Single-Round Proofs of Quantumness from Knowledge Assumptions

Petia Arabadjieva
Institute for Theoretical Physics, ETH Zurich, Switzerland

Alexandru Gheorghiu
Department of Computer Science and Engineering, Chalmers University of Technology, Göteborg, Sweden

Victor Gitton
Institute for Theoretical Physics, ETH Zurich, Switzerland

Tony Metger
Institute for Theoretical Physics, ETH Zurich, Switzerland

Abstract

A *proof of quantumness* is an efficiently verifiable interactive test that an efficient quantum computer can pass, but all efficient classical computers cannot (under some cryptographic assumption). Such protocols play a crucial role in the certification of quantum devices. Existing single-round protocols based solely on a cryptographic hardness assumption (like asking the quantum computer to factor a large number) require large quantum circuits, whereas multi-round ones use smaller circuits but require experimentally challenging mid-circuit measurements.

In this work, we construct efficient single-round proofs of quantumness based on existing *knowledge assumptions*. While knowledge assumptions have not been previously considered in this context, we show that they provide a natural basis for separating classical and quantum computation. Our work also helps in understanding the interplay between black-box/white-box reductions and cryptographic assumptions in the design of proofs of quantumness. Specifically, we show that multi-round protocols based on Decisional Diffie-Hellman (DDH) or Learning With Errors (LWE) can be "compiled" into single-round protocols using a knowledge-of-exponent assumption [7] or knowledge-of-lattice-point assumption [36], respectively. We also prove an *adaptive hardcore-bit* statement for a family of claw-free functions based on DDH, which might be of independent interest.

2012 ACM Subject Classification Theory of computation → Cryptographic protocols

Keywords and phrases Proofs of quantumness, Knowledge assumptions, Learning with errors, Decisional Diffie-Hellman

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.8

Related Version *Full Version*: `https://arxiv.org/abs/2405.15736` [5]

Funding P.A., V.G. and T.M. acknowledge support from the ETH Zurich Quantum Center and the Air Force Office for Scientific Research, grant No. FA9550-19-1-0202.
Alexandru Gheorghiu: Knut and Alice Wallenberg Foundation, through the Wallenberg Centre for Quantum Technology (WACQT).
Tony Metger: ETH Doc.Mobility Fellowship.

Acknowledgements We thank Alex Lombardi, Urmila Mahadev, Greg Kahanamoku-Meyer, Umesh Vazirani, John Wright, and Tina Zhang for helpful discussions. We are especially grateful to Vinod Vaikuntanathan to suggesting many of these ideas in the early stages of the project.

1 Introduction

Demonstrating quantum advantage, the point where a quantum computer can solve a problem that no existing classical computer can, is both a theoretical and technological challenge. It requires a problem that is plausibly intractable for classical algorithms and which admits an

efficient quantum algorithm, ideally one that can be performed with noisy intermediate-scale quantum (NISQ) devices [42]. Additionally, the problem's solution should be efficiently verifiable by a classical computer. This is necessary if one wishes to have a convincing and scalable way of proving quantum advantage. Currently, there are three main paradigms for demonstrating quantum advantage.

The most straightforward one is to solve a problem which is believed to be classically hard, such as integer factorization. A quantum computer could efficiently solve this task by running Shor's algorithm [43] and the solution can be efficiently verified by multiplying the factors reported by the prover. But while Shor's algorithm is efficient in the sense of only requiring a polynomial-size quantum circuit, the actual circuit for any reasonably-sized integer to be factored is too large to implement with NISQ devices [22, 25]. Another approach is based on the classical hardness of sampling problems such as *random circuit sampling* [12, 6, 44] or *boson sampling* [1, 48, 37]. While these experiments can be implemented with current hardware, they are not efficiently verifiable.

The work of Brakerski et al. [13] introduced a new approach towards testing for quantum advantage, referred to as a *proof of quantumness* protocol. Akin to the cryptographic notions of proof or argument systems [24, 15, 23], a proof of quantumness is an interactive protocol between a polynomial-time classical *verifier* and an ostensibly quantum polynomial-time *prover*. The verifier issues challenges to the prover and checks the correctness of the prover's responses. The key feature of such a protocol is that there should exist an efficient quantum strategy that allows the prover to correctly answer the verifier's challenges with high probability, whereas any efficient classical strategy can only succeed with low probability (under some plausible cryptographic hardness assumption such as the classical intractability of factoring, Decisional Diffie-Hellman (DDH), or the Learning With Errors (LWE) problem).

In contrast to other paradigms for proving quantum advantage, the cryptographic proofs of quantumness of Brakerski et al. and follow-up works do not require the quantum prover to break the underlying computational hardness assumption. Instead, they leverage the fact that the restriction imposed on the prover through mid-protocol interaction limits a classical prover's capacity for correctly responding to subsequent challenges from the verifier more strongly than it limits a quantum prover.

The advantage of these protocols over an integer factorization-based test is that they have much smaller circuits than Shor's algorithm, making them potentially more suitable for implementation on near-term devices [30, 29, 49, 26, 35, 4]. However, due to their interactive nature, they require the honest quantum prover to perform mid-circuit measurements for each of the verifier's challenges. Mid-circuit measurements on a subset of qubits are difficult to implement on existing quantum devices without disturbing neighboring qubits and can thus degrade the quality of the remaining computation. An experimental implementation of the two-round protocol by Brakerski et al. [49] directly compared the performance of an ion-trap quantum computer running the protocol with and without mid-circuit measurements, revealing a significant difference. Thus, implementing this proof of quantumness protocol at the scale required for quantum advantage seems especially challenging.[1] It would therefore be desirable to have proofs of quantumness with relatively small quantum circuits and which involve only one round[2] of interaction between the verifier and the prover. Beyond the practical motivation, this would also give us a better understanding of the structure required for demonstrating quantum advantage in a way that is efficiently verifiable.

[1] Note that the instance sizes of the underlying cryptographic problem in [49] are so small that they can easily be broken by a laptop. For a convincing demonstration of quantum advantage, one would have to use instance sizes that cannot be broken even by the fastest classical supercomputers.

[2] Throughout the paper, we use the convention that a one-round protocol refers to a protocol with two messages, one message from the verifier to the prover and one message back.

The recent breakthrough work of Yamakawa and Zhandry made progress in this direction by giving a single-round proof of quantumness protocol in the *random oracle model* (ROM) [45]. Prior to their work, Brakerski et al. constructed single-round proofs of quantumness in the ROM that additionally required a structured computational assumption, such as factoring or LWE [14]. However, with both of these approaches the circuits that the prover would have to perform are possibly larger than those in existing multi-round proofs of quantumness [13, 29].

Our main result is a single-round proof of quantumness that only requires the same small circuits as an existing multi-round interactive proof of quantumness and is based on *knowledge assumptions*, a cryptographic idea that turns out to be natural for proofs of quantumness. We achieve this by starting from the two-round protocol of [13] and removing one of the rounds of interaction through the use of a knowledge assumption [17, 40, 16, 36, 32]. As explained in [32], a knowledge assumption is a statement of the form:

> "If an algorithm $\mathcal{A}$ outputs an object of type X, it must know a corresponding witness of the type W, such that the output and the witness are in some relation $\mathcal{R} \subseteq X \times W$."

The rationale behind knowledge assumptions is that certain computational tasks, performed by some probabilistic algorithm $\mathcal{A}$, can only be performed efficiently by following a specific sequence of steps, thus obtaining a series of intermediate values. Informally, we say that $\mathcal{A}$ must have "known" the intermediate values for its specific output. This is made more precise by saying that there exists an efficient *extractor* $\mathcal{A}^*$ that receives as input $\mathcal{A}$'s random coins and outputs (or *extracts*) the relevant intermediate values of $\mathcal{A}$.

Knowledge assumptions have traditionally been used for the design of protocols that require both extractability (like in proof/argument of knowledge protocols) and succinctness [9, 32]. In our case, we are able to leverage knowledge assumptions to construct efficient, single-round proofs of quantumness. As far as we are aware, this is the first time knowledge assumptions are used in this context. We argue that their application here is natural and in some sense necessary, if one wishes to avoid multi-round interaction or working in the ROM. Intuitively, this is because some aspect of the proof of quantumness protocol has to differentiate between classical and quantum provers. In multi-round protocols, this distinction comes from the fact that classical provers can be *rewound*, but quantum provers generally cannot. In a ROM-based protocol, the distinction arises from the ability to *record* classical queries to a random oracle in a way that is not possible quantumly.

In our single-round protocols, the classical-quantum distinction is due to the knowledge assumption: for example, if $f : \mathcal{X} \to \mathcal{Y}$ is a one-way function with a sparse range (i.e. most values in $\mathcal{Y}$ are not in the range of f), it is plausible that a *classical* prover that produces a value in the range of f must have done so by evalating the function on some $x \in \mathcal{X}$, and must therefore also "know" the corresponding preimage; this can be captured formally as a knowledge assumption [16, 8]. However, a *quantum* prover can do the following: first prepare the state $\sum_x |x\rangle|f(x)\rangle$, then measure the second register in the computational basis to obtain an image $y \in \text{range}(f)$, and measure the first register in the Hadamard basis to "erase" any knowledge of the preimage x. Such a prover can be said to produce a value in the range of f without knowledge of a preimage. It is this distinction between classical and quantum computation that our single-round proofs of quantumness exploit. We note that a related idea of oblivious LWE sampling was recently explored in [18]; since a proof of quantumness only requires soundness against classical provers, we do not need to analyse the quantum prover in detail, but it might be interesting to explore the relation between knowledge assumptions used in our work and the (strong) notion of oblivious LWE sampling proposed in [18]. We discuss the classical-quantum distinction and the impossibility of single-round proofs of quantumness with black-box security reductions in more detail in Section 3.

Organsation

The rest of this work is structured as follows: in Section 1.1, we give a brief overview of the two-round proof of quantumness from [13]. This will form the basis for our single-round proof of quantumness. In Section 1.2, we recall two existing knowledge assumptions from [7, 36] that we use in our protocols. In Section 2, we explain how to make use of these knowledge assumptions to obtain different single-round proofs of quantumness. In Section 3 we argue that a white-box assumption (like a knowledge assumption) seems to be necessary for single-round proofs of quantumness. Finally, in Section 4 and Section 5 we discuss additional related works and open questions. For a formal description of our results and proofs, we refer to the full version of this paper [5].

1.1 A two-round proof of quantumness

To explain our results, we first need to outline the two-round (four-message) proof of quantumness protocol from [13]. At the heart of this protocol is a collection of functions known as Trapdoor Claw-free Functions (TCF). TCFs are a type of collision-resistant hash function – they are 2-to-1 functions for which it should be intractable to find a colliding pair of inputs (known as a *claw*), given the description of the function. Additionally, the functions are generated with a trapdoor that allows for efficient inversion. The specific TCFs used in [13] require an additional property known as the *adaptive hardcore-bit* (AHCB) property. Informally, this states that it is not only intractable to find collisions, but given any particular input x_0 and corresponding image under the TCF, denoted $y = f(x_0)$, it should be intractable to recover even a single bit of x_1, the other input with which x_0 forms a claw (i.e. $f(x_1) = f(x_0) = y$). Prior to our work, constructing TCFs with an AHCB had only been achieved from LWE and (non-standard) hardness assumptions of isogeny-based group actions [3]. One of our results shows an AHCB property for a TCF based on DDH.

Algorithm 1 The Proof of Quantumness of [13] (informal).

1. The verifier generates a description of a TCF f, together with its trapdoor, t. It then sends f to the prover.
2. The prover sends the verifier a point y in the image of f. Denote as x_0 and x_1 the associated preimages, so $y = f(x_0) = f(x_1)$.
3. With probability 1/2, the verifier sends the prover one of the following two challenges:
   - **(a) Preimage test.** The verifier asks the prover for a valid preimage of y. Denoting the prover's response as x, the verifier accepts iff $f(x) = y$.
   - **(b) Equation test.** The verifier asks the prover for a non-zero string d, such that $d \cdot (x_0 \oplus x_1) = 0$. The verifier uses the trapdoor, t, to recover x_0 and x_1 from y and then checks whether the equation is satisfied, accepting if it is and rejecting otherwise.

Given a family of TCFs, with the AHCB property, the Brakerski et al. protocol from [13] is described informally in Protocol 1. Brakerski et al. showed that, assuming the AHCB property, no polynomial-time classical prover makes the verifier accept with probability non-negligibly[3] larger than 3/4 (this is referred to as the *soundness* of the protocol). At the same time, they gave a simple quantum strategy which would allow the prover to succeed with probability 1 (referred to as the *completeness* of the protocol). This honest prover first creates an equal superposition over evaluations of f,

[3] A *negligible* function f, denoted by $f = \mathrm{negl}(\lambda)$, is a function $f : \mathbb{N} \to \mathbb{R}$ such that $f(\lambda) = o(\lambda^{-c})$ for every constant $c \in \mathbb{R}$.

$$\sum_x |x\rangle|f(x)\rangle\,.$$

The prover then measures the second register, resulting in a value $y = f(x_0) = f(x_1)$, and collapsing the state in the first register to

$$\frac{|x_0\rangle + |x_1\rangle}{\sqrt{2}}\,. \tag{1.1}$$

We can see that if this state is measured in the computational basis, it results in one of the two preimages of y, thus providing a valid response to the preimage test. Conversely, if the state is measured in the Hadamard basis (i.e. applying Hadamard gates to all qubits and measuring in the computational basis), the result will be a string d such that $d \cdot (x_0 \oplus x_1) = 0$, yielding a valid response to the equation test. Thus, a quantum prover implementing this strategy would make the verifier accept with probability 1.

The intuition for why it is classically intractable to succeed in this protocol is that a classical prover that answers both the preimage and equation tests correctly can be *rewound* so as to obtain both a valid preimage and a valid equation. However, having both would contradict the AHCB property. In contrast, a quantum prover cannot be rewound and can use the state from Equation (1.1) to answer *either of the two challenges* correctly, but not both at the same time. Communication between the two rounds is crucial for the security proof, as a reduction to the AHCB property requires that the prover holds a preimage and equation *for the same* y. This can only be guaranteed if the prover commits to a fixed choice of y before receiving the second challenge.

A natural question is whether the equation test on its own is already classically intractable. Unfortunately, simply removing the preimage test breaks the security proof of [13]: the AHCB property says that it is hard to produce an image, an equation, *and a preimage* together. Therefore, without the preimage test one cannot use a successful classical prover in the proof of quantumness to break the AHCB property, as one does not have access to a preimage[4]. Indeed, it can be shown that no *black-box* reduction can reduce the security of this "equation-test only" proof of quantumness to LWE ([39], see also Section 3).

Our results show that a variation of the "equation-test only" proof of quantumness can be made to work, provided we use a knowledge assumption to replace the role of the preimage test. Intuitively, the knowledge assumption can be used to "extract" a preimage from a successful classical prover without the need for an explicit preimage test. Since a knowledge assumption deals with the inner workings of a prover, our result can be interpreted as a *white-box* reduction from a version of the equation test to the AHCB property.

1.2 Knowledge assumptions

Knowledge assumptions posit the existence of an efficient extractor that is able to produce certain intermediate values that are in some relation with the output of an algorithm. A canonical example is the so-called *knowledge of exponent assumption* (KEA), introduced by Damgård in [17]. Informally, this says that given a generator g of some multiplicative

[4] We remark here that this is actually not the only problem that arises when removing the preimage test. In fact, Urmila Mahadev observed that for both constructions we consider, there exists a classical winning strategy in the equation test which involves evaluating f on an extended domain. In the original Brakerski et al. protocol [13], this strategy is excluded by the preimage test. In our protocols, this strategy is excluded by adding another type of test which does not increase the number of rounds of the protocol and is explained in Section 2.

group $\mathbb{G}$, as well g^α for some random power α, the only way to produce a new pair of the form (h, h^α) is by *exponentiating* the original pair. In other words, any efficient (randomized) algorithm $\mathcal{A}$ which outputs (h, h^α) given (g, g^α) must "know" an exponent k such that $h = g^k$. This is formalized by saying that for every such algorithm $\mathcal{A}$, there exists an efficient extractor $\mathcal{A}^*$ that, given (g, g^α) and the random coins of $\mathcal{A}$ for which $\mathcal{A}$ outputs (h, h^α), will output (or *extract*) the exponent k such that $h = g^k$ with overwhelming probability:

▶ **Assumption 1** (KEA (informal), from [17]). *Let $\mathcal{A}$ be an efficient probabilistic algorithm which receives as input a generator g of a multiplicative group $\mathbb{G}$ and a random group element g^α. Then there exists an efficient extractor $\mathcal{A}^*$, which uses the same input and random coins r as $\mathcal{A}$, and such that*

$$\Pr\left[\begin{array}{c}(h,h') \leftarrow \mathcal{A}(g, g^\alpha, r) \\ h^\alpha = h'\end{array} \wedge \begin{array}{c} k \leftarrow \mathcal{A}^*(g, g^\alpha, r) \\ g^k \neq h\end{array}\right] = \mathrm{negl}(\lambda)\,. \tag{1.2}$$

Variations of this assumption have also been considered, such as the t-KEA, in [7], in which the input to $\mathcal{A}$ is instead of the form $(g^{\mathbf{r}}, g^{\alpha\mathbf{r}})$, where $\mathbf{r}$ is a t-dimensional vector and exponentiation is element-wise.

The rationale behind knowledge of exponent assumptions is that the function that maps k to $(g^k, g^{\alpha k})$ has a sparse image (considered as a subset of $\mathbb{G} \times \mathbb{G}$). Thus, it seems implausible that $\mathcal{A}$ could come up with a valid image simply by cleverly sampling its output from $\mathbb{G} \times \mathbb{G}$ without exponentiating the input.[5]

More recently, similar knowledge assumptions were proposed for lattices. One example, introduced in [36], is known as the *knowledge of lattice point assumption* (denoted as LK-ϵ). This says the following: suppose there is an efficient randomized classical algorithm $\mathcal{A}$ which takes as input a lattice basis $\mathbf{A}$ and outputs a point $\mathbf{c}$ that is ϵ-close to the lattice $\mathcal{L}(\mathbf{A})$. Then $\mathcal{A}$ must "know" the lattice point closest to $\mathbf{c}$.

▶ **Assumption 2** (LK-ϵ (informal), from [36]). *Let ϵ be a fixed constant in $(0, 1/2)$. Denote by $\lambda(\mathcal{L})$ the norm of the shortest vector in a lattice $\mathcal{L}$. Let $\mathcal{A}$ be an efficient probabilistic algorithm which receives as input a basis $\mathbf{A}$ of a lattice $\mathcal{L}$ and outputs a vector $\mathbf{y}$. Then there exists an efficient algorithm $\mathcal{A}^*$ which uses the same input and random coins r as $\mathcal{A}$, and outputs a lattice point $\mathbf{p} \in \mathcal{L}$ such that*

$$\Pr\left[\begin{array}{c}\mathbf{y} \leftarrow \mathcal{A}(\mathbf{A}, r) \\ \exists\ \mathbf{c} \in \mathcal{L}(\mathbf{A})\ s.t.\ ||\mathbf{c} - \mathbf{y}|| \leq \epsilon \cdot \lambda(\mathcal{L}(\mathbf{A}))\end{array} \wedge \begin{array}{c}\mathbf{p} \leftarrow \mathcal{A}^*(\mathbf{A}, r) \\ \mathbf{p} \neq \mathbf{c}\end{array}\right] = \mathrm{negl}(\lambda)\,. \tag{1.3}$$

The motivation for this assumption is similar to that of KEA – the set of points that are close to the lattice (for a suitable choice of ϵ) is sparse in the set of all points, and the only apparent efficient strategy to sample from this sparse set is to pick a random lattice point and perturb it. This lattice knowledge assumption and variants of it have been used to construct efficient FHE and SNARK schemes [19, 27].

For our results, we will use t-KEA and LK-ϵ, which we state formally in [5, Assumption 4] and [5, Assumption 5].

[5] An alternative way of coming up with an image would be if $\mathcal{A}$ could determine α, given g and g^α. However, this would entail solving the discrete logarithm problem, which is believed to be classically intractable.

2 Main results

Our main result is that combining knowledge assumptions with standard cryptographic assumptions, like LWE or DDH, leads to efficient single-round proof of quantumness protocols. To make our results modular, we first show how to construct a general single-round proof of quantumness from a cryptographic primitive that we call Doubly Extended Extractable Noisy Trapdoor Claw-free Functions (abbreviated e^{3}NTCF). Second, we give two constructions of e^{3}NTCF: one from the DDH and t-KEA assumptions, and one from the LWE and LK-ϵ assumptions.

2.1 Single-round proof of quantumness from e^{3}NTCF

Our starting point is the protocol from [13], which we explained in Section 1.1. Recall that there, one uses a TCF with the AHCB property, and argues that if a classical prover could succeed in the preimage and equation tests, by rewinding we could construct a tuple (x, y, d) of a preimage, image, and equation that would contradict the AHCB property.

Now suppose that the TCF family $\mathcal{F} = \{f : \mathcal{X} \to \mathcal{Y}\}$ used in this protocol had an additional extractability property: for any classical prover that produces an image $y \in \text{range}(f)$, there exists an extractor that produces a corresponding preimage x. This is, in essence, a knowledge assumption. With such an extractable TCF, we could simply remove the preimage test from Protocol 1: then, if we had a classical prover that succeeded in this modified protocol, we could use that prover to find an image $y \in \text{range}(f)$ and equation d, and use the extractor to find a corresponding preimage x, such that (x, y, d) break the AHCB property.

Unfortunately, we do not know how to construct such an extractable TCF with AHCB from existing knowledge assumptions such as t-KEA or LK-ϵ[6]. The key difficulty here is that the extractability and AHCB properties have to hold simultaneously.

One way to circumvent this issue would be to introduce an additional function family $\mathcal{H} = \{h : \mathcal{X} \to \mathcal{Y}\}$ that is indistinguishable from the TCF family $\mathcal{F}$ (i.e. given a description of a random choice of either kind of function, it is hard to tell which kind of function it is). This function family $\mathcal{H}$ can be constructed such that it has an extractability property, i.e. if a classical algorithm produces a value y in the image of h, an extractor can find a preimage x (under a standard knowledge assumption such as t-KEA). However, $\mathcal{H}$ itself is not a TCF and has no AHCB property.

We now want to combine the AHCB property of $\mathcal{F}$ with the extractability property of $\mathcal{H}$. For this we leverage that the two function families are computationally hard to distinguish[7]: if we send a description of either $f \in \mathcal{F}$ or $h \in \mathcal{H}$ to a classical prover, the prover cannot tell which kind of function it received. This suggests the following protocol.

1. The verifier generates a description of a TCF $f \in \mathcal{F}$ (with a trapdoor t) or an extractable function $h \in \mathcal{H}$ and sends this function description to the prover.

[6] To be precise, we found that using these knowledge assumptions directly on the known NTCF constructions based on DDH and LWE does not work to infer extractability.

[7] One may ask here why we do not use computational indistinguishability to directly transfer the extractability property of $\mathcal{H}$ to $\mathcal{F}$, given that the extractor is an efficient algorithm. This has to do with the exact form of the extractability property: successful extraction is only guaranteed under the condition $y \in \text{range}(h)$, which cannot be checked efficiently without the trapdoor. Thus, while we can exploit the computational indistinguishability of $\mathcal{F}$ and $\mathcal{H}$ in our protocol, it is not possible to use computational indistinguishability to infer an analogous extractability property for $\mathcal{F}$.

2. The verifier receives (y, d) from the prover.
   - If the verifier sent a TCF f, it uses the trapdoor to recover x_0 and x_1 such that $y = f(x_0) = f(x_1)$ and accepts iff $d \cdot (x_0 \oplus x_1) = 0$.
   - If the verifier sent an extractable function h, it accepts iff $y \in \text{range}(h)$.

Suppose that a classical prover succeeded in this protocol with high probability. We can use this prover (and the corresponding extractor for the extractable function family) to break the AHCB property of f as follows: given a description of $f \in \mathcal{F}$, run the prover to generate (y, d). Then use the extractor on this prover on input $f \in \mathcal{F}$ to generate a preimage x. We claim that (x, y, d) violates the AHCB property.

Note that this reduction runs the extractor for the function family $\mathcal{H}$ on an input $f \in \mathcal{F}$, i.e. it runs the extractor on an input for which it was not designed. However, recall that descriptions of f and h are computationally indistinguishable. On the other hand, the extractor as well as the function that checks whether the extractor produced a correct preimage are efficient. Therefore, since the protocol ensures that on input h, the classical prover produces $y \in \text{range}(h)$, it follows that when we run the extractor for this classical prover on an input $f \in \mathcal{F}$, it will still produce a valid preimage of $y = f(x_0) = f(x_1)$; otherwise, we could distinguish f from h.

When trying to construct such a pair $(\mathcal{F}, \mathcal{H})$ of function families from DDH and t-KEA, we are faced with one additional technical challenge: the only evident construction of an extractable function family $\mathcal{H}$ works by extending the functions $f \in \mathcal{F}$ to a larger domain $\mathcal{X}'$ (on which $\mathcal{F}$ no longer satisfies the AHCB property) and constructing extractable functions $h : \mathcal{X}' \to \mathcal{Y}$, i.e. the extractability property of h only holds with respect to the extended domain $\mathcal{X}'$. This means that the extractor may, on input $y \in f(\mathcal{X})$, produce a $x \in \mathcal{X}' \setminus \mathcal{X}$ that also maps to y. While this is a valid preimage to y, it is not useful for breaking the AHCB property: for that we need a preimage $x \in \mathcal{X}$.

We therefore introduce a third function family, $\mathcal{G} = \{g : \mathcal{X}' \to \mathcal{Y}\}$, that is computationally indistinguishable from both $\mathcal{H}$ and the extension of $\mathcal{F}$. The functions in $\mathcal{G}$ are injective even on the larger domain $\mathcal{X}'$. In our proof of quantumness, the verifier will check that when sent a function g, the prover returns an image $y \in g(\mathcal{X})$, i.e. the image y must be in the range of the restricted domain $\mathcal{X}$. This essentially forces the prover to evaluate any function it receives only on the restricted domain[8] $\mathcal{X}$, since if it evaluated the *injective* function g on an input $x \in \mathcal{X}' \setminus \mathcal{X}$, the verifier would reject the resulting image y.

In summary, our single-round proof of quantumness relies on a triplet of function families $(\mathcal{F}, \mathcal{G}, \mathcal{H})$ with the following properties:[9]

▶ **Definition 1** (e^{3}NTCF, informal)**.** *A tuple of function families $(\mathcal{F}, \mathcal{G}, \mathcal{H})$ is called a* Doubly Extended Extractable Noisy Trapdoor Claw-free Functions *(abbreviated e^{3}NTCF) if*

1. *$\mathcal{F}$ is a TCF family with an AHCB property.*
2. *$\mathcal{G}$ is an injective trapdoor one-way function family.*
3. *$\mathcal{H}$ is an extractable one-way function family, in the sense that for any algorithm that takes the description of $h \in \mathcal{H}$ as input and outputs $y \in \text{range}(h)$, there exists an extractor which outputs a preimage x such that $h(x) = y$.*
4. *The functions are computationally indistinguishable from each other. In other words, given a description of a function from one of the three families, no polynomial-time classical algorithm can determine the function's type with probability non-negligibly greater than $1/3$.*

[8] This also serves to exclude the "extended-domain" attack mentioned in Footnote 4.

[9] As usual, this primitive depends on a security parameter λ, which we suppress for readability.

For the formal statement, see [5, Definition 8].

The notion of an e^{3}NTCF function family is an extension of Injective Invariant Noisy TCFs, which were introduced in [38] to derive a protocol for verifying general quantum computations, and which only use the first two types of families $\mathcal{F}$ and $\mathcal{G}$.

The preceding discussion naturally leads to the single-round Protocol 2 (see [5, Protocol 4] for the formal protocol), based on an e^{3}NTCF. Our main result is that this protocol is a proof of quantumness, i.e. that an efficient quantum prover can succeed with high probability, but no efficient classical prover can.

Algorithm 2 Single-Round Proof of Quantumness based on e^{3}NTCF. (informal)

1. The verifier samples a challenge type $a \leftarrow_U \{\texttt{Eq}, \texttt{sIm}, \texttt{wIm}\}$.
   if $a = \texttt{Eq}$:
2. Verifier samples $f \in \mathcal{F}$ with trapdoor t_f and sends a description of f to the prover.
3. Verifier receives tuple (y, d) and accepts iff $d \cdot (x_0 \oplus x_1) = 0$ (and $d \neq 0$). Here, (x_0, x_1) are the two preimages of y, which the verifier can efficiently compute with t_f.
   else if $a = \texttt{sIm}$:
4. Verifier samples $g \in \mathcal{G}$ with trapdoor t_g and sends a description of g to the prover.
5. Verifier receives tuple (y, d) and accepts iff $y \in g(\mathcal{X})$. This check is efficient using t_g.
   else if $a = \texttt{wIm}$:
6. Verifier samples $h \in \mathcal{H}$ with trapdoor t_h and sends a description of h to the prover.
7. Verifier receives tuple (y, d) and accepts iff $y \in h(\mathcal{X}')$. This check is efficient using t_h.
   end if

▶ **Theorem 2** (informal). *Suppose that $(\mathcal{F}, \mathcal{G}, \mathcal{H})$ is an e^{3}NTCF. Then Protocol 2 is a proof of quantumness, i.e.,*

1. ***Completeness.*** *There exists an efficient quantum prover that succeeds with probability $1 - \text{negl}(\lambda)$.*
2. ***Soundness.*** *Every efficient classical prover succeeds with probability at most $5/6 + \text{negl}(\lambda)$.*

In the theorem, λ is the security parameter. As usual, the families $(\mathcal{F}, \mathcal{G}, \mathcal{H})$ implicitly depend on the security parameter and "efficient" classical or quantum provers are those whose runtime scales polynomially in λ. For the formal statement, see [5, Theorem 4].

Proof Sketch. We have already sketched the security proof when we explained why we need to introduce the three different function families $\mathcal{F}, \mathcal{G}$, and $\mathcal{H}$. We briefly summarize the role that each of these function families and the associated challenge types $\texttt{Eq}, \texttt{sIm}$, and $\texttt{wIm}$ play, and provide a high-level explanation of how the soundness constant $5/6$ is derived.

- The *weak image test* (challenge type $\texttt{wIm}$) uses the extractable function family $\mathcal{H}$. This test ensures that for any classical prover in the protocol, there exists an extractor that extracts a preimage $x \in \mathcal{X}'$ to the output $y = h(x)$ produced by the prover. Note that as discussed above, this preimage might in principle be from the extended domain $\mathcal{X}'$, not just $\mathcal{X}$.
- The *strong image test* (challenge type $\texttt{sIm}$) uses the injective function family $\mathcal{G}$. This test, combined with the indistinguishability of the function families $\mathcal{F}, \mathcal{G}$, and $\mathcal{H}$, ensures that the prover evaluates any function f, g, or h only on inputs in the restricted domain $\mathcal{X}$. Furthermore, from this we can prove that for any $y \in \mathcal{Y}$ that has a preimage in $\mathcal{X}$ under a given *TCF function* f, the extractor will output such a preimage (rather than a different preimage in $\mathcal{X}' \setminus \mathcal{X}$).

- The *equation test* (challenge type `Eq`) uses the TCF function family $\mathcal{F}$ that has the AHCB property. The verifier checks that the prover's output (y, d) satisfies the equation $d \cdot (x_0 \oplus x_1) = 0$, with (x_0, x_1) the preimages of y. From the strong image test, we also know that the extractor for a successful classical prover will produce a preimage $x \in \mathcal{X}$ to y. Together, (x, y, d) would break the AHCB property of f, so no classical prover can win with very high probability.

The soundness factor of 5/6 arises from the fact that the adaptive hardcore bit property only gives an upper bound of 1/2 for the event that the prover fulfills the conditions of all three tests simultaneously, given a function of the family $\mathcal{F}$. Using the computational indistinguishability of the function families $\mathcal{F}, \mathcal{G}, \mathcal{H}$, and the relation

$$\Pr[A] + \Pr[B] \leq 1 + \Pr[A \wedge B]$$

we show that the overall success probability of a classical prover $\mathcal{A}$ in the prover can roughly be bounded as

$$\frac{1}{3}\Big(\Pr_{f\leftarrow\mathcal{F}}\Big[S_{\mathcal{A}}^{\mathtt{Eq}}\Big] + \Pr_{g\leftarrow\mathcal{G}}\Big[S_{\mathcal{A}}^{\mathtt{sIm}}\Big] + \Pr_{h\leftarrow\mathcal{H}}\Big[S_{\mathcal{A}}^{\mathtt{wIm}}\Big]\Big) \leq \frac{1}{3}\Big(2 + \Pr_{f\leftarrow\mathcal{F}}\Big[\text{“}S_{\mathcal{A}}^{\mathtt{wIm}} \wedge S_{\mathcal{A}}^{\mathtt{sIm}} \wedge S_{\mathcal{A}}^{\mathtt{Eq}}\text{”}\Big]\Big) \leq \frac{5}{6}$$

where $S_{\mathcal{A}}^{\mathtt{T}}$ denotes the event that $\mathcal{A}$ passes the test T, and “$S_{\mathcal{A}}^{\mathtt{wIm}} \wedge S_{\mathcal{A}}^{\mathtt{sIm}} \wedge S_{\mathcal{A}}^{\mathtt{Eq}}$” denotes the event that the prover succeeds in the equation test, and is employing a strategy that would succeed in the image tests as well. This corresponds to the event that $\mathcal{A}, \mathcal{A}^*$ produce a valid equation-image pair, which cannot occur with probability higher than 1/2 due to the AHCB property of the function family $\mathcal{F}$. ◀

A key objective in the design of our single-round proofs of quantumness was that the required circuit sizes should not be larger than for multi-round proofs of quantumness. Protocol 2 achieves this: to pass in the protocol, a quantum prover can simply prepare the state $\sum |x\rangle|p(x)\rangle$ for a given function $p \in \mathcal{F} \cup \mathcal{G} \cup \mathcal{H}$, measure the first register in the Hadamard basis to get a string d, and measure the second register in the computational basis to get an image y. These are exactly the same actions as those of an honest prover in the equation test in Protocol 1,[10] except that now the honest prover can perform the entire measurement in one step without experimentally challenging mid-circuit measurements.

2.2 Instantiating e^{3}NTCF families from DDH and LWE

We show that e^{3}NTCF can be instantiated either from DDH and t-KEA, or from LWE and LK-ϵ. Here, we only state the main results and defer a more detailed discussion of the construction to [5, Section 4] and [5, Section 6], respectively.

For the LWE-based construction, we already know that the LWE-based TCF from [13] has an AHCB property. We can combine this with suitable injective and extractable one-way functions. In fact, it turns out that for the LWE-based construction, the roles of the injective and extractable functions in Protocol 2 can be played by the same function family, i.e. we only require two distinct families $(\mathcal{F}, \mathcal{G})$. This simplifies the analysis somewhat, as we explain in [5, Section 5]. Combined with Theorem 2, we obtain the following:

[10] In our constructions of e^{3}NTCF, the functions $g \in \mathcal{G}$ and $h \in \mathcal{H}$ are essentially as costly to evaluate as $f \in \mathcal{F}$ in Protocol 1, so introducing these additional function families does not increase the demands on an honest prover.

▶ **Theorem 3** (Informal). *Assuming the classical intractability of LWE and the lattice knowledge assumption LK-ϵ, there exists a single-round proof of quantumness with completeness* $1 - \mathrm{negl}(\lambda)$ *and soundness* $3/4 + \mathrm{negl}(\lambda)$. *The circuit sizes for an honest prover in this protocol are identical to the circuit sizes in the 2-round protocol from [13].*

While a family of TCFs based on DDH was considered in [29] to construct a 3-round proof of quantumness, it had not been shown that these functions have an AHCB property. We prove this through a lossy sampling technique similar to the proof in [13], showing a reduction from the AHCB property to the *matrix d-linear assumption* [41] (which is implied by standard DDH). We then again augment this TCF family with an injective and an extractable family to get the following result.

▶ **Theorem 4** (Informal). *Assuming the classical intractability of DDH and the knowledge of exponent assumption t-KEA, there exists a single-round proof of quantumness with completeness*[11] $0.99 - \mathrm{negl}(\lambda)$ *and soundness* $5/6 + \mathrm{negl}(\lambda)$. *The circuit sizes for an honest prover in this protocol are identical to the circuit sizes in the 3-round protocol from [29].*

3 Impossibility of black-box reductions

Knowledge assumptions are non-falsifiable cryptographic assumptions [40, 20], which makes their use somewhat undesirable. A natural question is whether our main results, single-round proofs of quantumness, can also be achieved using only standard falsifiable (also called game-based) assumptions such as DDH or LWE, or even weaker assumptions like the existence of one-way functions.

As was observed in [39], the security of a single-round proof of quantumness cannot be reduced to a *quantumly hard* problem like LWE in a black-box manner. The reason for this is simple: if there existed a black-box reduction from a successful classical prover to an algorithm for breaking LWE, we could also apply that reduction to the honest quantum prover in the proof of quantumness. Since the honest quantum prover is required to succeed with high probability, this would give a quantum algorithm for LWE. This rules out the existence of such a black-box reduction.[12]

An extension of this reasoning suggests that the use of knowledge assumptions is still justified even for single-round proofs of quantumness based on computational problems that are quantumly easy: now it is of course not a contradiction to have a quantum algorithm that breaks the underlying cryptographic assumption, but we can argue that implementing an honest quantum prover in such a protocol is no easier than breaking the assumption outright: Suppose we had a single-round proof of quantumness with a black-box reduction to factoring. Then we could again use that reduction with the honest quantum prover to construct a factoring algorithm whose only quantum component is repeatedly (and independently) running the honest prover. In this sense, implementing the honest prover is as hard as breaking factoring: if we had an honest prover that was much easier to implement than Shor's algorithm, we could use that prover and the black-box reduction to construct a much more

[11] The reason completeness here is 0.99 instead of 1 is inherited from [29]. There, the quantum prover does not create a superposition over the exact range of the TCF but over a superset. However, it can be shown that the exact range contains at least a 0.99 fraction of the elements in the superset.

[12] Note that this argument does not apply to interactive protocols, ROM-based protocols, or protocols that use whitebox assumptions like ours: in those cases, one cannot simply run the reduction on the quantum prover as the reduction may perform operations that only work for classical provers, e.g. rewinding or using knowledge extractors.

practical quantum factoring algorithm. In contrast, in the DDH-based proof of quantumness from [29] and also in our single-round version thereof, implementing the honest quantum prover does not require computing discrete logarithms.

4 Related work

As mentioned in Section 1.1, cryptographic proofs of quantumness were introduced in [13]. Since then, there have been a number of follow-up works aiming to understand the types of cryptographic constructions on which these protocols can be based, as well as how to make them more efficient [30, 14, 35, 26, 4, 31, 39]. So far, only small-scale demonstrations of these protocols have been performed [49, 33], though the hope is that by further reducing the sizes of the quantum circuits and the amount of interaction between the verifier and the prover, these protocols can be performed with NISQ devices.

The only existing constructions of single-round proofs of quantumness (apart from trivial ones like asking the prover to factor a large number) use the random oracle model [14, 45]. In both of these prior works, the quantum circuits required to succeed in these protocols do not have to perform mid-circuit measurements. However, instantiating the random oracle with plausible cryptographic hash functions, like SHA-2 or SHA-3, incurs a substantial overhead. The most efficient known quantum implementations of SHA-2 and SHA-3 [28] need a number of qubits on the order of 10^4 and have a quantum gate count on the order of 10^6.

On the other hand, the ROM-based single-round proof of quantumness in [14] only requires a TCF, without the need for the AHCB property. General TCFs have been constructed from a wider variety of cryptographic assumptions, achieving lower asymptotic complexities than known NTCF constructions. Thus, despite the fact that the use of a random oracles introduces a significant overhead, a resource comparison between our LWE-based single-round proof of quantumness and the ROM-based protocol of [14] is not straightforward. A more conclusive answer on which single-round protocol admits the smallest quantum circuit for a demonstration of quantum advantage requires a detailed finite-size analysis, which we leave for future work.

Since our construction of e^3NTCFs makes use of extractable one-way functions (EOWFs), it is natural to ask whether we could have used existing results concerning these functions [16, 8, 10]. In particular, in [8], Bitansky et al. prove the existence of a class of EOWFs from the subexponential hardness of LWE against adversaries with bounded auxilary input, without having to rely on a knowledge assumption. There are at least two obstacles towards applying those results to our setting. First, the extractable one-way functions we require have to be computationally indistinguishable from TCFs. It is unclear whether the EOWFs of [8] can be suitably adapted to satisfy this requirement. Second, since one of the motivations for our work is to devise more efficient single-round proofs of quantumness, we would like to ensure that the honest quantum prover's circuits are at most as large as in the multi-round protocols. However, constructing a family of e^3NTCFs with the EOWFs in [8] would likely require larger circuits.

Knowledge assumptions in the quantum setting have also been considered in [34] to derive *quantum money* and *quantum lightning* schemes. There, the assumptions are required to be sound against quantum adversaries. This introduces some challenges in formalizing the appropriate notion of a quantum knowledge assumption. The subsequent work of [46] also discusses this point. Interestingly, this latter work points out how certain knowledge assumptions can be generically broken by a quantum adversary. In some sense, this is exactly what we exploit to arrive at our single-round proofs of quantumness.

5 Discussion and open problems

We have shown how existing knowledge assumptions, together with standard cryptographic assumptions, lead to efficient single-round proofs of quantumness. Our work opens up several avenues for future research.

One such avenue is the possibility of a *white-box proof* for a single-round proof of quantumness based only on the classical intractability of, say, LWE. This would circumvent the impossibility discussed in Section 3. As mentioned in the previous section, one possible approach would be to use the extractable one-way functions based on subexponential LWE from [8], with the caveat that this would provide soundness against classical provers with bounded auxiliary input. We leave this as an interesting but challenging open problem.

Currently, our construction necessitates the use of a TCF which satisfies the AHCB property. This is a fairly strong requirement which so far has only been achieved from LWE [13], isogeny-based group actions [3], and, in this work, the DDH assumption. It would be desirable to use TCFs that are not known to have an AHCB property, such as those based on Ring-LWE or Rabin's function (x^2 mod N), since those require smaller circuits than the TCFs based on LWE or DDH, as outlined in [29, 30]. However, it is unclear if our construction can be modified so as to not require the AHCB property; in the equation test, the prover is allowed to output *any* valid equation, hence the need for an *adaptive* hardcore bit. The protocol in [29] is able to circumvent this and use TCFs without an AHCB by introducing an additional round of interaction between the verifier and the prover, leading to a 3-round protocol. Unfortunately, there does not seem to be an obvious way of employing knowledge assumptions to reduce the round complexity of the protocol from [29]. Thus, a fundamentally different approach would be needed to achieve a single-round protocol without relying on the AHCB property. Of course, for the purpose of having more practical protocols, proving an AHCB statement for the TCFs based on Ring-LWE or Rabin's function (even relying on knowledge assumptions) would be equally useful.

We remark that the LWE-based approach is expected to be more efficient than the DDH one. This is because, as is also discussed in [29], the DDH approach requires performing a large number of group exponentiations coherently. It would therefore seem that the DDH approach is less efficient than performing Shor's algorithm to solve the discrete logarithm problem. However, as [29] points out, the proof of quantumness protocol can benefit from certain optimizations that are not known to be possible for Shor's algorithm. As an example, the group exponentiation circuits for the DDH-based proof of quantumness need not be reversible, thereby halving the depth and number of gates compared to Shor's algorithm (see section III.D in [29] for more details).

From existing knowledge assumptions one is only able to prove soundness against classical adversaries, i.e. one is able to bound the success probability of a classical prover, but one cannot make a statement about the internal operations of a successful quantum prover. This is sufficient for a proof of quantumness, but other protocols, such as certifiable randomness generation [13], remote state preparation [21], or verification of quantum computations [38], require soundness against quantum adversaries, i.e. one needs to have at least a partial characterization of any quantum adversary in the protocol. As mentioned, one can derive single-round proofs of quantumness in the random oracle model, and these can be extended to single-round versions of quantum-sound protocols by proving security in the quantum random oracle model [11, 14, 2, 47]. However, with knowledge assumptions, it is unclear what the right quantum-sound analogue would be to derive these more advanced functionalities. As mentioned in the previous section, quantum knowledge assumptions have been considered

in [34, 46] to construct quantum money and quantum lightning schemes. These could serve as a useful starting point for constructing single-round quantum protocols for functionalities like single-device randomness expansion. We leave formalizing this for future work.

References

1 Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342, 2011. doi:10.1145/1993636.1993682.

2 Gorjan Alagic, Andrew M Childs, Alex B Grilo, and Shih-Han Hung. Non-interactive classical verification of quantum computation. In *Theory of Cryptography Conference*, pages 153–180. Springer, 2020. doi:10.1007/978-3-030-64381-2_6.

3 Navid Alamati, Giulio Malavolta, and Ahmadreza Rahimi. Candidate trapdoor claw-free functions from group actions with applications to quantum protocols. In *Theory of Cryptography Conference*, pages 266–293. Springer, 2022. doi:10.1007/978-3-031-22318-1_10.

4 Yusuf Alnawakhtha, Atul Mantri, Carl A Miller, and Daochen Wang. Lattice-based quantum advantage from rotated measurements. *arXiv preprint*, 2022. doi:10.48550/arXiv.2210.10143.

5 Petia Arabadjieva, Alexandru Gheorghiu, Victor Gitton, and Tony Metger. Single-round proofs of quantumness from knowledge assumptions, 2024. doi:10.48550/arXiv.2405.15736.

6 Frank Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. doi:10.1038/s41586-019-1666-5.

7 Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 326–349, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2090236.2090263.

8 Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 505–514, 2014. doi:10.1145/2591796.2591859.

9 Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Omer Paneth, and Rafail Ostrovsky. Succinct non-interactive arguments via linear interactive proofs. In *Theory of Cryptography: 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings*, pages 315–333. Springer, 2013. doi:10.1007/s00145-022-09424-4.

10 Nir Bitansky, Noa Eizenstadt, and Omer Paneth. Weakly extractable one-way functions. In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part I 18*, pages 596–626. Springer, 2020. doi:10.1007/978-3-030-64375-1_21.

11 Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology–ASIACRYPT 2011: 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings 17*, pages 41–69. Springer, 2011. doi:10.1007/978-3-642-25385-0_3.

12 Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15(2):159–163, 2019. doi:10.1038/s41567-018-0318-2.

13 Zvika Brakerski, Paul F. Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 320–331. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00038.

14 Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. Simpler proofs of quantumness, 2020. doi:10.48550/arXiv.2005.04826.

15 Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *Journal of computer and system sciences*, 37(2):156–189, 1988. doi:10.1016/0022-0000(88)90005-0.

16 Ran Canetti and Ronny Ramzi Dakdouk. Towards a theory of extractable functions. In *Theory of Cryptography Conference*, pages 595–613. Springer, 2009. doi:10.1007/978-3-642-00457-5_35.

17 Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In Joan Feigenbaum, editor, *Advances in Cryptology — CRYPTO '91*, pages 445–456, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg. doi:10.1007/3-540-46766-1_36.

18 Thomas Debris-Alazard, Pouria Fallahpour, and Damien Stehlé. Quantum oblivious LWE sampling and insecurity of standard model lattice-based snarks. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 423–434. ACM, 2024. doi:10.1145/3618260.3649766.

19 Rosario Gennaro, Michele Minelli, Anca Nitulescu, and Michele Orrù. Lattice-based zk-SNARKs from square span programs. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, pages 556–573, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3243734.3243845.

20 Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 99–108, 2011. doi:10.1145/1993636.1993651.

21 Alexandru Gheorghiu and Thomas Vidick. Computationally-secure and composable remote state preparation. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1024–1033. IEEE, 2019. doi:10.1109/FOCS.2019.00066.

22 Craig Gidney and Martin Ekerå. How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits. *Quantum*, 5:433, 2021. doi:10.22331/q-2021-04-15-433.

23 Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Inf. Process. Lett.*, 67(4):205–214, 1998. doi:10.1016/S0020-0190(98)00116-1.

24 Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. COMPUT*, 18(1):186–208, 1989. doi:10.1145/22145.22178.

25 Élie Gouzien and Nicolas Sangouard. Factoring 2048-bit rsa integers in 177 days with 13 436 qubits and a multimode memory. *Physical review letters*, 127(14):140503, 2021. doi:10.1103/PhysRevLett.127.140503.

26 Shuichi Hirahara and François Le Gall. Test of quantumness with small-depth quantum circuits. *arXiv preprint*, 2021. doi:10.48550/arXiv.2105.05500.

27 Yuval Ishai, Hang Su, and David J. Wu. Shorter and faster post-quantum designated-verifier zkSNARKs from lattices. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, CCS '21, pages 212–234, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3460120.3484572.

28 Kyungbae Jang, Sejin Lim, Yujin Oh, Anubhab Baksi, Sumanta Chakraborty, and Hwajeong Seo. Quantum implementation and analysis of sha-2 and sha-3. Cryptology ePrint Archive, Paper 2024/513, 2024. URL: https://eprint.iacr.org/2024/513.

29 Gregory D. Kahanamoku-Meyer, Soonwon Choi, Umesh V. Vazirani, and Norman Y. Yao. Classically verifiable quantum advantage from a computational bell test. *Nature Physics*, 18(8):918–924, August 2022. doi:10.1038/s41567-022-01643-7.

30 Gregory D Kahanamoku-Meyer and Norman Y Yao. Fast quantum integer multiplication with zero ancillas. *arXiv preprint*, 2024. doi:10.48550/arXiv.2403.18006.

31 Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1617–1628, 2023. doi:10.1145/3564246.3585164.

32 Thomas Kerber, Aggelos Kiayias, and Markulf Kohlweiss. Composition with knowledge assumptions. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021*, pages 364–393, Cham, 2021. Springer International Publishing. doi:10.1007/978-3-030-84259-8_13.

33 Laura Lewis, Daiwei Zhu, Alexandru Gheorghiu, Crystal Noel, Or Katz, Bahaa Harraz, Qingfeng Wang, Andrew Risinger, Lei Feng, Debopriyo Biswas, et al. Experimental implementation of an efficient test of quantumness. *Physical Review A*, 109(1):012610, 2024. doi:10.1103/PhysRevA.109.012610.

34 Jiahui Liu, Hart Montgomery, and Mark Zhandry. Another round of breaking and making quantum money: How to not build it from lattices, and more. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 611–638. Springer, 2023. doi:10.1007/978-3-031-30545-0_21.

35 Zhenning Liu and Alexandru Gheorghiu. Depth-efficient proofs of quantumness. *Quantum*, 6:807, 2022. doi:10.22331/q-2022-09-19-807.

36 Jake Loftus, Alexander May, Nigel P. Smart, and Frederik Vercauteren. On cca-secure somewhat homomorphic encryption. In Ali Miri and Serge Vaudenay, editors, *Selected Areas in Cryptography*, pages 55–72, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. doi:10.1007/978-3-642-28496-0_4.

37 Lars S Madsen, Fabian Laudenbach, Mohsen Falamarzi Askarani, Fabien Rortais, Trevor Vincent, Jacob FF Bulmer, Filippo M Miatto, Leonhard Neuhaus, Lukas G Helt, Matthew J Collins, et al. Quantum computational advantage with a programmable photonic processor. *Nature*, 606(7912):75–81, 2022. doi:10.1038/s41586-022-04725-x.

38 Urmila Mahadev. Classical verification of quantum computations, 2018. doi:10.48550/arXiv.1804.01082.

39 Tomoyuki Morimae and Takashi Yamakawa. Quantum advantage from one-way functions. *arXiv preprint*, 2023. doi:10.48550/arXiv.2302.04749.

40 Moni Naor. On cryptographic assumptions and challenges. In *Annual International Cryptology Conference*, pages 96–109. Springer, 2003. doi:10.1007/978-3-540-45146-4_6.

41 Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. Cryptology ePrint Archive, Paper 2009/105, 2009. doi:10.1007/978-3-642-03356-8_2.

42 John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018. doi:10.22331/q-2018-08-06-79.

43 Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999. doi:10.1137/S0097539795293172.

44 Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, et al. Strong quantum computational advantage using a superconducting quantum processor. *Physical review letters*, 127(18):180501, 2021. doi:10.1103/PhysRevLett.127.180501.

45 Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 69–74. IEEE, 2022. doi:10.1109/FOCS54457.2022.00014.

46 Mark Zhandry. Quantum money from abelian group actions. *arXiv preprint*, 2023. doi:10.48550/arXiv.2307.12120.

47 Jiayu Zhang. Classical verification of quantum computations in linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 46–57. IEEE, 2022. doi:10.1109/FOCS54457.2022.00012.

48 Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020. doi:10.1126/science.abe8770.

49 Daiwei Zhu, Gregory D Kahanamoku-Meyer, Laura Lewis, Crystal Noel, Or Katz, Bahaa Harraz, Qingfeng Wang, Andrew Risinger, Lei Feng, Debopriyo Biswas, et al. Interactive cryptographic proofs of quantumness using mid-circuit measurements. *Nature Physics*, 19(11):1725–1731, 2023. doi:10.1038/s41567-023-02162-9.