Quantum Delegation with an Off-The-Shelf Device

Anne Broadbent
Department of Mathematics and Statistics, University of Ottawa, Canada
Nexus for Quantum Technologies, University of Ottawa, Canada

Arthur Mehta
Department of Mathematics and Statistics, University of Ottawa, Canada
Nexus for Quantum Technologies, University of Ottawa, Canada

Yuming Zhao
Institute for Quantum Computing, University of Waterloo, Canada
Department of Pure Mathematics, University of Waterloo, Canada

Abstract

Given that reliable cloud quantum computers are becoming closer to reality, the concept of delegation of quantum computations and its verifiability is of central interest. Many models have been proposed, each with specific strengths and weaknesses. Here, we put forth a new model where the client trusts only its classical processing, makes no computational assumptions, and interacts with a quantum server in a *single* round. In addition, during a set-up phase, the client specifies the size n of the computation and receives an untrusted, *off-the-shelf (OTS)* quantum device that is used to report the outcome of a single measurement.

We show how to delegate polynomial-time quantum computations in the OTS model. This also yields an interactive proof system for all of QMA, which, furthermore, we show can be accomplished in statistical zero-knowledge. This provides the first relativistic (one-round), two-prover zero-knowledge proof system for QMA.

As a proof approach, we provide a new self-test for n EPR pairs using only constant-sized Pauli measurements, and show how it provides a new avenue for the use of simulatable codes for local Hamiltonian verification. Along the way, we also provide an enhanced version of a well-known stability result due to Gowers and Hatami and show how it completes a common argument used in self-testing.

2012 ACM Subject Classification Theory of computation → Interactive proof systems; Theory of computation → Quantum complexity theory

Keywords and phrases Delegated quantum computation, zero-knowledge proofs, device-independence

Digital Object Identifier 10.4230/LIPIcs.TQC.2024.12

Related Version *Full Version*: `https://arxiv.org/abs/2304.03448` [9]

Funding This work was supported by the Mitacs Accelerate program IT24833 in collaboration with industry partner Agnostiq, online at `https://agnostiq.ai`.
Anne Broadbent: A.B. is supported by the Air Force Office of Scientific Research under award number FA9550-20-1-0375, NSERC, and the University of Ottawa's Research Chairs program.

Acknowledgements We thank Seyed Sajjad Nezhadi, Gregory Rosenthal, William Slofstra, Jalex Stark, and Henry Yuen for discussions on the model. We thank Alex Grilo for discussions on some of our proof techniques. We also thank Thomas Vidick for detailed discussions on the formulation and proof of Theorem 19. We thank the anonymous reviewers for suggesting the use case of an OTS for a classical proof system.

1 Introduction

In an interactive proof system, a computationally-bounded verifier interacts with a powerful prover in order to verify the truthfulness of an agreed-upon problem instance. Starting with QMA, and followed by QIP and QMIP (among others), *quantum* interactive proof system, (in which the verifier is *quantum* polynomial-time) were defined and studied [48, 49, 30].

Yet, these quantizations depend crucially on the tacit assumption that the verifier has access to *trusted* quantum polynomial-time verification. Given the current state-of-the-art in quantum computation development, the inherent difficulty at characterizing quantum systems, and the fact that there is no way to reliably verify the trace of a quantum computation, there is ample evidence that this assumption may be questionable. Indeed, despite impressive technological improvements, we may ultimately have to contend with a reality where quantum computers are never as trustworthy or reliable as classical devices. This prospect has motivated consideration of models where the verifier has access to very limited but trusted quantum functionality [1, 4, 18], or where the verifier is entirely classical and the prover is computationally bounded [31], while another class called MIP^* models an efficient classical verifier interacting with several isolated, unbounded quantum provers [14]. Each approach provides advantages and encounters challenges: early quantum servers will be expensive and thus all else equal, requiring a single prover is preferable; on the other hand, existing single-prover protocols either require a trusted device or make computational assumptions. Multi-prover protocols utilize powerful device-independence techniques which avoid these assumptions but at the high cost of requiring several powerful provers and requiring isolation.

The current zeitgeist in this field allows for imaginative considerations of how we describe and model tasks in a quantum world. These approaches have in common that instead of considering the straightforward quantum analog of classical protocols, we strive to make considerations that are naturally motivated in the quantum setting[1]. Here, we continue on this momentum and introduce a novel approach to proof verification, where the set-up itself can only be motivated in the quantum setting. To this end, we consider the following question:

▶ **Question 1.** What is the expressive power of the class of **relativistic**, interactive proof systems with a single quantum prover, and a classical verifier having access to an **off-the-shelf untrusted quantum device**?

Off-the-shelf Device. We call the above model the *off-the-shelf (OTS)* model since it models the fact that the verifier, in addition to interacting with a standard prover, has access to a device that is (1) generic (it does not depend on the instance of the problem to be solved, only on the instance size), (2) efficient (for completeness, polynomial resources suffice), (3) completely untrusted (for soundness, there are no assumptions on its computational power or inner-workings). Importantly, *relativistic* refers to a 1-round protocol; this is desirable for its relative ease in enforcing isolation[2].

Operationally, we imagine the OTS model as the prover providing the verifier with such a generic, off-the-shelf device ahead of the proof verification. In particular, the preparation of such a device in terms of its capabilities is independent of the particular problem instance, although we do allow dependence on its size. Once in possession of this device, the verifier

[1] See, for instance, the recent work on the complexity of preparing quantum states and unitaries [42].
[2] A relativistic protocol is highly desirable in the multi-prover scenario since isolation can be enforced using relative position and response times [10, 22].

may query the prover and simultaneously use a single measurement from the off-the-shelf device, which leads the verifier to *accept* or *reject*. The figures of merit for the interactive proof system are the usual *completeness* and *soundness*.

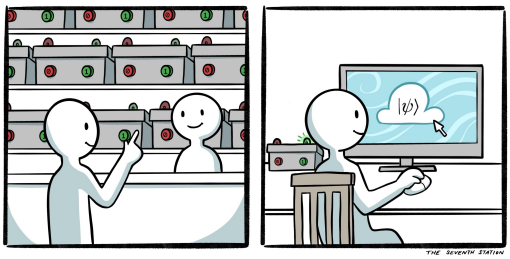

Figure 1 During the set-up the verifier selects an off-the-shelf device based on the required size of the problem instance. Afterward, the verifier is free to select any language and instance and interacts in a single round with both the prover and the off-the-shelf device, leading to the *accept/reject* output of the verifier.

Since the OTS scenario models aspects of near-term proof verification using untrusted quantum devices, we naturally wish to understand how it relates to some of the most relevant and studied properties of **interactive proof systems**:

▶ **Question 2.** Can the OTS model provide novel approaches to **zero-knowledge proof systems** and to **delegated quantum computation**?

Classical and Quantum Interactive Proof Systems. In the model of interactive proof systems (IP), an efficient classical verifier interacts with an all-powerful and untrusted prover in order to verify the correctness of a statement [20]. We note that class NP corresponds to a single-message interaction (with MA being in probabilistic version), while AM incorporates a single round (*i.e.*, *two* messages).

In *a multiprover interactive proof system* (MIP), a verifier interacts with multiple isolated provers [3]. Each of the models above has been *quantized*, *i.e.,* extended to the setting where some (or all) of the parties are quantum. This is captured, *e.g.* by the classes QMA (the quantum version of MA), QIP (the quantum version of IP) and MIP^* (a version of a multi-prover interactive proof system (MIP) where the unbounded provers share entanglement). Groundbreaking results have characterized some these quantum classes, *e.g.* $\mathsf{QIP} = \mathsf{PSPACE}$ [24] and $\mathsf{MIP}^* = \mathsf{RE}$ [27].

Zero-Knowledge Proof Systems. A strong motivation for the study of interactive proof systems is the connections to the counter-intuitive concept of a zero-knowledge proof system [19, 2]. Informally, a proof system is *zero-knowledge* when the verifier is unable to learn anything beyond the fact that the agreed-upon instance is true. This is more formally treated by establishing the existence of a *simulator* which can reproduce the transcript of the interaction.

Zero-knowledge proof systems were first extended to the quantum setting by Watrous [50], who considered the setting where the verifier has access to a trusted polynomial-time quantum device. Subsequently, it was shown that under certain cryptographic assumptions, all problems in QMA admit a zero-knowledge proof system [7, 8, 5] (again, assuming the verifier has trusted polynomial-time quantum computation). There have been several approaches in the case of a fully classical verifier. Vidick and Zhang showed that argument protocols

can be made to satisfy the zero-knowledge property [47]. Recent work by Crépeau and Stuart [17] provides a two-prover one-round zero-knowledge proof system for NP. The work of Chiesa, Forbes, Gur and Spooner provides a two-prover zero-knowledge proof system for NEXP [12], however, their work requires polynomially many rounds of interaction. Work due to Grilo, Yuen, and Slofstra [23] shows that any proof system for MIP^* can be made zero-knowledge at the cost of adding four additional provers. Although these works provide inspiration for studying zero-knowledge proof systems in the OTS model, as far as we are aware, they do not directly contribute to our main question on ZK. In fact, according to the current state-of-the-art, an implicit open question [22] is the following: *"Does there exists a* relativistic *zero-knowledge proof system for* QMA *with two provers and a classical verifier?"*. We emphasize that our OTS model takes this question further, by requiring one of the provers to operate generically and independently of the problem instance.

Delegated Quantum Computation. Delegated quantum computation allows a computationally-weak classical client to delegate a computational task to an untrusted, polynomial-time quantum server. Under certain conditions, an interactive proof system leads in a straightforward way to a protocol for delegated quantum computation. Typically, this is achieved if the interactive proof system captures *e.g.* QMA, and furthermore, given the witness, the prover is efficient; it is also relevant that the QMA witness is used in such a way that we can scale down the proof system in order to achieve a delegation protocol for BQP (*e.g.* [22])[3]. The sketch above is also applicable to the scenario of multiple servers. Note that because of the resemblance between the models of the interactive proof system and delegated quantum computation, we occasionally confound the two – using the complexity class acronym to refer to the interaction pattern between prover(s) and verifier – but we emphasize that in delegated quantum computation *protocols*, the server is always computationally bounded (as opposed to a prover in interactive *proof systems*).

Following Reichardt, Unger, and Vazirani [41], who showed a delegated quantum computation for the setting of MIP^*, much progress was made, aiming at improving parameters and techniques; despite these efforts, as far as we are aware, none of the existing works are applicable to our model. Notable here is the work of [16] which uses *quasilinear* resources for *both* servers, and achieves at best a constant round complexity, as well as [22] which is the first 2-server, 1-round (relativistic) protocol for delegated quantum computations, but uses the full polynomial-power of both servers.

1.1 Summary of results

In this work, we make important steps towards answering the above questions:

- We show that any language in QMA has a statistical ZK proof system in the OTS model.
- We show that the above OTS proof system can be adapted for delegated quantum computation for any problem in BQP, while remaining ZK and in the OTS model.

We now give more details and motivation for our model and an overview of our main contributions at the conceptual level.

Model. As introduced earlier, we are interested in modeling near-term proof verification and delegation of quantum computations. To this end, we propose a new paradigm that is particularly relevant to the quantum scenario: a verifier having access to an OTS device. To

[3] BQP is closed under complementation, hence this is sufficient for delegation

motivate the model, consider that the complexity class QMA models a verifier having access to fully-trusted polynomial-time quantum computation. While such a verifier is *skeptical* of the prover (and thus needs to verify the claimed proof independently), in the quantum case, a new level of skepticism is possible, namely that the verifier's quantum processing is untrusted. A common solution in this case is to postulate *two* (or more) untrusted and all-powerful devices together with a classical verifier; this is the realm of MIP*. In this work, we propose a new paradigm that treats the provers asymmetrically. Starting with a conventional two-prover interactive proof system, we ask that only *one* of the provers do the heavy lifting (via its unbounded computational capabilities), with the second prover becoming efficient and completely generic (for completeness, this prover need not even be given a description of the task at hand; soundness, however, is shown against *two* unbounded provers).

The inspiration for our model finds its roots in the elegant 2-prover, 1-round protocol introduced by Grilo [22]. The approach Grilo uses is a game that involves an energy test and the Pauli braiding test to verify the shared EPR pairs' integrity and the accuracy of Pauli-X/Z measurements on local terms H_i. Although in [22] the analysis assumes that players Alice and Bob are randomly assigned roles, it is natural to consider an asymmetric version of this game where one prover's functionality can be made independent of the problem instance. Grilo's work motivates the formal introduction and study of a new class which one may expect to be lower-bounded by QMA. As we outline in Section 1.2 there are substantial technical obstructions to extending this game to obtain a zero-knowledge protocol.

We denote OTS the set of all languages L that can be decided under a constant completeness-soundness gap, in the model that follows. Before the instance $x \in L$ is selected, the classical verifier is provided with an untrusted off-the-shelf device which only depends on a parameter n, indicating the size of the problem instance (without loss of generality, we can assume that the prover provides such OTS device). For completeness, such a device shares an entangled state $|\psi\rangle_n$ with a quantum prover and will be purported to perform efficient measurements from a predetermined list of available options[4]. The verifier may select any choice $x \in L$ provided $|x| \leq n$ and *simultaneously* uses a single question to the prover and to the device; the verifier then determines whether or not to accept based on the responses. We stress that OTS proof systems are sound against both an unbounded prover and unbounded OTS.

We observe that OTS is a refinement of and thus contained in MIP*, and is also a generalization of AM, where the otherwise classical verifier has additional 1-round query access to a small, off-the-shelf quantum device. In summary, we have AM $\subseteq$ OTS $\subseteq$ MIP* (see also Figure 2).

In a classical proof system, an OTS can be understood as an instance-independent hardware token. This device can be used to provide a commitment for a zero-knowledge proof system for NP [19]; what is more, the one-time property of the OTS can be used as an oblivious transfer device, which then yields a non-interactive zero-knowledge proof system for NP [28]. We note that in the quantum case, our model requires a fully classical verifier and hence the case of zero-knowledge for QMA [7, 5] in the OTS model is much more complex, and a classical-verifier analogue to the NP proof systems above is not directly applicable. Other approaches based on using the OTS as a one-time memory [6] also run into a roadblock due to the fact that we require a fully classical verifier.

[4] The entangled state $|\psi_n\rangle$ is consumed during the interactive proof, hence a new OTS must be obtained for subsequent evaluations (equivalently, the entanglement must eventually be replenished). This situation is entirely analogous to the case of shared randomness which is also consumed in an interactive proof system and must also eventually be replenished.

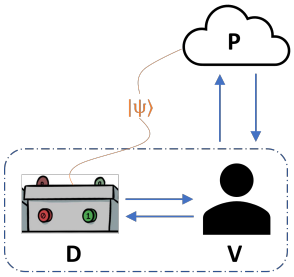

Figure 2 Off-the-shelf (OTS) proof system. P is the quantum prover, V is the classical verifier, and D a rudimentary off-the-shelf device which is entangled with the prover; each arrow represents a single classical message.

OTS Proof Systems for QMA. Our first result is that any language in QMA is also in OTS.

▶ **Theorem 3** (Restated as part of Theorem 31)**.** QMA $\subseteq$ OTS.

An interpretation of this result is that starting with a conventional proof system for QMA, we can exchange the unwavering trust of the verifier in its quantum verification process for a classical verifier with two new features: (1) the verifier has access to an untrusted, and instance-independent, off-the-shelf quantum device; and (2) the verifier interacts with the prover (and the device) in a single simultaneous round.

Zero-knowledge OTS Proof System for QMA. What is more, we show that the OTS proof system for QMA is also *statistical zero-knowledge*, meaning that we can simulate in classical polynomial time the verifier's transcript when interacting with the provers on a yes-instance.

▶ **Theorem 4** (Restated in Theorem 31)**.** *For every language L in* QMA*, there exists a statistical zero-knowledge OTS proof system for L.*

Delegated Quantum Computation in the OTS Model. As our final conceptual contribution, we show how our OTS proof system for QMA (Theorem 3) can be adapted to the setting of delegated quantum computation; note that the ZK property as described above also extends to the delegated quantum computation paradigm.

▶ **Theorem 5** (Restated as Theorem 32)**.** BQP *has a relativistic delegated quantum computation protocol in the OTS model with the statistical zero-knowledge property.*

We believe that this result is of particular impact since it addresses a new model for delegated quantum computation that has distinct conceptual benefits over existing protocols:

1. Comparing to single-server protocols, we note that we make an extra assumption of an off-the-shelf, isolated device. However, the benefits are:
   a. We achieve soundness against an unbounded server; existing single-server, classical-client delegation protocols require computational assumptions [31].
   b. The client does not trust *any* quantum device at all; existing single-server, statistically secure protocols require trust in a small quantum preparation device [4, 18].

2. Comparing to existing multiple-server (MIP*) protocols, we note that:
   a. Our approach only requires a single high-performance quantum server that handles the bulk of the computations; with a secondary efficient and generic device which need not even be given a description of the problem instance. This has practical advantages, especially when we consider that the off-the-shelf device can be acquired ahead of the verification stage (Figure 1).
   b. Our approach is a single round, which means that relativistic means to enforce isolation are possible. The only other known relativistic protocol requires full quantum computational power for both servers and is not ZK [22].

1.2 Proof approach and technical contributions

We first introduce two important techniques.

Self-testing. *Self-testing* (also called *device-independence*) is a ubiquitous and powerful technique in the study of MIP* and related delegation protocols. The concept was introduced by Mayers and Yao [33]. Informally, a protocol *self-tests* a particular state or measurement when this state/measurements (or an equivalent version thereof) are required for obtaining the maximal acceptance probability. The most well-known examples are the non-local games known as the CHSH game and the Magic Square game [13, 36, 40, 45]. Subsequently, numerous works have enriched our understanding of self-testing and its applications to delegated quantum computation, *e.g.*, [34, 35, 15, 11, 16, 37, 38]. Current approaches to formalizing self-testing use the theory of approximate representation theory of groups and C^*-algebras [43, 44, 32]. These formalisms, and especially their operationally-useful *approximate* versions utilize a key stability result due to Gowers and Hatami which allows one to relate approximate representations to exact representations [21].

Simulatable Codes. Recent works by Grilo, Yuen, and Slofstra [23], as well as Broadbent and Grilo [5] introduce the notion of simulatable codes as a tool for establishing zero-knowledge proof systems and protocols in the quantum setting. The idea is to use techniques from quantum error-correcting codes to create a "simulatable" witness or proof for use in the verification process. Here the witness is *simulatable* in the sense that there is an efficient classical algorithm which can reproduce the description of the local density matrix of the witness on any small enough subspace. This is a pivotal tool in establishing zero-knowledge, and the application of the technique consists in developing a verification protocol, (or verification circuit in the case of [5]) which verifies such simulatable witnesses; this can then be applied to the situation of encoding *e.g.*, a witness for QMA into a simulatable code [5].

1.2.1 Obstructions to the straightforward approach

In delegating quantum computations in two- or multi-server models, the classical verifier is able to command quantum provers [41] using two intertwined tests: (1) a computational test, with acceptance probability based on the required quantum computation (*e.g.*, computation-by-teleportation [41] or energy checking of a local Hamiltonian [26, 22]); (2) a rigidity test, ensuring provers' actions stay within a known range (*e.g.*, self-test via CHSH game or Pauli braiding test). In order to establish the ZK property, we must show that responses from the provers can be simulated using a classical probabilistic polynomial-time (PPT) device. Generally, approaches used for the rigidity test can be simulated in a straightforward

way, hence the difficulty in obtaining ZK in this setting is in simulating the energy test. Furthermore, even if both tests are simulatable in isolation, this does not guarantee the ZK property since a malicious verifier may form question pairs emanating from different tests, during a single round.

Grilo [22] presents a game $\mathcal{G}(H)$ determined by an "XZ-type"[5] Hamiltonian H. Honest provers for this game share suitably many EPR pairs, and one prover privately holds a ground state for H. The game $\mathcal{G}(H)$ combines an energy test with the Pauli braiding test [37, 46]. During the energy test, one prover reports measurement results of a randomly chosen term H_i on their side of EPR pairs, and the other provides teleportation keys from a Bell basis measurement on the other EPR pairs and the ground state. Combining the energy test with the Pauli braiding test allows the verifier to ensure that provers share n EPR pairs and that the required Pauli-X/Pauli-Z measurements are performed when measuring the local term H_i.

The straightforward approach to obtaining a two-prover ZK proof system would be to combine recent results on simulatable codes in order to make the measurement results in Grilo's energy test simulatable. More specifically, one could apply the well-known circuit-to-Hamiltonian construction using the family of simulatable verification circuits given in [5]. Given such a circuit V, it is shown that local measurements on the ground state of the corresponding Hamiltonian H_V are simulatable and thus this approach would make the results of the energy test simulatable. Unfortunately, this approach fails for two technical reasons.

The Choice of Encoding. Firstly, one cannot employ previously-known self-testing techniques to show the players perform the required measurements on the simulatable ground states given in [5]. On the one hand, previously-studied single-round self-testing techniques can only be used to show the players perform Pauli-X, and Pauli-Z measurements. On the other hand, the choice of physical gates used by Broadbent and Grilo during the encoding of logic gates may result in a local Hamiltonian that is not of XZ-type and thus local terms H_i may require measurements that have no known self-test.

The Size of the Measurement. The second obstruction arises from the fact that existing rigidity tests in this setting require both players to make large-sized measurements on their shared state. These large measurements can provide an avenue for attack by a malicious verifier which compromises the zero-knowledge property. In particular, since the Pauli braiding test allows for requests for measurements on all qubits, a malicious verifier may indicate to one player that an energy test is being played and simultaneously request Pauli-X and Pauli-Z measurements on a large number of qubits. Such a measurement result cannot be simulated using simulatable codes, which only protect against constant-sized measurements, and thus this compromises zero-knowledge.

1.2.2 Overview of proof and technical results

In order to correct for an appropriate choice of encoding, we prove that one can re-instantiate the verification circuit given by Broadbent and Grilo using an approach to simulatable codes given in [23]. This change allows us to encode logical gates of the verification circuit given

[5] These are Hamiltonians where each local term H_i is a real linear combination of tensor products of the Pauli-X and Pauli-Z operators.

by Broadbent and Grilo using a different set of physical gates and consequently, we show that the local Hamiltonian corresponding to the circuit is of XZ-type, while preserving simulatability.

▶ **Theorem 6** (Informal version of Theorem 27)**.** *For any language $L = (L_{yes}, L_{no})$ in* QMA*, there is a family of verification circuits V_x satisfying (1) the circuit-to-Hamiltonian construction applied to V_x produces a Hamiltonian H_x which is of XZ-type, and (2) if $x \in L_{yes}$ there exists a polynomial-time algorithm that can approximate the reduced density matrix obtained by tracing out all but 6 qubits of the ground state of H_x.*

In order to overcome the large measurement problem, we introduce a new self-test called the *low-weight Pauli braiding test* (LWPBT) which can self-test the low-weight tensor products of Pauli measurements and n EPR pairs but only requires the players to make measurements on a constant number of qubits.

▶ **Theorem 7** (Informal version of Theorem 18)**.** *The low-weight Pauli braiding test can self-test for n EPR pairs and 6-qubit Pauli measurements. This self-test is robust in the sense that any ε-perfect strategy must be $\text{poly}(n)\sqrt{\varepsilon}$ close to the canonical strategy.*

We use a group-theoretical approach to prove the rigidity of the LWPBT. It can be shown that the canonical perfect strategy $\widetilde{S}$ for LWPBT defines an irreducible representation of the Weyl-Heisenberg group H and every near-perfect strategy S for LWPBT forms an approximate homomorphism f of H. The well-known Gowers-Hatami theorem [21] and its variant [46] imply that the approximate homomorphism f of a finite group is close to a representation ϕ, so S must be close to $\widetilde{S}$. However, some subtle mathematical problems have come up in earlier approaches. In particular, one may need to discard some irreducible constituents of ϕ that do not correspond to $\widetilde{S}$. To tackle this problem, we make further improvements to the state-of-art understanding of the stability of groups. In particular, in Theorem 18 we state and prove an enhanced version of the Gowers-Hatami theorem that can be used for the stability analysis of the Weyl-Heisenberg group. Aside from our use case, this new version can simplify previous approaches to self-testing.

We use the above technical results to derive a modified version of [22] by interleaving the following tests: (1) a computational test consisting of an energy test in which a simulatable witness uses low-weight Pauli-X and Pauli-Z measurements and, (2) a rigidity test consisting of the LWPBT. The result of this modified Grilo protocol gives a ZK OTS protocol with an inverse polynomial completeness-soundness gap. Finally, we apply a threshold parallel repetition theorem to the above protocol to amplify the completeness-soundness gap to be constant, thus demonstrating both Theorem 3 and Theorem 4. We then show that the proof system is of a form that can be scaled down to yield a delegation protocol, yielding Theorem 5.

2 Preliminaries

We take $[n]$ to denote the set $\{1, \dots n\}$. Given two real valued functions $f, g : \mathbb{R} \to \mathbb{R}$, we write $f = O(g)$ (resp. $f = \Omega(g)$) if there exists a positive real number M and an $x_0 \in \mathbb{R}$ such that $|f(x)| \leq Mg(x)$ (resp. $|f(x)| \geq Mg(x)$) for all $x \geq x_0$. We call a function f negligible, and write $f = \text{negl}(n)$, if for all constants $c > 0$ we have $f = O(n^{-c})$. For two distributions P and Q on a finite set $\mathcal{X}$ the statistical differences of P and Q is given by $\sum_{x \in \mathcal{X}} |P(x) - Q(x)|$.

In this paper, all Hilbert spaces are finite-dimensional. Given a Hilbert space $\mathcal{H}$, we use $\mathcal{B}(\mathcal{H})$ to denote the set of bounded linear operators acting on $\mathcal{H}$, use $\mathcal{U}(\mathcal{H})$ to denote the group of unitary operators on $\mathcal{H}$, and use $\mathbb{1}_{\mathcal{H}}$ to denote the identity operator on $\mathcal{H}$. Given an operator $A \in \mathcal{B}(\mathcal{H})$ we take A^* to denote the adjoint operator (equivalently the conjugate transpose) and define the trace norm $\|A\|_{tr} := \text{Tr}\sqrt{A^*A}$.

2.1 Quantum information

A quantum state ρ on $\mathcal{H}$ is a positive operator in $\mathcal{B}(\mathcal{H})$ with $\mathrm{Tr}(\rho) = 1$. It induces a semi-norm $\|A\|_\rho := \sqrt{\mathrm{Tr}(A^*A\rho)}$ on $\mathcal{B}(\mathcal{H})$ which we call the ρ-norm. This norm is left unitarily invariant, meaning that $\|UA\|_\rho = \|A\|_\rho$ for all $U \in \mathcal{U}(\mathcal{H})$ and $A \in \mathcal{B}(\mathcal{H})$. Given two quantum states ρ and σ we define their trace distance $D(\rho,\sigma) = \frac{1}{2}\|\rho - \sigma\|_{tr} = \max_P \mathrm{Tr}(P(\rho-\sigma))$ where the max is taken over all projections $P \in \mathcal{B}(\mathcal{H})$.

We use $|\Phi_{\text{EPR}}\rangle$ to denote the EPR pair in $\mathbb{C}^2 \otimes \mathbb{C}^2$ and use $|\Phi_{\text{EPR}}^{\otimes n}\rangle$ to denote the n-qubit EPR pair. We also take σ_I, σ_X, and σ_Z to denote the following Pauli operators:

$$\sigma_I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \sigma_X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \text{ and } \sigma_Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{1}$$

For every $a \in \{0,1\}^n$ and $W \in \{I,X,Z\}^n$, we use $\sigma_W(a)$ to denote the operator $\otimes_{i\in[n]}\sigma_{W_i}^{a_i}$ on $(\mathbb{C}^2)^{\otimes n}$ where $\sigma_I^0 = \sigma_X^0 = \sigma_Z^0 = \sigma_I$. Definitions of these gates and other fundamental concepts from quantum computing can be found in [39].

Families of Quantum Circuits. A *unitary quantum circuit* is simply a unitary which can be written as a product of gates from some universal gate set $\mathcal{U}$. Unless otherwise specified we will assume the universal gate set is the following universal gate set $\{H, \Lambda(X), \Lambda^2(X)\}$, where H is the Hadamard gate, $\Lambda(X)$ is the controlled σ_X gate, and $\Lambda^2(X)$ is the Toffoli gate. A *general quantum circuit* or simply a *quantum circuit* is a unitary quantum circuit that can additionally apply non-unitary gates which, introduce qubits initialized in the 0 state, trace out qubits, or measure qubits in the standard basis.

▶ **Definition 8** (Polynomial-time uniform circuit family)**.** *We say a family of quantum circuits $\{Q_n\}_{n\in\mathbb{N}}$ is a polynomial-size family of quantum circuits if there exists polynomial r such that Q_n has size at most $r(n)$. A family of quantum circuits $\{Q_n\}$ is called* polynomial-time uniform *family if there exists a polynomial time Turing machine that on input 1^n outputs a description of Q_n. In this case, the family will also be a polynomial-size family of quantum circuits.*

Given a quantum circuit Q, we denote its size (number of gates and number of wires) by $|Q|$. The task of delegating the computation of Q is captured by the following promise problem:

▶ **Definition 9** (Q-CIRCUIT)**.** *The input is a quantum circuit Q on n qubits. The problem is to decide between the following two cases:*

- ***Yes.*** $\|((|1\rangle\langle 1| \otimes I_{n-1})Q|0^n\rangle\|^2 \geq 1-\gamma$
- ***No.*** $\|((|1\rangle\langle 1| \otimes I_{n-1})Q|0^n\rangle\|^2 \leq \gamma$

when we are promised that one of the two cases holds.

Problem in Definition 9 is known to be BQP-complete for $1 - 2\gamma > \frac{1}{\mathrm{poly}(n)}$.

2.2 Non-local games and rigidity

A two-player[6] one-round nonlocal game $\mathcal{G}$ is a tuple $\left(\lambda, \mu, \mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A, \mathcal{O}_B\right)$, where $\mathcal{I}_A, \mathcal{I}_B$ are finite input sets, and $\mathcal{O}_A, \mathcal{O}_B$ are finite output sets, μ is a probability distribution on $\mathcal{I}_A \times \mathcal{I}_B$, and $\lambda : \mathcal{O}_A \times \mathcal{O}_B \times \mathcal{I}_A \times \mathcal{I}_B \to \{0,1\}$ determines the win/lose conditions. A

[6] These two players are commonly called Alice and Bob.

quantum strategy $\mathcal{S}$ for $\mathcal{G}$ is given by finite-dimensional Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, a unit vector $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$, Alice's POVMs $\{E_a^x : a \in \mathcal{O}_A\}, x \in \mathcal{I}_A$ on $\mathcal{H}_A$, and Bob's POVMs $\{F_b^y : b \in \mathcal{O}_B\}, y \in \mathcal{I}_B$ on $\mathcal{H}_B$. The winning probability of $\mathcal{S}$ for game $\mathcal{G}$ is given by

$$\omega(\mathcal{G}, \mathcal{S}) := \sum_{a,b,x,y} \mu(x,y)\lambda(a,b|x,y) \langle\psi| E_a^x \otimes F_b^y |\psi\rangle .$$

A quantum strategy $\mathcal{S}$ for a non-local game $\mathcal{G}$ is said to be *perfect* if $\omega(\mathcal{G}, \mathcal{S}) = 1$. When the game is clear from the context we simply write $\omega(\mathcal{S})$ to refer to the winning probability. The *quantum value* of a non-local game $\mathcal{G}$ is defined as

$$\omega^*(\mathcal{G}) := \sup\{\omega(\mathcal{S}) : \mathcal{S} \text{ a quantum strategy for } G\}.$$

In this paper, we assume all measurements employed in a quantum strategy are PVMs. An m-outcome PVM $\{P_1, \cdots, P_m\}$ corresponds to an observable $\sum_{j\in[m]} \exp(\frac{2\pi i}{m} j) P_j$, so a quantum strategy for a game $\mathcal{G} = \big(\lambda, \mu, \mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A, \mathcal{O}_B\big)$ can also be specified by a triple

$$\mathcal{S} = (\tau^A, \tau^B, |\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B)$$

where $\tau^A(x)$, $x \in \mathcal{I}_A$ are $\mathcal{O}_A$-outcome observables on $\mathcal{H}_A$, and $\tau^B(y)$, $y \in \mathcal{I}_B$ are $\mathcal{O}_B$-outcome observables on $\mathcal{H}_B$.

Here we introduce the well-known Mermin-Peres Magic Square game, in which Alice and Bob are trying to convince the verifier that they have a solution to a system of equations over $\mathbb{Z}_2$. There are 9 variables $v_1, \ldots, v_9$ in a 3×3-array whose rows are labeled r_1, r_2, r_3 and columns are labeled c_1, c_2, c_3.

Each row or column corresponds to an equation: variables along the rows or columns in $\{r_1, r_2, r_3, c_1, c_2\}$ sum to 0; variables along the column c_3 sum to 1. In each round, Bob receives one of the 6 possible equations and he must respond with a satisfying assignment to the given equation. Alice is then asked to provide a consistent assignment to one of the variables contained in the equation Bob received. The following table describes an operator solution for this system of equations:

Table 1 Operator solution for Magic Square game.

$A_1 = \sigma_I \otimes \sigma_Z$	$A_2 = \sigma_Z \otimes \sigma_I$	$A_3 = \sigma_Z \otimes \sigma_Z$
$A_4 = \sigma_X \otimes \sigma_I$	$A_5 = \sigma_I \otimes \sigma_X$	$A_6 = \sigma_X \otimes \sigma_X$
$A_7 = \sigma_X \otimes \sigma_Z$	$A_8 = \sigma_Z \otimes \sigma_X$	$A_9 = \sigma_X\sigma_Z \otimes \sigma_Z\sigma_X$

2.3 Complexity classes and zero knowledge

▶ **Definition 10** (QMA). *A promise problem $L = (L_{yes}, L_{no})$ is in* QMA *if there exist polynomials p and q, and a polynomial-time uniform family of quantum circuits $\{Q_n\}$ where Q_n takes as input a string $x \in \Sigma^*$ with $|x| = n$, a $p(n)$-qubit quantum state $|\psi\rangle$, and $q(n)$ auxiliary qubits in state $|0\rangle^{\otimes q(n)}$, such that:*

- *(Completeness) if $x \in L_{yes}$, then there exists some $|\psi\rangle$ such that Q_n accepts $(x, |\psi\rangle)$ with probability at least $1 - \mathrm{negl}(n)$, and*
- *(Soundness) if $x \in L_{no}$, then for any state $|\psi\rangle$, Q_n accepts $(x, |\psi\rangle)$ with probability at most $\mathrm{negl}(n)$.*

We sometimes refer to the family of circuits $\{Q_n\}$ in Definition 10 simply as a family of *verification circuits*.

The following *local Hamiltonian problem* is QMA-complete for parameters $k = 5$ and $\beta - \alpha = \frac{1}{\mathrm{poly}(n)}$ [29].

▶ **Definition 11.** *Let $k \in \mathbb{N}$, $\alpha, \beta \in \mathbb{R}$ with $\alpha < \beta$, the* k-Local Hamiltonian *problem with parameters α and β is the following promise problem. Let n be the number of qubits of a quantum system. The input is a set of $m(n)$ Hamiltonians $H_1, \ldots, H_{m(n)}$ where m is a polynomial in n and each H_i acts on k qubits out of the n qubit system with $\|H_i\| \leq 1$. For $H = \sum_{j=1}^{m(n)} H_j$ the promise problem is to decide between the following.*

- ***Yes.*** *There exists an n-qubit state $|\varphi\rangle$ such that $\langle\varphi| H |\varphi\rangle \leq a \cdot m(n)$.*
- ***No.*** *For every n-qubit state $|\varphi\rangle$ it holds that $\langle\varphi| H |\varphi\rangle \geq b \cdot m(n)$.*

In this work, we also deal with MIP^* proof systems that involve two provers and one round.

▶ **Definition 12.** *A promise language $L = (L_{yes}, L_{no})$ is in $\mathsf{MIP}^*[2,1]_{c,s}$ if there exists a polynomial-time computable function that takes an instance $x \in L$ to a description of a non-local game $\mathcal{G}_x$ satisfying the following conditions.*

- *(Completeness) For every $x \in L_{yes}$ we have $\omega^*(\mathcal{G}_x) \geq c$.*
- *(Soundness) For every $x \in L_{no}$ we have $\omega^*(\mathcal{G}_x) < s$.*

We refer to the mapping, $x \mapsto \mathcal{G}_x$, as a $\mathsf{MIP}^*[2,1]_{c,s}$ proof system, or in some places a $\mathsf{MIP}^*[2,1]_{c,s}$ protocol. When the parameters are clear from the context we simply call it an MIP^* proof system.

Next, we discuss *zero knowledge*. In an interactive proof system, a malicious verifier $\widehat{V}$ is a probabilistic polynomial-time Turing machine which on input x and randomness θ samples question q_1 for either Alice or Bob. Given reply r_1, the malicious verifier samples question q_2 in a way that may depend on q_1 and r_1. For a given quantum strategy $\mathcal{S}$ and malicious verifier $\widehat{V}$, we take $View(\widehat{V}(x), \mathcal{S})$ to be the random variable corresponding to the transcript of questions and answers $(x, \theta, q_1, r_1, q_2, r_1)$. A protocol is zero-knowledge when for all "yes" instances a simulator can sample from the distribution above.

▶ **Definition 13.** *An* $\mathrm{MIP}^*[2,1]_{c,s}$ *proof system is* statistical zero-knowledge *if for every $x \in L_{yes}$ there exists an honest prover strategy $\mathcal{S}$ satisfying the following:*

1. $\omega^*(\mathcal{S}) \geq c$.
2. *For any PPT malicious verifier $\widehat{V}$ there exists a PPT simulator $Sim_{\widehat{V}}$ with output distribution that is ε-close to $View(\widehat{V}(x), S)$ in statistical distance for some negligible function $\varepsilon(|x|)$.*

2.4 Simulatable codes and encodings of gates

Recall that a quantum error-correcting code (QECC) $\mathcal{C} = [[n,k]]$ is a map $\mathrm{Enc} : (\mathbb{C}^2)^{\otimes k} \to (\mathbb{C}^2)^{\otimes n}$, which encodes a k-qubit state $|\psi\rangle$ into an n-qubit state $\mathrm{Enc}(|\psi\rangle)$ where $n \geq k$. The code is said to have distance d if the original state can be recovered from the encoded state that has transformed under any quantum operation which acts on at most $(d-1)/2$ qubits. Given an $[[m,1]]$ QECC with map Enc, we abuse notation and also write Enc for the corresponding $[[mn,n]]$ encoding that is obtained by applying Enc to each of the qubits in an n-qubit system.

We use $\underline{A}_n^k$ to denote the set of k distinct numbers between 1 and n through this section. Then $\underline{A}^k := \bigcup_{n \geq k} \underline{A}_n^k$ is the set of k distinct numbers. Given a k-qubit logical gate U and an element $\underline{a} = (a_1, \ldots, a_k) \in \underline{A}^k$, let $U(\underline{a})$ denote the gate U applied to qubits $a_1, \cdots a_k$.

Below we recall the definition of simulatable codes introduced in [23].

▶ **Definition 14.** *Given a k-qubit logical gate U and a quantum error-correcting code $\mathcal{C} = [[m,1]]$, let (σ_U, σ'_U) be a pair of states, and let ℓ be a positive integer. For each $1 \leq i \leq \ell$, let $\mathcal{O}_i$ be a mapping from elements $\underline{a} = (a_1, \cdots, a_k)$ in $\underline{A}^k$ to unitaries $\mathcal{O}_i(\underline{a})$ acting only on*

(i) *the physical qubits of codewords in* $\mathcal{C}$ *that corresponds to logical qubits* $a_1, \cdots, a_k$*, and*
(ii) *the register that holds* σ_U*.*

We say the tuple $(\sigma_U, \sigma'_U, \ell, \mathcal{O}_1, \cdots, \mathcal{O}_\ell)$ *is an encoding of* U *in code* $\mathcal{C}$ *if*

$$(\mathcal{O}_\ell(\underline{a}) \dots \mathcal{O}_1(\underline{a}))(Enc(\rho) \otimes \sigma_U)(\mathcal{O}_\ell(\underline{a}) \dots \mathcal{O}_1(\underline{a}))^* = Enc\big(U(\underline{a})\rho U(\underline{a})^*\big) \otimes \sigma'_U \tag{2}$$

for all $n \geq k$*, elements* $\underline{a} \in \underline{A}_n^k$*, and* n*-qubit states* ρ*. If in addition, the unitaries* $\mathcal{O}_1(\underline{a}), \cdots, \mathcal{O}_\ell(\underline{a})$ *are gates in some set* $\mathcal{U}$ *for all* $\underline{A} \in \underline{A}^k$*, then we say the encoding* $(\sigma_U, \sigma'_U, l, \mathcal{O}_1, \cdots, \mathcal{O}_\ell)$ *uses physical gates in* $\mathcal{U}$*.*

Given a circuit of logical gates $V = U_1 \dots U_k$ we refer to an encoding of V as the corresponding circuit of physical gates obtained by applying an encoding of each gate U_i.

▶ **Definition 15.** *An encoding* $(\sigma_U, \sigma'_U, \mathcal{O}_1, \dots, \mathcal{O}_\ell)$ *of a* k*-qubit logical gate* U *in a QECC* $\mathcal{C}$ *is called* s*-simulatable if for all* $0 \leq t \leq \ell$*,* n*-qubit states* ρ*, and subsets* S *of the physical qubits of* $Enc(\rho) \otimes \sigma_U$ *with* $|S| \leq s$*, the partial trace*

$$\mathrm{Tr}_{\overline{S}}\Big((\mathcal{O}_t(\underline{a}) \dots \mathcal{O}_1(\underline{a}))(Enc(\rho) \otimes \sigma_U)(\mathcal{O}_t(\underline{a}) \dots \mathcal{O}_1(\underline{a}))^*\Big)$$

is a $2^{|S|} \times 2^{|S|}$ *matrix whose entries are rational and can be computed in polynomial time from* t*,* $\underline{a}$ *and* S*. In particular, this matrix is independent of* ρ *if* $\mathcal{C}$ *can correct arbitrary errors on* s *qubits.*

▶ **Theorem 16** (Theorem 6 in [23])**.** *Let* $\mathcal{U} = \{H, \Lambda(X), \Lambda^2(X)\}$*. For every* $s \in \mathbb{N}$*, there exists a constant* $n \in \mathbb{N}$ *and a* $[[n, 1]]$ *QECC* $\mathcal{C}$ *such that any logical gate in* $\mathcal{U}$ *has an* s*-simulatable encoding in* $\mathcal{C}$ *using physical gates in* $\mathcal{U}$*.*

3 Low-weight Pauli braiding test and its rigidity

For any $a \in \{0,1\}^n$ and $W \in \{X, Z\}^n$, we use $W(a)$ to denote the sequence $W_1^{a_1} W_2^{a_2} \cdots W_n^{a_n}$ where $X^0 = Z^0 = I$. Let $\mathcal{I}_A := \{W(a) : W \in \{X, Z\}^n, a \in \{0,1\}^n \text{ such that } |a| \leq 6\}$ and let $\mathcal{I}_B := \{(W(a), W(a')) : W \in \{X, Z\}^n, a, a' \in \{0,1\}^n \text{ such that } |a|, |a'| \leq 6\}$ be the question sets for Alice and Bob respectively. We first describe the low-weight linearity test in Figure 3.

1. The verifier selects uniformly at random $W \in \{X, Z\}^n$ and strings $a, a' \in \{0,1\}^n$ satisfying $|a|, |a'| \leq 6$ (*i.e.* a, a' both have at most 6 non-zero entries).
2. The verifier sends $(W(a), W(a'))$ to Bob. If $a + a'$ has weight at most 6 then the verifier selects $W' \in \{W(a), W(a'), W(a+a')\}$ uniformly at random to send to Alice. Otherwise, the verifier uniformly at random sends $W' \in \{W(a), W(a')\}$ to Alice.
3. The verifier receives two bits (b_1, b_2) from Bob and one bit c from Alice.
4. If Alice receives $W(a)$ then the verifier requires $b_1 = c$. If Alice receives $W(a')$ then the verifier requires $b_2 = c$. If Alice receives $W(a + a')$ then the verifier requires $b_1 + b_2 = c$.

Figure 3 Low-weight linearity test.

Next, we introduce a natural version of the anti-commutation test in Figure 4. This test is built from the well-known Magic Square game which we described in Section 2.2.

Combining the low-weight linearity test and low-weight anti-commutation test, we now construct the low-weight Pauli braiding test and state its rigidity result.

1. The verifier samples uniformly at random a string $a \in \{0,1\}^n$ with exactly two non-zero entries $i < j$. The verifier also samples a row or column $q \in \{r_1, r_2, r_3, c_1, c_2, c_3\}$, and a variable v_k contained in q as in the Magic Square game.
2. Bob receives the question (q, a).
3. If $k \neq 9$ then Alice receives $W(a) = I^{i-1} W_i I^{j-i} W_j I^{n-j} \in \mathcal{I}_A$ with $\sigma_{W_i} \otimes \sigma_{W_j} = A_k$. If $k = 9$ then Alice receives question (v_9, a).
4. The players win if and only if Bob responds with a satisfying assignment to q and Alice provides an assignment to variable v_k that is consistent with Bob's.

Figure 4 Low-weight anti-commutation test.

▶ **Definition 17.** *The low-weight Pauli braiding test (LWPBT) is played by executing with probability* $1/2$ *either the low-weight anti-commutation test or the low-weight linearity test.*

The n-qubit LWPBT has a canonical perfect strategy $\widetilde{S}$ in which Alice and Bob share the n-qubit EPR pair $|\Phi_{EPR}^{\otimes}\rangle$ and Alice perform $\sigma_W(a) := \otimes_{i=1}^n \sigma_{W_i}^{a_i}$ on question $W(a) \in \mathcal{I}_A$. We have the following rigidity result for near-perfect strategies of LWPBT.

▶ **Theorem 18.** *There exists a constant* $C_{\mathrm{lw}} > 0$ *such that the following holds. For any* $\varepsilon > 0$, $n \in \mathbb{N}$, *and strategy* $\mathcal{S} = (\tau^A, \tau^B, |\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B)$ *for the* n*-qubit LWPBT with winning probability* $1 - \varepsilon$, *there are isometries* $V_A : \mathcal{H}_A \to (\mathbb{C}^2)^{\otimes n} \otimes \mathcal{H}_A^{aux}, V_B : \mathcal{H}_B \to (\mathbb{C}^2)^{\otimes n} \otimes \mathcal{H}_B^{aux}$ *and a unit vector* $|aux\rangle \in \mathcal{H}_A^{aux} \otimes \mathcal{H}_B^{aux}$ *such that*

$$\|(V_A \otimes V_B)\left(\tau^A(W(a)) \otimes Id_{\mathcal{H}_B} |\psi\rangle\right) - \left(\sigma_W(a) \otimes Id_{\mathbb{C}^{2^n}} |\Phi_{EPR}^{\otimes n}\rangle\right) \otimes |aux\rangle\| \leq C_{\mathrm{lw}} n^6 \varepsilon^{1/4}$$

for all $W(a) \in \mathcal{I}_A$.

The proof Theorem 18 uses a group-theoretical approach and can be found in the full version of our paper [9]. The idea is that we can round an approximate homomorphism (defined by a near-perfect stategy) of the Weyl-Heisenberg group to an exact representation using an enhanced Gowers-Hatami theorem.

3.1 An enhanced Gowers-Hatami theorem

For a finite group G, we use $\mathrm{Irr}(G)$ to denote the unique (up to unitary equivalence of elements) complete set of inequivalent irreducible representations. Given a finite group G, a function $f : G \to \mathcal{U}(\mathcal{H})$ from $\mathcal{G}$ to unitaries on a Hilbert space $\mathcal{H}$, and an irreducible representation $\phi : G \to \mathcal{U}(\mathbb{C}^d)$, the *Fourier transform of* f *at* ϕ is the operator

$$\widehat{f}(\phi) := \frac{1}{|G|} \sum_{g \in G} f(g) \otimes \overline{\phi(g)}, \tag{3}$$

where $\overline{\phi(g)}$ is the conjugate of the matrix $\phi(g) \in M_d(\mathbb{C})$ in the standard basis.

Let $f : G \to \mathcal{U}(\mathcal{H})$ be a function of a finite group G. Given a quantum state ρ on $\mathcal{H}$ and a positive real number ε, we say f is an (ε, ρ)-*homomorphism* provided that $f(g^{-1}) = f(g)^*$ and $\frac{1}{|G|} \sum_{h \in G} \|f(g)f(h) - f(gh)\|_\rho^2 \leq \varepsilon$ for all $g \in G$. In this case, by the well-known Gowers-Hatami theorem [21, 46], there is a Hilbert space $\mathcal{K}$, an isometry $V : \mathcal{H} \to \mathcal{K}$, and a representation $\phi : G \to \mathcal{U}(\mathcal{K})$ such that $\|f(g) - V^* \phi(g) V\|_\rho \leq \varepsilon$ for all $g \in G$. The following enhanced version of this theorem allows us to disregard all one-dimensional irreducible representations of the Weyl-Heisenberg group. Earlier works dealt with these

one-dimensional representations by invoking a truncation of the isometry given by the Gowers-Hatami theorem. Unfortunately, in general, truncation of an isometry can fail to be an isometry.

▶ **Theorem 19.** *For any (ε, ρ)-homomorphism $f : G \to \mathcal{U}(\mathcal{H})$ of a finite group G on a finite-dimensional space $\mathcal{H}$, there exists a finite-dimensional Hilbert space $\mathcal{K}$, an isometry $I : \mathcal{H} \to \mathcal{K}$, and a representation $\pi : G \to \mathcal{U}(\mathcal{K})$ such that*

(i) $\|f(g) - I^*\pi(g)I\|_\rho^2 \leq \varepsilon$ *for all* $g \in G$, *and*

(ii) $\xi \in \mathrm{Irr}(G)$ *is an irreducible constituent of* π *only if* $\widehat{f}(\xi) \neq 0$.

We refer the readers to the full paper [9] for the proof details.

4 Modified Hamiltonian game

In this section, we show that for some local Hamiltonian H, one can construct a nonlocal game $\mathcal{G}(H)$ whose winning probability is closely related to the ground state energy $\lambda_0(H)$ of H. Our game is based on the Hamiltonian game introduced by Grilo [22]. We employ LWPBT against dishonest quantum provers and then perform parallel repetition to achieve a constant completeness-soundness gap. To incorporate LWPBT in our modified Hamiltonian test, we consider Hamiltonians with specific structures:

▶ **Definition 20.** *We say a Hamiltonian H is of XZ-type if it can be decomposed as $H = \frac{1}{m}\sum_{\ell=1}^{m} \gamma_\ell H_\ell$ where each $\gamma_\ell \in [-1, 1]$ and each term H_ℓ is a tensor product of operators σ_X, σ_Z or σ_I.*

Next, we define the relevant energy test which is analogous to the energy test used in [22].

▶ **Definition 21** (Energy test)**.** *Given an n-qubit 6-local Hamiltonian $H = \frac{1}{m}\sum_{\ell=1}^{m} \gamma_\ell H_\ell$ of XZ-type we define the following energy test:*

1. *The verifier picks a term H_ℓ for $\ell \in [m]$ taken uniformly at random, and selects uniformly at random from the pairs $\{(W, r) \in \{X, Z\}^n \times \{0, 1\}^n : \sigma_W(r) = H_\ell\}$.*
2. *The verifier sends $W(r)$ to Alice, and tells Bob that the players are playing the energy test.*
3. *Alice responds with a single value $c \in \{-1, 1\}$ and Bob responds with $2n$ bits $a_1, \ldots a_n, b_1, \ldots b_n$.*
4. *The verifier next computes bit string d as follows. Take $d_i = (-1)^{a_i}$ if $r_i = 1$ and $W_i = X$, take $d_i = (-1)^{b_i}$ if $r_i = 1$ and $W_i = Z$, and take $d_i = 0$ in all other cases.*
5. *The verifier accepts if $c \cdot \prod_i d_i \neq sign(\gamma_l)$, and rejects with probability $|\gamma_l|$ otherwise.*

Combining the LWPBT and Energy test we define our modified Hamiltonian test:

▶ **Definition 22** (Hamiltonian test)**.** *Let $H = \frac{1}{m}\sum_{\ell=1}^{m} \gamma_\ell H_\ell$ be a k-local Hamiltonian of XZ-type and let $p \in (0, 1)$. We define the following game $\mathcal{G}(H, p)$: with probability $(1 - p)$ the players play LWPBT introduced in Section Section 3, and with probability p the players play energy test described in Definition 21.*

We refer to a strategy $\mathcal{S}$ for $\mathcal{G}(H, p)$ as a *semi-honest strategy* if the players employ the canonical perfect strategy when playing LWPBT. Hence in a semi-honest strategy Alice and Bob hold n EPR pairs and Alice must perform $\sigma_W(r)$ on question $W(r)$ since she cannot distinguish questions from LWPBT or energy test. We also define the *honest strategy* $\mathcal{S}_h$ for $\mathcal{G}(H, p)$ in which the players employ the canonical perfect strategy when playing LWPBT, and in the energy test, Bob honestly teleports the ground state of H to Alice and provides the verifier with the teleportation keys.

Below we analyze the players' ability to win the overall game $\mathcal{G}(H,p)$ assuming the players are using a semi-honest strategy.

▶ **Lemma 23** (Lower bound on semi-honest strategies). *Suppose $H = \frac{1}{m}\sum_{l=1}^{m}\gamma_l H_l$ is an n-qubit 6-local XZ Hamiltonian, and Alice and Bob are performing a semi-honest strategy $\mathcal{S}$ for $\mathcal{G}(H,p)$. Then*

$$\omega(\mathcal{S}) \leq \omega(\mathcal{S}_h) = 1 - p(\frac{1}{2m}\sum_{l}|\gamma_l| + \frac{1}{2}\lambda_0(H)).$$

Proof. Suppose the players are employing a semi-honest strategy $\mathcal{S} = (\tau^A, \tau^B, |\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B)$ for $\mathcal{G}(H,p)$. Let $\tau := \mathrm{Tr}_{\mathcal{H}_B}(|\psi\rangle\langle\psi|)$. Since the players win the LWPBT perfectly, they can only lose the overall game if they are playing an instance of the energy test. Let $a, b \in \{0,1\}^n$ be the answers Bob provides in the energy test, and let $\rho := \sigma_X^b\sigma_Z^a\tau\sigma_Z^a\sigma_X^b$.

Suppose in a round of the energy test, the verifier picks an $\ell \in [m]$ and selects a $W(r)$ for Alice. As discussed above, since Alice cannot distinguish questions from LWPBT and the energy test, she must perform $\sigma_W(r) = H_\ell$ on her registers. Hence $\mathbb{E}(c\cdot\prod_i d_i) = \mathrm{Tr}(H_\ell\sigma_X^b\sigma_Z^a\tau\sigma_Z^a\sigma_X^b) = \mathrm{Tr}(H_\ell\rho)$. Let p_ℓ be the probability of $c\prod_i d_i = sign(\gamma_\ell)$. Then $\mathbb{E}(c\prod_i d_i) = p_\ell sign(\gamma_\ell) - (1-p_\ell)sign(\gamma_\ell)$, or in other words, $\gamma_\ell\mathbb{E}(c\prod_i d_i) = (2p_\ell - 1)|\gamma_\ell|$. This implies the verifier rejects with probability

$$p_\ell|\gamma_\ell| = \frac{|\gamma_\ell| + \gamma_\ell\mathbb{E}(c\prod_i d_i)}{2} = \frac{|\gamma_\ell| + \gamma_\ell\mathrm{Tr}(H_\ell\rho)}{2}.$$

Thus by averaging over $\ell \in [m]$ we see that the players lose the energy test with probability

$$\frac{1}{m}\sum_{\ell\in[m]}\frac{|\gamma_\ell| + \gamma_\ell\mathrm{Tr}(H_\ell\rho)}{2} = \frac{1}{2m}\sum_{\ell\in[m]}|\gamma_\ell| + \frac{1}{2}\mathrm{Tr}(H\rho).$$

This probability is minimized if and only if ρ is indeed the density matrix of the ground state of H and in such case, the probability of winning the overall game is at most

$$1 - p\big(\frac{1}{2m}\sum_{\ell\in[m]}|\gamma_\ell| + \frac{1}{2}\lambda_0(H)\big).$$

This probability can be achieved if Bob teleports over the ground state and supplies the verifier with the verification keys in the energy test. ◀

▶ **Lemma 24** (Upper bound on dishonest strategies). *Let $H = \frac{1}{m}\sum_{\ell=1}^{m}\gamma_\ell H_\ell$ be a 6-local, n-qubit Hamiltonian of XZ-type. For any $\eta \in (0,1)$, let $p = \frac{4n^{-6}\eta^{3/4}}{(C_{\mathrm{lw}}+1)3^{3/4}}$ where C_{lw} is the constant given in Theorem 18. Then $\omega^*\big(\mathcal{G}(H,p)\big) \leq \omega(\mathcal{S}_h) + \eta$, where $\omega(\mathcal{S}_h) = 1 - p\big(\frac{1}{2m}\sum_{\ell\in[m]}|\gamma_\ell| + \frac{1}{2}\lambda_0(H)\big)$ as in Lemma 24.*

Proof. Suppose the provers are employing a strategy $\mathcal{S} = (\tau^A, \tau^B, |\psi\rangle)$ for $\mathcal{G}(H,p)$ that wins LWPBT with probability $1-\varepsilon$ and wins the energy test with probability $\delta + 1 - \frac{\sum_\ell|\gamma_\ell|}{2m} - \frac{\lambda_0(H)}{2}$. Theorem 18 implies $\delta \leq C_{\mathrm{lw}}n^6\varepsilon^{1/4}$. Then for $p := \frac{4n^{-6}\eta^{3/4}}{(C_{\mathrm{lw}}+1)3^{3/4}}$ with $\eta \in (0,1)$, we have $\eta + \varepsilon = \eta/3 + \eta/3 + \eta/3 + \varepsilon \geq 4(\frac{\eta^3\varepsilon}{3^3})^{1/4} = p(C_{\mathrm{lw}}+1)n^6\varepsilon^{1/4}$. It follows that

$$p\delta - (1-p)\varepsilon \leq pCn^6\varepsilon^{1/4} + p\varepsilon - \varepsilon \leq pC_{\mathrm{lw}}n^6\varepsilon^{1/4} + pn^6\varepsilon^{1/4} - \varepsilon = p(C_{\mathrm{lw}}+1)n^6\varepsilon^{1/4} - \varepsilon \leq \eta.$$

Hence the overall winning probability is given by

$$\omega(\mathcal{S}) = (1-p)(1-\varepsilon) + p(\delta + 1 - \frac{\sum_\ell|\gamma_\ell|}{2m} - \frac{\lambda_0(H)}{2}) = \omega(\mathcal{S}_h) + p\delta - (1-p)\varepsilon \leq \omega(\mathcal{S}_h) + \eta.$$

This completes the proof. ◀

In the rest of this paper, given an n-qubit, 6-local Hamiltonian H of XZ-type and parameters α and β with $\beta - \alpha \geq 1/poly(n)$, we use $\mathcal{G}(H)$ to denote the game $\mathcal{G}(H,p)$ with $p = \frac{32n^{-6}(\beta-\alpha)^{24}}{27(C_{\text{lw}}+1)^4}$.

▶ **Theorem 25.** *Given an n-qubit, 6-local Hamiltonian $H = \frac{1}{m}\sum_{\ell=1}^{m}\gamma_\ell H_\ell$ and parameters α, β with $\beta - \alpha \geq 1/poly(n)$, let ω_α (resp. ω_β) denote the maximum winning probability for $\mathcal{G}(H)$ when $\lambda_0(H) \leq \alpha$ (resp. $\lambda_0(H) \geq \beta$). Then $\omega_\alpha - \omega_\beta \geq 1/poly(n)$.*

Proof. Let $\eta := \frac{16(\beta-\alpha)^{32}}{27(C_{\text{lw}}+1)^4}$, and let $p := \frac{4n^{-6}\eta^{3/4}}{(C_{\text{lw}}+1)3^{3/4}}$. Then $p = \frac{32n^{-6}(\beta-\alpha)^{24}}{27(C_{\text{lw}}+1)^4}$, and hence $\mathcal{G}(H) = \mathcal{G}(H,p)$. By Lemma 23 and Lemma 24 we have $\omega(\mathcal{S}_h) \leq \omega^*\big(\mathcal{G}(H,p)\big) \leq \omega(\mathcal{S}_h) + \eta$. This implies $\omega_\alpha \geq 1 - p(\frac{1}{2m}\sum_\ell|\gamma_\ell| + \frac{1}{2}\alpha)$ and $\omega_\beta \leq 1 - p(\frac{1}{2m}\sum_\ell|\gamma_\ell| + \frac{1}{2}\beta) + \eta$. Since $n^6(\beta-\alpha)^7 \leq O(n^{-1})$, it follows that

$$\omega_\alpha - \omega_\beta \geq \frac{1}{2}p(\beta-\alpha) - \eta = \frac{16n^{-6}(\beta-\alpha)^{25}}{27(C_{\text{lw}}+1)^4}(1 - n^6(\beta-\alpha)^7) \geq \frac{16n^{-6}(\beta-\alpha)^{25}}{27(C_{\text{lw}}+1)^4}.$$

Hence $\omega_\alpha - \omega_\beta \geq 1/poly(n)$. ◀

We apply a threshold parallel repetition theorem due to Yuen for the gap amplification. For every $n \in \mathbb{N}$, let $\mathcal{G}(H), w_\alpha$ and w_β be as in Theorem 25. By [51, Theorem 41], there exists a $poly(n)$-computable transformation, called *anchoring*, that transforms $\mathcal{G}(H)$ to a two-player game $\mathcal{G}(H)_\perp$ with winning probability $1 - \frac{1-w^*(\mathcal{G}(H))}{2}$. So $w^*(\mathcal{G}(H)_\perp) = \begin{cases} 1-\varepsilon_\alpha/2 & \text{if } \lambda_0(H) \leq \alpha \\ 1-\varepsilon_\beta/2 & \text{if } \lambda_0(H) \geq \beta \end{cases}$, where $\varepsilon_\alpha := 1 - w_\alpha$ and $\varepsilon_\beta := 1 - w_\beta$. Then by [51, Theorem 42], there is a universal constant $C > 0$ such that for all integer $m \geq 1$, and $\gamma \geq 0$, the probability that in the game $\mathcal{G}(H)_\perp^m$ the players can win more than $\big(w^*(\mathcal{G}(H)_\perp) + \gamma\big)m$ games is at most $(1 - \gamma^9/2)^{Cm}$. Take $\gamma := \frac{\varepsilon_\alpha - \varepsilon_\beta}{4}$ and $m = \max\{4\gamma^{-2}, 2C\gamma^{-9}\}$. Let $\widehat{\mathcal{G}}(H) := \mathcal{G}(H)_\perp^m$ be the m parallel repeated anchoring version of $\mathcal{G}(H)$. We show that this nonlocal game has a constant completeness-soundness gap.

▶ **Theorem 26.** *Let $\widehat{\omega}_\alpha$ (resp. $\widehat{\omega}_\beta$) be the maximum winning probability for $\widehat{\mathcal{G}}(H)$ when $\lambda_0(H) \leq \alpha$ (resp. $\lambda_0(H) \geq \beta$). Then $\widehat{\omega}_\alpha - \widehat{\omega}_\beta \geq 1/4$.*

Proof. If $\lambda_0(H) \geq \beta$, then $\widehat{\omega}_\beta \leq (1 - \frac{\gamma^9}{2})^{Cm} \leq (1 - \frac{\gamma^9}{2})^{2/\gamma^9} < e^{-1} < 1/2$. Now suppose $\lambda_0(H) \leq \alpha$. An optimal strategy S for $\mathcal{G}(H)_\perp$ has winning probability $1 - \frac{\varepsilon_\alpha}{2}$. Let X be the random variable for the number of games the strategy S^m wins. Then $X \sim \text{Binomial}(m, 1 - \frac{\varepsilon_\alpha}{2})$, so $\mathbb{E}X = m(1-\frac{\varepsilon_\alpha}{2})$ and $\text{Var}X = m\frac{\varepsilon_\alpha}{2}(1-\frac{\varepsilon_\alpha}{2})$. Since $(1-\frac{\varepsilon_\alpha}{2}) - (1-\frac{\varepsilon_\beta}{2}+\gamma) = \frac{\varepsilon_\beta-\varepsilon_\alpha}{2} - \gamma = \gamma$, we obtain that

$$Pr(X \leq (1 - \tfrac{\varepsilon_\beta}{2} + \gamma)m) \leq Pr(|X - \mathbb{E}X| \geq \gamma m) \leq \frac{m(1-\frac{\varepsilon_\alpha}{2})\frac{\varepsilon_\alpha}{2}}{(\gamma m)^2} = \frac{1}{m\gamma^2} \leq 1/4.$$

This implies $\widehat{\omega}_\alpha \geq w(S^m) = 1 - Pr(X \leq (1 - \frac{\varepsilon_\beta}{2} + \gamma)m) \geq 3/4$, so the theorem follows. ◀

5 Zero-knowledge proof system

In this section, we show that the family of games described in Definition 22 provides a statistical zero-knowledge MIP*[2, 1] protocol for QMA with inverse polynomial completeness/soundness gap.

5.1 Simulation of history states for XZ-Hamiltonians

Before we introduce our MIP* protocol and proceed to our result on zero-knowledge, we reformulate a result, originally introduced by Broadbent and Grilo [5] (Lemma 3.5), so that it is more amenable to device-independent techniques.

▶ **Theorem 27** (Simulation of history states). *For any language $L = (L_{yes}, L_{no})$ in* QMA *and $s \in \mathbb{N}$, there is a family of verification circuits $V_x^{(s)} = U_T \dots U_1$ for L that acts on a witness of size $p(|x|)$ and on $q(|x|)$ ancillary qubits such that there exists a polynomial-time deterministic algorithm $Sim_{V^{(s)}}$ that takes as input an instance $x \in L$ and a subset $S \subseteq [T + p + q]$ with $|S| \leq 3s + 2$, then outputs a classical description of an $|S|$-qubit density matrix $\rho(x, S)$ with the following properties:*

1. *If $x \in L_{yes}$, then there exists a $p(|x|)$-qubit witness ψ^s such that $V_x^{(s)}$ accept with probability at least $1 - negl(n)$ on ψ^s and $\|\rho(x, S) - \mathrm{Tr}_{\overline{S}}(\rho)\|_{tr} \leq negl(|x|)$, where*

$$\rho = \frac{1}{T+1} \sum_{t,t' \in [T+1]} |unary(t)\rangle \langle unary(t)| \otimes U_t \dots U_1(\psi^s \otimes |0\rangle \langle 0|^{\otimes q}) U_1^* \dots U_{t'}^*$$

   is the history state of $V_x^{(s)}$ on witness ψ^s.
2. *Let H_i be one of the terms from the circuit-to-local Hamiltonian construction from $V_x^{(s)}$, and let S_i be the set of qubits on which H_i acts non-trivially. Then $\mathrm{Tr}(H_i \rho(x, S_i)) = 0$ for all $x \in L$.*
3. *The Hamiltonian H from the circuit-to-local Hamiltonian construction is a 6-local Hamiltonian of XZ type.*

The first two points were proven by Broadbent and Grilo using simulatable codes constructed from a different set of physical gates [5]. The last point follows from a similar approach in [25, Lemma 22]. A detailed proof can be found in the full version of our paper [9].

Below we only need to invoke Theorem 27 for the case of $s = 2$ in order to our zero-knowledge protocol. We use V_x to denote $V_x^{(2)}$ throughout the rest of this section.

5.2 A two prover zero-knowledge proof system for QMA

Let $L = (L_{yes}, L_{no})$ be a language in QMA. Figure 5 describes a two-prover one-round interactive proof system for L with a constant polynomial completeness-soundness gap.

$$x \xrightarrow{\text{Theorem 27}} V_x \xrightarrow{\text{circuit-to-Hamiltonian}} H_x \xrightarrow{\text{Definition 22}} \mathcal{G}(H_x) \xrightarrow{\text{Theorem 26}} \widehat{\mathcal{G}}_x := \widehat{\mathcal{G}}(H_x)$$

Figure 5 x is an instance in $L \in$ QMA. V_x is a $poly(|x|)$-size quantum circuit. H_x is a $poly(|x|)$-qubit 6-local Hamiltonian of XZ-type. $\widehat{\mathcal{G}}_x$ is a nonlocal game with $poly(|x|)$-bit questions and $poly(|x|)$-bit answers.

To prove the above interactive proof system for L has the statistical zero-knowledge property, we first establish that any malicious verifier $\widehat{V}$ and $x \in L_{yes}$, there exists a PPT simulator that can sample from $View(\widehat{V}(x), \mathcal{S}_h)$, where $\mathcal{S}_h$ is the honest strategy for $\mathcal{G}_x$ defined in Section 4.

▶ **Lemma 28.** *Suppose $x \in L_{yes}$ for some language $L = (L_{yes}, L_{no})$ in* QMA. *Let $\mathcal{G}_x$ be the corresponding nonlocal game described in Figure 5, and let $\mathcal{S}_h$ be the honest strategy for $\mathcal{G}_x$ defined in Section 4. For any malicious verifier $\widehat{V}$ there exists a PPT algorithm $Sim_{\widehat{V}}$ with output distribution negligibly close to $View(\widehat{V}(x), \mathcal{S}_h)$,*

The proof of this lemma can be found in the full version of our paper [9]. All that remains is to argue that the interactive protocol described in Figure 5 based on the scaled-up game $\widehat{\mathcal{G}}_x$ is statistically zero-knowledge.

▶ **Theorem 29.** *The protocol described in Figure 5 is statistical zero-knowledge and has a constant completeness-soundness gap.*

Proof. The constant completeness-soundness gap follows directly from Theorem 26. To show the statistical zero-knowledge, we first consider the anchoring procedure for the game $\mathcal{G}_x$. We can specify an honest strategy $\mathcal{S}_{h,\perp}$ for the anchored version of $\mathcal{G}_x$ by fixing a choice of output for either player who receives question $\perp$ in the honest strategy. Then, given any malicious verifier $\widehat{V}(x)$, the simulator given in Theorem 29 can be trivially modified to sample from a distribution which is negligibly close to $View(\widehat{V}(x), \mathcal{S}_{h,\perp})$.

In the case of the threshold parallel repeated game $\widehat{\mathcal{G}}_x$, the honest strategy $\mathcal{S}^m_{h,\perp}$ is taken to be the m-fold product of the honest strategy $\mathcal{S}_{h,\perp}$. Then, as commented in [23], since the protocol only queries each player once, a new simulator can be obtained by sampling according to the m-fold product of the simulator used in the above lemma. ◀

6 Off-the-shelf model

6.1 Formal description of the model

Here we provide a formal description of the OTS model. This model is defined as a refinement of MIP*, where the completeness condition is weakened, allowing only one of the provers to be "all-powerful", while the other has limited functionality determined independently of the problem instance.

Off-the-shelf device. We first formalize the definition of a family of off-the-shelf devices. A *verification device* $D = (|\psi\rangle, \{P^1_a\}_a, \ldots, \{P^q_a\}_a)$ consists of a state $|\psi\rangle$ on Hilbert spaces $\mathcal{K}_A \otimes \mathcal{K}_B$ and a collection of POVMs $\{P^1_a\}_a, \ldots, \{P^q_a\}_a$ on $\mathcal{K}_A$. We say that a quantum strategy $\mathcal{S} = (\{E^x_a\}, \{F^y_b\}, |\phi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B)$ can be *implemented* using D, if three conditions hold: (i) $\mathcal{H}_A = \mathcal{K}_A$ and $\mathcal{H}_B = \mathcal{K}_B \otimes \mathcal{K}_{B'}$ for some Hilbert space $\mathcal{K}_{B'}$. (ii) The set of measurements in $\mathcal{S}$ which act on $\mathcal{H}_A$ are a contained in D. (iii) The shared state $|\phi\rangle$ in $\mathcal{S}$ can be decomposed as $|\phi\rangle = |\psi\rangle \otimes |\phi'\rangle$ where $|\psi\rangle$ is the state in D and $|\phi'\rangle$ is some auxiliary state held on a Hilbert space $\mathcal{K}_{B'}$.

Given a collection of verification devices $\{D_n\}_{n\in\mathbb{N}}$, where each D_n consists of a state $|\psi_n\rangle$ and a sequence of POVMs, we say $\{D_n\}_{n\in\mathbb{N}}$ is an *efficient family of off-the-shelf devices* if there exists a polynomial-time uniform family of quantum circuits $\{Q_n\}_{n\in\mathbb{N}}$ satisfying the following: Q_n generates the state $|\psi_n\rangle$ from an all 0 state, and on input i measures $|\psi_n\rangle$ using the i-th POVM from D_n.

▶ **Definition 30.** *A promise language $L = (L_{yes}, L_{no})$ has an off-the-shelf (OTS) proof system if there exists an efficient family of off-the-shelf devices $\{D_n\}_{n\in\mathbb{N}}$, and a polynomial-time computable function that takes an instance x to the description of a non-local game $\mathcal{G}_x$ satisfying the following:*

1. ***Completeness using OTS devices.*** *For any $x \in L_{yes}$ with $|x| \leq n$, there exists a quantum strategy $\mathcal{S}_x$, which can be implemented using D_n, obtaining $\omega(\mathcal{G}_x, \mathcal{S}_x) \geq c$.*
2. ***Soundness.*** *For any $x \in L_{no}$ we have $\omega^*(\mathcal{G}_x) < s$.*

We use OTS *to denote the class of all languages L which admits an OTS proof system with a constant completeness-soundness gap.*

1. **Set-up:** The client sends a set-up parameter $k \in \mathbb{N}$ to the server who provides a verification device D_k from an efficient family of off-the-shelf devices $\{D_n\}_n$.
2. **Choice of computation:** The client sends a classical description of circuit Q, satisfying $|Q| \leq k$ to the server.
3. **Verifiable delegation:** The client plays a 1-round game $\widehat{\mathcal{G}}_Q$, using the server and device D_k as players. The client accepts if and only if the game is won.

Figure 6 A delegation protocol between a polynomial-time classical client and polynomial-time quantum server, who provides an untrusted verification device during set-up.

Any OTS proof system is described as a special instance of a 2-player, 1-round MIP^* proof system with additional constraints regarding the completeness condition and we say that an OTS proof system is statistically zero-knowledge if it is statistically zero-knowledge as an MIP^* proof system (see Definition 13 for details).

6.2 Applications to ZK and delegated computation

In this section, we show that any language in QMA admits a statistical zero-knowledge OTS proof system. We also consider how the OTS model can be scaled down to provide a protocol for verifiable delegated quantum computation.

▶ **Theorem 31.** *For every language L in QMA, there exists a statistical zero-knowledge OTS proof system for L with constant completeness and soundness gap.*

Proof. We will be working with the proof system sending an instance x to game $\widehat{\mathcal{G}}_x$, as described in Figure 5. Using the rigidity results in Section 3, we have already shown completeness and soundness of properties of the individual game $\widehat{\mathcal{G}}_x$ in Section 4. The ZK property of this game has also been shown in Section 5. All that remains to show is that this protocol further satisfies the extra restrictions of completeness using OTS devices outlined in Definition 30. That is, we need to show that there exists an efficient family of OTS devices $\{D_n\}$ which can implement the honest strategy $\mathcal{S}_x$ for all yes instances x.

For each $L \in \mathsf{QMA}$, there exists a polynomial f such that, for all $x \in L$ of size $|x| = n$ the corresponding Hamiltonian H_x is supported on at most $f(n)$ qubits. Next suppose $x \in L_{yes}$ with $|x| \leq n$. In the honest strategy for the game $\mathcal{G}_x$, Alice and Bob share at most $f(n)$-EPR pairs, additionally, Bob privately holds a ground state ρ for H_x. The measurements required by Alice always correspond to σ_X or σ_Z on up to 6 qubits of the shared EPR pairs, or $\sigma_X\sigma_Z \otimes \sigma_Z\sigma_X$ on two qubits. In the honest strategy $\mathcal{S}_x$ for the m-fold parallel repeated anchoring game $\widehat{\mathcal{G}}_x$, the players share $mf(n)$-EPR pairs and Alice's measures in σ_X or σ_Z on up to $6m$ qubits or measures with $\sigma_X\sigma_Z \otimes \sigma_Z\sigma_X$ on $2m$ qubits. Since $m = poly(n)$, we then satisfy the completeness condition required by specifying an efficient family of OTS devices $\{D_n\}$, where for each n the verification device D_n contains $mf(n)$-EPR pairs and all of the above required Pauli measurements on up to $6m$ qubits. ◀

In our application towards delegated quantum computation, we consider a novel type of interactive protocol, where in addition to exchanging classical messages, the server can send an untrusted verification device, as defined in Section 6.1, to the client (see Section 1). In Figure 6, we consider the case where in the first "message", called a set-up stage, the prover sends an untrusted verification device, which is followed by classical communication.

▶ **Theorem 32.** *For every language L in* BQP*, there is a statistical-zero-knowledge delegation protocol as outlined in Figure 6 for L with constant completeness and soundness gap.*

Proof (Sketch). We can view the BQP-complete problem from Definition 9 as a language in QMA. This allows us to apply the efficient mapping outlined in Figure 5 to obtain a corresponding game $\widehat{\mathcal{G}}_Q$. In this case, the ground state of the underlying Hamiltonian can be prepared by a polynomial-time quantum prover. Thus, as in the proof of Theorem 31, we can define the required polynomial-time uniform family of OTS devices $\{D_n\}_{n\in\mathbb{N}}$ by taking D_n to contain suitably many EPR pairs, as well as the required Pauli measurements. Since furthermore the required ground state can always be prepared by a polynomial-time quantum prover, an honest server can obtain the required completeness in Step 3 by generating this state and teleporting it to the verification device when required. We also have that the above delegation protocol inherits the ZK property via the results of Theorem 31. ◀

References

1 Dorit Aharonov, Michael Ben-Or, Elad Eban, and Urmila Mahadev. Interactive proofs for quantum computations, 2017. arXiv:1704.04487.

2 Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, and Phillip Rogaway. Everything provable is provable in zero-knowledge. In *Advances in Cryptology – CRYPTO '88*, pages 37–56, 1988. doi:10.1007/0-387-34799-2_4.

3 Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Widgerson. Multi-prover interactive proofs: how to remove intractability assumptions. In *STOC '88: Proceedings of the twentieth ACM symposium on Theory of computing*, pages 113–131, 1988. doi:10.1145/62212.62223.

4 Anne Broadbent. How to verify a quantum computation. *Theory of Computing*, 14(1):1–37, 2018. doi:10.4086/toc.2018.v014a011.

5 Anne Broadbent and Alex B. Grilo. QMA-hardness of consistency of local density matrices with applications to quantum zero-knowledge. *SIAM Journal on Computing*, 51(4):1400–1450, 2022. doi:10.1137/21M140729X.

6 Anne Broadbent, Gus Gutoski, and Douglas Stebila. Quantum one-time programs. In *Advances in Cryptology – CRYPTO 2013*, volume 2, pages 344–360, 2013. doi:10.1007/978-3-642-40084-1_20.

7 Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for QMA. In *2016 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2016)*, pages 31–40, 2016. doi:10.1109/FOCS.2016.13.

8 Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for QMA. *SIAM Journal on Computing*, 49(2):245–283, 2020. doi:10.1137/18M1193530.

9 Anne Broadbent, Arthur Mehta, and Yuming Zhao. Quantum delegation with an off-the-shelf device, 2023. arXiv:2304.03448.

10 André Chailloux and Anthony Leverrier. Relativistic (or 2-prover 1-round) zero-knowledge protocol for NP secure against quantum adversaries. In *Advances in Cryptology – EUROCRYPT 2017*, volume 3, pages 369–396, 2017. doi:10.1007/978-3-319-56617-7_13.

11 Rui Chao, Ben W. Reichardt, Chris Sutherland, and Thomas Vidick. Test for a large amount of entanglement, using few measurements. *Quantum*, 2:92, 2018. doi:10.22331/q-2018-09-03-92.

12 Alessandro Chiesa, Michael Forbes, Tom Gur, and Nicholas Spooner. Spatial isolation implies zero knowledge even in a quantum world. In *2018 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2018)*, pages 755–765, 2018. doi:10.1109/FOCS.2018.00077.

13 John F. Clauser, Michael A. Horne., Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880–884, 1969. doi:10.1103/PhysRevLett.23.880.

14 Richard Cleve, Peter Høyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *19th Annual Conference on Computational Complexity (CCC 2004)*, pages 236–249, 2004. doi:10.1109/CCC.2004.1313847.

15 Andrea Coladangelo. Parallel self-testing of (tilted) EPR pairs via copies of (tilted) CHSH and the magic square game. *Quantum Information & Computation*, 17(9&10):831–865, 2017. doi:10.26421/QIC17.9-10-6.

16 Andrea Coladangelo, Alex B Grilo, Stacey Jeffery, and Thomas Vidick. Verifier-on-a-leash: New schemes for verifiable delegated quantum computation, with quasilinear resources. In *Advances in Cryptology — EUROCRYPT 2019*, volume 3, pages 247–277, 2019. doi:10.1007/978-3-030-17659-4_9.

17 Claude Crépeau and John Stuart. Zero-knowledge MIPs using homomorphic commitment schemes, 2023. arXiv:2304.09784.

18 Joseph F. Fitzsimons and Elham Kashefi. Unconditionally verifiable blind quantum computation. *Physical Review A*, 96(1):012303, 2017. doi:10.1103/PhysRevA.96.012303.

19 Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(3):690–728, 1991. doi:10.1145/116825.116852.

20 Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989. doi:10.1137/0218012.

21 W. T. Gowers and O. Hatami. Inverse and stability theorems for approximate representations of finite groups. *Sbornik: Mathematics*, 208(12):1784–1817, 2017. doi:10.1070/sm8872.

22 Alex 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)*, pages 28:1–28:13, 2019. doi:10.4230/LIPIcs.ICALP.2019.28.

23 Alex B. Grilo, William Slofstra, and Henry Yuen. Perfect zero knowledge for quantum multiprover interactive proofs. In *2019 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2019)*, pages 611–635, 2019. doi:10.1109/FOCS.2019.00044.

24 Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Journal of the ACM*, 58(6):30, 2011. doi:10.1145/2049697.2049704.

25 Zhengfeng Ji. Classical verification of quantum proofs. In *STOC 2016: Proceedings of the 48th ACM SIGACT symposium on Theory of Computing*, pages 885–898, 2016. doi:10.1145/2897518.2897634.

26 Zhengfeng Ji. Compression of quantum multi-prover interactive proofs. In *STOC 2017: Proceedings of the 49th ACM SIGACT symposium on Theory of Computing*, pages 289–302, 2017. doi:10.1145/3055399.3055441.

27 Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*=RE, 2020. arXiv:2001.04383.

28 Joe Kilian. Founding cryptography on oblivious transfer. In *STOC '88: Proceedings of the twentieth ACM symposium on Theory of computing*, pages 20–31, 1988. doi:10.1145/62212.62215.

29 Alexei Yu. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 2002. doi:10.1090/gsm/047.

30 Hirotada Kobayashi and Keiji Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *Journal of Computer and System Sciences*, 66(3):429–450, 2003. doi:10.1016/S0022-0000(03)00035-7.

31 Urmila Mahadev. Classical verification of quantum computations. In *2018 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2018)*, pages 259–267, 2018. doi:10.1109/FOCS.2018.00033.

32 Laura Mančinska, Jitendra Prakash, and Christopher Schafhauser. Constant-sized robust self-tests for states and measurements of unbounded dimension, 2021. arXiv:2103.01729.

33 Dominic Mayers and Andrew Yao. Self testing quantum apparatus. *Quantum Information & Computation*, 4(4):273–286, 2004. doi:10.26421/QIC4.4-4.

34 M. McKague, T. H. Yang, and V. Scarani. Robust self-testing of the singlet. *Journal of Physics A*, 45(45):455304, 2012. doi:10.1088/1751-8113/45/45/455304.

35 Matthew McKague. Self-testing in parallel with CHSH. *Quantum*, 1:1, 2017. doi:10.22331/q-2017-04-25-1.

36 N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical Review Letters*, 65(27):3373–3376, 1990. doi:10.1103/PhysRevLett.65.3373.

37 Anand Natarajan and Thomas Vidick. A quantum linearity test for robustly verifying entanglement. In *STOC 2017: Proceedings of the 49th ACM SIGACT symposium on Theory of Computing*, pages 1003–1015, 2017. doi:10.1145/3055399.3055468.

38 Anand Natarajan and Thomas Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP for QMA. In *2018 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2018)*, pages 731–742, 2018. doi:10.1109/FOCS.2018.00075.

39 Michael A. Nielsen and Issac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

40 Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151(3):107–108, 1990. doi:10.1016/0375-9601(90)90172-K.

41 Ben W. Reichardt, Falk Unger, and Umesh Vazirani. Classical command of quantum systems. *Nature*, 496:456–460, 2013. doi:10.1038/nature12035.

42 Gregory Rosenthal and Henry Yuen. Interactive proofs for synthesizing quantum states and unitaries. In *13th Conference on Innovations in Theoretical Computer Science—ITCS 2022*, pages 112:1–112:4, 2022. doi:10.4230/LIPIcs.ITCS.2022.112.

43 William Slofstra. The set of quantum correlations is not closed. *Forum of Mathematics, Pi*, 7:e1:1–e1:41, 2019. doi:10.1017/fmp.2018.3.

44 William Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 33:1–56, 2020. doi:10.1090/jams/929.

45 Boris S. Tsirelson. Some results and problems on quantum Bell-type inequalities. *Hadronic Journal Supplement*, 8:329–345, 1993.

46 Thomas Vidick. An expository note on "a quantum linearity test for robustly verifying entanglement", 2018. URL: http://users.cms.caltech.edu/~vidick/notes/pauli_braiding_1.pdf.

47 Thomas Vidick and Tina Zhang. Classical zero-knowledge arguments for quantum computations. *Quantum*, 4:266, 2020. doi:10.22331/q-2020-05-14-266.

48 John Watrous. Succinct quantum proofs for properties of finite groups. In *41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*, pages 537–546, 2000. doi:10.1109/SFCS.2000.892141.

49 John Watrous. PSPACE has constant-round quantum interactive proof systems. *Theoretical Computer Science*, 292(3):575–588, 2003. doi:10.1016/S0304-3975(01)00375-9.

50 John Watrous. Quantum computational complexity. In *Encyclopedia of complexity and systems science*, pages 7174–7201. Springer, 2009. doi:10.1007/978-3-642-27737-5_428-3.

51 Henry Yuen. *Games, protocols, and quantum entanglement*. PhD thesis, Massachusetts Institute of Technology, 2016. URL: https://dspace.mit.edu/handle/1721.1/107364.