2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)

QMA-hardness of Consistency of Local Density Matrices with Applications to Quantum Zero-Knowledge

Anne Broadbent
Department of Mathematics and Statistics
University of Ottawa
Ottawa, Canada
abroadbe@uottawa.ca

Alex B. Grilo
LIP6
CNRS/Sorbonne Université
Paris, France
abgrilo@gmail.com

***Abstract*—We provide several advances to the understanding of the class of Quantum Merlin-Arthur proof systems (QMA), the quantum analogue of NP. Our central contribution is proving a longstanding conjecture that the *Consistency of Local Density Matrices* (CLDM) problem is QMA-hard under Karp reductions. The input of CLDM consists of local reduced density matrices on sets of at most k qubits, and the problem asks if there is an n-qubit global quantum state that is locally consistent with all of the k-qubit local density matrices. The containment of this problem in QMA and the QMA-hardness under Turing reductions were proved by Liu [APPROX-RANDOM 2006]. Liu also conjectured that CLDM is QMA-hard under Karp reductions, which is desirable for applications, and we finally prove this conjecture. We establish this result using the techniques of *simulatable codes* of Grilo, Slofstra, and Yuen [FOCS 2019], simplifying their proofs and tailoring them to the context of QMA.**

In order to develop applications of CLDM, we propose a framework that we call *locally simulatable proofs* for QMA: this provides QMA proofs that can be efficiently verified by probing only k qubits and, furthermore, the reduced density matrix of any k-qubit subsystem of a *good* witness can be computed in polynomial time, independently of the witness. Within this framework, we show several advances in zero-knowledge in the quantum setting. We show for the first time a commit-and-open computational zero-knowledge proof system for all of QMA, as a quantum analogue of a "sigma" protocol. We then define a *Proof of Quantum Knowledge*, which guarantees that a prover is effectively in possession of a quantum witness in an interactive proof, and show that our zero-knowledge proof system satisfies this definition. Finally, we show that our proof system can be used to establish that QMA has a quantum *non-interactive* zero-knowledge proof system in the secret parameter setting. [1]

***Keywords*-Quantum complexity theory; Quantum proofs; Quantum Zero-knowledge.**

I. Introduction

The complexity class QMA is the quantum analogue of NP, the class of problems whose solutions can be verified in deterministic polynomial time. More precisely, in QMA, an all-powerful prover produces a quantum proof that is verified by a quantum polynomially-bounded verifier. Given the probabilistic nature of quantum computation, we require that for true statements, there exists a quantum proof that makes the verifier accept with high probability (this is called *completeness*), whereas all "proofs" for false statements are rejected with high probability (which is called *soundness*).

The class QMA was first defined by Kitaev [1], who also showed that deciding if a k-local Hamiltonian problem has low-energy states is QMA-complete. The importance of this result is two-fold: first, from a theoretical computer science perspective, it is the quantum analogue of the Cook-Levin theorem, since it establishes the first non-trivial QMA-complete problem. Secondly, it shows deep links between physics and complexity theory, since the k-local Hamiltonian problem is an important problem in many-body physics. Thus, a better understanding of QMA would lead to a better understanding of the power of quantum resources in proof verification, as we well as the role of *quantum entanglement* in low-energy states.

Follow-up work strengthened our understanding of this important complexity class, *e.g.*, by showing that QMA is contained in the complexity class PP [2][2]; that it is possible to reduce completeness and soundness errors without increasing the length of the witness [3]; understanding the difference between quantum and classical proofs [4], [5], [6]; the possibility of perfect completeness [7]; and, more recently, the relation of QMA with non-local games [8], [9], [10].

Also, much follow-up work focused on understanding the complete problems for QMA, mostly by improving the parameters of the QMA-hard Local Hamiltonian problem, or making it closer to models more physically relevant [11], [12], [13], [14], [15], [16], [17]. In 2014, a survey of QMA-complete languages [18] contained a list of 21 general problems that are known to be QMA-complete[3], and since then, the situation has not drastically changed. This contrasts with the development of NP, where only a few years after

[1]The full version of this work can be found in https://arxiv.org/abs/1911.07782.

[2]PP is the complexity class of decision problems that can be solved by probabilistic polynomial-time algorithms with error strictly smaller than $\frac{1}{2}$.

[3]We remark that these problems can be clustered as variations of a handful of base problems.

2575-8454/20/$31.00 ©2020 IEEE
DOI 10.1109/FOCS46700.2020.00027

the developments surrounding 3–SAT, Karp published a theory of reducibility, including a list of 21 NP-complete problems [19]; while 7 years later, a celebrated book by Garey and Johnson surveyed over 300 NP-complete problems [20].[4]

Recently, the role of QMA in quantum cryptography has also been explored. For instance, several results used ideas of the QMA-completeness of the Local Hamiltonian problem in order to perform verifiable delegation of quantum computation [21], [22], [23]. Furthermore, another line of work studies *zero-knowledge protocols* for QMA [24], [25], [26]; which is extremely relevant, given the fundamental importance in cryptography of zero-knowledge protocols for NP.

Despite the multiple advances in our understanding of QMA and related techniques, a number of fundamental open questions remain. In this work, we solve some of these open problems by showing: (i) QMA-hardness of Consistency of Local Density Matrix (CLDM) problem under Karp reductions; (ii) "commit-and-open" Zero-Knowledge (ZK) proof of quantum knowledge (PoQ) protocols for QMA; and (iii) a non-interactive zero-knowledge (NIZK) protocol in the secret parameter scenario. Our main technical contribution consists in showing that every problem in QMA admits a verification algorithm whose history state[5] is *locally simulatable*, meaning that the reduced density matrices on any small set of qubits is efficiently computable (without knowledge of the quantum witness). In order to be able to explain our results in more details and appreciate their contribution to a better understanding of QMA, we first give an overview of these areas and how they relate to these particular problems.

A. Background

In this section, we discuss the background on the topics that are relevant to this work, summarizing their current state-of-the-art.

Consistency of Local Density Matrices (CLDM).: The Consistency of Local Density matrices problem (CLDM) is as follows: given the classical description of local density matrices $\rho_1, ..., \rho_m$, each on a set of at most k qubits and for a global system of n qubits, is there a state τ that is consistent with such reduced states? Liu [12] showed that this problem is in QMA and that it is QMA-hard under Turing reductions, *i.e.*, a deterministic polynomial time algorithm with access to an oracle that solves CLDM in unit time can solve any problem in QMA.

We remark that this type of reduction is rather troublesome for QMA, since the class is not known (nor expected) to be closed under complement, *i.e.*, it is widely believed that QMA $\neq$ coQMA. If this is indeed the case, then Turing reductions do not allow a black-box generalization of results regarding the CLDM problem to all problems in QMA. This highlights the open problem of establishing the QMA-hardness of the CLDM problem under Karp reductions, *i.e.*, to show an efficient mapping between yes- and no-instances of any QMA problem to yes- and no-instances of CLDM, respectively.

Zero-Knowledge (ZK) Proofs for QMA*.:* In an *interactive* proof, a limited party, the *verifier*, receives the help of some untrusted powerful party, the *prover*, in order to decide if some statement is true. This is a generalization of a *proof*, where we allow multiple rounds of interaction. As usual, we require that the *completeness* and *soundness* properties hold. For cryptographic applications, the *zero-knowledge (ZK)* property is often desirable: here, we require that the verifier learn nothing from the interaction with the prover. This property is formalized by showing the existence of an efficient *simulator*, which is able to reproduce (*i.e., simulate*) the output of any given verifier on a *yes* instance (without having direct access to the actual prover or witness)[6].

As paradoxical as it sounds, statistical zero-knowledge interactive proofs are known to be possible for a host of languages, including the Quadratic Non-Residuosity, Graph Isomorphism, and Graph Non-Isomorphism problems [27], [28]; furthermore, all languages that can be proven by multiple provers (MIP) admits perfect zero-knowledge MIPs [29]. What is more, by introducing computational assumptions, it was shown that all languages that admit an interactive proof system also admit a zero-knowledge interactive proof system [30]. Zero-knowledge interactive proof systems have had a profound impact in multiple areas, including cryptography [31] and complexity theory [32].

We now briefly review the zero-knowledge interactive proof system for the NP-complete problem of Graph 3-colouring (3-COL). This is a 3-message proof system, and has the additional property that, given a witness, the prover is efficient. As a first message, the prover *commits* to a *permutation* of the given 3-colouring (meaning that the prover randomly permutes the colours to obtain colouring c, and produces a list $(v_i, \mathsf{commit}(c(v_i)))$, using a cryptographic primitive commit which is a *commitment scheme*). In the second message, the verifier chooses uniformly at random an edge $\{v_i, v_j\}$ of the graph. The prover responds with the information that allows the verifier to open the commitments to the colouring of the vertices of this edge (and nothing more). The verifier *accepts* if and only if the revealed colours are different. It is easy to see that the protocol is complete and sound. For the zero-knowledge

[4]The first edition of Garey and Johnson [20] was published in 1979.

[5]See Equation (1).

[6]Different definitions of "reproduce" result in different definitions of zero-knowledge protocols. A protocol is *perfect zero-knowledge* if the distribution of the output of the simulator is exactly the same as the distribution of output of transcripts of the protocol. A protocol is *statistical zero-knowledge* if such distributions are statistically close. Finally, a protocol is *computational zero-knowledge* if no efficient algorithm can distinguish both distributions. The convention is that in the absence of such specification, we are considering the case of computational zero-knowledge.

property, the simulator consists in a process that *guesses* which edge will be requested by the verifier and commits to a colouring that satisfies the prover in case this guess is correct. If the guess is incorrect, the technique of *rewinding* allows the simulator to re-initialize the interaction until it is eventually successful. Protocols that follow the *commit-challenge-response* structure of this proof system are called Σ-*protocols*[7] and, due to their simplicity, they play a very important role, for instance in the celebrated Fiat-Shamir transformation [33].

In the cryptographic scenario, an important relaxation of zero-knowledge proof systems are zero-knowledge *argument* systems for NP. In this model, the prover is also bounded to polynomial-time computation, and, for positive instances, the prover is provided a witness to the NP instance. This model allows much more efficient protocols which enables it to be used in practice [34], [35], [36].

The foundations of zero-knowledge in the quantum world were established by Watrous, who showed a technique called *quantum rewinding* [37] which is used to show the security of some classical zero-knowledge proofs (including the protocol for 3-COL described above), even against quantum adversaries. The importance of this technique is that quantum measurements typically *disturb* the measured state. When we consider quantum adversaries, such difficulties concern even *classical* proof systems, due to the rewinding technique that is ubiquitous (see example in the case of 3-COL above). Indeed, in the quantum setting, intermediate measurements (such as checking if the guess is correct) may compromise the success of future executions, since it is not possible *a priori* to "rewind" to a previous point in the execution in a black-box way.

Another dimension where quantum information poses new challenges is in the study of interactive proof systems for *quantum* languages. We point out that Liu [12] observed very early on that the CLDM problem should admit a simple zero-knowledge proof system following the "commit-and-open" approach, as in the 3-COL protocol. Inspired by this observation, recent progress has established the existence of zero-knowledge protocols for all of QMA [24], [25]. We note that although the proof system used there is reminiscent of a Σ-protocol, there are a number of reasons why it is not a "natural" quantum analogue of a Σ protocol. These include: (i) the use of a coin-flipping protocol, which makes the communication cost higher than 3 messages; (ii) the fact that the verifier's message is not a random challenge; and (iii) the final answer from the prover is not only the opening of some committed values.

Recently, Vidick and Zhang [26] showed how to make classical all of the interaction between the verifier and the prover in [24], [25], by considering *argument systems* instead of proof systems. In their protocol, they compose the result of Mahadev [22] for verifiable delegation of quantum computation by classical clients with the zero-knowledge protocol of [24], [25].

Zero-Knowledge Proofs of Knowledge (PoK).: In a zero-knowledge proof, the verifier becomes convinced of the *existence* of a witness, but this a priori has no bearing on the prover actually having in her possession such a witness. In some circumstances, it is important to guarantee that the prover actually has a witness. This is the realm of a *zero-knowledge proof of knowledge (PoK)* [28], [38].

We give an example to depict this subtlety. Let us consider the task of anonymous credentials [39]. In this setting, Alice wants to authenticate into some online service using her private credentials. In order to protect her credentials, she could engage in a zero-knowledge proof; this, however would be unsatisfactory, since the verifier in this scenario would be become convinced of the *existence* of accepting credentials, which does not necessarily translate to Alice actually being in the *possession* of these credentials. To remedy this situation, the PoK property establishes an "if-and-only-if" situation: if the verifier accepts, then we can guarantee that the prover actually *knows* a witness. This notion is formally defined by requiring the existence of an *extractor*, which is polynomial-time process K that outputs a valid witness when given oracle access to some prover P^* that makes the verifier accept with high enough probability.

In the quantum case, there has been some positive results in terms of the security of classical proofs of knowledge for NP against quantum adversaries [40]. However, in the fully quantum case (that is, proofs of quantum knowledge for QMA), no scheme has been proposed. One of the possible reasons why no such proof of quantum knowledge protocols was proposed is the lack of a *simple* zero-knowledge proof for QMA.

Non-Interactive Zero-Knowledge Proofs (NIZK).: The interactive nature of zero knowledge proof systems (for instance, in Σ-protocols) means that in some situations they are not applicable since they require the parties to be simultaneously online. Therefore, another desired property of such proof systems is that they are *non-interactive*, which means the whole protocol consists in a single message from the prover to the verifier. *Non-interactive zero-knowledge proofs* (NIZK) is a fundamental construction in modern cryptography and has far-reaching applications, for instance to cryptocurrencies [34].

We note that NIZK is known to be impossible in the standard model [41], *i.e.*, without extra assumptions, and therefore NIZK has been considered in different models. In one of the models most relevant in cryptography, we assume a common reference string (CRS) [42], which can be seen as a trusted party sending a random string to both the prover and the verifier. In another model, the trusted party is allowed to send different (but correlated) messages to the prover and the verifier; this is called the secret parameter

[7]The Greek letter Σ visualizes the flow of the protocol.

setup [43]. Classically, this model has been shown to be very powerful, since even its *statistical* zero-knowledge version is equivalent to all of the problems in the complexity class AM (this is the class that contains problem that can be verified by public-coin polynomial-time verifiers). As mentioned in [43], this model encompasses another model for NIZK where the prover and the verifier perform an *offline* pre-processing phase (which is independent of the input) and then the prover provides the ZK proof [44]. This inclusion holds since the parties could perform secure multi-party computation to compute the trusted party's operations.

In the quantum case, very little is known on non-interactive zero-knowledge. Chailloux, Ciocan, Kerenidis and Vadhan studied this problem in a setup where the message provided by the trusted party can depend on the instance of the problem [45]. Recently, some results also showed that the Fiat-Shamir transformation for classical protocols is still safe in the quantum setting, in the quantum random oracle model [46], [47], [48]. One particular and intriguing open question is the possibility of NIZKs for QMA.

B. Results

As we have shown so far, the state-of-the-art in the study of QMA is that the body of knowledge is still developing, and that there are some specific goals that, if achieved, would help us better understand QMA and devise new protocols for quantum cryptography. Given this context, we present now our results in more detail.

Our first result is to show that the CLDM problem is QMA-hard under Karp reductions, solving the 14-year-old problem proposed by Liu [12].

Result 1: The CLDM problem is QMA-complete under Karp reductions.

We capture the techniques used in establishing the above into a new characterization of QMA that provides the best-of-both worlds in terms of two proof systems for QMA in an abstract way: we define SimQMA as the complexity class with proof systems that are (i) locally verifiable (as in the Local Hamiltonian problem), and (ii) every reduced density matrix of the witness can be efficiently computed (as in the CLDM problem). This results is the basis for our applications to quantum cryptography:

Result 2: SimQMA = QMA.

Next, we define a quantum notion of a classical Σ-protocol, which we call a Ξ-protocol[8] (please note, both a Σ and Ξ protocol is also referred to throughout as "commit-and-open" protocols.) Using our characterization given in Result 2, we show a QMA-complete language that admits a Ξ-protocol. Taking into account the importance of Σ protocols for zero-knowledge proofs, we are able to show a quantum analogue of the celebrated [27] paper:

Result 3: All problems in QMA admit a computational zero-knowledge Ξ-proof system.

We are also able to show that simple changes in the construction for the result above allow us to achieve the first *statistical* zero-knowledge *argument* system for QMA.

Result 4: All problems in QMA admit a statistical zero-knowledge Ξ-argument system.

Then we provide the definition of Proof of Quantum Knowledge (PoQ).[9] In short, we say that a proof system is a PoQ if there exists a quantum polynomial-time *extractor* K that has oracle access to a quantum prover which makes the verifier accept with high enough probability, and the extractor is able to output a sufficiently good witness for a "QMA-relation". We note that this definition for a PoQ is not a straightforward adaptation of the classical definition; this is because NP has many properties such as perfect completeness, perfect soundness and even that proofs can be copied, that are not expected to hold in the QMA case. More details are given in the full version of the paper [50]. We are then able to show that our Ξ protocols for QMA described in Results 3 and 4 are both PoQs. This is the first proof of knowledge for QMA.[10]

Result 5: All problems in QMA admit a zero-knowledge proof of quantum knowledge proof system and a statistical zero-knowledge proof of quantum knowledge argument system.

We remark that using techniques for post-hoc delegation of quantum computation [21], our PoQ for QMA may be understood as a *proof-of-work* for quantum computations, since it could be used to convince a verifier that the prover has indeed created the *history state* of some pre-defined computation. This is very relevant in the scenario of testing small-scale quantum computers in the most adversarial model possible: the zero-knowledge property ensures that the verifier learns nothing but the truth of the statement, while the PoQ property means that the prover has indeed prepared a ground state with the given properties. Comparatively, all currently known protocols either make assumptions on the devices, or certify only the answer of the computation, but not the knowledge of the prover.

Finally, using the techniques of Result 3, we show that every problem in QMA has a non-interactive *statistical* zero-knowledge proof in the secret parameter model. We are even able to strengthen our result to the complexity class QAM (recall that in a QAM proof system, the verifier first sends a random string to the prover, who answers with a quantum proof). Note that QAM trivially contains QMA.

Result 6: All problems in QAM have a non-interactive statistical zero-knowledge protocol in the secret parameter

[8]Besides being an excellent symbolic reminder of the interaction in a 3-message proof system, Ξ is chosen as a shorthand for what we might otherwise call a $q\Sigma$ protocol, due to the resemblance with the pronunciation as "**csi**gma".

[9]This definition is joint work with Coladangelo, Vidick and Zhang [49].

[10]See also independent and concurrent work by Coladangelo, Vidick and Zhang [49].

model.

Note that, as in the classical case [43], our result also implies a QNIZK protocol where the prover and the verifier run an offline (classical) pre-processing phase (independent of the witness) and then the prover sends the quantum ZK proof to the verifier. We note also that even though these models are less relevant to the cryptographic applications of NIZK, we think that our result moves us towards a QNIZK protocol for QMA in a more standard model.

C. Techniques

The starting point for our results are *locally simulatable codes*, as defined in [51]. We give now a rough intuition on the properties of such codes and leave the details to the full version of the paper [50].

First, a quantum error correcting code is s-*simulatable* if there exists an efficient classical algorithm that outputs the reduced density matrices of codewords on every subset of at most s qubits. Importantly, this algorithm is oblivious of the logical state that is encoded. We note that it was already known that the reduced density matrices of codewords hide the encoded information, since quantum error correcting codes can be used in secret sharing protocols [52], and in [51] they show that there exist codes such that the classical description of the reduced density matrices of the codewords can be efficiently computed. Next, [51] extends the notion of simulatability of *logical operations* on encoded data as follows. Recalling the theory of fault-tolerant quantum computation, according to which some quantum error-correcting codes allow computations over *encoded* data by using "transversal" gates and encoded magic states. The definition of s-simulatability is extended to require that the simulator also efficiently computes the reduced density matrix on at most s qubits of intermediate steps of the *physical* operations that implement a logical gate on the encoded data (again, by transversal gates and magic states).

Example 1: Let us suppose that the encoding map Enc admits transversal application of the one-qubit gate G, i.e., $G^{\otimes N}\mathsf{Enc}(|\psi\rangle) = \mathsf{Enc}(G|\psi\rangle)$. The simulatability property requires that the density matrices on at most s qubits of $(G^{\otimes t} \otimes I^{\otimes(N-t)})\mathsf{Enc}(|\psi\rangle)$ should be efficiently computed, for every $0 \leq t \leq N$.

In [51], the authors show that the concatenated Steane code is a locally simulatable code. With this tool, in [51], it is shown that every MIP* protocol[11] can be made zero-knowledge, thus quantizing the celebrated result of [29]. Here, we provide an alternative proof for the simulatability of concatenated Steane codes. Our new proof is much simpler than the proof provided in [51], but it holds for a slightly weaker statement (but which is already sufficient to derive the results in [51]). Then, for the first time, we apply the techniques of simulatable codes from [51] to QMA, which enables us to solve many open problems as previously described.

In order to explain our approach to achieving Result 1, we first recall the quantum Cook-Levin theorem proved by Kitaev [1]. In his proof, Kitaev uses the circuit-to-Hamiltonian construction [53], mapping an arbitrary QMA verification circuit $V = U_T ... U_1$ to a local Hamiltonian H_V that enforces that low energy states are *history states* of the computation, *i.e.*, a *superposition* of the snapshots of V for every timestep $0 \leq t \leq T$:

$$|\Phi\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0...T+1} |t\rangle \otimes U_t \dots U_1 |\psi_{init}\rangle. \quad (1)$$

In the above, the first register is called the *clock* register, and it encodes the timestep of the computation, while the second register contains the snapshot of the computation at time t, *i.e.*, the quantum gates $U_1, ..., U_t$ applied to the initial state $|\psi_{init}\rangle = |\phi\rangle|0\rangle^{\otimes A}$, that consists of the quantum witness and auxiliary qubits. The Hamiltonian H_V also guarantees that $|\psi_{init}\rangle$ has the correct form at $t = 0$, and that the final step *accepts*, *i.e.*, the output qubit is close to $|1\rangle$.

In [51], they note that an important obstacle to making a state similar to $|\Phi\rangle$[12] locally simulatable is its dependence on the witness state $|\phi\rangle$. The solution is to consider a different verification algorithm V' that implements V on *encoded data*, much like in the theory of fault-tolerant quantum computing. In more details, for a fixed locally simulatable code, V' expects the encoding of the original witness $\mathsf{Enc}(|\phi\rangle)$ and then, with her raw auxiliary states, she creates encodings of auxiliary states $\mathsf{Enc}(|0\rangle)$ and magic states $\mathsf{Enc}(|\mathsf{MS}\rangle)$, and then performs the computation V through transversal gates and magic state gadgets, and finally decodes the output qubit. This gives rise to a new history state:

$$|\Phi'\rangle = \frac{1}{\sqrt{T'+1}} \sum_{t=0...T'+1} |t\rangle \otimes U'_t \dots U'_1 |\psi'_{init}\rangle, \quad (2)$$

where $|\psi'_{init}\rangle = \mathsf{Enc}(|\phi\rangle)|0\rangle^{\otimes A'}$ and $U'_1, ..., U_{T'}$ are the gates of V' described above. Using the techniques from [51],[13] we can show that from the properties of the locally simulatable codes, the reduced density matrix on every set of 5 qubits of $|\Phi'\rangle$ can be efficiently computed. In this work, we prove that these reduced density matrices are in fact QMA-hard instances of CLDM. More concretely, we show that these reduced density matrices of a hypothetical history state of an accepting QMA-verification can always

[11] MIP* is the set of languages that admit a classical *multi-prover* interactive proof, where, in addition, the provers share entanglement

[12] In [51], they are simulating history states for MIP* computation and therefore they need to deal also with arbitrary Provers' operations.

[13] We remark that we also need to fix a small bug in their proof. The bug fix deals with technicalities regarding V' and $|\psi'\rangle$ that are beyond the scope of this overview. See the full version of the paper [50] for more details.

be computed, and there exists a global state (namely the history state) consistent with these reduced density matrices if and only if the original QMA verification accepts with overwhelming probability (and therefore we are in the case of a yes-instance).

Result 1 opens up a number of possible applications to cryptographic settings. However, we face a tradeoff: in CLDM, we have the description of the local density matrices, which would appear to yield a Ξ protocol which would be zero-knowledge. However, the QMA verification for such problem is rather complicated: we need multiple copies of the global state to perform tomography on the reduced states, instead of a single copy that is needed in the Local Hamiltonian problem.

In order to combine these two desired properties in a single object, we describe a powerful technique that we call *locally simulatable proofs*. In a locally simulatable proof system for some problem $A = (A_{yes}, A_{no})$, we require that: (i) the verification test performed by the verifier acts on at most k out of the n qubits of the proof, and (ii) for every $x \in A_{yes}$, there exists a locally simulatable witness $|\psi\rangle$, *i.e.*, a state $|\psi\rangle$ that passes all the local tests and such that for every $S \subseteq [n]$ with $|S| \leq k$, it is possible to compute the reduced state of the $|\psi\rangle$ on S efficiently (without the help of the prover). Notice that we have no extra restrictions on $x \in A_{no}$, since any quantum witness should make this verifier reject with high probability.

We then show that all problems in QMA admit a locally simulatable proof system. In order to achieve this, we use the local tests on the encoded version of the QMA verification algorithm that come from the Local Hamiltonian problem, together with the fact that the history state of such computation is a low-energy state and is simulatable (which is used to establish the QMA-hardness of CLDM).

We remark that a direct classical version of locally simulatable proofs as we define them is impossible. This is because, given the local values of a classical proof, it is always possible to reconstruct the full proof by gluing these pieces together. The fact that this operation is hard to perform quantumly is intrinsically related to entanglement: given the local density matrices, it is not a priori possible to know which parts are entangled in order to glue them together. As discussed in the next section, this allows us to achieve a type of simple zero-knowledge protocol that defies all classical intuition.

1) Locally Simulatable Proofs in Action: We now sketch how each of Result 3–Result 5 is obtained via the lens of locally simulatable proofs.

Zero Knowledge.: We use the characterization QMA = SimQMA to give a new zero-knowledge proof system for QMA. Our protocol is much simpler than previous results [24], [25], and it follows the "commit-challenge-response" structure of a Σ-protocol. Since our commitment is a quantum state (the challenge and response are classical), we call this type of protocol a "Ξ-*protocol*".

The main idea is to use the quantum one-time pad to split the first message in the protocol into a quantum and a classical part. More concretely, the prover sends $X^a Z^b |\psi\rangle$ and commitments to each bit of a and b to the verifier, where $|\psi\rangle$ is a locally simulatable quantum witness for some instance x and a and b are uniformly random strings. The verifier picks some $c \in [m]$, which corresponds to one of the tests of the simulatable proof system, and asks the prover to open the commitment of the encryption keys to the corresponding qubits. The honest prover opens the commitment corresponding to the one-time pad keys of the qubits involved in test c. The verifier then checks if: (i) the openings are correct and, (ii) the decrypted reduced state passes test c.

Assuming the existence of unconditionally binding and computationally hiding commitment schemes, we show that our protocol is a computational zero-knowledge proof system for QMA. Completeness and soundness follow trivially, whereas the zero-knowledge property is established by constructing a simulator that exploits the properties of the locally simulatable proof system and the rewinding technique of Watrous [37].

We also show that if the commitment scheme is computationally binding and unconditionally hiding, then our protocol is the first *statistical* zero-knowledge *argument* for QMA. This is because the protocol is secure only against polynomial-time malicious provers, since unbounded provers would be able to open the commitment to different values. On the other hand, we achieve statistical zero-knowledge since the commitments that are never opened effectively encrypt the witness in an information-theoretic sense.

To the best of our knowledge, this is the first time that quantum techniques are used in zero-knowledge to achieve a commit-and-open protocol that *requires no randomization of the witness*. Indeed, for reasons already discussed, all classical zero-knowledge Σ protocols require a *mapping* or *randomization* of the witness (*e.g.* in the $3-\mathsf{COL}$ protocol, this is the permutation that is applied to the colouring before the commitment is made). We thus conclude that quantum information enables a new level of encryption that is not possible classically: the "juicy" information is present in the global state, whose local parts are *fully* known [51].

Proof of Quantum Knowledge for QMA*.:* As discussed in Section I-B, our first challenge here is to define a Proof of Quantum Knowledge (PoQ). We recall that in the classical setting, we require an extractor that outputs some witness that passes the NP verification with probability 1, whenever the verifier accepts with probability greater than some parameter κ, known as the knowledge error.

In the quantum case, given: (i) that we are not able to clone quantum states and (ii) QMA is not known to be closed under perfect completeness, the best that we can hope

for is to extract some quantum state that would pass the QMA verification with some probability to be related to the acceptance probability in the interactive protocol, whenever this latter value is above some threshold κ.

To define a PoQ, we first fix the verification algorithm V_x for some instance of a problem in QMA. We also assume P^* to be a prover that makes the verifier accept with probability at least $\epsilon > \kappa$ in the Ξ protocol.[14] We assume that P^* only performs unitary operations on a private and message registers. We then define a quantum polynomial-time algorithm K that has oracle access to P^*, meaning that K can execute the unitary operations of P^*, their inverse operations and has access to the message register of P^*.[15] The protocol is said to be a Proof of Quantum Knowledge if K outputs, with non-negligible probability, some quantum state ρ that would make V_x accept with probability at least $q(\epsilon, n)$, where q is known as the quality function, or aborts otherwise.

The difficulty in showing that our Ξ protocols are PoQs lies in the fact that any measurement performed by the extractor disturbs the state held by P^*, and therefore when we rewind P^* by applying the inverse of his operation, we do not come back to the original state. We overcome this difficulty in the following way. We set κ to be some value very close to 1, namely $\kappa = 1 - \frac{1}{p(n)}$ for some large enough polynomial p. Our extractor starts by simulating P^* on the first message of the Ξ protocol, and then holds the (supposed) one-time-padded state and the commitments to the one-time-pad keys. K follows by iterating over all possible challenges of the Ξ protocol, runs P^* on this challenge, perform the verifier's check and then rewinds P^*. By the assumption that P^* has a very high acceptance probability, the measurements performed by K do not disturb the state too much, and in this case, K can retrieve the correct one-time pads for every qubit of the witness. If K is successful (meaning that k is able to open every committed bit), then K can decode the original one-time-padded state and it is a good witness for V_x with high probability.

We then analyse the sequential repetition of the protocol, that allows us to have a PoQ with *exponentially small* knowledge error κ, and extracts one good witness from P^* (out of the polynomially many copies that P^* should have in order to cause the verifier to accepted in the multiple runs of the protocol).

Non-Interactive zero knowledge proof for QMA in the secret parameter model.: Finally, we achieve our non-interactive statistical zero-knowledge protocol for QMA in the secret parameter setting using techniques similar to our Ξ protocol: the trusted party chooses the one-time pad key and a random (and small) subset of these values that are reported to the verifier. Since the prover does not know which are the values that were given to the verifier, he should act as in the Ξ-protocol, but now the verifier does not actually need to ask for the openings, since the trusted dealer has already sent them. Although this is a less natural model, we hope that this result will shed some light in developing QNIZK proofs for QMA in more commonly-used models.

[14]Note that we reserve the word "verifier" here for the Ξ protocol and refer to V_x as the QMA verification algorithm.

[15]This model is already considered by [40] in his work of quantum proofs of knowledge for NP.

D. Open problems

Further QMA*-complete languages.:* We note that a number of problems are currently known to be QMA-complete under Turing reductions, including the N-representability [54] and bosonic N-representability problems [55] as well as the universal functional of density function theory (DFT) [56]. It is an open question if these problems can be shown to be QMA-complete under Karp reductions using the techniques presented in our work.

Complexity of k CLDM for $k < 5$.: We prove in this work that 5-CLDM is QMA-hard under Karp reductions. We leave as an open problem proving if the problem is still QMA-complete for $k < 5$.

Marginal reconstruction problem.: We remark that the classical version of CLDM is defined as follows: given the description of m marginal distributions on sets of bits $C_1, \ldots, C_m$, such that $|C_i| \leq k$, decide if there is a probability distribution that is close to those marginals, or such a distribution does not exist. This problem was proven NP-complete by Pitowsky [57], and its containment in NP is proved by using the fact that such distribution can be seen as a point p in the *correlation polytope* in a polynomial-size Hilbert space. In this case, by Caratheodory's theorem, p is a convex combination of polynomially many vertices of such polytope, and therefore these vertices serve as the NP-proof and a linear program verifies if there is a convex combination of them that is consistent with the marginals of the problem's instance.

The difference here is that the proof and the marginals are different (but connected) objects. We leave as an open problem if we can extract a notion of a locally simulatable classical proof from this (or any other) problem, and its applications to cryptography and complexity theory. In particular, we wonder if there is a natural zero-knowledge protocol for this problem.

Applications of quantum ZK protocols.: In classical cryptography, ZK and PoK protocols are a fundamental primitive since they are crucial ingredients in a plethora of applications. We discussed in Section I-B that our quantum ZK PoQ for QMA could be used as a proof-of-work for quantum computations. An interesting open problem is finding other settings in which the benefits of our simple ZK protocols for QMA can be applied. We list now some possibilities that could be explored in future work: authentication with uncloneable credentials [58]; proof of quantum ownership [59]; or ZK PoQ verification for quantum money [60].

Practical ZK protocols for QMA.: Even if we reach a conceptually much simpler ZK protocol for QMA, the resources needed for it are still very far from practical. We leave as an open problem if one could devise other protocols that are more feasible from a physical implementation viewpoint, which could include classical communication protocols based on the protocols proposed by Vidick and Zhang [26], or device-independent ones based on the ideas of Grilo [23].

Non-interactive Zero-knowledge protocols for QMA in the CRS model.: In this work, we propose a QNIZK protocol where the information provided by the trusted dealer is asymmetric. We leave as an open problem if one could devise a protocol where the dealer distributes a common reference string (CRS)(or shared EPR pairs) to the prover and the verifier.

A possible way of achieving such non-interactive protocol would be to explore the properties of Ξ-protocols, as done classically with Σ-protocols. For instance, the well-known Fiat-Shamir transformation [33] allows us to make Σ-protocols non-interactive (in the Random Oracle model). We wonder if there is a version of this theorem when the first message can be quantum.

Witness indistinguishable/hiding protocols for QMA*.:* Classically, there are two weaker notions that can substitute for ZK in different applications. In Witness Indistinguishable (WI) proofs, we require that the verifier cannot distinguish if she is interacting with a prover holding a witness w_1 or w_2, for any $w_1 \neq w_2$. In Witness Hiding (WH), we require that the verifier is not able to cook-up a witness for the input herself. We note that zero-knowledge implies both such definitions, and we leave as an open problem finding WI/WH protocols for QMA with more desirable properties than the known ZK protocols.

Computational Zero-Knowledge proofs vs. Statistical Zero-Knowledge arguments.: In this work, we show that QMA admits quantum ZK proofs and statistical ZK arguments. We note that classically, it is known that the class of problems with computational ZK proofs is closely related to the class of problems with statistical ZK arguments [61]. We wonder if this relation is also true in the quantum setting.

E. Concurrent and subsequent works

Concurrently to this work, Bitansky and Shmueli [62] proposed the first quantum zero-knowledge argument system for QMA with constant rounds and negligible soundness. Their main building block is a *non black-box quantum extractor* for a post-quantum commitment scheme. Also concurrently to this work, Coladangelo, Vidick and Zhang [49] proposed a non-interactive argument of quantum knowledge for QMA with a quantum setup phase (that is independent of the witness) and a classical online phase.

Finally, Alagic, Childs, Grilo and Hung [63] proposed the first non-interactive zero-knowledge argument system for QMA where the communication is purely *classical*. Their protocol works in the random oracle model with setup.

Acknowledgements

We thank Dorit Aharonov, Thomas Vidick, and the anonymous reviewers for help in improving the presentation of this work. We thank Andrea Coladangelo, Thomas Vidick and Tina Zhang for discussions on the definition of proofs of quantum knowledge. A.G. thanks Christian Majenz for discussions on a suitable title for this work. A.B. is supported by the U.S. Air Force Office of Scientific Research under award number FA9550-17-1-0083, Canada's NFRF and NSERC, an Ontario ERA, and the University of Ottawa's Research Chairs program.

REFERENCES

[1] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*. American Mathematical Society, 2002.

[2] A. Kitaev and J. Watrous, "Parallelization, amplification, and exponential time simulation of quantum interactive proof systems," in *32nd Annual ACM Symposium on Theory of Computing—STOC 2000*, 2000, pp. 608–617.

[3] C. Marriott and J. Watrous, "Quantum Arthur–Merlin games," *Computational Complexity*, vol. 14, no. 2, pp. 122–152, 2005.

[4] S. Aaronson and G. Kuperberg, "Quantum versus classical proofs and advice," in *22nd Annual Conference on Computational Complexity—CCC 2007*, 2007, pp. 115–128.

[5] A. B. Grilo, I. Kerenidis, and J. Sikora, "QMA with subset state witnesses," *Chicago Journal of Theoretical Computer Science*, vol. 2016, no. 4, 2016.

[6] B. Fefferman and S. Kimmel, "Quantum vs. classical proofs and subset verification," in *43rd International Symposium on Mathematical Foundations of Computer Science—MFCS 2018*, 2018, p. 22.

[7] S. Aaronson, "On perfect completeness for QMA," *Quantum Information & Computation*, vol. 9, no. 1, pp. 81–89, 2009.

[8] A. Natarajan and T. Vidick, "A quantum linearity test for robustly verifying entanglement," in *49th Annual ACM Symposium on Theory of Computing—STOC 2017*, 2017, pp. 1003–1015.

[9] ——, "Low-degree testing for quantum states, and a quantum entangled games PCP for QMA," in *59th Annual Symposium on Foundations of Computer Science—FOCS 2018*, 2018, pp. 731–742.

[10] A. Coladangelo, A. B. Grilo, S. Jeffery, and T. Vidick, "Verifier-on-a-leash: New schemes for verifiable delegated quantum computation, with quasilinear resources," in *Advances in Cryptology—EUROCRYPT 2019*, vol. III, 2019, pp. 247–277.

[11] J. Kempe and O. Regev, "3-local Hamiltonian is QMA-complete," *Quantum Information & Computation*, vol. 3, no. 3, pp. 258–264, 2003.

[12] Y. Liu, "Consistency of local density matrices is QMA-complete," in *9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems—APPROX 2006 and 10th International Workshop on Randomization and Computation—RANDOM 2006*, 2006, pp. 438–449.

[13] J. Kempe, A. Kitaev, and O. Regev, "The complexity of the local Hamiltonian problem," *SIAM Journal on Computing*, vol. 35, no. 5, pp. 1070–1097, 2006.

[14] R. Oliveira and B. M. Terhal, "The complexity of quantum spin systems on a two-dimensional square lattice," *Quantum Information & Computation*, vol. 8, no. 10, pp. 900–924, 2008.

[15] T. S. Cubitt and A. Montanaro, "Complexity classification of local Hamiltonian problems," in *55th Annual Symposium on Foundations of Computer Science—FOCS 2014*, 2014, pp. 120–129.

[16] S. Hallgren, D. Nagaj, and S. Narayanaswami, "The local Hamiltonian problem on a line with eight states is QMA-complete," *Quantum Information & Computation*, vol. 13, no. 9&10, pp. 721–750, 2013.

[17] J. Bausch and E. Crosson, "Analysis and limitations of modified circuit-to-Hamiltonian constructions," *Quantum*, vol. 2, p. 94, 2018.

[18] A. D. Bookatz, "QMA-complete problems," *Quantum Information & Computation*, vol. 14, no. 5&6, pp. 361–383, 2014.

[19] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of Computer Computations*, 1972, pp. 85–103.

[20] M. R. Garey and D. S. Johnson, *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1990.

[21] J. F. Fitzsimons, M. Hajdušek, and T. Morimae, "Post hoc verification of quantum computation," *Physical Review Letters*, vol. 120, no. 4, p. 040501, 2018.

[22] U. Mahadev, "Classical verification of quantum computations," in *59th Annual Symposium on Foundations of Computer Science—FOCS 2018*, 2018, pp. 259–267.

[23] A. B. Grilo, "A simple protocol for verifiable delegation of quantum computation in one round," in *46th International Colloquium on Automata, Languages, and Programming—ICALP 2019*, 2019, p. 28.

[24] A. Broadbent, Z. Ji, F. Song, and J. Watrous, "Zero-knowledge proof systems for QMA," in *57th Annual Symposium on Foundations of Computer Science—FOCS 2016*, 2016, pp. 31–40.

[25] ——, "Zero-knowledge proof systems for QMA," *SIAM Journal on Computing*, vol. 49, no. 2, pp. 245–283, 2020.

[26] T. Vidick and T. Zhang, "Classical zero-knowledge arguments for quantum computations," 2019, available at https://arxiv.org/abs/1902.05217.

[27] O. Goldreich, S. Micali, and A. Wigderson, "Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems," *Journal of the ACM*, vol. 38, no. 3, pp. 690–728, 1991.

[28] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186–208, 1989.

[29] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Widgerson, "Multi-prover interactive proofs: how to remove intractability assumptions," in *20th Annual ACM Symposium on Theory of Computing—STOC 1988*, 1988, pp. 113–131.

[30] M. Ben-Or, O. Goldreich, S. Goldwasser, J. Håstad, J. Kilian, S. Micali, and P. Rogaway, "Everything provable is provable in zero-knowledge," in *Advances in Cryptology—CRYPTO 1988*, 1988, pp. 37–56.

[31] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game," in *19th Annual ACM Symposium on Theory of Computing—STOC 1987*, 1987, pp. 218–229.

[32] S. Vadhan, "The complexity of zero knowledge," in *27th International Conference on Foundation of Software Technology and Theoretical Computer Science—FSTTCS 2007*, 2007, pp. 52–70.

[33] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *Advances in Cryptology—CRYPTO 1986*, 1987, pp. 186–194.

[34] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from Bitcoin," in *2014 IEEE Symposium on Security and Privacy*, 2014, pp. 459–474.

[35] B. Parno, J. Howell, C. Gentry, and M. Raykova, "Pinocchio: nearly practical verifiable computation," *Communications of the ACM*, vol. 59, no. 2, pp. 103–112, 2016.

[36] E. Ben-Sasson, A. Chiesa, M. Riabzev, N. Spooner, M. Virza, and N. P. Ward, "Aurora: Transparent succinct arguments for R1CS," in *Advances in Cryptology—EUROCRYPT 2019*, vol. I, 2019, pp. 103–128.

[37] J. Watrous, "Zero-knowledge against quantum attacks," *SIAM Journal on Computing*, vol. 39, no. 1, pp. 25–58, 2009.

[38] M. Bellare and O. Goldreich, "On defining proofs of knowledge," in *Advances in Cryptology—CRYPTO 1992*, 1993, pp. 390–420.

[39] D. Chaum, "Blind signatures for untraceable payments," in *Advances in Cryptology—CRYPTO 1982*, 1983, pp. 199–203.

[40] D. Unruh, "Quantum proofs of knowledge," in *Advances in Cryptology—EUROCRYPT 2012*, 2012, pp. 135–152.

[41] O. Goldreich and Y. Oren, "Definitions and properties of zero-knowledge proof systems," *Journal of Cryptology*, vol. 7, no. 1, pp. 1–32, 1994.

[42] M. Blum, P. Feldman, and S. Micali, "Non-interactive zero-knowledge and its applications," in *20th Annual ACM Symposium on Theory of Computing—STOC 1988*, 1988, pp. 103–112.

[43] R. Pass and A. Shelat, "Unconditional characterizations of non-interactive zero-knowledge," in *Advances in Cryptology—CRYPTO 2005*, 2005, pp. 118–134.

[44] J. Kilian, S. Micali, and R. Ostrovsky, "Minimum resource zero knowledge proofs," in *30th Annual Symposium on Foundations of Computer Science—FOCS 1989*, 1989, pp. 474–479.

[45] A. Chailloux, D. F. Ciocan, I. Kerenidis, and S. Vadhan, "Interactive and noninteractive zero knowledge are equivalent in the help model," in *5th Theory of Cryptography Conference—TCC 2008*, 2008, pp. 501–534.

[46] Q. Liu and M. Zhandry, "Revisiting post-quantum Fiat-Shamir," in *Advances in Cryptology—CRYPTO 2019*, vol. II, 2019, pp. 326–355.

[47] J. Don, S. Fehr, C. M. Majenz, and C. Schaffner, "Security of the Fiat-Shamir transformation in the quantum random-oracle model," in *Advances in Cryptology—CRYPTO 2019*, vol. II, 2019, pp. 356–383.

[48] A. Chailloux, "Quantum security of the Fiat-Shamir transform of commit and open protocols," 2019, available at https://eprint.iacr.org/2019/699.

[49] A. Coladangelo, T. Vidick, and T. Zhang, "Non-interactive zero-knowledge arguments for QMA, with preprocessing," 2019, available at https://arxiv.org/abs/1911.07546.

[50] A. Broadbent and A. B. Grilo, "QMA-hardness of Consistency of Local Density Matrices
with Applications to Quantum Zero-Knowledge," 2019, available at https://arxiv.org/abs/1911.07782.

[51] A. B. Grilo, W. Slofstra, and H. Yuen, "Perfect zero knowledge for quantum multiprover interactive proofs," in *60th Annual Symposium on Foundations of Computer Science—FOCS 2019*, 2019, pp. 611–635.

[52] R. Cleve, D. Gottesman, and H.-K. Lo, "How to share a quantum secret," *Physical Review Letters*, vol. 83, no. 3, pp. 648–651, 1999.

[53] R. P. Feynman, "Simulating physics with computers," *International Journal of Theoretical Physics*, vol. 21, no. 6, pp. 467–488, 1982.

[54] Y.-K. Liu, M. Christandl, and F. Verstraete, "Quantum computational complexity of the N-representability problem: QMA complete," *Physical Review Letters*, vol. 98, no. 11, p. 110503, 2007.

[55] T.-C. Wei, M. Mosca, and A. Nayak, "Interacting boson problems can be QMA hard," *Physical Review Letters*, vol. 104, no. 4, p. 040501, 2010.

[56] N. Schuch and F. Verstraete, "Computational complexity of interacting electrons and fundamental limitations of density functional theory," *Nature Physics*, vol. 5, no. 10, p. 732, 2009.

[57] I. Pitowsky, "Correlation polytopes: their geometry and complexity," *Mathematical Programming*, vol. 50, no. 1–3, pp. 395–414, 1991.

[58] R. C. Cramer, I. Damgård, and B. Schoenmakers, "Proofs of partial knowledge and simplified design of witness hiding protocols," in *Advances in Cryptology—CRYPTO 1994*, 1994, pp. 174–187.

[59] C. Badertscher, D. Jost, and U. Maurer, "Agree-and-prove: Generalized proofs of knowledge and applications," 2019, available at https://eprint.iacr.org/2019/662.

[60] S. Aaronson and P. Christiano, "Quantum money from hidden subspaces," in *44th Annual ACM Symposium on Theory of Computing—STOC 2012*, 2012, pp. 41–60.

[61] S. J. Ong and S. Vadhan, "Zero knowledge and soundness are symmetric," in *Advances in Cryptology—EUROCRYPT 2007*, 2007, pp. 187–209.

[62] N. Bitansky and O. Shmueli, "Post-quantum zero knowledge in constant rounds," 2019, available at https://eprint.iacr.org/2019/1279.

[63] G. Alagic, A. M. Childs, A. B. Grilo, and S.-H. Hung, "Non-interactive classical verification of quantum computation," 2020, available at https://arxiv.org/abs/1911.08101.