Succinct Arguments in the Quantum Random Oracle Model

Alessandro Chiesa[1], Peter Manohar[2], and Nicholas Spooner[1]

[1] University of California, Berkeley
{alexch,nick.spooner}@berkeley.edu
[2] Carnegie Mellon University
pmanohar@cs.cmu.edu

Abstract. Succinct non-interactive arguments (SNARGs) are highly efficient certificates of membership in non-deterministic languages. Constructions of SNARGs in the random oracle model are widely believed to be post-quantum secure, provided the oracle is instantiated with a suitable post-quantum hash function. No formal evidence, however, supports this belief.

In this work we provide the first such evidence by proving that the SNARG construction of Micali is unconditionally secure in the *quantum* random oracle model. We also prove that, analogously to the classical case, the SNARG inherits the zero knowledge and proof of knowledge properties of the PCP underlying the Micali construction. We thus obtain the first zero knowledge SNARG of knowledge (zkSNARK) that is secure in the quantum random oracle model.

Our main tool is a new lifting lemma that shows how, for a rich class of oracle games, we can *generically* deduce security against quantum attackers by bounding a natural classical property of these games. This means that in order to prove our theorem we only need to establish *classical* properties about the Micali construction. This approach not only lets us prove post-quantum security but also enables us to prove explicit bounds that are tight up to small factors.

We additionally use our techniques to prove that SNARGs based on interactive oracle proofs (IOPs) with round-by-round soundness are unconditionally secure in the quantum random oracle model. This result establishes the post-quantum security of many SNARGs of practical interest.

Keywords: succinct arguments · quantum random oracle model · probabilistically checkable proofs

1 Introduction

The design and analysis of cryptographic primitives that are plausibly secure against quantum attackers is an increasingly important goal. The expected advent of quantum computers demands the cryptography community to be prepared well in advance, so much so that the National Institute of Standards and

Technology (NIST) is *already* in the process of selecting, from among many proposals, a new set of cryptography standards that are "post-quantum" [50]. The proposals involve schemes for key agreement, public-key encryption, and digital signatures, and are intended to eventually replace existing standards based on the hardness of factoring or discrete logarithms.

In this paper we study the post-quantum security of a cryptographic primitive that has recently received much attention across theoretical and applied communities: *succinct arguments* [33]. These are argument systems [19] for non-deterministic languages where the communication complexity between the prover and verifier is sublinear in the size of the non-deterministic witness.[3] This notion originates in seminal works of Kilian [43] and Micali [49], which construct succinct arguments for languages in $\mathsf{NTIME}(T(n))$ where communication complexity is $\mathsf{poly}(\lambda, \log T(n))$ and the time complexity of the verifier is $\mathsf{poly}(\lambda, n, \log T(n))$; here λ is the security parameter.

Researchers have studied many aspects of succinct arguments in the last two decades, leading to numerous constructions with different tradeoffs [62], efficient realizations in code [55, 18, 56, 24, 46, 57, 39], real-world deployments [29, 51], and standardization efforts [67]. A particularly useful feature is that many succinct arguments can be made zero knowledge with minimal overhead. At present, however, *most approaches to obtain efficient succinct arguments are "pre-quantum"*, since they rely on the discrete logarithm problem (and more).

A notable exception is a class of succinct arguments obtained by combining two ingredients: (a) probabilistic proof systems, which are unconditionally secure, and (b) cryptographic hash functions, for which we have post-quantum candidates. This class includes the succinct interactive argument of Kilian [43], which use probabilistically checkable proofs (PCPs) [7, 30, 5, 4] and collision-resistant hash functions. It also includes the succinct non-interactive argument (SNARG) of Micali [49], which uses PCPs and random oracles. More generally, by using random oracles one can construct a SNARG from a multi-round generalization of PCPs known as interactive oracle proofs (IOPs) [13, 54]. *All of these succinct arguments are widely believed to be post-quantum*, provided the hash function is suitably instantiated [11].[4]

There is, however, no formal evidence that supports the above widely-held belief. Since succinct arguments are a fundamental cryptographic primitive with both theoretical and real-world applications, it is important to prove guarantees on their post-quantum security.

1.1 SNARGs with random oracles

In this paper we focus our attention on the SNARG construction of Micali [49], which is unconditionally secure in the random oracle model [10, 53]. SNARGs in

[3] Achieving communication complexity that is sublinear in the witness size is known to *require* relaxing soundness from statistical to computational, provided one assumes standard complexity conjectures [34, 35].

[4] There is also a class of lattice-based succinct arguments that is plausibly post-quantum; see Section 1.3.

the random oracle model are not only plausibly post-quantum secure but also enjoy other desirable features. Namely, the random oracle can be heuristically instantiated via hash functions that avoid expensive public-key cryptographic operations. Moreover, the SNARG uses a transparent (public-coin) setup, because the only public parameter needed to produce/verify proofs is the choice of hash function.

We are thus interested in asking: can we establish formal evidence that the SNARG construction of Micali is post-quantum secure? One way to establish formal evidence is to prove security in a quantum analogue of the random oracle model, as we now explain. A quantum attacker can, among other things, evaluate a hash function in superposition when given the hash function's code. This enables the attacker, for instance, to find pre-images [38] or collisions [20] faster than a classical attacker. In light of this, Boneh et al. [15] have argued that, in the quantum setting, the correct way to model a random oracle is to allow the attacker to query the random oracle in superposition. The resulting model is known as the *quantum random oracle model* (QROM), and a line of work has established post-quantum security within this model for a variety of cryptographic primitives; see, e.g., [15, 64, 65, 58, 28].

Our goal is to study the SNARG construction of Micali in the quantum random oracle model. We also study the SNARG construction of BCS [13], which yields SNARGs of practical interest.

1.2 Our results

The main result of this paper is establishing that the SNARG construction of Micali [49] is unconditionally secure in the quantum random oracle model. This is the first formal evidence that supports the widely-held belief that this construction is post-quantum secure when the oracle is instantiated via a suitable post-quantum hash function.

Theorem 1 (informal) *The non-interactive argument of Micali, when based on a PCP with soundness error ϵ, has soundness error $O(t^2\epsilon + t^3/2^\lambda)$ against quantum attackers that make t queries to a random oracle with output size λ. This soundness error is tight up to small factors.*

A key step in our proof, of independent interest, is a *Lifting Lemma* that shows how, for a rich class of "oracle games", we can *generically* deduce security against quantum attackers by bounding a natural classical property of these games, instability, that we introduce. This means that to prove Theorem 1 we only need to bound the instability of the Micali construction. This approach not only yields the theorem but also enables us to prove explicit bounds that are tight up to small factors.

If we base the Micali construction on suitable PCPs, we obtain new statements about the existence of post-quantum non-interactive arguments. First, if the PCP achieves (honest-verifier) zero knowledge and proof of knowledge then through the Micali construction we obtain a zero knowledge non-interactive argument of knowledge that is *unconditionally* secure in the quantum random oracle model. This strengthens a result of Unruh [59], which assumes the existence of a post-quantum Σ-protocol for NP. Moreover, if the PCP has polylogarithmic query complexity and verifier running time then we obtain the first construction of a zero knowledge succinct non-interactive argument of knowledge (zkSNARK) that is secure in the quantum random oracle model.

Theorem 2 (informal) *There exists a zero knowledge non-interactive argument of knowledge for* NP *in the quantum random oracle model. Moreover, the non-interactive argument is succinct, in the sense that arguments have size* λ^c *and can be verified in time* $(\lambda \cdot n)^c$*, where* λ *is the random oracle's security parameter,* n *is instance size, and* $c > 0$ *is a universal constant.*

The above theorem is stated for NP only for simplicity. Analogously to the classical case, a more general statement holds for all non-deterministic time languages by relying on suitable PCPs for non-deterministic time. For example, the PCP in [7] achieves proof of knowledge, can be made (honest-verifier) zero knowledge [27, 44], and supports non-deterministic time.

The BCS construction. We conclude with a result that demonstrates how the tools in this paper can be used to study the post-quantum security of protocols that are of practical interest. Since known PCP constructions are expensive, efficient constructions of succinct arguments in the random oracle model are typically based on the BCS construction [13], which instead uses interactive oracle proofs (IOPs) [13, 54], a multi-round extension of PCPs. This extension additionally captures IPs [6, 37] and IPCPs [41] as special cases.

We prove that the BCS construction is unconditionally secure in the quantum random oracle model, if applied to public-coin IOPs that have round-by-round soundness [21]. The resulting argument inherits proof of knowledge and zero knowledge properties of the underlying IOP.

Theorem 3 (informal) *The non-interactive argument of BCS, when based on a public-coin IOP with round-by-round soundness error* ϵ*, has soundness error* $O(t^2\epsilon + t^3/2^\lambda)$ *against quantum attackers that make* t *queries to a random oracle with output size* λ*. Moreover, it is an argument of knowledge if the IOP has round-by-round proof of knowledge, and it is a (statistical) zero knowledge argument if the IOP is honest-verifier zero knowledge.*

Round-by-round proof of knowledge is a natural notion that we introduce, analogous to round-by-round soundness, and is satisfied by many natural protocols. In particular, Theorem 3 enables us to deduce the post-quantum security of succinct arguments based on well-known IPs such as the sumcheck protocol [48] and the GKR protocol [36], as well as zkSNARKs based on recent IOPs such as [11, 3, 12]. These protocols (among others) are of interest to practitioners, and our result can be used to guide parameter choices in practice.

1.3 Related work

Argument systems that use random oracles. Several works study the post-quantum security of zero knowledge non-interactive arguments of knowledge that use random oracles, most notably those obtained by applying the Fiat–Shamir transformation [31] to a post-quantum Σ-protocol. Such arguments are used to achieve post-quantum digital signatures [22, 42, 9], and underlie constructions submitted to the NIST call for post-quantum cryptography [50].

A security reduction for the Fiat–Shamir transformation in the quantum random oracle model has been recently achieved [26, 47]. Obtaining a security reduction had been elusive, as the classical approach of rewinding the adversary to reduce to special soundness of the Σ-protocol does not work for quantum adversaries.[5] Before, researchers were only able to prove security if the underlying Σ-protocol satisfies special properties [23, 60, 45], or resorted to proving security for alternative, less efficient, constructions such as the Unruh transformation [59].

The question that we study in this paper is complementary to these prior works. On the one hand, prior works study the security of the Fiat–Shamir transformation *given* that the underlying Σ-protocol is secure against efficient quantum attackers. On the other hand, we study protocols such as the Micali construction and BCS construction that can be viewed as applying the Fiat–Shamir transformation to specific public-coin protocols that are known to be unconditionally secure in the (classical) random oracle model. In particular, we establish *unconditional* security in the quantum random oracle model via an approach that considers the protocol as a whole (similarly to the classical analysis of these protocols).

The foregoing differences are reflected in a technical analysis that departs from prior works. Most of the effort in this paper is establishing *classical* security properties of the Micali and BCS constructions, which we then use to generically deduce their quantum security. This approach, besides being intuitive, yields tight bounds that can be used to guide parameter choices in practice.

Succinct arguments based on lattices. Several lattice problems are presumed to remain hard even against quantum adversaries, and researchers have relied on such problems to propose numerous cryptographic constructions that are plausibly post-quantum. A handful of works have used lattices to achieve various notions of succinct arguments that are plausibly post-quantum. Baum et al. [8] rely on the short integer solution (SIS) problem to obtain an argument system for arithmetic circuits where the communication complexity grows with the square-root of circuit size; the argument system is constant-round, public-coin, and honest-verifier zero knowledge. Boneh et al. [16, 17] and Gennaro et al. [32] rely on lattice knowledge assumptions to construct designated-verifier

[5] Rewinding quantum adversaries is a delicate matter [63] and, more importantly, special soundness does *not* imply post-quantum soundness (relative to some oracle) [2]. These difficulties have been circumvented by using new techniques that enable reducing directly to the (post-quantum) soundness of the underlying Σ-protocol.

SNARGs for boolean circuits, in the preprocessing model [14]. Whether one can use lattices to obtain public-coin argument systems with polylogarithmic communication complexity (as in the construction of Micali) remains an intriguing open problem.

2 Techniques

We discuss the main ideas behind our results. In Section 2.1 we recall the construction of Micali, and then in Section 2.2 we explain the challenges that arise when trying to prove its security in the quantum random oracle model. In Section 2.3 we outline our approach to obtain a proof of security for the Micali construction (Theorem 1); we elaborate on our approach in Sections 2.4 to 2.7. Finally, in Section 2.8 we discuss how to further establish zero knowledge and proof of knowledge; we thus obtain the first zkSNARK secure in the quantum random oracle model (Theorem 2).

We conclude in Section 2.9 by explaining how our techniques extend to establish post-quantum security for the BCS construction applied to many protocols of practical interest (Theorem 3).

Many of the proofs/sections have been omitted from this version of the paper due space limitations. We refer the reader to the full version of the paper for all relevant details.

2.1 The construction of Micali

The construction of Micali is a transformation that maps any *probabilistically checkable proof* (PCP) into a corresponding non-interactive argument in the random oracle model. (See Section 3.4 for the definition of a PCP, and Section 3.3 for that of a non-interactive argument.) The resulting non-interactive argument is *succinct*, i.e. a SNARG, provided the PCP has suitable parameters.

Let $(\mathbf{P}, \mathbf{V})$ be a PCP for a relation $\mathcal{R}$ with soundness error ϵ, proof length ℓ over alphabet Σ, and query complexity q. The honest prover $\mathbf{P}$ takes as input an instance-witness pair $(\mathbb{x}, \mathbb{w})$ and outputs a proof string $\Pi \colon [\ell] \to \Sigma$. The honest verifier $\mathbf{V}$ takes as input the instance $\mathbb{x}$, makes q probabilistic queries to a (possibly malicious) proof string $\tilde{\Pi} \colon [\ell] \to \Sigma$, and then accepts or rejects.

The PCP $(\mathbf{P}, \mathbf{V})$ for $\mathcal{R}$ is used to construct a SNARG $(\mathcal{P}, \mathcal{V})$ for $\mathcal{R}$, as follows.

The SNARG prover $\mathcal{P}$ takes as input an instance $\mathbb{x}$ and witness $\mathbb{w}$. First, $\mathcal{P}$ uses the random oracle h to commit to the proof string $\Pi := \mathbf{P}(\mathbb{x}, \mathbb{w})$ via a Merkle tree, obtaining a corresponding root rt. Second, $\mathcal{P}$ applies the random oracle h to the root rt in order to derive randomness r for the PCP verifier $\mathbf{V}$. Third, $\mathcal{P}$ simulates the PCP verifier $\mathbf{V}$ with the proof string Π, input $\mathbb{x}$, and randomness r, in order to deduce the queried locations of Π. Finally, $\mathcal{P}$ assembles a SNARG proof π that contains the root rt, answers to the queries, and an authentication path for each answer.

Observe that the SNARG proof π is succinct because it is small (it has size $|\pi| = O(q \cdot (\log |\Sigma| + \lambda \log \ell)) = O_\lambda(q)$ for $\ell, |\Sigma| = 2^{O(\lambda)}$) and it is cheap to validate via the algorithm described next.

The SNARG verifier $\mathcal{V}$ takes as input an instance $\mathbb{x}$ and a (possibly malicious) SNARG proof $\tilde{\pi}$. First, $\mathcal{V}$ uses the random oracle h to check that each answer in $\tilde{\pi}$ is certified by an authentication path relative to the claimed root $\tilde{\mathsf{rt}}$. Next, $\mathcal{V}$ applies the random oracle h to the root $\tilde{\mathsf{rt}}$ in order to derive randomness $\tilde{\mathsf{r}}$. Finally, $\mathcal{V}$ runs the PCP verifier $\mathbf{V}$ on the instance $\mathbb{x}$ and randomness $\tilde{\mathsf{r}}$, answering $\mathbf{V}$'s queries using the claimed answers in $\tilde{\pi}$.

The intuition behind the construction is that the soundness guarantee of a PCP holds only if the proof string $\tilde{\Pi}$ to be validated is *fixed* before the randomness $\tilde{\mathsf{r}}$ for the PCP verifier is known, and for this reason the SNARG prover must derive $\tilde{\mathsf{r}}$ by hashing a commitment $\tilde{\mathsf{rt}}$ to $\tilde{\Pi}$.

This construction is unconditionally secure in the random oracle model [49, 61, 13]:

Theorem 1. *The SNARG $(\mathcal{P}, \mathcal{V})$ has soundness error $O(t\epsilon + t^2/2^\lambda)$ against (classical) attackers that make at most t queries to the random oracle. This soundness error is tight up to small factors.*

A SNARG obtained via the Micali construction also inherits zero knowledge and proof of knowledge properties of the underlying PCP. We discuss these additional properties and how we establish them in the quantum setting later on in Section 2.8. We focus on soundness first.

2.2 Challenges in the quantum setting

Our goal is to show that the SNARG construction of Micali is unconditionally secure in the *quantum* random oracle model. Suppose that $\tilde{\mathcal{P}}$ is a t-query quantum prover that convinces the SNARG verifier $\mathcal{V}$ with probability δ (over the random oracle). We wish to construct a malicious PCP prover $\tilde{\mathbf{P}}$ that, using $\tilde{\mathcal{P}}$ as a subroutine, outputs a proof string $\tilde{\Pi} \colon [\ell] \to \Sigma$ that convinces the PCP verifier $\mathbf{V}$ with related probability $\epsilon(\delta, t)$ (here the probability is over the randomness of $\tilde{\mathbf{P}}$ and $\mathbf{V}$).

A natural approach to reduce the SNARG prover $\tilde{\mathcal{P}}$ to the PCP prover $\tilde{\mathbf{P}}$ would be to try to adapt to the quantum setting the reduction that is used for the classical setting. Below we recall the classical reduction, and then explain why adapting it to the quantum case is challenging.

The reduction for classical attackers. The reduction from a *classical* SNARG prover $\tilde{\mathcal{P}}$ to a PCP prover $\tilde{\mathbf{P}}$ relies on a *straightline extractor*, as we now explain.

While the SNARG prover $\tilde{\mathcal{P}}$ outputs a short proof π that contains a Merkle root and a few decommitted values, the PCP prover $\tilde{\mathbf{P}}$ must output a "long" proof string $\tilde{\Pi}$. How can $\tilde{\mathbf{P}}$ obtain all this information from seeing only π? The answer is that, when running $\tilde{\mathcal{P}}$ as a subroutine, $\tilde{\mathbf{P}}$ observes the queries that $\tilde{\mathcal{P}}$ makes to the oracle, and these queries reveal the proof string $\tilde{\Pi}$.

This is only a caricature of how $\tilde{\mathbf{P}}$ actually works, though. The reason is that $\tilde{\mathcal{P}}$ need not produce a query sequence from which $\tilde{\mathbf{P}}$ can just read off a proof string $\tilde{\Pi}$ consistent with the Merkle root in π. For example, $\tilde{\mathcal{P}}$ could try to commit to many possible proof strings "in its head", derive the corresponding randomness from each commitment, and then select which commitment to include in π. Even worse, $\tilde{\mathcal{P}}$ could try to commit to a *partial* proof string $\tilde{\Pi}$ via an incomplete Merkle tree and, because the PCP verifier inspects only a small fraction of a proof string, hope that queries will land to leaves of the Merkle tree that do exist.

The proof of Theorem 1 shows that, despite these complications, there is a way for $\tilde{\mathbf{P}}$ to observe all queries and answers of a single execution of the SNARG prover $\tilde{\mathcal{P}}$, and then run an algorithm on these to extract a suitable proof string $\tilde{\Pi}$.

How to deal with quantum attackers?. If we now return to the case where the SNARG prover $\tilde{\mathcal{P}}$ is a quantum attacker, we are immediately confronted with a severe problem. Since $\tilde{\mathcal{P}}$ can query the random oracle in superposition, how can $\tilde{\mathbf{P}}$ "observe" queries and answers to the oracle? If $\tilde{\mathbf{P}}$ were to just measure $\tilde{\mathcal{P}}$'s query register, $\tilde{\mathcal{P}}$ may detect this and stop working. This basic problem has made obtaining security reductions against quantum attackers that access random oracles exceedingly difficult when compared to the case of classical attackers. Papers that study the security of cryptographic primitives in the quantum random oracle model have had to develop clever techniques to somehow circumvent this problem in various settings of interest.

Most relevant to this paper is a work of Zhandry [66] that introduces *compressed oracles*, a set of notions and techniques that enables a quantum algorithm to simulate access to a random oracle for a quantum attacker. This is achieved by replacing a random oracle $h\colon \{0,1\}^m \to \{0,1\}^n$ with the action of a specially-crafted unitary $\mathcal{O}$ that implicitly keeps track of queries. This is a quantum analogue of when, in the classical setting, a simulator merely observes the queries made by the attacker and maintains a database of the query-answer pairs. Formally, the classical simulator keeps track of a database D, which is a partial function $D\colon \{0,1\}^m \rightharpoonup \{0,1\}^n$. The database represents the part of the random oracle that has been "revealed" to the attacker by answering its queries. In the quantum setting, the state space of the quantum attacker is augmented with registers to store the database, which (loosely) keep track of the database D in superposition, as it evolves from query to query. Thus, while the original oracle h operates on the state $|\psi_{\mathcal{A}}\rangle$ of the adversary, the unitary $\mathcal{O}$ operates on a bipartite state $|\psi_{\mathcal{A}}, \psi_D\rangle$. This extended state represents a purification of the mixed state of the adversary induced by choosing the oracle h at random.

One may conjecture that compressed oracles, by virtue of "exposing" a quantum attacker's queries, make proving the quantum security of the Micali construction, or indeed of any construction using random oracles, straightforward. This is, unfortunately, not the case.

For example, compressed oracles allow us to argue that, given an adversary that outputs a convincing SNARG proof π with high probability, if we measure the database D after the adversary terminates, then with high probability one can find a convincing SNARG proof π in the database D. This does not allow us to reduce to soundness of the underlying PCP, however, because to do that we need to argue that one can extract a PCP proof Π from D (that is much longer than the SNARG proof π) that convinces the PCP verifier with high probability.

Nevertheless, compressed oracles are a useful starting point for this work, and indeed a basic lemma about compressed oracles plays the role of a hybrid in our security proof.

2.3 Outline of our approach

The ideas that we use in this paper to analyze the Micali construction are almost entirely generic, and can be used to analyze any *oracle game*. Informally, given a "base game" $G \subseteq A^k \times B^k \times C$, an adversary with oracle access to a random oracle h wins the oracle game for G if it outputs a tuple $(\mathbf{a}, \mathbf{b}, c) \in G$ where $h(a_i) = b_i$ for each $i \in [k]$. Oracle games are a natural notion that captures many games of interest, such as finding pre-images or finding collisions. Producing a valid proof in the Micali construction can also be cast as an oracle game, and we shall view the soundness property as stating that the value (maximum winning probability) of this game is small.

Our proof of quantum security consists of two main parts. First, we *generically* reduce the value of any oracle game to the *instability* of the game, a purely classical property of the game that we introduce. Second, we analyze the instability of the oracle game induced by the Micali construction. The instability of this oracle game is not too difficult to analyze because it is a classical quantity, and the "hard work" is crisply, and conveniently, encapsulated within our generic reduction. We view bounding values of oracle games via instability as the main technical contribution of this paper.

We now elaborate on our approach: in Section 2.4 we recast prior work in the language of oracle games; in Section 2.5 we explain what is instability and how we use it to bound game values; in Section 2.6 we introduce conditional instability and use it to prove tighter bounds on oracle game values; and in Section 2.7 we outline the analysis of instability for the Micali construction.

2.4 From oracle games to database games

We begin with a sequence of three games whose values are closely related. These games play the role of hybrids in our analysis, and are all defined relative to the given base game $G \subseteq A^k \times B^k \times C$.

- **Oracle game.** This is the game defined earlier that is played in the real world, using a random oracle h. The adversary wins if it outputs a tuple $(\mathbf{a}, \mathbf{b}, c) \in G$ with $h(a_i) = b_i$ for each $i \in [k]$.
- **Simulated oracle game.** The simulator of Zhandry [66] is used to run the adversary and its final state is measured, leading to a tuple $(\mathbf{a}, \mathbf{b}, c)$ *and* a database D. The adversary wins if $(\mathbf{a}, \mathbf{b}, c) \in G$ and $D(a_i) = b_i$ for each $i \in [k]$. (The oracle $h\colon \{0,1\}^m \to \{0,1\}^n$ is now replaced by the database $D\colon \{0,1\}^m \rightharpoonup \{0,1\}^n$ stored by the simulator.)

- **Database game.** Again the simulator of Zhandry is used to run the adversary, leading to a tuple $(\mathbf{a}, \mathbf{b}, c)$ and a database D. However, now we ignore $(\mathbf{a}, \mathbf{b}, c)$ and only consider D. The adversary wins if there *exists* $(\mathbf{a}', \mathbf{b}', c') \in G$ such that $D(a_i) = b_i$ for each $i \in [k]$.

We let $\omega^*_{\mathsf{O}}(G, t)$, $\omega^*_{\mathsf{S}}(G, t)$, and $\omega^*_{\mathsf{D}}(G, t)$ denote the values of the oracle game, simulated oracle game, and database game against quantum adversaries that make at most t oracle queries.

A result of Zhandry [66, Lemma 5], when stated via the notions above, shows that $\sqrt{\omega^*_{\mathsf{O}}(G, t)} \leq \sqrt{\omega^*_{\mathsf{S}}(G, t)} + \sqrt{k/2^n}$. Moreover, $\omega^*_{\mathsf{S}}(G, t) \leq \omega^*_{\mathsf{D}}(G, t)$ holds trivially, because winning the simulated oracle game implies winning the database game, by taking $(\mathbf{a}', \mathbf{b}', c') := (\mathbf{a}, \mathbf{b}, c)$. In sum:

Lemma 1. *For any base game G,*

$$\sqrt{\omega^*_{\mathsf{O}}(G, t)} \leq \sqrt{\omega^*_{\mathsf{D}}(G, t)} + \sqrt{k/2^n} \ .$$

The above lemma is a conceptualization of prior work, and is the starting point for the technical contributions of this paper. In particular, the lemma tells us that in order to bound the maximum winning probability of a quantum adversary in an oracle game (played in the real world) it suffices to bound the maximum winning probability of the adversary in the corresponding database game.

See the full version of the paper for more details.

2.5 A basic lifting lemma for database games

We describe how we use a *classical* quantity $\mathbf{I}(\mathcal{P}_G, t)$ to bound $\omega^*_{\mathsf{D}}(G, t)$, the maximum winning probability of any t-query quantum algorithm in the database game of G. When combined with the hybrids in Section 2.4, this reduces the quantum security of oracle games to studying $\mathbf{I}(\mathcal{P}_G, t)$.

Given a base game G, we let $\mathcal{P}_G$ be the set of databases that win the database game of G. In the classical setting, a natural way to bound the maximum winning probability of the database game is to compute, for each possible database $D \notin \mathcal{P}_G$ (a database that is currently losing the game), the maximum probability that adding a query-answer pair to D puts D in $\mathcal{P}_G$. Assuming that the empty database is not in $\mathcal{P}_G$ (for otherwise one can win trivially), this quantity characterizes the probability that the adversary gets lucky and ends up with a winning database D.

We define the *instability* of $\mathcal{P}_G$ with query bound t, denoted $\mathbf{I}(\mathcal{P}_G, t)$, to be the maximum probability that, for any database D containing less than t queries, making one additional (classical) query changes whether or not D is in $\mathcal{P}_G$. *The foregoing argument explains that the classical value of the database game G is bounded by $t \cdot \mathbf{I}(\mathcal{P}_G, t)$.* Intuitively this is because each query can increase the probability that the database D is in $\mathcal{P}_G$ by at most $\mathbf{I}(\mathcal{P}_G, t)$.

We prove that an analogous result holds for quantum adversaries as well. We call this lemma a lifting lemma, because it enables us to use the *classical* quantity of instability to prove a bound on the maximum winning probability of *quantum* adversaries. The version below is a "basic" version, because we shall ultimately need a stronger statement, as we discuss in Section 2.6. The result below extends an idea of Zhandry sketched in [66, Section 4.3].

Lemma 2 (Basic lifting lemma). *For any base game G,*

$$\omega^*_{\mathsf{D}}(G,t) \leq O\big(t^2 \cdot \mathbf{I}(\mathcal{P}_G,t)\big) \ .$$

In particular, combining the above with Lemma 1, we get

$$\omega^*_{\mathsf{O}}(G,t) \leq O\big(t^2 \cdot \mathbf{I}(\mathcal{P}_G,t) + k/2^n\big) \ .$$

Even the above basic lifting lemma is a powerful tool. For example, suppose that G is the collision game, where the adversary wins if it outputs an oracle collision. Then $\mathbf{I}(\mathcal{P}_G,t) < t/2^n$, because if D is a database with no collisions and less than t entries, then making one more query produces a collision with probability less than $t/2^n$, and if D has collisions then it is not possible to make an additional query and remove collisions. Then (since $k = 2$ in the collision game) the lifting lemma immediately tells us that $\omega^*_{\mathsf{O}}(G,t) \leq O(t^3/2^n)$, which shows that the probability that a t-query quantum oracle algorithm finds a collision is bounded by $O(t^3/2^n)$. This further simplifies the analysis of this fact in [66] and matches the bound of [1] (which is tight [20]).

We now sketch the proof of the basic lifting lemma. The proof sketch differs slightly from the actual proof, as in the actual proof we do a slightly more complicated analysis that gives us smaller constants. The main ideas, however, remain the same.

We let P_G be the operator that projects onto databases that win the database game G: for any basis state $|D\rangle$ in the database register, $P_G|D\rangle = |D\rangle$ if $D \in \mathcal{P}_G$, and $P_G|D\rangle = 0$ if $D \notin \mathcal{P}_G$; P_G acts as the identity on other registers. If $|\phi\rangle$ is the final joint state of the quantum adversary and database, then $\|P_G|\phi\rangle\|^2$ is the probability that $D \in \mathcal{P}_G$ after measurement. We will assume that $\emptyset \notin \mathcal{P}_G$, i.e., that the empty database does not win the database game of G (or else the adversary can win by doing nothing).

We can represent any simulated quantum adversary making at most t queries as a sequence of unitary operators $U = A_t\mathcal{O}A_{t-1}\mathcal{O}\ldots A_1\mathcal{O}$ applied to an initial state $|\phi_0,\emptyset\rangle := |\phi_0\rangle \otimes |\emptyset\rangle$, where $\mathcal{O}$ is the compressed oracle and $|\emptyset\rangle$ is the state of the empty database. Each A_i acts non-trivially only on the registers of the adversary being simulated and P_G acts non-trivially only on the database registers, so P_G and A_i commute. So, if P_G and $\mathcal{O}$ were to also commute, then we could simply conclude that $P_G U|\phi_0,\emptyset\rangle = UP_G|\phi_0,\emptyset\rangle = 0$, i.e., that the adversary never wins. (Here we used the fact that $\emptyset \notin \mathcal{P}_G$.)

However, it is *not* the case that P_G and $\mathcal{O}$ commute. This should be expected because in general an adversary can win with some positive probability. However, if we could show that they *almost* commute, then we could apply the previous argument to show that $P_G U|\phi_0,\emptyset\rangle \approx UP_G|\phi_0,\emptyset\rangle = 0$; i.e., the adversary wins with small probability. The notion of "almost" commuting we use is that the operator norm $\|[P_G,\mathcal{O}]\|$ of the commutator $[P_G,\mathcal{O}] := P_G\mathcal{O} - \mathcal{O}P_G$ is small.

Unfortunately, for interesting games the operator norm $\|[P_G, \mathcal{O}]\|$ may not be small. For example, if G is the collision game and D is a database with a pre-image of every $y \in \{0,1\}^n$ but no collisions, then $\|[P_G, \mathcal{O}] \, |x, u, D\rangle\| = 1$. Generally, this norm may be large if D has many entries.

Query-bounded adversaries, however, cannot produce nonzero amplitudes on databases with more entries than the query bound. Hence, intuitively we should not consider states that correspond to large databases when bounding the operator norm of the aforementioned commutator. We follow this intuition by introducing the notion of a *projected oracle*, which acts as the compressed oracle except that it discards databases that do not belong to a certain subset.

Definition 1. *Let P be the operator that projects onto databases that belong to a given subset $\mathcal{P}$ of databases. A* **projected oracle** *is an operator of the form $P\mathcal{O}P$.*

We thus consider the projected oracle $P_t\mathcal{O}P_t$, where P_t is operator that projects onto databases containing at most t queries. For adversaries that make at most t queries, replacing $\mathcal{O}$ with $P_t\mathcal{O}P_t$ has no effect because the adversary cannot create a database that contains more than t entries. Moreover, $\|[P_G, P_t\mathcal{O}P_t] \, |D\rangle\| = 0$ if D contains more than t entries, so the operator norm of $[P_G, P_t\mathcal{O}P_t]$ accounts for the action of $\mathcal{O}$ *only* on databases containing at most t entries.

In sum, projected oracles allow us to cleanly compute the operator norm only over databases that are reachable by an adversary making a bounded number of queries. By carefully analyzing the action of $\mathcal{O}$, we show that

$$\|[P_G, P_t\mathcal{O}P_t]\|^2 \leq O\big(\mathbf{I}(\mathcal{P}_G, t)\big) \ .$$

We additionally prove that $\|P_G U \, |\phi_0, \emptyset\rangle - U P_G \, |\phi_0, \emptyset\rangle\| \leq t \|[P_G, P_t\mathcal{O}P_t]\|$. Combining these two inequalities yields the lifting lemma.

See Section 4.1 for more details.

2.6 Stronger lifting via conditional instability

The lifting lemma implies that to prove soundness of the Micali construction, it suffices to bound the instability of the Micali database game. Unfortunately, the instability of the Micali database game is actually large, even given the query bound. For example, suppose that D is a database containing Merkle trees for many different proof strings, but each of these Merkle trees has (miraculously) the same root due to collisions. Then, the probability that querying the root yields a good randomness for the underlying PCP verifier is large, because the answer to the query only needs to be a good random string for any one of the many proofs that D contains.

This counterexample, however, should not be of concern because it relies on the database having many collisions, and we have already argued that creating even a single collision in the database is difficult. To deal with this issue, we introduce the notion of *conditional instability*: $\mathbf{I}(\mathcal{P} \,|\, \mathcal{Q}, t)$. This is a refined notion of instability that allows us to condition on events, e.g., that the database has no collisions. Our main technical contribution is the following stronger variant of Lemma 2.

Definition 2. *A database property $\mathcal{P}$ is a set of databases. The complement of $\mathcal{P}$ is $\bar{\mathcal{P}}$.*

Lemma 3 (Lifting lemma). *For any base game G and database property $\mathcal{Q}$,*

$$\omega^*_{\mathsf{D}}(G,t) \le O\Big(t^2 \cdot \big(\mathbf{I}(\mathcal{P}_G \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)\big)\Big) \enspace .$$

In particular, combining the above with Lemma 1, we get

$$\omega^*_{\mathsf{O}}(G,t) \le O\Big(t^2 \cdot \big(\mathbf{I}(\mathcal{P}_G \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)\big) + k/2^n\Big) \enspace .$$

The above statement is an "instability analogue" of the standard fact that for any two events E_1 and E_2, $\Pr[E_1] \le \Pr[E_1 \cup E_2] \le \Pr\big[E_1 \mid \bar{E_2}\big] + \Pr[E_2]$.

The proof of Lemma 3 has three steps. First, we relax the database game $\mathcal{P}_G$ so that the adversary wins if the database is in $\mathcal{P}_G \cup \mathcal{Q}$. Clearly, winning the relaxed game is only easier than the original database game. Lemma 2 then implies that $\omega^*_{\mathsf{D}}(G,t) \le O\Big(t^2 \cdot \mathbf{I}(\mathcal{P}_G \cup \mathcal{Q}, t)\Big)$. Finally, we show that for any two database properties $\mathcal{P}$ and $\mathcal{Q}$ it holds that $\mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t) \le \mathbf{I}(\mathcal{P} \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)$, which completes the proof.

We remark that Lemma 3 cannot be proved by simply arguing that $\mathbf{I}(\mathcal{P}, t) \le \mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t)$ and then applying Lemma 2. This is because $\mathbf{I}(\mathcal{P}, t)$ and $\mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t)$ are in general *incomparable* (see Proposition 5 for examples).

See Section 4.2 for more details.

2.7 Instability of the Micali oracle game

Armed with our lifting lemma, establishing the quantum security of the Micali construction is now relatively straightforward. Let $\mathcal{P}_{\mathsf{Mic}}$ be the database property for the Micali game, and let $\bar{\mathcal{P}}_{\mathsf{col}}$ be the no-collision property (the set of databases that do not contain collisions). We show that, for a random oracle of the form $h\colon \{0,1\}^{2\lambda} \to \{0,1\}^{\lambda}$,

$$\mathbf{I}(\mathcal{P}_{\mathsf{col}}, t) < t/2^{\lambda} \quad \text{and} \quad \mathbf{I}(\mathcal{P}_{\mathsf{Mic}} \mid \bar{\mathcal{P}}_{\mathsf{col}}, t) < \varepsilon + O(t/2^{\lambda}) \enspace .$$

Proving each of these inequalities is merely a classical argument.

- $\mathbf{I}(\mathcal{P}_{\mathsf{col}}, t)$: If D is a database containing less than t entries and has a collision, then adding an entry to D cannot remove the collision, so the probability that adding a new entry to D makes D have no collisions is 0. Let D be a database containing less than t entries and no collisions. For any new query x, adding the query-answer pair (x, y) to D for a random y will contain a collision with probability less than $t/2^{\lambda}$. Thus, $\mathbf{I}(\mathcal{P}_{\mathsf{col}}, t) < t/2^{\lambda}$.

- $\mathbf{I}(\mathcal{P}_{\mathsf{Mic}} \,|\, \bar{\mathcal{P}}_{\mathsf{col}}, t)$: It is impossible to go from a database D in $\mathcal{P}_{\mathsf{Mic}}$ to a database D not in $\mathcal{P}_{\mathsf{Mic}}$ by adding entries. Let D be a database not in $\mathcal{P}_{\mathsf{Mic}}$ containing less than t entries that contains no collisions. There are two ways to make D in $\mathcal{P}_{\mathsf{Mic}}$: either the new query is for the randomness of the PCP verifier in the Micali construction, in which case this finds a good choice of randomness with probability at most ε, or the new query extends one of the Merkle trees that the adversary is constructing. To extend the Merkle tree the adversary must find a pre-image, which happens with probability less than $O(t/2^{\lambda})$. Hence, $\mathbf{I}(\mathcal{P}_{\mathsf{Mic}} \,|\, \bar{\mathcal{P}}_{\mathsf{col}}, t) < \varepsilon + O(t/2^{\lambda})$, completing the proof.

Combining these bounds on instability with the lifting lemma completes the proof of soundness, and completes a proof sketch for Theorem 1. See the full version of the paper for more details.

2.8 zkSNARKs in the QROM

We have so far discussed how to establish soundness of the Micali construction in the quantum setting. We now discuss how to further establish zero knowledge and proof of knowledge, obtaining the first zkSNARKs secure in the quantum random oracle model (and thereby proving Theorem 2).

Zero knowledge. In the classical setting, the Micali construction achieves statistical zero knowledge provided the underlying PCP is (honest-verifier) statistical zero knowledge (and leaves in the Merkle tree are suitably salted to ensure statistical hiding of unrevealed leaves) [40, 13]. In the quantum setting, an analogous statement is immediate simply because the zero knowledge property holds against computationally unbounded verifiers *that make an unbounded number of queries to the random oracle*, and any quantum verifier can be simulated by an unbounded verifier.

Proof of knowledge. In the classical setting, the Micali construction achieves proof of knowledge provided the underlying PCP is a proof of knowledge [61]. The quantum analogue of this statement, however, does *not* immediately follow from our soundness analysis. Recall that our strategy was to bound the instability of the Micali property for $x \notin \mathcal{L}$, conditioned on no collisions. But when $x \in \mathcal{L}$ this approach will not work, because the instability of the Micali property even conditioned on the absence of collisions is 1 (as witnessed by the existence of the honest prover).

Nevertheless, the tools that we develop in this work are flexible enough that we can apply them to also establish proof of knowledge. We consider the following natural extractor strategy: run the prover until completion, and measure the database. Then, for each entry in the database, try to extract a PCP proof rooted at that entry, and then run the PCP extractor on this proof.

Let $\mathcal{P}$ be the set of databases D where there exists a root rt such that D wins the Micali game with a SNARG proof rooted at rt, but the PCP extractor does not extract a valid witness from the PCP proof rooted at rt. If the prover wins the Micali game but the extractor fails, then D must be in $\mathcal{P}$. We then argue that $\mathbf{I}(\mathcal{P} \,|\, \bar{\mathcal{P}}_{\mathsf{col}}, t)$ is at most $\mathsf{k} + O(t/2^{\lambda})$, where k is the knowledge error of the underlying PCP. Intuitively, this is because if the PCP extractor fails to extract a witness from the PCP proof Π rooted at rt, then Π convinces the verifier with probability at most k, and hence the probability of finding good randomness for Π is at most k. Combining this with Lemma 3 implies that the probability that the prover wins the Micali game but the extractor fails is at most $O(t^2\mathsf{k}+t^3/2^\lambda)$. Hence, if μ is the probability that the prover wins the Micali game, then the probability that the extractor succeeds is at least $\Omega(\mu - t^2\mathsf{k} - t^3/2^\lambda)$.

See the full version of the paper for more details.

2.9 The BCS construction: succinct arguments beyond Micali

We apply our techniques to prove post-quantum security of the BCS construction [13], when the underlying public-coin IOP satisfies a notion of soundness achieved by many protocols of practical interest. The notion is *round-by-round soundness*, and was introduced for IPs in [21] for the purposes of facilitating proofs of security of the Fiat–Shamir transformation for correlation-intractable hash functions. The notion can be extended in a straightforward way to any IOP, and this is the notion that we consider in this work. We further show that if the underlying IOP is honest-verifier zero knowledge and/or has round-by-round proof of knowledge, then the BCS argument inherits these properties. Round-by-round proof of knowledge is a type of knowledge property that is analogous to round-by-round soundness (and is also achieved by many protocols of practical interest). Below we sketch our analysis; see the full version of the paper for details.

Soundness. An IOP has round-by-round soundness if, for any partial transcript tr of the protocol, one can tell if tr is "doomed", i.e., that it is highly unlikely to be accepted by the verifier when completed to a full transcript; a doomed full transcript is never accepted by the verifier.

By the lifting lemma, in order to prove the post-quantum security of the BCS construction it suffices to bound the conditional instability of the database property $\mathcal{P}$, where $D \in \mathcal{P}$ if D contains a partial transcript where the last verifier message has flipped the transcript from "doomed" to "not doomed". We argue that $\mathbf{I}(\mathcal{P} \mid \bar{\mathcal{P}}_{\mathsf{col}}, t) < \epsilon + O(t/2^\lambda)$, where ϵ is the round-by-round soundness error of the IOP. The proof is similar to the proof for the Micali construction. If $D \notin \mathcal{P}$, there are two ways to add an entry and make $D \in \mathcal{P}$: either the new query is for the randomness of the next verifier message in the IOP for some doomed transcript tr, in which case we find a message that makes tr not doomed with probability ϵ; or the new query extends one of the Merkle trees that the adversary is constructing, which happens with probability less than $O(t/2^\lambda)$ as this implies finding a pre-image. Hence, $\mathbf{I}(\mathcal{P} \mid \bar{\mathcal{P}}_{\mathsf{col}}, t) < \epsilon + O(t/2^\lambda)$, which completes the proof.

Zero knowledge. As in the case of Micali, zero knowledge is straightforward, as the BCS construction classically achieves *statistical* zero knowledge when the IOP is honest-verifier zero knowledge.

Proof of knowledge. Analogously to our analysis of the Micali construction, we define a property $\mathcal{Q}$, where $D \in \mathcal{Q}$ if D contains a partial transcript that is in $\mathcal{P}$ but the BCS extractor fails to extract a valid witness. We then argue that $\mathbf{I}(\mathcal{Q} \mid \bar{\mathcal{P}}_{\mathsf{col}}) < \mathsf{k} + O(t/2^\lambda)$, where k is the round-by-round knowledge error of the IOP; the proof of this fact is similar to the proof of soundness. We conclude that if the prover causes the verifier to accept with probability at least μ, then the probability that the extractor succeeds is at least $\Omega(\mu - t^2\mathsf{k} - t^3/2^\lambda)$.

3 Preliminaries

We denote by $\mathcal{R}$ a binary relation of instance-witness pairs $(\mathbb{x}, \mathbb{w})$, and by $\mathcal{L}(\mathcal{R})$ its corresponding language, which is the set $\{\mathbb{x} \mid \exists\, \mathbb{w} \text{ s.t. } (\mathbb{x}, \mathbb{w}) \in \mathcal{R}\}$. We denote by $f\colon X \to Y$ a function from a set X to a set Y; similarly, we denote by $f\colon X \rightharpoonup Y$ a *partial* function from a set X to a set Y, i.e., a function $f\colon X \to Y \cup \{\bot\}$, where $\bot \notin Y$ is a special symbol indicating that $f(x)$ is undefined.

3.1 Quantum notation

We briefly recall standard quantum notation. We let $|\phi\rangle$ denote an arbitrary quantum state, and let $|x\rangle$ denote an element of the standard (computational) basis. The norm of a state $|\phi\rangle$ is $\||\phi\rangle\| := \sqrt{\langle\phi|\phi\rangle}$. In general, the states that we consider will have norm 1. The operator norm of an operator A is $\|A\|$, defined to be $\max_{|\phi\rangle : \||\phi\rangle\|=1} \|A\,|\phi\rangle\|$. Note that if A is unitary then $\|A\| = 1$. The commutator of two operators A and B is $[A, B] := AB - BA$. The following proposition relates operator norms and commutators.

Proposition 1. *Let A, B, C be operators with $\|B\|, \|C\| \leq 1$. Then*

$$\|[A, BC]\| \leq \|[A, B]\| + \|[A, C]\| \ .$$

Proof. By definition, $[A, BC] = ABC - BCA = ABC - BAC + BAC - BCA = [A, B]C + B[A, C]$. Therefore, $\|[A, BC]\| \leq \|[A, B]C\| + \|B[A, C]\| \leq \|[A, B]\| + \|[A, C]\|$, as $\|B\|, \|C\| \leq 1$.

A projector P is an idempotent linear operator (i.e., $P^2 = P$). Throughout, we will only consider orthogonal projectors of the form $P_S := \sum_{x \in S} |x\rangle\langle x|$, where S is a set of binary strings. Measuring a state $|\phi\rangle$ in the standard basis results in an output that is in S with probability equal to $\|P_S\,|\phi\rangle\|^2$. Since all P_S are diagonal in the same basis, they commute with each other. Note that for any non-zero orthogonal projector P it holds that $\|P\| = 1$. In particular, since $\|AB\| \leq \|A\|\|B\|$, we see that if A is the product of projectors and unitaries then $\|A\| \leq 1$.

3.2 Oracle algorithms

Let $f\colon \{0,1\}^m \to \{0,1\}^n$ be a function. The standard way to model oracle access to f in the quantum setting is via a unitary operator O_f that acts as $|x,y\rangle \mapsto |x, y \oplus f(x)\rangle$ for all $x \in \{0,1\}^m$ and $y \in \{0,1\}^n$. We label the input and output registers X and Y, respectively.

A *t-query quantum oracle algorithm* $\mathcal{A}$ is specified via $m, n \in \mathbb{N}$, t unitary operators $A_1, \ldots, A_t$ and an initial state $|\phi_0\rangle$ on four registers $\mathsf{X}, \mathsf{Y}, \mathsf{S}, \mathsf{T}$. The register X is on m qubits and is for queries to the oracle; the register Y is on n qubits and is for answers from the oracle; the register S is for the output of $\mathcal{A}$; and the register T is for scratch space of $\mathcal{A}$. The initial state $|\phi_0\rangle$ and unitary operators A_i need not be efficiently computable.

We write $\left|\mathcal{A}^f\right\rangle$ to denote $A_t O_f A_{t-1} O_f \cdots A_1 O_f |\phi_0\rangle$, the final state of the adversary before measurement. (We implicitly extend O_f to act as the identity on S, T.) We write $\mathcal{A}^f$ to denote the random variable which is the outcome of measuring the register S of $\left|\mathcal{A}^f\right\rangle$ in the computational basis. This is the output of $\mathcal{A}$ when accessing the oracle f.

A *random oracle* is a function $h\colon \{0,1\}^m \to \{0,1\}^n$ sampled from $\mathcal{U}(m,n)$, the uniform distribution over functions from $\{0,1\}^m$ to $\{0,1\}^n$. We write $h \leftarrow \mathcal{U}(m,n)$ to say that h is sampled from $\mathcal{U}(m,n)$. In the quantum random oracle model [15], we study $\mathcal{A}^h$ for $h \leftarrow \mathcal{U}(m,n)$.

3.3 Non-interactive arguments in the quantum random oracle model

Let $(\mathcal{P}, \mathcal{V})$ be two polynomial-time (classical) algorithms, known as the prover and verifier. We say that $(\mathcal{P}, \mathcal{V})$ is a *non-interactive argument in the quantum random oracle model* (QROM) with soundness error ϵ for a relation $\mathcal{R}$ if it satisfies the following properties.

- **Completeness.** For every $(\mathbbm{x}, \mathbbm{w}) \in \mathcal{R}$ and function $h \in \mathcal{U}(2\lambda, \lambda)$, $\mathcal{P}^h(\mathbbm{x}, \mathbbm{w})$ outputs a (classical) proof string π for which $\mathcal{V}^h(\mathbbm{x}, \pi) = 1$.
- **Soundness.** For every $\mathbbm{x} \notin \mathcal{L}(\mathcal{R})$ and t-query quantum oracle algorithm $\tilde{\mathcal{P}}$, the probability over a function $h \leftarrow \mathcal{U}(2\lambda, \lambda)$ and (classical) proof string $\tilde{\pi} \leftarrow \tilde{\mathcal{P}}^h$ that $\mathcal{V}^h(\mathbbm{x}, \tilde{\pi}) = 1$ is at most $\epsilon(t, \lambda)$.

We say that $(\mathcal{P}, \mathcal{V})$ has *argument size* s if a proof π output by $\mathcal{P}^h(\mathbbm{x}, \mathbbm{w})$ consists of $s(|\mathbbm{x}|)$ bits.

We also consider non-interactive arguments that additionally achieve *proof of knowledge* and *zero knowledge*. The first property will hold against query-bounded adversaries (that are otherwise all-powerful), while the second property will hold against unbounded adversaries (and in particular need not refer to quantum algorithms). We define both of these properties below.

Knowledge. The non-interactive argument $(\mathcal{P}, \mathcal{V})$ is an *argument of knowledge* with extraction probability κ if there exists a polynomial-time quantum extractor $\mathcal{E}$ such that, for every instance $\mathbbm{x}$ and t-query quantum oracle algorithm $\tilde{\mathcal{P}}$, if, over a random oracle $h \leftarrow \mathcal{U}(2\lambda, \lambda)$, for $\pi := \tilde{\mathcal{P}}^h$ it holds that $\mathcal{V}^h(\mathbb{x}, \pi) = 1$ with probability μ, the probability that $\mathcal{E}^{\tilde{\mathcal{P}}}(\mathbb{x}, 1^t, 1^\lambda)$ outputs a valid witness for $\mathbb{x}$ is at least $\kappa(t, \mu, \lambda)$. Here the notation $\mathcal{E}^{\tilde{\mathcal{P}}}$ denotes that $\mathcal{E}$ has black-box access to $\tilde{\mathcal{P}}$ as defined by Unruh [60]. Informally, this means that if $\tilde{\mathcal{P}} = (A_1, \dots, A_t)$ with initial state $|\phi_0\rangle$, then $\mathcal{E}$ is given an auxiliary register containing $|\phi_0\rangle$ and may apply, in addition to any efficient quantum operation, any A_i to any of its registers.

Zero knowledge. The non-interactive argument $(\mathcal{P}, \mathcal{V})$ has (statistical) zero knowledge if there exists a probabilistic polynomial-time simulator $\mathcal{S}$ such that for every instance-witness pair $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$ the distributions below are statistically close (as a function of λ):

$$\left\{ (h, \pi) \;\middle|\; \begin{array}{c} h \leftarrow \mathcal{U}(2\lambda, \lambda) \\ \pi \leftarrow \mathcal{P}^h(\mathbb{x}, \mathbb{w}) \end{array} \right\} \quad \text{and} \quad \left\{ (h[\mu], \pi) \;\middle|\; \begin{array}{c} h \leftarrow \mathcal{U}(2\lambda, \lambda) \\ (\mu, \pi) \leftarrow \mathcal{S}^h(\mathbb{x}) \end{array} \right\} .$$

Above, $h[\mu]$ is the function that, on input x, equals $\mu(x)$ if μ is defined on x, or $h(x)$ otherwise. This definition uses explicitly-programmable random oracles [10]. (Non-interactive zero knowledge with non-programmable random oracles is impossible for non-trivial languages [52, 13].)

Succinctness for non-deterministic time. A zkSNARK for $\mathsf{NTIME}(T(n))$ in the QROM is a non-interactive argument for $\mathsf{NTIME}(T(n))$ in the QROM such that: (a) it has (statistical) zero knowledge; (b) it has extraction probability $\mathsf{poly}(\mu, 1/t) - \mathsf{poly}(\mu, t)/2^\lambda$; (c) arguments have size $\mathsf{poly}(\lambda, \log T(n))$, the prover runs in time $\mathsf{poly}(\lambda, n, T(n))$, and the verifier runs in time $\mathsf{poly}(\lambda, n, \log T(n))$.

3.4 Probabilistically checkable proofs

A *probabilistically checkable proof* (PCP) for a relation $\mathcal{R}$ with soundness error ϵ, proof length ℓ, and alphabet Σ is a pair of polynomial-time algorithms $(\mathbf{P}, \mathbf{V})$ for which the following holds.

- **Completeness.** For every instance-witness pair $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$, $\mathbf{P}(\mathbb{x}, \mathbb{w})$ outputs a proof string $\Pi \colon [\ell] \to \Sigma$ such that $\Pr\big[\mathbf{V}^\Pi(\mathbb{x}) = 1\big] = 1$.
- **Soundness.** For every instance $\mathbb{x} \notin \mathcal{L}(\mathcal{R})$ and proof string $\Pi \colon [\ell] \to \Sigma$, $\Pr\big[\mathbf{V}^\Pi(\mathbb{x}) = 1\big] \le \epsilon$.

The quantities ϵ, ℓ, Σ can be functions of the instance size $|\mathbb{x}|$. Probabilities are taken over the randomness of $\mathbf{V}$. The *randomness complexity* is the number of random bits used by $\mathbf{V}$, and the *query complexity* q is the number of locations of Π read by $\mathbf{V}$. (Both can be functions of $|\mathbb{x}|$.)

We also consider PCPs that achieve *proof of knowledge* and (honest-verifier) *zero knowledge*. We define both of these properties below.

Proof of knowledge. The PCP $(\mathbf{P}, \mathbf{V})$ has knowledge error k if there exists a polynomial-time extractor $\mathbf{E}$ such that for every instance $\mathbb{x}$ and proof string $\Pi \colon [\ell] \to \Sigma$ if $\Pr[\mathbf{V}(\mathbb{x}, \Pi) = 1] > \mathsf{k}$ then $\mathbf{E}(\mathbb{x}, \Pi)$ outputs a valid witness for $\mathbb{x}$.

Zero knowledge. The PCP $(\mathbf{P}, \mathbf{V})$ is (perfect) honest-verifier zero knowledge if there exists a probabilistic polynomial-time simulator $\mathbf{S}$ such that for every instance-witness pair $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$ the view of $\mathbf{V}(\mathbb{x})$ when given access to a proof string sampled as $\Pi \leftarrow \mathbf{P}(\mathbb{x}, \mathbb{w})$ equals the view of $\mathbf{V}(\mathbb{x})$ when given access to $\mathbf{S}(\mathbb{x})$. In the latter case, $\mathbf{S}(\mathbb{x})$ adaptively answers queries received from $\mathbf{V}(\mathbb{x})$.

3.5 Databases

A *database* mapping X to Y is a partial function $D \colon X \rightharpoonup Y$. The support of a database D is $\mathrm{supp}(D) := \{x \in X \colon D(x) \neq \bot\}$ and its image $\mathrm{im}(D)$ is $\{D(x) \colon x \in \mathrm{supp}(D)\}$. The size of a database is the size of its support: $|D| := |\mathrm{supp}(D)|$. Given two databases D and D', we write $D \subseteq D'$ if $\mathrm{supp}(D) \subseteq \mathrm{supp}(D')$ and $D(x) = D'(x)$ for every $x \in \mathrm{supp}(D)$.

We define two operations on databases, corresponding to deletions and insertions. Given a database D, input values $x, x' \in X$, and output value $y \in Y$, we define the two databases

$$(D - x)(x') := \begin{cases} \bot & \text{if } x = x' \\ D(x') & \text{if } x \neq x' \end{cases} \quad \text{and} \quad (D + [x \mapsto y])(x') := \begin{cases} y & \text{if } x = x' \\ D(x') & \text{if } x \neq x' \end{cases} .$$

For $D \colon \{0,1\}^m \rightharpoonup \{0,1\}^n$ and $t \in \mathbb{N}$ with $|D| \leq t \leq 2^m$, we define the pure quantum state

$$|D_t\rangle := \left|x_1, y_1, \ldots, x_{|D|}, y_{|D|}\right\rangle \otimes |\bot, 0^n\rangle^{\otimes(|D|-t)}$$

where $x_1, \ldots, x_{|D|}$ is the lexicographic ordering of $\mathrm{supp}(D)$ and $y_i := D(x_i)$ for each $i \in [|D|]$. We will write $|D\rangle$ for $|D_t\rangle$ when the bound t is clear from context.

3.6 Compressed phase oracle

The standard method to encode a function $h \colon \{0,1\}^m \to \{0,1\}^n$ as a quantum operation is the unitary matrix O_h defined in Section 3.2, which acts as $|x, y\rangle \mapsto |x, y \oplus h(x)\rangle$. Another method is to encode h in the *phase* of a quantum state, via the unitary matrix O'_h that acts as $|x, u\rangle \mapsto (-1)^{u \cdot h(x)} |x, u\rangle$. These two encodings are equivalent under an efficient change of basis: $O_h = (I^m \otimes H^n) O'_h (I^m \otimes H^n)$ where I^m is the identity on the first m qubits and H^n is the Hadamard transformation on the other n qubits. Thus, choosing between the *standard oracle* O_h or the *phase oracle* O'_h is a matter of convenience. For example, the Deutsch–Josza algorithm [25] is easier to describe with a standard oracle, while Grover's algorithm [38] is easier with a phase oracle.

In this paper it is more convenient to *always work with phase oracles*. All quantum query algorithms will thus have an oracle phase register U instead of the oracle answer register Y. Moreover, since h is sampled at random from the set of all functions from m bits to n bits, we follow Zhandry [66] and extend the adversary's initial state with a random superposition of all functions h, which represents a purification of the adversary's mixed state relative to the random oracle.

In fact, instead of considering a superposition of functions h, we will consider a superposition of databases D, according to the *compressed oracle* formalism of [66]. Specifically, throughout this paper we will only deal with the *compressed phase oracle* with m input bits and n output bits, which we denote by $\mathcal{O}$. We fix the database query bound of the compressed oracle to be t in advance. For the purposes of this paper, we will only use the fact that $\mathcal{O}$ is a certain unitary matrix, indistinguishable from a real random oracle, whose action is given by the following lemma. We refer the reader to [66] for more details.

Lemma 4 ([66]). *The compressed phase oracle $\mathcal{O}$ (with query bound t) acts on a quantum state $|x, u, z, D\rangle$, where $x \in \{0,1\}^m$, $u \in \{0,1\}^n$, $z \in \{0,1\}^*$, and $D\colon \{0,1\}^m \rightharpoonup \{0,1\}^n$ is a database with $|D| \leq t$, as follows.*

- *If $|D| = t$ or $u = 0^n$, then $\mathcal{O}\,|x, u, z, D\rangle = (-1)^{u \cdot D(x)}\,|x, u, z, D\rangle$, where $u \cdot \perp := 0$.*
- *If $D(x) = \perp$, $|D| < t$, and $u \neq 0^n$, then $\mathcal{O}\,|x, u, z, D\rangle = |x, u, z\rangle \otimes |\phi\rangle$ where*

$$|\phi\rangle \;:=\; \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{u \cdot y}\,|D + [x \mapsto y]\rangle \quad .$$

- *If $D(x) \neq \perp$, $|D| < t$, and $u \neq 0^n$, then $\mathcal{O}\,|x, u, z, D\rangle = |x, u, z\rangle \otimes |\phi\rangle$ where*

$$\begin{aligned}|\phi\rangle \;:=\; &(-1)^{u \cdot D(x)}\,|D\rangle + \frac{(-1)^{u \cdot D(x)}}{\sqrt{2^n}}\,|D - x\rangle \\ &+ \frac{1}{2^n} \sum_{y \in \{0,1\}^n} \Big(1 - (-1)^{u \cdot y} - (-1)^{u \cdot D(x)}\Big)\,|D - x + [x \mapsto y]\rangle \quad .\end{aligned}$$

Given a quantum algorithm $\mathcal{A}$ described by unitaries $A_1, \ldots, A_t$ and initial state $|\phi_0\rangle$, we write $|\mathsf{Sim}^*(\mathcal{A})\rangle$ to represent the final state of $\mathcal{A}$ before measurement when simulated using $\mathcal{O}$ as the oracle. Formally, $|\mathsf{Sim}^*(\mathcal{A})\rangle = A_t\mathcal{O} \cdots A_1\mathcal{O}|\phi_0, \emptyset\rangle$, where $\emptyset$ denotes that the D register holds the empty database with t slots, and we implicitly extend each A_i to act as the identity on D.

The following lemma of [66] shows simulating $\mathcal{A}$ by using $\mathcal{O}$ as the oracle is perfectly indistinguishable from running $\mathcal{A}$ with access to a random oracle.

Lemma 5 ([66, Lemma 4]). *For any quantum oracle algorithm $\mathcal{A}$ making at most t queries,*

$$\mathrm{Tr}_{\mathsf{D}}(|\mathsf{Sim}^*(\mathcal{A})\rangle\!\langle\mathsf{Sim}^*(\mathcal{A})|) = \frac{1}{(2^n)^{2^m}} \sum_{h\colon \{0,1\}^m \to \{0,1\}^n} |\mathcal{A}^h\rangle\!\langle\mathcal{A}^h| \quad .$$

I.e., $|\mathsf{Sim}^*(\mathcal{A})\rangle$ purifies the mixed state of $\mathcal{A}$ when interacting with a random oracle $h \leftarrow \mathcal{U}(m, n)$.

The notation Tr_{D} denotes the partial trace over the D (database) register, defined as the unique linear operator such that $\mathrm{Tr}_{\mathsf{D}}(|a\rangle\!\langle a|_{\mathsf{Z}} \otimes |b\rangle\!\langle b|_{\mathsf{D}}) := \langle b|b\rangle\,|a\rangle\!\langle a|_{\mathsf{Z}}$ for all vectors $|a\rangle, |b\rangle$. Here Z denotes all the registers of the adversary.

4 A lifting lemma for database games

In this section we show how to bound the value of a (classical or quantum) database game via the *instability* of the game, a purely classical quantity that we introduce in this paper. As we will see shortly, it is straightforward to argue that for any base game G (Section 2.4), the value $\omega_{\mathsf{D}}(G,t)$ is at most t times the instability of G. The goal of this section is to to prove that the (quantum) value $\omega^*_{\mathsf{D}}(G,t)$ is at most t^2 times the instability of G. In particular, we enable *lifting* a bound on the (classical) instability of G to a bound on the (quantum) value $\omega^*_{\mathsf{D}}(G,t)$. Combining the lifting lemma with the fact that oracle games can be generically reduced to database games (Lemma 1), we are able to establish the post-quantum security of the Micali construction solely by analyzing classical properties of it.

4.1 Database properties and the basic lifting lemma

A database property is a more general notion of a database game.

Definition 3. *A* **database property** $\mathcal{P}$ *is a set of databases* $D \colon X \rightharpoonup Y$*. The negation of* $\mathcal{P}$*, denoted* $\bar{\mathcal{P}}$*, is the set* $(X \rightharpoonup Y) \setminus \mathcal{P}$*.*

Given a base game, we define a corresponding database property as follows.

Definition 4. *The database property of a base game* $G \subseteq A^k \times B^k \times C$ *is*

$$\mathcal{P}_G := \{D : \exists\, (\mathbf{a},\mathbf{b},c) \in G \text{ with } D(a_i) = b_i \;\forall\, i \in [k]\} \enspace .$$

For a base game G, the database property $\mathcal{P}_G$ is closely related to the database game of G. This is because winning the database game is equivalent to the database outputted by $\mathsf{Sim}^*(\mathcal{A})$ being in $\mathcal{P}_G$. In particular, the following proposition holds.

Proposition 2. *For every base game* $G \subseteq A^k \times B^k \times C$ *and quantum algorithm* $\mathcal{A}$*,*

$$\Pr[\mathcal{A} \text{ wins } G^*_{\mathsf{D}}] = \Pr\left[D \in \mathcal{P}_G \;\middle|\; \big((\mathbf{a},\mathbf{b},c),D\big) \leftarrow \mathsf{Sim}^*(\mathcal{A})\right] \enspace .$$

We define the *flip probability* of a pair of database properties.

Definition 5. *The* **flip probability** $\mathrm{flip}(\mathcal{P} \to \mathcal{Q}, t)$ *from property* $\mathcal{P}$ *to property* $\mathcal{Q}$ *is the quantity*

$$\mathrm{flip}(\mathcal{P} \to \mathcal{Q}, t) := \max_{\substack{D \colon \{0,1\}^m \rightharpoonup \{0,1\}^n \\ |D|<t\,,\, D \in \mathcal{P}}} \; \max_{x \notin \mathrm{supp}(D)} \Pr_y\left[D + [x \mapsto y] \in \mathcal{Q}\right] \enspace ,$$

and $\mathrm{flip}(\emptyset \to \mathcal{Q}, t) := 0$.

Intuitively, this is the maximum probability over all databases $D \in \mathcal{P}$ with less than t entries that making an additional query puts $D \in \mathcal{Q}$. The following properties can be obtained easily from the above definition.

Proposition 3 (Properties of the flip probability). *Let $\mathcal{P}, \mathcal{P}', \mathcal{Q}, \mathcal{Q}'$ be database properties.*

(i) *If* $\mathcal{P} \subseteq \mathcal{P}'$ *and* $\mathcal{Q} \subseteq \mathcal{Q}'$ *then* $\operatorname{flip}(\mathcal{P} \to \mathcal{Q}) \leq \operatorname{flip}(\mathcal{P}' \to \mathcal{Q}')$.
(ii) $\operatorname{flip}(\mathcal{P} \cup \mathcal{P}' \to \mathcal{Q}) = \max\big(\operatorname{flip}(\mathcal{P} \to \mathcal{Q}), \operatorname{flip}(\mathcal{P}' \to \mathcal{Q})\big)$.
(iii) $\operatorname{flip}(\mathcal{P} \to \mathcal{Q} \cup \mathcal{Q}') \leq \operatorname{flip}(\mathcal{P} \to \mathcal{Q}) + \operatorname{flip}(\mathcal{P} \to \mathcal{Q}')$.

The instability of a database property is the following classical quantity.

Definition 6. *The* **instability** $\mathbf{I}(\mathcal{P}, t)$ *of a database property* $\mathcal{P}$ *with query bound* t *is the maximum probability that, for any database* D *containing less than* t *queries, making one additional (classical) query changes whether or not* D *has the property* $\mathcal{P}$*. Formally, we let*

$$\mathbf{I}(\mathcal{P}, t) := \max\{\operatorname{flip}(\bar{\mathcal{P}} \to \mathcal{P}, t), \operatorname{flip}(\mathcal{P} \to \bar{\mathcal{P}}, t)\} \ .$$

Note that instability is symmetric: $\mathbf{I}(\mathcal{P}, t) = \mathbf{I}(\bar{\mathcal{P}}, t)$. There is a direct argument that shows that $\omega_{\mathsf{D}}(G, t)$ is bounded by $t\mathbf{I}(\mathcal{P}_G, t)$.[6] Similarly, our basic lifting lemma shows that $\omega^*_{\mathsf{D}}(G, t)$ is bounded by the instability of the database property $\mathcal{P}_G$. Thus, it lifts a *classical* notion to prove a bound on the *quantum* value of a database game.

Lemma 6 (Basic lifting lemma). *For any base game* G*,*

$$\omega^*_{\mathsf{D}}(G, t) \leq t^2 \cdot 6\mathbf{I}(\mathcal{P}_G, t) \ .$$

Before we proceed to the proof of Lemma 6, we first introduce some quantum notation. Recall that we let $|\mathsf{Sim}^*(\mathcal{A})\rangle$ denote the final quantum state of the simulated adversary. Using the definition of measurement, we can express the probability that the final measured database D is in a database property $\mathcal{P}$ in terms of the state $|\mathsf{Sim}^*(\mathcal{A})\rangle$.

Proposition 4. *For every database property* $\mathcal{P}$ *and quantum adversary* $\mathcal{A}$*,*

$$\Pr\left[D \in \mathcal{P} \,\middle|\, \big((\mathbf{a}, \mathbf{b}, c), D\big) \leftarrow \mathsf{Sim}^*(\mathcal{A})\right] = \|P|\mathsf{Sim}^*(\mathcal{A})\rangle\|^2 \ ,$$

where $P := I \otimes \sum_{D \in \mathcal{P}} |D\rangle\!\langle D|$ *is the projector that maps all basis states of the form* $|x, u, z\rangle \otimes |D\rangle$ *to* 0 *if* $D \notin \mathcal{P}$*, and is otherwise the identity.*

We learn that in order to bound $\omega^*_{\mathsf{D}}(G, t)$ it suffices to bound $\|P_G\,|\mathsf{Sim}^*(\mathcal{A})\rangle\|$ for every $\mathcal{A} \in \mathcal{C}^*_t$.

Next, define $P_t := I \otimes \sum_{D : |D| \leq t} |D\rangle\!\langle D|$ to be the projector that maps all basis states of the form $|x, u, z\rangle \otimes |D\rangle$ to 0 if $|D| > t$, and is otherwise the identity.

The proof of Lemma 6 follows from two lemmas. The first lemma shows that $\|P\,|\mathsf{Sim}^*(\mathcal{A})\rangle\|$ is bounded by $t\|P(P_t \mathcal{O} P_t)\bar{P}\|$. Intuitively, this is because if P and

[6] Let $\mathcal{A}$ be a classical adversary, and let $\mathcal{A}_i$ be the adversary obtained by stopping $\mathcal{A}$ immediately before its i-th query. Then $|\Pr[\mathcal{A}_{i+1} \text{ wins } G_{\mathsf{D}}] - \Pr[\mathcal{A}_i \text{ wins } G_{\mathsf{D}}]| \leq \mathbf{I}(\mathcal{P}, t)$ holds for each $i \in [t]$ by definition of instability, and $\Pr[\mathcal{A}_1 \text{ wins } G_{\mathsf{D}}] = 0$ since $\emptyset \notin \mathcal{P}_G$. Therefore, $\Pr[\mathcal{A} \text{ wins } G_{\mathsf{D}}] \leq t\mathbf{I}(\mathcal{P}, t)$.

$P_t \mathcal{O} P_t$ almost commute (i.e., P and $\mathcal{O}$ almost commute when acting on databases with at most t entries) then each oracle query cannot change the probability that the database is in $\mathcal{P}$ by too much. The second lemma shows that $\|P(P_t \mathcal{O} P_t)\bar{P}\|^2$ is bounded by $\mathbf{I}(\mathcal{P}, t)$. Combining the two lemmas with Proposition 4 completes the proof of Lemma 6.

Lemma 7. *Let $\mathcal{P}$ be a database property with $\emptyset \notin \mathcal{P}$. For every $\mathcal{A} \in \mathcal{C}_t^*$,*

$$\|P|\mathsf{Sim}^*(\mathcal{A})\rangle\| \leq t \cdot \|P(P_t \mathcal{O} P_t)\bar{P}\| \ .$$

Lemma 8. *For any database property $\mathcal{P}$,*

$$\|P(P_t \mathcal{O} P_t)\bar{P}\|^2 \leq 6\mathbf{I}(\mathcal{P}, t) \ .$$

Lemmas 7 and 8 strengthen the proof sketch outlined in Section 2.5. This is because for any operator A and projector P, $[P, A] = PA - AP = (PAP + PA\bar{P}) - (PAP + \bar{P}AP) = PA\bar{P} - \bar{P}AP$, and so $\|[P, A]\|^2 = \|PA\bar{P}\|^2 + \|\bar{P}AP\|^2$. Hence, Lemma 7 implies that $\|P|\mathsf{Sim}^*(\mathcal{A})\rangle\| \leq t \cdot \|[P, P_t \mathcal{O} P_t]\|$ and Lemma 8 implies that $\|[P, P_t \mathcal{O} P_t]\|^2 \leq 12\mathbf{I}(\mathcal{P}, t)$.

We now prove Lemma 7; the proof of Lemma 8 can be found in the full version of the paper.

Proof (Proof of Lemma 7). Recall that the quantum algorithm $\mathcal{A}$ is described by some unitaries $(A_1, \ldots, A_t)$ and initial state $|\phi_0\rangle$. We can thus describe the quantum algorithm $\mathsf{Sim}^*(\mathcal{A})$ via the cumulative unitary $U := A_t \mathcal{O} A_{t-1} \cdots \mathcal{O} A_1 \mathcal{O}$ acting on the initial state $|\phi_0, \emptyset\rangle$ where $\emptyset$ denotes the empty database. (We abuse notation and implicitly extend A_i to act as the identity on the database register.) The final state is $|\mathsf{Sim}^*(\mathcal{A})\rangle := U|\phi_0, \emptyset\rangle$.

Let $U' := A_t(P_t \mathcal{O} P_t)A_{t-1} \cdots (P_t \mathcal{O} P_t)A_1(P_t \mathcal{O} P_t)$. We have that $U'|\phi_0, \emptyset\rangle = U|\phi_0, \emptyset\rangle$, as applying each P_t has no effect, since the database can only have at most t queries when P_t is applied.

For any operators $C_1, \ldots, C_t$ and projector P, we have that

$$C_t \cdots C_1 = \bar{P} C_t \bar{P} C_{t-1} \bar{P} \cdots C_1 \bar{P} + \sum_{i=0}^{t} (C_t \cdots C_{i+1}) \cdot P \cdot (C_i \bar{P} \cdots C_1 \bar{P}) \ . \quad (1)$$

To see this, we observe that

$$C_t \cdots C_1 = (C_t \cdots C_2)(C_1 \bar{P}) + (C_t \cdots C_1) \cdot P \ ,$$

which implies Eq. (1) by induction.

Let $C_i = A_i(P_t \mathcal{O} P_t)$. Then we have that

$$\begin{aligned}
&\|P\,|\mathsf{Sim}^*(\mathcal{A})\rangle\| = \|PU'|\phi_0, \emptyset\rangle\| \\
&= \|\Big(P\bar{P}C_t\bar{P}C_{t-1}\bar{P}\cdots C_1\bar{P} + \sum_{i=0}^{t} P(C_t\cdots C_{i+1})\cdot P\cdot(C_i\bar{P}\cdots C_1\bar{P})\Big)|\phi_0, \emptyset\rangle\| \\
&\leq \sum_{i=0}^{t} \|P(C_t\cdots C_{i+1})\cdot P\cdot(C_i\bar{P}\cdots C_1\bar{P})|\phi_0, \emptyset\rangle\| \\
&\leq \|P(C_t\cdots C_1)\cdot P|\phi_0, \emptyset\rangle\| + \sum_{i=1}^{t} \|P(C_t\cdots C_{i+1})\| \cdot \|P\cdot(C_i\bar{P}\cdots C_1\bar{P})|\phi_0, \emptyset\rangle\| \\
&\leq 0 + \sum_{i=1}^{t} \|PC_i\bar{P}\| \cdot \|(C_i\bar{P}\cdots C_1\bar{P})|\phi_0, \emptyset\rangle\| \\
&\leq \sum_{i=1}^{t} \|PA_i(P_t\mathcal{O}P_t)\bar{P}\| \ ,
\end{aligned}$$

where we use the fact that the operator norm of a product of unitaries/projectors is at most 1, and that $\emptyset \notin \mathcal{P}$. Since P and A_i commute for every i, we get that $\|PA_i(P_t\mathcal{O}P_t)\bar{P}\| = \|A_iP(P_t\mathcal{O}P_t)\bar{P}\| \leq \|A_i\|\|P(P_t\mathcal{O}P_t)\bar{P}\| = \|P(P_t\mathcal{O}P_t)\bar{P}\|$. Hence, $\|P\,|\mathsf{Sim}^*(\mathcal{A})\rangle\| \leq t\|P(P_t\mathcal{O}P_t)\bar{P}\|$.

4.2 Conditional instability and the lifting lemma

Lemma 6 is not quite sufficient to analyze the database game that corresponds to the Micali construction. In fact, the instability of this game is high because we take a maximum over all bounded databases, including those which contain collisions. If we were to only take the maximum over databases that do not contain collisions, then the instability would be low. Moreover, the instability of the "no collision" property is itself low.

In this section, we strengthen the results of the previous section by introducing the notion of *conditional* instability, which allows us to analyze the value $\omega^*_{\mathsf{D}}(G, t)$ by splitting its database property $\mathcal{P}_G$ into subproperties and analyzing the subproperties separately, analogous to conditioning in probability. In particular, we can then analyze the Micali game by analyzing the no collision property and the instability of the Micali database property conditioned on the no collision property.

For the entirety of this section we will let $\mathcal{P}$ and $\mathcal{Q}$ be database properties, and we will analyze quantities about $\mathcal{P}$ conditioned on $\mathcal{Q}$. These results strengthen the results of Section 4.1, as the previous results can be recovered by setting $\mathcal{Q}$ to be the database property containing all databases.

Definition 7. *Let $\mathcal{P}$ and $\mathcal{Q}$ be two database properties, and let t be a query bound. We define*

$$\mathrm{flip}(\mathcal{P} \mid \mathcal{Q}, t) := \mathrm{flip}(\bar{\mathcal{P}} \cap \mathcal{Q} \to \mathcal{P} \cap \mathcal{Q}, t) \ .$$

The **conditional instability** $\mathbf{I}(\mathcal{P} \mid \mathcal{Q}, t)$ *is defined as*

$$\mathbf{I}(\mathcal{P} \mid \mathcal{Q}, t) := \max\{\mathrm{flip}(\mathcal{P} \mid \mathcal{Q}, t),\ \mathrm{flip}(\bar{\mathcal{P}} \mid \mathcal{Q}, t)\} \ .$$

Before we state the lifting lemma, we observe the following properties of instability.

Proposition 5. *Let $\mathcal{P}$ and $\mathcal{Q}$ be two database properties. Then*

1. $\mathbf{I}(\mathcal{P}, t)$ *and* $\mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t)$ *are incomparable.*
2. $\mathrm{flip}(\mathcal{P} \mid \mathcal{Q}, t) \leq \mathrm{flip}(\bar{\mathcal{P}} \to \mathcal{P}, t)$, *and therefore* $\mathbf{I}(\mathcal{P} \mid \mathcal{Q}, t) \leq \mathbf{I}(\mathcal{P}, t)$.
3. $\mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t) \leq \mathbf{I}(\mathcal{P} \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)$.

Proof. To show Item 1, we give database properties $\mathcal{P}, \mathcal{Q}$ such that $\mathbf{I}(\mathcal{P}, t) > \mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t)$ and properties $\mathcal{P}', \mathcal{Q}'$ such that $\mathbf{I}(\mathcal{P}', t) < \mathbf{I}(\mathcal{P}' \cup \mathcal{Q}', t)$. Let $\mathcal{P}$ be the property that $D \neq \emptyset$. Then clearly $\mathbf{I}(\mathcal{P}, t) \geq \mathrm{flip}(\bar{\mathcal{P}} \to \mathcal{P}, t) = 1$. Let $\mathcal{Q}$ be the property that $D = \emptyset$. Now $\mathcal{P} \cup \mathcal{Q}$ is the set of all databases, so $\mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t) = 0$.

On the other hand, let $\mathcal{P}' = \emptyset$ be the empty property, and let $\mathcal{Q}'$ be the property that $D = \emptyset$. Then, $\mathbf{I}(\mathcal{P}', t) = 0$, and $\mathbf{I}(\mathcal{P}' \cup \mathcal{Q}', t) = \mathbf{I}(\mathcal{Q}', t) = 1$.

Item 2 holds since

$$\mathrm{flip}(\mathcal{P} \mid \mathcal{Q}, t) = \mathrm{flip}(\bar{\mathcal{P}} \cap \mathcal{Q} \to \mathcal{P} \cap \mathcal{Q}, t) \leq \mathrm{flip}(\bar{\mathcal{P}} \to \mathcal{P}, t) \ .$$

Finally, for Item 3 we observe that

$$\begin{aligned} \mathrm{flip}(\overline{\mathcal{P} \cup \mathcal{Q}} \to \mathcal{P} \cup \mathcal{Q}, t) &= \mathrm{flip}(\bar{\mathcal{P}} \cap \bar{\mathcal{Q}} \to \mathcal{P} \cup \mathcal{Q}, t) \\ &\leq \mathrm{flip}(\bar{\mathcal{P}} \cap \bar{\mathcal{Q}} \to \mathcal{P} \cap \bar{\mathcal{Q}}, t) + \mathrm{flip}(\bar{\mathcal{P}} \cap \bar{\mathcal{Q}} \to \mathcal{Q}, t) \\ &\leq \mathrm{flip}(\mathcal{P} \mid \bar{\mathcal{Q}}, t) + \mathrm{flip}(\bar{\mathcal{Q}} \to \mathcal{Q}, t) \ . \end{aligned}$$

On the other hand,

$$\begin{aligned} \mathrm{flip}(\mathcal{P} \cup \mathcal{Q} \to \overline{\mathcal{P} \cup \mathcal{Q}}, t) &= \mathrm{flip}(\mathcal{P} \cup \mathcal{Q} \to \bar{\mathcal{P}} \cap \bar{\mathcal{Q}}, t) \\ &= \max(\mathrm{flip}(\mathcal{P} \cap \bar{\mathcal{Q}} \to \bar{\mathcal{P}} \cap \bar{\mathcal{Q}}, t), \mathrm{flip}(\mathcal{Q} \to \bar{\mathcal{P}} \cap \bar{\mathcal{Q}}, t)) \\ &\leq \max(\mathrm{flip}(\bar{\mathcal{P}} \mid \bar{\mathcal{Q}}, t), \mathrm{flip}(\mathcal{Q} \to \bar{\mathcal{Q}}, t)) \ . \end{aligned}$$

Therefore, we get that $\mathbf{I}(\mathcal{P} \cup \mathcal{Q}) \leq \mathbf{I}(\mathcal{P} \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)$.

We now state the lifting lemma.

Lemma 9 (Lifting lemma). *Let G be a base game. Then for any database property $\mathcal{Q}$,*

$$\omega^*_{\mathsf{D}}(G, t) \leq t^2 \cdot 6 \left(\mathbf{I}(\mathcal{P}_G \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)\right) \ .$$

Proof. Let $\mathcal{P}$ and $\mathcal{Q}$ be two database properties. We show that for every $\mathcal{A} \in \mathcal{C}^*_t$ it holds that

$$\|P|\mathsf{Sim}^*(\mathcal{A})\rangle\|^2 \leq t^2 \cdot 6 \left(\mathbf{I}(\mathcal{P} \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)\right) \ .$$

Let $\mathcal{R} = \mathcal{P} \cup \mathcal{Q}$. Then by Lemmas 7 and 8 we have that

$$\|P|\mathsf{Sim}^*(\mathcal{A})\rangle\|^2 \leq \|R\,|\mathsf{Sim}^*(\mathcal{A})\rangle\|^2 \leq t^2 \cdot \|[R, P_t \mathcal{O} P_t]\|^2 \leq t^2 \cdot 6\mathbf{I}(\mathcal{R}, t) \ ,$$

where the first inequality holds since $\mathcal{P} \subseteq \mathcal{R}$. Finally, we use the fact that $\mathbf{I}(\mathcal{R}, t) = \mathbf{I}(\mathcal{P} \cup \mathcal{Q}, t) \leq \mathbf{I}(\mathcal{P} \mid \bar{\mathcal{Q}}, t) + \mathbf{I}(\mathcal{Q}, t)$, which completes the proof.

Acknowledgments

We thank Chinmay Nirkhe for taking part in early stages of this research, and for providing valuable feedback. This research was supported in part by: a Google Faculty Award; the UC Berkeley Center for Long-Term Cybersecurity; the NSF Graduate Research Fellowship Program; the ARCS Foundation; and donations from the Ethereum Foundation and the Interchain Foundation.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1745016. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Aaronson, S., Shi, Y.: Quantum lower bounds for the collision and the element distinctness problems. Journal of the ACM **51**(4), 595–605 (2004)
2. Ambainis, A., Rosmanis, A., Unruh, D.: Quantum attacks on classical proof systems: The hardness of quantum rewinding. In: Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science. pp. 474–483. FOCS '14 (2014)
3. Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Ligero: Lightweight sublinear arguments without a trusted setup. In: Proceedings of the 24th ACM Conference on Computer and Communications Security. pp. 2087–2104. CCS '17 (2017)
4. Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M.: Proof verification and the hardness of approximation problems. Journal of the ACM **45**(3), 501–555 (1998), preliminary version in FOCS '92.
5. Arora, S., Safra, S.: Probabilistic checking of proofs: a new characterization of NP. Journal of the ACM **45**(1), 70–122 (1998), preliminary version in FOCS '92.
6. Babai, L.: Trading group theory for randomness. In: Proceedings of the 17th Annual ACM Symposium on Theory of Computing. pp. 421–429. STOC '85 (1985)
7. Babai, L., Fortnow, L., Levin, L.A., Szegedy, M.: Checking computations in polylogarithmic time. In: Proceedings of the 23rd Annual ACM Symposium on Theory of Computing. pp. 21–32. STOC '91 (1991)
8. Baum, C., Bootle, J., Cerulli, A., Pino, R.d., Groth, J., Lyubashevsky, V.: Sublinear lattice-based zero-knowledge arguments for arithmetic circuits. In: Proceedings of the 38th Annual International Cryptology Conference. pp. 669–699. CRYPTO '18 (2018)
9. Baum, C., Nof, A.: Concretely-efficient zero-knowledge arguments for arithmetic circuits and their application to lattice-based cryptography. Cryptology ePrint Archive, Report 2019/532 (2019)
10. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security. pp. 62–73. CCS '93 (1993)
11. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable zero knowledge with no trusted setup. In: Proceedings of the 39th Annual International Cryptology Conference. pp. 733–764. CRYPTO '19 (2019)

12. Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., Ward, N.P.: Aurora: Transparent succinct arguments for R1CS. In: Proceedings of the 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 103–128. EUROCRYPT '19 (2019), full version available at `https://eprint.iacr.org/2018/828`
13. Ben-Sasson, E., Chiesa, A., Spooner, N.: Interactive oracle proofs. In: Proceedings of the 14th Theory of Cryptography Conference. pp. 31–60. TCC '16-B (2016)
14. Bitansky, N., Chiesa, A., Ishai, Y., Ostrovsky, R., Paneth, O.: Succinct non-interactive arguments via linear interactive proofs. In: Proceedings of the 10th Theory of Cryptography Conference. pp. 315–333. TCC '13 (2013)
15. Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Proceedings of the 17th International Conference on the Theory and Application of Cryptology and Information Security. pp. 41–69. ASIACRYPT '11 (2011)
16. Boneh, D., Ishai, Y., Sahai, A., Wu, D.J.: Lattice-based SNARGs and their application to more efficient obfuscation. In: Proceedings of the 36th Annual International Conference on Theory and Applications of Cryptographic Techniques. pp. 247–277. EUROCRYPT '17 (2017)
17. Boneh, D., Ishai, Y., Sahai, A., Wu, D.J.: Quasi-optimal SNARGs via linear multi-prover interactive proofs. In: Proceedings of the 37th Annual International Conference on Theory and Application of Cryptographic Techniques. pp. 222–255. EUROCRYPT '18 (2018)
18. Bowe, S.: bellman: a zk-snark library (2015), `https://github.com/zkcrypto/bellman`
19. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. Journal of Computer and System Sciences **37**(2), 156–189 (1988)
20. Brassard, G., Høyer, P., Tapp, A.: Quantum cryptanalysis of hash and claw-free functions. In: Proceedings of the 3rd Latin American Symposium on Theoretical Informatics. pp. 163–169. LATIN '98 (1998)
21. Canetti, R., Chen, Y., Holmgren, J., Lombardi, A., Rothblum, G.N., Rothblum, R.D.: Fiat–Shamir from simpler assumptions. Cryptology ePrint Archive, Report 2018/1004 (2018)
22. Chase, M., Derler, D., Goldfeder, S., Orlandi, C., Ramacher, S., Rechberger, C., Slamanig, D., Zaverucha, G.: Post-quantum zero-knowledge and signatures from symmetric-key primitives. In: Proceedings of the 24th ACM Conference on Computer and Communications Security. pp. 1825–1842. CCS '17 (2017)
23. Dagdelen, Ö., Fischlin, M., Gagliardoni, T.: The Fiat-Shamir transformation in a quantum world. In: Proceedings of the 19th International Conference on the Theory and Application of Cryptology and Information Security. pp. 62–81. ASIACRYPT '13 (2013)
24. dalek cryptography: A pure-Rust implementation of Bulletproofs using Ristretto (2018), `https://github.com/dalek-cryptography/bulletproofs`
25. Deutsch, D., Jozsa, R.: Rapid solution of problems by quantum computation. Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**(1907), 553–558 (1992)
26. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Security of the Fiat–Shamir transformation in the quantum random-oracle model. In: Proceedings of the 39th Annual International Cryptology Conference. pp. 356–383. CRYPTO '19 (2019)
27. Dwork, C., Feige, U., Kilian, J., Naor, M., Safra, S.: Low communication 2-prover zero-knowledge proofs for NP. In: Proceedings of the 11th Annual International Cryptology Conference. pp. 215–227. CRYPTO '92 (1992)

28. Eaton, E.: Leighton-Micali hash-based signatures in the quantum random-oracle model. In: Proceedings of the 24th International Conference on Selected Areas in Cryptography. pp. 263–280. SAC '17 (2017)
29. Electric Coin Company: Zcash Cryptocurrency (2014), `https://z.cash/`
30. Feige, U., Goldwasser, S., Lovász, L., Safra, S., Szegedy, M.: Interactive proofs and the hardness of approximating cliques. Journal of the ACM **43**(2), 268–292 (1996), preliminary version in FOCS '91.
31. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Proceedings of the 6th Annual International Cryptology Conference. pp. 186–194. CRYPTO '86 (1986)
32. Gennaro, R., Minelli, M., Nitulescu, A., Orrù, M.: Lattice-based zk-SNARKs from square span programs. In: Proceedings of the 25th ACM Conference on Computer and Communications Security. pp. 556–573. CCS '18 (2018)
33. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing. pp. 99–108. STOC '11 (2011)
34. Goldreich, O., Håstad, J.: On the complexity of interactive proofs with bounded communication. Information Processing Letters **67**(4), 205–214 (1998)
35. Goldreich, O., Vadhan, S., Wigderson, A.: On interactive proofs with a laconic prover. Computational Complexity **11**(1/2), 1–53 (2002)
36. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: Delegating computation: Interactive proofs for muggles. Journal of the ACM **62**(4), 27:1–27:64 (2015)
37. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM Journal on Computing **18**(1), 186–208 (1989), preliminary version appeared in STOC '85.
38. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the 28th Annual ACM Symposium on Theory of Computing. pp. 212–219. STOC '96 (1996)
39. iden3: websnark: A fast zkSNARK proof generator written in native web assembly. (2019), `https://github.com/iden3/websnark`
40. Ishai, Y., Mahmoody, M., Sahai, A., Xiao, D.: On zero-knowledge PCPs: Limitations, simplifications, and applications (2015), available at `http://www.cs.virginia.edu/~mohammad/files/papers/ZKPCPs-Full.pdf`
41. Kalai, Y., Raz, R.: Interactive PCP. In: Proceedings of the 35th International Colloquium on Automata, Languages and Programming. pp. 536–547. ICALP '08 (2008)
42. Katz, J., Kolesnikov, V., Wang, X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. In: Proceedings of the 25th ACM Conference on Computer and Communications Security. pp. 525–537. CCS '18 (2018)
43. Kilian, J.: A note on efficient zero-knowledge proofs and arguments. In: Proceedings of the 24th Annual ACM Symposium on Theory of Computing. pp. 723–732. STOC '92 (1992)
44. Kilian, J., Petrank, E., Tardos, G.: Probabilistically checkable proofs with zero knowledge. In: Proceedings of the 29th Annual ACM Symposium on Theory of Computing. pp. 496–505. STOC '97 (1997)
45. Kiltz, E., Lyubashevsky, V., Schaffner, C.: A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. In: Proceedings of the 37th Annual International Conference on Theory and Application of Cryptographic Techniques. pp. 552–586. EUROCRYPT '17 (2018)
46. libstark: libstark: a C++ library for zkSTARK systems (2018), `https://github.com/elibensasson/libSTARK`

47. Liu, Q., Zhandry, M.: Revisiting post-quantum Fiat–Shamir. In: Proceedings of the 39th Annual International Cryptology Conference. pp. 326–355. CRYPTO '19 (2019)
48. Lund, C., Fortnow, L., Karloff, H.J., Nisan, N.: Algebraic methods for interactive proof systems. Journal of the ACM **39**(4), 859–868 (1992)
49. Micali, S.: Computationally sound proofs. SIAM Journal on Computing **30**(4), 1253–1298 (2000), preliminary version appeared in FOCS '94.
50. NIST: Post-quantum cryptography (2016), `https://csrc.nist.gov/Projects/Post-Quantum-Cryptography`
51. O(1) Labs: Coda Cryptocurrency (2017), `https://codaprotocol.com/`
52. Pass, R.: On deniability in the common reference string and random oracle model. In: Proceedings of the 23rd Annual International Cryptology Conference. pp. 316–337. CRYPTO '03 (2003)
53. Pointcheval, D., Stern, J.: Security proofs for signature schemes. In: Proceedings of the 14th Annual International Conference on Theory and Application of Cryptographic Techniques. pp. 387–398. EUROCRYPT '96 (1996)
54. Reingold, O., Rothblum, R., Rothblum, G.: Constant-round interactive proofs for delegating computation. In: Proceedings of the 48th ACM Symposium on the Theory of Computing. pp. 49–62. STOC '16 (2016)
55. SCIPR Lab: libsnark: a C++ library for zkSNARK proofs (2014), `https://github.com/scipr-lab/libsnark`
56. SCIPR Lab: Dizk: Java library for distributed zero knowledge proof systems (2018), `https://github.com/scipr-lab/dizk`
57. SCIPR Lab: libiop: C++ library for IOP-based zkSNARKs (2019), `https://github.com/scipr-lab/libiop`
58. Targhi, E.E., Unruh, D.: Post-quantum security of the Fujisaki–Okamoto and OAEP transforms. In: Proceedings of the 14th Theory of Cryptography Conference. pp. 192–216. TCC '16-B (2016)
59. Unruh, D.: Non-interactive zero-knowledge proofs in the quantum random oracle model. In: Proceedings of the 34th Annual International Conference on Theory and Application of Cryptographic Techniques. pp. 755–784. EUROCRYPT '15 (2015)
60. Unruh, D.: Post-quantum security of Fiat-Shamir. In: Proceedings of the 23rd International Conference on the Theory and Applications of Cryptology and Information Security. pp. 65–95. ASIACRYPT '17 (2017)
61. Valiant, P.: Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In: Proceedings of the 5th Theory of Cryptography Conference. pp. 1–18. TCC '08 (2008)
62. Walfish, M., Blumberg, A.J.: Verifying computations without reexecuting them. Communications of the ACM **58**(2), 74–84 (Jan 2015)
63. Watrous, J.: Zero-knowledge against quantum attacks. SIAM Journal on Computing **39**(1), 25–58 (2009), preliminary version appeared in STOC '06.
64. Zhandry, M.: Secure identity-based encryption in the quantum random oracle model. In: Proceedings of the 32nd Annual International Cryptology Conference. pp. 758–775. CRYPTO '12 (2012)
65. Zhandry, M.: A note on the quantum collision and set equality problems. Quantum Information & Computation **15**(7&8), 557–567 (2015)
66. Zhandry, M.: How to record quantum queries, and applications to quantum indifferentiability. In: Proceedings of the 39th Annual International Cryptology Conference. pp. 239–268. CRYPTO '19 (2019)
67. ZKP Standards: Zero knowledge proof standardization (2017), `https://zkproof.org/`