CONJURA · OPEN PROBLEM

Classical Presampling Against Quantum Queries

Can an auxiliary-input random oracle be replaced by a classically bit-fixed one, against a quantum distinguisher?

Status. Statement: AI-written from Siyao Guo, Qian Li, Qipeng Liu and Jiapeng Zhang, *Unifying Presampling via Concentration Bounds*, IACR ePrint 2020/1589; TCC 2021, LNCS 13042, pp. 177–208. Not yet checked by a human. Open, and known to be hard in a precise sense: the source paper proves that this conjecture implies the folklore conjecture that polynomial quantum speedups need structure, a central open problem in quantum computing related to the Aaronson-Ambainis conjecture. No refutation is known either.

Category. *Quantum Cryptography*, *Idealized Models*.

Conjecture (informal). *Classical non-uniform security proofs in the random oracle model rest on one reduction: an adversary holding S bits of arbitrary advice about a random oracle is no better off than an adversary facing an oracle whose values have been fixed in advance at $P \approx ST$ points of someone's choosing and are uniform elsewhere. That is presampling, and it turns hard auxiliary-input questions into easy bit-fixing ones. This conjecture says the same reduction survives when the online adversary queries the oracle in superposition — with a bound so generous that any dependence on the advice length at all is permitted. The interest is not that anyone expects to prove it. The source paper shows that even this weak form would settle a long-standing open problem about whether polynomial quantum speedups require structured inputs, which is why the paper's own quantum presampling theorem fixes coordinates "quantumly" rather than classically.*

1 The setting

In the auxiliary-input random oracle model (AI-ROM) of Unruh [Unr07] an adversary $(\mathcal{A}_1, \mathcal{A}_2)$ attacks a game in two stages: a computationally unbounded $\mathcal{A}_1$ sees the entire oracle H and

outputs S bits of advice; $\mathcal{A}_2$ receives the advice and makes T oracle queries. The model is the right one for non-uniform attackers, and it is awkward to analyse directly, because the advice can encode any global property of H whatsoever.

Presampling is the technique that makes it tractable. In the P-bit-fixing model the oracle is instead drawn from a P-bit-fixing source: an adversary fixes at most P input-output pairs of its choice, and the oracle is uniform on the remaining points. Bounds there are usually easy — for one-way functions, "if the challenge y is not in the list L, to invert y in the P-BF-ROM is as difficult as that in the ROM", giving $(P+T)/\min\{N,M\}$ in a line. Coretti, Dodis, Guo and Steinberger [CDGS18] proved the optimal transfer between the two models, quoted as Lemma 1 below: the two distributions are indistinguishable to a classical T-query distinguisher up to $(S+\log 1/\gamma)T/P+\gamma$.

The question. For post-quantum security the online adversary should be allowed superposition queries, giving the AI-QROM. The source paper opens its Section 3 by asking whether Lemma 1 carries over: "we ask the question: is it possible to leverage Lemma 3 (and Theorem 3) to the quantum world? The following conjecture formally states that the presampling technique could reduce security proofs in AI-QROM to those in the simpler 'P-BF-QROM' ". Conjecture 1 below is that statement, and it deliberately asks for much less than the classical lemma delivers: "Note that this conjecture is weaker than Section 2.5 in the sense that the dependency on S can be arbitrary, but Lemma 3 is polynomial in S". The function h is unconstrained and the error term carries no additive slack parameter γ.

Why one would want it. Because the ideal version makes the hard bounds fall out. The source paper says what it would buy for function inversion: "If the ideal presampling holds, we can get the lower bound of function inversion in the AI-QROM easily, without using any involved techniques" (p. 5), the two cases being, over the page break, "either the challenge image is in one of the fixed coordinates (with probability ST/N), or it is outside the fixed coordinates, in which we argue the success probability by simply using the existing lower bound of Grover's search" (p. 6). That is

a two-case argument reproving a result of Chung, Guo, Liu and Qian [CGLQ20] whose original proof needs compressed oracles and a multi-instance reduction.

Why it is hard: a named barrier. The source paper's first contribution is that this cannot be cheap. Its Theorem 4 shows Conjecture 1 implies an equivalent form (its Conjecture 3) of the folklore conjecture that "any quantum algorithm can be approximated on most inputs by an efficient classical algorithm" — a question "dating back to (according to [AA11]) 1999 or earlier", included twice in Aaronson's list of ten semi-grand challenges, and the target of the Aaronson-Ambainis conjecture [AA11] on influences of low-degree functions. The paper records the status of both: "both Conjecture 1 and Conjecture 2 are still quite open, and they are proven only for some class of functions". Its own summary of the consequence is blunt: "As discussed in the last section, the natural extension of Lemma 3 does not work in the quantum world; otherwise, we can prove AA conjecture".

The proof of Theorem 4 needs only the $S = 1$ case: a single advice bit saying whether the distinguisher's acceptance probability on the sampled oracle is high. So no weakening in the advice length escapes the barrier, which is why the conjecture below lets h be arbitrary.

What the source did instead, and what remains. Faced with the barrier, the paper changes the target model rather than the technique: it re-characterizes the classical P-bit-fixing model as "conditioned on a P-query algorithm accepting", which generalizes verbatim to a P-query *quantum* algorithm, and proves an optimal presampling theorem into that quantum bit-fixing model (its Theorem 1). The gap between the two is exactly the content of this page: "To overcome the barrier, we may need to 'quantumly' fix P input-output pairs and avoid the AA conjecture barrier. However, it is not clear how to 'fix quantumly' or 'fix in superposition' ". The quantum bit-fixing model is the one all later work uses; the classically bit-fixed statement below has stayed where the paper left it.

2 Notation and parameters

- N is the size of the oracle's domain and M the size of its range, so an oracle is an element of $[M]^N$ where $[n] = \{1, \ldots, n\}$.
- $S \in \mathbb{N}$ is the advice length in bits and $T \in \mathbb{N}$ the number of online queries; $P \in \mathbb{N}$ is the number of coordinates the bit-fixing source may fix.
- $f : [M]^N \to \{0,1\}^S$ is an arbitrary (unbounded) function: the advice-generating first stage.
- X is uniform on $[M]^N$; $\mathcal{A}^Y(z)$ denotes the distinguisher run with input z and oracle access to an oracle drawn from the source Y, and $\Pr[\mathcal{A}^Y(z) = 1]$ averages over the draw of the oracle, the input and the algorithm's measurement.
- Quantum queries are the standard superposition queries to the oracle; the distinguisher is computationally unbounded and only its query count is bounded.
- $h : \mathbb{N} \to \mathbb{R}^+$ and the constant C in Conjecture 1 are existentially quantified: the conjecture asserts that *some* universal C and *some* function h make the bound hold. See Remark 1.

3 The conjecture

Definition 1 ((N,M)-source, P-bit-fixing). An (N,M)-source is a random variable on $[M]^N$; since an oracle $O : [N] \to [M]$ is an element of $[M]^N$, a distribution over oracles is such a source. An (N,M)-source is *P-bit-fixing* if it is fixed on at most P coordinates and uniform on the rest.

Lemma 1 (Classical presampling, Coretti-Dodis-Guo-Steinberger, quoted as the source's Lemma 3). *Let X be uniform over $[M]^N$ and $Z := f(X)$ for an arbitrary $f : [M]^N \to \{0,1\}^S$. For any $\gamma > 0$ and $P \in \mathbb{N}$ there is a family $\{Y_z\}_{z \in \{0,1\}^S}$ of convex combinations Y_z of P-bit-fixing (N,M)-sources such that for any classical distinguisher $\mathcal{D}$ taking an S-bit input and querying at most $T < P$ coordinates of its*

oracle,

$$\left|\Pr[\mathcal{D}^X(f(X))=1]-\Pr[\mathcal{D}^{Y_{f(X)}}(f(X))=1]\right| \leq \frac{(S+\log 1/\gamma)\cdot T}{P}+\gamma.$$

Conjecture 1 (Guo-Li-Liu-Zhang, Conjecture 4). *Let X be distributed uniformly over $[M]^N$ and $Z := f(X)$, where $f : [M]^N \to \{0,1\}^S$ is an arbitrary function. For any $P \in \mathbb{N}$, there exists a family $\{Y_z\}_{z\in\{0,1\}^S}$ of convex combinations Y_z of P-bit-fixing (N,M)-sources such that for any quantum distinguisher $\mathcal{A}$ taking an S-bit input and making T quantum queries of its oracle,*

$$\left|\Pr[\mathcal{A}^X(f(X))=1]-\Pr[\mathcal{A}^{Y_{f(X)}}(f(X))=1]\right| \leq h(S)\cdot T\cdot\left(\frac{\log M}{P}\right)^C.$$

Here C is a universal constant and $h : \mathbb{N} \to \mathbb{R}^+$ can be any function.

Theorem 1 (The barrier, the source's Theorem 4). *Conjecture 1 implies the source's Conjecture 3, and hence its Conjecture 1: that for every quantum algorithm $\mathcal{A}$ making T queries to a Boolean input and every $\varepsilon,\delta > 0$ there is a deterministic classical algorithm making $\mathrm{poly}(T,1/\varepsilon,1/\delta)$ queries that approximates $\mathcal{A}$'s acceptance probability to within ε on a $(1-\delta)$ fraction of inputs. The proof uses Conjecture 1 only for $S = 1$.*

Remark 1 (Reading the quantifiers). The source prints C and h after the displayed inequality — "Here C is a universal constant and $h : \mathbb{N} \to \mathbb{R}^+$ can be any function" — rather than in the statement, so their quantification has to be read off the surrounding text. Both are read here as existential, since that is what makes the conjecture weaker than Lemma 1 in the way the paper says it is: h arbitrary relaxes the classical linear dependence on S, and a smaller C makes the bound larger and so easier to achieve. Note that T appears to the first power and outside the C-th power, exactly as printed. A reader who takes C to be universally quantified is looking at a different, stronger statement.

Remark 2 (What a refutation would have to look like). Theorem 1 says a proof is at least as hard as a well-known open problem; it says nothing about a refutation, and the source offers no evidence in either direction. A counterexample would be an advice function

f, a parameter P, and a quantum distinguisher that separates the auxiliary-input distribution from every convex combination of P-bit-fixing sources by more than any $h(S) \cdot T \cdot (\log M/P)^C$ permits — that is, a quantum distinguisher exploiting global structure in a random oracle that no classical pre-fixing of coordinates can imitate. Refuting the conjecture would not touch the Aaronson-Ambainis conjecture; it would instead confirm that the quantum bit-fixing route the source took is the necessary one.

Remark 3 (What has and has not happened since). Liu's quantum bit-fixing framework [Liu23] and the subsequent tight quantum time-space tradeoff for permutation inversion [ABCGY25] both work in the *quantum* bit-fixing model that the source introduced as its way around this barrier, and neither addresses the classically bit-fixed statement above. On the barrier's far side, the Aaronson-Ambainis conjecture itself remains open: it is known for completely bounded block-multilinear forms [BSW22] and for a non-negligible fraction of random restrictions [Bha24], and the best bound for general functions is still far from the conjectured one. This is a targeted check made on 23 August 2026, not an exhaustive sweep.

4 Bibliography

[AA11] Scott Aaronson and Andris Ambainis. *The need for structure in quantum speedups.* Innovations in Theoretical Computer Science (ITCS), pp. 338–352, 2011.

[ABCGY25] Akshima, Tyler Besselman, Kai-Min Chung, Siyao Guo and Tzu-Yi Yang. *Tight quantum time-space tradeoffs for permutation inversion.* arXiv:2510.12112, October 2025.

[Bha24] Sreejata Kishor Bhattacharya. *Aaronson-Ambainis conjecture is true for random restrictions.* arXiv:2402.13952, February 2024; ITCS 2025.

[BSW22] Nikhil Bansal, Makrand Sinha and Ronald de Wolf. *Influence in completely bounded block-multilinear forms and classical simulation of quantum algorithms.* arXiv:2203.00212, March 2022.

[CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo and John P. Steinberger. *Random oracles and non-uniformity.* EUROCRYPT 2018, Part I, LNCS 10820, pp. 227–258. Springer, 2018.

[CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu and Luowen Qian. *Tight quantum time-space tradeoffs for function inversion.* arXiv:2006.05650, 2020.

[GLLZ21] Siyao Guo, Qian Li, Qipeng Liu and Jiapeng Zhang. *Unifying presampling via concentration bounds.* IACR Cryptology ePrint Archive, Report 2020/1589; TCC 2021, LNCS 13042, pp. 177–208. Springer, 2021.

[Liu23] Qipeng Liu. *Non-uniformity and quantum advice in the quantum random oracle model.* EUROCRYPT 2023, Part I, LNCS 14004, pp. 117–143. Springer, 2023; arXiv:2210.06693.

[Unr07] Dominique Unruh. *Random oracles and auxiliary input.* CRYPTO 2007, LNCS 4622, pp. 205–223. Springer, 2007.