Quantum Free Games

Anand Natarajan
anandn@mit.edu
MIT
Cambridge, MA, USA

Tina Zhang
tinaz@mit.edu
MIT
Cambridge, MA, USA

ABSTRACT

The complexity of free games with two or more classical players was essentially settled by Aaronson, Impagliazzo, and Moshkovitz (CCC'14). In the quantum world, there are two complexity classes that can be considered quantum analogues of classical free games: (1) AM*, the multiprover interactive proof class corresponding to free games with entangled players, and, somewhat less obviously, (2) BellQMA(2), the class of quantum Merlin-Arthur proof systems with two unentangled Merlins, whose proof states are separately measured by Arthur. In this work, we make significant progress towards a tight characterization of both of these classes. (1) We show a BellQMA(2) protocol for 3SAT on n variables, where the total amount of communication is $\tilde{O}(\sqrt{n})$. This answers an open question of Chen and Drucker (2010) and also shows, conditional on ETH, that the algorithm of Brandão, Christandl and Yard (STOC'11) for optimizing over separable states is tight up to logarithmic factors. (2) We show that AM* with $n_{\text{provers}} = 2$, question length $O(1)$, and answer-length $\text{poly}\log(n)$ is equal to RE, i.e. that free entangled games with constant-sized questions are as powerful as general entangled games. (In contrast, Aaronson, Impagliazzo and Moshkovitz show that classical free games are much weaker than general classical games.) We show this using a question "hypercompression" theorem that iteratively applies the introspection technique of Ji et al. (2020). Our result is a significant improvement over the headline result of Ji et al., whose MIP* protocol for the halting problem has $\text{poly}(n)$-sized questions and answers. (3) By the same techniques, we obtain a zero-gap AM* protocol for a Π_2 complete language with constant-size questions and almost logarithmically ($O(\log n \cdot \log^* n)$) large answers, improving on the headline result of Mousavi, Nezhadi and Yuen (STOC'22). (4) Using a connection to the nonuniform complexity of the halting problem we show that any MIP* protocol for RE requires $\Omega(\log n)$ bits of communication. It follows that our results in item 3 are optimal up to an $O(\log^* n)$ factor, and that the gapless compression theorems of Mousavi, Nezhadi and Yuen are asymptotically optimal. We conjecture that these bounds can be saturated in the gapped case as well.

CCS CONCEPTS

• **Theory of computation** → **Quantum complexity theory**; **Interactive computation**; *Turing machines*; **Complexity classes**.

STOC '23, June 20–23, 2023, Orlando, FL, USA

ACM ISBN 978-1-4503-9913-5/23/06.
https://doi.org/10.1145/3564246.3585208

KEYWORDS

nonlocal games, interactive proofs, MIP*, QMA(2), entanglement, self-testing

ACM Reference Format:
Anand Natarajan and Tina Zhang. 2023. Quantum Free Games. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3564246.3585208

1 INTRODUCTION

The 1991 work of Babai, Fortnow and Lund which showed that MIP = NEXP [5] remains one of the most important achievements of complexity theory to date. The techniques used in the proof provided a springboard toward several other important results, including the proof that PCP = NP [3, 4] and, more recently, MIP* = RE [22]. In short, history shows that the study of *multiplayer games*, in which an honest, polynomial-time verifier referees a game involving two or more potentially dishonest and unbounded provers, has yielded some of the most fruitful avenues of research in the field of complexity theory.

In a multiplayer game, the verifier, since it is computationally bounded, is inherently at a disadvantage, and must find clever ways to force the unbounded provers with which it interacts to do computations on its behalf, even though it cannot necessarily replicate those computations to check if they were done accurately. One of the most useful powers at its disposal is the ability to ask the provers correlated questions. For example, a common paradigm in multiprover proof design is the *consistency test*, in which the verifier asks one prover ('Alice') to provide an answer to some small subproblem of the overall problem which it is trying to decide, and asks another prover ('Bob') to provide *part of* the answer to the same subproblem. An example of this paradigm is the *clause-variable game*, in which the verifier—who is attempting to decide whether some instance of a constraint satisfaction problem (CSP), such as 3SAT or graph colouring, is satisfiable—asks Alice to provide a satisfying assignment to a single constraint in that CSP, and asks Bob to provide an assignment to just one of the variables participating in that constraint. For example, if the verifier is trying to decide an instance of 3SAT, then it may ask Alice for a satisfying assignment to a single clause in the 3SAT formula, and ask Bob for an assignment to one of the variables in that clause. The verifier then checks that the assignment which Alice provided is indeed a satisfying assignment, and moreover that Bob's answer was consistent with Alice's assignment. Note that a single clause in a 3SAT formula always has a satisfying assignment, even if the entire formula does not. The key element that prevents Alice and Bob from exploiting this to convince the verifier that an unsatisfiable formula is satisfiable is the fact that they cannot communicate. Therefore, the only way Alice can be consistent with Bob is if they have agreed beforehand

to answer consistently with a *global* satisfying assignment to the 3SAT formula. Otherwise, if Alice reports assignments that depend on the clause which she was given by the verifier, she will not be consistent with Bob, because Bob does not know which clause she was given.

The crucial advantage which the consistency test paradigm enjoys over single-prover NP-style verification, in which the verifier simply asks for an assignment which satisfies all the clauses in an instance of 3SAT and checks that this is the case, is that a consistency test allows the verifier to efficiently check the satisfiability of CSPs which have far more clauses than the verifier could efficiently read, if it were to read all of them. Specifically, in only polynomial time, the consistency test paradigm allows a polynomially-bounded verifier to check whether an *exponentially* long CSP is satisfiable or far from satisfiable, since the verifier only needs to send polynomially long questions to Alice and Bob in order to specify which constraint and which variable it wants to know about in an exponentially long CSP. This is the basic but necessary verification framework at the heart of results such as MIP = NEXP.

We may then ask: what if we take away this power of the verifier, essential to the design of many multiprover interactive proof systems, which allows it to ask the provers correlated questions (and therefore perform tests like the consistency test)? How is the computational power of the multiplayer game model altered if we demand that the verifier's questions to the provers be *independently sampled*? The model in which a computationally efficient verifier referees a game involving two or more potentially dishonest and unbounded provers, and the verifier's questions to the provers must consist of independently sampled and uniformly random bits, is known as the *free game model*. In this paper we will be primarily concerned with the free game model.

The power of classical free games was considered in 2014 by Aaronson, Impagliazzo and Moshkovitz [2]. They defined the complexity class $\mathsf{AM}(k)$, which is exactly the class of problems that can be decided by a computationally efficient verifier interacting with k potentially dishonest and unbounded provers, under the restriction that the verifier's questions to the provers must consist of independently sampled and uniformly random bits. Aaronson, Impagliazzo and Moshkovitz also showed that, for any $k = \mathrm{poly}(n)$, where n is the length of the verifier's input, $\mathsf{AM}(k) = \mathsf{AM}$.

This tells us that, relative to the full power of classical multiplayer games, classical free games are very weak. Babai, Fortnow and Lund [5] showed that $\mathsf{MIP}(2)$ contains NEXP, where $\mathsf{MIP}(2)$ corresponds to the 'unrestrained' multiplayer two-player game model; on the other hand, Aaronson, Impagliazzo and Moshkovitz showed that, when we force the verifier to ask independently sampled questions, it can only decide problems in AM—a class which is equal to NP under plausible complexity-theoretic assumptions [23]. That is, in the classical world, placing the free-game restriction on the verifier seems to result in an exponential decrease in its deciding power!

Intuitively, we can understand this relationship as follows. The best paradigm known for converting $\mathsf{MIP}(2)$ protocols into $\mathsf{AM}(2)$ protocols is what we will call *birthday repetition*. Suppose that there is a one-round $\mathsf{MIP}(2)$ protocol P with constant completeness-soundness gap which allows the verifier to decide some language L of interest, and in which the verifier samples correlated questions (x, y) for the two provers from a set $\mathcal{X} \times \mathcal{Y}$. For simplicity, let us suppose that the verifier samples (x, y) uniformly at random from a set $\mathcal{S} \subseteq \mathcal{X} \times \mathcal{Y}$. (This is true, for instance, in the clause-variable example we considered earlier.) We produce a free version of this protocol, P_{free}, by simply having the free verifier sample k questions $(x_1, \ldots, x_k)$ from $\mathcal{X}$, and ℓ questions $(y_1, \ldots, y_\ell)$ from $\mathcal{Y}$, independently at random. The verifier then checks whether there exists (x_i, y_j) for $i \in [k], j \in [\ell]$ such that $(x_i, y_j) \in \mathcal{S}$. If there exists such a pair (x_i, y_j), then the free verifier acts as the MIP verifier would; otherwise, it automatically accepts.

We can represent the set $\mathcal{X} \times \mathcal{Y}$ as a bipartite graph $\mathcal{G}$, with a vertex corresponding to every $x \in \mathcal{X}$ on the left, and a vertex corresponding to every $y \in \mathcal{Y}$ on the right. We can also imagine that there is an edge between x and y if and only if $(x, y) \in \mathcal{S}$. If we assume that every vertex in $\mathcal{G}$ participates in at least one edge (i.e. that every Alice question $x \in \mathcal{X}$ has a nonzero chance of being asked, and similarly for every Bob question), the probability that $(x, y) \in \mathcal{S}$ for x chosen uniformly at random from $\mathcal{X}$ and y chosen independently and uniformly from $\mathcal{Y}$ is at least $\Omega\left(\frac{|\mathcal{X}|+|\mathcal{Y}|}{|\mathcal{X}||\mathcal{Y}|}\right)$. In the case where $|\mathcal{X}| \approx |\mathcal{Y}| := N$, this probability is $\Omega(1/N)$. Therefore, using the birthday paradox, we expect to set $k \approx \ell \approx O(\sqrt{N})$ in order to ensure that P_{free}, our free version of the MIP protocol P, still has constant soundness. (We call this procedure to turn a non-free game into a free game *birthday repetition* because of the link to the birthday paradox.)

For the MIP protocols which are sufficiently powerful to capture all of NEXP, as we mentioned earlier, it is typically the case that $|\mathcal{X}|$ and $|\mathcal{Y}|$ are both exponentially large in the input length. In other words, in order to convert an $\mathsf{MIP}(2)$-complete protocol into an $\mathsf{AM}(2)$ protocol using birthday repetition, we would need the verifier to send the provers questions of length 2^{cn}, where c is some constant and n is the length of the input. This is clearly computationally infeasible. Birthday repetition is only computationally feasible when the question sets $\mathcal{X}$ and $\mathcal{Y}$ are polynomially large—or, in other words, when the CSP the verifier wants to decide is only polynomially long. This brings us back into the NP setting.

However, building on a line of previous work [1, 7], Aaronson, Impagliazzo and Moshkovitz also identify something which a free verifier interacting with two noncommunicating provers can do that no polynomially-bounded verifier interacting with a single prover can: the former verifier can decide 3SAT *using only* $O(\sqrt{N} \cdot \mathrm{poly}\log N)$ *bits of communication with its provers*, where N is the number of clauses in the 3SAT instance. Assuming the Exponential Time Hypothesis (ETH), i.e. that 3SAT cannot be solved in $2^{o(N)}$ time[1], a polynomially bounded verifier interacting with a single prover cannot verify 3SAT so efficiently: any such verifier who could would lead to a subexponential-time algorithm for 3SAT. As such, though it may well be the case that $\mathsf{AM}(2) = \mathsf{NP}$, a verifier who referees a two-player free game may still have capabilities that a polynomially bounded verifier interacting with a single prover does not.

[1] Actually $2^{o(n)}$ whee n is the number of variables in the 3SAT instance, but this is linearly related to N for hard instances.

1.1 BellQMA(2)

1.1.1 Background and Previous Work. Aaronson, Impagliazzo and Moshkovitz's original motivation [2, Section 4] in studying classical free games was this latter application of deciding 3SAT in sublinear communication, the central ideas in which arose first not from the classical literature but from the study of a quantum class known as QMA(2). Informally, if NP is the class of problems which can be efficiently decided by a deterministic classical verifier who is provided with a classical witness, and QMA is the class of problems which can be efficiently decided by a quantum verifier given a quantum witness, then QMA(2) is the class of problems which can be decided by a quantum verifier given *two* quantum witnesses. In the NP world, drawing a distinction between one and two witnesses is clearly absurd: any two classical witnesses can be concatenated into one witness, and any one classical witness can be split arbitrarily into two. In the QMA(2) world, however, the distinction between one and two witnesses is given meaning by requiring that any 'two' witness states are *unentangled* (or, equivalently, that they come from two noncommunicating provers who cannot share entanglement). If unentanglement were a property that an efficient quantum verifier could check for itself, then any QMA(2) protocol could be converted into a QMA protocol; however, this is *not* known to be the case, and it remains unknown if QMA(2) = QMA. We can, of course, also define QMA(k) for $k \geq 2$, in which the quantum verifier receives k unentangled witnesses from k separate provers.

Our most compelling example of an application for the QMA(2) setup (which cannot be instantiated in the QMA setup, conditioned on the Exponential Time Hypothesis) is deciding NP problems in sublinear communication. That is, we know of a QMA(2) protocol in which each of the two provers sends only $O(\sqrt{N} \cdot \log N)$ qubits to the quantum verifier, where N is the number of clauses in a 3SAT formula ϕ, and that verifier can subsequently decide ϕ with constant probability of error. (Note that, in the QMA(2) model, the quantum verifier does not send any challenges to the two provers, unlike in the AM(k) model—all the communication in a QMA(2) protocol happens in a single quantum message from provers to verifier.) A protocol with sublinear communication to decide 3SAT was firstly proposed for the QMA($\sqrt{N}\operatorname{poly}\log N$) setup [1], in which a quantum verifier interacts with $\sqrt{N}\operatorname{poly}\log N$ separate provers, and it was subsequently shown [19] that there is a QMA(2) (two-prover) protocol achieving the same purpose with similar overall communication length.

The pervasive $\sqrt{N}$, which also appeared in the communication complexity of the [2] AM(2) protocol for the same purpose, is not a coincidence—in fact, the [2] AM(2) protocol for deciding 3SAT in sublinear communication draws close inspiration from protocols originally designed for QMA(2). The $\sqrt{N}$ factor does have some motivation: it originates from a clever application of the birthday paradox [1]. So far, nobody has thought of any other technique that might do better. It is natural, then, to wonder whether $\sqrt{N}$ qubits of communication is unavoidable. Is it the case that *any* QMA(2) protocol for 3SAT must use at least $O(\sqrt{N})$ qubits of communication?

Unfortunately, our provable communication lower bounds in this case fail to match the upper bounds exactly. The best known lower bound on the communication complexity of a QMA(2) protocol to decide 3SAT originates from [8], which shows that, if there is any QMA(2) protocol *of a certain restricted form* that can decide 3SAT with constant probability of error, then that protocol must involve at least $O(\sqrt{N})$ qubits of communication. [8] shows that any such protocol with smaller communication complexity implies a subexponential-time algorithm for 3SAT, and therefore contradicts the Exponential Time Hypothesis.

The restricted model which [8] consider in their lower bound is one in which the quantum QMA(2) verifier acts as if it consisted of two separate parties—call them Arthur and Lancelot—who each receive one of the two unentangled witnesses provided by the all-powerful provers. Arthur and Lancelot can then perform separate measurements on their respective witness states and communicate classically. After they communicate their measurement outcomes to each other classically, they are allowed to perform more measurements, and then communicate classically again, *ad infinitum*; however, they cannot perform any entangling measurements which straddle the two witness states. At the end of many rounds of separate measurements and classical communication, Arthur and Lancelot output a joint decision. This model is called the *local operations and classical communication* (LOCC) model, and the version of QMA(2) in which the verifier is restricted to behaving in this way is known as LOCC-QMA(2). [8] shows that any LOCC-QMA(2) protocol for 3SAT must use at least $O(\sqrt{N})$ qubits of communication between provers and verifier.

What do we know about upper bounds on the communication necessary to solve 3SAT in the LOCC-QMA(2) model? Can we at least get a tight characterisation of that class, if not of general QMA(2) protocols for NP? The LOCC-QMA model of course encompasses a model in which Arthur and Lancelot measure their separate witnesses exactly once, and then perform joint classical computations on the measurement results in order to determine their decision. This latter model is known as the BellQMA model. In 2010, building on work by Blier and Tapp [7], Chen and Drucker [10] exhibited a remarkably clean BellQMA version of the original [1] QMA($\sqrt{N}\operatorname{poly}\log N$) protocol for 3SAT. In the Chen-Drucker protocol, every separate quantum witness is measured separately, and the classical measurement results from these measurements are post-processed classically by the verifier in order to determine the final decision. The total communication complexity of the Chen-Drucker protocol is $O(\sqrt{N} \cdot \log N)$ qubits. The Chen-Drucker protocol would therefore appear to answer our desire for an LOCC-QMA protocol for NP whose communication complexity matches (up to $\log N$ factors) the lower bound on the communication complexity of LOCC-QMA protocols which was proven by Brandão, Christandl and Yard. Unfortunately, the Chen-Drucker protocol requires $\sqrt{N}$ unentangled provers, not only two, so it cannot show us that the communication lower bound from the [8] algorithm is optimal.[2]

Nonetheless, the Chen-Drucker protocol for NP illuminates the close connection that exists between BellQMA(k) and AM(k), a connection which may not be obvious at first glance. The key

[2] We remark that the original [1] protocol also required $\Omega(\sqrt{N})$ unentangled provers, and it was 'compiled down' to a two-prover protocol by [19]; however, the 'compilation' technique required entangling measurements across witness states. We also remark that in the multipartite setting, the Chen-Drucker protocol was proven optimal by Brandão and Harrow [9].

ingredient in the Chen-Drucker protocol is a quantum test called the *uniformity test*, which, broadly speaking, involves measuring certain registers of the witness states provided by the provers in the *Fourier basis*, and requiring the measurement outcomes to be zeroes. The zero Fourier state is the uniform superposition in the standard basis. As such, the uniformity test can (morally speaking) act as a substitute for the uniformly random challenges generated by the verifier in the $\mathsf{AM}(k)$ model, because the uniformity test in a sense *forces* the provers to generate their own uniformly sampled challenges. More specifically, if the prover provides us (the verifier) with a state of the form

$$|\psi\rangle = \sum_{q \in Q} \alpha_q |q\rangle_Q |a(q)\rangle_A \tag{1}$$

where Q is a set of questions and $a(q)$ represents an answer to a given question q, and we can somehow certify that the α_qs are all equal to each other (i.e. if we can certify that the question register Q is in a uniform superposition after we—somehow!—'disentangle' it from the answer register A), then measuring this state $|\psi\rangle$ in the standard basis is just as good as sampling a uniformly random question $q \in Q$, sending it to the prover, and receiving the prover's answer $a(q)$. Therefore, using the uniformity test to replace uniformly generated challenges, we can—sweeping all the inevitable caveats under the rug—simulate $\mathsf{AM}(k)$ protocols in $\mathsf{BellQMA}(k)$.

Reality, of course, is not quite as clean: this approach to simulating $\mathsf{AM}(k)$ in $\mathsf{BellQMA}(k)$ is not as general as our vague exposition just now made it out to be. In particular, for the uniformity test technique to work (for an *honest* strategy to exist that passes the uniformity test), it is vital that the answers $a(q)$ are short—constant sized, or at most logarithmically sized. The reason is that, in order to truly get the question register Q into the zero Fourier state, which we define as $|0\rangle_{\mathcal{F}} = \sum_{q \in Q} \frac{1}{\sqrt{|Q|}} |q\rangle$, we need to 'disentangle' it from the answer register first, and this operation involves performing a measurement on the answer register and post-selecting on a measurement outcome which occurs with negligible probability if the answers are long.

In the case where the questions q represent constraints in a CSP, however, and the answers $a(q)$ represent assignments to the variables involved in those constraints, the skies are clear. In essentially all well-studied CSPs, such as 3SAT and graph colouring, any single constraint and any assignment to a single variable in a constraint can be described in constantly many bits! As such, an $\mathsf{AM}(k)$ protocol for 3SAT in the clause-variable paradigm can indeed, at least morally, be 'compiled down' into $\mathsf{BellQMA}(k)$ in this way—which is the starting point for the Chen-Drucker protocol.

1.1.2 Our Results About $\mathsf{BellQMA}(2)$. In this work, we resolve the question of whether or not the lower bound proven by [8] is tight for $\mathsf{LOCC\text{-}QMA}(2)$ by exhibiting a $\mathsf{BellQMA}(2)$ (two-prover) protocol which has communication complexity $O(\sqrt{N} \cdot \log N)$ and decides 3SAT instances with constant probability of error. This question was raised by Chen and Drucker in 2010 [10] after they published their protocol, and raised or mentioned several times since then by others [11, Question 1] [9] [2], but despite this has remained open for more than 10 years. In resolving this question, we present an (arguably) simpler analysis of the Chen-Drucker uniformity test, as well as a more modular analysis of a $\mathsf{QMA}(2)$ protocol for 3SAT with sublinear communication than any other one we know of, which we hope may be conceptually useful.

As a consequence, we show that the runtime of the [8] algorithm (for approximating the value of a $\mathsf{LOCC\text{-}QMA}(2)$ protocol up to constant additive error) is optimal up to logarithmic factors, assuming the Exponential Time Hypothesis. This is because any improvement to their algorithm would, in combination with our protocol, result in a subexponential-time algorithm for 3SAT, which would contradict the ETH.

Our protocol is very similar to the Chen-Drucker protocol. The key changes we make are in the analysis, and these changes hinge on the observation that unentanglement is actually *not* necessary to the soundness of the uniformity test. Chen and Drucker assume that the $k = O(\sqrt{N})$ honest provers in their protocol provide the verifier with $k = O(\sqrt{N})$ unentangled copies of a state of the form in Equation (1), and they define a k-state uniformity test (implementable, of course, using separate measurements on the separate states) with the following properties:

(1) Completeness: honest provers providing k copies of a state of the form in Equation (1), with $\alpha_q = \alpha^*$ for some constant α^* for all q, will pass the k-state uniformity test with probability $1 - 2^{-\Omega(k)}$.
(2) Soundness: any k unentangled states passing the k-state uniformity test with high probability will be such that sufficiently many states among the k states have the form in Equation (1) with $\alpha_q \approx \alpha^*$ for some constant α^* for all q.

Our essential observation is as follows: it is not necessary for the input state to the k-state uniformity test to lie in k unentangled registers for a certain form of soundness, which we shall shortly define, to hold. Informally, the soundness guarantee we prove is as follows:

Lemma 1 (informal version of Lemma 11). *Let $|\psi\rangle$ be any state passing the Chen-Drucker k-state uniformity test with high probability. $|\psi\rangle$ is divided, without loss of generality, into k 'question registers' and k 'answer registers', à la Equation (1), which may be entangled. $|\psi\rangle$ is such that measuring all the question registers in the standard basis (approximately) yields a uniformly random string on sufficiently many registers and junk elsewhere.*

Lemma 1 is the main technical lemma in this part of our work. We will now explain why Lemma 1 yields a two-prover $\mathsf{BellQMA}$ protocol for 3SAT with $O(\sqrt{N} \cdot \log N)$ communication.

[2] exhibits an $\mathsf{AM}(2)$ (classical two-prover free game) protocol with $O(\sqrt{N} \cdot \log N)$ communication complexity which decides 3SAT instances with constant probability of error. This protocol (since it is inspired by the Chen-Drucker protocol) happens to be a clause-variable game which can be 'simulated' by a $\mathsf{BellQMA}$ protocol, in the way that we described at the end of Section 1.1.1. In particular, the verifier Arthur's challenge to the first prover Alice consists of k constraints in a CSP, and for each constraint Arthur expects an answer consisting of a constant-sized assignment to the variables involved in that constraint; while Arthur's challenge to the second prover Bob consists of k variables from the same CSP, and for each variable he sends to Bob, Arthur expects to receive a constant-sized assignment to that variable. Leveraging the intuition which we

described at the end of Section 1.1.1, therefore, we can 'compile' the [2] AM(2) protocol into a BellQMA(2) protocol: the honest strategy for either prover consists of providing k copies of a state of the form in Equation (1), with $\alpha_q = \alpha^*$ for some constant α^* for all q. We can then use the Chen-Drucker k-state uniformity test to enforce uniformly sampled questions. Completeness holds because the answer to any given question is constantly sized, and the form of soundness which we prove in Lemma 1 is sufficient to induce the soundness guarantees from [2].

The main technical observation which leads to the proof of Lemma 1 is as follows. The Chen-Drucker k-state uniformity test has, informally speaking, the following structure:

(1) Given an input state $|\psi\rangle$ on k question and k corresponding answer registers: measure all the answer registers in the Fourier basis. If some 'large number' of the resulting measurement outcomes were zeroes, we continue; otherwise, we reject. (We will not be precise about what 'large number' means here. For details, see Figure 2.)
(2) For every $i \in [k]$: if the ith answer register measured to zero in step 1, measure the ith question register in the Fourier basis as well. If the answer is not zero for any such i, reject; otherwise, if the answer is zero for all i such that the ith answer register measured to zero in step 1, accept.

Intuitively, this test is trying to leverage the intuition we explained at the end of Section 1.1.1 to guarantee that as many question registers as possible are in a uniform superposition. Step 1 is necessary because we must 'disentangle' the answer registers from the question registers first. We will not explain the completeness property of this test in detail, since it is analysed in [10, Section 3.1]. Instead, we will sketch how we prove Lemma 1.

Assume that we have some state $|\psi\rangle$ which passes this test with probability 1. Then the measurement in step 1 will yield a 'large' set of indices $\mathcal{S} \subseteq [k]$ such that, for all $i \in \mathcal{S}$, the ith answer register measured to zero in the Fourier basis. Denote the post-measurement state after the measurement in step 1 has been performed by ρ_1. Because $|\psi\rangle$ passes the k-state uniformity test with probability 1, we know that ρ_1 must be such that a 'large number' of its question registers are in the zero Fourier state (i.e., all the question registers of ρ_1 indexed by $i \in \mathcal{S}$ must be in the zero Fourier state—otherwise, step 2 would reject). Therefore, if we hypothetically measured the question registers of ρ_1 in the *standard* basis, we would get uniformly random outcomes on a 'large number' of the question registers, and junk elsewhere.

The key observation is that this latter hypothetical standard basis measurement on the question registers and the Fourier basis measurement which we performed in step 1 on the answer registers *commute*—because they are performed on different registers. As such, even if we do not firstly measure the answer registers of $|\psi\rangle$ as the test prescribes, and instead directly measure the question registers of $|\psi\rangle$ in the standard basis, we will get uniformly random outcomes on a 'large number' of the question registers, and junk elsewhere, just as if we had measured the question registers of ρ_1. Lemma 1 follows.

1.2 $\mathsf{AM}^*(2)$

1.2.1 Background and Previous Work. The close connection between BellQMA(2) and AM(2) which we explained at the end of Section 1.1.1 suggests that BellQMA(2) should be considered a 'quantum analogue' of AM(2). However, there is another quantum class which has equally strong claims upon the title. This is the class of problems which can be decided by a *classical* verifier who referees a free game with two unbounded provers *who are allowed to share entanglement*. Following [2], we denote this class by $\mathsf{AM}^*(2)$. As far as we know, Aaronson, Impagliazzo and Moshkovitz were the first ones to define this class [2, Section 8], and they left characterising its power relative to AM(2) as an open problem.

Studying the power of 'entangled versions' of classical multiprover classes has a long and fruitful history [12, 17, 20, 21, 29], and has recently led to some surprising and deep results [22] with connections to pure mathematics. It is not *a priori* clear whether allowing entanglement between the two provers increases or decreases the deciding power of the verifier. On the one hand, the entanglement might allow the provers to help the verifier more effectively, but on the other hand, it might also allow them to cheat more effectively. This is a familiar story: we have seen the same question of whether entanglement helps or hurts arise and be resolved several times already in the history of $\mathsf{MIP}^*(2)$ (the entangled version of MIP) and variants of that class. Ji, Natarajan, Vidick, Wright, and Yuen recently showed that $\mathsf{MIP}^*(2) = \mathsf{RE}$ [22], which clearly indicates that, in the plain multiplayer game model, allowing entanglement increases the deciding power of the verifier (from NEXP to RE!). On the other hand, it is far from a foregone conclusion that allowing entanglement makes any given multiprover proof system more powerful. For example, it is known that the entangled version of $\oplus$MIP, a version of MIP in which the verifier's decision is simply the XOR of two one-bit answers from the two provers, is inside EXP, even though $\oplus$MIP itself is equal to NEXP [31]. We can conclude from these examples only that it is not clear *a priori* how $\mathsf{AM}^*(2)$ ought to relate to AM(2).

1.2.2 Our Results About $\mathsf{AM}^*(2)$. In this work, we resolve the open question about the power of $\mathsf{AM}^*(2)$ which was posed by Aaronson, Impagliazzo and Moshkovitz, by showing that, in fact, $\mathsf{AM}^*(2) = \mathsf{MIP}^*(2) = \mathsf{RE}$. In other words, quantum free games are just as powerful as general quantum multiplayer games, even though in the classical world the free-game restriction results in a significant decrease in the verifier's deciding power!

We note that the best lower bound on $\mathsf{AM}^*(k)$ prior to our work, due to Brandão and Harrow [9, Corollary 4], was $\mathsf{NP} \subseteq \mathsf{AM}^*(\sqrt{N})$. In particular, Brandão and Harrow showed that there is an $\mathsf{AM}^*(\sqrt{N})$ protocol (analogous to the [2] AM(2) protocol) with $\sqrt{N}$ provers and $O(\sqrt{N} \cdot \log N)$ total communication that decides N-clause 3SAT with constant probability of error. Our result subsumes this result: explicitly, we show that there is an $\mathsf{MIP}^*(2)$ protocol with *constant*-sized questions, and answer sizes growing as $\mathrm{poly}\log(n)$, that is capable of deciding all of RE with constant probability of error, where n is the size of the problem instance being decided. Since free games with constant-sized questions are equivalent to general games with constant-sized questions, we obtain a very communication-efficient $\mathsf{AM}^*(2)$ protocol for RE. We prove this result by using the powerful machinery developed by

Ji, Natarajan, Vidick, Wright and Yuen which was not available in 2013 to Brandão and Harrow.

The key difference between quantum and classical free games, which allows $\mathsf{AM}^*(2) = \mathsf{MIP}^*(2)$ even though $\mathsf{AM}(2)$ and MIP are significantly different in power, is that allowing the provers to share entanglement opens up access to the tools provided by the self-testing literature, which allows us to get around the 'birthday repetition barrier' we identified in the first section of this introduction. In particular, self-testing allows us to use relatively little communication to force the two provers to *introspect*, namely to generate their own (long) questions, when they play an entangled game, and thus allows us to avoid having to send very large questions in order to achieve a constant probability of free collisions. The machinery of introspection was introduced in [27] in order to prove $\mathsf{MIP}^* \supseteq \mathsf{NEEXP}$, and is at the heart of the *compression theorems* which led to $\mathsf{MIP}^* = \mathsf{RE}$. Compression theorems are transformations that take as input some multiplayer game with long questions and answers and large verifier complexity, and output a new multiplayer game with (usually exponentially) smaller questions and answers and verifier complexity, *that has about the same value as the original game*: in particular, the fact of whether the value of the original game was $= 1$ or $\leq 1/2$ should be preserved. That compression theorems can exist at all for entangled games is testament to the marvellous power of self-testing theorems. In particular, the proof that $\mathsf{MIP}^* = \mathsf{RE}$ follows by (in a sense) recursively applying a compression theorem.

In this work we take the compression theorems that were used to prove $\mathsf{MIP}^* = \mathsf{RE}$ and 'bootstrap' them to prove that MIP^* is equal to MIP^* with constantly sized questions. Specifically, we prove what we term a *hypercompression theorem*, which is also the result of recursively applying compression theorems, but in a slightly different way from the way that appears in the proof of $\mathsf{MIP}^* = \mathsf{RE}$. Our hypercompression theorem starts with any MIP^* protocol with polynomially long questions and answers, and applies a general compression theorem once in order to turn it into a protocol with polylogarithmically long questions and answers, before recursively applying a *question reduction* theorem to bring the question size down to constant while more or less preserving the answer size. (The efficacy of this recursive application procedure is dependent on the structure of the question reduction theorem—in particular, we cannot reduce the answer size in quite the same way, for reasons related to the fact that the efficacy of answer reduction depends on the running time of the verifier in the original game while question reduction does not.) The question reduction procedure that we use is similar to the one in [22], although we believe that, by incorporating recent improvements to the analysis of question reduction made by de la Salle [15], one would be able to prove a better question reduction theorem that might be a stepping stone towards an MIP^* protocol for RE with constantly sized questions and (truly, or up to a factor of $O(\log^* n)$) logarithmically sized answers. Due to time constraints, we leave this improvement for a future version of the paper. We discuss this possibility in more detail in Section 1.5.

Two remarks about this result are in order for the benefit of the interested reader.

- We also prove the *gapless* version of this result—namely, that zero-gap MIP^* is equal to zero-gap MIP^* with *constantly* sized questions and $O(\log n \cdot \log^* n)$ sized answers. Zero-gap MIP^* is the same as normal MIP^* except that, for no-instances, the verifier's acceptance probability is only required to be strictly less than 1 instead of $\leq \frac{1}{2}$. In order to get gapless analogues of the question reduction and answer reduction theorems of [22], we look to [24], which proves that zero-gap MIP^* is equal to Π_2.
- A natural corollary of our gapless hypercompression theorem is that there is a (non-robust) two-prover test for n EPR pairs that uses only *constantly* sized questions (and $O(n)$ sized answers). This result arises from applying the gapless hypercompression theorem to the 'question sampling game' of [24], which self-tests for n EPR pairs, because hypercompression also preserves entanglement bounds. (We believe that it may be possible to improve the $O(n)$-sized answers to $\mathrm{poly}\log(n)$ by applying a round of gapless answer reduction to the question sampling game before we apply hypercompression.) As far as we know, this is the first nonlocal game[3] in the literature which achieves a self-test for a growing number of EPR pairs using constantly sized questions (see [30, Table 1]). We leave obtaining an analogous result in the gapped case, which would result in a *robust* two-prover test for n EPR pairs with constant sized questions, as an open problem; see Section 1.5 for more discussion.

1.3 Lower Bounds on MIP^* Protocols from Kolmogorov Complexity

Our previous result shows that MIP^* with constant question complexity and polylogarithmic answer complexity is equal to general MIP^* (with polynomial question, answer and decision complexity). (We also prove that zero-gap MIP^* with constant question complexity and almost-logarithmic answer complexity is equal to general zero-gap MIP^*.) It is natural to ask how far we can push in this direction. For example, is MIP^* with constantly sized questions and (truly) $O(\log n)$ sized answers equal to general MIP^*? What about MIP^* with constantly sized questions and, say, $O(\log\log n)$ sized answers?

Our final set of results shows that the parameters we can achieve by using hypercompression (see the previous section of this introduction) are in fact almost tight. Specifically, we prove that any MIP^* protocol deciding all of RE—in fact, any MIP^* protocol deciding all of EEXP—must have $q(n) + a(n) \geq \frac{1}{2}\log n$, where n is the instance size and $q(n)$ and $a(n)$ are the question and answer sizes (for a single prover) in the protocol respectively. An identical lower bound holds on question and answer sizes for gapless MIP^* protocols deciding all of EEXP. In particular, the latter shows that we have essentially already achieved a tight characterisation of zero-gap MIP^* as far as question and answer complexity are concerned: [24] exhibited a zero-gap MIP^* protocol for Π_2 with $O(\log n)$ question complexity and $O(1)$ answer complexity, and we exhibit a zero-gap MIP^* protocol for Π_2 with $O(1)$ question complexity and $O(\log n \cdot \log^* n)$ answer complexity, the former of

[3] There are *nonlocal correlations* with constant-sized questions [13], and indeed constant-sized questions and answers [18] that self-test maximally entangled states of arbitrarily high dimension. However, this is a different notion of self-testing, where one requires not just the winning probability to be close to optimal, but the entire distribution of answers given questions to be close to a target distribution.

which matches the lower bound up to constant factors, and the latter of which matches the lower bound up to a factor of $O(\log^* n)$. In the gapped case, some degree of leeway remains between the upper and the lower bound—in particular, the lower bound has $q(n) + a(n) \geq \frac{1}{2}\log n$, but the best upper bound that we believe current techniques could prove only has $q(n) + a(n) = \text{poly}\log(n)$. We think that the upper bound is the one that can be tightened, and we leave closing the gap as an interesting open problem whose resolution may have other significant implications. (See our open problems section, Section 1.5, for more discussion of this.)

We prove this lower bound by observing a connection between the sizes of questions and answers in an MIP* protocol deciding a computational problem and the size of the *advice* that a deterministic Turing machine must take to solve the same problem (or, equivalently, the size of the description of a Turing machine that solves the same problem). More specifically, we show a way to convert any MIP* protocol with questions of size $q(n)$, answers of size $a(n)$, and verifier time complexity $t(n)$ deciding a language L into a deterministic Turing machine running in time roughly $2^{t(n)}$ and taking advice of length roughly $2^{2^{2(q(n)+a(n))}}$ which also decides L. We then observe that one can use techniques from time-bounded Kolmogorov complexity theory to show that EEXP cannot be decided by any Turing machine running in time $2^{\text{poly}(n)}$ and taking $\varepsilon 2^{cn}$ advice for any $\varepsilon + c < 1$.[4] Combining the two statements shows our claimed lower bound, since an MIP* protocol for RE with very small questions and answers would result in a Turing machine to decide RE that takes comparatively little advice, which would contradict the lower bound on EEXP.

1.4 Related Work

We have already addressed most of the related literature above; for further discussion of some relationships between our work and [11] and [9], see the arXiv version of this paper.

[4] One may ask why we had to prove this, i.e. why we did not use known circuit lower bounds for large time classes. The answer is that, because the *advice complexity* of the Turing machine which we obtain from 'converting' the MIP* protocol is more sensitive to $q(n)+a(n)$ than the *running time* of the same Turing machine, we wanted a lower bound which treated running time and advice separately. In particular, $t(n)$ (the verifier's time complexity in the MIP* protocol) could be any arbitrary polynomial in n, e.g. n^{100}, and may not depend explicitly on $q(n) + a(n)$ (which here could be sub-logarithmic). Because any language is decidable by circuits of size 2^n, and the running time of the Turing machine M which comes out of our 'conversion' process is $2^{t(n)}$, we would not be able to prove any substantial conclusions about $q(n)+a(n)$ by comparing the circuit version of M with known circuit lower bounds if $t(n)$ happened to be n^{100}, since then the complexity of the circuit version of M would already be large enough to decide any language even if we only counted the $2^{t(n)} \approx 2^{n^{100}}$ gates that came from encoding the tableau of M's execution. On the other hand, since the *advice complexity* of the Turing machine M depends sharply on $q(n) + a(n)$, a lower bound on RE (or EEXP $\subseteq$ RE) that has a precise dependence on advice and a looser dependence on time complexity serves our purposes well.

We remark that another lower bound for RE with a precise dependence on advice and a looser dependence on time complexity is the bound which states that no finite-time Turing machine can solve the halting problem with fewer than $\approx n$ bits of advice. However, this bound is 'too much in the other direction', i.e. the lower bound on the advice is very weak because the running time is allowed to be any finite time, and therefore potentially much larger than $2^{\text{poly}(n)}$. We wanted a bound which captured a trade-off between running time and advice that would allow us to derive a logarithmic lower bound on $q(n) + a(n)$, and so we proved the bound stated in the main text.

1.5 Open Questions

(1) **Putting MA or AM in QMA(k) with small communication complexity.** There is now a long line of works about QMA(2) protocols for NP with sublinear communication. It is natural then to ask: is there a QMA(k) protocol (or an AM(k) protocol) with sublinear communication for AM (or MA)? The main obstacle here is that we do not have a 'PCP theorem' for MA or AM (in the sense that we have one for NP), unless MA = NP (resp. AM = NP), but the birthday paradox trick which puts NP in QMA(2) with sublinear communication complexity relies centrally on having a very short (short in terms of proof length) PCP for NP. Alternatively, could we prove that, if MA (or AM) is in QMA(k) with sublinear communication complexity, then MA = NP (or AM = NP)?

(2) **A tight gapped MIP* protocol for RE.** As we mentioned in Section 1.3, there is some leeway between our lower bound on the communication complexity of any MIP* protocol for EEXP and the upper bound which we believe we can achieve by applying hypercompression to the [22] MIP* protocol for RE. One way to prove a tight upper bound would be to show that there are *entanglement-sound PCPPs for* NP—that is, to prove the result which [26] claimed to prove, but whose proof subsequently turned out to have a bug. Such a result would yield a *gapped answer reduction theorem* that has similar question and answer size parameters to those of the gapless answer reduction theorem which we make use of in our gapless results.

(3) **A *rigid* self-test for n EPR pairs with constant-sized questions.** The techniques in this work can only yield a bound on the *dimension* of the Hilbert space shared by any pair of provers who pass in our so-called 'self-test for n EPR pairs' with constantly sized questions. Is there a self-test for n EPR pairs with constantly sized questions (and, say, poly(n) sized answers) which guarantees that any two provers who pass in the self-test must be using a particular strategy, up to local isometries? This is a very powerful and useful property of most self-tests for EPR pairs in the literature which is known as rigidity.

(4) **A *robust* test for Schmidt rank n with constant-sized questions.** In this work, the only entanglement bounds we can obtain are in the *gapless* case (i.e. for perfect strategies). Can we show a game with constant-sized questions where any strategy achieving value $\geq 1/2$ must have *Schmidt rank* at least n? (The Schmidt rank is the number of nonzero Schmidt coefficients, and is a relatively loose characterization of entanglement.) Such a bound was obtained by [22] for their compression theorems, but we cannot use it for an interesting technical reason: in our work, in order to perform parallel repetition for games with large answer sizes, we must use the analysis of [16], rather than that of [6]. However, this analysis does not preserve entanglement bounds, since the reduction from the parallel repeated strategy to a single-round strategy requires adding a large amount of entanglement. We refer the reader to the discussion in [22, Section 11] for more details on this point.

(5) **A communication lower bound for a self-test for n EPR pairs.** We get a lower bound on the communication complexity of MIP* protocols for RE which almost matches the upper bound we achieve by using hypercompression. As we mention in Section 1.2.2, another consequence of (gapless) hypercompression is a non-robust self-test for n EPR pairs with $O(1)$ sized questions and (probably, using answer reduction) poly log(n) sized answers. Is there a way to lower bound the communication complexity of a self-test for n EPR pairs using computational arguments, as we did the communication complexity of MIP* protocols for RE? It is not at once obvious how to do this, since such a self-test does not directly solve any well-understood computational problem.

(6) **Infinite randomness expansion with two provers.** Hypercompression yields a self-test for n EPR pairs with *constantly* sized questions. It is tempting then to ask: can we do *infinite randomness expansion* [14] using only 2 provers by using this self-test? The naïve approach does not work because, in the 'question sampling game' of [25] and the 'introspection game' of [22], the probability that the provers are asked the 'introspect' questions which cause them to generate randomness (as opposed to being asked questions that test their consistency with each other) decreases by a constant factor every time one applies question reduction, and we need to apply question reduction approximately $\log^*(n)$ times in order to make the questions constant sized. Can this obstacle be gotten around?

1.6 Omitted Material from This Proceedings Version

Due to space constraints we have omitted the technical statements and proofs of all results on AM* in this proceedings version, and we refer readers to the arXiv version of this paper.

2 PRELIMINARIES

2.1 Probability Basics

We can represent a probability distribution $\mu : \Omega \to [0,1]$ over a finite sample space Ω as a vector $\vec{\mu}$ of length $|\Omega|$ such that the ith entry of the vector $\vec{\mu}$ is exactly $\mu(i)$. For two probability distributions μ, ν over sample spaces Ω and Ω', we then denote by $\mu \otimes \nu$ the probability distribution over $\Omega \times \Omega'$ whose vector representation is the vector $\vec{\mu} \otimes \vec{\nu}$.

2.2 Quantum Information Basics

Definition 2. *For $K \in \mathbb{N}$, the* quantum Fourier transform $\mathcal{F}_K$ *is the unitary map over $\mathbb{C}^K$ defined by*

$$\mathcal{F}_K|s\rangle = \frac{1}{\sqrt{K}} \sum_{t=0}^{K-1} \omega_K^{s \cdot t} |t\rangle, \tag{2}$$

where $\omega_K = \exp(2\pi i/K)$. This map defines the Fourier basis *consisting of the states $|\bar{s}\rangle = \mathcal{F}_K|s\rangle$ for $s \in \{0, \ldots, K-1\}$. In particular,*

$$|\bar{0}\rangle = \frac{1}{\sqrt{K}} \sum_t |t\rangle. \tag{3}$$

Definition 3. *A* Bell measurement *is a two-outcome measurement $\{M, \mathbf{1} - M\}$ on a bipartite Hilbert space $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ that can be implemented by separately measuring the A and B registers with a POVM measurement, and then applying a classical Boolean function to the measurement outcomes. In other words, there exist POVMs $\{A_a\}$ acting on $\mathcal{H}_A$ and $\{B_b\}$ acting on $\mathcal{H}_B$, and a Boolean function f such that*

$$M = \sum_{a,b:f(a,b)=1} A_a \otimes B_b.$$

2.3 QMA(2) and related classes

For a fuller treatment of this class, we refer the reader to [19].

Definition 4. QMA(2) *is the class of quantum Merlin-Arthur proof systems where Arthur is a polynomial-time quantum machine and receives a witness $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle$ from Merlin that is guaranteed to be in tensor product across a fixed cut. A* QMA(2) *protocol decides a language L if for any input $x \in L$, there is a witness state $|\psi_1\rangle \otimes |\psi_2\rangle$ that Arthur accepts with probability at least c (the* completeness probability*), and for any input $x \notin L$, no witness state in tensor product form makes Arthur accept with probability greater than s (the* soundness probability*); when not otherwise specified, we assume $c = 2/3$ and $s = 1/3$.*

Definition 5. BellQMA(2) *is the class of* QMA(2) *proof systems where the POVM element corresponding to the accepting measurement outcome of verifier is a Bell measurement (see Definition 3).*

3 A BellQMA(2) PROTOCOL FOR 3SAT

3.1 The Protocol

Definition 6 (Generalised K-colouring). *Let $K \in \mathbb{N}$, let $\mathcal{G} = (V, E)$ be a graph, and let $R : E \times [K] \times [K] \to \{0,1\}$ be a function. We say that $\mathcal{G}$ is generalised K-colourable with respect to R if there exists an assignment function $c : V \to [K]$ such that, for all edges $e = (v_1, v_2) \in E$, $R(e, c(v_1), c(v_2)) = 1$.*

Definition 7 (δ-GAP-KCOL). *δ-GAP-KCOL is a promise problem. An instance of δ-GAP-KCOL consists of a graph $\mathcal{G} = (V, E)$, a number $K \in \mathbb{N}$, and a function $R : E \times [K] \times [K] \to \{0,1\}$.*

- *$(\mathcal{G}, K, R)$ is a YES-instance of δ-GAP-KCOL if $\mathcal{G}$ is generalised K-colourable with respect to R.*
- *$(\mathcal{G}, K, R)$ is a NO-instance of δ-GAP-KCOL if, for all possible assignments $c : V \to [K]$, there exist at least $\delta|E|$ edges $e \in E$ such that $R(e, c(v_1), c(v_2)) = 0$.*

Theorem 8. *There is a reduction $f : \{0,1\}^* \to \{0,1\}^*$ from 3SAT to δ-GAP-KCOL with constant $\delta > 0$ and constant $K > 0$ such that, if x is an N-clause 3SAT instance, $f(x)$ is an instance $(\mathcal{G}, K, R)$ of δ-GAP-KCOL such that $|V| = O(N \cdot \operatorname{poly}\log N)$ and $|E| = O(N \cdot \operatorname{poly}\log N)$.*

Proof. See Theorem 2 of [10]. □

We now present a BellQMA(2) protocol for δ-GAP-KCOL. The protocol relies on two sub-tests: the uniformity test (Figure 2) and the consistency test (Figure 3). We call the verifier in this protocol Arthur, and the two provers Alice and Bob.

Fix an instance $(\mathcal{G} = (V, E), K, R)$ of δ-GAP-KCOL.

Input: Let $n = |V|, m = |E|$. All parties in the protocol receive the instance $(\mathcal{G} = (V, E), K, R)$ as input, along with an integer $k = O(\sqrt{n})$, and a constant $0 < \eta < 1$ to use in the uniformity test (Figure 2). Honest provers also receive as input a generalised K-colouring of $\mathcal{G}$, described as a function $c : V \to [K]$.

The protocol is as follows:

(1) Alice and Bob both send Arthur a state; Alice's state is $k(\log m + 2\log K)$ qubits long, and Bob's state is $k(\log n + \log K)$ qubits long. Let the states that they send be $|\psi_1\rangle$ and $|\psi_2\rangle$. Honest provers send the states

$$|\psi_1\rangle = \left(\frac{1}{\sqrt{m}} \sum_{e=(v_1,v_2)\in E} |e\rangle|c(v_1), c(v_2)\rangle\right)^{\otimes k}$$

$$|\psi_2\rangle = \left(\frac{1}{\sqrt{n}} \sum_{v\in V} |v\rangle|c(v)\rangle\right)^{\otimes k}.$$

(2) Arthur flips a single coin. If it lands heads, he performs the uniformity test (Figure 2) on both $|\psi_1\rangle$ and $|\psi_2\rangle$, setting η to be the choice of η that was provided to him as input. The uniformity test also takes two natural number parameters, K' and Q. For the uniformity test on $|\psi_1\rangle$, he sets

$$K' = K^2, Q = m,$$

and for the uniformity test on $|\psi_2\rangle$, he sets

$$K' = K, Q = n.$$

If it lands tails, Arthur performs the consistency test (Figure 3) on $|\psi_1\rangle \otimes |\psi_2\rangle$, setting $\mathcal{G}, K, R$ to be the choices which were provided to him as input.

Figure 1: The BellQMA(2) protocol for δ-GAP-KCOL.

Lemma 9 (Completeness). *If $\mathcal{G}$ is generalised K-colourable with respect to R, then the honest strategy outlined in Figure 1 is accepted with probability $1 - \exp(-\Omega(\sqrt{n}))$.*

Proof. The consistency test accepts with probability 1 when $\mathcal{G}$ is generalised K-colourable and the two provers are honest. According to the analysis in [10, Section 3.1], the uniformity test on $|\psi_1\rangle$ and the uniformity test on $|\psi_2\rangle$ each pass with probability $1 - \exp(-\Omega(\sqrt{n}))$ when the provers are honest. A union bound gives the desired conclusion. □

In the following sections, we analyse the soundness of the protocol.

3.2 Soundness of Uniformity Test

For illustrative purposes, we begin with a zero-error analysis of the uniformity test. In the proof of Lemma 11, we will show how the argument presented below generalises to the case of nonzero error.

Lemma 10. *Suppose $|\psi\rangle$ passes the uniformity test with certainty. Then there exists a collection $\mathcal{S}$ of subsets of $[k]$ such that,*

(1) for all $T \in \mathcal{S}$, it holds that $|T| \geq \frac{k}{K'}(1 - \eta)$, and
(2) the distribution μ_Q which results from measuring the question registers of $|\psi\rangle$ in the standard basis can be decomposed as a mixture

$$\mu_Q = \sum_{T\in\mathcal{S}} p(T)\mu_T^{unif} \otimes \mu_{\overline{T}}^{junk},$$

Input: Two numbers $K', Q \in \mathbb{N}$, another number $0 < \eta < 1$, and a state $|\psi\rangle_{\mathsf{Q}_1\mathsf{A}_1\ldots\mathsf{Q}_k\mathsf{A}_k}$ on registers $\mathsf{Q}_1\mathsf{A}_1 \ldots \mathsf{Q}_k\mathsf{A}_k$. The registers Q_i are called the *question* registers and the registers A_i are called the *answer* registers.

(1) Perform a Fourier transform $\mathcal{F}_{K'}$ on each answer register and then measure it in the standard basis.
(2) Let $Z = \{i :$ the answer register A_i measured to $0\}$. If $\frac{|Z|}{k} < (1 - \eta)\frac{1}{K'}$, reject; otherwise, continue.[1]
(3) For each answer register A_i that measured to 0 in step 1, perform a Fourier transform $\mathcal{F}_Q$ on the ith question register Q_i, and measure it in the standard basis. If any non-zero measurement outcome is obtained at this step, reject. Otherwise, accept.

[1] η is necessary because even honest Merlins will not always pass in the uniformity test; instead, they will only be able to achieve an *average* of $|Z| = \frac{k}{K'}$, so η is necessary to be able to perform a Chernoff bound and achieve $1 - \exp(-k)$ completeness. See [10, Section 3.1] for more details.

Figure 2: The uniformity test.

Input:

- A state $|\psi\rangle_{\mathsf{Q}_1\mathsf{A}_1\ldots\mathsf{Q}_k\mathsf{A}_k} \otimes |\psi'\rangle_{\mathsf{Q}'_1\mathsf{A}'_1\ldots\mathsf{Q}'_\ell\mathsf{A}'_\ell}$ on registers $\mathsf{Q}_1\mathsf{A}_1 \ldots \mathsf{Q}_k\mathsf{A}_k\mathsf{Q}'_1\mathsf{A}'_1 \ldots \mathsf{Q}'_\ell\mathsf{A}'_\ell$. The registers Q_i and Q'_j, $i \in [k], j \in [\ell]$, are called the *question* registers, and the registers A_i and A'_j are called the *answer* registers.
- A graph $\mathcal{G} = (V, E)$.
- A number $K \in \mathbb{N}$.
- A relation $R : [K] \times [K] \to \{0, 1\}$.

(1) Measure all the registers in the standard basis. Interpret each measurement outcome in a register Q_i, $i \in [k]$, as a question for Alice, and interpret the measurement outcome coming from the associated answer register A_i as her answer to that question. (Therefore, Alice receives k questions and answers each one.) Interpret each measurement outcome in a register Q'_j, $j \in [\ell]$, as a question for Bob, and interpret the measurement outcome coming from the associated answer register A'_j as his answer to that question.
(2) Interpret each Alice question as an edge $e \in E$, and interpret the corresponding Alice answer as a pair of colours in $[K]$ for the vertices that form the endpoints of e. Interpret each Bob question as a vertex $v \in V$, and interpret the corresponding Bob answer as a colour in $[K]$ for v. Let $A \subseteq E$ be the set of all edges obtained as Alice questions and $B \subseteq V$ be the set of all vertices obtained as Bob questions.
(3) For every edge $e \in A$ and vertex $v \in B$ such that $v \in e$, check that Alice's and Bob's colorings agree and that the two endpoints of e are assigned colours that satisfy the function $R(e, \cdot, \cdot)$.

Figure 3: The consistency test.

where $p : \mathcal{S} \to [0,1]$ is a distribution over $\mathcal{S}$, μ_T^{unif} is the uniform distribution over $[Q]^{|T|}$ on the indices in T, and $\mu_{\overline{T}}^{junk}$ is an arbitrary distribution on the indices in $[k] - T$.

Proof. Suppose we perform the first step of the uniformity test (Figure 2) on $|\psi\rangle$, i.e., we measure all the answer registers of $|\psi\rangle$ in the Fourier basis. Let $\rho_{\vec{r}}$ denote the post-measurement state after this measurement conditioned on getting outcome $\vec{r}$. Assuming that $|\psi\rangle$ passes the uniformity test with certainty, $\vec{r} = r_1, \ldots, r_k$ must be such that $r_i = 0\ \forall i \in T$ for some subset $T \subset [k]$ with $|T| \geq \frac{k}{K'}(1-\eta)$. Moreover, the probability that $\rho_{\vec{r}}$ now passes step 3 of the uniformity test is still 1. Therefore,

$$\rho_{\vec{r}} = (|\bar{0}\rangle\langle\bar{0}|)^{\otimes T} \otimes \rho_{\overline{T}},$$

where the notation $(|\bar{0}\rangle\langle\bar{0}|)^{\otimes T}$ means that the registers with indices in T are in the all-zero state in the Fourier basis and in tensor product with the other registers.

Thus, measuring $\rho_{\vec{r}}$ in the standard basis will yield uniformly random iid outcomes on the registers in T and some arbitrary distribution on the other registers.

Finally, to get the lemma, observe that (letting ρ denote the post-measurement state after the Fourier measurement of step 1 with no conditioning)

$$\rho = \sum_{\vec{r}} q_{\vec{r}}\ \rho_{\vec{r}},$$

for some distribution $q_{\vec{r}}$. Thus, the conclusion follows. □

We now proceed to the main technical lemma in this section, which is a version of Lemma 10 that tolerates constant error.

Lemma 11. *Suppose $|\psi\rangle$ passes the uniformity test with probability $1 - \varepsilon > 0$. Then there exists a collection $\mathcal{S}$ of subsets of $[k]$ such that,*

(1) for all $T \in \mathcal{S}$, it holds that $|T| \geq \frac{k}{K'}(1-\eta)$, and
(2) the distribution μ_Q which results from measuring the question registers of $|\psi\rangle$ in the standard basis can be decomposed as a mixture

$$\mu_Q \simeq_{\delta(\varepsilon)} \sum_{T \in \mathcal{S}} p(T) \mu_T^{unif} \otimes \mu_{\overline{T}}^{junk},$$

where

(a) $p : \mathcal{S} \to [0,1]$ is a distribution over $\mathcal{S}$,
(b) μ_T^{unif} is the uniform distribution over $[Q]^{|T|}$ on the indices in T,
(c) $\mu_{\overline{T}}^{junk}$ is an arbitrary distribution on the indices in $[k] - T$,
(d) the notation $\simeq_\delta$ indicates that the two sides are a distance of δ apart in total variational distance, and
(e) $\delta(\varepsilon) = O(\varepsilon^{1/4})$.

Proof. Suppose we perform the first step of the uniformity test (Figure 2) on $|\psi\rangle$, i.e., we measure all the answer registers of $|\psi\rangle$ in the Fourier basis. Let $\rho_{\vec{r}}$ denote the post-measurement state after this measurement conditioned on getting outcome $\vec{r}$, and let ρ denote the overall post-measurement state after this measurement without conditioning on any particular outcome. Let $q_{\vec{r}}$ denote the probability of obtaining any given outcome $\vec{r}$.

Let $p_{success,\vec{r}}$ be a function mapping density matrices to $[0,1]$ such that $p_{success,\vec{r}}(\sigma)$ gives the probability that a given mixed state σ passes when it is subjected to step 3 of the uniformity test and $\vec{r}$ was the outcome obtained in step 1 of the uniformity test. Let $\mathbf{1}_{\vec{r}}$ be an indicator function which indicates whether or not a given vector $\vec{r}$ passes step 2 of the uniformity test (i.e. whether or not $\vec{r}$ is such that there exists $T \subseteq [k]$, with $|T| \geq \frac{k}{K'}(1-\eta)$, for which $r_i = 0\ \forall i \in T$). Using this notation, the probability that $|\psi\rangle$ passes in the uniformity test can then be expressed as

$$\sum_{\vec{r}} q_{\vec{r}} \cdot \mathbf{1}_{\vec{r}} \cdot p_{success,\vec{r}}(\rho_{\vec{r}}) \geq 1 - \varepsilon.$$

Rewrite as

$$\sum_{\vec{r}} q_{\vec{r}} (1 - \mathbf{1}_{\vec{r}} \cdot p_{success,\vec{r}}(\rho_{\vec{r}})) \leq \varepsilon.$$

Therefore (using a Markov bound), with probability at least $1 - \frac{1}{\alpha}$, $\vec{r}$ obtained in step 1 is such that

$$(1 - \mathbf{1}_{\vec{r}} \cdot p_{success}(\rho_{\vec{r}})) \leq \alpha\varepsilon. \tag{4}$$

Let us set $\alpha = \frac{1}{\sqrt{\varepsilon}}$, and define any such $\vec{r}$ to be *good*. With this definition of α, $\vec{r}$ is good with probability at least $1 - \sqrt{\varepsilon}$. Note that, for any good $\vec{r}$, step 2 of the uniformity test passes with certainty (or else $\mathbf{1}_{\vec{r}} = 0$ and Equation (4) would become $1 \leq \sqrt{\varepsilon}$), and step 3 of the uniformity test applied to $\rho_{\vec{r}}$ passes with probability at least $1 - \sqrt{\varepsilon}$.

Let $\rho_{question|\vec{r}}$ be $\rho_{\vec{r}}$ restricted to its question registers. For any fixed good $\vec{r}$ (for which step 3 of the uniformity test applied to $\rho_{\vec{r}}$ passes with probability at least $1 - \sqrt{\varepsilon}$), we have

$$\mathrm{tr}[((|\bar{0}\rangle\langle\bar{0}|)^{\otimes T} \otimes I_{\overline{T}})\rho_{question|\vec{r}}] \geq 1 - \sqrt{\varepsilon},$$

where the notation $(|\bar{0}\rangle\langle\bar{0}|)^{\otimes T}$ means that the registers with indices in T are in the zero Fourier state and in tensor product with the other registers. Thus, by the Gentle Measurement Lemma [32, Lemma 9.4.1], it holds that

$$\|\rho_{question|\vec{r}} - \underbrace{(|\bar{0}\rangle\langle\bar{0}|)^{\otimes T} \otimes \sigma(\vec{r})_{\overline{T}}}_{\sigma(\vec{r})}\|_1 \leq 2\varepsilon^{1/4}. \tag{5}$$

By construction, measuring $\sigma(\vec{r})$ in the standard basis will yield a distribution $\mu^{\vec{r}}$ that is uniformly random iid outcomes on the registers in T and some arbitrary distribution on the other registers.

Thus, measuring $\rho_{question|\vec{r}}$ in the standard basis (for any good $\vec{r}$) will yield a distribution that is $O(\varepsilon^{1/4})$-close to $\mu^{\vec{r}}$ in total variational distance, by the relation between variational distance and trace distance [28, Theorem 9.1].

Let $\rho_{question}$ denote the state ρ (defined in the first paragraph of this proof) restricted to its question registers. To argue about the distribution we obtain by measuring $\rho_{question}$ without the conditioning on a fixed good $\vec{r}$, observe that

$$\begin{aligned}\rho_{question} &= \sum_{\vec{r}} q_{\vec{r}}\ \rho_{question|\vec{r}} \\ &= \sum_{\vec{r} \in BAD} q_{\vec{r}}\ \rho_{question|\vec{r}} + \sum_{\vec{r} \in GOOD} q_{\vec{r}}\ \rho_{question|\vec{r}},\end{aligned} \tag{6}$$

and recall that $\vec{r}$ is good with probability at least $1 - \sqrt{\varepsilon}$. Given this, there exists a state with no weight on $\rho_{question|\vec{r}}$s with $\vec{r}$ in *BAD*

which is at most $O(\sqrt{\varepsilon})$ from $\rho_{question}$ in trace distance. Formally, if we define a new state

$$\rho'_{question} = \sum_{\vec{r}\in GOOD} q_{\vec{r}}\, \rho_{question|\vec{r}} + \left(\sum_{\vec{r}\in BAD} q_{\vec{r}}\right) \rho_{question|\vec{r}^*},$$

where $\vec{r}^*$ is an arbitrary (for concreteness, the lexicographically first) $\vec{r}$ in $GOOD$, we have that

$$\|\rho'_{question} - \rho_{question}\|_1 = O(\varepsilon^{1/2}). \tag{7}$$

Meanwhile, note that $\rho'_{question}$ can be expressed as a sum

$$\rho'_{question} = \sum_{\vec{r}\in GOOD} q'_{\vec{r}}\, \rho_{question|\vec{r}}, \tag{8}$$

where $q'_{\vec{r}}$ is some distribution over $\vec{r}$. For any good $\vec{r}$, let $T(\vec{r})$ denote a set such that $T \subseteq [k]$, $|T| \geq \frac{k}{K'}(1-\eta)$, $r_i = 0 \forall i \in T$. By the strong convexity of the trace distance [28, Theorem 9.3] and Equation (5), we have that

$$\left\|\rho'_{question} - \sum_{\vec{r}\in GOOD} q'_{\vec{r}} \left((|\bar{0}\rangle\langle\bar{0}|)^{\otimes T(\vec{r})} \otimes \sigma(\vec{r})_{\overline{T(\vec{r})}}\right)\right\|_1 \leq \sum_{\vec{r}\in GOOD} q'_{\vec{r}} \cdot 2\varepsilon^{1/4}. \tag{9}$$

Therefore,

$$\left\|\rho'_{question} - \sum_{\vec{r}\in GOOD} q'_{\vec{r}} \left((|\bar{0}\rangle\langle\bar{0}|)^{\otimes T(\vec{r})} \otimes \sigma(\vec{r})_{\overline{T(\vec{r})}}\right)\right\|_1 \leq 2\varepsilon^{1/4}. \tag{10}$$

By the triangle inequality, then,

$$\left\|\rho_{question} - \sum_{\vec{r}\in GOOD} q'_{\vec{r}} \left((|\bar{0}\rangle\langle\bar{0}|)^{\otimes T(\vec{r})} \otimes \sigma(\vec{r})_{\overline{T(\vec{r})}}\right)\right\|_1 = O(\varepsilon^{1/4}) + O(\varepsilon^{1/2}) = O(\varepsilon^{1/4}). \tag{11}$$

Finally, by the contractivity of the trace distance under completely positive trace-preserving maps [28, Theorem 9.2], measuring both $\rho_{question}$ and $\sum_{\vec{r}\in GOOD} q'_{\vec{r}} \left((|\bar{0}\rangle\langle\bar{0}|)^{\otimes T(\vec{r})} \otimes \sigma(\vec{r})_{\overline{T(\vec{r})}}\right)$ in the standard basis will not increase the trace distance between them. Measuring the latter in the standard basis manifestly results in a distribution of the form

$$\sum_{T\in\mathcal{S}} p(T)\mu_T^{unif} \otimes \mu_{\overline{T}}^{junk}$$

for $\mathcal{S}$ the set $\{T : \exists \vec{r} \in GOOD \text{ s.t. } T = T(\vec{r})\}$. Thus, the conclusion follows.

□

3.3 Soundness of Consistency Test and Soundness of Main Protocol

We begin by making a few definitions.

Definition 12 (Consistency game). *For any given graph $\mathcal{G} = (V, E)$, natural number K, and relation $R : [K] \times [K] \to \{0, 1\}$, we define the* (k, ℓ) consistency game, *denoted $G^{k,\ell}(\mathcal{G}, K, R)$ or simply $G^{k,\ell}$ when the parameters are clear from context, to be the following* classical *two player free game. (Note that this game is identical to the k, ℓ birthday repetition game from [2].)*

- *Alice receives a uniformly random size-k subset A of the set of edges E, and Bob receives a uniformly random size-ℓ subset B of the set of vertices V.*
- *Alice responds with a colouring of all the vertices that are at the endpoints of edges in A (i.e. Alice gives a number in $[K]$ for every vertex that is at the end of some edge in A), and Bob responds with a colouring of all the vertices in B.*
- *For every edge $e \in A$ and vertex $v \in B$ such that $v \in e$, Arthur checks that Alice and Bob's colorings agree and that the colours assigned to the two endpoints of e satisfy the relation $R(e, \cdot, \cdot)$.*

Definition 13 (Free game with special question distribution). *Let G be a two-player free game where question pairs are uniformly sampled from a question set $X \times Y$, and let $\mathcal{D}$ be a distribution over $X \times Y$. Then $G|_{\mathcal{D}}$ denotes the game G where the question pairs are sampled according to $\mathcal{D}$.*

We would like to prove the soundness of the protocol from section 3.1 by reducing its soundness to that of the consistency game from Definition 12, which was already analysed as the 'birthday game' in [2]. This means that, given a strategy for our QMA(2) protocol (i.e. a pair of witness states) from section 3.1, we would like to construct a strategy for the consistency game. The statement we want to prove is formalised in the following lemma.

Lemma 14. *Let $(\mathcal{G}, K, R)$ be an instance of δ-GAP-KCOL. Suppose the two states $|\psi_1\rangle, |\psi_2\rangle$ are accepted in the protocol of Figure 1 with probability at least $1 - \varepsilon$. Then there exists a strategy for $G^{k',k'}(\mathcal{G}, K, R)$, $k' = \frac{k}{K'}(1-\eta)$, with value $1 - O(\varepsilon^{1/4})$.*

We delay the proof of Lemma 14 until we have proven Lemmas 16, 17, and 18. It is clear, however, that Lemma 14 taken together with the following lemma, Lemma 15, yields the desired constant soundness for the protocol of section 3.1.

Lemma 15. *Suppose $\mathcal{G} = (V, E)$ is a graph with n vertices and $\Theta(n)$ edges, $K \in \mathbb{N}$ is a constant, and $R : E \times [K] \times [K] \to \{0, 1\}$ is a function. Suppose that any generalised K-coloring of $\mathcal{G}$ with respect to R violates at least δ-fraction of the edges for some constant $\delta > 0$. Then the classical value of the consistency game $G^{k,\ell}(\mathcal{G})$ for $k = \ell = \Omega(\sqrt{n})$ is at most $1 - c$ for a constant $c > 0$.*

Proof. This follows from Theorem 26 of [2]. □

From now on in this section, we will fix an instance $(\mathcal{G}, K, R)$ of δ-GAP-KCOL, and omit the parameters in our notation for the consistency game $G^{k,\ell}(\mathcal{G}, K, R)$.

Lemma 16. *For any pair of states $|\psi_1\rangle, |\psi_2\rangle$ that are accepted by the* QMA(2) *verifier Arthur in the consistency test from Figure 3 with probability $1 - \nu$, there exists a product distribution $\mathcal{D}(\psi_1, \psi_2) = \mathcal{D}^A \otimes \mathcal{D}^B$ over question pairs in $G^{k,\ell}$ and a randomized classical strategy achieving value $1 - \nu$ on $G^{k,\ell}|_{\mathcal{D}(\psi_1,\psi_2)}$.*

Proof. To obtain $\mathcal{D}$ and the classical strategy, perform a standard basis measurement of $|\psi_1\rangle$ and $|\psi_2\rangle$. □

The following is a restatement of Lemma 11.

Lemma 17. *For $|\psi_1\rangle, |\psi_2\rangle$ each passing the uniformity test (Figure 2) with probability $1 - \nu$, the distribution $\mathcal{D}(\psi_1, \psi_2)$ obtained by*

measuring the question registers of $|\psi_1\rangle$ and $|\psi_2\rangle$ in the standard basis is of the form

$$\mathcal{D}(\psi_1,\psi_2) = \mathcal{D}^A \otimes \mathcal{D}^B,$$

where, for $W \in \{A,B\}$, $\mathcal{D}^W$ has a decomposition of the form

$$\mathcal{D}^W = \sum_{T:T\subseteq[k],|T|\geq \frac{k}{K'}(1-\eta)} p^W(T)\mathcal{D}_T^{unif} \otimes \mathcal{D}_{\overline{T}}^{junk,W} + \mathcal{D}^{error,W},$$

where p^W is a distribution mapping subsets $T \subseteq [k]$ to $[0,1]$, and $\|\mathcal{D}^{error,W}\|_1 \leq \delta(\nu) = O(\nu^{1/4})$.

Intuitively, Lemma 17 says that, if Alice and Bob provide states $|\psi_1\rangle, |\psi_2\rangle$ which pass the uniformity test (Figure 2) with high probability, then the distribution over questions for $G^{k,\ell}$ which is obtained by measuring the question registers of $|\psi_1\rangle$ and $|\psi_2\rangle$ in the standard basis can be expressed (on each of Alice's and Bob's sides) as a convex mixture of distributions, such that most of the distributions making up this convex mixture are uniform on some significant fraction of their indices, and the rest of the distributions (the $\mathcal{D}^{error}$ ones) are arbitrary.

We now prove a lemma which will allow us to reduce the soundness of the consistency game with questions sampled in such a way to the soundness of a smaller instance of the consistency game with uniformly random questions.

Lemma 18. *Let $\mathcal{D} = \mathcal{D}^A \otimes \mathcal{D}^B$, where $\mathcal{D}^A$ is of the form*

$$\sum_{T:T\subseteq[k],|T|=k'} p^A(T)\mathcal{D}_T^{unif} \otimes \mathcal{D}_{\overline{T}}^{junk,A}$$

and $\mathcal{D}^B$ is of the form

$$\sum_{T:T\subseteq[\ell],|T|=\ell'} p^B(T)\mathcal{D}_T^{unif} \otimes \mathcal{D}_{\overline{T}}^{junk,B}$$

with $k' \leq k, \ell' \leq \ell$. Then

$$\omega(G^{k,\ell}|_{\mathcal{D}}) \leq \omega(G^{k',\ell'}).$$

Proof. Fix S, a strategy for the game $G^{k,\ell}|_{\mathcal{D}}$. Any strategy S for $G^{k,\ell}|_{\mathcal{D}}$ automatically induces a strategy S' for $G^{k',\ell'}$. Concretely, this induced strategy works as follows: given an Alice question x' from $G^{k',\ell'}$, Alice samples a set T according to the distribution $p(T)$, samples an x'' from $\mathcal{D}_{\overline{T}}^{junk}$, and sets $x = x'\|x''$. She then samples an answer a for the question x using her strategy for $G^{k,\ell}|_{\mathcal{D}}$. a will necessarily assign colours to all the endpoints of edges in the set of edges represented by x'; Alice responds with these colours only, which we will denote by a'. Bob does likewise, receiving question y', sampling question y'' to form $y = y'\|y''$, obtaining answer b to question y, and returning b', the restriction of b to that part which is relevant to y'.

We can analyse the success probability of S' relative to that of S through a series of hybrids.

(1) In the first hybrid, the strategy S is played in $G^{k,\ell}|_{\mathcal{D}}$. Suppose that S has a success probability of p in $G^{k,\ell}|_{\mathcal{D}}$.
(2) In the second hybrid, we define a new strategy R for $G^{k,\ell}|_{\mathcal{D}}$. The new strategy works as follows: Alice samples a question x' from the question distribution of $G^{k',\ell'}$, samples a set T according to the distribution $p(T)$, samples an x'' from $\mathcal{D}_{\overline{T}}^{junk}$, and sets $x = x'\|x''$. She then plays strategy S on question x. Bob does likewise, sampling question y', sampling question y'' to form $y = y'\|y''$, and playing strategy S on y. The form of $\mathcal{D}_W$ for $W \in \{A,B\}$ means that the questions x and y in this hybrid are distributed exactly as they would be in $G^{k,\ell}|_{\mathcal{D}}$. Therefore, the success probability of R is still p.
(3) In the third hybrid, Alice and Bob play the 'induced strategy' S' outlined in the first paragraph of this proof in the game $G^{k',\ell'}$. Note that, for all possible questions (x',y') in $G^{k',\ell'}$ and all possible embeddings of (x',y') into questions (x,y) in $G^{k,\ell}|_{\mathcal{D}}$, the checks that the rules of $G^{k',\ell'}$ require Arthur to perform on $a' \subseteq a, b' \subseteq b$ form a subset of the checks that the rules of $G^{k,\ell}|_{\mathcal{D}}$ require Arthur to perform on a, b. The latter means that, for any valid question (x',y') in $G^{k',\ell'}$ and any (x,y) induced by (x',y') according to the procedure in the first paragraph of this proof, a winning answer to question pair (x,y) induces a winning answer to (x',y'). Therefore, the success probability can only increase between the last hybrid and this one.

We conclude that, if there is a strategy S for $G^{k,\ell}|_{\mathcal{D}}$ which wins with probability p, then the strategy S' for $G^{k',\ell'}$ induced by S also wins with probability at least p. □

Now we are ready to prove Lemma 14. For convenience, we restate Lemma 14 below as Lemma 19.

Lemma 19. *Let $(\mathcal{G},K,R)$ be an instance of δ-GAP-KCOL. Suppose the two states $|\psi_1\rangle, |\psi_2\rangle$ are accepted in the protocol of Figure 1 with probability at least $1-\varepsilon$. Then there exists a strategy for $G^{k',k'}(\mathcal{G},K,R)$, $k' = \frac{k}{K'}(1-\eta)$, with value $1 - O(\varepsilon^{1/4})$.*

Proof. Let $\mathcal{D} = \mathcal{D}^A \otimes \mathcal{D}^B$ and $\mathcal{S}$ be, respectively, the product distribution over questions in $G^{k,k}$ and the randomised classical strategy succeeding with probability $1-2\varepsilon$ in $G^{k,k}|_{\mathcal{D}}$ which arise from applying Lemma 16 to $|\psi_1\rangle, |\psi_2\rangle$.

Now, applying Lemma 17 to $\mathcal{D}^W$ for $W \in \{A,B\}$, we obtain decompositions

$$\mathcal{D}^W = \sum_{T:T\subseteq k,|T|\geq \frac{k}{K'}(1-\eta)} p^W(T)\mathcal{D}_T^{unif} \otimes \mathcal{D}_{\overline{T}}^{junk,W} + \mathcal{D}^{error,W},$$

where the error terms have bounded 1-norm at most $\delta(\varepsilon) = O(\varepsilon^{1/4})$.

We can, without loss of generality, rewrite such a decomposition as

$$\mathcal{D}^W = \sum_{T:T\subseteq k,|T|=\frac{k}{K'}(1-\eta)} p^W(T)\mathcal{D}_T^{unif} \otimes \mathcal{D}_{\overline{T}}^{junk,W} + \mathcal{D}^{error,W}.$$

Since the error terms in $\mathcal{D}^W$ have probability mass at most $\delta(\varepsilon)$, the probability mass of the subgames in $G^{k,k}|_{\mathcal{D}}$ where either the Alice or the Bob questions are drawn from $\mathcal{D}^{error,W}$ is at most $2\cdot\delta(\varepsilon)$. We may thus discard these terms from $\mathcal{D}^W$ at the cost of changing the value of the game $G^{k,k}|_{\mathcal{D}}$ by $O(\delta(\varepsilon))$, obtaining a new game $G^{k,k}|_{\mathcal{D}'}$ where the question distribution $\mathcal{D}' = (\mathcal{D}^A)' \otimes (\mathcal{D}^B)'$ is such that $(\mathcal{D}^A)'$ and $(\mathcal{D}^B)'$ are mixtures over only the "good" terms, and which has value at least $1-2\varepsilon-O(\delta(\varepsilon))$.

To proceed, let us apply Lemma 18 to $G^{k,k}|_{\mathcal{D}'}$. This tells us that $\omega(G^{k,k}|_{\mathcal{D}'}) \leq \omega(G^{k',k'})$, with $k' = \frac{k}{K'}(1-\eta)$. Therefore,

$$1 - 2\varepsilon - O(\delta(\varepsilon)) \leq \omega(G^{k',k'}).$$

This concludes the proof of the lemma. □

Moreover, by Lemma 15, $\omega(G^{k',k'})$ is at most $1-c$ when $k' = \Omega(\sqrt{n})$. Putting this together with Lemma 14, we obtain that the success probability of a cheating Merlin in the protocol of section 3.1 is at most $1-c'$ for some constant $c' > 0$: the soundness of the consistency game $\omega(G^{k',k'})$, as expressed in Lemma 15, requires that

$$1 - O(\varepsilon^{1/4}) \leq 1 - c,$$

and therefore

$$\varepsilon \geq \Omega(c^4).$$

3.4 Lower Bounds for h_{Sep} Conditional on ETH

Definition 20 (h_{Sep}). *Let $\mathcal{H}_{AB} := \mathcal{H}_A \otimes \mathcal{H}_B$, where $\mathcal{H}_A$ and $\mathcal{H}_B$ are finite dimensional Hilbert spaces. Let $\mathcal{S}_{AB}$ be the set of separable states on $\mathcal{H}_{AB}$, i.e. the set of density matrices ρ_{AB} such that ρ_{AB} can be written as*

$$\rho_{AB} = \sum_k p_k \, \rho_{A,k} \otimes \rho_{B,k}$$

where, for all k, $\rho_{A,k}$ is a density matrix on $\mathcal{H}_A$, $\rho_{B,k}$ is a density matrix on $\mathcal{H}_B$, and p_k is a probability. Given a Hermitian matrix M on $\mathcal{H}_{AB}$, $h_{\text{Sep}}(\varepsilon, M, \mathcal{H}_{AB})$ is the problem of estimating

$$\max_{\rho \in \mathcal{S}_{AB}} \operatorname{tr}(M\rho)$$

up to additive error ε.

Theorem 21. *Let $|\mathcal{H}|$ denote the dimension of a Hilbert space $\mathcal{H}$. Suppose $\mathscr{A}$ is an algorithm to solve $h_{\text{Sep}}(\varepsilon, M, \mathcal{H}_{AB})$ for constant ε and M such that $\{M, \mathbf{1} - M\}$ is a Bell measurement (see Definition 3). If $\mathscr{A}$ has time complexity at most*

$$\exp(O(\log^{1-\nu} |\mathcal{H}_A| \log^{1-\mu} |\mathcal{H}_B|))$$

for $\nu + \mu > 0$, then 3SAT with N clauses has an algorithm taking time $2^{N^{1-(\nu+\mu)/2} \cdot \text{poly}\log N}$.

Proof. Suppose $\mathcal{G} = (V, E)$ is a graph with n vertices and m edges, and let $(\mathcal{G}, K, R)$ be an instance of δ-GAP-KCOL (Definition 7). We show a BellQMA(2) protocol in Section 3.1 to decide any such instance of δ-GAP-KCOL, which has a constant completeness-soundness gap if δ and K are both constants, and where the witness is on two unentangled registers of size $O(\sqrt{n} \log m)$ and $O(\sqrt{n} \log n)$. By Theorem 8, we can reduce any N-clause instance of 3SAT to an instance of δ-GAP-KCOL where δ and K are constants and $m, n = N \operatorname{poly}\log N$. As such, we can set $|\mathcal{H}_A| = |\mathcal{H}_B| = 2^{\sqrt{N} \cdot \text{poly}\log N}$. Applying the hypothetical algorithm $\mathscr{A}$ to the measurement $\{M, \mathbf{1} - M\}$ induced by our protocol, we get that $\mathscr{A}$ can solve 3SAT in time

$$\exp\left(O(\log^{1-\nu} |\mathcal{H}_A| \log^{1-\mu} |\mathcal{H}_B|)\right) \leq \exp(N^{1-(\nu+\mu)/2} \cdot \operatorname{poly}\log N).$$

□

Assuming the Exponential Time Hypothesis (that N-clause 3SAT does not have any algorithm taking time $2^{o(N)}$), Theorem 21 shows that there does not exist an algorithm for h_{Sep} on Bell measurements taking time at most

$$\exp(O(\log^{1-\nu} |\mathcal{H}_A| \log^{1-\mu} |\mathcal{H}_B|))$$

for any constant $\nu + \mu > 0$. Therefore, Theorem 21 shows that the algorithm given by [8] for h_{Sep} on LOCC measurements, a superclass of Bell measurements, is optimal (possibly up to factors doubly logarithmic in $|\mathcal{H}_A|$ and $|\mathcal{H}_B|$).

ACKNOWLEDGEMENTS

We thank Aram Harrow, Hamoon Mousavi, Chris Umans, and Ryan Williams for helpful conversations. TZ was supported in part by an Akamai Presidential Fellowship.

REFERENCES

[1] Scott Aaronson, Salman Beigi, Andrew Drucker, Bill Fefferman, and Peter Shor. 2009. The Power of Unentanglement. *Theory Of Computing* 5 (2009), 1–42. arXiv:0804.0802

[2] S. Aaronson, R. Impagliazzo, and D. Moshkovitz. 2014. AM with Multiple Merlins. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*. 44–55. https://doi.org/10.1109/CCC.2014.13 arXiv:1401.6848

[3] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. 1998. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)* 45, 3 (1998), 501–555.

[4] Sanjeev Arora and Shmuel Safra. 1998. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM (JACM)* 45, 1 (1998), 70–122.

[5] László Babai, Lance Fortnow, and Carsten Lund. 1991. Non-deterministic exponential time has two-prover interactive protocols. *Computational complexity* 1, 1 (1991), 3–40.

[6] Mohammad Bavarian, Thomas Vidick, and Henry Yuen. 2017. Hardness amplification for entangled games via anchoring. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 303–316. arXiv:1509.07466

[7] Hugue Blier and Alain Tapp. 2009. All languages in NP have very short quantum proofs. In *2009 Third International Conference on Quantum, Nano and Micro Technologies*. IEEE, 34–37. arXiv:0709.0738

[8] Fernando Brandão, Matthias Christandl, and John Yard. 2010. A quasipolynomial-time algorithm for the quantum separability problem. (2010). arXiv:1011.2751 [quant-ph]

[9] Fernando GSL Brandão and Aram W Harrow. 2013. Quantum de Finetti theorems under local measurements with applications. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 861–870.

[10] J. Chen and A. Drucker. 2010. Short Multi-Prover Quantum Proofs for SAT without Entangled Measurements. arXiv:1011.0716

[11] Alessandro Chiesa and Michael J. Forbes. 2011. Improved Soundness for QMA with Multiple Provers. (2011). arXiv:1108.2098

[12] Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. 2004. Consequences and limits of nonlocal strategies. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.* IEEE, 236–249. arXiv:quant-ph/0404076

[13] Andrea Coladangelo, Koon Tong Goh, and Valerio Scarani. 2017. All pure bipartite entangled states can be self-tested. *Nature Communications* 8, 1 (2017), 15485. arXiv:1611.08062

[14] Matthew Coudron and Henry Yuen. 2013. Infinite Randomness Expansion and Amplification with a Constant Number of Devices. (2013). arXiv:1310.6755

[15] Mikael de la Salle. 2022. Spectral gap and stability for groups and non-local games. (2022). arXiv:2204.07084 [math.OA]

[16] Irit Dinur, David Steurer, and Thomas Vidick. 2015. A parallel repetition theorem for entangled projection games. *Computational Complexity* 24, 2 (2015), 201–254. arXiv:1310.4113

[17] Tobias Fritz, Tim Netzer, and Andreas Thom. 2014. Can you compute the operator norm? *Proc. Amer. Math. Soc.* 142, 12 (2014), 4265–4276. arXiv:1207.0975

[18] Honghao Fu. 2022. Constant-sized correlations are sufficient to self-test maximally entangled states with unbounded dimension. *Quantum* 6 (2022), 614. arXiv:1911.01494

[19] Aram W Harrow and Ashley Montanaro. 2013. Testing product states, quantum Merlin-Arthur games and tensor optimization. *Journal of the ACM (JACM)* 60, 1 (2013), 1–43. arXiv:1001.0017

[20] Tsuyoshi Ito and Thomas Vidick. 2012. A multi-prover interactive proof for NEXP sound against entangled provers. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. IEEE, 243–252. arXiv:1207.0550

[21] Zhengfeng Ji. 2017. Compression of quantum multi-prover interactive proofs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 289–302. arXiv:1610.03133

[22] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. 2020. $\mathsf{MIP}^* = \mathsf{RE}$. (2020). arXiv:2001.04383 [cs.CC]

[23] Peter Bro Miltersen and N Variyam Vinodchandran. 2005. Derandomizing Arthur–Merlin games using hitting sets. *computational complexity* 14, 3 (2005), 256–279.

[24] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. 2020. On the complexity of zero gap MIP^*. (2020). arXiv:2002.10490 [cs.CC]

[25] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. 2021. Nonlocal Games, Compression Theorems and the Arithmetical Hierarchy. (2021). arXiv:2110.04651 [cs.CC]

[26] Anand Natarajan and Thomas Vidick. 2018. Retracted: Two-Player Entangled Games are NP-Hard. (2018). https://doi.org/10.4230/LIPICS.CCC.2018.20

[27] Anand Natarajan and John Wright. 2019. $\mathsf{NEEXP} \subseteq \mathsf{MIP}^*$. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 510–518. arXiv:1904.05870

[28] Michael A Nielsen and Isaac Chuang. 2002. Quantum computation and quantum information.

[29] Ben W Reichardt, Falk Unger, and Umesh Vazirani. 2013. Classical command of quantum systems. *Nature* 496, 7446 (2013), 456. arXiv:1209.0448

[30] Ivan Šupić and Joseph Bowles. 2020. Self-testing of quantum systems: a review. *Quantum* 4 (sep 2020), 337. https://doi.org/10.22331/q-2020-09-30-337

[31] Stephanie Wehner. 2006. Entanglement in Interactive Proof Systems with Binary Answers. In *STACS 2006*. Springer Berlin Heidelberg, 162–171. https://doi.org/10.1007/11672142_12

[32] Mark M Wilde. 2011. From classical to quantum Shannon theory. *arXiv preprint arXiv:1106.1445* (2011). arXiv:1106.1445

Received 2022-11-07; accepted 2023-02-06