CONJURA · OPEN PROBLEM

A Tight Quantum Time-Space Tradeoff for Two-Block Merkle-Damgård Collisions

Is the known quantum preprocessing attack optimal at two blocks?

Status. Statement: AI-written from Akshima, Siyao Guo and Qipeng Liu, *Time-Space Lower Bounds for Finding Collisions in Merkle-Damgård Hash Functions*, CRYPTO 2022 (IACR ePrint 2022/885), not yet checked by a human. The classical analogue is settled — $\tilde{\Theta}(ST/N + T^2/N)$ at two blocks, by the source paper and by [ACDW20] — while in the quantum setting the source paper reports that "no matching bounds are known, even for $B = 2$ and $B = T$". Nothing is settled here; the source paper poses this as a question and takes no position on the answer.

Category. *Quantum*, *Symmetric-Key Cryptography*, *Idealized Models*.

Conjecture (informal). *Classically, an attacker who preprocesses a compression function into S bits of advice and then makes T online queries finds a* two-block *salted Merkle-Damgård collision with probability about ST over the digest space, and no better — that much is proved, and it is strictly less than the roughly ST^2 available when the collision may be arbitrarily long. Quantumly, where the advice may be S qubits and the T queries are superposition queries, no such matching pair of bounds is known at any block length. The best attack anyone knows succeeds with probability about $ST^2 + T^3$ over the digest space, and it achieves that for every block bound from two upwards; the best proved bound is a factor of about S away. The conjecture is that at two blocks the attack is the truth: that the quantum two-block advantage is about $ST^2 + T^3$ over the digest space. Note what this would mean — unlike the classical case, the quantum two-block problem would then be no harder than the unbounded-length one, so the classical jump at two blocks would have no quantum counterpart.*

1 The setting

Merkle-Damgård iterates a compression function over message blocks starting from a salt. Against a preprocessing (non-uniform)

attacker the salt is what makes collision-finding non-trivial; the model is the auxiliary-input random-oracle model of Unruh [Unr07], in which an unbounded first stage sees the whole random oracle H and outputs S bits of advice, and a second stage receives the advice and a uniformly random salt and makes T oracle queries.

Classically, the picture is now sharp at both ends. With no bound on the length of the collision the advantage is $\tilde{\Theta}(ST^2/N)$ [CDGS18]; at two blocks it drops to $\tilde{\Theta}(ST/N + T^2/N)$, proved by Akshima, Cash, Drucker and Wee [ACDW20] and reproved more simply as the source paper's Theorem 2. That drop — a factor of T between $B = 2$ and $B = T$ — is the "security jump" the source paper's classical results establish, and it is the reason short collisions are studied separately at all.

The quantum question is the same question with quantum resources: the advice is S qubits and the online stage makes T quantum queries to H. The source paper's own summary of the state of the art, in its discussion of open problems (p. 10), is that "[m]otivated by analyzing post-quantum non-uniform security, several recent works [CGLQ20, GLLZ21] studied the same question in the quantum setting, in which the adversary is given S-(qu)bit of advice and T quantum oracle queries. However, unlike the classical setting, no matching bounds are known, even for $B = 2$ and $B = T$."

What is known. The source paper reports two numbers, and they are the only two:

- An attack achieving, in the paper's own notation, $O(ST^2/N + T^3/N)$ "for every $2 \le B \le T$" — so the best known quantum attack does not distinguish two blocks from unbounded length at all.
- A security bound the paper attributes to Guo, Li, Liu and Zhang [GLLZ21] and writes as $\Omega(ST^3/N)$, which it reads as suggesting "that the optimal attack may speed up the trivial quantum collision finding by a factor of S".

The two are a factor of about S apart — ST^2 against ST^3 with $T^2 \le N$ — and no case of any B is known to be tight. See Remark 3 on the paper's use of O and Ω here, which is the reverse of the usual convention and is reproduced rather than repaired.

What the source paper asks. "Is there a security jump for finding 2-block collisions and unbounded collisions in the quantum setting? Can we leverage our new proof for $B = 2$ to prove a tight security bound in the quantum setting?" (p. 10). Conjecture 1 below is the second of those two questions, in the affirmative direction; the first is related but is not the same statement, and Remark 4 explains why.

Why it is hard. The classical $B = 2$ proof that the paper offers to "leverage" works by reducing auxiliary-input security to *sequential* multi-instance security and then isolating a small list of "high knowledge gaining" events on the offline queries — more than S distinct salts carrying a one-block collision, more than S^2 colliding pairs of queries, more than S distinct salts carrying a self-loop — and showing each is unlikely even conditioned on the adversary having won every earlier round. Two things obstruct transplanting this. The events are counting statements about a *recorded* set of offline queries, and a quantum offline stage has no such transcript to count; the compressed-oracle machinery that substitutes for one does not obviously support conditioning on having won all previous rounds. And the target advantage would have to be quantum-specific: classically the sequential-multi-instance route recovers ST^2/N for unbounded length, whereas here the conjectured two-block answer already contains ST^2/N, so the same route cannot separate $B = 2$ from $B = T$ in the quantum setting the way it does classically. The source paper records the difficulty as the bare fact that no matching bound is known at any B, and offers no partial result in the quantum setting.

Why settling it matters. It would be the first tight quantum time-space tradeoff for collision-finding in a real hash-function mode, and the first at any block length. It also decides a structural question: the classical theory's organizing fact is that bounding the collision length buys real security, and if the conjecture holds then that fact has no quantum analogue at two blocks, so post-quantum non-uniform security would not inherit the classical short-collision advantage.

2 Notation and parameters

- N is the size of the chaining-value space and $M > N$ that of the block space; the compression function is $H : [N]\times[M] \to [N]$.
- $S \in \mathbb{N}$ is the number of advice qubits and $T \in \mathbb{N}$ the number of online quantum queries. Only queries are counted; the offline stage is computationally unbounded.
- Asymptotics are in N, uniformly over M, S and T; $\tilde{O}$, $\tilde{\Omega}$ and $\tilde{\Theta}$ suppress factors polylogarithmic in N. As is standard here, $T^2 \le N$ may be assumed, since otherwise a collision is found by the birthday attack.
- Definition 2 is not stated in the source paper, which is classical throughout and refers to [CGLQ20] and [GLLZ21] for the quantum model. See Remark 2.

3 The conjecture

Definition 1 (Merkle-Damgård). For $H : [N] \times [M] \to [N]$, a salt $a \in [N]$ and a message $m = (m_1, \ldots, m_\ell)$ with each $m_i \in [M]$, set $\mathsf{MD}_H(a, m) := H(a, m_1)$ if $\ell = 1$ and

$$\mathsf{MD}_H(a, m) := H\big(\mathsf{MD}_H(a, (m_1, \ldots, m_{\ell-1})),\ m_\ell\big)$$

if $\ell > 1$. A message with ℓ blocks is an *ℓ-block message.*

Definition 2 (Two-block collision resistance against quantum auxiliary input). An *(S,T)-quantum-AI adversary* is a pair $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in which $\mathcal{A}_1$ is a computationally unbounded map from the whole function table of H to an S-qubit state ρ_H, and $\mathcal{A}_2$ is a quantum algorithm that takes ρ_H and a salt $a \in [N]$, makes at most T queries to the standard quantum-accessible oracle $|x, y\rangle \mapsto |x, y \oplus H(x)\rangle$, and outputs a pair of messages m_1, m_2.
The game returns 0 if m_1 or m_2 consists of more than two blocks; it returns 1 if $m_1 \neq m_2$ and $\mathsf{MD}_H(a, m_1) = \mathsf{MD}_H(a, m_2)$;

otherwise 0. Write

$$\mathsf{Adv}^{\mathsf{qAICR}}_{\text{2-MD}}(S,T) = \max_{\mathcal{A}} \Pr[\text{ the game returns } 1\,],$$

the maximum over all (S,T)-quantum-AI adversaries, the probability taken over uniform $H \leftarrow \{f : [N] \times [M] \to [N]\}$, uniform and independent $a \leftarrow [N]$, and the measurement outcomes of $\mathcal{A}_2$.

Conjecture 1 (tight quantum two-block tradeoff).

$$\mathsf{Adv}^{\mathsf{qAICR}}_{\text{2-MD}}(S,T) \;=\; \tilde{\Theta}\!\left(\frac{ST^2}{N} + \frac{T^3}{N}\right).$$

The lower half is the attack the source paper reports as best known for every $2 \le B \le T$; the open content is the matching security bound,

$$\mathsf{Adv}^{\mathsf{qAICR}}_{\text{2-MD}}(S,T) \;\le\; \tilde{O}\!\left(\frac{ST^2}{N} + \frac{T^3}{N}\right),$$

which would improve the ST^3/N figure the paper attributes to [GLLZ21] by a factor of T at two blocks.

Remark 1 (the source poses a question; the direction is a reading). The source paper asks whether a tight quantum bound can be proved at $B = 2$; it does not say what the tight value is, and it does not predict the answer. Conjecture 1 fixes the direction by taking the best known attack to be optimal, which is the only value the paper's own two reported numbers make available as a candidate. The complementary possibility — that a better quantum attack exists at two blocks, anywhere up to the ST^3/N figure — is left equally open by the source, and refuting Conjecture 1 is as much a resolution of the paper's question as proving it. A reviewer who thinks the statement should not be committed to a direction at all should say so.

Remark 2 (the model is supplied here, not taken from the source). The source paper is classical: it defines the classical auxiliary-input game (its Definitions 2 and 3, which Definition 2

follows block for block on the classical parts) and describes the quantum setting only as the one "in which the adversary is given S-(qu)bit of advice and T quantum oracle queries", citing [CGLQ20] and [GLLZ21]. Definition 2 spells out the standard reading of that phrase: quantum advice rather than classical, a single unbounded preprocessing stage, standard XOR-type superposition queries, and no bound on the online stage's time or memory beyond the query count. Each of those is a choice, and a different reading — classical advice, or a bounded-memory online stage — would be a different statement. This is the part of the statement most in need of a check against [CGLQ20] and [GLLZ21] themselves.

Remark 3 (the source's O and Ω are reproduced as printed). The source paper writes "[t]he $\Omega(ST^3/N)$ security bound by [GLLZ21]" and "the best-known attack achieves $O(ST^2/N + T^3/N)$", which inverts the usual convention: a security bound is an upper bound on advantage and an attack is a lower bound. The figures are quoted above in the paper's symbols without reinterpretation, and Conjecture 1 is stated in the conventional direction instead of relying on them. A reviewer should confirm both figures against [GLLZ21] directly rather than through this page.

Remark 4 (the security-jump question is a different statement). The paper's other quantum question — "[i]s there a security jump for finding 2-block collisions and unbounded collisions in the quantum setting?" — asks whether $\mathsf{Adv}^{\mathsf{qAICR}}_{\text{2-MD}}(S,T)$ is asymptotically smaller than its unbounded-length counterpart. Conjecture 1 does not settle it: even granting the conjecture, a jump would still exist if the unbounded-length quantum advantage turned out to be $\tilde{\Theta}(ST^3/N)$. The two are kept apart deliberately rather than merged into one statement with two clauses.

4 Bibliography

[ACDW20] Akshima, David Cash, Andrew Drucker, and Hoeteck Wee. *Time-space tradeoffs and short collisions in Merkle-Damgård hash functions.* In Daniele Micciancio and Thomas Ristenpart, editors, Advances in Cryptology — CRYPTO 2020, Part I, volume 12170 of Lecture Notes in Computer Science, pp. 157–186. Springer, 2020.

[AGL22] Akshima, Siyao Guo, and Qipeng Liu. *Time-space lower bounds for finding collisions in Merkle-Damgård hash functions.* In Yevgeniy Dodis and Thomas Shrimpton, editors, Advances in Cryptology — CRYPTO 2022, Part III, volume 13509 of Lecture Notes in Computer Science, pp. 192–221. Springer, 2022. (The source paper of this statement; IACR ePrint 2022/885.)

[CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. *Random oracles and non-uniformity.* In Jesper Buus Nielsen and Vincent Rijmen, editors, Advances in Cryptology — EUROCRYPT 2018, Part I, volume 10820 of Lecture Notes in Computer Science, pp. 227–258. Springer, 2018.

[CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. *Tight quantum time-space tradeoffs for function inversion.* In Sandy Irani, editor, 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, pp. 673–684. IEEE, 2020.

[GLLZ21] Siyao Guo, Qian Li, Qipeng Liu, and Jiapeng Zhang. *Unifying presampling via concentration bounds.* In Theory of Cryptography — 19th International Conference, TCC 2021, Part I, pp. 177–208, 2021.

[Unr07] Dominique Unruh. *Random oracles and auxiliary input.* In Alfred Menezes, editor, Advances in Cryptology — CRYPTO 2007, volume 4622 of Lecture Notes in Computer Science, pp. 205–223. Springer, 2007.