Quantum Proofs of Knowledge

Dominique Unruh
University of Tartu, Estonia

April 14, 2011

Abstract

We motivate, define and construct quantum proofs of knowledge, proofs of knowledge secure against quantum adversaries. Our constructions are based on a new quantum rewinding technique that allows us to extract witnesses in many classical proofs of knowledge. We give criteria under which a classical proof of knowledge is a quantum proof of knowledge. Combining our results with Watrous' results on quantum zero-knowledge, we show that there are zero-knowledge quantum proofs of knowledge for all languages in NP (assuming quantum one-way permutations).

Contents

I value comments on this paper. If you find mistakes or have other suggestions, please send them to unruh@ut.ee.

1 Introduction

Cryptographic protocols, with few exceptions, are based on the assumption that certain problems are computationally hard. Typical examples include specific number-theoretic problems such as the difficulty of finding discrete logarithms, and general problems such as inverting one-way functions. It is well-known, however, that many such problems would become easy in the advent of quantum computers. Shor's algorithm [Sho94], e.g., efficiently solves the discrete logarithm problem and allows to factor large integers. While quantum computers do not exist today, it is not unreasonable to expect quantum computers to be available in the future. To meet this threat, we need cryptographic protocols that are secure even in the presence of an adversary with a quantum computer. We stress that this does not necessarily imply that the protocol itself should make use of quantum technology; instead, it is preferable that the protocol itself can be easily implemented on today's readily-available classical computers.

Finding such quantum-secure protocols, however, is not trivial. Even when we have found suitable complexity-theoretic assumptions such as the hardness of certain lattice problems, a classical protocol based on these assumptions may fail to be secure against quantum computers. The reason for this is that many cryptographic proofs use a technique called rewinding. This technique requires that it is possible, when simulating some machine, to make snapshots of the state of that machine and then later to go back to that snapshot. As first observed by van de Graaf [vdG98], classical rewinding-based proofs do not carry over to the quantum case. Two features unique to the quantum setting prohibit (naive) rewinding: The no-cloning theorem [WZ82] states that quantum-information cannot be copied, so we cannot make snapshots. Furthermore, measurements destroy information, so interacting with a simulated machine may destroy information that would be needed later.

This leads to the following observation: Even if a classical protocol is proven secure based on the hardness of some problem, and even if that problem is hard even for quantum computers, we have no guarantee that the protocol is secure against quantum computers. The reduction of the protocol's security to the problem's hardness may be based on inherently classical features such as the possibility of rewinding.

An example of a protocol construction that suffers from this difficulty are zero-knowledge proofs. Zero-knowledge proofs are interactive proofs with the special property that the verifier does not learn anything except the validity of the proven statement. Zero-knowledge proofs are inherently based on rewinding (at least as long as we do not assume additional trusted setup such as so-called common-reference strings). Yet, zero-knowledge proofs are one of the most powerful tools available to the cryptographer; a multitude of protocol constructions use zero-knowledge proofs. These protocol constructions cannot be proven secure without using rewinding. To resolve this issue, Watrous [Wat09] introduced a quantum rewinding technique. This technique allows to prove the quantum security of many common zero-knowledge proofs. One should note, however, that Watrous' technique is restricted to a specific type of rewinding: If we use Watrous' technique, whenever some

machine rewinds another machine to an earlier point, the rewinding machine forgets everything it learned after that point (we call this oblivious rewinding). That is, we can only use Watrous' technique to backtrack if the rewinding machine made a mistake that should be corrected, but it cannot be used to collect and combine information from different branches of an execution.

Constructing quantum zero-knowledge proofs solves, however, only half of the problem. In many, if not most, applications of zero-knowledge proofs one needs zero-knowledge *proofs of knowledge*. A proof of knowledge [GMR85, BG93] is a proof system which does not only show the truth of a certain statement, but also that the prover knows a witness for that statement. This is made clearer by an example: Assume that Alice wishes to convince Bob that she (the prover) is in possession of a signature issued by some certification authority. For privacy reasons, Alice does not wish to reveal the signature itself. If Alice uses a zero-knowledge *proof*, she can only show the statement "there exists a signature with respect to the CA's public key". This does not, however, achieve anything: A signature always exists in a mathematical sense, even if it has never been computed. What Alice wishes to say is: "I *know* a signature with respect to the CA's public key." To prove such a statement, Alice needs a zero-knowledge *proof of knowledge*; a proof of knowledge would convince Bob that Alice indeed knows a witness, i.e., a signature. Very roughly, the definition of a proof of knowledge is the following: Whenever the prover can convince the verifier, one can extract the witness from the prover given oracle access to the prover. Here oracle access means that one can interact with the prover and *rewind* him. Thus, we have the same problem as in the case of quantum zero-knowledge proofs: To get proofs of knowledge that are secure against quantum adversaries, we need to use quantum rewinding. Unfortunately, Watrous' *oblivious* rewinding does not work here; proofs of knowledge use rewinding to produce two (or more) different protocol traces and compute the witness by combining the information from both traces. Thus, we are back to where we started: to make classical cryptographic protocols work in a quantum setting, we need (in many cases) quantum zero-knowledge *proofs of knowledge*, but we only have constructions for quantum zero-knowledge *proofs*.

Our contribution. We define and construct quantum proofs of knowledge. Our protocols are classical (i.e., honest parties do not use quantum computation or communication) but secure against quantum adversaries. Our constructions are based on a new quantum rewinding technique (different from Watrous' technique) that allows us to extract witnesses in many classical proofs of knowledge. We give criteria under which a classical proof of knowledge is a quantum proof of knowledge. Combining our results with Watrous' results on zero-knowledge, we can show that there are zero-knowledge quantum proofs of knowledge for all languages in NP (assuming quantum one-way permutations).

Organization. In Section 1.1, we give an overview over the techniques underlying our results. In Section 2 we present and discuss the definition of quantum proofs of knowledge (QPoKs). In Section 3, we give criteria under which a proof system is a QPoK. In Section 4, we show that zero-knowledge QPoKs exist for all languages in NP.

1.1 Our techniques

Defining proofs of knowledge. In the classical setting, proofs of knowledge are defined as follows:[1] A proof system consisting of a prover P and a verifier V is a proof of knowledge (PoK) with knowledge error κ if there is a polynomial-time machine K (the extractor) such that the following holds: For any prover P^*, if P^* convinces V with probability $\Pr_{\mathsf{V}} \geq \kappa$, then $\mathsf{K}^{\mathsf{P}^*}$ (the extractor K with rewinding black-box access to P^*) outputs a witness with probability $\Pr_{\mathsf{K}} \geq \frac{1}{p}(\Pr_{\mathsf{V}} - \kappa)^d$ for some polynomial p and constant $d > 0$. In order to transfer this definition to the quantum setting, we need to specify what it means that K has quantum rewinding black-box access to P^*. We choose the following definition: Let U denote the unitary transformation describing one activation of P^* (if P^* is not unitary, this needs to work for *all* purifications of P^*). K may invoke U (this corresponds to running P^*), he may invoke the inverse $U^\dagger$ of U (this corresponds to rewinding P^* by one activation), and he may read/write a shared register N for exchanging messages with P^*. But K may not make snapshots of the state of P^*. Allowing K to invoke $U^\dagger$ is justified by the fact that all quantum circuits are reversible; given a circuit for U, we can efficiently apply $U^\dagger$. Note that previous black-box constructions such as Watrous' rewinding technique and Grover's algorithm [Gro96] make use of this fact. We can now define quantum proofs of knowledge: (P, V) is a quantum proof of knowledge (QPoK) with knowledge error κ iff there is a polynomial-time quantum algorithm K such that for all malicious provers P^*, $\mathsf{K}^{\mathsf{P}^*}$ (the extractor K with quantum rewinding black-box access) outputs a witness with probability $\Pr_{\mathsf{K}} \geq \frac{1}{p}(\Pr_{\mathsf{V}} - \kappa)^d$ for some polynomial p and constant $d > 0$.

We illustrate that QPoKs according to this definition are indeed useful for analyzing cryptographic protocols. Assume the following toy protocol: In phase 1, a certification authority (CA) signs the pair $(\texttt{Alice}, a)$ where a is Alice's age. In phase 2, Alice uses a zero-knowledge QPoK with negligible knowledge error κ to prove to Bob that she possesses a signature σ on $(\texttt{Alice}, a')$ for some $a' \geq 21$. That is, a witness in this QPoK would consist of an integer $a' \geq 21$ and a signature σ on $(\texttt{Alice}, a')$ with respect to the CA's public key. We can now show that, if Alice is underage, i.e., if $a < 21$, Bob accepts the QPoK only with negligible probability: Assume that Bob accepts with non-negligible probability ν. Then, by the definition of QPoKs, $\mathsf{K}^{\texttt{Alice}}$ will, with probability $\frac{1}{p}(\nu-\kappa)^d$, output an integer $a' \geq 21$ and a (forged) signature σ on $(\texttt{Alice}, a')$ with respect to the CA's public key (given the information learned in phase 1 as auxiliary input). Notice that $\frac{1}{p}(\nu - \kappa)^d$ is non-negligible. However, the CA only signed $(\texttt{Alice}, a)$ with $a < 21$. This implies that $\mathsf{K}^{\texttt{Alice}}$ can produce with non-negligible probability a valid signature of a message that has never been signed by the CA. This contradicts the security of the signature scheme (assuming, e.g., existential unforgeability [GMR88]). This shows the security of our toy protocol.

Amplification. Our toy example shows that QPoKs with negligible knowledge error can

[1] This is one of different possible definitions, loosely following [HM98]. It permits us to avoid the use of expected polynomial-time. We discuss alternatives in Section 2.2.

be used to show the security of protocols. But what about QPoKs with non-negligible knowledge error? In the classical case, we know that the knowledge error of a PoK can be made exponentially small by sequential repetition. Fortunately, this result carries over to the quantum case; its proof follows the same lines.

Elementary constructions. In order to understand our constructions of QPoKs, let us first revisit a common method for constructing classical PoKs. Assume a protocol that consists of three messages: the commitment (sent by the prover), the challenge (picked from a set C and sent by the verifier), and the response (sent by prover). Assume that there is an efficient algorithm K_0 that computes a witness given two conversations with the same commitment but different challenges; this property is called special soundness. Then we can construct the following (classical) extractor K: $\mathsf{K}^{\mathsf{P}^*}$ runs P^* using a random challenge ch. Then $\mathsf{K}^{\mathsf{P}^*}$ rewinds P^* to the point after it produced the commitment, and then $\mathsf{K}^{\mathsf{P}^*}$ runs P^* with a random challenge ch'. If both executions lead to an accepting conversation, and $ch \neq ch'$, K_0 can compute a witness. The probability of getting two accepting conversations can be shown to be $\Pr_{\mathsf{V}}^2$, where $\Pr_{\mathsf{V}}$ is the probability of the verifier accepting P^*'s proof. From this, a simple calculation shows that the knowledge error of the protocol is $1/\#C$.

If we directly translate this approach to the quantum setting, we end up with the following extractor: K runs one step of P^*, measures the commitment com, provides a random challenge ch, runs the second step of P^*, measures the response, runs the inverse of the second step of P^*, provides a random challenge ch', runs the second step of P^*, and measures the response $resp'$. If $ch \neq ch'$, and both $(com, ch, resp)$ and $(com, ch', resp')$ are accepting conversations, then we get a witness using K_0. We call this extractor the canonical extractor. The problem is to bound the probability F of getting two accepting conversations. In the classical setting, one uses that the two conversations are essentially independent (given a fixed commitment), and each of them is, from the point of view of P^*, the same as an interaction with the honest verifier V. In the quantum setting, this is not the case. Measuring $resp$ disturbs the state of P^*; we hence cannot make any statement about the probability that the second conversation is accepting.

How can we solve this problem? Note that we cannot use Watrous' oblivious rewinding since we need to remember both responses $resp$ and $resp'$ from two different execution paths of P^*. Instead, we observe that, the more information we measure in the first conversation (i.e., the longer $resp$ is), the more we destroy the state of P^* used in the second conversation. Conversely, if would measure only one bit, the disturbance of P^*'s state would be small enough to still get a sufficiently high success probability. But if $resp$ would contain only one bit, it would clearly be too short to be of any use for K_0. Yet, it turns out that this conflict can be resolved: In order not to disturb P^*'s state, we only need that the $resp$ information-theoretically contains little information. For K_0, however, even an information-theoretically determined $resp$ is still useful; it might, for example, reveal a value which P^* was already committed to. To make use of this observation, we introduce an additional condition on our proof systems, strict soundness. A proof system has strict

soundness if for any commitment and challenge, there is at most one response that makes the conversation accepting. Given a proof system with special and strict soundness, we can show that measuring *resp* does not disturb P^*'s state too much; the canonical extractor is successful with probability approximately $\Pr_{\mathsf{V}}^3$. A precise calculation shows that a proof system with special and strict soundness has knowledge error $1/\sqrt{\#C}$.

QPoKs for all languages in NP. Blum [Blu86] presents a classical zero-knowledge PoK for showing the knowledge of a Hamiltonian cycle. Using a suitable commitment scheme (it should have the property that the opening information is uniquely determined by the commitment), the proof system is easily seen to have special and strict soundness, thus it is a QPoK. By sequential repetition, we get a QPoK for Hamiltonian cycles. Using the Watrous' results, we get that the QPoK is also zero-knowledge. Using the fact that the Hamiltonian cycle problem is NP-complete, we get zero-knowledge QPoKs for all languages in NP (assuming quantum one-way permutations).

1.2 Preliminaries

General. A non-negative function μ is called negligible if for all $c > 0$ and all sufficiently large k, $\mu(k) < k^{-c}$. A non-negative function μ is non-negligible if it is not negligible. $\oplus$ denotes the XOR operation on bitstrings. $\mathrm{E}[X]$ denotes the expected value of X. $\#C$ is the cardinality of the set C.

Quantum systems. We can only give a terse overview over the formalism used in quantum computing. For a thorough introduction, we recommend the textbook by Nielsen and Chuang [NC00, Chap. 1–2]. A (pure) state in a quantum system is described by a unit vector $|\Phi\rangle$ in some Hilbert space $\mathcal{H}$. We always assume a designated orthonormal basis for each Hilbert space, called the computational basis. The tensor product of several states (describing a joint system) is written $|\Phi\rangle \otimes |\Psi\rangle$. We write $\langle\Psi|$ for the linear transformation mapping $|\Phi\rangle$ to the scalar product $\langle\Psi|\Phi\rangle$. The norm $\||\Phi\rangle\|$ is defined as $\sqrt{\langle\Phi|\Phi\rangle}$. A unit vector is a vector with $\||\Phi\rangle\| = 1$. The Hermitean transpose of a linear operator A is written $A^\dagger$. A is called positive if $A = A^\dagger$ and $\langle\Phi|A|\Phi\rangle \geq 0$ for all $|\Phi\rangle$. The operator norm of A is $|||A||| := \sup_{|\Phi\rangle} \|A|\Phi\rangle\|$ with $|\Phi\rangle$ ranging over unit vectors; we call A bounded if $|||A|||$ exists.

2 Quantum Proofs of Knowledge

2.1 Definitions

Interactive machines. Intuitively, an *interactive quantum machine* M (machine, for short) is a machine that maintains two quantum registers, a register S for the internal state of M, and a register N for sending and receiving messages (the network register). Upon each activation, M expects some message in N, and the state of the preceding invocation

in S. After the activation, S contains the new state of M, and N contains the message that M sends. A machine M can get both a classical input x and a quantum input $|\Phi\rangle$. For simplicity, we assume that the number of messages a machine sends and receives is determined by the classical input. The quantum input is initially stored in S. More formally, a quantum machine is described by a family of quantum circuits $(\mathsf{M}_x)_{x\in\{0,1\}^*}$ and a family of integers $(r_x^{\mathsf{M}})_{x\in\{0,1\}^*}$. M_x determines the unitary operation that is performed on the quantum registers S and N, and r_x^{M} determines the number of messages. Note that all our machines perform only unitary operations. This does not, however, constitute a restriction since a machine with measurements can be transformed into a unitary machine by a standard purification argument. We call a machine M *polynomial-time* if the circuit M_x has polynomial-size in $|x|$, r_x^{M} is polynomially-bounded in $|x|$, and the circuit's description can be computed in deterministic polynomial time given x.

Execution of interactive machines. Given a pair of machines M and M', a pair of quantum states $|\Phi\rangle$ and $|\Phi'\rangle$, and a pair of classical bitstrings $x, x' \in \{0,1\}^*$, we define the execution $\langle \mathsf{M}(x,|\Phi\rangle), \mathsf{M}'(x',|\Phi'\rangle)\rangle$ by the following process: Initialize quantum registers S, S', N with $|\Phi\rangle, |\Phi'\rangle, |0\rangle$, respectively. Alternatingly, apply the circuit M_x to S, N and the circuit $\mathsf{M}'_{x'}$ to S', N. Stop applying M_x after r_x^{M} applications and stop applying $\mathsf{M}_{x'}$ after $r_{x'}^{\mathsf{M}'}$ applications.[2] Then measure S' in the computational basis. The random variable $\langle \mathsf{M}(x,|\Phi\rangle), \mathsf{M}'(x',|\Phi'\rangle)\rangle$ denotes the result of that measurement. In other words, $\langle \mathsf{M}(x,|\Phi\rangle), \mathsf{M}'(x',|\Phi'\rangle)\rangle$ is the classical output of M' in an interaction where M is activated first. Often, we will omit the quantum input $|\Phi\rangle$ or $|\Phi'\rangle$. In this case, we assume the input $|0\rangle$.

Oracle algorithms with rewinding. An *quantum oracle algorithm* A is an algorithm that has access to a quantum interactive machine that is given as an oracle. Besides its own (classical) input x, the algorithm gets access to an interactive quantum machine M running on classical input x' and quantum input $|\Phi\rangle$. We allow A to provide messages to and read messages from $\mathsf{M}_{x'}$ and to execute the (unitary) quantum circuit $\mathsf{M}_{x'}$ that describes M. Furthermore, A may execute the inverse of $\mathsf{M}_{x'}$, this corresponds to the classical notion of rewinding the machine M. We also allow that A is in a superposition between executing $\mathsf{M}_{x'}$ and not executing it.[3] We will not, however, allow A to directly access the state of M or to its quantum input. (I.e., A has no access to the internal state and the quantum input of the prover. Any access to this information is done by communicating with M.) Formally, a quantum oracle algorithm A is described by a family of circuits $(\mathsf{A}_x)_{x\in\{0,1\}^*}$ operating on three quantum registers S_{A}, N and S_{M}. (S_{A} and S_{M} contain the states of A and M, respectively. N is used for communication between A and M.) The circuit A_x may contain normal gates (from some fixed universal set of gates) operating on S_{A} and N (but not

[2] If r_x^{M} and $r_{x'}^{\mathsf{M}'}$ do not match, it may happen that the circuit of one machine is executed several times in a row after the other machine already stopped.

[3] The ability of A to execute M_x in superposition is not, however, necessary for the results presented in this work.

S_{M}), as well as two special gates $\square$ and $\square^\dagger$. (These represent an application of the oracle given to A.) Both operate on one qubit of S_{A} (the control qubit) and on the whole of N, S_{M}. We define an execution $\mathsf{A}^{\mathsf{M}(x',|\Phi\rangle)}(x)$ as follows: Initialize $S_{\mathsf{A}}, N, S_{\mathsf{M}}$ with $|0\rangle, |0\rangle, |\Phi\rangle$. Execute the circuit A_x. When the gate $\square$ is to be applied on C, N, S_{M} where C is a qubit in S_{A}, apply the unitary transformation U defined by $U(|0\rangle \otimes |\psi\rangle \otimes |\varphi\rangle) := |0\rangle \otimes |\psi\rangle \otimes |\varphi\rangle$ and $U(|1\rangle \otimes |\psi\rangle \otimes |\varphi\rangle) := |1\rangle \otimes \mathsf{M}_{x'}(|\psi\rangle \otimes |\varphi\rangle)$ where $\mathsf{M}_{x'}$ is the unitary transformation describing one activation of M. (Intuitively, $\mathsf{M}_{x'}$ is applied if C contains $|1\rangle$.) The gate $\square^\dagger$ is treated analogously, except that we use $\mathsf{M}_{x'}^\dagger$ instead of $\mathsf{M}_{x'}$. Finally, we measure S_{A} in the computational basis. The random variable $\mathsf{A}^{\mathsf{M}(x',|\Phi\rangle)}(x)$ describes the outcome of that measurement. We call an algorithm A *polynomial-time* if the circuit A_x has polynomial-size in $|x|$ and its description can be computed in deterministic polynomial time given x.

Proof systems. A *quantum proof system* for a relation R is a pair of two machines (P, V). We call P the prover and V the verifier. The prover expects a classical input (x, w) with $(x, w) \in R$, the verifier expects only the input x. We call (P, V) *complete* if there is a negligible function μ such that for all $(x, w) \in R$, we have that $\Pr[\langle \mathsf{P}(x, w), \mathsf{V}(x)\rangle = 1] \geq 1 - \mu(|x|)$. (Remember that, if we do not explicitly specify a quantum input, we assume the quantum input $|0\rangle$.) Although we allow P and V to be quantum machines, and in particular to send and receive quantum messages, we will not need this property in the following; all protocols constructed in this paper will consist of classical machines. We call a (P, V) *sound* with soundness error s iff for all malicious prover P^*, all auxiliary inputs $|\Phi\rangle$, and all x with $\nexists w : (x, w) \in R$, we have $\Pr[\langle \mathsf{P}^*(x, |\Phi\rangle), \mathsf{V}(x)\rangle = 1] \leq s(|x|)$. A proof system is computational zero-knowledge iff for all polynomial-time verifiers V^* there is a polynomial-time machine S (the simulator) such that for all auxiliary inputs $|\Phi\rangle$, and all $(x, w) \in R$, we have that the quantum state of V^* after an interaction $\langle \mathsf{P}(x, w), \mathsf{V}^*(x, |\Phi\rangle)\rangle$ is computationally indistinguishable from the output of $\mathsf{S}(x, |\Phi\rangle)$; we refer to [Wat09] for details.

Quantum Proofs of Knowledge. We can now define quantum proofs of knowledge (QPoKs). Roughly, a quantum proof system (P, V) is a QPoK if there is a quantum oracle algorithm K (the extractor) that achieves the following: Whenever some malicious prover P^* convinces V that a certain statement holds, the extractor $\mathsf{K}^{\mathsf{P}^*}$ with oracle access to P^* is able to return a witness. Here, we allow a certain knowledge error κ; if P^* convinces V with a probability smaller than κ, we do not require anything. Furthermore, we also do not require that the success probability of $\mathsf{K}^{\mathsf{P}^*}$ is as high as the success probability of P^*; instead, we only require that it is polynomially related. Finally, to facilitate the use of QPoKs as subprotocols, we give the malicious prover an auxiliary input $|\Phi\rangle$. We get the following definition:

Definition 1 (Quantum Proofs of Knowledge) *We call a proof system* (P, V) *for a relation* R quantum extractable with knowledge error κ *if there exists a constant* $d > 0$*, a polynomially-bounded function* $p > 0$*, and a polynomial-time quantum oracle machine* K

such that for any interactive quantum machine P^**, any state* $|\psi\rangle$*, and any* $x \in \{0,1\}^*$*, we have that*

$$\Pr[\langle \mathsf{P}^*(x,|\psi\rangle), \mathsf{V}(x)\rangle = 1] \geq \kappa(|x|) \implies$$
$$\Pr[(x,w) \in R : w \leftarrow \mathsf{K}^{\mathsf{P}^*(x,|\psi\rangle)}(x)] \geq \tfrac{1}{p(|x|)} \Big(\Pr\big[\langle \mathsf{P}^*(x,|\psi\rangle), \mathsf{V}(x)\rangle = 1\big] - \kappa(|x|) \Big)^d.$$

A quantum proof of knowledge *for* R *with knowledge error* κ *(QPoK, for short) is a complete quantum extractable proof system for* R *with knowledge error* κ*.*

Note that by quantifying over all unitary provers P^*, we implicitly quantify over *all* purifications of *all* possible non-unitary provers. Note that extractability with knowledge error κ implies soundness with soundness error κ. We thus do not need to explicitly require soundness in Definition 1. The knowledge error κ can be made exponentially small by sequential repetition:

2.2 Discussion

In this section, we motivate various design choices made in the definition of QPoKs.

Access to the black-box prover's state and input. The extractor has no access to the prover's state nor to its quantum input. (This is modeled by the fact that an oracle algorithm may not apply any gates except for $\Box, \Box^\dagger$ to the register containing the oracle's state and quantum input.) In this, we follow [BG93] who argue in Section 4.3 that a proof of knowledge is supposed to "capture the knowledge of the prover *demonstrated by the interaction*" and that thus the extractor is not supposed to see the internal state of the prover. We stress, however, that our results are independent of this issue; they also hold if we allow the extractor to access the prover's state directly.

Unitary & invertible provers – technical view. Probably the most important design choice in our definition is to require the prover to be a unitary operation, and to allow the extractor to also execute the inverse of this operation. We begin with a discussion of this design choice from a technical point of view. First, we stress that seems that these assumptions are necessary: Since in a quantum world, making a snapshot/copy of a state is not possible or even well-defined, we have to allow the extractor to run the prover "backwards". But the inverse of a non-unitary quantum operation does not, in general, exist. Thus rewinding seems only possible with respect to unitary provers. Second, the probably most important question is: Does the definition, from an operational point of view, make sense? That is, does our definition behave well in cryptographic, reduction-based proofs? A final answer to this question can only be given when more protocols using QPoKs have been analyzed. However, the toy protocol discussed on page 4 gives a first indication that our definition can be used in a similar fashion to classical proofs of knowledge. Third, we would like to remind the reader that any non-unitary prover can be transformed into a

unitary one by purification before applying the definition of QPoKs. Thus allowing only unitary malicious provers does not seem to be a restriction in practice.

Unitary & invertible provers – philosophical view. Intuitively, a QPoK should guarantee that a prover that convinces the verifier "knows" the witness.[4] The basic idea is that if an extractor can extract the witness without using anything that is not available to the prover, then the prover "knew" the witness (or could have computed it). In particular, we may allow the extractor to run a purified (unitary) version of the prover because the prover himself could have done so. Similarly for the inverse of that operation. Of course, this leaves the question why we give these two capabilities to the extractor but not others (e.g., access to the circuit of the prover)? We would like to stress that analogous questions are still open (from a philosophical point) even in the classical case: Why is it natural to allow an extractor to rewind the prover? Why is it natural to give a trapdoor for a common reference string to the extractor? We would like to point out one justification for the assumption that the prover is unitary, though: [BG93] suggests that we "capture the knowledge of the prover *demonstrated by the interaction*". A prover that performs non-unitary operations is identical in terms of its interaction to one that is purified. Thus, by restricting to unitary provers, we come closer to only capturing the interaction but not the inner workings of the prover.

On the success probability of the extractor. We require the extractor to run in polynomial-time and to succeed with probability $\frac{1}{p}(\Pr_V - \kappa)^d$ where $\Pr_V$ is the probability that the prover convinces the verifier. (We call this an A-style definition.) In classical PoKs, a more common definition is to require the extractor to have expected runtime $\frac{p}{\Pr_V - \kappa}$ and to succeed with probability 1. (We call this a B-style definition.) This definition is known to be equivalent to the definition in which the extract runs in expected polynomial-time and succeeds with probability $\frac{1}{p}(\Pr_V - \kappa)$. (We call this a C-style definition.) The advantage of an A-style definition (which follows [HM98]) is that we can consider polynomial-time extractors (instead of expected polynomial-time extractors). To get extractors for B-style and C-style definitions, one has to increase the success probability of an extractor by repeatedly invoking it until it outputs a correct witness. In the quantum case, however, this does not work directly: If the invoked extractor fails once, the auxiliary input of the prover is destroyed. The oblivious rewinding technique by Watrous' would seem to help here, but when trying to apply that technique one gets the requirement that the invoked extractors' success probability must be independent of the auxiliary input. This condition is not necessarily fulfilled. To summarize, all three styles of definitions have their advantages, but it is not clear how one could fulfil B- and C-style definitions in the quantum case. This is why we chose an A-style definition. There are, however, applications that would benefit from a proof system fulfilling a C-style definition. For example, general multi-party computation protocols such as [GMW87] use extractors as part of the construction

[4] We believe, though, that this issue is secondary to the technical suitability; it is much more important that a QPoK is useful as a cryptographic subprotocol.

of the simulator for the multi-party computation; these extractors must then succeed with probability close to 1. Another example of a protocol needing C-style extractors is the proof for graph-non-isomorphism [GMW91], see the discussion in [HM98]. We leave the construction of C-style QPoKs as an open problem.

2.3 Amplification

In some cases, elementary constructions only yield QPoKs with constant knowledge error κ. Yet, in most cases we need QPoKs with negligible knowledge error. One possibility to construct these is to sequentially iterate a QPoK with constant knowledge error, the knowledge error of the resulting QPoK then becomes exponentially small. This result is well-known in the classical case [BG93]; the proof in the quantum case follows the same lines.

Theorem 2 *Let n be a polynomially bounded and efficiently computable function. Let (P,V) be extractable with knowledge error κ. Let $(\mathsf{P}',\mathsf{V}')$ be the proof system consisting of n-sequential executions of (P,V). Then $(\mathsf{P}',\mathsf{V}')$ is extractable with knowledge error κ^n.*

Proof. We call (P,V) the atomic proof and $(\mathsf{P}',\mathsf{V}')$ the composed proof. Fix a malicious prover P^* (that is supposed to interact with V'), a statement x, and an auxiliary input $|\Phi\rangle$ for P^*. In the execution of the composed proof with prover P^*, we call each execution of the atomic proof a round. Without loss of generality, we can assume that P^* consists of n sequentially executed provers P_i^* such that P_i^* executes the i-th round of the composed proof. For $i \geq 2$, P_i^* expects as quantum input the state that was output by P_{i-1}^*. Let K be the knowledge extractor for the atomic proof. We construct a knowledge extractor K' for the composed proof as follows: First, K' picks a random $i \in \{1,\dots,n\}$. Then K' internally simulates the first $i-1$ rounds of the composed proof (with provers $\mathsf{P}_1^*,\dots,\mathsf{P}_{i-1}^*$). Let $|\Phi'\rangle$ denote the state output by P_{i-1}^*. (And $|\Phi'\rangle := |\Phi\rangle$ if $i = 1$.) Then K' runs $w \leftarrow \mathsf{K}^{\mathsf{P}_i^*(x,|\Phi'\rangle)}(x)$ and outputs w.[5]

We fix the following notation: a_i is the probability that the first i rounds of the composed proof succeed (with prover P^*). We stress that a_{i-1} is also the probability that in an execution of K', the internal simulation of the first $i-1$ rounds succeeds. Let c_i denote the probability that the i-th round of the composed proof succeeds, conditioned on the event that the first $i-1$ rounds succeed. We have $a_0 = 1$ and $a_i = c_i a_{i-1}$ for $i = 1,\dots,n$.

Let $\Pr_{\mathsf{K}'}$ denote the probability that K' succeeds (i.e., returns a witness), and let $\Pr_{\mathsf{V}'}$ denote the probability that the composed proof succeeds. Fix some i. Let $\Pr_{\mathsf{K}'}^{(i)}$ denote the probability that K' succeeds, conditioned on the fact that K' chooses that i. Then, by construction of K', we have that $\Pr_{\mathsf{K}'} = \sum_{i=1}^n \frac{1}{n} \Pr_{\mathsf{K}'}^{(i)} \geq \max_i \frac{1}{n} \Pr_{\mathsf{K}'}^{(i)}$. We will show that

[5]Note that K as defined and analyzed here is not a unitary algorithm, but instead performs random choices and measurement. Since any such K can be converted into a unitary one by purification, we can use a non-unitary K without loss of generality.

there exists an i (dependent on P^*, $|\Phi\rangle$, and x), as well as a polynomially-bounded $p > 0$ and an integer $d > 0$ (independent of i, P^*, $|\Phi\rangle$, and x) such that $\Pr^{(i)}_{\mathsf{K}'} \geq \frac{1}{p}(\Pr_{\mathsf{V}'} - \kappa^n)^d$. This implies that $\Pr_{\mathsf{K}'} \geq \frac{1}{pn}(\Pr_{\mathsf{V}'} - \kappa^n)^d$. Thus $(\mathsf{P}', \mathsf{V}')$ has knowledge error κ.

We proceed to bound $\Pr^{(i)}_{K'}$ in terms of a_{i-1} and c_i. Let $\mathcal{D}_{i-1}$ denote the probability distribution of the output state of P^*_{i-1} conditioned on the event that the first $i-1$ rounds of the composed proof succeed. Let $\Pr^{(i)}_{\mathsf{K}}(|\Phi'\rangle)$ denote the probability that $\mathsf{K}^{\mathsf{P}^*(x,|\Phi'\rangle)}(x)$ succeeds (outputs a witness), and $\Pr^{(i)}_{V}(|\Phi'\rangle)$ the probability that the atomic proof with prover P^* and auxiliary input $|\Phi'\rangle$ succeeds. Then, the probability that K' succeeds, conditioned on the event that the first $i-1$ rounds succeed, is $\mathrm{E}[\Pr^{(i)}_{\mathsf{K}}(|\Phi'\rangle)]$ where $|\Phi'\rangle$ is distributed according to $\mathcal{D}_{i-1}$. Hence $\Pr^{(i)}_{\mathsf{K}'} = a_{i-1}\,\mathrm{E}[\Pr^{(i)}_{\mathsf{K}}(|\Phi'\rangle)]$. Since the atomic proof has knowledge error κ, there are a polynomially-bounded $p > 0$ and an integer $d > 0$ such that $\Pr^{(i)}_{\mathsf{K}}(|\Phi'\rangle) \geq \frac{1}{p}\big(\Pr^{(i)}_{\mathsf{V}}(|\Phi'\rangle) - \kappa\big)^d$ for all $|\Phi'\rangle$. We stress that p and d are independent of i, P^*, $|\Phi\rangle$, and x. It follows that

$$\begin{aligned}\Pr^{(i)}_{\mathsf{K}'} &= a_{i-1}\,\mathrm{E}\big[\Pr^{(i)}_{\mathsf{K}}(|\Phi'\rangle)\big] \geq a_{i-1}\,\mathrm{E}\big[\tfrac{1}{p}\big(\Pr^{(i)}_{\mathsf{V}}(|\Phi'\rangle) - \kappa\big)^d\big] \\ &\overset{(*)}{\geq} a_{i-1}\tfrac{1}{p}\big(\mathrm{E}[\Pr^{(i)}_{\mathsf{V}}(|\Phi'\rangle)] - \kappa\big)^d = a_{i-1}\tfrac{1}{p}\big(c_i - \kappa\big)^d.\end{aligned}$$

Here $(*)$ uses Jensen's inequality [Jen06].

Summarizing, at this point we know that $\Pr_{\mathsf{K}'} \geq \max_i \frac{1}{n}\Pr^{(i)}_{\mathsf{K}'} \geq \max_i \frac{a_{i-1}}{pn}\big(c_i - \kappa\big)^d$, that $a_i = c_i a_{i-1}$ for all i, and that $\Pr_{\mathsf{V}'} = a_n$.

Let $\delta := \Pr_{\mathsf{V}'} - \kappa^n$. Assume that $\delta > 0$, otherwise nothing need to be shown. Since $a_0 = 1$ and $a_n = \Pr_{\mathsf{V}'}$, we have that for some $i \in \{1, \ldots, n\}$, $a_{i-1} < \kappa^{i-1} + \frac{(i-1)\delta}{n}$ and $a_i \geq \kappa^i + \frac{i\delta}{n}$. For that i, we have

$$a_{i-1}(c_i - \kappa) = a_i - a_{i-1}\kappa \geq (\kappa^i + \tfrac{i\delta}{n}) - (\kappa^i + \tfrac{(i-1)\delta}{n}) = \tfrac{\delta}{n}$$

and hence

$$\Pr_{\mathsf{K}'} \geq \max_i \frac{a_{i-1}}{pn}\big(c_i - \kappa\big)^d \geq \max_i \frac{1}{pn} a_{i-1}^d \big(c_i - \kappa\big)^d \geq \max_i \frac{1}{pn}(\tfrac{\delta}{n})^d = \frac{1}{pn^{d+1}}(\Pr_{\mathsf{V}'} - \kappa^n)^d.$$

Since pn^{d+1} is polynomially-bounded, it follows that the composed proof $(\mathsf{P}', \mathsf{V}')$ has knowledge error κ^n. □

3 Elementary constructions

In this section, we show that under certain conditions, a classical PoK is also a QPoK (i.e., secure against malicious quantum provers). The first condition refers to the outer form of the protocol; we require that the proof systems is a protocol with three messages (commitment, challenge, and response) with a public-coin verifier. Such protocols

are called Σ-protocols. Furthermore, we require that the proof system has special soundness. This means that given two accepting conversations between prover and verifier that have the same commitment but different challenges, we can efficiently compute a witness. Σ-protocols with special soundness are well-studied in the classical case; many efficient classical protocols with these properties exist. The third condition (strict soundness) is non-standard. We require that given the commitment and the challenge of a conversation, there is at most one response that would make the verifier accept. We require strict soundness to ensure that the response given by the prover does not contain too much information; measuring it will then not disturb the state of the prover too much. Not all known protocols have strict soundness (the proof for graph isomorphism [GMW91] is an example). Fortunately, many protocols do satisfy strict soundness; a slight variation of the proof for Hamiltonian cycles [Blu86] is an example (see Section 4).

Definition 3 (Σ-protocol) *A proof system (P,V) is called a Σ-protocol if P and V are classical, the interaction consists of three messages $com, ch, resp$ (sent by P, V, and P, respectively, and called* commitment*,* challenge*, and* response*), and ch is uniformly chosen from some set C_x (the* challenge space*) that may only depend on the statement x. Furthermore, the verifier decides whether to accept or not by a deterministic polynomial-time computation on $x, com, ch, resp$. (We call $(com, ch, resp)$ an* accepting conversation for x *if the verifier would accept it.) We also require that it is possible in polynomial time to sample uniformly from C_x, and that membership in C_x should be decidable in polynomial time.*

Definition 4 (Special soundness) *We say a Σ-protocol (P,V) for a relation R has* special soundness *if there is a deterministic polynomial-time algorithm K_0 (the* special extractor*) such that the following holds: For any two accepting conversations $(com, ch, resp)$ and $(com, ch', resp')$ for x such that $ch \neq ch'$ and $ch, ch' \in C_x$, we have that $w := \mathsf{K}_0(x, com, ch, resp, ch', resp')$ satisfies $(x, w) \in R$.*

Definition 5 (Strict soundness) *We say a Σ-protocol (P,V) has* strict soundness *if for any two accepting conversations $(com, ch, resp)$ and $(com, ch, resp')$ for x, we have that $resp = resp'$.*

Canonical extractor. Let (P,V) be a Σ-protocol with special soundness and strict soundness. Let K_0 be the special extractor for that protocol. We define the *canonical extractor* K for (P,V). K will use measurements, even though our definition of quantum oracle algorithms only allows for unitary operations. This is only for the sake of presentation; by purifying K one can derive a unitary algorithm with the same properties. Given a malicious prover P^*, $\mathsf{K}^{\mathsf{P}^*(x,|\Phi\rangle)}(x)$ operates on two quantum registers N, S_{P^*}. N is used for communication with P^*, and S_{P^*} is used for the state of P^*. As described in the definition of quantum oracle machines, the registers N, S_{P^*} are initialized with $|0\rangle, |\Phi\rangle$. Let P^*_x

denote the unitary transformation describing a single activation of P. First, K applies P^*_x to N, S_{P^*}. (This can be done using the special gate $\square$.) This corresponds to running the first step of P^*; in particular, N should now contain the commitment. Then K measures N in the computational basis; call the result *com*. Then K initializes N with $|0\rangle$. Then K chooses uniformly random values $ch, ch' \in C_x$. Let U_{ch} denote the unitary transformation operating on N such that $U_{ch}|x\rangle = |x \oplus ch\rangle$. Then K applies $\mathsf{P}^*_x U_{ch}$. (Now N is expected to contain the response for challenge *ch*.) Then K measures N in the computational basis; call the result *resp*. Then K applies $(\mathsf{P}^*_x U_{ch})^\dagger$ (we rewind the prover). Then $\mathsf{P}^*_x U_{ch'}$ is applied. (Now N is expected to contain the response for challenge ch'.) Then N is measured in the computational basis; call the result *resp*$'$. Then $(\mathsf{P}^*_x U_{ch'})^\dagger$ is applied. Finally, K outputs $w := \mathsf{K}_0(x, com, ch, resp, ch', resp')$.

Analysis of the canonical extractor. In order to analyze the canonical extractor (Theorem 8 below), we first need a lemma that bounds the probability that two consecutive binary measurements P_{ch} and $P_{ch'}$ with random $ch \neq ch'$ succeed in terms of the probability that a single such measurement succeeds. In a classical setting (or in the case of commuting measurements), the answer is simple: the outcomes of the measurements are independent; thus the probability that two measurements succeed is the square of the probability that a single measurement succeeds. In the quantum case, however, the first measurement may disturb the state; this makes the analysis considerably more involved.

Lemma 6 *Let C be a set with $\#C = c$. Let $(P_i)_{i\in C}$ be orthogonal projectors on a Hilbert space $\mathcal{H}$. Let $|\Phi\rangle \in \mathcal{H}$ be a unit vector. Let $V := \sum_{i\in C} \frac{1}{c}\|P_i|\Phi\rangle\|^2$ and $F := \sum_{i,j\in C} \frac{1}{c^2}\|P_iP_j|\Phi\rangle\|^2$. Then $F \geq V^3$.*

Proof. To prove the lemma, we first show two simple facts:

Claim 1 *For any positive operator A on $\mathcal{H}$ and any unit vector $|\Phi\rangle \in \mathcal{H}$, we have that $(\langle\Phi|A|\Phi\rangle)^3 \leq \langle\Phi|A^3|\Phi\rangle$.*

Since A is positive, it is diagonalizable. Thus we can assume without loss of generality that A is diagonal (by applying a suitable basis transform to A and $|\Phi\rangle$). Let a_i be the i-th diagonal element of A, and let f_i be the i-th component of $|\Phi\rangle$. Then

$$(\langle\Phi|A|\Phi\rangle)^3 = \Big(\sum_i |f_i|^2 a_i\Big)^3 \overset{(*)}{\leq} \sum_i |f_i|^2 a_i^3 = \langle\Phi|A^3|\Phi\rangle.$$

Here $(*)$ uses Jensen's inequality [Jen06] and the facts that $a_i \geq 0$, that $a_i \mapsto a_i^3$ is a convex function on nonnegative numbers, and that $\sum_i |f_i|^2 = 1$. This concludes the proof of Claim 1.

Claim 2 *For vectors $|\Psi_1\rangle, \ldots, |\Psi_c\rangle \in \mathcal{H}$, it holds that $\|\frac{1}{c}\sum_i |\Psi_i\rangle\|^2 \leq \frac{1}{c}\sum_i \||\Psi_i\rangle\|^2$.*

To show the claim, let $|\bar{\Psi}\rangle := \sum_i \frac{1}{c}|\Psi_i\rangle$. Then

$$\begin{aligned}
\sum_i \Big(\||\Psi_i\rangle\|^2 - \||\bar{\Psi}\rangle\|^2 \Big) &= \sum_i \Big(\||\Psi_i\rangle\| - \||\bar{\Psi}\rangle\| \Big) \Big(\||\Psi_i\rangle\| - \||\bar{\Psi}\rangle\| + 2\||\bar{\Psi}\rangle\| \Big) \\
&= \sum_i \Big(\||\Psi_i\rangle\| - \||\bar{\Psi}\rangle\| \Big)^2 + 2\||\bar{\Psi}\rangle\| \sum_i \Big(\||\Psi_i\rangle\| - \||\bar{\Psi}\rangle\| \Big) \\
&\geq 2\||\bar{\Psi}\rangle\| \sum_i \Big(\||\Psi_i\rangle\| - \||\bar{\Psi}\rangle\| \Big) = 2\||\bar{\Psi}\rangle\| \Big(\sum_i \||\Psi_i\rangle\| - \|n|\bar{\Psi}\rangle\| \Big) \qquad (1) \\
&= 2\||\bar{\Psi}\rangle\| \Big(\sum_i \||\Psi_i\rangle\| - \Big\| \sum_i |\Psi_i\rangle \Big\| \Big) \qquad (2)
\end{aligned}$$

From the triangle inequality, it follows that $\sum_i \||\Psi_i\rangle\| \geq \|\sum_i |\Psi_i\rangle\|$, hence with (2), we have $\sum_i \Big(\||\Psi_i\rangle\|^2 - \||\bar{\Psi}\rangle\|^2 \Big) \geq 0$. Since $\frac{1}{c}\sum_i \||\Psi_i\rangle\|^2 - \|\frac{1}{c}\sum_i |\Psi_i\rangle\|^2 = \frac{1}{c}\sum_i \Big(\||\Psi_i\rangle\|^2 - \||\bar{\Psi}\rangle\|^2 \Big) \geq 0$, Claim 2 follows.

We proceed to prove Lemma 6. Let $A := \sum_i \frac{1}{c} P_i$, let $|\Psi_{ij}\rangle := P_j P_i |\Phi\rangle$. Then A is positive. Furthermore,

$$\begin{aligned}
V^3 &= \Big(\sum_i \tfrac{1}{c} \langle \Phi | P_i | \Phi \rangle \Big)^3 = \big(\langle \Phi | A | \Phi \rangle \big)^3 \overset{(*)}{\leq} \langle \Phi | A^3 | \Phi \rangle = \sum_{i,j,k} \tfrac{1}{c^3} \langle \Phi | P_i P_j P_k | \Phi \rangle \\
&= \sum_{i,j,k} \tfrac{1}{c^3} \langle \Psi_{ij} | \Psi_{kj} \rangle = \sum_j \tfrac{1}{c} \Big(\sum_i \tfrac{1}{c} \langle \Psi_{ij} | \Big) \Big(\sum_k \tfrac{1}{c} |\Psi_{kj}\rangle \Big) = \sum_j \tfrac{1}{c} \Big\| \sum_i \tfrac{1}{c} |\Psi_{ij}\rangle \Big\|^2 \\
&\overset{(**)}{\leq} \sum_j \tfrac{1}{c} \sum_i \tfrac{1}{c} \||\Psi_{ij}\rangle\|^2 = F.
\end{aligned}$$

Here $(*)$ uses Claim 1 and $(**)$ uses Claim 2. Thus we have $F \geq V^3$ and Lemma 6 follows.

Lemma 7 *Let C be a set with $\#C = c$. Let $(P_i)_{i \in C}$ be orthogonal projectors on a Hilbert space $\mathcal{H}$. Let $|\Phi\rangle \in \mathcal{H}$ be a unit vector. Let $V := \sum_{i \in C} \frac{1}{c} \|P_i|\Phi\rangle\|^2$ and $E := \sum_{i,j \in C, i \neq j} \frac{1}{c^2} \|P_i P_j |\Phi\rangle\|^2$. Then, if $V \geq \frac{1}{\sqrt{c}}$, $E \geq V(V^2 - \frac{1}{c})$.*

Proof. Let F be as in Lemma 6. Then

$$\begin{aligned}
E &= \sum_{\substack{i,j \in C \\ i \neq j}} \frac{1}{c^2} \|P_i P_j |\Phi\rangle\|^2 = \sum_{i,j \in C} \frac{1}{c^2} \|P_i P_j |\Phi\rangle\|^2 - \sum_{i \in C} \frac{1}{c^2} \|P_i P_i |\Phi\rangle\|^2 \\
&\overset{(*)}{=} \sum_{i,j \in C} \frac{1}{c^2} \|P_i P_j |\Phi\rangle\|^2 - \sum_{i \in C} \frac{1}{c^2} \|P_i |\Phi\rangle\|^2 = F - \frac{V}{c} \overset{(**)}{\geq} V^3 - \frac{V}{c} = V(V^2 - \tfrac{1}{c})
\end{aligned}$$

Here $(*)$ uses that $P_i = P_i P_i$ since P_i is a projection, and $(**)$ uses Lemma 6. □

Theorem 8 *A Σ-protocol (P,V) for a relation R with special and strict soundness and challenge space C_x is extractable with knowledge error $\frac{1}{\sqrt{\#C_x}}$.*

Proof. To show that (P,V) is extractable, we will use the canonical extractor K. Fix a malicious prover P^*, a statement x, and an auxiliary input $|\Phi\rangle$. Let $\Pr_{\mathsf{V}}$ denote the probability that the verifier accepts when interacting with P^*. Let $\Pr_{\mathsf{K}}$ denote the probability that $\mathsf{K}^{\mathsf{P}^*(x,|\Phi\rangle)}(x)$ outputs some w with $(x,w)\in R$. We will show that $\Pr_{\mathsf{K}} \geq \Pr_{\mathsf{V}}\cdot(\Pr_{\mathsf{V}}^2 - \frac{1}{\#C_x})$. For $\Pr_{\mathsf{V}} \geq \frac{1}{\sqrt{\#C_x}}$, we have that $\Pr_{\mathsf{V}}(\Pr_{\mathsf{V}}^2 - \frac{1}{\#C_x}) \geq (\Pr_{\mathsf{V}} - \frac{1}{\sqrt{\#C_x}})^3$. Since furthermore K is polynomial-time, this implies that (P,V) is a extractable with knowledge error $\frac{1}{\sqrt{\#C_x}}$.

In order to show $\Pr_{\mathsf{K}} \geq \Pr_{\mathsf{V}}\cdot(\Pr_{\mathsf{V}}^2 - \frac{1}{\#C_x})$, we will use a short sequence of games. Each game will contain an event Succ, and in the first game, we will have $\Pr[\mathsf{Succ} : \text{Game } 1] = \Pr_{\mathsf{K}}$. For any two consecutive games, we will have $\Pr[\mathsf{Succ} : \text{Game } i] \geq \Pr[\mathsf{Succ} : \text{Game } i+1]$, and for the final game, we will have $\Pr[\mathsf{Succ} : \text{Game } 7] \geq \Pr_{\mathsf{V}}\cdot(\Pr_{\mathsf{V}}^2 - \frac{1}{\#C_x})$. This will then conclude the proof. The description of each game will only contain the changes with respect to the preceding game.

Game 1. An execution of $\mathsf{K}^{\mathsf{P}^*(x,|\Phi\rangle)}(x)$. Succ denotes the event that K outputs a witness for x. By definition, $\Pr_{\mathsf{K}} = \Pr[\mathsf{Succ} : \text{Game } 1]$.

Game 2. Succ denotes the event that $(com, ch, resp)$ and $(com, ch', resp')$ are accepting conversations for x and $ch \neq ch'$. (The variables $(com, ch, resp)$ and $(com, ch', resp')$ are as in the definition of the canonical extractor.) Since (P,V) has special soundness, if Succ occurs, K outputs a witness. Thus $\Pr[\mathsf{Succ} : \text{Game } 1] \geq \Pr[\mathsf{Succ} : \text{Game } 2]$.

Game 3. Before K measures $resp$, it first measures whether measuring $resp$ would yield an accepting conversation. More precisely, it measures N with the orthogonal projector P_{ch} projecting onto $V_{ch} := \operatorname{span}\{|resp\rangle : (com, ch, resp) \text{ is accepting}\}$. Analogously for the measurement of $resp'$ (using the projector $P_{ch'}$.) Since a complete measurement (of $resp$ and $resp'$, respectively) is performed on N after applying the measurement P_{ch} and $P_{ch'}$, introducing the additional measurements does not change the outcomes $resp$ and $resp'$ of these complete measurements, nor their post-measurement state. Thus $\Pr[\mathsf{Succ} : \text{Game } 2] = \Pr[\mathsf{Succ} : \text{Game } 3]$.

Game 4. Succ denotes the event that both measurements P_{ch} and $P_{ch'}$ succeed. By definition of these measurements, this happens iff $(com, ch, resp)$ and $(com, ch', resp')$ are accepting conversations. Thus $\Pr[\mathsf{Succ} : \text{Game } 3] = \Pr[\mathsf{Succ} : \text{Game } 4]$.

Game 5. We do not execute K_0, i.e., we stop after applying $(\mathsf{P}_x^* U_{ch'})^\dagger$. Since at that point, Succ has already been determined, $\Pr[\mathsf{Succ} : \text{Game } 4] = \Pr[\mathsf{Succ} : \text{Game } 5]$.

Game 6. We remove the measurements of $resp$ and $resp'$. Note that the outcomes of these measurements are not used any more. Since (P,V) has strict soundness, $V_{ch} = \operatorname{span}\{|resp_0\rangle\}$ for a single value $resp_0$ (depending on com and ch, of course). Thus if the measurement P_{ch} succeeds, the post-measurement state in N is $|resp_0\rangle$. That is, the state

in N is classical at this point. Thus, measuring N in the computational basis does not change the state. Hence, the measurement of *resp* does not change the state. Analogously for the measurement of *resp'*. It follows that $\Pr[\mathsf{Succ} : \text{Game 5}] = \Pr[\mathsf{Succ} : \text{Game 6}]$.

Game 7. First, N and S_{P^*} are initialized with $|0\rangle$ and $|\Phi\rangle$. Then the unitary transformation P^*_x is applied. Then *com* is measured (complete measurement on N), and N is initialized to $|0\rangle$. Random $ch, ch' \in C_x$ are chosen. Then $\mathsf{P}^*_x U_{ch}$ is applied. Then the measurement P_{ch} is performed. Then $(\mathsf{P}^*_x U_{ch})^\dagger$ is applied. Then $\mathsf{P}^*_x U_{ch'}$ is applied. Then the measurement $P_{ch'}$ is performed. Then $(\mathsf{P}^*_x U_{ch'})^\dagger$ is applied. The event Succ holds if both measurements succeed. Games 6 and 7 are identical; we have just recapitulated the game for clarity. Thus, $\Pr[\mathsf{Succ} : \text{Game 6}] = \Pr[\mathsf{Succ} : \text{Game 7}]$.

In Game 7, for some value d, let p_d denote the probability that $com = d$ is measured. Let $|\Phi_d\rangle$ denote the state of N, S_{P^*} after measuring $com = d$ and initializing N with $|0\rangle$. (I.e., the state directly before applying $\mathsf{P}^*_x U_{ch}$.) Let K_d denote the probability that starting from state $|\Phi_d\rangle$, both measurements P_{ch} and $P_{ch'}$ succeed. Let $c := \#C_x$. Then we have that $\Pr[\mathsf{Succ} : \text{Game 7}] = \sum_d p_d K_d$ and

$$K_d = \sum_{ch,ch' \in C_x} \frac{1}{c^2} \|(\mathsf{P}^*_x U_{ch'})^\dagger P_{ch'} (\mathsf{P}^*_x U_{ch'}) (\mathsf{P}^*_x U_{ch})^\dagger P_{ch} (\mathsf{P}^*_x U_{ch}) |\Phi_d\rangle\|^2 = \sum_{ch,ch' \in C_x} \frac{1}{c^2} \|P^*_{ch'} P^*_{ch} |\Phi_d\rangle\|^2$$

where $P^*_{ch} := (\mathsf{P}^*_x U_{ch})^\dagger P_{ch} (\mathsf{P}^*_x U_{ch})$. Since P_{ch} is an orthogonal projector and $\mathsf{P}^*_x U_{ch}$ is unitary, P^*_{ch} is an orthogonal projector. Let $\varphi(v) := v(v^2 - \frac{1}{c})$ for $v \in [\frac{1}{\sqrt{c}}, 1]$ and $\varphi(v) := 0$ for $v \in [0, \frac{1}{\sqrt{c}}]$. Then, by Lemma 7, $K_d \geq \varphi(V_d)$ for $V_d := \sum_{ch \in C_x} \frac{1}{c} \|P^*_{ch} |\Phi_d\rangle\|^2$.

Furthermore, by construction of the honest verifier V, we have that

$$\Pr_{\mathsf{V}} = \sum_d p_d \sum_{ch \in C_x} \tfrac{1}{c} \|P_{ch} \mathsf{P}^*_x U_{ch} |\Phi_d\rangle\|^2 \overset{(*)}{=} \sum_d p_d \sum_{ch \in C_x} \tfrac{1}{c} \|(\mathsf{P}^*_x U_{ch})^\dagger P_{ch} (\mathsf{P}^*_x U_{ch}) |\Phi_d\rangle\|^2 = \sum_d p_d V_d$$

where $(*)$ uses that $(\mathsf{P}^*_x U_{ch})^\dagger$ is unitary. Finally, we have

$$\Pr_{\mathsf{K}} = \Pr[\mathsf{Succ} : \text{Game 1}] \geq \Pr[\mathsf{Succ} : \text{Game 7}] = \sum_d p_d K_d \geq \sum_d p_d \varphi(V_d) \overset{(*)}{\geq} \varphi(\Pr_{\mathsf{V}}).$$

Here $(*)$ uses Jensen's inequality [Jen06] and the fact that φ is convex on $[0, 1]$. As discussed in the beginning of the proof, $\Pr_{\mathsf{K}} \geq \varphi(\Pr_{\mathsf{V}}) = \Pr_{\mathsf{V}} \cdot (\Pr_{\mathsf{V}}^2 - \frac{1}{c})$ for $\Pr_{\mathsf{V}} \geq \frac{1}{\sqrt{c}}$ implies that (P, V) is a QPoK with knowledge error $1/\sqrt{\#C_x}$.

3.1 On using existing bounds from the literature

Lemma 7 says that if a random measurement (out of some fixed family of measurements) is very likely to succeed, then two repeated random measurements are also likely to succeed. Intuitively, the reason is that if a measurement is very likely to succeed, it will not change

the state too much, and then a second measurement will still be likely to succeed. In this, Lemma 7 is similar to the "Almost As Good As New Lemma" [Aar05, Lemma 2.2] and the "Tender Measurement Lemma" [Win99, Lemma 1.5]. These lemmas show that if the outcome of a measurement is almost deterministic, then this measurement will not change the state very much. In fact, for the case $\#C = 2$, we can easily derive Lemma 7 (even with a slightly different bound) from the "Almost As Good As New Lemma":

Lemma 9 *Let C be a set with $\#C = c = 2$. Let $(P_i)_{i\in C}$ be orthogonal projectors on a Hilbert space $\mathcal{H}$. Let $|\Phi\rangle \in \mathcal{H}$ be a unit vector. Let $V := \sum_{i\in C} \frac{1}{c}\|P_i|\Phi\rangle\|^2$ and $E := \sum_{i,j\in C, i\neq j} \frac{1}{c^2}\|P_iP_j|\Phi\rangle\|^2$. Then, if $V \geq \frac{3}{4}$, we have $E \geq \frac{1}{4}(V - \frac{3}{4})$.*

Proof. Without loss of generality, $C = \{0, 1\}$. Let $p_i := \|P_i|\Phi\rangle\|^2$ and $|\Psi_i\rangle := P_i|\Phi\rangle/\sqrt{p_i}$ (the state after measuring with P_i).

The "Almost As Good As New Lemma" [Aar05, Lemma 2.2] guarantees that if a measurement succeeds with probability at least $1-\varepsilon$ for some ε, then the trace distance between the state before measuring and the state after measuring successfully is at most $\sqrt{\varepsilon}$. Thus, in our case, the trace distance between $|\Phi\rangle$ (the state before measuring) and $|\Psi_0\rangle$ (the state after measuring successfully) is at most $\sqrt{1-p_0}$.

Thus the difference in probability that measurement P_1 succeeds given $|\Phi\rangle$ and $|\Psi_0\rangle$, respectively, is at most $\sqrt{1-p_0}$. Hence $\|P_1|\Psi_0\rangle\|^2 \geq \|P_1|\Phi\rangle\| - \sqrt{1-p_0} = p_1 - \sqrt{1-p_0}$. Thus

$$\tfrac{1}{4}\|P_0P_1|\Psi_0\rangle\|^2 = \tfrac{1}{4}\|P_0|\Phi\rangle\|^2\|P_1|\Psi_0\rangle\|^2 \geq \tfrac{1}{4}p_0(p_1 - \sqrt{1-p_0}) =: B_0.$$

Analogously we get

$$\tfrac{1}{4}\|P_1P_0|\Psi_0\rangle\|^2 \geq \tfrac{1}{4}p_1(p_0 - \sqrt{1-p_1}) =: B_1.$$

Hence $E \geq B_0 + B_1$.

Since $V \geq \frac{3}{4}$, we have $p_0, p_1 \geq \frac{1}{2}$. Thus

$$E \geq B_0 + B_1 \geq \tfrac{1}{8}(p_1 - \tfrac{1}{\sqrt{2}}) + \tfrac{1}{8}(p_0 - \tfrac{1}{\sqrt{2}}) = \tfrac{1}{4}(V - \tfrac{1}{\sqrt{2}}) \geq \tfrac{1}{4}(V - \tfrac{3}{4}).$$

(Note that this bound for $B_0 + B_1$ is not tight, we could easily get a better knowledge error that $\frac{3}{4}$. However, since this lemma is for illustration purposes only, we opted for the simpler calculation.) □

Notice that using the "Almost As Good As New Lemma" (and the approach used in the proof of Lemma 9), we cannot significantly improve over the bound given in Lemma 9, even if $\#C$ is large: The proof idea of Lemma 9 for large $\#C$ is that the first measurement P_i will, on average, succeed with probability V. Thus the trace distance between the original state $|\Phi\rangle$ and the state $|\Psi_i\rangle$ after measuring P_i will be bounded by $\sqrt{1-V}$. Thus the difference between the probability that the second measurement P_j succeeds given $|\Phi\rangle$ and given $|\Psi_i\rangle$ is at most $\sqrt{1-V}$. Thus the probability for the second measurement to succeed

is at least $V - \sqrt{1-V}$. But this lower bound becomes negative for $V \leq \frac{\sqrt{5}-1}{2}$. Thus, we only get a positive lower bound for E if $V > \frac{\sqrt{5}-1}{2}$. In the setting of proofs of knowledge, this would mean that we end up with a constant knowledge error of at least $\frac{\sqrt{5}-1}{2}$, even for large $\#C$ (as opposed to $1/\sqrt{\#C}$ in Lemma 7). Similar problems apply if we try to use the "Tender Measurement Lemma" [Win99, Lemma 1.5].

4 QPoKs for all languages in NP

In the preceding section, we have seen that complete proof systems with strict and special soundness are QPoKs. The question that remains to be asked is: do such proof systems, with the additional property of being zero-knowledge, exist for interesting languages? In this section, we will show that for any language in NP (more precisely, for any NP-relation), there is a zero-knowledge QPoK. (Assuming the existence of quantum one-way permutations.) Here and in the following, by zero-knowledge we mean quantum computational zero-knowledge.

The starting point for our construction will be the Blum's zero-knowledge PoK for Hamiltonian cycles [Blu86]. In this Σ-protocol, the prover's commits to the vertices of a graph using a perfectly binding commitment scheme. In the prover's response, some of these commitments are opened. That is, the response contains the opening information for some of the commitments. The problem is that standard definitions of commitment schemes do not guarantee that the opening information is unique; only the actual content of the commitment has to be determined by the commitment. This means that the prover's response is not unique. Thus, with a standard commitment scheme we do not get strict soundness. Instead we need a commitment scheme such that the sender of the commitment scheme is committed not only to the actual content of the commitment, but also to the opening information.

Definition 10 (Strict binding) *A commitment scheme* COM *is a deterministic polynomial-time function taking two arguments a, y, the opening information a and the message y. We say* COM *is* strictly binding *if for all a, y, a', y' with $(a, y) \neq (a', y')$, we have that* $\mathrm{COM}(a, y) \neq \mathrm{COM}(a', y')$.

Furthermore, in order to get the zero-knowledge property, we will need that our commitment schemes are quantum computationally concealing. We refer to [Wat09] for a precise definition of this property. In [AC02], an unconditionally binding, quantum computationally concealing commitment scheme based on quantum one-way permutations is presented. Their definitions differ somewhat from those of [Wat09], but as mentioned in [Wat09], their proof carries over to the definitions from [Wat09]. Furthermore, in the scheme from [AC02], the commitment contains the image of the opening information under a quantum one-way permutation. Thus the strict binding property is trivially fulfilled. Thus strictly binding,

Inputs: A directed graph x (the statement) with vertices W, and a Hamiltonian cycle w in x (the witness).
Subprotocols: A strictly binding, quantum computationally concealing commitment scheme COM.
Protocol:

1. P picks a random permutation π on W. Let A be the adjacency matrix of the graph $\pi(x)$. Let $H := \{(\pi(i), \pi(j)) : (i,j) \in w\}$. Using COM, P commits to π, H, and to each entry A_{ij} of A. P sends the resulting commitments to V.
2. V picks $ch \in \{0,1\}$ and sends ch to P.
3. If $ch = 0$, P opens the commitments to π and A. If $ch = 1$, P opens the commitments to H and to all A_{ij} with $(i,j) \in H$.
4. If $\mathsf{ch} = 0$, V checks that the commitments are opened correctly, that π is a permutation, and that A is the adjacency matrix of $\pi(x)$. If $ch = 1$, V checks that the commitments are opened correctly, that H is a cycle, that exactly the A_{ij} with $(i,j) \in H$ are opened, and that $A_{ij} = 1$ for all $(i,j) \in H$. If all checks succeed, V outputs 1.

Figure 1: A QPoK (P, V) for Hamiltonian cycles.

quantum computationally concealing commitment schemes exist under the assumption that quantum one-way permutations exist.

Given such a commitment scheme COM, we construct the proof system (P, V) presented in Figure 1. Besides using a strictly binding commitment, (P, V) differs in one other aspect from the proof system in [Blu86]: The prover does not only commit to the vertices in the graph $\pi(x)$, but also to the permutation π and the cycle H. Without these additional commitments, we would not get strict soundness; there might be several permutations leading to the same graph, or the graph might contain several Hamiltonian cycles.

Theorem 11 *Let $(x,w) \in R$ iff w is a Hamiltonian cycles of the graph x. Assume that* COM *is a strictly binding, quantum computationally concealing commitment scheme. Then the proof system (P, V) is a zero-knowledge QPoK for R with knowledge error $\frac{1}{\sqrt{2}}$.*

Proof. We need to show completeness, extractability (via special and strict soundness), and zero-knowledge. Completeness is straightforward by inspection of the protocol.

Special soundness. Let $(com, ch, resp)$ and $(com, ch', resp')$ be two accepting conversations for x with $ch \neq ch'$. Without loss of generality, $ch = 0$ and $ch' = 1$. Then $resp$ contains a permutation π and the adjacency matrix A of $\pi(x)$. And $resp'$ contains a cycle H such that $\tilde{A}_{ij} = 1$ for all $(i,j) \in H$ where $\tilde{A}_{ij}$ are the committed values opened in $resp'$. Since ch is strictly binding, $1 = \tilde{A}_{ij} = A_{ij}$ for all $(i,j) \in H$, thus H is a Hamiltonian cycle of $\pi(x)$. Then $w := \mathsf{K}_0(x, com, ch, resp, ch', resp') := \pi^{-1}(H)$ is a Hamiltonian of x, i.e., $(x,w) \in R$.

Strict soundness. Fix an accepting conversation $(com, ch, resp)$. If $ch = 0$, *resp* consists only of opening of commitments. Since COM has strict binding, it follows that *resp* is uniquely determined by com, ch. If $ch = 1$, COM consists of an opening of the commitment to H, and of the commitments to A_{ij} with $(i, j) \in H$. Hence H and its opening information are uniquely determined since COM has strict binding, and thus it is also determined which A_{ij} are opened. Again by strict binding, the values A_{ij} and corresponding opening information are uniquely determined. Thus *resp* is uniquely determined by com, ch.

Extractability. Since (P, V) has special and strict soundness, and a challenge space of size 2, by Theorem 8, we have that (P, V) is extractable with knowledge error $\frac{1}{\sqrt{2}}$.

Zero-knowledge. We first describe an intermediate simulator S_1. Fix a malicious verifier V^*, some $(x, w) \in R$, and an auxiliary input $|\Phi\rangle$. $\mathsf{S}_1^{\mathsf{V}^*(x,|\Phi\rangle)}(x)$ first picks a random $ch^* \in \{0, 1\}$. If $ch^* = 0$, S_1 chooses a random permutation π, computes the adjacency matrix A of $\pi(x)$, and picks an arbitrary H. If $ch^* = 1$, S_1 chooses an arbitrary permutation π, sets A to be the all-one matrix, and lets H be a random cycle. Then S_1 sends the commitments to π, A, H to V^*. If V^* does not answer with the challenge ch^*, S_1 aborts. Otherwise, S_1 sends the response as specified in Figure 1 to V^*. Finally, S_1 outputs the (quantum) output of V^*.

Furthermore, let S_2 be the simulator that additionally gets w as input, and then behaves like S_1 except that it constructs π, A, H honestly (i.e., as specified in Figure 1). The probability that S_2 aborts is $\frac{1}{2}$. Furthermore, the quantum state output by S_2, conditioned on not aborting, is the same as the output of V^* in an interaction with honest P. Furthermore, since all commitments that are opened are constructed in the same way in S_1 and S_2, and since COM is quantum computationally concealing, the outputs of S_1 and S_2 are quantum computationally indistinguishable. Thus the probability that S_2 aborts is $\frac{1}{2} \pm \mu$ for some negligible μ, and the output of S_2, conditioned on not aborting, is quantum computationally indistinguishable from the output of V^* in an interaction with honest P.

Applying [Wat09, Lemma 9] (oblivious rewinding) to S_1, we get that there is a polynomial-time simulator S such that the output S is statistically indistinguishable from the output of S_1 conditioned on not aborting. Thus the output of S is quantum computationally indistinguishable from the output of V^* in an interaction with honest P. □

Corollary 12 (QPoKs for all languages in NP) *Let R be an NP-relation.*[6] *Then there is a zero-knowledge QPoK for R with negligible knowledge error.*

Proof. Using the fact that the Hamiltonian cycle problem is NP-complete, from Theorem 11 it follows that there is a zero-knowledge QPoK for R with knowledge error $\frac{1}{\sqrt{2}}$. By sequential repetition, we get a QPoK for R with negligible knowledge error (Theorem 2). Sequential repetition preserves the zero-knowledge property (see [Wat09]). □

[6] An NP-relation is a relation R such that $(x, w) \in R$ is decidable in deterministic polynomial time, and there is a polynomial p such that for all $(x, w) \in R$, $|w| \leq p(|x|)$.

Acknowledgements. We thank the anonymous referees and Märt Põldvere for suggestions on how to significantly simplify the proof of Lemma 7. This research was supported by the Cluster of Excellence "Multimodal Computing and Interaction" and by European Social Fund's Doctoral Studies and Internationalisation Programme DoRa.

References

[Aar05] Scott Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1(1):1–28, 2005. Online available at `http://www.theoryofcomputing.org/articles/v001a001`.

[AC02] Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In *STACS 2002*, volume 2285 of *LNCS*, pages 323–334, Berlin, Heidelberg, 2002. Springer.

[BG93] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In Ernest F. Brickell, editor, *Advances in Cryptology, Proceedings of CRYPTO '92*, number 740 in Lecture Notes in Computer Science, pages 390–420. Springer-Verlag, 1993. Extended version online available at `http://www-cse.ucsd.edu/users/mihir/papers/pok.ps`.

[Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. In *Proceedings of the International Congress of Mathematicians*, pages 1444–1451, Berkeley, 1986.

[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 291–304. ACM Press, 1985.

[GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, 1988. Online available at `http://theory.lcs.mit.edu/~rivest/GoldwasserMicaliRivest-ADigitalSignatureSchemeSecureAgainstAdaptiveChosenMessageAttacks.ps`.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game – or – a completeness theorem for protocols with honest majority. In *STOC 87*, pages 218–229, 1987.

[GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(3):690–728, 1991. Online available at `http://www.wisdom.weizmann.ac.il/~oded/X/gmw1j.pdf`.

[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC*, pages 212–219, 1996.

[HM98] Shai Halevi and Silvio Micali. More on proofs of knowledge. IACR ePrint 1998/015, 1998.

[Jen06] Johan L. W. V. Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(1):175–193, 1906. In French.

[NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[Sho94] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science, Proceedings of FOCS 1994*, pages 124–134. IEEE Computer Society, 1994.

[vdG98] Jeroen van de Graaf. *Towards a formal definition of security for quantum protocols*. PhD thesis, Départment d'informatique et de r.o., Université de Montréal, 1998. Online available at `http://www.cs.mcgill.ca/~crepeau/PS/these-jeroen.ps`.

[Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, 2009.

[Win99] Andreas Winter. *Coding Theorems of Quantum Information Theory*. PhD thesis, Universität Bielefeld, 1999. arXiv:quant-ph/9907077v1.

[WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.

Index