IACR Communications in Cryptology
ISSN 3006-5496, Vol. 2, No. 1, 35 pages.

https://doi.org/10.62056/a6n56c0kr

The Round Complexity of Proofs in the Bounded Quantum Storage Model

Alex B. Grilo[1] and Philippe Lamontagne[2,3]

[1] Sorbonne Université, CNRS, LIP6, France
[2] National Research Council Canada, Canada
[3] Université de Montréal, Canada

Abstract. The round complexity of interactive proof systems is a key question of practical and theoretical relevance in complexity theory and cryptography. Moreover, results such as QIP = QIP(3) (STOC'00) show that quantum resources significantly help in such a task.

In this work, we initiate the study of round compression of protocols in the bounded quantum storage model (BQSM). In this model, the malicious parties have a bounded quantum memory and they cannot store the all the qubits that are transmitted in the protocol.

Our main results in this setting are the following:

1. There is a non-interactive (statistical) witness indistinguishable proof for any language in NP (and even QMA) in BQSM in the plain model. We notice that in this protocol, only the memory of the verifier is bounded.
2. Any classical proof system can be compressed in a two-message quantum proof system in BQSM. Moreover, if the original proof system is zero-knowledge, the quantum protocol is zero-knowledge too. In this result, we assume that the prover has bounded memory.

Finally, we give evidence towards the "tightness" of our results. First, we show that NIZK in the plain model against BQS adversaries is unlikely with standard techniques. Second, we prove that without the BQS model there is no 2–message zero-knowledge quantum interactive proof, even under computational assumptions.

1 Introduction

The round complexity of interactive proof systems[1] is a central question in complexity theory and cryptography. For example, while it is expected that not all interactive proof systems can be compressed to a constant number of rounds, showing such a result would have major implications in complexity theory such as P ≠ PSPACE. In cryptographic settings, the round complexity is very relevant to the practical applications of protocols, specially in the setting of zero-knowledge (ZK) proof systems [2]. While there exist 4-messages ZK protocols for NP [FS90b], it is known that 2-messages ZK protocols for NP

E-mail: Philippe.Lamontagne2@cnrc-nrc.gc.ca (Philippe Lamontagne)

[1] In an interactive proof system, an all-powerful prover wants convince a computationally bounded verifier that $x \in L$ for some language L by exchanging polynomially many messages. We want such interactive protocols such that the prover can convince the verifier if $x \in L$, whereas if $x \notin L$ the prover cannot convince the verifier except with negligible probability

[2] In a zero-knowledge interactive proof system, the verifier "learns nothing" from the interaction with the prover. This is formally defined as requiring the existence of a simulator which can produce the same output as the verifier, but without the help of the prover. Zero-knowledge proofs are extremely useful in building other cryptographic primitives, such as a maliciously secure multiparty computation [GMW86], IND-CCA encryption [BFM88], identification and digital signatures schemes [FS87].

Received: 2024-10-07 Accepted: 2025-03-11

are impossible [GO01]; and in specific settings such as black-box zero-knowledge, even 3-message private-coin protocols and constant-round public-coin protocols are known to be impossible [GK96]. These negative results on the round complexity can often be circumvented through additional resources. For example, in the random oracle model, any public-coin zero-knowledge proof can be made non-interactive through the use of the Fiat-Shamir heuristic [FS87]. While it has been shown that such a heuristic cannot be implemented in a black-box way [GK03; BDG$^+$13], it is possible to instantiate it in specific settings and achieve, for example, non-interactive ZK for NP in the common reference string (CRS) model from standard cryptographic assumptions [BFM88; PS19].

With the development of quantum computing, the notion of interactive protocols has been also extended to the quantum setting. Here, the prover and verifier are now allowed to exchange quantum messages back-and-forth to prove that $x \in L$. One of the first results in this direction already indicated that quantum resources are useful in reducing the rounds of protocols: it was shown that any quantum interactive protocol can be compressed to a 3-message protocol [KW00]. The natural question raised by such a result is the power of two-messages quantum interactive proof systems. More concretely, can we compress any quantum (or less ambitiously classical) protocol into a one-round protocol with quantum communication? This is tightly connected with the question of instantiation of Fiat-Shamir with quantum resources which was recently shown black-box impossible in [DLS12].

In this work, we make progress in this direction by studying round compression of protocols in the *bounded quantum storage model* (BQSM). In this model, we assume that the malicious parties have a bounded quantum memory. In particular, the quantum messages are transmitted in a sequential manner, qubit by qubit, and at every instant, the memory bound holds. We notice that in our protocols, the honest parties do not need quantum memory at all: they measure the qubits as soon as they are received. This model has been shown very powerful, allowing the implementation of several important cryptographic primitives with information-theoretic security [DFSS08; DFR$^+$07; DFSS07; BS06]. In this work, we show that such a powerful tool is also relevant for round-efficient interactive protocols. More concretely, we show the following:

1. There is a non-interactive (statistical) witness indistinguishable proof for any language in NP (and even QMA) in the *plain model* against BQS adversaries. We notice that in this protocol, only the memory of the verifier is bounded.

2. Any classical proof system can be compressed in a two-message quantum proof system in BQSM. Moreover, if the original proof system is zero-knowledge, the quantum protocol is zero-knowledge too. In this result, we assume that the prover has bounded memory.

We present now our results in more detail and give a brief overview on the techniques to prove them.

1.1 Our Results

As previously mentioned, in this work, we investigate the round complexity of proof systems in the bounded quantum storage (BQS) model. It is based on the physical assumption that the adversary has a bounded-size quantum memory of $q(\lambda)$ qubits where λ is the security parameter. Our main results are two compilers for reducing the round complexity of proofs in the BQS model. Each one operates differently and has its own applications. The bounded party differs in each of our main results; either the verifier or the prover has bounded quantum memory, but never both. The memory bound q on the malicious party is independent of the underlying proof system and a larger bound can be tolerated by increasing the size of the quantum messages.

Non-interactive proof for NP. In our first result, we provide a compiler NIP that takes a 3–message public-coin interactive proof system with 1-bit challenges and compresses it to one message. The main idea of the compiler is to use non-interactive oblivious transfer (OT) in non-interactive proofs, an idea which was introduced by [KMO90] and first appeared in writing in [BM90].

More concretely, the starting point of our protocol is the non-interactive quantum oblivious transfer protocol of [DFR+07] which is secure against BQS receivers. We can construct a non-interactive proof by having the prover send its first message[3] a in the clear and input the responses r_0, r_1 to both possible challenges $c \in \{0,1\}$ as its inputs to OT. Our compiler preserves the soundness of the underlying interactive proof, and it can be amplified through parallel repetition. Intuitively, the security of BQS-OT implies that a quantum memory bounded verifier will only receive one of the two transcripts, which reveals no information since accepting transcripts can be simulated if the underlying Σ-protocol is honest-verifier zero-knowledge.

While we manage to prove that the protocol satisfies the witness indistinguishable property, proving zero-knowledge is challenging since it is hard for the simulator to "decode" the measurements of a potentially malicious verifier. In particular, we prove in Section 3.1 that a "natural" simulation technique cannot work.

Result 1. *Let Π be a Σ–protocol. Then $\mathsf{NIP}[\Pi]$ preserves soundness and preserves witness indistinguishability against BQS verifiers.*

Our compiler can be extended in a trivial way to Σ–protocols with logarithmic challenge length (by using a 1-out-of-2^p OT with $p \in O(\lg(\lambda))$). Furthermore, the first message of the prover may be quantum, so our compiler can be applied to Ξ–protocols as long as the verifier is receive-and-measure. Our result thus implies a NIWI for QMA based on the Ξ–protocol from [BG22] which has short challenges and is receive-and-measure for the verifier.

This compiler allows us to achieve a non-interactive (statistically) witness indistinguishable proof for all languages in NP in the BQS model without any prior setup.

Result 2. *For any $L \in \mathsf{NP}$, there is a quantum non-interactive proof system for L with unconditional soundness and witness indistinguishability against BQS verifiers.*

To obtain Result 2, we apply our compiler to the typical proof system for the NP–complete language of graph Hamiltonicity. This would normally introduce a computational assumption on either the prover or the verifier since the proof uses a commitment scheme, however we can instead use a quantum bit commitment, which only needs to satisfy a very weak notion of binding.

A stronger notion than witness indistinguishability (yet still weaker than zero-knowledge) is witness hiding. We show that witness hiding can be preserved by our compiler in a regime where the soundness error is inverse polynomial. See Appendix B for details.

A Round Collapse Theorem in the BQSM. We show that under the BQS assumption, the round complexity of proof systems essentially collapses to two messages (one round). We present a *round reduction* compiler RR that takes as input a $\mathsf{poly}(\lambda)$ rounds interactive proof Π and produces a single round (2 messages) proof for the same language with the following properties.

Result 3. *Let Π be a $\mathsf{poly}(\lambda)$–round public-coin[4] interactive proof system, then there is a 1–round quantum interactive proof $\mathsf{RR}[\Pi]$ such that*

[3] The first message may be classical in a Σ–protocol or quantum, in which case we call it a Ξ–protocol [BG22].

[4] We actually only require that the verifier messages at each round are independent from the previously exchanged messages and do not use fact that they are uniformly distributed.

1. *soundness is preserved against BQS provers;*
2. *zero-knowledge is preserved.*

Our compiler RR is conceptually very simple. It relies on a distinctive property of the original bit commitment in the BQSM, in that the committer commits to a bit b by measuring a state it gets from the receiver. This allows us to remove one round of interaction by having the verifier send a state $|\psi\rangle$ for the commitment at the same time as its next challenge c_i (we assume this challenge is sampled independently of the prior messages). The prover commits to its message a_i by measuring $|\psi\rangle$, then receives c_i, and can respond with its next message a_{i+1}. Since the prover has bounded quantum memory, it will have to perform a (partial) measurement on $|\psi\rangle$ before receiving the verifier's challenge c_i. By the binding property of the commitment against BQS provers, this implies that a_i is independent of c_i, and thus any attack against this protocol is also an attack against Π. By repeating this technique for every round in protocol Π, we end up with a protocol with one round that has the same soundness error, plus a negligible term from the BQS-BC binding theorem.

By using the correspondence $\mathsf{IP} = \mathsf{PSPACE}$ [Sha92], we obtain the following.

Result 4. $\mathsf{PSPACE} = \mathsf{QIP}(2)^{\mathsf{BSQM}}$*, i.e., there exists a 2–message quantum protocol for every problem in* PSPACE *if the computationally unbounded prover has a bounded quantum memory.*

Furthermore applying our compiler to the doubly efficient protocols for delegation of classical computation [GKR15; RRR21], we achieve the following application.

Result 5. *In the BQSM, there is a quantum interactive protocol for any language in* P *such that the honest prover runs in polynomial time, the verifier runs in linear time and logarithmic space, and there is a single round of communication.*

By applying our compiler to a concrete scheme, we get the first 1–round interactive proof for NP that is both statistically sound (against BQS provers) and statistically ZK against arbitrary verifiers.

Other Contributions. We give evidence towards the "tightness" of our results. We show that NIZK in the plain model against BQS adversaries is unlikely with standard techniques. We also show that an assumption such as the BQSM is necessary for our round compression result by proving that there is no 2–message zero-knowledge quantum interactive proof system when the prover is not memory-bounded. This result is an extension of the impossibility of Goldreich and Oren [GO01] to the quantum case and is presented in Appendix C.

Our round reduction transform uses a string commitment built by parallel composition of the original BQS-BC scheme. To commit to $n(\lambda) \in O(\lambda)$ bit strings requires sending $\lambda \cdot n(\lambda) \in O(\lambda^2)$ qubits against a $O(\lambda)$–bounded adversary. Thus, the memory bound is sublinear in the number of transmitted qubits. In Appendix A, we propose a new string commitment where the length of committed strings, the number of transmitted qubits and the memory bound are all linear in λ. While we were unable to prove that this new commitment meets the definition of binding required by our RR transform, we can show that it is sum-binding, so it might be useful in improving the efficiency of other BQSM schemes.

1.2 Related Work

Classical non-interactive witness indistinguishable proof systems can be built from strong computational assumptions such as a derandomization circuit complexity assumption [BOV03; BP15] and the decision linear assumption on bilinear groups [GOS06].

Quantum NIZK for QMA can be achieved in the following models: in the secret parameter model [BG22], in the QROM with quantum preprocessing [MY22], in the designated verifier model [Shm21], using pre-shared EPR pairs and subexponential assumptions [BKS23], and with a CRS with an instance-dependent quantum message from the verifier to the prover [CVZ20]. While we call the bounded quantum storage assumption a "model", our results do not rely on any prior setup.

The bounded quantum storage model was introduced in [DFSS08] as a physical assumption upon which information theoretically secure two-party cryptographic primitives such as oblivious transfer (OT) and bit-commitment (BC) could be built. The BQSM has found further application to quantum key distribution [DFR+07; DFSS07] and to secure identification [DFSS07]. The noisy quantum storage model [WST05; STW09; KWW03] (NQSM) is a generalization of the BQSM, where the adversary's quantum memory is subject to noise, that enables protocols for OT and BC. There are OT protocols in the BQSM and NQSM where the tolerated bound or noise level is an arbitrary large fraction of the number of exchanged qubits [DFW02]. The model was recently exploited to achieve strong primitives such as one-time programs [BS06]. Composability frameworks have been proposed for the BQSM [Unr11; WW08; FS09]. These results and that of [BS06] require extracting the malicious party's input, which in general is inefficient. This is not a problem when the class of adversary is quantum memory bounded and computationally unbounded, but it doesn't work when simulation needs to be efficient, as in ZK proofs. Finally, post-quantum zero-knowledge against BQS adversaries was recently studied [AG22] in the context where all information exchanged by the parties is classical, but the adversaries may be quantum.

2 Preliminaries

For a set $\mathcal{S}$ we write $2^{\mathcal{S}}$ to denote the powerset, or set of subsets, of $\mathcal{S}$. We let $x \in_R \mathcal{S}$ denote that x is chosen uniformly at random in $\mathcal{S}$. We let $\Delta : \{0,1\}^n \times \{0,1\}^n \to [0,1]$ denote the relative Hamming distance between two n–bit strings. It is a well-known fact that for any $x \in \{0,1\}^n$, $|\{x' : \Delta(x,x') < \delta\}| \leq 2^{H(\delta)n}$ where H is the binary Shannon entropy. A *universal* set of hash function is functions $\mathcal{H}$ mapping n–bit strings to ℓ–bit strings such that for any $a, b \in \{0,1\}^n$, $\Pr_{h \in_R \mathcal{H}}[h(a) = h(b)] \leq 2^{-\ell}$.

We let $\{|0\rangle, |1\rangle\}$ denote the *computational* basis states for a single qubit register. The *Hadamard* basis is denoted $\{H|0\rangle, H|1\rangle\}$ where $H = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}$ is the Hadamard transform. We often specify the basis using a bit and write $|x\rangle_\theta := H^\theta|x\rangle$ for $x, \theta \in \{0,1\}$. We also use the "+" and "×" notation to refer to the computational and Hadamard basis, respectively (i.e. $|b\rangle_+ = |b\rangle$ and $|b\rangle_\times = H|b\rangle$ for $b \in \{0,1\}$).

A matrix $U \in \mathbb{C}^{n \times n}$ is *unitary* if $U^*U = UU^* = \mathbb{I}$ where U^* is the conjugate transpose of U. We refer to a general quantum state as a *density operator*, i.e. a positive semidefinite matrix $\rho \in \mathbb{C}^{n \times n}$ with trace equal to 1. Quantum transformations are modeled as completely positive trace-preserving (CPTP) maps, i.e. transforms that map density operators to density operators.

Throughout this paper, $\|\cdot\|$ denotes the trace norm $\|A\| = \mathrm{tr}(\sqrt{A^*A})$ when its argument is an operator and the Euclidean norm $\||\psi\rangle\| = \sqrt{\langle\psi|\psi\rangle}$ when its argument is a vector. The trace norm has the following properties. Let $\mathcal{E}$ be a quantum operation modeled as a completely positive trace-preserving (CPTP) map, then $\|\mathcal{E}(A)\| \leq \|A\|$. Let ρ, σ be two density operators, the maximal probability with which ρ can be distinguished from σ is $\frac{1}{2} + \frac{1}{4}\|\rho - \sigma\|$.

For a string $a \in \{0,1\}^n$, we let a_i^j for $1 \leq i < j \leq n$ denote the substring of a composed of the bits $a_i, \ldots, a_j$. When $a_1, \ldots, a_k$ are Boolean strings, we let $(a_i)_j$ denote the jth bit of the ith string.

The *min-entropy* of a classical random variable X conditioned on an event Ψ is $H_\infty(X|\Psi) = -\log \max_x \Pr[X = x|\Psi]$. Conditioned on a random variable Z, it is defined as $H_\infty(X|Z) := \min_z H_\infty(X|Z = z)$. The *max-entropy* of a quantum or classical register A in state ρ is $H_0(A)_\rho = \log \operatorname{rank}(\rho_A)$. A trivial upper-bound on $H_0(A)$ is $\dim A$. The *min-entropy splitting lemma* will also be useful. For a proof of this lemma, please refer to the full version of [DFR+07].

Lemma 1 (Min-entropy splitting)**.** *Let X_0, X_1 and Z be random variables with $H_\infty(X_0X_1|Z) \geq \alpha$. Then there exists a random variable C with support over $\{0,1\}$ such that $H_\infty(X_{1-C}|ZC) \geq \alpha/2 - 1$.*

2.1 The Bounded Quantum Storage Model

In the BQSM, the adversary has access to a quantum memory of at most q qubits. The assumption is that the bound q holds at every stage in the protocol. Quantum messages are transmitted sequentially (qubit by qubit). There are no other restrictions on the adversary, in particular they can store an unbounded number of classical bits and can perform arbitrarily long computations. The BQSM allows for the honest parties to have a (smaller than the adversary) quantum memory, and some recent works [BS06] exploit this. However, in this paper only the malicious parties are assumed to have a (bounded size) quantum memory. Our protocols are *prepare-and-measure* and only require the honest participants to prepare, send and measure qubits from the set of states $\{|0\rangle, |1\rangle, H|0\rangle, H|1\rangle\}$.

We assume that the parties share a single bidirectional *quantum* communication channel. Qubits arrive at their destination in the order that they are sent. A classical message is transmitted over this channel by encoding it in the computational basis ($|0\rangle$ and $|1\rangle$). This model of communication[5] is consistent with known methods for quantum communication via optical fibers (e.g. QKD). For the BSQM, it has the advantage that it removes timing issues with a separate classical channel where security does not hold if a classical message arrives prior to the quantum transmission.

We review two important protocols in the BQSM and state their security properties below.

2.1.1 Bit Commitment in the BQSM.

We begin by discussing the bit commitment scheme of [DFSS08]. One of the unique features of this protocol is that the *committer* commits to a bit through the measurement of a received quantum state, and does not need to send any message back to the receiver of the commitment. The opening phase consists of the transmission of the committed bit and along with measurement outcome, which will enable the consistency verification. This commitment scheme is perfectly hiding since no information is sent to the receiver prior to the reveal phase. The original [DFSS08] bit commitment protocol in the BQSM proceeds as described below.

Protocol DFSS-BC

Input: a bit $b \in \{0,1\}$ for the committer.

Commit phase:

[5] One could consider a different model where qubits don't necessarily arrive in order. Each qubit would need a "classical header" indicating their order, otherwise the receiver could not reorder them, but then the classical header would be transmitted in a separate channel of the quantum data. Thus, synchronizing them would actually be harder than sending qubits sequentially.

1. V sends $|x\rangle_\theta$ for $x \in_R \{0,1\}^n$ and $\theta \in_R \{+,\times\}^n$ to the committer.
2. C commits to bit b by measuring all qubits in basis $+$ if $b = 0$ and in basis $\times$ if $b = 1$, obtaining a measurement outcome x'.

Reveal phase:

1. To open the commitment, C sends b and x' to V who checks that $x'_i = x_i$ whenever $\theta_i = b$.

Binding Property of BC in the BQSM. Since unconditionally secure bit commitment is impossible in the quantum setting [LC01; May28], binding relies on the quantum storage bound of the malicious committer. A malicious committer $\tilde{\mathsf{C}}$ is bound to a single value by the fact that it is forced to perform a partial measurement on the register it receives. This notion is formalized by the following definition. We first introduce some notation. Let W be $\tilde{\mathsf{C}}$'s classical register, E be $\tilde{\mathsf{C}}$'s q–qubit quantum register and V be the receiver's state. The joint state ρ_{EWV} of the committer and receiver at the end of the commit phase after the memory bound is applied can be expressed as

$$\rho_{EWV} = \sum_{w,v} P_{WV}(w,v) \cdot \rho_E^{w,v} \otimes |w\rangle\langle w| \otimes |v\rangle\langle v| \tag{1}$$

where P_{WV} is some probability distribution and $\rho_E^{w,v}$ are density operators that may depend arbitrarily on w and v.

Definition 1. A commitment scheme in the bounded-quantum-storage model is called *ϵ-binding*, if for every (dishonest) committer $\tilde{\mathsf{C}}$, inducing a joint state ρ_{EWV} of the form of (1) after the commit phase, there exists a classical random variable B' with support in $\{0,1\}^n$, given by its conditional distribution $P_{B'|WV}$, such that for any $b' \in \{0,1\}^n$, the state

$$\rho_{EWV}^{b'} = \sum_{w,v} P_{WV|B'}(w,v|b') \cdot \rho_E^{w,v} \otimes |w\rangle\langle w| \otimes |v\rangle\langle v| \tag{2}$$

satisfies the following condition. When executing the opening phase on the state $\rho_{EWV}^{b'}$, for any strategy of $\tilde{\mathsf{C}}$, the honest verifier accepts an opening to $b \neq b'$ with probability at most ϵ.

It was shown in [DFR+07] that DFSS-BC satisfies the above definition (for $b \in \{0,1\}$). This implies a string commitment protocol where C commits bit-wise to b_i using protocol DFSS-BC.

Theorem 1 (Security of DFSS-BC)**.** *The quantum bit commitment scheme DFSS-BC is ϵ–binding according to Definition 1 against q–bounded committers where $\epsilon(n)$ is negligible in n if $n/4 - q \in \Omega(n)$.*

2.1.2 Oblivious Transfer in the BQSM.

The original OT protocol in the BQSM was a Rabin OT (where the sender has one input and the receiver gets to see it with probability $\frac{1}{2}$). We use the $\binom{2}{1}$–OT from [DFR+07] which is presented below. It is a *non-interactive* protocol which consists of a single message with quantum and classical parts from the sender to the receiver. The memory bound is applied after the transmission of the quantum state.

Protocol DFRSS-OT

Input: two bits $s_0, s_1 \in \{0,1\}^\ell$ for the sender. A bit $c \in \{0,1\}$ for the receiver.

Sender:

- Pick $x \in \{0,1\}^n$ and $\theta \in \{+,\times\}^n$.
- Pick two universal hash functions $h_0, h_1 \in \mathcal{H}$ and set $m_0 = s_0 \oplus h_0(x_0)$ and $m_1 = s_1 \oplus h_1(x_1)$ where x_0 (resp. x_1) is the substring of x for which $\theta_i = +$ (resp. $\times$).
- Prepare the quantum state $|x\rangle_\theta$ and the classical message $(\theta, h_0, h_1, m_0, m_1)$.
- Send $|x\rangle_\theta |\theta, h_0, h_1, m_0, m_1\rangle$ to the receiver.

Receiver:

- Measure each qubit of the first register in basis $[+,\times]_c$ to get a result x'. Measure the remaining registers in the computational basis.
- Compute x'_c using θ and output $m_c \oplus h_c(x'_c)$.

Correctness of the protocol follows from the fact that $x'_c = x_c$ if both parties follow the protocol. The security is established by the following result.

Theorem 2 (Security of DFRSS-OT [DFR+07])**.** *Let R be a malicious q-bounded receiver against ℓ–bit DFRSS-OT and let $\rho_{M_0 M_1 H_0 H_1 E}$ be the state of R right after the classical message from the sender (where* $\dim E \leq 2^q$*). Then there exists a random variable C such that*

$$\left\| \rho_{M_{1-C} M_C C H_0 H_1 E} - \frac{\mathbb{I}_{M_{1-C}}}{2^\ell} \otimes \rho_{M_C C H_0 H_1 E} \right\| \leq 2^{-\frac{n}{4}+\ell+q} \tag{3}$$

Parallel repetition of DFRSS-OT. While protocol DFRSS-OT does not generally compose in parallel[6], it does compose in the case where the same party is the sender in every instance. By parallel repetition, we mean the protocol where the sender prepares and sends a quantum state of the form $\bigotimes_i |x^i\rangle_{\theta^i}$ followed by $\bigotimes_i |\theta^i, h_0^i, h_1^i, m_0^i, m_1^i\rangle$ to the receiver (as opposed to the alternating quantum and "classical" messages that would occur in sequential repetition).

Corollary 1 (Parallel repetition of DFRSS-OT)**.** *Let R be a malicious q-bounded receiver against k parallel repetitions of ℓ–bit DFRSS-OT. Let $\rho_{\vec{M}_0 \vec{M}_1 \vec{H}_0 \vec{H}_1 E}$ be the state of R right after the sender's transmission (where* $\dim E \leq 2^q$*). Then there exist random variables $\vec{C} = C^1, \ldots, C^k$ such that*

$$\left\| \rho_{\vec{M}_{\neg\vec{C}} \vec{M}_{\vec{C}} \vec{C} \vec{H}_0 \vec{H}_1 E} - \frac{\mathbb{I}_{\vec{M}_{\neg\vec{C}}}}{2^{k\cdot\ell}} \otimes \mathrm{tr}_{\vec{M}_{\neg\vec{C}}} \left(\rho_{\vec{M}_{\vec{C}} \vec{C} \vec{H}_0 \vec{H}_1 E} \right) \right\| \leq k \cdot 2^{-\frac{n}{4}+\ell+q} \tag{4}$$

where $\vec{M}_{\neg\vec{C}}$ denotes registers $M^i_{1-C^i}$ for $i \in [k]$ and $M_{\vec{C}}$ denotes registers $M^i_{C^i}$.

Proof. Consider the purified variant of the scheme, where the sender sends halves of EPR pairs in the first step and measures its halves in basis θ. Consider k parallel executions of this purified scheme. Let X^i and Θ^i be the measurement result and basis for the ith repetition. The distribution (X^i, Θ^i) is independent from that of (X^j, Θ^j) for $j \neq i$. Since

[6] See [WW08] for a counter-example. The issue occurs when Alice acts as the receiver of an OT instance while simultaneously acting as the sender in another. If she receives a commitment to b from Bob, she can commit to b to Charlie by forwarding Bob's message.

$H_\infty(X^i|\Theta^i) \geq (\frac{1}{2} - \epsilon)n$, by the min-entropy splitting lemma there exists C^i such that M_{1-C^i} is indistinguishable from uniform. Note that the min-entropy bound holds even if we condition on the random variables from other executions and on the receiver's registers:

$$H_\infty(X^i|(\Theta^j)_j(X^j)_{j\neq i}ZE) \geq H_\infty(X^i|\Theta^i) - H_0(E) \geq (\frac{1}{2} - \epsilon)n - q$$

Also note that the random variable C^i depends only on the conditional distribution $P_{X^i|\Theta^i}$, so the C^is are simultaneously well-defined for each $i \in [k]$. We have that for each i,

$$\left\| \rho_{M^i_{1-C^i} M^i_{C^i} C^i H^i_0 H^i_1 E} - \frac{\mathbb{I}_{M^i_{1-C^i}}}{2^{k\cdot\ell}} \otimes \mathrm{tr}_{M^i_{1-C^i}}\left(\rho_{M^i_{C^i} C^i H^i_0 H^i_1 E}\right) \right\| \leq 2^{-\frac{n}{4}+\ell+q} \tag{5}$$

and, by starting with $\rho_{\vec{M}_{\neg\vec{C}}\vec{M}_{\vec{C}}\vec{C}\vec{H}_0\vec{H}_1E}$ and invoking the triangle inequality k times (where each time we replace $M^i_{1-C^i}$ with the completely mixed state), we get the corollary's statement. □

2.2 Quantum Interactive Proofs and Quantum Zero-Knowledge

An interactive proof system is a protocol between two participants, a prover P and a verifier V. We consider proofs of language membership where each participant receives a common input x, and the prover may receive an additional input w, such as a witness that x is a member of a NP language. A proof system is *classical* if the message exchanged are classical, but P and V are allowed to be quantum. We say that a classical or quantum proof system is *public coin* if the verifier's messages are uniformly and independently distributed.

We denote by $\mathsf{P}(x) \leftrightarrows \mathsf{V}(x)$ the output of the verifier after the interactive proof. An interactive proof system for a language L is δ–correct if for all $x \in L$,

$$\Pr[\mathsf{P}(x) \leftrightarrows \mathsf{V}(x) = 1] \geq \delta \ . \tag{6}$$

It is (computationally) ϵ–sound if for all (QPT) malicious prover $\tilde{\mathsf{P}}$, for all $x \notin L$,

$$\Pr[\tilde{\mathsf{P}}(x) \leftrightarrows \mathsf{V}(x) = 1] \leq \epsilon \ . \tag{7}$$

We now define quantum zero-knowledge [Wat01].

Definition 2 (Indistinguishability of Quantum States)**.** Let L be an infinite set of strings and let $\psi = \{\psi_x\}_{x\in L}$ and $\phi = \{\phi_x\}_{x\in L}$ be two families of quantum states. We say that ψ and ϕ are *computationally indistinguishable* if for all $x \in L$ for every $\mathsf{poly}(|x|)$–time quantum algorithm D and for all state σ over $\mathcal{H}^{\otimes \mathsf{poly}(|x|)}$,

$$\|\mathsf{D}(\psi_x \otimes \sigma) - \mathsf{D}(\phi_x \otimes \sigma)\| \leq \mathsf{negl}(|x|) \ .$$

ψ and ϕ are *statistically indistinguishable* the above holds with respect to all D and all states σ.

Definition 3 (Indistinguishability of Quantum Channels)**.** Let L be an infinite set of strings and let $\Psi = \{\Psi_x\}_{x\in L}$ and $\Phi = \{\Phi_x\}_{x\in L}$ be two families of CPTP maps agreeing on their input and output spaces: $\Psi_x, \Phi_x : \mathcal{H}^{\otimes n(|x|)} \to \mathcal{H}^{\otimes m(|x|)}$. We say that Ψ and Φ are *computationally indistinguishable* if for all $x \in L$, for every $\mathsf{poly}(|x|)$–time quantum algorithm $\mathsf{D} : \mathcal{H}^{\otimes m(|x|)+k(|x|)} \to \mathcal{H}$ and for every $\sigma \in \mathcal{H}^{\otimes m(|x|)+k(|x|)}$,

$$\|\mathsf{D}(\Psi_x \otimes \mathbb{I}^{\otimes k(|x|)}(\sigma)) - \mathsf{D}(\Phi_x \otimes \mathbb{I}^{\otimes k(|x|)}(\sigma))\| \leq \mathsf{negl}(|x|) \tag{8}$$

where $m(|x|)$, $n(|x|)$ and $k(|x|)$ are $\mathsf{poly}(|x|)$. Ψ and Φ are *statistically indistinguishable* if the above holds with respect to all CPTP map D and all states σ (for unbounded $k(|x|)$).

Definition 4 (Quantum Zero-Knowledge)**.** An interactive proof system $\Pi = \langle \mathsf{P}, \mathsf{V} \rangle$ for a language L is *computationally quantum zero-knowledge* (qZK) if for every $\mathsf{poly}(|x|)$–time verifier V^* receiving the common input $x \in L$, there exists a $\mathsf{poly}(|x|)$–time simulator $\mathsf{Sim}_{\mathsf{V}^*}$ that receives the same inputs and such that the quantum channel families $\{\mathsf{P} \leftrightarrows \mathsf{V}(x, \cdot))\}_{x \in L}$ and $\{\mathsf{Sim}_{\mathsf{V}^*}(x, \cdot)\}_{x \in L}$ are computationally indistinguishable. We say that Π is *statistically quantum zero-knowledge* if the two channel families are statistically indistinguishable. We say it is (computationally or statistically) *quantum honest verifier zero-knowledge* (qHVZK) if indistinguishability holds with respect to the honest verifier $\mathsf{V}^* = \mathsf{V}$.

We point out that the concept of quantum zero-knowledge can be extended to quantum memory bounded verifiers (see, e.g. [AG22]). When the verifier is limited in some way (in memory or computation time), the goal is for the simulator to mimic the verifier's actions with comparable resources. For our results of Section 4, we only need to bound the memory of adversaries in the soundness against malicious provers, while the zero-knowledge property holds even against verifiers that have an unbounded quantum memory. We notice that in this case, it is natural that Sim is allowed to have unbounded quantum memory as well.

Definition 5 (Ξ–protocols)**.** A Ξ–protocol for a language L is an interactive proof system $\Pi = (\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ with the following structure.

1. The prover receives as input x and a witness $|w\rangle$, computes $|\phi\rangle_{AB} \leftarrow \mathsf{P}_1(x)$ and sends ϕ_A to the verifier.
2. The verifier chooses a uniformly random challenge $c \in \{0,1\}^\ell$ and sends c to the prover.
3. The prover computes $r \leftarrow \mathsf{P}_2(x, \phi_B, c)$ and sends r to the verifier.
4. The verifier accepts if $\mathsf{V}(x, \phi_A, c, r) = 1$ and rejects otherwise.

A Σ–protocol is a Ξ–protocol where $|w\rangle$ and $|\phi\rangle_{AB}$ are classical. A Ξ–protocol is *prepare-and-measure* for the verifier if the verifier measures ϕ_A upon reception in a basis chosen by c and the predicate V is applied on the measurement outcome.

3 Non-Interactive Proofs in the BQSM

We present a generic transform to turn arbitrary Σ–protocols with small challenge space to non-interactive proofs. We actually consider a slight generalization of Σ–protocols where the first message send by the prover can be a quantum state, while the challenge by the verifier should be uniformly random bits and the third message by the prover is classical. Broadbent and Grilo [BG22] called this type of protocols as Ξ–protocols (Definition 5), and we will use their notation to stress that the first message can be quantum.

The soundness of our transform does not rely on any setup assumption. We will show later that while we cannot show zero-knowledge for such a transform, we can prove some weaker notions. We notice that since we are working in the bounded storage model, we consider Ξ protocols where an honest verifier measures the qubits of the first message as they arrive based on the chosen challenge.

Protocol $\mathsf{NIP}[\Pi]$ for a Ξ–protocol Π

Prerequisite: A 3–message, 1-bit public coin, interactive proof $\Xi = (\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$.

Prover:

1. For $i \in [k]$,
   - 1.1 Compute $|\phi_i\rangle \leftarrow \mathsf{P}_1$ with fresh randomness each time
   - 1.2 Prepare the n–qubit state $|x_i\rangle_{\theta_i}$
   - 1.3 Compute responses r_i^c using P_2 for $c \in \{0,1\}$
   - 1.4 Sample two universal$_2$ hash functions h_i^0 and h_i^1
   - 1.5 Compute $m_i^0 = r_i^0 \oplus h_i^0(x_i^+)$ and $m_i^1 = r_i^1 \oplus h_i^1(x_i^\times)$ where x_i^+ (resp. $x_i^\times$) is the substring of x_i corresponding to the $+$ (resp. $\times$) basis.
2. Send $\bigotimes_i |\phi_i\rangle|x_i\rangle_{\theta_i}|\theta_i, h_i^0, h_i^1, m_i^0, m_i^1\rangle$ to the verifier.

Verifier:

3. Pick a k random selection bits $c_1, \ldots c_k$
4. For $i \in [k]$,
   - 4.1 Measure $|\phi_i\rangle$ according to V to get an outcome a_i
   - 4.2 Measure $|x_i\rangle_{\theta_i}$ on basis c_i getting x_i'
   - 4.3 Compute $x_i^{c_i}$ from θ_i and x_i'; and compute r^{c_i} from $x_i^{c_i}$, $h_i^{c_i}$ and $m_i^{c_i}$
   - 4.4 Check that for all $i \in [k]$ $\mathsf{V}(a_i, c_i, r_i^{c_i}) = 1$, otherwise abort

The soundness of the protocol follows from the fact the prover is oblivious to which response the verifier has learned. Since BQS-OT is secure against unbounded senders, the soundness of $\mathsf{NIP}[\Pi]$ is unconditionally reducible to the soundness of Π.

We notice that this technique can be used to compress $\log n$ rounds protocols with $\log n$ bit challenges by using $\mathsf{poly}(n)$ instances of OT. Let's say for simplicity that we have a k rounds protocol with 1 bit challenges with $k \in O(\log n)$, and have access to a $\binom{2^k}{1}$–OT. Then for each of the 2^k inputs $0 \leq j < 2^k$, the prover sends the transcript it would produce if j's bits were the challenges. We can extend this to $m = O(\log n)$ bit challenges by considering $j \in \{0,1\}^{2^{m+k}}$.

Theorem 3. *Let Π be a 1-bit challenge Ξ–protocol with soundness $\frac{1}{2}$ against quantum adversaries. Then $\mathsf{NIP}[\Pi]$ is an unconditionally sound quantum non-interactive proof with soundness error $\frac{1}{2^k}$.*

Proof. Let $\mathcal{A}$ be a malicious prover against $\mathsf{NIP}[\Pi]$. We construct a reduction $\mathcal{R}$ against (the k-wise parallel repetition of) Π that has the same success probability as $\mathcal{A}$. The reduction simulates the OT instances while being able to recover both messages sent by $\mathcal{A}$ by having a sufficiently large quantum memory. Thus we reduce to the soundness of Π against quantum adversaries.

When $\mathcal{A}$ sends the quantum state $\rho_{A_1 X_1 \ldots A_k X_k}$ in the first step, where A_i is the register that is supposed to have the state $|\phi_i\rangle$ and X_i is the register supposed to have $|x_i\rangle_{\theta_i}$, $\mathcal{R}$ stores the qubits. When $\mathcal{A}$ sends the classical message $(\theta_i, h_i^0, h_i^1, m_i^0, m_i^1, a_i)_{i\in[k]}$, $\mathcal{R}$ measures the register X_i in basis θ_i and compute the responses to each possible challenge r_i^0 and r_i^1.

Now $\mathcal{R}$ acts as the sender in protocol Π^k. It sends the state $\rho_{A_1 \ldots A_k}$ as the first message of the Π protocol. Upon reception of the challenges $c_1, \ldots, c_k \in \{0,1\}$ from the verifier, $\mathcal{R}$ replies with $r_1^{c_1}, \ldots r_k^{c_k}$.

It remains to argue that the verifier accepts in Π^k against $\mathcal{R}$ with the same probability that the verifier accepts in $\mathsf{NIP}[\Pi]$ against $\mathcal{A}$. This follows from the observation that a_i, c_i and $r_i^{c_i}$ are identically distributed in both cases. Therefore, the success probability of $\mathcal{R}$ against Π^k is exactly that of $\mathcal{A}$ against $\mathsf{NIP}[\Pi]$. □

Remark 1. We notice that the soundness of $\mathsf{NIP}[\Pi]$ actually follows from a weaker notion of soundness that we call *oblivious soundness*, which intuitively says that the Prover cannot simultaneously answer the two challenges. More concretely, $\mathsf{NIP}[\Pi]$ is sound if Π has the following property: for any no instance $x \notin L$ and first message ρ, no prover can create, at the same time, a valid answer for $c = 0$ *and* a valid answer for $c = 1$. More concretely, for all possible values (r_0, r_1)

$$\sup_{M} \sum_{b\in\{0,1\}} \sum_{r_0,r_1} \mathrm{tr}\left((M_{r_0,r_1} \otimes V(x, b, r_b))\rho_{AB}\right) \leq 1 + \mathsf{negl}(n). \tag{9}$$

where M_{r_0,r_1} consists of a measurement made by the prover to answer r_0 to the first challenge and r_1 to the second challenge. While this property is implied by standard soundness, we will see a protocol later in this section that only satisfies oblivious soundness.

3.1 Security Against Malicious Verifier

We now turn to the security against malicious verifiers of $\mathsf{NIP}[\Pi]$. A verifier with an arbitrarily large quantum memory may postpone its measurement and learn both transcripts of Π. If for example Π is special-sound, it would allow them to recover an NP witness. Thus, we focus the security against bounded quantum storage verifiers. The question remains as to exactly what properties can be proven in this setting.

In this section, we give evidence that proving zero-knowledge for $\mathsf{NIP}[\Pi]$ (or variations of it) might be out of reach for non-interactive proofs in the BQSM. However, we show that this protocol preserves some properties of the Π such as Witness Indistinguishability and Witness Hiding properties.

3.1.1 Impossibility of "Black-Box" Non-Interactive Zero-Knowledge in the BQSM.

In order to achieve (computational) zero-knowledge, one would need to construct a simulator that can produce an output that is indistinguishable from the output in the real protocol by polynomially bounded distinguishers. For that, the simulator should have minimal access to the verifier's state and be able to run its program. We show here that only looking at the state of the verifier after a partial measurement is not sufficient to prove zero-knowledge. To overcome such an impossibility, we would need a "white-box" simulator that take advantage from the *code* of the verifier.

We notice that we will show the impossibility result for Σ protocols (i.e. the first message is classical), and that the verifier does not have access to quantum auxiliary input. These two cases usually makes proving quantum zero-knowledge much simpler, making our no-go result stronger.

We define an adversarial verifier strategy as a pair of unitaries $V = (V_1, V_2)$ where V_1 maps $|\phi\rangle|x\rangle_\theta$ and an auxiliary register initialized in state $|0\rangle$ (and potentially an auxiliary quantum input) to registers E and Z where $\dim E \leq 2^q$ and the register Z is measured in the computational basis to enforce the memory bound. The unitary V_2 acts on registers EZ and a register T containing the rest of the prover's transmission (in state $|\theta, h^0, h^1, m^0, m^1\rangle$), and produces the verifier's output.

We define a special type of black-box simulator for the NIP scheme, which we call "BQS-BB", as a QPT algorithm Sim that has black-box access to the unitaries V_1, V_1^*, V_2 and V_2^*. In particular, the simulator is allowed to "look" at the state of the verifier between application of these unitaries, and can even purify the verifier's action (i.e. without the measurement on Z). We show that this simulation technique cannot be used to prove zero-knowledge. Intuitively, the reason why simulation is impossible in this setting is that the simulator cannot (efficiently) retrieve which challenge the verifier could have information about. This prevents, for instance, the simulator from applying the rewinding technique.

This impossibility holds regardless of whether or not the verifier receives an auxiliary input.

Lemma 2. *Let* Π *be an arbitrary 1–bit challenge special-sound* Σ*–protocol for a language* $L \in \mathsf{NP} \setminus \mathsf{BQP}$ *(assuming such a language exists) and let* $V = (V_1, V_2)$ *be an adversarial verifier strategy. If post-quantum one-way functions exists, then non-interactive proof* $\mathsf{NIP}[\Pi]$ *is not zero-knowledge with BQS-BB simulation.*

Proof. We assume that we are running $\mathsf{NIP}[\Pi]$ on a single instance of Π, i.e. with $k = 1$. Let $(\mathsf{Enc}, \mathsf{Dec})$ be a symmetric encryption scheme with semantic security against quantum adversaries [ABF+16], which are implied by the existence of one-way functions. We consider a family of malicious verifiers that collude with the distinguisher in order to thwart any simulation attempt. Let $\{(\mathsf{D}^k, V^k = (V_1^k, V_2^k))\}_{k \in \{0,1\}^\lambda}$ be described as follows.

The unitary V_1^k does the following in a purified manner.

1. Initialize register $Z = (\Theta, X, P)$ in state $|0\rangle_Z$.
2. Upon reception of an n-qubit state $|\Psi\rangle$, move it to register P.
3. Apply H to register Θ to obtain a uniform superposition over $\{0, 1\}$.
4. Perform a *purified measurement* on each qubit of register P in basis Θ to get x, i.e. apply the unitary $|\theta\rangle_\Theta |\psi\rangle_P |0\rangle_X \mapsto |\theta\rangle \sum_x |x\rangle\langle x|(H^\theta)^{\otimes n}|\psi\rangle |x\rangle_X$.
5. Encrypt all registers in Z *in place* using the unitary $|m\rangle \mapsto |\mathsf{Enc}_k(m)\rangle$ (which is possible if Enc_k is perfectly correct).
6. The state of register Z is now

$$\frac{1}{\sqrt{2}} \sum_{\theta, x} \langle x|(H^\theta)^{\otimes n}|\psi\rangle \cdot |\mathsf{Enc}_k(\theta, x, x)\rangle_Z \tag{10}$$

We set V_2^k as the identity, i.e. it just outputs everything it receives: the classical memory register Z from V_1^k and the classical message M from P.

The distinguisher D_k receives register Z containing the encryption (under key k) of the verifier's classical memory and a register M containing the prover message in protocol $\mathsf{NIP}[\Pi]$. It decrypts z to get the verifier's measurement basis θ and outcome x, uses it to recover one of the two transcripts contained in M and outputs 1 if the transcript obtained is accepting and 0 otherwise.

We now argue that it is impossible to simulate such a verifier efficiently. First, we notice that in a real execution (where V interacts with P), D_k always outputs 1 assuming perfect correctness of Π. In a simulated execution, we can use the fact that the language is hard (so that Sim cannot produce two accepting transcripts) and that the verifier's memory is encrypted (so that Sim cannot guess the verifier's challenge) to show that the distinguisher outputs 0 with probability close to $\frac{1}{2}$.

By the semantic security of Enc [ABF+16], for any QPT Sim there exists a QPT simulator Sim' that, whenever Sim calls V_1^k and receives register Z which contains the encryption of the memory of an honest verifier, Sim' ignores register Z, but is still able to produce an output indistinguishable from Sim. Then, for any $i \in L$ with corresponding witness w,

$$\begin{aligned}
&|\Pr[\mathsf{D}_k(\langle \mathsf{P}(i,w) \to \mathsf{V}(i)\rangle) = 1] - \Pr[\mathsf{D}_k(\mathsf{Sim}(i)) = 1]| \\
&= 1 - \Pr[\mathsf{D}_k(\mathsf{Sim}(i)) = 1] \\
&\leq 1 - \Pr[\mathsf{D}_k(\mathsf{Sim}'(i)) = 1] + \|\mathsf{Sim}'(i) - \mathsf{Sim}(i)\| \\
&\leq 1 - \frac{1}{2} + \Pr[(i,w') \in R_L \mid w' \leftarrow \mathsf{Sim}'(i)] + \mathsf{negl}(n)
\end{aligned}$$

In the last inequality, we used the special soundness of Π which says that producing two accepting transcript for the same commitment a is as hard as producing a witness for $i \in L$. By the assumed quantum hardness of L, the probability of this happening is negligible, and if the output of Sim' contains only one accepting transcripts, then D_k outputs 0 with probability $\frac{1}{2}$.

□

Lemma 2 indicates that techniques restricted to evaluating V and V^* will not suffice for proving zero-knowledge of $\mathsf{NIP}[\Pi]$. White-box techniques exist for "looking inside" the verifier to infer the index $\bar{c}$ of the OT message on which it has uncertainty. See Section 1.2 for examples. These techniques rely on computing the exact probability distributions induced by the adversary's actions and inferring the random variable $\bar{C}$ whose existence is established by the min-entropy splitting lemma (Lemma 1). Extraction is therefore inefficient, which makes these results inapplicable in the context of zero-knowledge.

Nevertheless, it would be surprising if the verifier could learn anything from the non-interactive proof that it could not learn in the Σ–protocol. The security of BQS-OT ensures that the response to one of the two possible challenges is hidden *information theoretically*. The impossibility of zero-knowledge appears to be more due to a lack of ways in which the simulator can "cheat" than to an actual leakage of information. We can therefore show that other security properties against malicious verifiers – e.g. witness hiding and witness indistinguishability – are preserved by our transformation.

3.1.2 Honest Verifier Zero-Knowledge.

It is trivial to show that a simulator able to read the honest verifier's memory after the honest measurement is able to produce a valid proof. The simulator acts as both the prover and the honest verifier: for each $i \in [k]$, it prepares the states $|x_i\rangle_{\theta_i}$, picks a bit $c_i \in \{0,1\}$ at random and measures the state in basis c_i. After the measurements with outcomes $x'_1, \ldots, x'_k$, the simulator uses the HVZK simulator for Σ on input c_i to produce a valid transcript $(a_i, c_i, r_i^{c_i})$. For the classical prover message, the simulator chooses h_i^0, h_i^1 at random and sets $m_i^{c_i} = r_i^{c_i} \oplus h_i^{c_i}(x_i^{c_i})$ and $m_i^{1-c_i}$ uniformly random. The simulator runs V on the message $(\theta_i, m_i^0, m_i^1, a_i, h_i^0, h_i^1)_{i\in[k]}$ outputs whatever V outputs.

3.1.3 Witness Indistinguishability.

Witness indistinguishability was introduced in [FS90a] as a relaxation of zero-knowledge. We adapt the definition to quantum proof systems.

Definition 6 (Witness Indistinguishability)**.** Let R be an NP relation and let Π be a quantum proof system for R. We say that Π is computationally (resp. statistically) *witness indistinguishable* (BQS-WI) if for any V', for any instance x and witnesses w_1, w_2, and any auxiliary input y, the quantum states

$$\langle P(x,w_1), V'(x,y)\rangle \text{ and } \langle P(x,w_2), V'(x,y)\rangle$$

are computationally (resp. statistically) trace-indistinguishable. We say Π is WI in the BQSM (BQS-WI) if indistinguishability holds for any q–bounded V'.

Theorem 4. *If* Π *is a (computational/statistical) witness indistinguishable proof system, then* $\mathsf{NIP}[\Pi]$ *is (computational/statistical) witness indistinguishable in the BQSM against* q*–bounded verifiers for* $n/4 - q \in \Omega(n)$*.*

Proof. Let w, w' be two witnesses for $x \in L$. Let ρ_{TZE} be the state of the q–bounded verifier after interacting with $\mathsf{P}(x, w)$ where E is the q-qubit quantum memory of V, Z is its classical partial measurement outcome and $T = (\Theta^{(i)}, H_0^{(i)}, H_1^{(i)}, M_0^{(i)}, M_1^{(i)}, A^{(i)})$ is the classical register sent by the prover. Let σ_{TZE} be the state where the prover uses the witness w' instead.

By the security of BQS-OT (Theorem 2), for each i there exists a random variable C_i such that $M_{1-C_i}^{(i)}$ is statistically close to independently and uniformly random. Let $c = c_1 \dots c_k$ and let ρ^c denote the state where $M_{1-c_i}^{(i)}$ is replaced with the completely mixed state for each i:

$$\rho^c = \frac{1}{2^\ell} \mathbb{I}_{M_{\bar{c}}} \otimes \operatorname{tr}_{M_{\bar{c}}}(\rho) \quad . \tag{11}$$

We define σ^c in the same way. By Theorem 2, $\rho \approx^\epsilon \sum_{c \in \{0,1\}^k} p_c \rho^c$ where $p_c = \Pr[C = c]$ for a negligible ϵ as long as $n/4 - q$ is linear in n. By the witness indistinguishability of Π, we have that

$$\begin{aligned} \|\mathsf{D}(\rho) - \mathsf{D}(\sigma)\| &\leq \|\mathsf{D}(\sum_c p_c \rho^c) - \mathsf{D}(\sum_c p_c \sigma^c)\| + 2\epsilon \\ &\leq \sum_c p_c \|\mathsf{D}(\rho^c) - \mathsf{D}(\sigma^c)\| + 2\epsilon \\ &\leq \nu + 2\epsilon \end{aligned}$$

since the distinguishing advantage between ρ^c and σ^c is at most the advantage to distinguish between a transcript for Π with challenge c and witness w and one with witness w'.

□

3.2 Non-interactive statistical WI proofs for NP

We describe now the application of our protocol for (statistical) WI non-interactive proofs for NP. Before discussing such a protocol, we first describe a new non-interactive weak bit-commitment, which may have independent interest.

3.2.1 A new non-interactive weak BC.

The previous protocols for bit commitment in the BQSM had the weird property that the sender commits by measuring a quantum state created by the receiver. For example, in [DFSS08], in order to commit to a message $m \in \{0, 1\}$, the sender would get a message $|x\rangle_\theta$, measure it in basis m and take note of the outcome x'. To open its commitment, it would send m and x' to the receiver who could check that $x'_i = x_i$ whenever $\theta_i = m$. This quirk of DFSS-BC is actually the reason why our round compression transform can go down to two messages as we will see in Section 4. But in the context of our non-interactive proof using DFSS-BC applied to commit-and-open protocols, we cannot replace the classical commitments with DFSS-BC since it would introduce communication form the verifier to the sender.

Intuitively, it does not matter who prepares the state and who measures it since by a purification argument, the state preparation of DFSS-BC can be seen as measuring halves of EPR pairs. Formally proving that this is still secure is more difficult, and the tools to do so were only discovered a couple of years later in [DFLS16], which can show that it is still sum-binding. For our purpose, we actually need a weaker security notion than sum-binding. We first present the "reversed" protocol and then describe and prove the security notion it needs to satisfy.

Protocol weak-BC

Commit Phase

- **Committer(b):** Choose $x \in_R \{0,1\}^n$. Send $|x\rangle_b$.
- **Receiver:** Measure qubits upon reception in a random basis $\theta \in_R \{0,1\}^n$, gets outcome x'.

Open Phase

- **Committer(b):** Send x and b. Receiver checks that $x_i = x'_i$ whenever $\theta_i = b$.
- **Receiver:** Check that $x_i = x'_i$ whenever $\theta_i = b$.

The usual sum-binding criteria asks that, for a fixed commitment ρ_{AB}, if the sender succeeds in opening b with probability p_b, then $p_0 + p_1 \leq 1 + \mathsf{negl}(n)$. In this context, the malicious sender can measure its part of the state A adaptively based on the knowledge of the target bit b. We consider a weaker task where the sender must provide *both* openings simultaneously, and does not know which will be tested. This is strictly weaker than sum-binding since, as the following theorem shows, this is achieved unconditionally by the above protocol.

Theorem 5. *The above weak-BC protocol is perfectly hiding and is binding according to the following. Let ρ_{AB} be an arbitrary density operator describing the joint state of the committer and the receiver after the commit phase. Let $\{V_{acc}^{x,b}, V_{rej}^{x,b}\}$ be the verifier's measurement for opening (x, b). Then*

$$\sup_M \sum_{b\in\{0,1\}} \sum_{x_0,x_1} \operatorname{tr}\left((M_{x_0,x_1} \otimes V_{acc}^{x_b,b})\rho_{AB}\right) \leq 1 + 2^{-\frac{n}{2}+2h(\delta)n} + 2^{-\delta n+1} \tag{12}$$

where $h(\cdot)$ is the binary entropy and $\delta > 0$ is an arbitrary constant.

Proof. Hiding follows from the fact that

$$\sum_{x\in\{0,1\}^n} |x\rangle\langle x| = \sum_{x\in\{0,1\}^n} H^{\otimes n}|x\rangle\langle x|H^{\otimes} = \frac{\mathbb{I}}{2^n}$$

We will bound the weak binding criteria through a series of hybrids which each negligibly change the success probability. Let $p_b = \sup_M \sum_{x_0,x_1} \operatorname{tr}((M_{x_0,x_1} \otimes V_{acc}^{x_b,b})\rho_{AB})$ be the probability of acceptance when the opening to bit b is checked, where of course M cannot depend on b.

Hybrid 1. The receiver holds on to the qubits in the commit phase and waits for the committer to send its opening before measuring in a random basis Θ. The trace in (12) is unchanged by this modification.

Hybrid 2. As Hybrid 1, but instead of choosing Θ at random and measuring in basis Θ, the receiver measures all the qubits in the basis b sent by the committer. Then, the receiver chooses a subset $T \subseteq [n]$ uniformly at random and rejects if for any $i \in T$, the result x'_i is different from x_i. The probability distributions are also unchanged as this is equivalent to the checking procedure with $\Theta_i = b$ if $i \in T$ and $\Theta_i = 1 - b$ if $i \notin T$. The marginal distribution of Θ is still uniform.

Hybrid 3. As Hybrid 2, but instead of comparing the positions for a random subset T, the receiver rejects if the measurement outcome x' is at Hamming distance greater than δn from x. The receiver will reject more often in this hybrid. The probability that the verifier rejects in Hybrid 3 and not in Hybrid 2 is the probability that $\Delta(x', x) > \delta n$, yet $x'_i = x_i$ for all $i \in T$. Since T is chosen uniformly at random, this probability is at most $2^{-\delta n}$. Let p'_b be the probability that the receiver accepts an opening to b in Hybrid 3, then $p_b \leq p'_b + 2^{-\delta n}$.

We now bound the sum of probabilities for Hybrid 3. Let $\sum_{x' \approx x} |x'\rangle\langle x'|_b$ be the projector onto accepting outcomes for the opening of $b \in \{0, 1\}$ in Hybrid 3. We have that

$$\begin{aligned}
p'_0 + p'_1 &= \sup_M \sum_{b \in \{0,1\}} \sum_{x_0, x_1} \operatorname{tr}\left((M_{x_0,x_1} \otimes \sum_{x \approx x_b} |x'\rangle\langle x'|_b) \rho_{AB} \right) \\
&\leq \sup_{x_0, x_1} \operatorname{tr}\left(\sum_{x \approx x_0} |x\rangle\langle x|_0 \cdot \rho \right) + \operatorname{tr}\left(\sum_{y \approx x_1} |y\rangle\langle y|_1 \cdot \rho \right) \\
&\leq \left\| \sum_{x \approx x_0} |x\rangle\langle x|_0 + \sum_{y \approx x_1} |y\rangle\langle y|_1 \right\|_\infty \\
&\leq 1 + \left\| \sum_{x \approx x_0} |x\rangle\langle x|_0 \cdot \sum_{y \approx x_1} |y\rangle\langle y|_1 \right\|_\infty \\
&\leq 1 + 2^{2h(\delta)n - n/2}
\end{aligned}$$

where we use the inequality $\|A + B\| \leq 1 + \|A \cdot B\|$ for projectors A and B (a fact whose proof can be found in [BFGS13]), the fact that there are at most $2^{h(\delta)n}$ strings at distance δn from x_b and that $\langle x|_0 |y\rangle_1 = 2^{-\frac{n}{2}}$ for any x, y.

Compiling the error introduced with Hybrid 3, we have that

$$(12) = p_0 + p_1 \leq 1 + 2^{2h(\delta)n - n/2} + 2 \cdot 2^{-\delta n}$$

□

3.2.2 A non-interactive statistical WI proof for NP in the BQSM.

We now consider the following Ξ protocol for the NP-complete L_{ham} corresponding to graphs that have a Hamiltonian cycle. It consists of the original Σ protocol for this problem, but using weak BC as the commitment.

Protocol Ξ protocol Π_{ham} for Hamiltonian cycle

1. **Prover:** Using weak BC, commits to the adjacency matrix of a random permutation σ of the graph G
2. **Verifier:** Send a random bit b
3. **Prover:** If $b = 0$ open the whole adjacency matrix and provide the permutation σ. If $b = 1$, open the edges corresponding to the Hamiltonian cycle.
4. **Verifier:** Check the consistency of the Prover's opening.

The completeness of the protocol follows directly from the completeness of the original protocol, and zero-knowledge follows from [Wat01].

However, since we use weak BC, this protocol does not satisfy the standard soundness definition (in particular, the Prover can answer the two challenges by keeping the purification of the commitment and measuring it accordingly).

However, we prove now that it satisfies the *oblivious soundness* property that we mentioned in Remark 1.

Lemma 3. Π_{ham} *satisfies oblivious soundness.*

Proof. Let $G \notin L_{ham}$. Let G' be the graph corresponding to the answer r_0. If G' has a Hamiltonian cycle, it cannot be a permutation of G, therefore the first check will fail with probability 1. Moreover, if r_1 does not open to a Hamiltonian cycle, the second check will fail with probability 1. In this case, for the two checks to pass, there is at least one entry i, j of the adjacency matrix whose opening $o_{i,j}$ is b in r_0 and whose opening $o'_{i,j}$ is $\neg b$ in r_1.

Therefore, in order to make the verifier accept, the prover has to provide values (r_0, r_1) such that Equation (9) holds, which is upper-bounded by the probability that the prover can provide simultaneously two different openings to the commitment, which is impossible by Theorem 5. □

By observing that Π_{ham} is perfectly witness indistinguishable because weak-BC is perfectly hiding, and combining Lemma 3 and Remark 1, we obtain the following result.

Corollary 2. *There is a non-interactive quantum proof system for* L_{ham} *which is unconditionally sound and witness indistinguishable against BQS verifiers.*

4 A General Round-Compression Transform in the BQSM

In this section, we present and prove the soundness of the general transform mapping k–round interactive proofs for $k = \mathsf{poly}(\lambda)$ to 2–message quantum proofs.

We assume for simplicity that all the prover messages are of length $\ell = \ell(\lambda)$ and all the verifier challenges are of length $m = m(\lambda)$ for some polynomials $\ell, m : \mathbb{N} \to \mathbb{N}$, and that the prover sends the first and last messages. We let $a_1, \ldots, a_{k+1}$ denote the $k+1$ prover messages and $c_1, \ldots, c_k$ the k verifier challenges, where a_{i+1} responds to challenge c_i. Let P_i^{Π} denote the next-message function of the prover in protocol Π that takes as input the partial transcript so far and outputs a_i. The RR transform is presented below.

Protocol $\mathsf{RR}[\Pi]$

Parameter: A k–round interactive proof system $\Pi = (\mathsf{P}^{\Pi}, \mathsf{V}^{\Pi})$ for a language L.

Verifier message:

1. For $i \in [k]$:
   1.1 V runs the commit phase of the DFSS-BC string commitment to get a quantum register P_i.
   1.2 V picks $c_i \in_R \{0,1\}^m$ to initialize a register C_i in state $|c_i\rangle$
2. V sends the registers $P_1 C_1 \ldots P_k C_k$ to P.

Prover message:

3. On input $x \in L$, P first computes $a_1 = \mathsf{P}_1^{\Pi}(x)$.
4. For $i \in [k]$,
   4.1 On reception of register P_i, P commits to a_i as in the commit phase of DFSS-BC.
   4.2 P measures register C_i in the computational basis to obtain c_i. P computes $a_{i+1} = \mathsf{P}_{i+1}^{\Pi}(a_1, \ldots, a_i, c_1, \ldots, c_i, x)$.
5. P runs the reveal phase of DFSS-BC, sending every a_i and opening string to V.

Verification:

6. V performs the verification for every instance of DFSS-BC. It accepts if every opening is valid and if $a_1, \ldots, a_{k+1}, c_1, \ldots, c_k$ is an accepting transcript for Π on input x. Otherwise, it rejects.

Theorem 6. *Let DFSS-BC be the δ–binding BQS-BC from Section 2.1.1. If Π is a k–round public-coin interactive proof with soundness error ϵ against unbounded (resp. QPT) provers, then $\mathsf{RR}[\Pi]$ is a 1–round quantum interactive proof (resp. argument) with soundness error*

$$\epsilon + k^2 \cdot \delta \tag{13}$$

against q–bounded adversaries where δ is negligible if $n/4 - q \in \Omega(\lambda)$ where $n = n(\lambda)$ is the number of qubits sent in DFSS-BC and $q = q(\lambda)$ is the quantum memory bound on the prover.

Proof. We use a hybrid argument to prove the soundness of $\mathsf{RR}[\Pi]$. Consider the following hybrid protocols where in Hyb i the round-compression transform is applied up to the ith prover message, and the rest of protocol is interactive.

- Hyb 0: same as protocol Π
- Hyb i: apply transformation RR to the messages of Π up to round i.
  1. V prepares i registers $P_1 \ldots P_i$ and i random values $c_1, \ldots, c_i$ in registers $C_1 \ldots C_i$ and sends $\bigotimes_{j=1}^{i} P_j C_j$.
  2. On reception of a message $(a_1 \ldots a_{i+1}, z_1 \ldots z_i)$ from the prover, V checks that (a_j, z_j) is valid opening for $j \in [i]$ and rejects if any are invalid.
  3. V and P continue as in protocol Π: V sends c_j and P responds with a_{j+1} for $j = i+1, \ldots, k$. V checks that $(a_1 \ldots a_{k+1}, c_1 \ldots c_k)$ is an accepting transcript for Π.
- Hyb k: same as in $\mathsf{RR}[\Pi]$

The difference between two hybrids $i-1$ and i is that in hybrid $i-1$, $A_1 \ldots A_i$ are sent to V before it sends C_i whereas in hybrid i, the adversary receives C_i before opening its commitments to $A_1 \ldots A_i$. We will show that this only confers a negligible advantage to an adversary.

Consider a q–bounded adversary $\mathcal{A}^i$ against Hyb i. By the definition of binding for BQS-BC (Definition 1), after the commit phase of the jth commitment (i.e. after the transmission of register P_j for $j \leq i$), there is a random variable A'_j such that conditioned on $A'_j = a'_j$, $\mathcal{A}^i$ has negligible probability of opening the jth commitment to $a_j \neq a'_j$. This random variable is defined by the partial measurement $\mathcal{A}^i$ is forced to make on register P_j before V begins transmission of register C_j, so it is independent of C_j.

This independence means that learning C_i before sending $A_1 \dots A_i$ does not give a noticeable advantage to the adversary. We make this formal by constructing, from the adversary $\mathcal{A}^i$ that has success probability ϵ_i against Hyb i, an adversary $\mathcal{A}^{i-1}$ against hybrid $i-1$ that has success probability at least $\epsilon_i - \mathsf{negl}(\lambda)$. $\mathcal{A}^{i-1}$ performs the same strategy as $\mathcal{A}^i$ on reception of the registers $P_1C_1 \dots P_{i-1}C_{i-1}$. For producing the next value a_i, $\mathcal{A}^{i-1}$ simulates the verifier in the ith commitment, i.e. creates the register P_i just as V would in hybrid i, again applying $\mathcal{A}^i$'s strategy, and checking that the resulting opening is valid.

Adversary $\mathcal{A}^{i-1}$

1. While receiving registers $\bigotimes_{j=1}^{i-1} P_jC_j$ from the verifier, forward them to $\mathcal{A}^i$. For the last registers P_iC_i that $\mathcal{A}^i$ expects, $\mathcal{A}^{i-1}$ simulates the verifier, i.e. constructs register P_i from the commit phase and sends P_i followed by a random challenge c to $\mathcal{A}^i$.
2. $\mathcal{A}^{i-1}$ now receives a message $(a_1 \dots a_{i+1}, z_1 \dots z_i)$ from $\mathcal{A}^i$. It checks (a_i, z_i) is a valid opening of the commitment and aborts if the check fails. It discards a_{i+1} and sends $(a_1 \dots a_i, z_1 \dots z_{i-1})$ to the verifier.
3. After receiving the challenge c_j for $i \leq j \leq k$ from the verifier, it computes and sends a_{j+1} using the same strategy as $\mathcal{A}^i$.

Observe the following facts about $\mathcal{A}^{i-1}$:

- The quantum memory required to perform attack $\mathcal{A}^{i-1}$ against Hyb $i-1$ is the same as attack $\mathcal{A}^i$ against Hyb i.
- $\mathcal{A}^i$ cannot distinguish whether it is interacting with V in Hyb i or with $\mathcal{A}^{i-1}$ in Hyb $i-1$.
- The random variables $A'_1, \dots, A'_i$ have the same distribution in both experiments ($\mathcal{A}^i$ against Hyb i and $\mathcal{A}^{i-1}$ against Hyb $i-1$).

Let us fix some arbitrary values $a'_1 \dots a'_i$ for $A'_1 \dots A'_i$. Assume for now that $a_1 \dots a_i = a'_1 \dots a'_i$. Since these values are independent of C_i, they would remain unchanged for any value c_i that $\mathcal{A}^{i-1}$ had given to $\mathcal{A}^i$. And since $\mathcal{A}^{i-1}$ answers the rest of the challenges exactly as $\mathcal{A}^i$ would, the whole transcript is identically distributed in both experiments, thus the probability of the verifier accepting is the same.

Now for the other case (there is some $j \leq i$ such that $a_j \neq a'_j$), the verifier will reject the opening to the ith commitment with overwhelming probability. By the definition of δ–binding for BQS-BC schemes (Definition 1), the probability that $\mathcal{A}^{i+1}$ can announce a basis $A_j \neq A'_j$ is upper-bounded by δ. By a union bound, the probability that there is a $1 \leq j \leq i$ such that $A_j \neq A'_j$ is at most $i \cdot \delta$. Therefore if we let ϵ_j denote the soundness error of Hyb j for $j = 0 \dots k$, then

$$\epsilon_i \leq \epsilon_{i-1} + i \cdot \delta \ . \tag{14}$$

Since by assumption, Π is ϵ–sound, then $\mathsf{RR}[\Pi]$ is ϵ'–sound for

$$\epsilon' \leq \epsilon + k^2 \cdot \delta \tag{15}$$

where the k^2 comes from the fact that when going from hybrid i to hybrid $i+1$, we introduce the negligible term $i \leq k$ times, and there are k hybrids. If δ is negligible, then the above is arbitrarily close to ϵ. By Theorem 1, this happens if the memory bound on the prover satisfies $n/4 - q \in \Omega(\lambda)$

□

4.1 Application: Two-Message Zero-Knowledge in the BQSM

The goal of this section is to construct a two-message zero-knowledge proof for any NP language in the BQSM. We begin by proving that our transform produces a zero-knowledge 2–message quantum proof when applied to proof systems that satisfy the following notion of honest-verifier zero-knowledge, which is a generalization of special HVZK to multi-round protocols.

Definition 7. We say that a Π protocol is special qHVCZK (special HVSZK) if for any given challenge $(c_1, ..., c_k)$, there is an efficient simulator $\mathcal{S}(c_1, ..., c_k)$ such that for every QPT (unbounded) distinguisher $\mathcal{D}$,

$$|\Pr[\mathcal{D}^{\mathcal{S}(x,\cdot)}(1^\lambda) = 1] - \Pr[\mathcal{D}^{\mathsf{P}\leftrightharpoons\mathsf{V}(x,\cdot)}(1^\lambda) = 1]| \leq \mathsf{negl}(\lambda),$$

where $\mathcal{D}$ can query its oracle with (classical) values $c_1, ..., c_k$. In the first term, it receives $\mathcal{S}(x, c_1, ..., c_k)$ and in the second term, it receives the transcript $\mathsf{P} \leftrightharpoons \mathsf{V}(x, c_1, \ldots, c_k)$ that come from the real protocol when the challenges are fixed.

We now show that if RR is applied to a special-HVZK k–round protocol for $k = \mathsf{poly}(\lambda)$, then the resulting scheme is zero-knowledge against quantum verifiers. To prove zero-knowledge, instead of producing a simulator for the malicious verifier, we show that there exists a simulator which does not interact with the prover and that can simulate the actions of P. One can easily see that this implies (auxiliary-input) zero-knowledge by running any malicious verifier $\widetilde{\mathsf{V}}$ with this simulated prover.

At first glance, the existence of this simulator appears to be at odds with the soundness of our transform. For example, if the prover relies on the knowledge of a witness w that $x \in L$ for $L \in \mathsf{NP}$, then the simulator can convince the verifier that $x \in L$ without access to w. This matter is resolved by observing that the quantum memory of the simulator is not bounded, unlike the (malicious) prover. This fact is crucial, as we show in Section C that there are no 2–message quantum proof systems for hard languages that are both sound and zero-knowledge when the quantum memory of the prover is not bounded.

Furthermore, the existence of a fully quantum simulator for a BQS adversary appears vacuous, since the simulator might be more powerful than the adversary it simulates. However, we need to emphasize that the party we are simulating – the verifier – is *not* assumed to be quantum memory bounded, only the prover is. Thus, we show that zero-knowledge holds against fully quantum verifiers using a fully quantum simulator.

Theorem 7. *If* $\Pi = (\mathsf{P}^\Pi, \mathsf{V}^\Pi)$ *is a special qHVZK* Σ*–protocol for a language* L*, then* $\mathsf{RR}[\Pi]$ *is qZK. The type of zero-knowledge (computational or statistical) is preserved by* RR.

Proof. We construct a simulator for the prover instead of the verifier; i.e. this simulator mimics the actions of the prover from the verifier's point of view and does not have access to the real prover. Turning this simulator into one for the verifier is then just a question of making the verifier interact with this simulated prover.

First observe that from the verifier's point of view, the action of the quantum memory-less honest prover P is perfectly indistinguishable from the action of a "semi-honest" prover P^* that *does* have a quantum memory and that delays its commitment to a_i using P_i until after every challenge c_i was measured.

Now since the prover messages a_i are committed to after every challenge is learned, we can employ the simulator Sim_Π for the Σ-protocol to obtain a simulated transcript $(a_1, \ldots, a_{k+1})$ indistinguishable from a real transcript. In more details, we construct the simulator Sim for $\mathsf{RR}[\Pi]$ as follows:

1. Receive the registers $P_1, C_1, ..., P_k, C_k$ from $\tilde{\mathsf{V}}$, delaying any measurement
2. Measure registers $C_1, ..., C_k$ in the computational basis and get outcomes $c_1, ..., c_k$
3. Compute $\mathsf{Sim}_\Pi(c_1, ..., c_k) = (a_1, ..., a_{k+1})$, where Sim_Π is the special qHVZK simulator
4. Perform the commitment phase of BQS-BC on register P_i by committing to a_i and get the opening string z_i
5. Return $(a_1, ..., a_{k+1}, z_1, ..., z_k)$ to $\tilde{\mathsf{V}}$

We now show that this simulator indistinguishable from P. For that, let us assume towards a contradiction that there exists a distinguisher $\mathcal{D}$ and a state ρ_{QE}, where $Q = P_1C_1 \ldots C_kP_k$ is sent to the prover/simulator and E is kept by the distinguisher, such that

$$\|\mathcal{D}(\mathsf{P} \otimes \mathbb{I}_E(\rho)) - \mathcal{D}(\mathsf{Sim} \otimes \mathbb{I}_E(\rho))\| \geq \lambda^{-d} \tag{16}$$

for $d \in O(1)$. Then, we can construct a distinguisher $\mathcal{D}_\Pi^{\mathcal{C}}$ that can break the special qHVZK property of Π with probability at least λ^{-d}, where $\mathcal{C}$ is an oracle for either $\mathsf{P}_\Pi \leftrightarrows \mathsf{V}_\Pi(x, \cdot)$ or $\mathsf{Sim}_\Pi(x, \cdot)$. It works as follows:

1. Compute the state ρ_{QE} which allows to distinguish Sim and P
2. Measure registers $C_1, \ldots, C_k$ of ρ_Q and get outcome $(c_1, ..., c_k)$
3. Query $\mathcal{C}(c_1, ..., c_k)$ and get the output $(a_1, ..., a_{k+1})$
4. Commit to a_i using register P_i and get opening string z_i
5. Output $\mathcal{D}(a_1, ..., a_{k+1}, z_1, ..., z_k)$

Notice that when $\mathcal{C} = \mathsf{Sim}_\Pi(x, \cdot)$, then the output of $\mathcal{D}_\Pi^{\mathcal{C}}$ is $\mathcal{D}(\mathsf{Sim} \otimes \mathbb{I}_E(\rho))$. Moreover, when $\mathcal{C} = \mathsf{P}_\Pi \leftrightarrows \mathsf{V}_\Pi(x, \cdot)$, we have that $\mathcal{D}_\Pi^{\mathcal{C}}$ is $\mathcal{D}(\mathsf{P}^* \otimes \mathbb{I}_E(\rho))$ where P^* is the semi-honest prover introduced earlier. In this case, we have that

$$\|\mathcal{D}_\Pi^{\mathsf{P}_\Pi \leftrightarrows \mathsf{V}_\Pi(x,\cdot)}(1^\lambda) - \mathcal{D}_\Pi^{\mathsf{Sim}_\Pi(x,\cdot)}(1^\lambda)\| = \|\mathcal{D}(\mathsf{P}^* \otimes \mathbb{I}_E(\rho)) - \mathcal{D}(\mathsf{Sim} \otimes \mathbb{I}_E(\rho))\| \geq \lambda^{-d},$$

which contradicts the qHVZK of Π by recalling that the actions of P^* and P are perfectly indistinguishable. Therefore we conclude that the CPTP maps P and Sim are (computationally or statistically) indistinguishable if Π is (computationally or statistically) qHVZK. □

4.1.1 Quantum statistical zero-knowledge proofs.

In this section, we show that using the statistically binding and hiding BQS-BC scheme of Section 2.1.1, we can achieve 2–message quantum statistical zero-knowledge proofs in the BQSM.

In the previous subsection, we showed that special qHVZK Σ protocols can be converted into 2-messages QZK protocols in the BQSM. However, (honest verifier) ZK proofs for NP-complete languages rely on computational assumptions, usually to implement commitment schemes. Since we are in BQSM, we can instead use quantum commitment schemes with perfect hiding and statistical binding and achieve statistical ZK proofs in the BQSM.

For simplicity, we will prove the result for a single-shot run of 3-coloring, but the result follows analogously with the parallel repetition of the protocol.

2–message perfect zero-knowledge proof

Input: Graph $G = (V, E)$ with $|V| = n$.

Verifier message:

1. For $i = 1, ..., n$, V runs the commit phase of the DFSS-BC string commitment to get a quantum register P_i.
2. V picks $c \in_R E$ to initialize a register C in state $|c\rangle$
3. V sends the registers $P_1 \dots P_n C$ to P.

Prover message:

3. P first computes a random 3-coloring of the graph G. Let w_1,...,w_n be the color of each vertex of the graph. P commits to each of the colors independently: for $i \in [n]$,
4. On reception of register P_i, P commits to w_i as in the commit phase of DFSS-BC.
5. P measures register C in the computational basis to obtain $\{i, j\} \in E$.
6. P runs the reveal phase of DFSS-BC for w_i and w_j.

Verification:

6. V runs the verification of DFSS-BC for w_j and w_i and checks that $w_j \neq w_i$.
7. If verification or the check failed, it aborts. Otherwise, it accepts.

Theorem 8. *The protocol described above is a two-message perfect zero-knowledge proof for* 3*-coloring which is statistically sound against q–bounded provers with $n/4 - q \in \Omega(n)$.*

Proof. Completeness follows straightforwardly if the P follows the honest strategy.

To prove soundness, we use the ϵ-binding property of the commitment scheme. For that, let $w'_1, ..., w'_n$ be values of the the random variables $b'_1, ..., b'_n$ that come from Definition 1 corresponding to the commitment of the color of each node. We notice that since the graph is not 3-colorable, there exists at least one edge $\{i, j\} \in E$ such that $w'_i = w'_j$. We also have that the V's challenge is $\{i, j\}$ with probability $\frac{1}{m}$, and let us consider this case.

If P opens the commitments to the values w'_i and w'_j, V rejects with probability 1. If P opens the commitments to values $\tilde{w}_i \neq w'_i$ or $\tilde{w}_j \neq w'_j$, V rejects except with probability ϵ.

In this case, if the graph is not 3-colorable, V rejects with probability at least $\frac{1-\epsilon}{m}$.

The simulator and the zero-knowledge proofs follow closely the proof of Theorem 7. The fact that the flavour of zero-knowledge is perfect comes from the fact that the commitment scheme has perfect hiding since no information of non-open values is sent to V.

□

4.2 Applications: Two-Message Interactive Proof for PSPACE

In this section, we describe applications of our round compression transform RR presented in 4 when applied to a specific interactive proof system.

4.2.1 Sum-Check Protocol.

The sum-check protocol is the key ingredient of several fundamental results in complexity theory and cryptography. In this protocol, the prover aims to prove that

$$\sum_{x_1,\dots,x_n\in\{0,1\}} f(x_1,\dots,x_n)=B,$$

for some given value B and function f an n-variate polynomial of degree at most d. The idea of the sum-check protocol is to consider a field $\mathbb{H}$, where $\mathbb{F}_2\subseteq\mathbb{H}$ and $|\mathbb{H}|\gg d$,

Sum-check Protocol

Prover 1st message: P computes $g_1(x_1)=\sum_{x_2,\dots,x_n\in\{0,1\}} f(x_1,\dots,x_n)$ and sends g_1 to V, who checks that g_1 is an univariate polynomial of degree at most d and that $g_1(0)+g_1(1)=B$. If any of the checks failed, reject.

Verifier 1st message: V sends a uniformly random $r_1\in\mathbb{H}$ to P.

Prover ith message: P computes

$$g_i(x_i)=\sum_{x_{i+1},\dots,x_n\in\{0,1\}} f(r_1,\dots,r_{i-1},x_i,x_{i+1},\dots,x_n)$$

and sends g_i to V, who checks that g_i is an univariate polynomial of degree at most d and that $g_i(0)+g_i(1)=g_{i-1}(r_{i-1})$. If any of the checks failed, reject.

Verifier ith message: V sends a uniformly random $r_i\in\mathbb{H}$ to P.

Prover last message: P computes $g_n(x_n)=f(r_1,\dots,r_{n-1},x_n)$ and sends g_n to V, who checks that g_n is an univariate polynomial of degree at most d and that $g_n(0)+g_n(1)=g_{n-1}(r_{n-1})$. Moreover, V also checks that $g_n(r_1,\dots,r_n)=f(r_1,\dots,r_n)$, for a random $r_n\in\mathbb{H}$. If either of these tests do no pass, reject.

If all tests passed, V accepts.

The main result regarding the sum-check protocol is the following [LFKN92; Sha92].

Theorem 9. *The sum-check protocol presented above has the following properties:*

- **Completeness:** *If $\sum_{x_1,\dots,x_n\in\{0,1\}} f(x_1,\dots,x_n)=B$, there is a strategy for P such that V accepts with probability 1.*
- **Soundness:** *If $\sum_{x_1,\dots,x_n\in\{0,1\}} f(x_1,\dots,x_n)\neq B$, for any strategy for P, V accepts with probability at most $\frac{nd}{|\mathbb{H}|}$.*
- **Complexity:** *The honest prover runs in time $\mathsf{poly}(|\mathbb{H}|^n)$, the verifier runs in time $\mathsf{poly}(|\mathbb{H}|,n)$ and space $O(n\log|\mathbb{H}|)$. The communication complexity is $\mathsf{poly}(|\mathbb{F}|,n)$ and the number of bits sent by the verifier is $O(m\log|\mathbb{H}|)$. Moreover, the protocol is public-coin.*

We notice that the sum-check protocol is a multi-round interactive proof where V only sends random coins (interpreted as field elements) as messages. In this case, we can apply the RR transformation to it to achieve a one-round quantum protocol with similar guarantees of the classical sum-check protocol in the quantum bounded storage model.

Corollary 3 (2–Message Quantum Sum-Check Protocol)**.** *For any $q \in \mathbb{N}$, there is a 2–message quantum proof for the sum-check problem in the bounded storage model with negligible soundness error against q–bounded provers where the communication grows as a polynomial in q.*

Proof. Apply the RR compiler to the sum-check protocol. There are $n + 1$ messages exchanged each of whom is $\mathsf{poly}(n)$ in length. Committing to an ℓ–bit string using DFSS-BC requires $\ell \times \lambda$ qubits for security against q–bounded committers with $\lambda/4 - q \in \Omega(\lambda)$. □

The sum-check protocol is a crucial tool in results in complexity theory and cryptography, especially regarding delegation of (classical) computation. We can easily replace the classical sum-check protocol by its quantum version to to achieve round-efficient protocols, that we describe below.

Corollary 4. *Every language in* PSPACE *has a 2–message quantum protocol in the bounded storage model.*

Notice that if we do not consider provers with bounded memory, we have that $\mathsf{PSPACE} = \mathsf{QIP}(3)$, and if we define $\mathsf{QIP}(2)^{\mathsf{BQSM}}$ as the class of problems with two–message quantum interactive proof systems where the prover has bounded quantum memory (but unbounded computational power), we have that $\mathsf{PSPACE} = \mathsf{QIP}(2)^{\mathsf{BQSM}}$.

More recently, the sum-check protocol has been also used to achieve protocols for doubly-efficient delegation of computation. In this setting, the goal is to achieve a protocol where V interacts with a P in order to delegate the computation of an arithmetic circuit with the following properties:

- The honest prover's computation should not be much more costly than running the original circuit.
- The running time of the verifier should be linear in the input size of the circuit.

Such a protocol was originally proposed by Goldwasser, Kalai and Rothblum [GKR15] and later improved by Reingold, Rothblum and Rothblum [RRR21]

Lemma 4 (Corollary 1.4 of [GKR15])**.** *Let L be a language in* P*, that is, one that can be computed by a deterministic Turing machine in time poly(n). There is an interactive proof for L where:*

- *the honest prover runs in time* $\mathsf{poly}(n)$ *and the verifier in time* $\mathsf{poly}(n)$ *and space $O(log(n))$;*
- *the protocol has perfect completeness and soundness* $1/2$*; and*
- *the protocol is public coin, with communication complexity* $\mathsf{poly}(n)$*.*

Again here, the interaction between the verifier and the prover consists of multiple instances of the classical sum-check protocol. Therefore, using Corollary 3, we achieve the following.

Corollary 5. *Let L be a language in* P*. L has a quantum interactive proof in the bounded storage model where:*

- *the honest prover runs in time* $\mathsf{poly}(n)$ *and the verifier in time* $\mathsf{poly}(n)$ *and space $O(log(n))$;*
- *the protocol has perfect completeness and soundness* $1/2$*; and*
- *there is one round of communication.*

Acknowledgments

We are grateful to Serge Fehr for enlightening discussions on the all-but-one uncertainty relation (Theorem 10) and to Louis Salvail for discussions on simulating bounded storage adversaries. ABG is supported by ANR JCJC TCS-NISQ ANR-22-CE47-0004, by the PEPR integrated project EPiQ ANR-22-PETQ-0007 part of Plan France 2030, and by the European Union's Horizon Europe Framework Programme under the Marie Sklodowska Curie Grant No. 101072637, Project Quantum-Safe Internet (QSI). This work was done in part while ABG was visiting the Simons Institute for the Theory of Computing.

References

[ABF$^+$16] Gorjan Alagic, Anne Broadbent, Bill Fefferman, Tommaso Gagliardoni, Christian Schaffner, and Michael St. Jules. Computational Security of Quantum Encryption. In Anderson C.A. Nascimento and Paulo Barreto, editors, *Information Theoretic Security*, pages 47–71, Cham. Springer International Publishing, 2016. ISBN: 978-3-319-49175-2. DOI: 10.1007/978-3-319-49175-2_3.

[AG22] Prabhanjan Ananth and Alex B. Grilo. Post-Quantum Zero-Knowledge with Space-Bounded Simulation. 2022. DOI: 10.48550/arXiv.2210.06093. arXiv: 2210.06093 [quant-ph]. URL: http://arxiv.org/abs/2210.06093 (visited on 10/13/2022). preprint.

[BDG$^+$13] Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why "fiat-shamir for proofs" lacks a proof. In Amit Sahai, editor, *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings*, volume 7785 of *Lecture Notes in Computer Science*, pages 182–201. Springer, 2013. DOI: 10.1007/978-3-642-36594-2_11. URL: https://doi.org/10.1007/978-3-642-36594-2%5C_11.

[BFGS13] Niek J. Bouman, Serge Fehr, Carlos González-Guillén, and Christian Schaffner. An all-but-one entropic uncertainty relation, and application to password-based identification. In Kazuo Iwama, Yasuhito Kawano, and Mio Murao, editors, *Theory of Quantum Computation, Communication, and Cryptography*, Lecture Notes in Computer Science, pages 29–44, Berlin, Heidelberg. Springer, 2013. ISBN: 978-3-642-35656-8. DOI: 10.1007/978-3-642-35656-8_3.

[BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive Zero-knowledge and Its Applications. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing* (Chicago, Illinois, USA), STOC '88, pages 103–112, New York, NY, USA. ACM, 1988. ISBN: 978-0-89791-264-8. DOI: 10.1145/62212.62222. URL: http://doi.acm.org/10.1145/62212.62222.

[BG22] Anne Broadbent and Alex Bredariol Grilo. Qma-hardness of consistency of local density matrices with applications to quantum zero-knowledge. *SIAM Journal on Computing*, 51(4):1400–1450, 2022. DOI: 10.1137/21M140729X.

[BKS23] James Bartusek, Dakshita Khurana, and Akshayaram Srinivasan. Secure Computation with Shared EPR Pairs. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 224–257, Cham. Springer Nature Switzerland, 2023. ISBN: 978-3-031-38554-4. DOI: 10.1007/978-3-031-38554-4_8.

[BM90] Mihir Bellare and Silvio Micali. Non-interactive oblivious transfer and applications. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, Lecture Notes in Computer Science, pages 547–557, New York, NY. Springer, 1990. ISBN: 978-0-387-34805-6. DOI: 10.1007/0-387-34805-0_48.

[BOV03] Boaz Barak, Shien Jin Ong, and Salil Vadhan. Derandomization in Cryptography. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, pages 299–315, Berlin, Heidelberg. Springer, 2003. ISBN: 978-3-540-45146-4. DOI: 10.1007/978-3-540-45146-4_18.

[BP15] Nir Bitansky and Omer Paneth. ZAPs and Non-Interactive Witness Indistinguishability from Indistinguishability Obfuscation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography*, pages 401–427, Berlin, Heidelberg. Springer, 2015. ISBN: 978-3-662-46497-7. DOI: 10.1007/978-3-662-46497-7_16.

[BS06] Mohammed Barhoush and Louis Salvail. Powerful primitives in the bounded quantum storage model, 2023-06-06. DOI: 10.48550/arXiv.2302.05724. arXiv: 2302.05724[quant-ph]. URL: http://arxiv.org/abs/2302.05724.

[CVZ20] Andrea Coladangelo, Thomas Vidick, and Tina Zhang. Non-interactive zero-knowledge arguments for QMA, with preprocessing. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020*, Lecture Notes in Computer Science, pages 799–828, Cham. Springer International Publishing, 2020. ISBN: 978-3-030-56877-1. DOI: 10.1007/978-3-030-56877-1_28.

[DFLS16] Frédéric Dupuis, Serge Fehr, Philippe Lamontagne, and Louis Salvail. Adaptive versus non-adaptive strategies in the quantum setting with applications. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016*, Lecture Notes in Computer Science, pages 33–59, Berlin, Heidelberg. Springer, 2016. ISBN: 978-3-662-53015-3. DOI: 10.1007/978-3-662-53015-3_2.

[DFR$^+$07] Ivan B. Damgård, Serge Fehr, Renato Renner, Louis Salvail, and Christian Schaffner. A tight high-order entropic quantum uncertainty relation with applications. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007*, Lecture Notes in Computer Science, pages 360–378, Berlin, Heidelberg. Springer, 2007. ISBN: 978-3-540-74143-5. DOI: 10.1007/978-3-540-74143-5_20.

[DFSS07] Ivan B. DamgÅrd, Serge Fehr, Louis Salvail, and Christian Schaffner. Secure identification and QKD in the bounded-quantum-storage model. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, pages 342–359, Berlin, Heidelberg. Springer Berlin Heidelberg, 2007. ISBN: 978-3-540-74143-5. DOI: 10.1007/978-3-540-74143-5_19.

[DFSS08] Ivan B. DamgÅrd, Serge Fehr, Louis Salvail, and Christian Schaffner. Cryptography in the bounded-quantum-storage model. *SIAM Journal on Computing*, 37(6):1865–1890, 2008. DOI: 10.1137/060651343. eprint: https://doi.org/10.1137/060651343. URL: https://doi.org/10.1137/060651343.

[DFW02] Frédéric Dupuis, Omar Fawzi, and Stephanie Wehner. Entanglement sampling and applications. *IEEE Transactions on Information Theory*, 61(2):1093–1112, 2015-02. ISSN: 1557-9654. DOI: 10.1109/TIT.2014.2371464. Conference Name: IEEE Transactions on Information Theory.

[DLS12] Frédéric Dupuis, Philippe Lamontagne, and Louis Salvail. Fiat-shamir for proofs lacks a proof even in the presence of shared entanglement, 2022-09-12. DOI: 10.48550/arXiv.2204.02265. arXiv: 2204.02265[quant-ph]. URL: http://arxiv.org/abs/2204.02265 (visited on 09/20/2022).

[FS09] Serge Fehr and Christian Schaffner. Composing quantum protocols in a classical environment. In Omer Reingold, editor, *Theory of Cryptography*, Lecture Notes in Computer Science, pages 350–367, Berlin, Heidelberg. Springer, 2009. ISBN: 978-3-642-00457-5. DOI: 10.1007/978-3-642-00457-5_21.

[FS87] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology — CRYPTO' 86*, pages 186–194, Berlin, Heidelberg. Springer Berlin Heidelberg, 1987. ISBN: 978-3-540-47721-1. URL: https://link.springer.com/content/pdf/10.1007/3-540-47721-7_12.pdf.

[FS90a] U. Feige and A. Shamir. Witness indistinguishable and witness hiding protocols. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing - STOC '90*. the twenty-second annual ACM symposium, pages 416–426, Baltimore, Maryland, United States. ACM Press, 1990. ISBN: 978-0-89791-361-4. DOI: 10.1145/100216.100272. URL: http://portal.acm.org/citation.cfm?doid=100216.100272.

[FS90b] U. Feige and A. Shamir. Zero Knowledge Proofs of Knowledge in Two Rounds. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, Lecture Notes in Computer Science, pages 526–544, New York, NY. Springer, 1990. ISBN: 978-0-387-34805-6. DOI: 10.1007/0-387-34805-0_46.

[GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the fiat-shamir paradigm. In *44th Symposium on Foundations of Computer Science *FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 102–113. IEEE Computer Society, 2003. DOI: 10.1109/SFCS.2003.1238185. URL: https://doi.org/10.1109/SFCS.2003.1238185.

[GK96] Oded. Goldreich and Hugo. Krawczyk. On the Composition of Zero-Knowledge Proof Systems. *SIAM Journal on Computing*, 25(1):169–192, 1996. ISSN: 0097-5397. DOI: 10.1137/S0097539791220688. URL: https://epubs.siam.org/doi/abs/10.1137/S0097539791220688 (visited on 10/17/2019).

[GKR15] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. *J. ACM*, 62(4):27:1–27:64, 2015. DOI: 10.1145/2699436.

[GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 174–187, 1986. DOI: 10.1109/SFCS.1986.47.

[GO01] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, 1994-12-01. ISSN: 1432-1378. DOI: 10.1007/BF00195207. URL: https://doi.org/10.1007/BF00195207 (visited on 10/17/2019).

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive Zaps and New Techniques for NIZK. In Cynthia Dwork, editor, *Advances in Cryptology - CRYPTO 2006*, pages 97–111, Berlin, Heidelberg. Springer, 2006. ISBN: 978-3-540-37433-6. DOI: 10.1007/11818175_6.

[KMO90] Joe Kilian, Silvio Micali, and Rafail Ostrovsky. Minimum resource zero-knowledge proofs. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, Lecture Notes in Computer Science, pages 545–546, New York, NY. Springer, 1990. ISBN: 978-0-387-34805-6. DOI: 10.1007/0-387-34805-0_47.

[KW00] Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 608–617, 2000. DOI: 10.1145/335305.335387.

[KWW03] R. Konig, S. Wehner, and J. Wullschleger. Unconditional security from noisy quantum storage. *IEEE Transactions on Information Theory*, 58(3):1962–1984, 2012-03. ISSN: 1557-9654. DOI: 10.1109/TIT.2011.2177772. Conference Name: IEEE Transactions on Information Theory.

[LC01] Hoi-Kwong Lo and H. F. Chau. Why quantum bit commitment and ideal quantum coin tossing are impossible. *Physica D: Nonlinear Phenomena*. Proceedings of the Fourth Workshop on Physics and Consumption, 120(1):177–187, 1998-09-01. ISSN: 0167-2789. DOI: 10.1016/S0167-2789(98)00053-0. URL: https://www.sciencedirect.com/science/article/pii/S0167278998000530.

[LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992. DOI: 10.1145/146585.146605.

[May28] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78(17):3414–3417, 1997-04-28. DOI: 10.1103/PhysRevLett.78.3414. URL: https://link.aps.org/doi/10.1103/PhysRevLett.78.3414 (visited on 03/20/2023). Publisher: American Physical Society.

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Classically verifiable NIZK for QMA with preprocessing. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology - ASIACRYPT 2022 - 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5-9, 2022, Proceedings, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 599–627. Springer, 2022. DOI: 10.1007/978-3-031-22972-5_21. URL: https://doi.org/10.1007/978-3-031-22972-5%5C_21.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692, pages 89–114. Springer, 2019. DOI: 10.1007/978-3-030-26948-7_4.

[RRR21] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. *SIAM J. Comput.*, 50(3), 2021. DOI: 10.1137/16M1096773.

[Sha92] Adi Shamir. Ip = pspace. *J. ACM*, 39(4), 1992. DOI: 10.1145/146585.146609.

[Shm21] Omri Shmueli. Multi-theorem Designated-Verifier NIZK for QMA. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021*, Lecture Notes in Computer Science, pages 375–405, Cham. Springer International Publishing, 2021. ISBN: 978-3-030-84242-0. DOI: 10.1007/978-3-030-84242-0_14.

[STW09] Christian Schaffner, Barbara M. Terhal, and Stephanie Wehner. Robust cryptography in the noisy-quantum-storage model. *Quantum Information & Computation*, 9(11):963–996, 2009. URL: http://www.rintonpress.com/xxqic9/qic-9-1112/0963-0996.pdf.

[Unr11] Dominique Unruh. Concurrent composition in the bounded quantum storage model. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, Lecture Notes in Computer Science, pages 467–486, Berlin, Heidelberg. Springer, 2011. ISBN: 978-3-642-20465-4. DOI: 10.1007/978-3-642-20465-4_26.

[Wat01] John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009-01-01. ISSN: 0097-5397. DOI: 10.1137/060670997. URL: https://epubs.siam.org/doi/abs/10.1137/060670997 (visited on 12/09/2020). Publisher: Society for Industrial and Applied Mathematics.

[WST05] Stephanie Wehner, Christian Schaffner, and Barbara M. Terhal. Cryptography from noisy storage. *Physical Review Letters*, 100(22):220502, 2008-06-05. DOI: 10.1103/PhysRevLett.100.220502. URL: https://link.aps.org/doi/10.1103/PhysRevLett.100.220502 (visited on 08/06/2020). Publisher: American Physical Society.

[WW08] Stephanie Wehner and Jürg Wullschleger. Composable security in the bounded-quantum-storage model. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming*, Lecture Notes in Computer Science, pages 604–615, Berlin, Heidelberg. Springer, 2008. ISBN: 978-3-540-70583-3. DOI: 10.1007/978-3-540-70583-3_49.

A A New String Commitment Scheme in the BQSM

The starting point of our new protocol is a more powerful uncertainty relation found in [BFGS13] and described below. We present our protocol, which we call ABO-BC for the all-but-one uncertainty relation it crucially relies on, and prove its security.

A.1 All-but-one Uncertainty Relation

We use an uncertainty relation from [BFGS13]. It states that for a given quantum state ρ and a family of bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ that have a small overlap, there exists a basis J' (defined as a random variable whose distribution depends on ρ) such that for any other basis $J \neq J'$, the uncertainty of the measurement outcome in basis J is high.

Formally, let $\mathcal{B}_i := \{|x\rangle_i \mid x \in \{0,1\}^N\}$ and define the maximal overlap of bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ as $c := \max\{\langle x|_i |y\rangle_j \mid x, y \in \{0,1\}^N, i \neq j\}$. Let $\delta := -\frac{1}{n}\log c^2$. The uncertainty relation is as follows.

Theorem 10 (Theorem 9 of [BFGS13]). *Let ρ be an arbitrary N–qubit state, let J be a random variable over $[n]$ with distribution P_J, and let X be the outcome of measuring ρ in basis $\mathcal{B}_J$. Then for any $0 < \epsilon < \delta/4$, there exists a random variable J' with joint distribution $P_{JJ'X}$ such that*

- *J and J' are independent and*
- *there exists an event Ψ with $\Pr[\Psi] \geq 1 - 2 \cdot 2^{-\epsilon n}$ such that*

$$H_\infty(X|J=j, J'=j', \Psi) \geq \left(\frac{\delta}{2} - 2\epsilon\right) N - 1 \tag{17}$$

for all $j, j' \in [n]$ with $j \neq j$ and $P_{JJ'|\Psi}(j, j') > 0$.

As emphasized in [BFGS13], the distribution of J does not need to be set for J' to be well defined. In particular, the distribution of J' is fully determined by ρ.

We now present how to efficiently construct a family of bases with large overlap δ. Let G be the generator matrix of a linear $[N, n, d]$–error correcting code. Then for the family of bases defined by

$$\mathcal{B}_j := \{(H^{c_1} \otimes \cdots \otimes H^{c_N})|x\rangle \mid x \in \{0,1\}^N, c = G \cdot j\} \tag{18}$$

for $j \in \{0,1\}^n$ satisfies $\delta = \frac{d}{N}$.

A.2 The Commitment Scheme

Our new bit commitment scheme is presented below. The intuition behind the scheme is that the basis used by the committer to commit to a string a should be far from the basis of $a' \neq a$. Therefore, we can use code words of an error correcting code as the bases to ensure this distance holds. The original DFSS-BC scheme (presented in Section 2.1.1) can be seen as employing the repetition code (where one commits to a bit b by measuring in basis $bb \ldots b$).

Protocol ABO-BC

Setup: The generator matrix G of a $[N, n, d]$ linear error correcting code.

Commit phase:

1. V sends $|x\rangle_\theta$ for $x \in \{0,1\}^N$ and $\theta \in \{+, \times\}^N$ to the committer.
2. C commits to a string $a \in \{0,1\}^n$ by measuring each qubit i in basis $(G \cdot a)_i$, obtaining a measurement outcome $z \in \{0,1\}^N$.

Reveal phase:

3. To open the commitment, C sends a and z to V who checks that $z_i = x_i$ whenever $\theta_i = (G \cdot a)_i$.

Intuitively, we would like the basis J' from Theorem 10 to define the value to which the sender is committed in the sense of Definition 1. The proof would have the verifier purify its actions and perform the measurement in basis a when the sender opens the commitment. Theorem 10 would ensure the existence of an a' such that the sender is committed to a'. There is a subtle issue that prevents us from applying this argument: the random variable J' whose existence is stated by Theorem 10 exists in the probability space of X, the measurement outcome of the receiver in the opening phase. Therefore, we cannot assert that J' exists and that the sender is committed to it in the sense of Definition 1. Nevertheless, the techniques from [BFGS13] allows us to prove a weaker statement, namely that the commitment scheme is sum-binding.

Theorem 11. *The string commitment protocol ABO-BC is sum–biding:*

$$\sum_a p_a \leq 1 + \mathsf{negl}(n) \tag{19}$$

Proof. We consider an equivalent protocol (from the committer's point of view) where the verifier purifies its actions:

1. **Commit phase:** V prepares N EPR pairs $\bigotimes_{i=1}^{N} \frac{1}{\sqrt{2}}(|00\rangle_{P_iV_i} + |11\rangle_{P_iV_i})$ and sends registers $P_1 \dots P_N$ to C.
2. **Reveal phase:** After receiving $(a, z) \in \{0,1\}^{2N}$ from C, V measures its register V in basis a and checks that the result x matches z for each position i in a random sample $I \subseteq [N]$.

Let $\mathcal{E}_{P\to EW}$ be the CPTP map describing the partial measurement of $\tilde{\mathsf{C}}$ after the commit phase, where $\dim E \leq 2^q$. The joint state of V and $\tilde{\mathsf{C}}$ is the density operator

$$\rho_{EWV} := \sum_w P_W(w)|w\rangle\langle w| \otimes \rho_{EV}^w = (\mathcal{E}_P \otimes \mathbb{I}_V)(|EPR\rangle_{PV}^{\otimes N}) \ . \tag{20}$$

In general, $\tilde{\mathsf{C}}$ may perform a measurement on its quantum register E to decide which string a to announce in the reveal phase. The most general strategy for $\tilde{\mathsf{C}}$ is a POVM $\mathcal{M} = \{M_{EW}^{a,z}\}_{(a,z)\in\{0,1\}^{2N}}$ where $\operatorname{tr}(M^{a,z} \cdot \rho_{EW})$ gives the probability that $\tilde{\mathsf{C}}$ sends (a, z) in the reveal phase. The probability that $\tilde{\mathsf{C}}$ successfully decommits to a is given by

$$\Pr[A = a \wedge \mathsf{V} \text{ accepts}] = \sum_z \operatorname{tr}(M_{EW}^{a,z} \otimes \mathbb{V}_V^{a,z} \rho_{EWV}) \tag{21}$$

where $\mathbb{V}^{a,z}$ is the projective measurement operator corresponding to V's check in the reveal phase.

Consider a fixed $W = w$ and the reduced state ρ_{EV}^w. For $a \in \{0,1\}^n$, let $\mathcal{S}^a := \{x \mid \langle x|_a \rho_V^w |x\rangle_a \leq 2^{-\epsilon N}\}$ be the set of outcomes x that have small probability of being observed and let $\mathcal{L}^a = \{0,1\}^N \setminus \mathcal{S}^a$ its complement. Let $Q^a(x) = \langle x|_a \rho_V^w |x\rangle_a$ and $Q^a(\mathcal{X}) = \sum_{x\in\mathcal{X}} Q^a(x)$ for $\mathcal{X} \subseteq \{0,1\}^N$. By Theorem 7 of [BFGS13],

$$\sum_{a\in\{0,1\}^n} Q^a(\mathcal{L}^a) \leq 1 + c \cdot 2^n \cdot \max_{a\neq a'} \sqrt{|\mathcal{L}^a||\mathcal{L}^{a'}|} \tag{22}$$

where $c = \max_{a\neq a',x,y} \langle x|_a |y\rangle_{a'} \leq 2^{-\frac{d}{2}}$. Since $Q^a(x)$ forms a probability distribution over x and $Q^a(x) > 2^{-\epsilon N}$ for all $x \in \mathcal{L}^a$, we have that $|\mathcal{L}^a| < 2^{\epsilon N}$. We thus have that (22) is bounded above by $1 + 2^{n-d/2+(1-\epsilon)N}$. Let $\eta = 2^{n-d/2+\epsilon N}$.

Define $\mathbb{L}_a$ and $\mathbb{S}_a$ the projectors onto $\mathcal{L}_a$ and $\mathcal{S}_a$, respectively. Observe that $\mathbb{L}_a + \mathbb{S}_a = \mathbb{I}$. The probability of successful opening to *any* a is at most

$$\begin{aligned}
&\sum_z \operatorname{tr}(M_E^{a,z} \otimes \mathbb{V}_V^{a,z} \rho_{EV}^w)\\
&= \sum_{a,z} \operatorname{tr}(M_E^{a,z} \otimes \mathbb{V}_V^{a,z} (\mathbb{L}_a + \mathbb{S}_a)\rho_{EV}^w)\\
&\leq \sum_a \operatorname{tr}(\mathbb{L}_a \cdot \rho_V^w) + \sum_{a,z} \operatorname{tr}(M_E^{a,z} \otimes \mathbb{V}_V^{a,z} \cdot \mathbb{S}_a \cdot \rho_{EV}^w)
\end{aligned}$$

The first operand in the sum above corresponds to $\sum_a Q^a(\mathcal{L}^a)$ which is bounded above by $1 + \eta$. The second operand can be upper-bounded by

$$2^q \max_{a,z} \operatorname{tr}(\mathbb{V}_V^{a,z} \cdot \mathbb{S}_a \cdot \rho_V^w) \lesssim 2^{q-\frac{\epsilon}{2}N}$$

since the trace corresponds to the probability of guessing a random subset of a low-probability ($2^{-\epsilon N}$) outcome.

The term η can be made $\mathsf{negl}(n)$ with an appropriate choice of parameters[7], if we let p_a denote the probability that $\tilde{\mathsf{C}}$ successfully opens string a, we have that

$$\sum_a p_a \leq 1 + \mathsf{negl}(n) \tag{23}$$

□

Observe that to commit to a string of length n, protocol DFSS-BC above requires sending n^2 qubits from the verifier to the committer.

B Witness hiding of $\mathsf{NIP}[\Pi]$

Definition 8 (Witness Hiding)**.** Let R be an NP relation, let G be a hard instance generator for R and let Σ be a proof system for R. We say that Σ is *witness hiding* (WH) if there exists a PPT witness extractor M such that for any non-uniform PPT V', for any instance x,

$$\Pr[(x, w') \in R \mid w' = \langle P(x, w), V'(x)\rangle] \leq \Pr[(x, w') \in R \mid w' = M^{V',G}(x)] + \mathsf{negl}(n)$$

where the probability is (in part) over $x = G(1^n)$.

We can show that if a Σ–protocol Π is witness hiding, so is $\mathsf{NIP}[\Pi]$. We notice that this could also be extended to a Ξ–protocol with an inverse polynomial multiplicative factor on the success of the extractor M.

Theorem 12. *If* Π *is a witness hiding* Σ*–protocol with* $O(\lg \lambda)$*–bit challenges, then* $\mathsf{NIP}[\Pi]$ *is witness hiding*[8]*.*

Proof. We want to reduce the witness hiding property of $\mathsf{NIP}[\Pi]$ to that of Π. That is, given a malicious BQS verifier V against $\mathsf{NIP}[\Pi]$ that produces a witness with some probability, we construct a verifier V_Π against Π that produces a witness with essentially the same probability. For simplicity, we assume challenges are single bits $c \in \{0, 1\}$. The proof for logarithmic length challenges is almost identical.

Verifier V_Π is constructed as follows: in its interaction with the prover P_Π, it selects its challenge c uniformly at random. After the interaction with P_Π, V_Π is left with a transcript (a, c, r). Now to produce a witness, V_Π acts as the prover in an interaction with V. It prepares and sends the quantum state for the oblivious transfers as P would. For its classical message, V_Π uses a and r_c from the transcript received from P for and sets r_{1-c} to a uniformly random value. By Theorem 2, there exists a random variable C such that the value of r_{1-C} is statistically hidden from V. With probability $\Pr[C = c]$, the view of V in its interaction with V_Π will be indistinguishable to its view in an interaction with P. If V produces a valid witness with some probability p, the probability that V_Π outputs w is at least $\Pr[C = c] \cdot p$.

At this point, an issue occurs if C never takes value c, i.e. $\Pr[C = c] = 0$ for the particular choice of c by V_Π. This can easily be fixed by having the prover in protocol $\mathsf{NIP}[\Pi]$ randomize the transcript order. With equal probability, the prover uses either (r_0, r_1) or (r_1, r_0) as inputs for the OT. The transcript that V receives is now uniformly random, such that $\Pr[C = c] = \frac{1}{2}$

In the context of witness hiding, there is no auxiliary input to the verifier, so V_Π can run V again with the same transcript multiple times such that with overwhelming probability,

[7] If we pick $0 < \epsilon \ll \frac{1}{2}$ and a code with $N = c \cdot n$ for big enough c and optimal distance $O(N/2)$, then $\eta = \mathsf{negl}(n)$. For concreteness, pick $\epsilon = \frac{1}{16}$ and $N = 8n$, we have that $\eta \leq 2^{-\frac{n}{2}}$.

[8] With a slight modification explained in the proof.

at least one of the runs will provide V with the correct view (i.e. it will obtain the transcript (a, c, r) that V_Π received from P_Π), in which case it will produce a witness with probability p. The strategy of V_Π is to simulate V k times and if any of the simulations V produces a witness w, V_Π outputs w. Using this strategy, we have

$$\begin{aligned}
&\Pr[(x, w') \in R \mid w' \leftarrow \langle P_\Pi(x, w), V_\Pi(x)\rangle] \\
&= \Pr[(x, w') \in R \mid w' \leftarrow \langle V_\Pi(x), V(x)\rangle] \\
&\geq \Pr[(x, w') \in R \mid w' \leftarrow \langle V_\Pi(x), V(x)\rangle \mid \exists i : C_i = c] \cdot \Pr[\exists i : C_i = c] \\
&\geq \Pr[(x, w') \in R \mid w' \leftarrow \langle P(x, w), V(x)\rangle] - 2^{-k} - 2^{-\frac{n}{4}+\ell+q}
\end{aligned}$$

where the last inequality follows from the fact that conditioning on $C_i = c$, the view of V in the i simulated execution has trace distance at most $2^{-\frac{n}{4}+\ell+q}$ from the view in the real execution by Theorem 2.

□

C Triviality of Quantum 2–Message Zero-Knowledge Proofs

In this section, we present a quantum version of the impossibility of zero-knowledge 2–message *quantum* proof systems for hard languages. This generalizes the impossibility of [GO01] to quantum protocols.

Theorem 13. *Let $\Pi = \langle \mathsf{P}, \mathsf{V}\rangle$ be a 2–message quantum proof system for a language L. If Π is computationally ϵ–sound for $\epsilon < \frac{1}{3}$ and computationally zero-knowledge, then $L \in \mathsf{BQP}$.*

We assume, without loss of generality, that the the verifier is purified, i.e., we assume that the general structure of the two-message protocol is as follows:

1. V prepares a state $|\psi\rangle_{PV}$ and sends register P to P.
2. P applies some transform on register P and returns a register P' to V.
3. V applies a binary-outcome measurement $\{V_0^x, V_1^x\}$ on registers $P'V$ and accepts iff outcome is 0.

Let us assume that this protocol is auxiliary-input quantum ZK, i.e., there exits a polynomial time quantum simulator Sim such that for any $\tilde{\mathsf{V}}$ the output of $\tilde{\mathsf{V}}$ on input x and ρ in a real interaction is indistinguishable from $\mathsf{Sim}_{\tilde{\mathsf{V}}}(x, \rho)$.

Consider the cheating verifier V^* that

1. On common input x and auxiliary input register E (of same dimension as P), sends register E as the first message.
2. On reception of the prover message in quantum register P', output this register P'.

This verifier runs in polynomial time, and so does its simulator.

Then consider the following QPT machine M_L for deciding if $x \in L$. Lemmas 5 and 6 below show that this is indeed a QPT algorithm for deciding L which errs with probability at most $\frac{1}{3}$.

BQP algorithm M_L

1. Run the first message function of the honest verifier V on input x to get a register P and an internal register V.

2. Run the simulator for V^* on input x and register P. Let P' be the output register.
3. Run the verification circuit of V on registers $P'V$. Output "yes" if V accepts and "no" otherwise.

Lemma 5 (BQP Completeness)**.** *If* Π *is an* $\frac{2}{3}$*–correct quantum auxiliary-input zero-knowledge proof of language membership for* L*, then for all* $x \in L$*,* M_L *accepts on input* x *with probability at least* $\frac{2}{3}$*.*

Proof. Since Π is zero-knowledge, for any cheating verifier V^*, there exists a BQP machine $\mathsf{Sim}_{\mathsf{V}^*}$ such that the quantum map induced by the interaction of P and V^* on the auxiliary input of V^* is indistinguishable from the quantum map $\mathsf{Sim}_{\mathsf{V}^*}(x, \cdot)$.

Let $\psi_{PV} = \mathsf{V}(x)$ and let $\mathsf{D}(\rho) := \mathrm{tr}(V_0^x \rho)$. Let $\Psi_x := \mathsf{P}(x, w) \leftrightharpoons \mathsf{V}^*(x, \cdot)$ and $\Phi_x := \mathsf{Sim}_{\mathsf{V}^*}(x, \cdot)$ be the real and simulated maps acting on the auxiliary information of the verifier. Observe that the quantity $\mathsf{D}(\Psi_x \otimes \mathbb{I}_V(\psi_{PV}))$ corresponds to the probability that the verifier accepts in the real protocol and $\mathsf{D}(\Phi_x \otimes \mathbb{I}_V(\psi_{PV}))$ is the probability that M_L accepts on input $x \in L$. By the assumption that the scheme is zero-knowledge,

$$\|\mathsf{D}(\Psi_P^x \otimes \mathbb{I}_V(\psi_{PV})) - \mathsf{D}(\Phi_P^x \otimes \mathbb{I}_V(\psi_{PV}))\| \leq \mathsf{negl}(n) \ . \tag{24}$$

This means that M_L accepts with essentially the same probability with which V accepts in the interactive proof, which is at least $\frac{2}{3}$.

□

Lemma 6 (BQP Soundness)**.** *If* $x \notin L$*, then* M_L *rejects input* x *with probability* $\epsilon > \frac{2}{3}$*.*

Proof. Consider the cheating prover P^* that acts as follows: on common input x and register P received from V, compute $P' = \mathsf{Sim}_{\mathsf{V}^*}(x, P)$ and reply P' to V. Then the probability that V accepts in this interaction with a cheating prover is equal to the probability that M_L accepts, which by soundness of the interactive proof is at most ϵ. □