Quantum Merkle Trees

Lijie Chen[1] and Ramis Movassagh[2]

[1]Miller Institute for Basic Research in Science, University of California, Berkeley, CA, 94720, USA
[2]Google Quantum AI, Venice, CA, 90291, USA
IBM Quantum, MIT-IBM Watson AI Lab, Cambridge, MA, 02142, USA

arXiv:2112.14317v4 [quant-ph] 14 Jun 2024

Committing to information is a central task in cryptography, where a party (typically called a prover) stores a piece of information (*e.g.*, a bit string) with the promise of not changing it. This information can be accessed by another party (typically called the verifier), who can later learn the information and verify that it was not meddled with. Merkle trees [1] are a well-known construction for doing so in a succinct manner, in which the verifier can learn any part of the information by receiving a short proof from the honest prover. Despite its significance in classical cryptography, there was no quantum analog of the Merkle tree. A direct generalization using the Quantum Random Oracle Model (QROM) [2] does not seem to be secure. In this work, we propose the *quantum Merkle tree*. It is based on what we call the *Quantum Haar Random Oracle Model* (QHROM). In QHROM, both the prover and the verifier have access to a *Haar* random quantum oracle $\mathcal{G}$ and its inverse.

Using the quantum Merkle tree, we propose a succinct quantum argument for the Gap-k-Local-Hamiltonian problem. Assuming the Quantum PCP conjecture is true, this succinct argument extends to all of QMA. This work raises a number of interesting open research problems.

1 Introduction

A commitment scheme [3] is a cryptographic primitive that allows a party (*i.e.*, a prover) to (1) commit to a piece of information such as a bit string while keeping it hidden from others and (2) reveal the information they have committed to later. Commitment schemes are designed to ensure that a party cannot change the information after they have committed to it. Commitment schemes have numerous applications in cryptography, such as the construction of protocols for secure coin flipping, zero-knowledge proofs, and secure computation.

The Merkle tree [1] is an efficient example of commitment schemes, which captures the following scenario: There are two parties, the prover $\mathcal{P}$ and the verifier $\mathcal{V}$. $\mathcal{P}$ first computes a short string called the commitment which is denoted by $\mathsf{commit}(x)$ from a long input string x and sends $\mathsf{commit}(x)$ to $\mathcal{V}$. Then $\mathcal{V}$ asks $\mathcal{P}$ to reveal a subset of bits of x together with a short message that would enable $\mathcal{V}$ to verify that the string x has not been altered. The security promise is that after $\mathcal{P}$ sends $\mathsf{commit}(x)$ to $\mathcal{V}$, then upon $\mathcal{V}$'s request of any subset of bits, a computational bounded $\mathcal{P}$ can only reveal those bits faithfully. Namely, if $\mathcal{P}$ claims that the i-th bit of x is the wrong value $1-x_i$, then her claim will be rejected by $\mathcal{V}$ with high probability.

Lijie Chen: lijiechen@berkeley.edu
Ramis Movassagh: q.eigenman@gmail.com

The Merkle tree has wide applications in cryptography since it allows $\mathcal{P}$ to *delegate* a potentially very long string to $\mathcal{V}$ (*i.e.*, a database) while enabling $\mathcal{V}$ to maintain an *efficient verifiable* random access to that string (say to any subset of the bits of the string). A well-known application of the Merkle tree is the succinct arguments for NP from probabilistically checkable proofs [4, 5] or interactive oracle proofs [6], where by succinctness one means that the total communication between the prover and verifier constitutes a small number of bits, say polylog(n) bits of communication.

Despite being very influential in (classical) cryptography, there is no known quantum analog of the Merkle tree that allows committing to quantum states. Such a quantum analog is appealing since it would allow a party to commit to a large quantum state σ while maintaining verifiable access to individual qubits.

Protocols based on the classical Merkle tree are often analyzed in the random oracle model. There are also quantum models such as the Quantum Random Oracle Model [2] (QROM) for analyzing the quantum attacks against the classical Merkle tree. There are works showing that classical Merkle-tree-based protocols are secure against quantum attacks [7, 8]. These works showed that commitment to classical bit strings by the Merkle tree cannot be broken by quantum adversaries. Here we hope to obtain a quantum analog of the Merkle tree that can be used to commit to quantum states.

In this work, we propose a new random oracle model which we call the Quantum Haar Random Oracle Model (QHROM). We then use it in our construction of the *Quantum Merkle tree*. We then use it to propose a quantum analog of Kilian's succinct argument for NP and conjecture its security.

1.1 The Merkle Tree Algorithm

Our definition of QHROM is motivated by our adaptation of the Merkle tree to the quantum setting, so it is instructive to recall the standard Merkle tree algorithm.

Let $b \in \mathbb{N}$ be the block-length parameter. We assume that both $\mathcal{P}$ and $\mathcal{V}$ have access to a random oracle function $h\colon \{0,1\}^{2b} \to \{0,1\}^b$. For simplicity of the argument, we will first focus on the simplest non-trivial case of a Merkle tree with two leaves and depth one, and take $n = 2b$ to be the length of the string that $\mathcal{P}$ wishes to commit to. Here the string x resides on the leaves and $\mathsf{commit}(x)$ string resides on the root. As we will see shortly, a straightforward adaption of the Merkle tree to the quantum setting is not secure even in this simple setting.

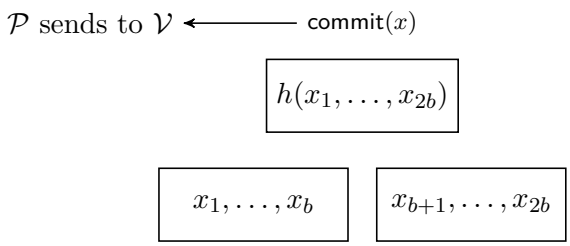

Figure 1: An illustration of the toy example for the classical Merkle tree

In this simplified setting, the protocol starts by $\mathcal{P}$ simply sending the hash value $\tilde{h} = h(x)$ of $x \in \{0,1\}^{2b}$ as the $\mathsf{commit}(x)$ of length b to $\mathcal{V}$ (see Figure 1 for an illustration). Then $\mathcal{V}$ requests the values of a subset of bits in x, for which the honest $\mathcal{P}$ simply responds by revealing the whole string x to $\mathcal{V}$. Then $\mathcal{V}$ checks that the string has the same hash value $\tilde{h}$. If a (dishonest) $\mathcal{P}$ can first commit to x and later convince $\mathcal{V}$ that its i-th bit is $1 - x_i$, then $\mathcal{P}$ has found two strings $x \neq \tilde{x}$ with $h(x) = h(\tilde{x})$. This requires at least $2^{b/2}$ queries to the random oracle h due

to the birthday paradox, which is infeasible.

1.2 A Failed Attempt to Adapt Merkle Tree in the Quantum Setting

Let us see how one might directly try to adapt the special case above of the Merkle tree algorithm to the quantum setting. An immediate idea is that, given a $2b$-qubit quantum state $|\psi\rangle = \sum_z \alpha_z |z\rangle$ in the register denoted by **data**, $\mathcal{P}$ treats h as a quantum oracle O_h[1], creates b qubits initialized to $|0^b\rangle$ in register **com**, applies O_h to both **data** and **com** to obtain $\sum_z \alpha_z |z\rangle|h(z)\rangle$, and sends the **com** register to $\mathcal{V}$; see Figure 2 for an illustration.

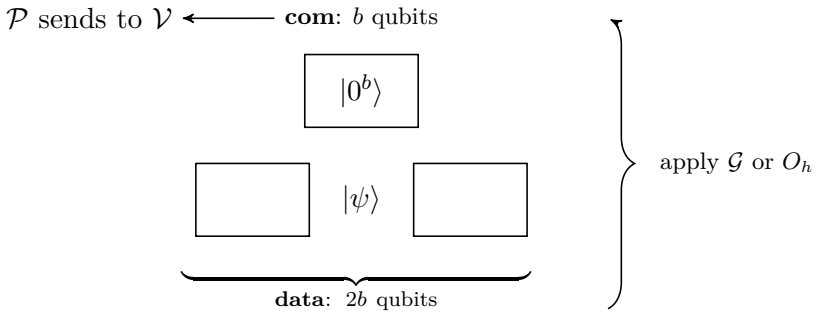

Figure 2: An illustration of the toy example for the quantum Merkle tree

To reveal qubits in $|\psi\rangle$, $\mathcal{P}$ simply sends the **data** register to $\mathcal{V}$ as well, and $\mathcal{V}$ applies O_h again to both **data** and **com**, and measures **com** in the computational basis to check if it is 0^b and rejects immediately otherwise. However, this is not secure against *phase attack*. After sending **com** to $\mathcal{V}$, for every Boolean function $f\colon \{0,1\}^{2b} \to \{0,1\}$, $\mathcal{P}$ can apply the unitary $|z\rangle \mapsto (-1)^{f(z)}|z\rangle$ to **data**, and then sends it to $\mathcal{V}$. One can see that $\mathcal{V}$ still accepts this state with probability 1, but $\mathcal{P}$ has cheated by changing the state from $\sum_z \alpha_z |z\rangle$ to $\sum_z (-1)^{f(z)} \alpha_z |z\rangle$, which can be entirely a different state for some function f.

The issue above is that the mapping O_h, $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus h(x)\rangle$, has too much structure to be exploited by the attacker. This immediately suggests to us to consider a more random choice of quantum oracles which indeed we take to be *the most* random choice of quantum oracles: a Haar random quantum oracle.

Comment: One way to address the phase attack above is to make O_h more complicated. For example, instead of applying O_h once to the registers **data** and **com**, we can repeatedly apply $O_h \mathsf{H}^{\otimes 2b}$ several times (H denotes the Hadamard gate). We found such a construction more cumbersome and even harder to analyze compared to a Haar random unitary. Moreover, it is conjectured [9, Section 6] that similar constructions may already be indistinguishable from a Haar random unitary (see Section 3.3 for more discussions). Hence, it seems more natural to directly work with a Haar random unitary.

1.3 The Quantum Haar Random Oracle Model (QHROM) and Quantum Merkle Tree

We are now ready to introduce the Quantum Haar Random Oracle Model (QHROM). In QHROM both $\mathcal{P}$ and $\mathcal{V}$ have access to a Haar random quantum oracle $\mathcal{G}$ and its inverse $\mathcal{G}^\dagger$ that act on $3b$ qubits (see Definition 2.1 for a precise definition). The protocol between $\mathcal{P}$ and $\mathcal{V}$ remains the same for the special case $n = 2b$ except for replacing O_h by $\mathcal{G}$. It is easy to see

[1] That is, $O_h|x\rangle|y\rangle = |x\rangle|y \oplus h(x)\rangle$, where $x \in \{0,1\}^{2b}$, $y \in \{0,1\}^b$, and $\oplus$ denotes the entry-wise addition over $\mathsf{GF}(2)$.

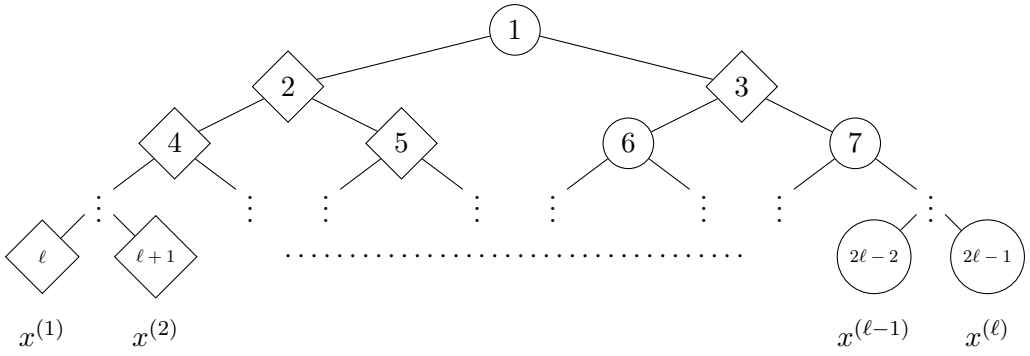

Figure 3: An illustration of the quantum Merkle tree with $\ell = 2^d$ input blocks; when block $x^{(2)}$ is requested by $\mathcal{V}$, $\mathcal{P}$ sends all the diamond shape nodes.

that since $\mathcal{G}$ completely *obfuscates* the original state $|\psi\rangle$, the phase attack described above no longer applies.

Next, we describe the quantum Merkle tree in the general setting in which n can be arbitrarily large and denotes the number of qubits in the state that $\mathcal{P}$ wishes to commit to. Given a quantum state σ on $n = b \cdot \ell$ qubits for some $\ell = 2^d$ and $d \in \mathbb{N}$,[2] we partition x into ℓ consecutive blocks of length b as $x^{(1)}, x^{(2)}, \ldots, x^{(\ell)}$. Then, we build a perfect binary tree with ℓ leaves (see Figure 4), where each leaf corresponds to a block of the input. Next, from the leaves to the root, we assign to each node α a b-qubit register $\mathbf{com}_\alpha$ as follows: (1) if α is a leaf, then $\mathbf{com}_\alpha$ is simply the qubits of the assigned block and (2) if α is an intermediate node with two children β and γ, then we initialize $\mathbf{com}_\alpha$ to $|0^b\rangle$, and apply $\mathcal{G}$ to the three registers $\mathbf{com}_\beta$, $\mathbf{com}_\gamma$, and $\mathbf{com}_\alpha$. Finally, $\mathcal{P}$ sends the register $\mathbf{com}_{\mathsf{rt}}$ to $\mathcal{V}$, where rt is the root of the binary tree.

Suppose $\mathcal{V}$ requests the state of the i-th block $x^{(i)}$ of the quantum state. To reveal the i-th block $x^{(i)}$ on a leaf (which we denote μ) of the tree $\mathcal{P}$ sends all the $\mathbf{com}_\alpha$ for nodes α that are the (1) ancestors of μ, (2) siblings of an ancestor of μ, or (3) μ or the sibling of μ. $\mathcal{V}$ then "undoes" all the applied $\mathcal{G}$ in the exact reverse order by applying $\mathcal{G}^\dagger$ to the registers sent by $\mathcal{P}$ starting from the register $\mathbf{com}_{\mathsf{rt}}$, and then from the root downwards to the leaves. After that, for every ancestor α of μ, $\mathcal{V}$ checks that $\mathbf{com}_\alpha$ is $|0^b\rangle$ by measuring it in the computational basis. To illustrate, if $\mathcal{V}$ asks for the block $x^{(2)}$, then $\mathcal{P}$ sends the corresponding $\mathbf{com}_\alpha$ registers for all diamond shape nodes in Figure 3.

Comment: So how might one heuristically instantiate a Haar-Random unitary? One might use a random quantum circuit that well approximates the behavior of a Haar unitary. For example, one might use a polynomially deep circuit. One way to formalize the degree of approximation is via the ideas in k-design [10].

1.4 A Candidate for Succinct Quantum Argument for Gap-k-LH in QHROM

Similar to Kilian's succinct argument for NP, the quantum Merkle tree naturally suggests a succinct argument Π_{succinct} for the Gap Local Hamiltonian Problem. We first recall its definition below.

Definition 1.1. (Gap-k-Local Hamiltonian Problem) Given α, β with $0 < \alpha < \beta \leq 1$ and a k-local Hamiltonian with m local terms $\{H_i\}_{i\in[m]}$ such that $0 \leq H_i \leq I$, decide whether

[2] We can always pad any quantum state to such length by adding dummy qubits. This at most doubles the number of qubits.

$\lambda_{\mathsf{min}}(\sum_{i=1}^m H_i)$ is at most αm or at least βm. Below we abbreviate this problem by (α, β)-k-LH.

Formally, in Π_{succinct} the honest prover $\mathcal{P}$ applies the quantum Merkle tree to a ground state σ of $\sum_{i=1}^m H_i$, and sends $\mathbf{com}_{\mathsf{rt}}$ to $\mathcal{V}$. Then $\mathcal{V}$ draws an integer i from $\{1, 2, \ldots, m\}$ uniformly at random and asks $\mathcal{P}$ to reveal the qubits in the support of the term H_i. $\mathcal{V}$ does the decommitment from the root towards the qubits in the support of H_i as described above. If in this decommitment phase the ancestors of the qubits in the support of H_i all result in 0^b it proceeds to the last step. In the last step, it measures the POVM $\{H_i, 1 - H_i\}$ on the qubits in the support of H_i and rejects if it sees H_i. Indeed, this is the natural analog of Kilian's succinct argument [4] in the quantum setting.

We prove that if $\mathcal{P}$ follows the protocol, then (1) when $\lambda_{\mathsf{min}}(\sum_{i=1}^m H_i) \leq \alpha \cdot m$, $\mathcal{P}$ can make $\mathcal{V}$ accept with probability at least $1 - \alpha$, and (2) when $\lambda_{\mathsf{min}}(\sum_{i=1}^m H_i) \geq \beta \cdot m$, $\mathcal{P}$ cannot force $\mathcal{V}$ to accept with a probability greater than $1 - \beta < 1 - \alpha$ (See Theorem 3.1 for details). By a sequential repetition argument, the completeness $1 - \alpha$ and the soundness $1 - \beta$ can be boosted to $1 - n^{-\omega(1)}$ and $n^{-\omega(1)}$ respectively where $\omega(1)$ means super constant.

However, a malicious $\mathcal{P}$ may not follow the protocol, but instead come up with some arbitrary states for the different nodes that are not the result of the quantum Merkle tree algorithm and send those to $\mathcal{V}$ instead. We currently do not know how to analyze such an arbitrary attack, but we conjecture the following:

Conjecture 1.2. *For the constants $k \in \mathbb{N}$ and $0 < \alpha < \beta \leq 1$, Π_{succinct} (with sequential repetition) for (α, β)-k-LH has completeness $1 - n^{-\omega(1)}$ and soundness $n^{-\omega(1)}$ in QHROM.*

1.5 Open Questions and Follow-up Works

We believe our inability to prove Conjecture 1.2 is mainly due to the lack of tools available for analyzing this new QHROM setting. We remark that only two years ago [7] managed to prove that the succinct argument for NP [4, 5] is secure in the QROM model by using the recently proposed *compressed oracles* technique introduced in [11] which gives a nice way to analyze QROM. To prove the security of our succinct argument for Gap-k-LH, one likely needs similar advances for analyzing the QHROM. We now list some specific open problems:

Open Problem 1. *Is there an analog of the compressed oracle technique in [11] for the QHROM?*

Above we generalized Kilian's constant-round succinct argument [4] to the quantum setting and conjectured its soundness. A natural open question is whether we can generalize Micali's non-interactive succinct argument for NP [5] to the quantum settings as well.

Open Problem 2. *Is there an analog of Micali's non-interactive succinct argument for Gap-k-LH?*

A particularly useful feature of previous succinct arguments for NP [4, 5] is that they can be made *zero-knowledge* with minimal overhead. A natural open question is whether we can make our proposed succinct argument for Gap-k-LH zero-knowledge as well.

Open Problem 3. *Is there a zero-knowledge succinct argument for Gap-k-LH in QHROM?*

Subsequent work. Finally, we remark that this paper formed the basis of the ideas in an exciting subsequent work [12]. They proved the security of a tree commitment similar to what is presented here but from standard (quantum) cryptographic assumptions. Note that it is not a priori clear what "security" even means for commitments to quantum states[3], and a major

[3] Although such commitment scheme has appeared implicitly in several quantum cryptographic protocols [13, 14, 15].

contribution of [12] is to formally define the security of commitments to quantum states. We refer readers to [12] for an overview of more prior works on quantum commitment schemes. As far as we know, the security of the precise protocol (in the QHROM) given in this paper remains open.

2 Preliminaries

2.1 Notation

We always denote by log the logarithm in base 2. We denote by $[n]$ the set of integers $\{1, 2, \ldots, n\}$. Let $\mathbf{reg}$ be a register of n qubits. For each $i \in [n]$, $\mathbf{reg}(i)$ denotes the i-th qubit in $\mathbf{reg}$, and $\mathbf{reg}[\ell, r]$ denotes the qubits from $\mathbf{reg}(\ell)$ to $\mathbf{reg}(r)$. The corresponding Hilbert space is denoted by $\mathcal{H}_{\mathbf{reg}}$. For k pairwise-disjoint sets $S_1, \ldots, S_k$, we use $\bigsqcup_{i\in[k]} S_i$ to denote their union. We say a function $\alpha\colon \mathbb{N} \to [0,1]$ satisfies $\alpha(n) \leq \mathsf{negl}(n)$ (*i.e.*, α is *negligible*), if for all constant $k \geq 1$, $\lim_{n\to\infty} \alpha(n) \cdot n^k = 0$ (*i.e.*, $\alpha(n) = o(n^{-k})$ for every $k \in \mathbb{N}$).

For a quantum state σ on n qubits and a subset $S \subseteq [n]$, $\sigma_S := \mathrm{Tr}_{[n]\setminus S}[\sigma]$ is the reduced density matrix. For a quantum state $|\psi\rangle \in \mathcal{H}_{\mathbf{reg}}$, for simplicity we sometimes use ψ to denote the corresponding density matrix $\psi = |\psi\rangle\langle\psi|$. Given a unitary sequence $U_1, \ldots, U_T$, we write $U_{[\ell,r]}$ to denote the product $U_r U_{r-1} \cdots U_\ell$ for ease of notation.

For two quantum states σ and ρ, we use $\|\sigma - \rho\|_1$ to denote their trace distance. We also write $x \in_{\mathsf{R}} A$ to mean that x is drawn from the set A uniformly at random.

2.2 The Quantum Haar Random Oracle Model

We will consider the *Quantum Haar random oracle model* (QHROM), in which every agent (prover and verifier) gets access to a Haar random oracle $\mathcal{G}$ acting on λ qubits and its inverse $\mathcal{G}^\dagger$, where λ is the so-called *security parameter*.

We denote by $\mathbb{U}(N)$ the set of all $N \times N$ unitaries. By $\mathcal{G} \in_{\mathsf{R}} \mathbb{U}(N)$ we mean that $\mathcal{G}$ is an $N \times N$ unitary drawn from the Haar measure.

Definition 2.1. An interactive protocol Π between the prover $\mathcal{P}$ and verifier $\mathcal{V}$ is a proof system for a promise problem $L = (L_{\mathsf{yes}}, L_{\mathsf{no}})$ with completeness $c(n, \lambda)$ and soundness $s(n, t, \lambda)$ in QHROM, if the following holds:

$\mathcal{P}$ and $\mathcal{V}$: $\mathcal{P}$ and $\mathcal{V}$ are both given an input $x \in L_{\mathsf{yes}} \cup L_{\mathsf{no}}$. $\mathcal{V}$ is polynomial-time and outputs a classical bit indicating acceptance or rejection of x, and $\mathcal{P}$ is unbounded. Both $\mathcal{V}$ and $\mathcal{P}$ are given access to a Haar random quantum oracle $\mathcal{G}$ and its inverse $\mathcal{G}^\dagger$ that act on λ qubits (that is, $\mathcal{G} \in_{\mathsf{R}} \mathbb{U}(2^\lambda)$). Let $n = |x|$.

Completeness: If $x \in L_{\mathsf{yes}}$,

$$\mathop{\mathbb{E}}_{\mathcal{G}\in_{\mathsf{R}}\mathbb{U}(2^\lambda)} \Pr[(\mathcal{V}^{\mathcal{G},\mathcal{G}^\dagger} \leftrightarrows \mathcal{P}^{\mathcal{G},\mathcal{G}^\dagger})(x) = 1] \geq c(n, \lambda).$$

Soundness: If $x \in L_{\mathsf{no}}$, for every $t \in \mathbb{N}$ and any unbounded prover $\mathcal{P}^*$ making at most t total queries to $\mathcal{G}$ and $\mathcal{G}^\dagger$, we have that

$$\mathop{\mathbb{E}}_{\mathcal{G}\in_{\mathsf{R}}\mathbb{U}(2^\lambda)} [(\mathcal{V}^{\mathcal{G},\mathcal{G}^\dagger} \leftrightarrows (\mathcal{P}^*)^{\mathcal{G},\mathcal{G}^\dagger})(x) = 1)] \leq s(n, t, \lambda).$$

In the above "$\leftrightarrows$" denotes the interactive nature of the protocol between $\mathcal{P}$ and $\mathcal{V}$.

We remark that in the soundness part, the only restriction on a malicious prover $\mathcal{P}^*$ is the number of queries it can make to $\mathcal{G}$ and $\mathcal{G}^\dagger$. In particular, this means that even if $\mathcal{P}^*$ has unbounded computational power, as long as it makes a small number of queries to $\mathcal{G}$ and $\mathcal{G}^\dagger$, it cannot fool the verifier.

2.3 Quantum Local Proofs

Next, we provide formal definitions of LocalQMA. For a reader familar with QMA in the following definition, one can think of x as the classical description of the local Hamiltonian problem, and $m(n)$ as the number of terms in it (*i.e.*, $H = \sum_{i=1}^{m} H_i$).

Definition 2.2 ((k,γ)-LocalQMA)**.** For two constants $k, \gamma \in \mathbb{N}$, a promise problem $L = (L_{\mathsf{yes}}, L_{\mathsf{no}})$ is in the complexity class (k,γ)-LocalQMA with soundness $s(n)$ and completeness $c(n)$ if there are polynomials m and p such that the following hold:

- (**A k-local verifier** V_L) Let $n = |x|$. There is a verifier V_L that acts as follows:
  1. V_L gets access to a $p(n)$-qubit proof σ for L and draws $i \in_{\mathsf{R}} [m(n)]$. V_L then computes in $\mathrm{poly}(n,k,\gamma)$ time a k-size subset $S_i \subseteq [p(n)]$ and a γ-size quantum circuit C_i that is over the Clifford + T gate-set and acts on k qubits. C_i may use γ ancilla qubits, with the first ancilla qubit being the output qubit.
  2. V_L next applies C_i to the restriction of σ on qubits in S_i and measures the first ancilla qubit. V_L accepts if the outcome is 1 and rejects otherwise.
- (**Completeness**) If $x \in L_{\mathsf{yes}}$, there is a $p(n)$-qubit state σ such that V_L accepts σ with probability at least $c(n)$.
- (**Soundness**) If $x \in L_{\mathsf{no}}$, V_L accepts every $p(n)$-qubit state σ with probability at most $s(n)$.
- (**Strongly explicit**) Moreover, we say that V_L is strongly explicit, if V_L computes S_i and C_i in $\mathrm{poly}(\log n, k, \gamma)$ time instead of $\mathrm{poly}(n,k,\gamma)$ time.

We will use (k,γ)-$\mathsf{LocalQMA}_{s,c}$ to denote the class above for notational convenience.

2.4 The Quantum PCP Conjecture

We first recall the quantum PCP conjecture [16, 17].

Conjecture 2.3 (QPCP conjecture)**.** *There are constants $k \in \mathbb{N}$ and α, β satisfying $0 < \alpha < \beta \leq 1$ such that (α,β)-k-LH is QMA-complete.*

In particular, the following corollary is immediate from the definition of (α,β)-k-LH.

Corollary 2.4. *If QPCP holds, then there are constants $k, \gamma \in \mathbb{N}$ and $c, s \in [0,1]$ satisfying that $s < c$, such that*

$$\mathsf{QMA} \subseteq (k,\gamma)\text{-}\mathsf{LocalQMA}_{s,c}.$$

3 A Candidate Succinct Argument for LocalQMA in QHROM

In this section, we present a candidate succinct argument for LocalQMA in QHROM. Assuming QPCP, this succinct argument also works for all of QMA.

3.1 The Succinct Protocol Π_{succinct}

Notation. Let $L = (L_{\mathsf{yes}}, L_{\mathsf{no}}) \in (k, \gamma)\text{-}\mathsf{LocalQMA}_{s,c}$ for two integers $k, \gamma \in \mathbb{N}$ and two reals $s, c \in [0, 1]$ such that $s < c$. Let m_L and p_L be the polynomials and V_L be the k-local verifier in Definition 2.2. Throughout this section, we will always use n to denote the length of the input to L, $N = p_L(n)$ to denote the number of qubits in a witness for V_L, and λ to denote the security parameter.

We set $b = \lambda/3$, and $\ell = N/b$. We assume that b is an integer and ℓ is a power of 2 for simplicity and without loss of generality since one can always add dummy qubits to the witness.

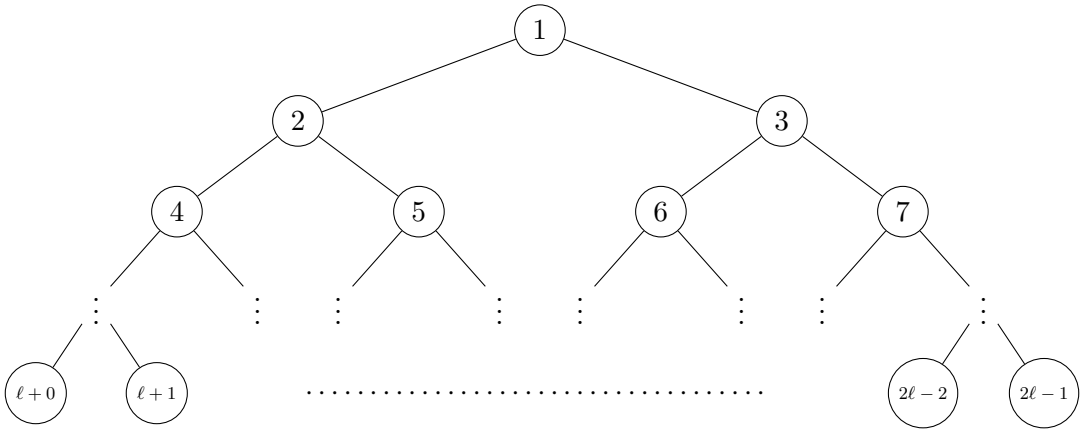

Figure 4: An illustration of the labeling of the nodes in the tree T_ℓ with ℓ leaves

The perfect binary tree T_ℓ. We will consider a perfect binary tree T_ℓ of ℓ leafs (see Figure 4 for an illustration). Note that T_ℓ has $\log \ell$ layers. We label the nodes in T_ℓ first from root to leaves and then from left to right, starting with 1.

For a node u in T_ℓ, we observe that u's parent is $\lfloor u/2 \rfloor$ if u is not the root (*i.e.*, $u \neq 1$) and u's two children are $2u$ and $2u + 1$ if u is not a leaf (*i.e.*, $u < \ell$). We use P_u to denote the set of nodes consisting of u and all ancestors of u. Formally, we have

$$P_u = \begin{cases} \{u\} & \text{if } u = 1, \\ \{u\} \cup P_{\lfloor u/2 \rfloor} & \text{if } u > 1. \end{cases}$$

We also define R_u as follows:

$$R_u = \{v \in P_u \text{ or } \lfloor v/2 \rfloor \in P_u : v \in [2\ell - 1]\}.$$

That is, a node v belongs to R_u if either v is in P_u or the parent of v is in P_u. Also, for a set of nodes $S \subseteq [2\ell - 1]$, we set $R_S = \bigcup_{u \in S} R_u$.

Given an N-qubit state σ, we define the following commitment algorithm (Algorithm 1) and the corresponding local decommitment algorithm (Algorithm 2).

Algorithm 1: Algorithm for committing to an N-qubit quantum state

Function $\mathsf{commit}^{\mathcal{G}}(\sigma, N, \lambda)$

Input: σ is an N-qubit quantum state, λ is the security parameter (recall that $\lambda = 3b$)

Let $\ell = N/b$;

For each node u in T_ℓ, create a b-qubit register $\mathbf{state}^{(u)}$;

Store σ in registers $\mathbf{state}^{(\ell)}, \mathbf{state}^{(\ell+1)}, \ldots, \mathbf{state}^{(2\ell-1)}$;

for u *from* $\ell - 1$ *down to* 1 **do**

  Initialize $\mathbf{state}^{(u)}$ as $|0^b\rangle$;

  Apply $\mathcal{G}$ on $\mathbf{state}^{(2u)}$, $\mathbf{state}^{(2u+1)}$, and $\mathbf{state}^{(u)}$;

return all qubits in $\{\mathbf{state}^{(u)}\}_{u\in[2\ell-1]}$; `// Here` $\mathbf{state}^{(1)}$ `is the commitment to be sent to the verifier, while the prover keeps all other states` $\{\mathbf{state}^{(u)}\}_{u\in\{2,\ldots,2\ell-1\}}$

Algorithm 2: Algorithm for recovering part of the original quantum state

Function $\mathsf{decommit}^{\mathcal{G}^\dagger}(N, \lambda, S, \{\eta_u\}_{u\in R_S})$

Input: $S \subseteq \{\ell, \ldots, 2\ell - 1\}$ is a subset of leafs in T_ℓ, for each $u \in R_S$, η_u is a b-qubit quantum state, λ is the security parameter. (We remark that $\{\eta_u\}_{u\in(R_S\setminus\{1\})}$ are the states provided by the prover to the verifier.)

Let $\ell = N/b$;

For each node u in R_S, create a b-qubit register $\mathbf{state}^{(u)}$, and store η_u at $\mathbf{state}^{(u)}$;

for $u \in R_S \cap [\ell - 1]$, *from smallest to the largest* **do**

  Apply $\mathcal{G}^\dagger$ on $\mathbf{state}^{(2u)}$, $\mathbf{state}^{(2u+1)}$, and $\mathbf{state}^{(u)}$;

  Measure $\mathbf{state}^{(u)}$ in the computational basis to obtain an outcome $z \in \{0,1\}^b$;

  if $z \neq 0^b$ **then**

    return $\perp$ `// ` $\perp$ `means the check fails`

return all qubits in $\{\mathbf{state}^{(u)}\}_{u\in S}$;

Finally, we are ready to specify the following candidate succinct argument for $L \in (k, \gamma)$-$\mathsf{LocalQMA}_{s,c}$.

The candidate succinct argument Π_{succinct} for $L \in (k, \gamma)$-$\mathsf{LocalQMA}_{s,c}$

- Both prover ($\mathcal{P}$) and verifier ($\mathcal{V}$) get access to a Haar random quantum unitary $\mathcal{G}$ acting on $3b = \lambda$ qubits and its inverse $\mathcal{G}^\dagger$. They also both get an input $x \in \{0,1\}^n$ to L. The goal for the prover is to convince the verifier that $x \in L_{\mathsf{yes}}$.

  Let $\ell = N/b$ and we assume that $\ell = 2^d$ for $d \in \mathbb{N}$.

- (**First message:** $\mathcal{P} \to \mathcal{V}$) The honest prover $\mathcal{P}$ acts as follows: If $x \in L_{\mathsf{no}}$, $\mathcal{P}$ aborts immediately. Otherwise, $\mathcal{P}$ finds an N-qubit state σ such that V_L accepts with probability at least c, and runs $\mathsf{commit}(\sigma, N, \lambda)$ to obtain qubits $\{\eta_u\}_{u\in[2\ell-1]}$.

  $\mathcal{P}$ then sends η_1 to $\mathcal{V}$.

- (**Second message:** $\mathcal{V} \to \mathcal{P}$) $\mathcal{V}$ now simulates the local verifier V_L: $\mathcal{V}$ first draws $i \in_{\mathsf{R}} [m_L(n)]$ and sends i to $\mathcal{P}$, and then computes a k-size subset $S_i \subseteq [N]$ and a γ-size circuit C_i acting on k qubits, according to Definition 2.2.

Let W_i be the set of leaves in T_ℓ that contain the qubits indexed by S_i. That is,

$$W_i = \{\ell + \lfloor (u-1)/b \rfloor : u \in S_i\}.$$

- (**Third message:** $\mathcal{P} \to \mathcal{V}$) The honest prover $\mathcal{P}$ sends $\{\eta_u\}_{u \in R_{W_i}, u \neq 1}$ to $\mathcal{V}$. $\mathcal{V}$ then runs $\mathsf{decommit}(N, \lambda, W_i, \{\eta_u\}_{u \in R_{W_i}})$ (note that $\mathcal{V}$ already has η_1). If $\mathsf{decommit}$ returns $\perp$, $\mathcal{V}$ rejects immediately.

  Otherwise, $\mathcal{V}$ continues the simulation of V_L by running C_i using $\{\eta_u\}_{u \in W_i}$, and $\mathcal{V}$ accepts if and only if V_L accepts.

3.2 Analysis of Π_{succinct}

We say a prover $\mathcal{P}$ is *semi-honest*, if $\mathcal{P}$ commits to an arbitrary N-qubit state (as opposed to the true ground state) σ in the first message but indeed follows Π_{succinct}. We remark that, unlike an honest prover, a semi-honest prover may not necessarily commit to a state that makes V_L accepts with probability at least c.[4]

Now we prove the completeness and succinctness of Π_{succinct}. We also show Π_{succinct} is sound against semi-honest provers.[5]

Theorem 3.1. *Let Π_{succinct} be the protocol between $\mathcal{P}$ and $\mathcal{V}$ for the promise language $L \in (k, \gamma)\text{-}\mathsf{LocalQMA}_{s,c}$. For every $x \in \{0,1\}^n$, the following hold:*

Completeness: *If $x \in L_{\mathsf{yes}}$, then for every $\mathcal{G} \in \mathbb{U}(2^\lambda)$,*

$$\Pr[(\mathcal{V}^{\mathcal{G},\mathcal{G}^\dagger} \leftrightarrows \mathcal{P}^{\mathcal{G},\mathcal{G}^\dagger})(x) = 1] \geq c.$$

Soundness against semi-honest provers: *If $x \in L_{\mathsf{no}}$, then for every $\mathcal{G} \in \mathbb{U}(2^\lambda)$ and every semi-honest prover $\mathcal{P}$,*

$$\Pr[(\mathcal{V}^{\mathcal{G},\mathcal{G}^\dagger} \leftrightarrows (\mathcal{P})^{\mathcal{G},\mathcal{G}^\dagger})(x) = 1] \leq s.$$

Succinctness: *$\mathcal{P}$ and $\mathcal{V}$ communicate at most $O(k \cdot \lambda \cdot \log n)$ qubits in total.*

Efficiency: *$\mathcal{V}$ runs in $\mathrm{poly}(n, k, \gamma)$ time. If V_L is strongly explicit, then $\mathcal{V}$ runs in $O(k \cdot \lambda \cdot \log n + \mathrm{poly}(\log n, k, \gamma))$ time.*

Proof. We first establish the succinctness part. Examining the protocol Π_{succinct}, one can see that the first message takes $O(\lambda)$ qubits, the second messages takes $O(\log m_L(n)) = O(\log n)$ classical bits, and the third message takes $O(|R_{W_i}| \cdot \lambda)$ qubits. Note that $|R_{W_i}| \leq |W_i| \cdot O(\log \ell) \leq k \cdot O(\log N) \leq O(k \cdot \log n)$, the total communication complexity is thus bounded by $O(k \cdot \lambda \cdot \log n)$.

For the running time of $\mathcal{V}$, one can see that its running time is dominated by the running time of $\mathsf{decommit}(N, \lambda, W_i, \{\eta_u\}_{u \in R_{W_i}})$ and the running time of V_L computing W_i and C_i, which are at most $O(k \cdot \lambda \cdot \log N)$ and $\mathrm{poly}(n, k, \gamma)$ respectively. The latter becomes $(\mathrm{poly}(\log n, k, \gamma)$ if V_L is strongly explicit.

[4] In particular, a semi-honest prover may still commit to some state even when $x \in L_{\mathsf{no}}$, while an honest prover would abort when $x \in L_{\mathsf{no}}$.

[5] We remark that semi-honest prover security is more of a sanity check than a solid contribution since it is not hard to construct a trivial protocol that satisfies this semi-honest prover security.

Now we prove the completeness. Let $\mathcal{G}_{(u)}$ be a $\mathcal{G}$ gate applying on registers $\mathbf{state}^{(2u)}$, $\mathbf{state}^{(2u+1)}$, and $\mathbf{state}^{(u)}$. Then we know for the honest prover $\mathcal{P}$, when $x \in L_{\mathsf{yes}}$, it prepares an N-qubit state σ that makes V_L accept with probability at least c, and then applies $U_{\mathsf{com}} \coloneqq \mathcal{G}_{(\ell-1)} \cdot \ldots \cdot \mathcal{G}_{(1)}$ to $\sigma \otimes |0\rangle\langle 0|_{\mathbf{state}^{(1)},\ldots,\mathbf{state}^{(\ell-1)}}$.

Let $U_{\mathsf{decom}} \coloneqq U_{\mathsf{com}}^\dagger = \mathcal{G}_{(1)}^\dagger \cdot \ldots \cdot \mathcal{G}_{(\ell-1)}^\dagger$. Recall that verifier $\mathcal{V}$ at the end simulates the quantum circuit C_i only on registers in $\{\mathbf{state}^{(u)}\}_{u \in W_i}$. We now argue that $\mathcal{V}$ is effectively simulating C_i on

$$U_{\mathsf{decom}}^\dagger U_{\mathsf{com}} \sigma \otimes |0\rangle\langle 0|_{\mathbf{state}^{(1)},\ldots,\mathbf{state}^{(\ell-1)}} = \sigma \otimes |0\rangle\langle 0|_{\mathbf{state}^{(1)},\ldots,\mathbf{state}^{(\ell-1)}}.$$

The reason is that $\mathsf{decommit}(N, \lambda, W_i, \{\eta_u\}_{u \in R_{W_i}})$ performs all gates in U_{decom} that reside in the lightcone of the registers $\{\mathbf{state}^{(u)}\}_{u \in W_i}$ in the chronological order (see Line 4 of Algorithm 2). Also, since $\mathcal{P}$ starts with the state $\sigma \otimes |0\rangle\langle 0|_{\mathbf{state}^{(1)},\ldots,\mathbf{state}^{(\ell-1)}}$, $\mathsf{decommit}$ never outputs $\perp$. Therefore, $\mathcal{V}$ is simulating V_L faithfully on σ, meaning that it accepts with probability at least c.

Finally, we establish the soundness against semi-honest provers. The argument above for completeness indeed established that whenever the prover commits to a state σ in the first message and follows Π_{succinct} (*i.e.*, the prover is semi-honest), for every possible $\mathcal{G}$, the acceptance probability of $\mathcal{V}$ equals the acceptance probability of the simulated V_L on σ. Hence, when $x \in L_{\mathsf{no}}$, for every semi-honest prover and every possible $\mathcal{G}$, the acceptance probability of $\mathcal{V}$ is at most s. □

We conjecture that the soundness also holds more generally.

Conjecture 3.2 (Π_{succinct} is sound in QHROM)**.** *Let Π_{succinct} be the protocol between $\mathcal{P}$ and $\mathcal{V}$ for the promise language $L \in (k,\gamma)$-$\mathsf{LocalQMA}_{s,c}$. For every $x \in \{0,1\}^n$, the following hold:*

Soundness: *If $x \in L_{\mathsf{no}}$, then for every $t \in \mathbb{N}$ and all (potentially malicious) $\mathcal{P}^*$ that makes at most t total queries to $\mathcal{G}$ and $\mathcal{G}^\dagger$, for some $\delta = \delta(t,\lambda) = \mathrm{poly}(t)/2^{\Omega(\lambda)}$, it holds that*

$$\Pr_{\mathcal{G} \in_R \mathbb{U}(2^\lambda)} \left[\Pr[(\mathcal{V}^{\mathcal{G},\mathcal{G}^\dagger} \leftrightarrows (\mathcal{P}^*)^{\mathcal{G},\mathcal{G}^\dagger})(x) = 1)] \geq s + \delta \right] \leq \delta.$$

Comment: The difference between this conjecture and our main theorem (Theorem 3.1) is that in the conjecture we require soundness against all unbounded provers instead of only against semi-honest provers.

3.3 Discussions

We remark that (1) the constant soundness in Conjecture 3.2 and the constant completeness in Theorem 3.1 can be easily amplified to $n^{-\omega(1)}$ and $1 - n^{-\omega(1)}$ by repeating the protocols $\log^2 n$ times, and (2) assuming QPCP, the protocol works for all languages in QMA.

Corollary 3.3. *Assuming Conjecture 3.2, there is a protocol for $L \in (k,\gamma)$-$\mathsf{LocalQMA}_{s,c}$ with $\lambda \cdot \mathrm{polylog}(n)$ communication complexity, completeness $1 - n^{-\omega(1)}$ and soundness $n^{-\omega(1)}$ in QHROM. Also, if V_L is strongly explicit, then the verifier running time of the protocol is also bounded by $\lambda \cdot \mathrm{polylog}(n)$.*

Moreover, if we further assume that QPCP holds, then the aforementioned succinct protocol exists for every $L \in \mathsf{QMA}$.

How easy is it for the prover to cheat after having sent the commitment to the verifier in the quantum Merkle tree construction? We believe (but cannot yet prove; see Conjecture 3.2) that a computationally bounded prover will not be able to make the verifier accept. However, a computationally unbounded prover can. We now demonstrate this by the application of Schrödinger–HJW theorem [18] to the simple toy example of a Merkle tree with depth one. Suppose $|\psi\rangle$ is the $2b$-qubit state $\mathcal{P}$ initially committed to, and $|\phi\rangle$ is another $2b$-qubit state that $\mathcal{P}$ wishes to cheat by switching $|\psi\rangle$ with. Mathematically the process of committing, switching the initial state and lastly decommitting writes

$$\mathcal{G}^\dagger(W \otimes I)\mathcal{G}(|\psi\rangle \otimes |0^b\rangle) \approx (|\phi\rangle \otimes |0^b\rangle),$$

where in the above we think of $|0^b\rangle$ as the parent and see that the initially committed $2b$-qubit state $|\psi\rangle$ can be changed to another completely different $2b$-qubit state $|\phi\rangle$ by applying Schrödinger-HJW theorem (*i.e.*, application of a $W \otimes I$) for a suitable unitary W that acts on the first $2b$ qubits. We note that such W exists, because the reduced density matrix of the last b qubits of both $\mathcal{G}(|\psi\rangle \otimes |0^b\rangle)$ and $\mathcal{G}(|\phi\rangle \otimes |0^b\rangle)$ are very close to the maximally mixed state, for any two fixed states $|\phi\rangle$ and $|\psi\rangle$. However, we conjecture that finding W requires computationally unbounded prover. For example, in the foregoing equation a direct way to solve for W would require solving a linear system of equations that is exponentially large. Moreover, the oracle $\mathcal{G}$ is fully random and does not afford any structure we can utilize to reduce the computation.

This is exacerbated by the fact that finding a unitary W that makes the two sides approximately equal can also make the verifier accept with sufficiently high probability. We leave this resolution as a mathematical challenge.

Acknowledgments

L.C. would like to thank Jiahui Liu and Qipeng Liu for helpful discussions and pointing out many related works. This work was done while L.C. did an internship at IBM Quantum Research. L. C. is supported by NSF CCF-2127597 and an IBM Fellowship.

References

[1] Ralph C. Merkle. "A digital signature based on a conventional encryption function". In Advances in Cryptology - CRYPTO '87, A Conference on the Theory and Applications of Cryptographic Techniques, Santa Barbara, California, USA, August 16-20, 1987, Proceedings. Volume 293 of Lecture Notes in Computer Science, pages 369–378. Springer (1987).

[2] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. "Random oracles in a quantum world". In Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings. Volume 7073 of Lecture Notes in Computer Science, pages 41–69. Springer (2011).

[3] Gilles Brassard, David Chaum, and Claude Crépeau. "Minimum disclosure proofs of knowledge". J. Comput. Syst. Sci. **37**, 156–189 (1988).

[4] Joe Kilian. "A note on efficient zero-knowledge proofs and arguments (extended abstract)". In S. Rao Kosaraju, Mike Fellows, Avi Wigderson, and John A. Ellis, editors, Proceedings of the 24th Annual ACM Symposium on Theory of Computing, May 4-6, 1992, Victoria, British Columbia, Canada. Pages 723–732. ACM (1992). url: https://doi.org/10.1145/129712.129782.

[5] Silvio Micali. "Computationally sound proofs". SIAM J. Comput. **30**, 1253–1298 (2000). url: https://doi.org/10.1137/S0097539795284959.

[6] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. "Interactive oracle proofs". In Martin Hirt and Adam D. Smith, editors, Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II. Volume 9986 of Lecture Notes in Computer Science, pages 31–60. (2016). url: https://doi.org/10.1007/978-3-662-53644-5_2.

[7] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. "Succinct arguments in the quantum random oracle model". In Dennis Hofheinz and Alon Rosen, editors, Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part II. Volume 11892 of Lecture Notes in Computer Science, pages 1–29. Springer (2019). url: https://doi.org/10.1007/978-3-030-36033-7_1.

[8] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. "Post-quantum succinct arguments: Breaking the quantum rewinding barrier". In 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022. Pages 49–58. IEEE (2021).

[9] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. "Pseudorandom quantum states". In Hovav Shacham and Alexandra Boldyreva, editors, Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III. Volume 10993 of Lecture Notes in Computer Science, pages 126–152. Springer (2018).

[10] Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. "Local random quantum circuits are approximate polynomial-designs". Communications in Mathematical Physics **346**, 397–434 (2016). url: https://doi.org/10.1007/s00220-016-2706-8.

[11] Mark Zhandry. "How to record quantum queries, and applications to quantum indifferentiability". In Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II. Volume 11693 of Lecture Notes in Computer Science, pages 239–268. Springer (2019).

[12] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. "Commitments to quantum states". In Barna Saha and Rocco A. Servedio, editors, Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023. Pages 1579–1588. ACM (2023).

[13] Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. "Zero-knowledge proof systems for QMA". SIAM J. Comput. **49**, 245–283 (2020).

[14] 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 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III. Volume 12172 of Lecture Notes in Computer Science, pages 799–828. Springer (2020).

[15] Anne Broadbent and Alex Bredariol Grilo. "Qma-hardness of consistency of local density matrices with applications to quantum zero-knowledge". SIAM J. Comput. **51**, 1400–1450 (2022).

[16] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh V. Vazirani. "The detectability lemma and quantum gap amplification". In Proceedings of the 41st Annual ACM Sympo-

sium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009. Pages 417–426. ACM (2009).

[17] Dorit Aharonov, Itai Arad, and Thomas Vidick. "Guest column: the quantum PCP conjecture". SIGACT News **44**, 47–79 (2013). url: https://doi.org/10.1145/2491533.2491549.

[18] Lane P Hughston, Richard Jozsa, and William K Wootters. "A complete classification of quantum ensembles having a given density matrix". Physics Letters A **183**, 14–18 (1993). url: https://doi.org/10.1016/0375-9601(93)90880-9.