{fresh}}$	$3\log FT$
paths $\mathsf{P}_1, \mathsf{P}_2, \mathsf{P}_3, \mathsf{P}_4$	$2\log B^2 m_0^{2(B-2)}$

Add the tuple to the list L_{6d4}. Define the following tuple.

Element	Size in bits
the smallest index of q_5 in $\mathsf{Q}_{\mathsf{fresh}}$	$\log FT$
paths P_5	$\log Bm_0^{(B-2)}$

Add the tuple to the set S_{6d4}. Remove the answers corresponding to queries q_1, q_3, q_4, q_5 from $\tilde{h}$ and set $\mathsf{Cases}[j] = 6d4$.

− Continue to next iteration.

39

- MulCol: If there is a query in this mouse structure whose answer is the center of the multi-collision, and this query was in an earlier mouse structure, do nothing. Else if there is a query in this mouse structure whose answer is the center of the multi-collision, and this query was not in any earlier mouse structure. If the center multi-collision of the multi-collision is s and $(s, \mathsf{Q}_s) \in S_{\mathsf{MC}}$, add the first index in $\mathsf{Q}_{\mathsf{fresh}}$ of all the queries in the large multi-collision into the set that were not already in Q_s to it. Otherwise, set s to be the center of the multi-collisions and Q_s to be the set of the first index in $\mathsf{Q}_{\mathsf{fresh}}$ of the queries involved in the large multi-collision., and remove the answer of all these queries from $\tilde{h}$. Add (s, Q_s) to S_{MC}, set $\mathsf{Cases}[j] \leftarrow m$.

5. Add all the evaluations of $h(a, \alpha)$ to $\tilde{h}$ in lexicographical order for all (a, α) that were never queried.

6. Output F, U', P, L_1, L_2, L_3, L_4, L_5, S_{MC}, L_{6a}, L_{6b}, L_{6c1}, L_{6c2}, L_{6d1}, L_{6d2}, L_{6d3}, L_{6d4}, S_{6c1}, S_{6c2}, S_{6d1}, S_{6d2}, S_{6d3}, S_{6d4}, $\tilde{h}$.

A.2 The Decoding Procedure

This section gives the decoding procedure.

The decoding algorithm Decode:

1. Parse U' as $\{a'_1, \ldots, a'_F\}$ where $a'_1, \ldots, a'_F$ are in lexicographical order.

2. Set $j \leftarrow 0$.

3. Initialize the set of nodes in the query graph to $[N]$ and the set of edges to empty.

4. For $j = 1, \ldots, F$ do:

   (i) Run $\mathcal{A}_2$ on (σ, a'_j), answering its queries to h.

   (ii) Answer the i-th query made by $\mathcal{A}_2$ to h on (a, α) as follows.

   (iii) If for some $(a', \mathsf{Q}_{a'}) \in S_{\mathsf{MC}}$, $((j-1) \cdot T + i) \in \mathsf{Q}_a$, then set $h(a, \alpha) \leftarrow a'$. Return a'.

   (iv) If $((j-1) \cdot T + i) \in \mathsf{P}$ then add a to U.

   (v) If $((j-1) \cdot T + i) \in S_4$ then set $h(a, \alpha) = a$ and return a.

   (vi) If $((j-1) \cdot T + i)$ appears as q_4 in any tuple of L_{6b}, then let a' be the answer of the query q_3 listed in the tuple. Set $h(a, \alpha) \leftarrow a'$ and return a'.

   (vii) If $((j-1) \cdot T + i)$ appears as q_3 in any tuple of S_{6c1}, then let the path P_2 in the tuple starting from a end in a'. Set $h(a, \alpha) \leftarrow a'$ and return a'.

   (viii) If $((j-1) \cdot T + i)$ appears as q_4 in any tuple of S_{6c2}, then let the path P_2 in the tuple starting from a end in a'. Set $h(a, \alpha) \leftarrow a'$ and return a'.

   (ix) If $((j-1) \cdot T + i)$ appears as q_3 in any tuple of S_{6d1}, then let the path P_2 in the tuple starting from a end in a'. Check for which tuple in the list L_{6d1} the path P_1 from the answer of the query q_2 of the tuple ends in a'. Let a'' be the answer of the query q_2 of that tuple. Set $h(a, \alpha) \leftarrow a''$ and return a''.

   (x) If $((j-1) \cdot T + i)$ appears as q_3 in any tuple of S_{6d2}, then let a' be the answer of the query q_2 of that tuple. Set $h(a, \alpha) \leftarrow a'$ and return a'.

   (xi) If $((j-1) \cdot T + i)$ appears as q_4 in any tuple of S_{6d2}, then let the path P_3 in the tuple starting from a end in a'. Check for which tuple in the list L_{6d2} the path P_2 from the answer of the query q_2 of the tuple ends in a'. Let a'' be the answer of the query q_2 of that tuple. Let the path P_1 in the tuple starting from a'' end in a'''. Set $h(a, \alpha) \leftarrow a'''$ and return a'''.

(xii) If $((j-1) \cdot T + i)$ appears as q_3 in any tuple of S_{6d3}, then let the path P_4 in the tuple starting from a end in a'. Check for which tuple in the list L_{6d2} the path P_2 from the answer of the query q_2 of the tuple ends in a'. Let a'' be the answer of the query q_2 of that tuple. Set $h(a, \alpha) \leftarrow a''$ and return a''.

(xiii) If $((j-1) \cdot T + i)$ appears as q_4 in any tuple of L_{6d3}, then let a' be the first body node of the mouse structure located when answering the query q_3 corresponding to the tuple. Let the path P_1 in the tuple starting from a' end in a''. Set $h(a, \alpha) \leftarrow a''$ and return a''.

(xiv) If $((j-1) \cdot T + i)$ appears as q_3 in any tuple of L_{6d4}, then let a' be the answer of the query q_2 of the tuple. Set $h(a, \alpha) \leftarrow a'$ and return a'.

(xv) If $((j-1) \cdot T + i)$ appears as q_5 in any tuple of S_{6d4}, then let the path P_5 in the tuple starting from a end in a'. Check for which tuple in the list L_{6d2} the path P_2 from the answer of the query q_3 of the tuple ends in a'. Let a'' be the input of the query q_2 of that tuple. Let the path P_1 in the tuple starting from a'' end in a'''. Set $h(a, \alpha) \leftarrow a'''$ and return a'''.

(xvi) If $((j-1) \cdot T + i)$ appears as q_4 in any tuple of L_{6d4}, then let a' be the first body node of the mouse structure located when answering the query q_5 corresponding to the tuple. Let the path P_3 in the tuple starting from a' end in a''. Set $h(a, \alpha) \leftarrow a''$ and return a''.

(xvii) If $\mathsf{Cases}[j] = 1$ and i is the front element of the list L_1, then remove i from L_1, set $h(a, \alpha) \leftarrow a$ and return a.

(xviii) If $\mathsf{Cases}[j] = 2$ and (k, i) is the front element of the list L_2 for some $k < i$, let a' be the answer of the k-th query. Remove (k, i) from L_2, set $h(a, \alpha) \leftarrow a'$ and return a'.

(xix) If $\mathsf{Cases}[j] = 3$ and (k, o, i) is the front element of the list L_3 for some $k < j$, let a' be the salt that has lexicographic order o in $\mathsf{MouseSlts}_k$-th query. Remove (k, o, i) from L_3, set $h(a, \alpha) \leftarrow a'$ and return a'.

(xx) If $\mathsf{Cases}[j] = 5a$ and (k, o, i) is the front element of the list L_3 for some $k < j$, let a' be the salt that has lexicographic order o in $\mathsf{MouseSlts}_k$-th query. Remove (k, o, i) from L_5, set $h(a, \alpha) \leftarrow a'$ and return a'.

(xxi) If $\mathsf{Cases}[j] = 6a$ and if the tuple which is at the front of list L_{6a} has i as the index of q_2, then suppose the query listed as q_3 had answer a', set $h(a, \alpha) \leftarrow a'$ and return a'. Remove the tuple from the list L_{6a} if both the queries q_1, q_2 in the tuple have been answered.

(xxii) If $\mathsf{Cases}[j] = 6a$ and the tuple which is at the front of list L_{6a} has i as the index of q_1, then suppose the query listed as q_3 had answer a'. Let the path P in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6a} if both the queries q_1, q_2 in the tuple have been answered.

(xxiii) If $\mathsf{Cases}[j] = 6b$ and if the tuple which is at the front of list L_{6b} has i as the index of q_1, then suppose the query listed as q_3 had answer a'. Let the path P_1 in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6b} if both the queries q_1, q_2 in the tuple have been answered.

(xxiv) If $\mathsf{Cases}[j] = 6b$ and if the tuple which is at the front of list L_{6b} has i as the index of q_2, then suppose the query listed as q_4 had answer a'. Let the path P_2 in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6b} if both the queries q_1, q_2 in the tuple have been answered.

(xxv) If $\mathsf{Cases}[j] = 6c1$ and if the tuple which is at the front of list L_{6c1} has i as the index of q_1, then suppose the query listed as q_3 had answer a'. Let the path P_1 in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6c1}.

(xxvi) If $\mathsf{Cases}[j] = 6c2$ and if the tuple which is at the front of list L_{6c2} has i as the index of q_1, then suppose the query listed as q_3 had answer a'. Let the path P_1 in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6c2}.

(xxvii) If $\mathsf{Cases}[j] = 6d1$ and if the tuple which is at the front of list L_{6d1} has i as the index of q_1, then let a' be the salt a' computed when answering the query q_3 that corresponded to this tuple. Set $h(a, \alpha) \leftarrow a'$ and return a'. Remove the tuple from the list L_{6d1}.

(xxviii) If $\mathsf{Cases}[j] = 6d2$ and if the tuple which is at the front of list L_{6d2} has i as the index of q_1, then then let a' be the salt a' computed when answering the query q_3 that corresponded to this tuple. Set $h(a, \alpha) \leftarrow a'$ and return a'. Remove the tuple from the list L_{6d3}.

(xxix) If $\mathsf{Cases}[j] = 6d3$ and if the tuple which is at the front of list L_{6d3} has i as the index of q_1, then suppose the path P_2 from the input salt of the query listed as q_3 ended in salt a'. Set $h(a, \alpha) \leftarrow a'$ and return a'. Remove the tuple from the list L_{6d3}.

(xxx) If $\mathsf{Cases}[j] = 6d3$ and if the tuple which is at the front of list L_{6d3} has i as the index of q_1, then suppose the query listed as q_4 had answer a'. Let the path P_3 in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6d3}.

(xxxi) If $\mathsf{Cases}[j] = 6d4$ and if the tuple which is at the front of list L_{6d3} has i as the index of q_1, then suppose the query listed as q_4 had answer a'. Let the path P_4 in the tuple from a' in the query graph end in salt a''. Set $h(a, \alpha) \leftarrow a''$ and return a''. Remove the tuple from the list L_{6d4}.

5. Populate the rest of h by remaining elements in $\tilde{h}$ in lexicographical order.

6. Output (U, h).

B   Changes for Theorem 6.1

In the encoding procedure, all the checks for multi-collisions are replaced by the check for the number of possible destination salts at the specified distance d from the source salt: if there are more than $e^3 u^2/2$ possible salts such that their shortest path to the destination salt in the query graph so far has d edges, we encode the set of multi-collisions. For cases $6c$ and $6d$, the indices of queries q_1, q_2, q_3 are all stored in lists (i.e. we do not need the sets which thereby reduces the number of paths). The code for the decoding procedure is modified accordingly.