Tight Quantum Time-Space Tradeoffs for Function Inversion

Kai-Min Chung*, Siyao Guo†, Qipeng Liu‡ and Luowen Qian§
*Institute of Information Science, Academia Sinica
Taipei, Taiwan
kmchung@iis.sinica.edu.tw
† Department of Computer Science, New York University Shanghai
Shanghai, China
siyao.guo@nyu.edu
‡ Department of Computer Science, Princeton University & NTT Research
Princeton, USA
qipengl@cs.princeton.edu
§Department of Computer Science, Boston University
Boston, USA
luowenq@bu.edu

***Abstract*—In function inversion, we are given a function $f : [N] \mapsto [N]$, and want to prepare some advice of size S, such that we can efficiently invert any image in time T. This is a well studied problem with profound connections to cryptography, data structures, communication complexity, and circuit lower bounds. Investigation of this problem in the quantum setting was initiated by Nayebi, Aaronson, Belovs, and Trevisan (2015), who proved a lower bound of $ST^2 = \tilde{\Omega}(N)$ for random permutations against classical advice, leaving open an intriguing possibility that Grover's search can be sped up to time $\tilde{O}(\sqrt{N/S})$. Recent works by Hhan, Xagawa, and Yamakawa (2019), and Chung, Liao, and Qian (2019) extended the argument for random functions and quantum advice, but the lower bound remains $ST^2 = \tilde{\Omega}(N)$.**

In this work, we prove that even with quantum advice, $ST + T^2 = \tilde{\Omega}(N)$ is required for an algorithm to invert random functions. This demonstrates that Grover's search is optimal for $S = \tilde{O}(\sqrt{N})$, ruling out any substantial speed-up for Grover's search even with quantum advice. Further improvements to our bounds would imply new classical circuit lower bounds, as shown by Corrigan-Gibbs and Kogan (2019).

To prove this result, we develop a general framework for establishing quantum time-space lower bounds. We further demonstrate the power of our framework by proving the following results.

- **Yao's box problem: We prove a tight quantum time-space lower bound for classical advice. For quantum advice, we prove a first time-space lower bound using shadow tomography. These results resolve two open problems posted by Nayebi et al (2015).**
- **Salted cryptography: We show that "salting generically provably defeats preprocessing," a result shown by Coretti, Dodis, Guo, and Steinberger (2018), also holds in the quantum setting. In particular, we prove quantum time-space lower bounds for a wide class of salted cryptographic primitives in the quantum random oracle model. This yields the first quantum time-space lower bound for salted collision-finding, which in turn implies that $\mathsf{PWPP}^{\mathcal{O}} \not\subseteq \mathsf{FBQP}^{\mathcal{O}}/\mathsf{qpoly}$ relative to a random oracle $\mathcal{O}$.**

***Keywords*-time-space tradeoffs; quantum computation; quantum query complexity; quantum advice; post-quantum cryptography; function inversion**

I. Function Inversion

The task of function inversion asks that given a function $f : [N] \mapsto [N]$ and a point y, find an x such that $f(x) = y$. It is easy to show that a classical inversion algorithm requires $\Omega(N)$ queries (which trivially lower bounds time) to succeed with constant probability in inverting a random function f, even with the help of randomness. Grover [Gro96] considered the same problem in the context of database search, and showed that quantum computers can invert any function in time only $\tilde{O}(\sqrt{N})$, which was subsequently shown to be tight [BBBV97].

The situation becomes intriguing when preprocessing is allowed. Namely, we allow the algorithm to take the entire truth table of f and arbitrarily preprocess an S-bit advice string $\alpha = \alpha(f)$, and then we give the algorithm a random $f(x)$ and ask the algorithm to invert it using at most T queries. Understanding the tradeoff between S and T is referred to as time-space tradeoffs for function inversion. This tradeoff is an important problem in cryptography. Recent works [GGH+19, KP19, CK19] showed its connections to well studied problems in data structures, communication complexity, and circuit lower bounds.

For classical algorithms, the heuristic algorithm proposed by Hellman [Hel80], and subsequently rigorously analyzed by Fiat and Naor [FN99], uses S bits of advice and T queries to invert a random function with high probability for every S, T satisfying $S^2T \geq \tilde{O}(N^2)$. For the lower bound, Yao [Yao90] and De et al. [DTT10] proved that any preprocessing algorithm that uses S bits of advice and T queries must satisfy $ST = \tilde{\Omega}(N)$, which remains the best general lower bound we know today. Corrigan-Gibbs and Kogan [CK19] recently investigated possible improvements

2575-8454/20/$31.00 ©2020 IEEE
DOI 10.1109/FOCS46700.2020.00068

on the lower bound, and showed that any improvements on Yao's lower bound will lead to improved circuit lower bounds.

The study of time-space tradeoffs in the quantum setting was initiated by Nayebi, Aaronson, Belovs and Trevisan [NABT15]. When the preprocessing algorithm is quantum, it is natural to distinguish the cases of quantum versus classical advice, as analogous to the complexity classes of BQP/qpoly versus BQP/poly and QMA versus QCMA. Nayebi et al. [NABT15] showed that any quantum preprocessing algorithm that uses S bits of *classical* advice and T queries to invert a random permutation with a constant probability must satisfy $ST^2 = \tilde{\Omega}(N)$. Note that the T^2 term is necessary given Grover's search algorithm. Recently, with motivations from post-quantum cryptography, Hhan, Xagawa, and Yamakawa [HXY19] extended the lower bound to handle general function inversion (and other cryptographic primitives). For the more challenging case of algorithms with *quantum* advice, Hhan el al. [HXY19] and Chung et al. [CLQ19] proved lower bounds for inverting random permutation and a restricted class of random functions, specifically, the lower bound only holds for functions with roughly the same domain and image size. However, these lower bounds remain $ST^2 = \tilde{\Omega}(N)$. As pointed out by Nayebi et al. [NABT15], this leaves open the following intriguing possibility:

> *Could a piece of preprocessed advice help speed up Grover's search algorithm?*

In this work, we prove the following quantum time-space lower bound for function inversion, which shows that even quantum advice of size $S = O(\sqrt{N})$ does not help speed up Grover's search algorithm.

Theorem 1. *Let $f : [N] \mapsto [N]$ be a random function. For any quantum oracle algorithm $\mathcal{A}$ with S-qubit oracle-dependent advice $\alpha = \alpha(f)$ and T queries to f,*

- *if α is classical (i.e. an S-bit string), then*

$$\Pr\left[\mathcal{A}^f(\alpha, f(x)) \in f^{-1}(f(x))\right] = \tilde{O}\left(\frac{ST + T^2}{N}\right);$$

- *if α is quantum (i.e. the general case), then*

$$\Pr\left[\mathcal{A}^f(\alpha, f(x)) \in f^{-1}(f(x))\right] = \tilde{O}\left(\sqrt[3]{\frac{ST + T^2}{N}}\right),$$

where both probabilities are over f, a uniformly random x from $[N]$, and randomness of $\mathcal{A}$.

Our lower bound implies that for a quantum preprocessing algorithm to invert a random function with a constant probability, it must satisfy $ST + T^2 = \tilde{\Omega}(N)$ even for the case of quantum advice. This further shows that Grover's search is optimal for $S = \tilde{O}(\sqrt{N})$, ruling out any substantial speed-up for Grover's search even with quantum advice. For $S = \tilde{\omega}(\sqrt{N})$, our lower bound matches Yao's lower bound for classical algorithms and is the best provable lower bound in light of the above-mentioned barrier results of [CK19].

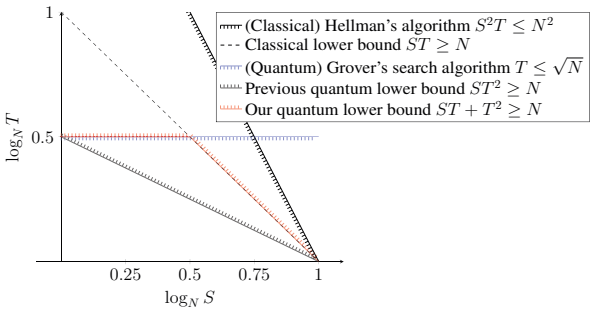

Figure 1. Time-space tradeoffs for inverting a random function $f : [N] \mapsto [N]$ with a constant successful probability. For the classical setting, the best upper bound is given by Hellman's algorithm [Hel80, FN99], and the best lower bound is given by [Yao90, DGK17, CDGS18]. For the quantum setting, the best upper bound is given either by Grover's search algorithm [Gro96] or (classical) Hellman's algorithm, and the previous best lower bound is given by [HXY19, CLQ19]. For simplicity, logarithmic terms and constant factors are omitted.

Furthermore, in the context of cryptography, a typical complexity measure to define the security is $S + T$ (corresponding to the program length and the running-time), with respect to which our lower bound implies a lower bound $S + T \ge \tilde{\Omega}(\sqrt{N})$. This matches Grover's search algorithm and gives a tight characterization for function inversion.

A comparison of our bounds with known upper and lower bounds is included in Figure 1. We remark that our lower bounds can be extended to general functions $f : [N] \mapsto [M]$.

II. Yao's Box Problem

Yao's box problem [Yao90] is another basic problem investigated in the literature of time-space tradeoffs that is closely related to the function inversion problem and non-uniform security of pseudorandom generators. In this problem, a preprocessing algorithm $\mathcal{A}$ can compute an S-bit advice for a function $f : [N] \mapsto \{0, 1\}$ in the preprocessing phase, and then in the online phase, it is required to compute the bit $f(x) \in \{0, 1\}$ for a random point $x \in [N]$ by making at most T queries but without querying f on x. Similar formulations of such problem in the classical setting are recently studied in the context of circuit complexity [ST18, MW19], which led to a new result on depth 3 circuits.

In the classical setting, Yao [Yao90] proved that for $\mathcal{A}$ to succeed with probability $2/3$, it must satisfy $ST = \Omega(N)$, which is known to be optimal. In the quantum setting, Nayebi et al. [NABT15] showed an $ST^2 = \Omega(N)$ lower bound for solving Yao's box problem with *classical* advice, which does not rule out the possibility that a preprocessing algorithm with S bits of advice and $T = O(\sqrt{N/S})$ queries. Subsequently, Hhan et al. [HXY19, Lemma 6] refined the analysis and showed that for any S, T, N, any quantum algorithm with classical advice can only succeed with probability

$1/2+\tilde{O}(ST^2/N)^{1/6}$. However, the following two problems posted by Nayebi et al. remained open.

Open Problem 2 ([NABT14, Section 3.4, Section 5])**.** *Is there a quantum algorithm that solves Yao's box in time $T = O(\sqrt{N/S})$? Or equivalently, prove or disprove the optimality of the lower bound $ST^2 = \tilde{\Omega}(N)$ for Yao's box.*

Open Problem 3 ([NABT14, Section 5])**.** *Extend the lower bound for Yao's box to the setting where the advice can be an arbitrary quantum state.*

We prove the following theorem for Yao's box problem, answering both open problems above.

Theorem 4. *Let $f : [N] \mapsto \{0,1\}$ be a random function. For any quantum oracle algorithms $\mathcal{A}$ with S-qubit oracle-dependent advice $\alpha = \alpha(f)$ and T queries to f except on the given challenge point x,*

- *if α is classical, then*

$$\Pr[\mathcal{A}^f(\alpha, x) = f(x)] \leq \frac{1}{2} + \tilde{O}\left(\sqrt[3]{\frac{ST}{N}}\right);$$

- *if α is quantum, then*

$$\Pr[\mathcal{A}^f(\alpha, x) = f(x)] \leq \frac{1}{2} + \tilde{O}\left(\sqrt[19]{\frac{S^5T}{N}}\right),$$

where both probabilities are over f, a uniformly random x from $[N]$, and randomness of $\mathcal{A}$.

In particular, this theorem implies that any algorithm achieving success probability $2/3$ has to satisfy $ST = \tilde{\Omega}(N)$ for classical advice, which shows that the power of quantum query does not help solving Yao's box problem, and $S^5T = \tilde{\Omega}(N)$ for quantum advice.

We note that while in the classical setting, time-space lower bounds for Yao's box problem and the function inversion problem can be proved using the same techniques, such as compression [DTT10, DGK17] or presampling [Unr07, CDGS18], proving quantum time-space lower bounds for Yao's box problem seems to be more challenging for quantum advice. This is also the case for our framework since in Yao's box problem, the algorithm cannot verify the answer on its own. For the case of quantum advice, we employ a novel use of online shadow tomography [AR19], which enables us to prove the first time-space lower bound for algorithms with quantum advice. We discuss the issue more carefully in Section IV-C.

Finally, we note that while we did not explicitly prove this, we believe that our techniques are sufficient for proving a lower bound $ST = \tilde{\Omega}(N)$ for quantum advice, if we restrict the algorithm such that it has to recover $f(x)$ for any $x \in [N]$ with probability $2/3$. We discuss this further in Remark 7.

III. Post-Quantum Non-Uniform Security

It turns out that our techniques for proving the two results above are fairly general, and we can use them to prove a variety of quantum non-uniform lower bounds. We now turn our focus to proving non-uniform lower bounds in cryptographic ideal models, which is a topic that has gained a lot of momentum recently.

Random oracle methodology and concrete security. While theoretically we can instantiate a lot of private-key cryptography assuming only the existence of any one-way function [Lev87], the constructions are almost always way too inefficient to be any useful for practical purposes. Practical cryptographic schemes, instead are usually designed under an ideal model and are proven secure under that model. A popular model of choice is random oracle model (ROM) [BR93], which "heuristically" models a function in question $f : [N] \mapsto [M]$, say a SHA3 hash function, as a truly random function that everyone has only oracle access to. To determine the security parameter to use in the construction, usually concrete security bounds are derived in the ideal model, and we work out the calculations to make sure any adversary in the ideal model, using at most a reasonable amount of resources, can only succeed with small enough probability, usually 2^{-32} [LR10, Arc20].

While the random oracle model "heuristically" captures all attacks that do not employ the structure of any particular instantiation of the random oracle, the model itself captures neither preprocessing attackers nor quantum attackers. Moreover, security bounds obtained in the random oracle model are inaccurate or do not apply at all once preprocessing or quantum computation are allowed.

Quantum random oracle model. Quantum algorithms are known to achieve various nontrivial speedups compared to classical algorithms, sometimes even an exponential speedup. For example, Grover's search algorithm achieves $\sqrt{N}$ speedup over classical algorithms, and Shor's algorithm solves factoring in quantum polynomial time.

Motivated by assessing the post-quantum security of constructions in ROM, Boneh et al. [BDF+11] introduced the quantum random oracle model (QROM) where the attacker can make superposition queries to the oracle, as given a classical description of any function, a quantum algorithm can trivially perform such superposition queries. Extending security proofs in classical ROM into the stronger QROM has been an active area of investigation [BDF+11, Zha12, BZ13, Unr15, TU16, Unr17, KLS18, Zha19, AHU19, LZ19b, DFMS19].

Auxiliary-input random oracle model. Consider the example where we use random oracle to instantiate an one-way function (OWF). As shown by Hellman's algorithm [Hel80], a preprocessing attacker can achieve a non-trivial saving of resources from N to $N^{2/3}$ even in the random oracle model. Because the instantiated function is usually public, there is

no way of preventing the adversary from performing a heavy precomputation to speed up the online attack.

To address this mismatch, Unruh [Unr07] introduced the auxiliary-input random oracle model (AI-ROM) where the adversary is allowed to obtain a bounded length advice about the random oracle before attacking the system. Several works [Unr07, DGK17, CDGS18, CDG18] have developed various techniques for analyzing the security in this model.

Hhan et al. [HXY19] first considered the auxiliary-input quantum random oracle model (AI-QROM) and quantum auxiliary-input quantum random oracle model (QAI-QROM), which is the natural generalization of the model above for quantum adversaries. Despite the progress in QROM and AI-ROM respectively, proving post-quantum non-uniform security remains quite challenging. Only a few security bounds have been proven against classical advice, i.e. under AI-QROM [HXY19], and *no* security bounds are known against quantum advice, i.e. QAI-QROM, except for OWFs [HXY19, CLQ19]. The only known technique under these models is the (quantum) compression technique, introduced by Nayebi et al. [NABT15], and its variants. We note that this technique is only somewhat generic, and its effectiveness seems limited under QAI-QROM.

New generic framework for security bounds in (Q)AI-QROM. We generalize our techniques for function inversion and Yao's box into a general framework for proving concrete post-quantum non-uniform security in AI-QROM and QAI-QROM. In particular, we show the following theorem.

Theorem (Informal)**.** *For any security game G, adversary time bound T, and any $g > 0$, consider its multi-instance game $G^{\otimes g}$, which requires the adversary to break g independent challenges* sequentially*, and each instance is given time T.*

If the best winning probability for $G^{\otimes g}$ is δ^g in the QROM, meaning that the best adversary can only win the game with probability at most δ^g, then for adversaries with $S = g$ (qu)bits of advice:

1) *G is roughly $\tilde{O}(\delta)$-secure in AI-QROM;*
2) *If G is publicly verifiable, G is $\mathsf{negl}(N)$-secure in QAI-QROM if $\delta = \mathsf{negl}(N)$;*
3) *If G is a decision game, G is $1/2 + \mathsf{negl}(N)$-secure in QAI-QROM if $\delta = 1/2 + \mathsf{negl}(N)$.*

We note that our bound for function inversion we show in Section I directly translates to the concrete security bound for OWFs, and our techniques for OWFs can be additionally used to show security for pseudorandom generators (PRGs) with little efforts. The concrete bounds, are summarized into the table below.

	OWFs	PRGs
AI-ROM	$\frac{ST}{\alpha}$	$\left(\frac{ST}{N}\right)^{1/2} + \frac{T}{N}$
AI-QROM [HXY19]	$\left(\frac{ST^2}{\alpha}\right)^{1/2}$	$\left(\frac{ST^4}{N} + \frac{T^4}{N}\right)^{1/6}$
AI-QROM (ours)	$\frac{ST}{\alpha} + \frac{T^2}{\alpha}$	$\left(\frac{ST}{N} + \frac{T^2}{N}\right)^{1/3}$
QAI-QROM (ours)	$\left(\frac{ST}{\alpha} + \frac{T^2}{\alpha}\right)^{1/3}$	$\left(\frac{S^5T}{N} + \frac{S^4T^2}{N}\right)^{1/19}$

Table I
ASYMPTOTIC SECURITY BOUNDS ON THE SECURITY OF OWFS AND PRGS CONSTRUCTED FROM A RANDOM FUNCTION $f : [N] \mapsto [M]$ AGAINST (S, T)-ALGORITHMS, WHERE $\alpha := \min\{N, M\}$.

1) Salting defeats preprocessing in the quantum setting. Instead of proving more concrete security bounds using the framework, we show that salting, a common mechanism used in practice for defeating auxiliary input, generically extends the security of applications proven in the QROM to the (Q)AI-QROM. A similar statement was first shown to hold in the classical world by Coretti et al. [CDGS18], where they showed that the security in the ROM can be generically extended into salted security in the AI-ROM.

In this work, we prove the hardness of salted multi-instance game under QROM.

Lemma (Informal)**.** *For any security game G with security δ in the QROM against adversary running in time T, let G_S denote its salted version with salt space $[K]$. Then the best winning probability for multi-instance salted game $G_S^{\otimes g}$ is at most $(\delta + gT/K)^g$, which is tight.*

Combining this lemma with our reduction theorem above, we conclude that salting generically defeats preprocessing in the AI-QROM, and defeats preprocessing against publicly-verifiable and decision games in the QAI-QROM. We believe our techniques are sufficient and both the reduction theorem and salted multi-instance bound can be generalized for proving quantum non-uniform salted security bounds under other idealized cryptographic models, e.g. the ideal-cipher model (which is by definition a salted permutation family), and the salted generic group model. However, in order to simplify presentation, in this work we only focus on the ROM.

Using this generic bootstrapping theorem, we can also easily obtain security bounds for a wide class of salted cryptographic primitives in the AI-QROM and QAI-QROM. In particular, we use it to give the first bound for salted collision-resistant hash (CRH) in AI-QROM and QAI-QROM, resolving an open problem raised by Hhan et al. [HXY19]. This in turn implies that $\mathsf{PWPP}^{\mathcal{O}} \not\subseteq \mathsf{FBQP}^{\mathcal{O}}/\mathsf{qpoly}$ relative to a random oracle $\mathcal{O}$.

Finally, we note that our proof can also be naturally downgraded to a classical reduction, which gives a new proof of the classical result first proven by Coretti et al. [CDGS18]

IV. Technical Overview

In this section, we give an overview of our techniques for proving the results. Please refer to the full version for the formal proof.

A. From Non-uniform Algorithms to Multi-Instance Games

We start by considering applying a general approach in the literature for proving lower bounds for non-uniform algorithms in the context of classical function inversion, and then generalize it to quantum non-uniformity. In particular, the following argument is based on the work of Aaronson [Aar05].

Let $f : [N] \mapsto [N]$ be a function. Consider a classical (randomized) algorithm $\mathcal{A}$ with S-bit of advice $\alpha = \alpha(f)$, T queries to f (which is a lower bound on its running time), that can invert a random image for *any* function[1] f with probability at least δ, i.e.

$$\delta := \Pr_{\mathcal{A},x}[\mathcal{A}^f(\alpha, f(x)) \in f^{-1}(f(x))],$$

where the probability is taken over the measurement randomness of the algorithm $\mathcal{A}$ and the random choice of x, and we want to give an upper bound on δ. At a high level, the approach is to reduce it to proving lower bound for the multi-instance version of the problem for algorithms *without* advice. Specifically, we reduce it to bound the success probability of an algorithm $\mathcal{B}$ with gT queries to invert random g inputs $f(x_1), \ldots, f(x_g)$ simultaneously, for some parameter $g \in [N]$ on the number of instances. For function inversion, it is easy to show that for any $g > 0$, the best success probability drops exponentially fast in g. Specifically, for any (randomized) algorithm $\mathcal{B}$ and random functions f,

$$\Pr_{\mathcal{B},f,x_1,\ldots,x_g}[\mathcal{B}^f(f(x_1),\ldots,f(x_g)) \text{ inverts } f(x_1),\ldots,f(x_g)] \leq O(gT/N)^g. \quad (1)$$

The reduction proceeds in two simple steps. We first use $\mathcal{A}$ to construct an algorithm $\mathcal{B}'$ using *only one copy of the original advice* for the multi-instance problem (with decent success probability), and then get rid of the advice by simply guessing a uniformly random bitstring. The algorithm $\mathcal{B}'^f(\alpha, f(x_1), \ldots, f(x_g))$ simply invokes $\mathcal{A}$ to invert each $f(x_i)$, and succeeds when $\mathcal{A}^f(f(x_i))$ succeeds on all g instances. By independence, it is easy to see that

$$\Pr_{\mathcal{B}',x_1,\ldots,x_g}[\mathcal{B}'^f(\alpha, f(x_1),\ldots,f(x_g)) \text{ inverts } f(x_1),\ldots,f(x_g)] \geq \delta^g.$$

Next we remove the advice by guessing. Consider an algorithm $\mathcal{B}$ first guesses a *random* advice $\alpha \in \{0,1\}^S$ and then runs $\mathcal{B}'$. Clearly, $\mathcal{B}$ guesses the advice correctly with probability 2^{-S}, in which case $\mathcal{B}$ simulates $\mathcal{B}'$ perfectly. Hence,

$$\Pr_{\mathcal{B},x_1,\ldots,x_g}[\mathcal{B}^f(f(x_1),\ldots,f(x_g)) \text{ inverts } f(x_1),\ldots,f(x_g)] \geq 2^{-S}\delta^g.$$

As this statement is for any function, the same conclusion holds for random functions. Combining this with the above upper bound (1) on the success probability of $\mathcal{B}$ with $g = S$ shows that $\delta \leq O(ST/N)$, which matches the best known classical bound [Yao90, DGK17, CDGS18].

It should be clear from the above example that this approach is fairly general and reduces the non-uniform lower bound problems to analyze the success probability of the corresponding multi-instance games. Indeed, this approach and its variants have implicitly appeared in various contexts under the name of direct product theorems. We discuss this in more details in Section V.

First attempt using multi-instance problems. Seeing this as a promising start, we now consider the setting of quantum algorithms with *classical* advice. Indeed, the argument works out similarly, except that now we need to consider the success probability of the best *quantum* algorithm that queries gT locations and succeed inverting g independently random images. We could hope that analyzing the best success probability for quantum algorithms solving the multi-instance problem would lead us to the desired bound.

Unfortunately, it turns out that this approach is destined to fail to achieve any $\delta \ll ST^2/N$. With q quantum queries, the famous Grover's search can find one element in η fraction of all elements with probability $\approx \eta q^2$. Consider the following algorithm $\tilde{\mathcal{B}}$, where it tries to find one pre-image among all the g images using the first T queries, which succeeds with probability gT^2/N. If it succeeds in finding the first pre-image, it then use the next T queries to find one pre-image among all remaining $(g-1)$ images, which succeeds with probability $(g-1)T^2/N$ and so on. Using this algorithm, for any function, we can find all g pre-images with probability at least roughly

$$(gT^2/N) \cdot ((g-1)T^2/N) \cdot \ldots \cdot (T^2/N) \approx (T^2/N)^g \cdot g! \approx (gT^2/N)^g.$$

This implies that the best bound we can hope to achieve for the multi-instance problem would be $O(gT^2/N)^g$, which would in turn imply the bound $O(ST^2/N)$, which is the same bound as what we already had before.

Bypassing the barrier via multi-instance games. A natural question arises that whether we can go beyond ST^2/N. We claim that this is actually the case.

To see this point, we first recall the high level ideas of the argument above – we first bootstrap the best algorithm $\mathcal{A}$ with advice α for computing f^{-1} with success probability δ, into a multi-instance algorithm $\mathcal{B}$ with advice α with success

[1] Since we are requiring the algorithm to invert any function, the lower bound we present here is slightly weaker. We intentionally make this omission for the overview to highlight the more important ideas in our proofs. Interested readers should refer to the formal proofs.

probability δ^g, and then remove the advice and incur a loss of 2^{-S}, which we amortize into δ^g. The problem essentially reduces to proving a lower bound for the success probability of the resulting algorithm $\mathcal{B}$ that solves the multi-instance problem.

While it may seem like that the $(gT^2/N)^g$ algorithm above could be an upper bound for the new problem, we observe that this iterated Grover's search algorithm $\tilde{\mathcal{B}}$ actually never arises from our reduction from $\mathcal{A}$ to $\mathcal{B}$. In particular, $\mathcal{B}$ always solves each instance one by one, while $\tilde{\mathcal{B}}$ in some sense solves *all* g instances at once. To view this issue in terms of quantum queries, the first T queries by $\mathcal{B}$ are only searching pre-image of $f(x_1)$, but the first T queries by $\tilde{\mathcal{B}}$ are searching for all g pre-images.

We formalize this intuition by strengthening the multi-instance problem into what we call a "multi-instance game," where the algorithm (or adversary) is instead *interacting* with a verifier (or challenger). For each round $i \in [g]$, the challenger samples a new image $f(x_i)$, and the adversary is given T queries to f, before producing an output x'_i, which the challenger checks whether $f(x_i) = f(x'_i)$. The main change we make to the multi-instance *games*, compared with multi-instance *problems*, is that the adversary gets challenges one-by-one or *sequentially*, rather than getting all challenges at once or in *parallel*.

Observe that this multi-instance game seems to rule out the algorithm $\tilde{\mathcal{B}}$ above, as the adversary does not have any information about $f(x_i)$ until he has issued $(i-1)T$ queries.

If we assume that the probability that any quantum adversary wins such multi-instance game for function inversion is $O\left(\frac{gT+T^2}{N}\right)^g$, using the same reduction as before, we would reach the conclusion that the best success probability for function inversion with S bits of classical advice and T quantum queries is $O(ST + T^2)/N$. It turns out that we can indeed prove this assumption, but we will defer the discussion and consider quantum advice first.

Beyond classical preprocessing. The reduction above requires the algorithm to solve multiple instances using only a single copy of the advice, which is problematic in the quantum setting due to no-cloning theorem. We resolve this problem by constructing $\mathcal{B}$ similarly as before, and add gentle measurement to solve multiple instances. To fill in the details, $\mathcal{B}$ does the following.

1) Prepare $k = \Theta(\log g)$ copies of the quantum advice $\beta := \alpha^{\otimes k}$.
2) Boost $\mathcal{A}$'s success probability from constant to $1 - o(1)$, by running $\mathcal{A}$ on each copy of the advice α k times, and identify the correct answer using one additional query.
3) To solve g instances simultaneously, $\mathcal{B}$ simply runs boosted $\mathcal{A}$ for each instance, and applies measurement. As we have boosted the success probability high enough, the measurement will be "gentle" and we can recover an almost-as-good-as-new quantum advice for the next instance.
4) Finally, to remove the quantum advice, we replace β with a maximally mixed state, which gives us a multiplicative loss of 2^{-Sk} in success probability.

However, for function inversion, this idea seems to fail as function inversion problem can have non-unique correct answers. In particular, if $f^{-1}(f(x)) = \{x_1, x_2, ..., x_\ell\}$, and the algorithm somehow prepares the answer $|x'_i\rangle$ that is a uniform superposition over all the answers $x_1, ..., x_\ell$, while it succeeds with probability 1, performing a gentle measurement on this answer seems very difficult. This difficulty of performing gentle measurements for function inversion was also acknowledged by the work of Hhan et al. [HXY19], although under a different context.

We claim that using multi-instance *games* (as opposed to multi-instance problems), there is actually a very elegant solution to this. Instead of having the adversary submitting a classical answer x'_i, we will allow the adversary to submit a *quantum state* $|x'_i\rangle$, and the challenger can compute in superposition whether $f(|x'_i\rangle) = f(x_i)$, and measure her decision (which will be gentle, as we boosted its success probability), and send back $|x'_i\rangle$. On a high level, the idea is basically having the adversary and the challenger "jointly" perform this gentle measurement. As the adversary cannot control challenger's behavior, this change should not impact the best winning probability of the multi-instance game.

B. Analyzing Multi-Instance Game via "Compressed Oracles"

To complete our time-space tradeoff for function inversion, the only remaining step is to bound the best winning probability of the multi-instance game for function inversion. We extends the techniques of Zhandry's compressed oracles [Zha19] and combine with a new indistinguishability lemma to give a tight bound for this problem.

Compressed oracles. In the classical setting, there is a commonly used technique for arguing random functions called the lazy sampling of a random oracle. The idea is that a simulator will maintain a partial truth table about the random function. Upon an oracle query x, the simulator looks up x in the table D and returns it as the answer. If not found, the simulator freshly samples a new y as the output of x and inserts the pair (x, y) into the table D.

Zhandry observes that, if care is taken to implement the oracle correctly, a quantum analogy of the classical on-the-fly simulation is possible. Unlike the classical simulation, they simulate a random oracle as a superposition of tables, each of which partially instantiates a random function. Below is the high level idea of their simulation. The table is initialized as an empty table. Upon a quantum query is made by an algorithm, the simulator updates the database in superposition: for a query $|x\rangle$ and a table $|D\rangle$, the simulator look up x in the table D; if not found, it

initializes a superposition of all possible output $\sum_y |y\rangle$ (up to a normalization) as the value of $D(x)$ and updates D in superposition to get $|D \cup (x, y)\rangle$; it then returns $D(x)$ as the output.

A major difference between quantum setting and classical setting is an algorithm may forget some query it made before. As an example, an algorithm can query the same input twice to un-compute everything and thus completely lose the information about the output. Therefore, to perfectly simulate a quantum random oracle, the simulator also checks after every query that if the algorithm loses all information about the query. With all these operations above, Zhandry shows a quantum random oracle can be efficiently simulated on-the-fly.

Analyzing multi-instance game. To prove the success probability of multi-instance function inversion, we consider a stronger statement regarding the success probability of inverting for any round: assume an algorithm already makes $(i-1)T$ queries for the first $(i-1)$ rounds, and conditioned on it having passed the first $(i-1)$ rounds, what is the probability of succeeding in the i-th round by making T queries?

Lemma 5. *For any quantum algorithm making $q_0 + q$ queries to a random function $f : [N] \to [N]$, if $f(x)$ is sampled and given after the q_0-th query, conditioned on arbitrary outcomes (with non-zero probability) of the algorithm's measurement during the first q_0 queries, the probability of inverting $f(x)$ is at most $O((q_0 + q^2)/N)$.*

We consider the lemma above as remarkable and perhaps even surprising, as intuitively, it is saying that quantum power can achieve a quadratic speed up for search *only* if you know what you are looking for, and there is no classical analogue of this.

With the above lemma, the probability of succeeding in the i-th round is at most $O((iT + T^2)/N)$. Therefore, the probability of succeeding in inverting all random images is at most $O((gT + T^2)/N)^g$.

To prove this lemma, let us start by assuming that a uniformly random image y is instead given as a challenge and without conditioning on the intermediate measurements. By Zhandry's techniques, after making q_0 queries to a random oracle, the knowledge of an algorithm about the random oracle can be viewed as a superposition of some tables with at most q_0 entries which specifies partial random functions over at most q_0 inputs. Since y is sampled uniformly, the amplitude (square root of probability) of database containing y is about $\sqrt{q_0/N}$. After given the image y, each query can increase the amplitude by at most $\sqrt{1/N}$. Therefore, the final amplitude of of database containing y is about $(\sqrt{q_0} + q)/\sqrt{N}$ which gives us the lemma above.

There are two challenges we need to overcome for proving the full lemma.

The first challenge comes from the fact that our lemma statement requires an algorithm that is conditioned on some fixed measurement outcomes. To address this, we extend Zhandry's techniques to such settings, which is a very natural but non-trivial extension and crucial for analyzing the multi-instance game.

The second challenge seems even more difficult. For function inversion, the image is sampled by first sampling a random pre-image x and computing $f(x)$. The natural application of compressed oracle only gives us a probability bound when a random y is sampled instead of $f(x)$. This is a nontrivial issue as with high probability, the distribution of $f(x)$ is only supported on a constant fraction of $[N]$, so the statistical difference of the two distributions is significant. The issue is more prominent when we consider the general random functions $f : [N] \mapsto [M]$ where $N \ll M$, where a uniform sample of y is with high probability not an image of any x.

Towards this challenge, we prove the following indistinguishability lemma to bridge the gap. The proof of the lemma does not require the compressed oracle technique and we believe that both lemmas are of independent interest.

Lemma 6 (Indistinguishability)**.** *For any quantum algorithm making q_0+q queries to a random function $f : [N] \to [N]$, if $f(x)$ or a uniformly random y is sampled and given after the q_0-th query, conditioned on arbitrary outcomes (with non-zero probability) of the algorithm's measurement during the first q_0 queries, the advantage of distinguishing is at most $O((q_0 + q^2)/N)$.*

We can then assume a random y is sampled instead of a random $f(x)$ with only an additive loss of $O((gT+T^2)/N)$ in each round, giving us the bound that we desire.

Proving the indistinguishability lemma. We first convert the problem from distinguishing samples (random $f(x)$ or random y) into distinguishing oracles. Assume the oracle is sampled as follows: first, a uniformly random input x is sampled; then a random function f_{-x} defined on all inputs except x is sampled, together with two independently sampled y_0, y_1; define $f_{-x}||y$ as a function that outputs $f_{-x}(x')$ on all inputs that are not x and y on input x; the distinguisher is ask to given either oracle access to $f_{-x}||y_0$ or $f_{-x}||y_1$ and the same challenge y_0 after the q_0-th query, distinguish which oracle is given (without knowing x). When the function $f_{-x}||y_0$ and challenge y_0 is given, it corresponds to the case a random $f(x)$ is given; while giving the second function corresponds to the second case that a random y is given. It can be shown that the two problems have the exact same difficulty.

Intuitively, every quantum query to a function entangles a quantum algorithm with one of the output of that function in superposition. By making q_0 queries, the algorithm is entangled with at most q_0 outputs of the function in superposition. When y_0 is given, the only way to tell which oracle is given is by already making a query on x and

entangling the algorithm with either y_0 or y_1 respectively. Since y_0 has not been given during the first q_0 queries, x is perfectly hidden and completely uniformly random from the algorithm's view. Thus, such entanglement only happens with probability q_0/N. For the remaining q queries, knowing the information y_0 does help build the entanglement faster. One strategy is to use Grover's search to check if y_0 is an image of the function since with constant probability, y_0 is not an image of $f_{-x}||y_1$. By using Grover's search, the advantage of distinguish is about q^2/N and we show such advantage is the best one can get for these q queries. Combining with these two separate analysis, we conclude the indistinguishability lemma.

C. Yao's Box Problem

We now focus our attention on Yao's box problem. Assume that an algorithm $\mathcal{A}$, given any function $f : [N] \mapsto \{0, 1\}$, prepares an S-qubit advice $\alpha = \alpha(f)$ such that $\mathcal{A}$ can recover $f(x)$ for a random x using α without ever querying $f(x)$ in time T with probability $1/2 + \varepsilon$.

We claim that for classical advice, our reduction from function inversion with advice to multi-instance game can also be generalized to Yao's box, with a more careful amortizing analysis; and the multi-instance game for Yao's box is fairly straightforward to argue using similar techniques as for function inversion. Intuitively, for Yao's box, the only non-trivial strategy is that the adversary's first $(i-1)T$ quantum queries predicted the challenge x_i, so the best advantage of any algorithm, i.e. the best winning probability minus $1/2$, can only be $O(\sqrt{gT/N})$ instead of $O(\sqrt{gT^2/N})$. This leads us to the final bound $1/2 + \tilde{O}(ST/N)^{1/3}$, where the additional exponent loss comes from the new amortizing argument.

For quantum advice, things are a lot trickier. While Aaronson [Aar05] proved a similar lower bound for a different problem against quantum advice, the techniques there only allow us to prove lower bounds against algorithms that find the correct answer with probability at least $2/3$ for *all* f and x (and indeed under this setting, our techniques combined with [Aar05] are sufficient to give query lower bound $ST \geq \tilde{\Omega}(N)$ even for quantum advice). However, this is insufficient in our settings where we need to consider the stronger lower bound where x is sampled randomly, and the algorithm can only predict some $f(x)$ and output a random guess for others.

We first revisit the idea of Aaronson [Aar05], which is to prepare $\mathsf{poly}(\log g)$ copies of the advice, and use majority vote[2] to boost the success probability from $2/3$ to $1-o(1)$ to make the measurement gentle. We show that majority vote cannot possibly boost success probability under the average-case (over x) setting, by considering the following example: the algorithm $\mathcal{A}$ outputs the correct answer with probability 1 on 40% of the inputs, and with probability 0.45 on the other 60% of the inputs. Overall, the success probability is 67%, but with majority vote, the success probability goes down, and can go down arbitrarily close to 40%, which is much worse than random guessing!

One way to resolve this is to instead "gently" measure the success probability of $\mathcal{A}$ for each instance $f(x)$, and throw a biased random coin according to this distribution as our answer. We observe that the problem of "gently" measuring this probability can be reduced to shadow tomography, which is a problem introduced by Aaronson [Aar18]. As our multi-instance game requires that each challenge is given sequentially, we also require the shadow tomography to be able to handle online queries. Aaronson and Rothblum [AR19] showed that online shadow tomography indeed can be done using $S^2 \log^2 g$ copies of the advice, which leads us to the final bound $S^5T = \tilde{\Omega}(N)$ for constant $\varepsilon > 0$.

Remark 7 (Improving the lower bound for Yao's box with quantum advice)**.** *We note that if we make the restriction so that for most functions f, $\mathcal{A}$ succeeds to output $f(x)$ for* every *input x with probability at least $2/3$, then the idea of using majority vote still works, since in this case, boosting will give us the correct outcome with overwhelming probability. Using the same reduction, we can prove that $\mathcal{A}$ has to satisfy $ST = \tilde{\Omega}(N)$. Therefore, we think that $ST = \tilde{\Omega}(N)$ should be the optimal lower bound.*

V. Related Works

In this section, we compare our techniques with other related works.

The approach of reducing time-space tradeoff lower bounds to multi-instance *problems*, which is outlined in Section IV-A, has appeared implicitly in various works [Bea91, Kla03, KŠW07], where they usually refer to the (exponential) hardness of multi-instance problems as "(strong) direct product theorems." While the different approaches presented in different works share some similar high-level ideas, the context and details in each work are slightly different. In this work, to avoid confusion, we use the term "multi-instance problem" instead of direct products. Recently Hamoudi and Magniez [HM20] independently applied the similar technique along with Zhandry's compressed oracles to prove quantum time-space tradeoffs for finding multiple collisions.

Classically, this approach has also been considered in various works [GKL93, Imp11] for proving non-uniform lower bounds. Aaronson [Aar05] first showed how to employ such ideas when the non-uniform lower bounds need to hold even against quantum advice. While the problem they consider is quite simple and somewhat arbitrary, the starting point we outline in Section IV-A is based on this work. However, to the best of our knowledge this technique has not been explored in the AI-ROM literature (possibly due to the

[2]This idea was implicitly given, where they called it "boosting" under the context of randomized algorithms in complexity theory.

fact that same bounds can already be achieved using other – possibly more complicated – techniques), and the reductions presented in this work can also be easily "dequantized" to the classical setting.

As far as we are concerned, our work is the first one to consider the stronger variant "multi-instance games" and show a separation of the two variants for function inversion in Section IV-A. Additionally, we present our reduction under a general framework for the stronger variant of multi-instance games.

The idea of using gentle measurements is almost ubiquitous for proving lower bounds against quantum advice. To the best of our knowledge, Aaronson [Aar05] first showed how to combine boosting and gentle measurements for quantum advice lower bounds, which we briefly discuss in Section IV-C. In particular, this technique was also employed by Hhan et al. [HXY19] to prove an asymptotic lower bound of $ST^2 \geq N$ for inverting random permutations, although under a different context.

VI. Open Problems

Quantum time-space tradeoff lower bounds for permutation inversion. While our work provides substantial evidence that the quantum time-space tradeoff bound for inverting permutations is $ST + T^2 = \tilde{\Theta}(N)$, which would have resolved the open problem posted by Nayebi et al. [NABT15, Section 4.3], we are not able to formally prove this due to a lack of "compressed permutation oracles". This is especially interesting, considering that it is easier to argue about permutations than functions using the compression argument, which is used in prior works [NABT15, HXY19, CLQ19].

Nontrivial quantum speed ups for function inversion. While our lower bound is best possible, it still leaves open the possibility that some nontrivial quantum speed ups exist under the following asymptotic regime:

$$\begin{cases} ST \gg N, \\ T^2 \ll N, \\ S^2T \ll N^2. \end{cases}$$

An especially interesting case is that there might be a quantum algorithm with advice under this regime, but it might seem extremely hard to be "dequantized."

Our paper also gives many lower bounds, improving any of the bounds that are not tight (for example, the query bound for average-case Yao's box against quantum advice), or showing new non-trivial attacks, are all interesting possibilities.

Suboptimal exponent on success probability for quantum advice, or is it? All the bounds we have achieved for quantum advice are not tight in terms of the exponent on the success probability for quantum advice. In particular, we note that for function inversion, our classical advice lower bound has exponent 1, which is tight – while the quantum advice bound has exponent $1/3$. Can this loss be avoided, or is there any speed up in terms of S and T for sub-constant success probability?

We make the observation that the loss in exponent ultimately comes from the use of gentle measurements. Looking back at the literature, all the quantum advice lower bound techniques [Aar05, HXY19, CLQ19] we have seen so far always require "reusing" of the advice, which in turns require gentle measurements. Is there a way to avoid reusing the quantum advice to escape this cost? We suspect that this is the case, as in the work of Chung et al. [CLQ19], they presented a reduction to quantum random access code, which by definition seems to avoid this issue, albeit in the end, the advice reusing issue somehow kicks back in.

Acknowledgements

Kai-Min Chung is partially supported by the Academia Sinica Career Development Award under Grant no. 23-17, and MOST QC project under Grant no. MOST 108-2627-E-002-001-.

Siyao Guo is supported by Shanghai Eastern Young Scholar Program.

Qipeng Liu is supported in part by NSF and DARPA. Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF or DARPA.

Luowen Qian is supported by DARPA under Agreement No. HR00112020023.

References

[Aar05] Scott Aaronson. "Limitations of Quantum Advice and One-Way Communication". In: *Theory of Computing* 1.1 (2005), pp. 1–28. DOI: 10.4086/toc.2005.v001a001. URL: https://doi.org/10.4086/toc.2005.v001a001.

[Aar18] Scott Aaronson. "Shadow tomography of quantum states". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 325–338.

[AHU19] Andris Ambainis, Mike Hamburg, and Dominique Unruh. "Quantum security proofs using semi-classical oracles". In: *Annual International Cryptology Conference*. Springer. 2019, pp. 269–295.

[AR19] Scott Aaronson and Guy N. Rothblum. "Gentle measurement of quantum states and differential privacy". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*. Ed. by Moses Charikar and Edith Cohen. ACM, 2019, pp. 322–333. DOI: 10.1145/3313276.3316378. URL: https://doi.org/10.1145/3313276.3316378.

[Arc20] Scott Arciszewski. *XChaCha: eXtended-nonce ChaCha and AEAD_XChaCha20_Poly1305*. Internet-Draft draft-irtf-cfrg-xchacha-03. Work in Progress. Internet Engineering Task Force, Jan. 2020. 18 pp. URL: https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-xchacha-03.

[AS04] Scott Aaronson and Yaoyun Shi. "Quantum lower bounds for the collision and the element distinctness problems". In: *Journal of the ACM (JACM)* 51.4 (2004), pp. 595–605.

[BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. "Strengths and weaknesses of quantum computing". In: *SIAM journal on Computing* 26.5 (1997), pp. 1510–1523.

[BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. "Random Oracles in a Quantum World". In: *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Vol. 7073. Lecture Notes in Computer Science. Springer, 2011, pp. 41–69. DOI: 10.1007/978-3-642-25385-0_3. URL: https://doi.org/10.1007/978-3-642-25385-0%5C_3.

[Bea91] Paul Beame. "A General Sequential Time-Space Tradeoff for Finding Unique Elements". In: *SIAM J. Comput.* 20.2 (1991), pp. 270–277. DOI: 10.1137/0220017. URL: https://doi.org/10.1137/0220017.

[BR93] Mihir Bellare and Phillip Rogaway. "Random Oracles are Practical: A Paradigm for Designing Efficient Protocols". In: *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993*. 1993, pp. 62–73. DOI: 10.1145/168588.168596. URL: http://doi.acm.org/10.1145/168588.168596.

[BV97] Ethan Bernstein and Umesh V. Vazirani. "Quantum Complexity Theory". In: *SIAM J. Comput.* 26.5 (1997), pp. 1411–1473. DOI: 10.1137/S0097539796300921. URL: https://doi.org/10.1137/S0097539796300921.

[BZ13] Dan Boneh and Mark Zhandry. "Secure Signatures and Chosen Ciphertext Security in a Quantum Computing World". In: *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*. Ed. by Ran Canetti and Juan A. Garay. Vol. 8043. Lecture Notes in Computer Science. Springer, 2013, pp. 361–379. DOI: 10.1007/978-3-642-40084-1_21. URL: https://doi.org/10.1007/978-3-642-40084-1%5C_21.

[CDG18] Sandro Coretti, Yevgeniy Dodis, and Siyao Guo. "Non-Uniform Bounds in the Random-Permutation, Ideal-Cipher, and Generic-Group Models". In: *Annual International Cryptology Conference*. Springer. 2018, pp. 693–721.

[CDGS18] Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. "Random Oracles and Non-uniformity". In: *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10820. Lecture Notes in Computer Science. Springer, 2018, pp. 227–258. DOI: 10.1007/978-3-319-78381-9_9. URL: https://doi.org/10.1007/978-3-319-78381-9%5C_9.

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

[CLQ19] Kai-Min Chung, Tai-Ning Liao, and Luowen Qian. "Lower Bounds for Function Inversion with Quantum Advice". In: *arXiv preprint arXiv:1911.09176* (2019).

[DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. "Security of the fiat-shamir transformation in the quantum random-oracle model". In: *arXiv preprint arXiv:1902.07556* (2019).

[DGK17] Yevgeniy Dodis, Siyao Guo, and Jonathan Katz. "Fixing Cracks in the Concrete: Random Oracles with Auxiliary Input, Revisited". In: *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10211. Lecture Notes in Computer Science. 2017, pp. 473–495. DOI: 10.1007/978-3-319-56614-6_16. URL: https://doi.org/10.1007/978-3-319-56614-6%5C_16.

[DTT10] Anindya De, Luca Trevisan, and Madhur Tulsiani. "Time Space Tradeoffs for Attacks against One-Way Functions and PRGs". In: *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*. 2010, pp. 649–665. DOI: 10.1007/978-3-

642-14623-7_35. URL: http://dx.doi.org/10.1007/978-3-642-14623-7_35.

[FN99] Amos Fiat and Moni Naor. "Rigorous Time/Space Trade-offs for Inverting Functions". In: *SIAM J. Comput.* 29.3 (1999), pp. 790–803. DOI: 10.1137/S0097539795280512. URL: http://dx.doi.org/10.1137/S0097539795280512.

[GGH+19] Alexander Golovnev, Siyao Guo, Thibaut Horel, Sunoo Park, and Vinod Vaikuntanathan. "3SUM with Preprocessing: Algorithms, Lower Bounds and Cryptographic Applications". In: *CoRR* abs/1907.08355 (2019). arXiv: 1907.08355. URL: http://arxiv.org/abs/1907.08355.

[GKL93] Oded Goldreich, Hugo Krawczyk, and Michael Luby. "On the existence of pseudorandom generators". In: *SIAM Journal on Computing* 22.6 (1993), pp. 1163–1175.

[Gro96] Lov K Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. 1996, pp. 212–219.

[Hel80] Martin Hellman. "A cryptanalytic time-memory trade-off". In: *IEEE transactions on Information Theory* 26.4 (1980), pp. 401–406.

[HM20] Yassine Hamoudi and Frédéric Magniez. "Quantum Time-Space Tradeoffs by Recording Queries". In: *arXiv preprint arXiv:2002.08944* (2020).

[HXY19] Minki Hhan, Keita Xagawa, and Takashi Yamakawa. "Quantum Random Oracle Model with Auxiliary Input". In: *AsiaCrypt*. Springer. 2019.

[Imp11] Russell Impagliazzo. "Relativized separations of worst-case and average-case complexities for NP". In: *2011 IEEE 26th Annual Conference on Computational Complexity*. IEEE. 2011, pp. 104–114.

[Kla03] Hartmut Klauck. "Quantum time-space tradeoffs for sorting". In: *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*. 2003, pp. 69–76.

[KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. "A Concrete Treatment of Fiat-Shamir Signatures in the Quantum Random-Oracle Model". In: *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10822. Lecture Notes in Computer Science. Springer, 2018, pp. 552–586. DOI: 10.1007/978-3-319-78372-7_18. URL: https://doi.org/10.1007/978-3-319-78372-7%5C_18.

[KP19] Tsvi Kopelowitz and Ely Porat. "The Strong 3SUM-INDEXING Conjecture is False". In: *CoRR* abs/1907.11206 (2019). arXiv: 1907.11206. URL: http://arxiv.org/abs/1907.11206.

[KŠW07] Hartmut Klauck, Robert Špalek, and Ronald de Wolf. "Quantum and classical strong direct product theorems and optimal time-space tradeoffs". In: *SIAM Journal on Computing* 36.5 (2007), pp. 1472–1493.

[Lev87] Leonid A. Levin. "One-way functions and pseudorandom generators". In: *Combinatorica* 7.4 (1987), pp. 357–363. DOI: 10.1007/BF02579323. URL: https://doi.org/10.1007/BF02579323.

[LR10] Gregory M. Lebovitz and Eric Rescorla. "Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)". In: *RFC* 5926 (2010), pp. 1–15. DOI: 10.17487/RFC5926. URL: https://doi.org/10.17487/RFC5926.

[LZ19a] Qipeng Liu and Mark Zhandry. "On finding quantum multi-collisions". In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2019, pp. 189–218.

[LZ19b] Qipeng Liu and Mark Zhandry. "Revisiting Post-Quantum Fiat-Shamir." In: *IACR Cryptology ePrint Archive* 2019 (2019), p. 262.

[MW19] Or Meir and Avi Wigderson. "Prediction from Partial Information and Hindsight, with Application to Circuit Lower Bounds". In: *Comput. Complex.* 28.2 (2019), pp. 145–183. DOI: 10.1007/s00037-019-00177-4. URL: https://doi.org/10.1007/s00037-019-00177-4.

[NABT14] Aran Nayebi, Scott Aaronson, Aleksandrs Belovs, and Luca Trevisan. "Quantum lower bound for inverting a permutation with advice". In: *CoRR* abs/1408.3193 (2014). arXiv: 1408.3193. URL: http://arxiv.org/abs/1408.3193.

[NABT15] Aran Nayebi, Scott Aaronson, Aleksandrs Belovs, and Luca Trevisan. "Quantum lower bound for inverting a permutation with advice". In: *Quantum Information & Computation* 15.11&12 (2015), pp. 901–913. URL: http://www.rintonpress.com/xxqic15/qic-15-1112/0901-0913.pdf.

[ST18] Alexander V. Smal and Navid Talebanfard. "Prediction from partial information and hindsight, an alternative proof". In: *Inf. Process. Lett.* 136 (2018), pp. 102–104. DOI: 10.1016/j.ipl.2018.04.011. URL: https://doi.org/10.1016/j.ipl.2018.04.011.

[TU16] Ehsan Ebrahimi Targhi and Dominique Unruh. “Post-Quantum Security of the Fujisaki-Okamoto and OAEP Transforms”. In: *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II*. Ed. by Martin Hirt and Adam D. Smith. Vol. 9986. Lecture Notes in Computer Science. 2016, pp. 192–216. DOI: 10.1007/978-3-662-53644-5_8. URL: https://doi.org/10.1007/978-3-662-53644-5%5C_8.

[Unr07] Dominique Unruh. “Random Oracles and Auxiliary Input”. In: *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*. Ed. by Alfred Menezes. Vol. 4622. Lecture Notes in Computer Science. Springer, 2007, pp. 205–223. DOI: 10.1007/978-3-540-74143-5_12. URL: https://doi.org/10.1007/978-3-540-74143-5%5C_12.

[Unr15] Dominique Unruh. “Non-Interactive Zero-Knowledge Proofs in the Quantum Random Oracle Model”. In: *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*. Ed. by Elisabeth Oswald and Marc Fischlin. Vol. 9057. Lecture Notes in Computer Science. Springer, 2015, pp. 755–784. DOI: 10.1007/978-3-662-46803-6_25. URL: https://doi.org/10.1007/978-3-662-46803-6%5C_25.

[Unr17] Dominique Unruh. “Post-quantum Security of Fiat-Shamir”. In: *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*. Ed. by Tsuyoshi Takagi and Thomas Peyrin. Vol. 10624. Lecture Notes in Computer Science. Springer, 2017, pp. 65–95. DOI: 10.1007/978-3-319-70694-8_3. URL: https://doi.org/10.1007/978-3-319-70694-8%5C_3.

[Yao90] AC-C Yao. “Coherent functions and program checkers”. In: *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. 1990, pp. 84–94.

[Zha12] Mark Zhandry. “Secure Identity-Based Encryption in the Quantum Random Oracle Model”. In: *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*. Ed. by Reihaneh Safavi-Naini and Ran Canetti. Vol. 7417. Lecture Notes in Computer Science. Springer, 2012, pp. 758–775. DOI: 10.1007/978-3-642-32009-5_44. URL: https://doi.org/10.1007/978-3-642-32009-5%5C_44.

[Zha19] Mark Zhandry. “How to record quantum queries, and applications to quantum indifferentiability”. In: *Annual International Cryptology Conference*. Springer. 2019, pp. 239–268.