

Short Multi-Prover Quantum Proofs for SAT without Entangled Measurements

Jing Chen ^{*} Andrew Drucker [†]

Abstract

BellQMA protocols are a subclass of multi-prover quantum Merlin-Arthur protocols in which the verifier is restricted to perform nonadaptive, unentangled measurements on the quantum states received from each Merlin. In this paper, we prove that m -clause 3-SAT instances have BellQMA proofs of satisfiability with constant soundness gap, in which $\tilde{O}(\sqrt{m})$ Merlins each send $O(\log m)$ qubits to Arthur. Our result answers a question of Aaronson et al., who gave a protocol with similar parameters that used entangled measurements; the analysis of our protocol is significantly simpler than that of Aaronson et al. Our result also complements recent work of Brandao, Christandl, and Yard, who showed upper bounds on the power of multi-prover quantum proofs with unentangled but adaptive (LOCC) measurements.

1 Introduction

In quantum Merlin-Arthur (QMA) proof systems, a computationally unbounded but untrusted prover Merlin tries to convince a polynomial-time quantum verifier Arthur that a given statement is true, by sending to Arthur a quantum state as a “proof”. We desire that the protocol have two properties. The first is “completeness”: if the statement is true, then there should exist a proof which makes Arthur accept with at least some high probability c . The second is “soundness”: if the statement is false, then for any proof received, Arthur should accept with at most some lower probability $s < c$. In general, the complexity class $\text{QMA}_{\ell,c,s}$ consists of all languages whose membership can be proved by a quantum Merlin-Arthur proof system using ℓ -qubit proofs, with completeness c and soundness s . The complexity class QMA is defined to be $\text{QMA}_{\text{poly}(n),2/3,1/3}$ where n is the input length.

The generalized multi-prover version of $\text{QMA}_{\ell,c,s}$, denoted $\text{QMA}(k)_{\ell,c,s}$, was introduced by Kobayashi, Matsumoto, and Yamakami in [KMY03]. In such a proof system, k Merlins are trying to convince a single Arthur that a given statement is true, by each sending Arthur a quantum state with ℓ qubits, and these k states are assumed to be unentangled with each other. The class $\text{QMA}(k)$ is defined to be $\text{QMA}(k)_{\text{poly}(n),2/3,1/3}$.

One piece of evidence for the power of multiple quantum provers was given by Blier and Tapp [BT09], who showed that every language in NP has a 2-prover proof system with extremely short proofs, of $\ell = O(\log n)$ qubits each. Unfortunately, the soundness gap in their proof system (i.e., the quantity $c - s$) is very small: their protocol has $c = 1, s = 1 - 1/\text{poly}(n)$. A related but

^{*}MIT. Email: jingchen@csail.mit.edu.

[†]MIT. Email: adrucker@mit.edu. Supported by a DARPA YFA grant. Supported during part of this work by an Akamai Presidential Graduate Fellowship.

incomparable result was shown by Aaronson et al. [ABD⁺09]: they showed that m -clause 3-SAT instances can be proved satisfiable by a proof system using $\tilde{O}(\sqrt{m})$ Merlins, each sending $O(\log m)$ bits, and with an improved soundness gap $c = 1, s = 1 - \Omega(1)$. This still gives an almost-quadratic improvement in total proof length compared to known classical proofs, at least in the regime where the number n of variables satisfies $n = \Theta(m)$.

A recent paper by Harrow and Montanaro [HM10] answers several important questions about $\text{QMA}(k)$. They prove that soundness amplification for $\text{QMA}(k)$ is possible and that $\text{QMA}(k) = \text{QMA}(2)$, for any $k = O(\text{poly}(n))$. Building on the result of [ABD⁺09], Harrow and Montanaro also show that there exists a 2-prover proof system with proof length $\tilde{O}(\sqrt{m})$ for m -clause 3-SAT instances.

In all results mentioned above, Arthur uses the so-called *swap test* measurement [BCWdW01] as an important step in the protocol. This is an efficient method to test whether two unentangled states are approximately equal. A natural question thus arises: how crucial is the swap test to the power of multi-prover quantum proof systems? The swap test is an example of an *entangled measurement*, in which the states may become entangled by the measurement process; so more generally, how crucial are entangled measurements to these proof systems? To make such questions formal, [ABD⁺09] defined the complexity classes $\text{LOCCQMA}(k)$ and $\text{BellQMA}(k)$. The class $\text{LOCCQMA}(k)$ consists of all languages whose membership can be proved by a k -prover proof system where Arthur is constrained to make unentangled measurements on the states provided by the Merlins, but is allowed to make these measurements adaptively based on the outcome of previous measurements. $\text{BellQMA}(k)$ is the subclass of $\text{LOCCQMA}(k)$ in which we additionally require that no choice of measurement depends on the outcomes of other measurements. (For more precise definitions of LOCCQMA and BellQMA protocols, see Sec. 2.1.) Brandao [Bra08] showed that $\text{BellQMA}(k) = \text{QMA}$ for constant k . Quite recently Brandao, Christandl, and Yard [BCY10] made a breakthrough in the study of entanglement, and used their techniques to show that $\text{LOCCQMA}(k) = \text{QMA}$ for constant k . The situation for growing values of k remains unclear.¹

Our contribution. In this paper, we exhibit a BellQMA proof system for 3-SAT, which essentially matches the parameters of the earlier protocol of [ABD⁺09]. Formally, we prove the following theorem:

Theorem 1. *There is a BellQMA proof system which, given a 3-SAT instance with m clauses, uses $\tilde{O}(\sqrt{m})$ Merlins, each of which sends $O(\log m)$ qubits. The proof system has completeness $1 - \exp\{-\Omega(\sqrt{m})\}$ and soundness $1 - \Omega(1)$.*

Our result shows that entangled measurement is not necessary for short proofs of membership in 3-SAT, and thus answers a question raised in [ABD⁺09, Sec. 6.3]. Our system (just barely) loses perfect completeness as achieved by the protocol of [ABD⁺09], but retains the constant soundness gap of that protocol. The analysis of our protocol is also significantly simpler than that of [ABD⁺09], which may be viewed as another contribution of this work.

Our protocol also complements a negative result from [BCY10] (although our work was independent of theirs). Corollary 5 of [BCY10] implies that if there exists a 2-prover LOCCQMA

¹On the one hand, the ideas of [HM10] rely on the swap test and do not apply to LOCCQMA and BellQMA . On the other hand, both the proof for $\text{BellQMA}(k) = \text{QMA}$ and that for $\text{LOCCQMA}(k) = \text{QMA}$ for constant k blow up the total length of the proofs to $n^{\exp(\Omega(k))}$, and thus cannot be used for $k = \omega(1)$. We can at least say that, if $\text{QMA}(2) = \text{QMA}$, then all classes here collapse to QMA .

protocol for 3-SAT with proof length $o(\sqrt{m})$ and with constant soundness gap, then there exists a deterministic algorithm solving 3-SAT in subexponential time.² This result seems to pose a significant barrier to achieving shorter proof length using unentangled measurements. Our positive result nearly reaches this barrier, except for the fact that we use more than 2 provers. If our protocol could be converted to a 2-prover BellQMA or LOCCQMA protocol with similar proof length, then (under the plausible assumption that 3-SAT requires exponential time) we would obtain a nearly tight understanding of the power of these restricted quantum proof systems for 3-SAT (and of many other NP languages, via standard reductions).

Our Techniques. The construction of our proof system, which we sketch next, adapts techniques used by Blier and Tapp in their proof system for 3-colorability from [BT09], and combines them with sampling and PCP ideas similar to those used by Aaronson et al. [ABD⁺09]. In the Blier-Tapp protocol, Arthur receives two states $|\Psi^1\rangle, |\Psi^2\rangle$ of form

$$|\Psi^i\rangle = \sum_{v,c} \alpha_{v,c}^i |v\rangle |c\rangle, \quad i = 1, 2,$$

where $v \in \{0, 1, \dots, n-1\}$ indexes a vertex in a graph G to be properly 3-colored and $c \in \{0, 1, 2\}$ is a color for v . The protocol randomly performs one of three tests on $|\Psi^1\rangle, |\Psi^2\rangle$:

- An “Equality Test” uses the swap test to check that $|\Psi^1\rangle, |\Psi^2\rangle$ are nearly equal. This is the only entangled measurement.
- A “Uniformity Test” uses the quantum Fourier transform to check that each state has amplitudes which are almost uniformly spread over the n vertices.
- A “Consistency Test” directly measures the vertex and color registers on each proof, rejecting if it sees two adjacent vertices with the same color or two differently-colored copies of the same vertex.

This proof is extremely succinct—only two states of $\log n$ qubits each—but as mentioned, its soundness gap is only inverse-polynomial in n . Intuitively³ this is because, if the Merlins send proofs uniformly spread over the vertices, each equipped with a coloring violating only one edge constraint, then the Consistency test can only succeed if the two vertices sampled come from this edge, which happens with probability $2/n^2$.

We modify this protocol as follows. First, we ask for $O(\sqrt{n})$ proofs instead of 2. The “birthday paradox” then ensures that the Consistency Test will turn up pairs of equal vertices. These will cause rejection unless almost all vertices are nearly-unanimous in their colorings across the supplied proofs. With this added assurance, we simply omit the Equality Test. Our modified Uniformity Test ensures that there are enough states in which the amplitudes are almost uniformly spread

²This is not quite made explicit in [BCY10], so we elaborate. [BCY10, Corollary 5] gives an explicit construction of certain mappings called “approximate disentanglers” (for LOCC measurements). Following the use of such mappings as described in [ABD⁺09] (where they were defined), a 2-prover LOCCQMA proof system for 3-SAT of proof length $o(\sqrt{m})$ would imply a single-prover QMA protocol for 3-SAT of proof length $\ell = o(m)$. The maximum acceptance probability of such a protocol can be approximately determined using semidefinite-programming solvers in time $\text{poly}(2^\ell) = 2^{o(m)}$.

³The actual soundness gap shown in [BT09] is $\Omega(n^{-6})$, even worse than our sketch would suggest; this was improved to $\Omega(n^{-3-\epsilon})$ by Beigi in [Bei10] by a modified protocol that still uses two $O(\log n)$ -sized proofs, but loses perfect completeness.

over the vertices, although to improve the soundness of this test, we are forced to sacrifice perfect completeness.

In the Consistency Test, we now also expect to sample pairs of vertices adjacent in G . However, this will only lead to rejection with noticeable probability if the constraint problem is “highly unsatisfiable”, in the sense that every coloring violates an $\Omega(1)$ fraction of the edge constraints. To ensure this, we apply the size-efficient PCP reduction of Dinur [Din07] to our original 3-coloring problem (or 3-SAT instance), which incurs only a polylogarithmic blowup in the instance size. This completes the sketch of our protocol and the basic ideas of the analysis; the full proof of correctness is slightly more involved.

Open Problems. Some questions raise from our result and those mentioned above. The most immediate one is whether the number of provers in our system can be further reduced without expanding the total proof length much. In particular, is there a 2-prover LOCCQMA proof for 3-SAT with length $\tilde{O}(\sqrt{m})$? As we discussed earlier, if 3-SAT requires exponential time, this would be a nearly tight result in terms of proof length. Whether general 2-prover protocols can achieve even shorter proofs of satisfiability remains an interesting question.

It also seems promising to see whether the entanglement theory ideas of [BCY10] can be extended to give a fuller understanding of entanglement between more than 2 quantum states. As just one benefit, this could yield new information about the power of $\text{LOCCQMA}(k)$ and $\text{BellQMA}(k)$ for superconstant k .

2 Preliminaries

We assume familiarity with (uniform) polynomial-time quantum algorithms. Such algorithms are describable by a polynomial-size quantum circuit with polynomially many auxiliary qubits; the circuit is required to be constructible by a classical logarithmic-space algorithm.

2.1 BellQMA and LOCCQMA protocols

We now more formally define the restricted multi-prover proof systems called BellQMA and LOCCQMA protocols. The complexity classes $\text{BellQMA}(k)$, $\text{LOCCQMA}(k)$ are defined in perfect analogy with $\text{QMA}(k)$, using these restricted protocols.⁴ In BellQMA protocols, Arthur performs a so-called “Bell test” upon the quantum proofs; in LOCCQMA protocols, Arthur performs a test involving only “local operations and classical communication” (LOCC) upon the proofs. This motivates the terminology.

Our definition of BellQMA protocols is slightly broader than that given in [Bra08], and we discuss the difference below. The more general class of LOCCQMA protocols will not be important in this paper, but we choose to provide a definition since previous discussions presume familiarity with the framework of LOCC tests (see [BBPS96, Nie99]).

In a k -prover QMA protocol, the verifier (Arthur) receives a classical input $x \in \{0, 1\}^n$, as well as k “proof” states $|\Psi_1\rangle, \dots, |\Psi_k\rangle$ from k provers (Merlins). The k proofs are required to be unentangled. Arthur performs some quantum-polynomial time test on the proofs, after which he either accepts or rejects. We say that a QMA protocol for a language $L \subseteq \{0, 1\}^*$ possesses *completeness* c and *soundness* $s < c$ if:

⁴In [BCY10], the notations $\text{QMA}_{\text{LO}}(k) = \text{BellQMA}(k)$ and $\text{QMA}_{\text{LOCC}}(k) = \text{LOCCQMA}(k)$ are used.

1. If $x \in L$, some collection $|\Psi_1\rangle, \dots, |\Psi_k\rangle$ causes Arthur to accept with probability at least c ;
2. If $x \notin L$, any collection $|\Psi_1\rangle, \dots, |\Psi_k\rangle$ causes Arthur to accept with probability at most s .

In a k -prover BellQMA protocol, we restrict the form of Arthur’s test as follows: Arthur first performs a polynomial-time quantum computation on x alone. The workspace is then measured fully, yielding a description of measurements M_1, \dots, M_k described by polynomial-size quantum circuits; the i -th measurement, which may output more than one bit, is required to act locally on the i -th proof. The measurements are then performed, and we let y_i denote the output of the i -th measurement. Finally, Arthur performs a quantum polynomial-time computation on (x, y_1, \dots, y_k) to decide whether to accept or reject.

Since the measurements M_1, \dots, M_k are fully determined by the intermediate measurement and act separately on the k unentangled proof states, the proof states remain unentangled after the M_i are performed. A second observation about BellQMA protocols is that the identities of the measurements M_1, \dots, M_k can be random variables, and these random variables need not be independent.⁵

In [ABD⁺09], the definition of BellQMA protocols was informal and slightly ambiguous. In Brandao’s thesis [Bra08], the definition of BellQMA protocols required the final computation on (x, y_1, \dots, y_k) to be performed by a *classical* polynomial-time algorithm. We feel that, since Arthur is allowed to use arbitrary polynomial-time quantum measurements M_i on the k proofs, it is natural to allow polynomial-time quantum computations in the final stage. Indeed, Brandao’s proof in [Bra08] that $\text{BellQMA}(k) = \text{BellQMA}$ works equally well if this final computation is allowed to be quantum. The BellQMA protocol that we give in this paper actually obeys Brandao’s more restrictive definition.

In LOCCQMA protocols, Arthur is allowed to repeatedly and adaptively choose measurements to perform on the proofs. However, these measurements are required to act locally on a single proof state, and they must be performed when Arthur’s workspace is in a computational basis state. This forces the proofs to remain unentangled throughout the computation.

Formally, k -prover LOCCQMA protocols can be defined as follows. Arthur’s verification algorithm consists of a polynomial number $p(n)$ of stages. Each stage $t \leq p(n)$ has the following form:

1. Arthur first performs a polynomial-time quantum computation acting on his workspace qubits alone. Arthur’s full workspace is then measured, yielding a tuple (i_t, M_t, z_t) . Here M_t describes a polynomial-time quantum measurement to be performed locally on the i_t -th proof, and z_t is an auxiliary memory string.
2. M_t is then performed, yielding an outcome y_t of one or more bits. Arthur then begins the $(t + 1)$ -st stage with his workspace initialized to the computational basis state $|y_t, z_t\rangle$.

Finally, Arthur accepts or rejects based upon the first bit of $y_{p(n)}$. We remark that Arthur is allowed to measure individual proof states more than once.

Note that BellQMA protocols can be defined as LOCCQMA protocols in which all measurements to be performed on the k proof states are determined in the first computation phase and described by the string z_1 , then nonadaptively performed in the following phases.

⁵In our protocol, the measurements will be chosen in a dependent fashion; however, it is not hard to modify our protocol to make these choices independent (for a fixed input x), with only a constant-factor increase in the number of provers.

2.2 Dinur’s PCP reduction

The recent version of the PCP Theorem given by Dinur [Din07] is a reduction from the Boolean Satisfiability problem to a so-called *constraint graph* problem, or 2-CSP. A constraint graph is an undirected graph (possibly with self-loops) along with a set Σ of “colors”. For each edge $e = (u, v) \in E$ the constraint graph has an associated constraint $R_e : \Sigma \times \Sigma \rightarrow \{0, 1\}$. A coloring $\tau : V \rightarrow \Sigma$ *satisfies* the constraint R_e if $R_e(\tau(u), \tau(v)) = 1$. We say that G is satisfiable if there exists a mapping τ that satisfies all constraints. We say that G is $(1 - \eta)$ -unsatisfiable if for all mappings $\tau : V \rightarrow \Sigma$, the fraction of constraints satisfied by τ is at most $(1 - \eta)$.

Theorem 2. [Din07, Thm. 8.1 and its proof] *There exists a reduction T from 3-SAT instances to 2-CSP instances, with the following properties:*

1. *Completeness: If φ is a satisfiable formula, $T(\varphi)$ is a satisfiable 2-CSP instance;*
2. *Soundness: There exists an absolute constant $\eta > 0$ such that if φ is unsatisfiable, $G = T(\varphi)$ is $(1 - \eta)$ -unsatisfiable;*
3. *Size-Efficiency: If φ has m clauses, then $|V(G)| = n = O(m \cdot \text{polylog } m)$ and also $|E(G)| = O(m \cdot \text{polylog } m)$;*
4. *Alphabet Size: $|\Sigma| = K = O(1)$;*
5. *Regularity: G is d -regular (with self-loops), where $d = O(1)$.*

The last point is not quite explicit in the main statement of Dinur’s result, but can be readily extracted from her proof: simply apply the “preprocessing” transformation of [Din07, Lemma 1.9] to the graph output by her main reduction. Also, Dinur’s main reduction takes as input a constraint graph, not a formula, but we can simply begin by transforming any 3-SAT instance of m clauses into an equivalent instance of an NP-hard 2-CSP such as 3-Colorability, yielding a constraint graph whose number of edges is $O(m)$.

In our protocol, Arthur first performs the above reduction, yielding a 2-CSP G on $n = \tilde{O}(m)$ vertices that is either satisfiable or $(1 - \eta)$ -unsatisfiable. We now describe our BellQMA protocol for the problem, starting directly from the constraint graph G .

3 The BellQMA protocol

Given a constraint graph G , let the proof states Arthur receives be denoted $|\Psi_1\rangle, \dots, |\Psi_{C\sqrt{n}}\rangle$, with C a constant (to be determined later). Each $|\Psi_i\rangle$ consists of a “vertex” register with base states $|0\rangle, \dots, |n-1\rangle$ (describable by $\lceil \log n \rceil$ qubits) and a “color” register with base states $|0\rangle, \dots, |K-1\rangle$ (describable with $\lceil \log K \rceil = O(1)$ qubits). Let $\mu \triangleq C\sqrt{n}/K$. The verifier’s protocol is given below.

Verifier \mathcal{V} :

- Flip a fair coin. If Heads, do the Uniformity Test; if Tails, do the Consistency Test.
- **Uniformity Test:**

1. For each $|\Psi_i\rangle$, perform a Fourier transform F_K on the color register and measure it. Let $Z = \{i : \text{the color register of } |\Psi_i\rangle \text{ is measured } 0\}$. If $|Z| < \frac{99\mu}{100}$, *reject*; otherwise continue.
2. For each $|\Psi_i\rangle$ such that the measurement in Step 1 gets 0, perform a Fourier transform F_n on the vertex register and measure it. If there exists a $|\Psi_i\rangle$ such that the measurement doesn't get 0, *reject*; otherwise *accept*.

• **Consistency Test:**

1. For each $|\Psi_i\rangle$, measure it and denote the value in the two registers as (v_i, c_i) .
2. If there exists two states $|\Psi_i\rangle$ and $|\Psi_j\rangle$ such that $e = (v_i, v_j) \in E$ but $R_e(c_i, c_j) = 0$, *reject*. Also *reject* if $v_i = v_j$ but $c_i \neq c_j$. Otherwise, *accept*.

Note that, since $n = \tilde{O}(m)$, we have $\tilde{O}(m)$ proofs, each consisting of $\log n + O(1) = O(\log m)$ qubits, as needed. The verifier is clearly polynomial-time and performs only nonadaptive, unentangled measurements, so it defines a valid BellQMA protocol.

3.1 Completeness of our protocol

In the rest of the paper, we use \hat{i} to denote the square root of -1 , and reserve the symbol i as an index of states sent by the provers. We first consider the case where the 3-SAT instance φ is satisfiable, so that the constraint graph G is also satisfiable.

Lemma 1. *If G is satisfiable, then there exist (unentangled) states $|\Psi_1\rangle, \dots, |\Psi_{C\sqrt{n}}\rangle$ such that \mathcal{V} accepts with probability at least $1 - \exp(-\mu/(2 \cdot 10^4)) = 1 - \exp(-\Omega(\sqrt{m}))$.*

Proof. Let $|\Psi_i\rangle = |\Psi\rangle \triangleq \frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} |v\rangle |\tau(v)\rangle$ for all $i \leq C\sqrt{n}$, where τ is a coloring satisfying the constraint graph G . Since τ is satisfying, the Consistency Test will accept with probability 1. Below we analyze the probability that the Uniformity Test will accept if that test is performed.

Observe that a Fourier transform on the color register changes $|\Psi\rangle$ into

$$(I_n \otimes F_K) \frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} |v\rangle |\tau(v)\rangle = \frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} |v\rangle \frac{1}{\sqrt{K}} \sum_{k=0}^{K-1} \exp\left(\frac{2\pi \hat{i} \tau(v) k}{K}\right) |k\rangle. \quad (1)$$

Therefore for each $|\Psi_i\rangle$, the measurement in Step 1 of the Uniformity Test will see 0 with probability $n(1/\sqrt{n})^2(1/\sqrt{K})^2 = 1/K$. Accordingly, $\mathbb{E}[|Z|] = C\sqrt{n}/K = \mu$. Since the $|\Psi_i\rangle$'s are unentangled, their measurement outcomes are independent. By a Chernoff bound, the probability that the Uniformity Test passes Step 1 is

$$1 - \Pr\left[|Z| < \frac{99\mu}{100}\right] > 1 - \exp\left(-\frac{\mu}{2 \cdot 10^4}\right).$$

Further notice that according to Eq. 1, conditioned on the color register measuring to 0 in Step 1 of the Uniformity Test, the state in the vertex register of $|\Psi\rangle$ becomes $\frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} |v\rangle$, and a Fourier transform F_n will change this state into

$$F_n \frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} |v\rangle = \frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} \frac{1}{\sqrt{n}} \sum_{u=0}^{n-1} \exp\left(\frac{2\pi \hat{i} v u}{n}\right) |u\rangle = |0\rangle.$$

Thus for each $|\Psi_i\rangle$ which is measured 0 in Step 1 of the Uniformity Test, Step 2 of this test will measure 0 with probability 1. Accordingly, if the Uniformity Test passes Step 1, it will accept in Step 2 with probability 1.

Thus the probability that \mathcal{V} accepts is at least

$$\frac{1}{2} \cdot \left(1 - 2 \exp\left(-\frac{\mu}{2 \cdot 10^4}\right)\right) \cdot 1 + \frac{1}{2} \cdot 1 = 1 - \exp\left(-\frac{\mu}{2 \cdot 10^4}\right).$$

□

4 Soundness of our protocol

Now we consider the case where the 3-SAT instance φ is unsatisfiable, so that the constraint graph G is $(1 - \eta)$ -unsatisfiable. We show that for any sequence of proof states $|\Psi_1\rangle, \dots, |\Psi_{C\sqrt{n}}\rangle$, \mathcal{V} will reject with probability $\Omega(1)$. The proof depends on three lemmas, corresponding to three cases that cover all possible sequences of states sent by the Merlins.

First, we can assume without loss of generality that the states Arthur receives are pure states, since by convexity some sequence of pure states maximizes Arthur's acceptance probability. For each $i \in [C\sqrt{n}]$, we can express $|\Psi_i\rangle$ as

$$|\Psi_i\rangle = \sum_{v=0}^{n-1} \alpha_v^i |v\rangle \sum_{j=0}^{K-1} \beta_{v,j}^i |j\rangle,$$

where $\sum_{v=0}^{n-1} |\alpha_v^i|^2 = 1$ for each i , and $\sum_{j=0}^{K-1} |\beta_{v,j}^i|^2 = 1$ for each i, v .

Let p_0^i be the probability that the color register of $|\Psi_i\rangle$ is measured 0 after the Fourier transform in Step 1 of the Uniformity Test (conditioned on our performing that test). Let

$$Z' \triangleq \left\{ i : p_0^i \geq \frac{1}{4K} \right\}.$$

We claim:

Lemma 2. *If $|Z'| \leq \frac{\mu}{2}$, then Step 1 of the Uniformity Test rejects with probability $\Omega(1)$.*

Proof. Let $Z_1 = Z \cap Z'$ and $Z_2 = Z \setminus Z'$. We have $|Z_1| \leq |Z'| \leq \mu/2$, and $\Pr[i \in Z_2] < 1/(4K)$ independently for every $i \in [C\sqrt{n}]$. Let W be a random subset of $[C\sqrt{n}]$ such that $\Pr[i \in W] = 1/(4K)$ independently for every i . Then $|W|$ stochastically dominates $|Z_2|$ and we have $\mathbb{E}[|W|] = C\sqrt{n}/(4K) = \mu/4$. The probability that Step 1 of the Uniformity Test accepts is

$$\begin{aligned} \Pr\left[|Z| \geq \frac{99\mu}{100}\right] &= \Pr\left[|Z_1| + |Z_2| \geq \frac{\mu}{2} + \frac{\mu}{4} + \frac{24\mu}{100}\right] \\ &\leq \Pr\left[|Z_2| \geq \frac{\mu}{4} + \frac{24\mu}{100}\right] \\ &\leq \Pr\left[|W| \geq \frac{\mu}{4} + \frac{24\mu}{100}\right] \\ &\leq \exp\left(-\frac{24^2}{25^2 \cdot 2} \cdot \frac{\mu}{4}\right) = o(1), \end{aligned}$$

where we used a Chernoff bound. □

Let $\varepsilon < \eta/20$ be a constant (recall that η is the soundness constant in Dinur's PCP reduction), and for each $i \in [C\sqrt{n}]$, let

$$R_i \triangleq \{v : v \in V, |\alpha_v^i|^2 < 1/(8Kn)\}.$$

The next lemma considers the case where one of the sets R_i is noticeably large.

Lemma 3. *Suppose there exists $i \in Z'$ such that $|R_i| \geq \varepsilon n$. Then the Uniformity Test rejects with probability $\Omega(1)$.*

Proof. We focus on any such index i . After the Fourier transform on the color register in Step 1 of the Uniformity Test, $|\Psi_i\rangle$ becomes the state $|\Phi_i\rangle$ defined by

$$\begin{aligned} |\Phi_i\rangle &\triangleq (I_n \otimes F_K) \sum_{v=0}^{n-1} \alpha_v^i |v\rangle \sum_{j=0}^{K-1} \beta_{v,j}^i |j\rangle = \sum_{v=0}^{n-1} \alpha_v^i |v\rangle \sum_{j=0}^{K-1} \beta_{v,j}^i \frac{1}{\sqrt{K}} \sum_{k=0}^{K-1} \exp\left(\frac{2\pi i j k}{K}\right) |k\rangle \\ &= \frac{1}{\sqrt{K}} \sum_{k=0}^{K-1} \left(\sum_{v=0}^{n-1} \alpha_v^i \left(\sum_{j=0}^{K-1} \beta_{v,j}^i \exp\left(\frac{2\pi i j k}{K}\right) \right) |v\rangle \right) |k\rangle. \end{aligned} \quad (2)$$

Let $|\gamma\rangle = \sum_{v=0}^{n-1} \gamma_v^i |v\rangle$ with $\sum_{v=0}^{n-1} |\gamma_v^i|^2 = 1$ be the state left in the vertex register of $|\Phi_i\rangle$, after conditioning on the color register of $|\Phi_i\rangle$ measuring to 0, which occurs with probability p_0^i by definition. For each $v \in \{0, 1, \dots, n-1\}$, let $P_{0,v}^i$ be the probability that the color register of $|\Phi_i\rangle$ is measured 0 and the vertex register of $|\Phi_i\rangle$ is measured v . We have that

$$P_{0,v}^i = p_0^i \cdot |\gamma_v^i|^2.$$

On the other hand, by Eq. 2 we have

$$P_{0,v}^i = \left| \frac{\alpha_v^i}{\sqrt{K}} \sum_{j=0}^{K-1} \beta_{v,j}^i \right|^2 = \frac{|\alpha_v^i|^2}{K} \left| \sum_{j=0}^{K-1} \beta_{v,j}^i \right|^2 \leq \frac{|\alpha_v^i|^2}{K} \cdot K \sum_{j=0}^{K-1} |\beta_{v,j}^i|^2 = |\alpha_v^i|^2,$$

where we used the Cauchy-Schwarz inequality and the fact that $\sum_{j=0}^{K-1} |\beta_{v,j}^i|^2 = 1$. Combining the above two equations, we find $p_0^i \cdot |\gamma_v^i|^2 \leq |\alpha_v^i|^2$. Because $i \in Z'$, we have $p_0^i \geq 1/(4K)$. Thus $|\gamma_v^i|^2 \leq 4K|\alpha_v^i|^2$ for each v . Accordingly, for each $v \in R_i$,

$$|\gamma_v^i|^2 \leq \frac{4K}{8Kn} = \frac{1}{2n}.$$

Define

$$|\psi\rangle \triangleq F_n^{-1}|0\rangle = \frac{1}{\sqrt{n}} \sum_{v=0}^{n-1} |v\rangle.$$

For each $v \in R_i$, $|\gamma_v^i - 1/\sqrt{n}|^2 \geq |1/\sqrt{2n} - 1/\sqrt{n}|^2 = (1 - 1/\sqrt{2})^2/n$. Then, using unitarity of F_n , we have

$$\|F_n|\gamma\rangle - |0\rangle\|_2^2 = \|\gamma - |\psi\rangle\|_2^2 \geq \sum_{v \in R_i} (1 - 1/\sqrt{2})^2/n \geq \varepsilon n \cdot (1 - 1/\sqrt{2})^2/n = \Omega(1).$$

Since $\|F_n|\gamma\rangle\|_2 = 1$, it follows that the amplitude of $|0\rangle$ in $F_n|\gamma\rangle$ is of norm $\leq 1 - \Omega(1)$. Thus if the color register of $|\Phi^i\rangle$ measures 0 in the Uniformity Test (as happens with probability $p_0^i \geq 1/(4K)$ since $i \in Z'$), the vertex register measures to some $v \neq 0$ with probability $\Omega(1)$. The Uniformity Test's rejection probability is therefore $\Omega(1/K) = \Omega(1)$ as claimed. \square

In light of Lemmas 2 and 3, we need only to address the case when $|Z'| > \mu/2$ and $|R_i| < \varepsilon n$ for all $i \in Z'$. We show that in this case, the Consistency Test rejects with probability $\Omega(1)$.

Consider an arbitrary state index $i \in Z'$. Let D_i denote the distribution on vertex/color pairs when $|\Psi_i\rangle$ is measured by the Consistency Test. We can equivalently generate each D_i as $D_i = g_i(U_i)$, where each U_i is a uniform, independent value from $[0, 1]$, and $g_i : [0, 1] \rightarrow V(G) \times \Sigma$ is a function such that each preimage $g_i^{-1}((v, c))$ is an interval of length equal to $\Pr[D_i = (v, c)]$. Then for each $v \notin R_i$, the set $g_i^{-1}(v, \star)$ is of measure $|\alpha_v^i|^2 \geq 1/(8Kn)$. Select $J_{i,v} \subseteq g_i^{-1}(v, \star)$ of measure *exactly* $1/(8Kn)$ for each such v , and let $J_i = \bigcup_{v \notin R_i} J_{i,v}$. Observe the following: first, J_i has measure greater than $(1 - \varepsilon)/(8K)$. Second, conditioned on $U_i \in J_i$, the posterior distribution of the vertex v_i that g_i outputs is now *uniform* over $S_i \triangleq \{0, 1, \dots, n-1\} \setminus R_i$.

So let us consider the Consistency Test applied to a sequence of states satisfying $|Z'| > \mu/2$ and $|R_i| < \varepsilon n$ for all $i \in Z'$. Letting the measurement outcomes be generated as described above, define the random set

$$Z'' \triangleq \{i : i \in Z', U_i \in J_i\}.$$

Notice that Z'' is itself a random variable determined by the U_i 's. Notice also that for each $i \in Z'$, the probability that $i \in Z''$ is at least $(1 - \varepsilon)/(8K)$, and these events are independent from each other. Therefore we have that

$$\mathbb{E}[|Z''|] \geq \frac{(1 - \varepsilon)|Z'|}{8K} > \frac{(1 - \varepsilon)\mu}{16K} = \frac{(1 - \varepsilon)C\sqrt{n}}{16K^2}.$$

Since $|Z''|$ never exceeds $C\sqrt{n}$, the total number of proof states, we find that with probability $\Omega(1)$,

$$|Z''| \geq \frac{C\sqrt{n}}{32K^2}. \quad (3)$$

The following lemma tells us that if C is chosen as a suitably large constant, then conditioned on Eq. 3 holding, the Consistency Test rejects with $\Omega(1)$ probability.

Lemma 4. *Let $(G, \{R_e\})$ be an n -vertex, d -regular constraint graph (possibly with self-loops, and $d > 1$) with alphabet K , such that G is $(1 - \eta)$ -unsatisfiable. Let $D_1, \dots, D_{m'}$ be independent distributions on $V(G) \times \Sigma$, with (v_i, c_i) denoting the output of D_i . Suppose for each $i \leq m'$ there exists an $S_i \subseteq V(G)$ of size at least $(1 - \varepsilon)n$, such that v_i is uniformly distributed over S_i , where $\varepsilon < \eta/20$.*

Then we can set $m' = O(\sqrt{n})$ large enough so that with probability at least .99 there exists $i < j \leq m'$ such that: either $e = (v_i, v_j)$ is an edge of G and $R_e(c_i, c_j) = 0$; or $v_i = v_j$, $c_i \neq c_j$. (The constant in the $O()$ notation depends on d and η , but not K .)

To apply Lemma 4 to our Consistency Test when $|Z'| > \mu/2$ and $|R_i| < \varepsilon n$ for all $i \in Z'$, choose $C = O(1)$ such that $C\sqrt{n}/(32K^2) \geq m'$. Then conditioned on $|Z''| \geq m'$, which occurs with probability $\Omega(1)$, we can select $D_1, \dots, D_{m'}$ from the distributions of $|\Psi_i\rangle$'s (when measured in the Consistent Test) such that $i \in Z''$. By definition of Z'' , these $D_1, \dots, D_{m'}$ satisfy the hypothesis of Lemma 4, so the Consistency Test rejects with probability .99. Thus in this case \mathcal{V} also rejects with probability $\Omega(1)$. This completes the proof that \mathcal{V} possesses soundness $1 - \Omega(1)$, proving Theorem 1.

Proof of Lemma 4. For $i < j \leq m'$, let $V_{i,j}$ be the indicator random variable for the event that either $e = (v_i, v_j) \in E(G)$ and $R_e(c_i, c_j) = 0$, or $v_i = v_j$ and $c_i \neq c_j$. Let $V = \sum_{i < j} V_{i,j}$. To

prove Lemma 4 it is enough to show that $\Pr[V = 0] \leq .01$. We show this using the second moment method.

A first observation is that we can generate D_i in the following way: first, randomly select a coloring τ_i according to some distribution H_i ; next select v_i uniformly from S_i , and set $c_i = \tau_i(v_i)$. To be explicit, each H_i independently chooses colors according to the rule $\Pr[\tau_i(v) = c] = \Pr[c_i = c | v_i = v]$. It is easily verified that this process yields D_i .

Next we lower-bound $\mathbb{E}[V] = \sum_{i < j} \mathbb{E}[V_{i,j}]$. Fix any pair i, j , $1 \leq i < j \leq m'$. Condition on any values of the colorings τ_i, τ_j ; we'll show that $\mathbb{E}[V_{i,j} | \tau_i, \tau_j] \geq \varepsilon/n$. Let $P_{i,j} \subseteq V(G)$ be the subset of vertices v for which $\tau_i(v) = \tau_j(v)$. Suppose first that $|P_{i,j}| \leq (1 - 3\varepsilon)n$. In this case there are at least εn vertices contained in $S_i \cap S_j \cap \overline{P_{i,j}}$, and $V_{i,j} = 1$ whenever a vertex in this set is selected as both v_i and v_j . Thus in this case $\mathbb{E}[V_{i,j} | \tau_i, \tau_j] \geq \varepsilon n \cdot |S_i|^{-1} \cdot |S_j|^{-1} \geq \varepsilon/n$.

For our second case, suppose $|P_{i,j}| > (1 - 3\varepsilon)n$. Consider the induced subgraph $G[S_i \cap S_j \cap P_{i,j}]$, which contains at least $n - 2\varepsilon n - 3\varepsilon n = (1 - 5\varepsilon)n$ vertices. Since G has maximum degree d , $|E(G)| = dn/2$, and the set $\overline{S_i} \cup \overline{S_j} \cup \overline{P_{i,j}}$ is incident on at most $d(5\varepsilon n)$ edges, we have that $|E(G[S_i \cap S_j \cap P_{i,j}])| \geq dn/2 - 5d\varepsilon n = (1 - 10\varepsilon)(dn)/2 = (1 - 10\varepsilon)|E(G)|$. By $(1 - \eta)$ -unsatisfiability of $(G, \{R_e\})$, the coloring τ_i satisfies at most a $(1 - \eta)/(1 - 10\varepsilon)$ fraction of the edge constraints in $G[S_i \cap S_j \cap P_{i,j}]$. Thus the fraction of these constraints which are *violated* by τ_i is at least $1 - \frac{1 - \eta}{1 - 10\varepsilon} = \frac{1 - 10\varepsilon - (1 - \eta)}{1 - 10\varepsilon} > \frac{\eta}{2(1 - 10\varepsilon)} > \frac{\eta}{2}$, since $\eta > 20\varepsilon$.

Now $\tau_j \equiv \tau_i$ on $P_{i,j}$. We can thus lower-bound $\mathbb{E}[V_{i,j} | \tau_i, \tau_j]$ by the probability that $v_i, v_j \in S_i \cap S_j \cap P_{i,j}$ and that v_i, v_j form (in either order) an edge violated by the color assignment $(\tau_i(v_i), \tau_j(v_j))$. Note that some edges are self-loops and so may only be chosen in one way. We get

$$\begin{aligned} \mathbb{E}[V_{i,j} | \tau_i, \tau_j] &\geq (1 - 5\varepsilon)^2 \cdot \frac{\frac{\eta}{2}|E(G[S_i \cap S_j \cap P_{i,j}])|}{n^2} \\ &\geq (1 - 5\varepsilon)^2 \cdot \frac{\eta(1 - 10\varepsilon)|E(G)|}{2n^2} = \frac{\eta(1 - 5\varepsilon)^2(1 - 10\varepsilon)d}{4n}. \end{aligned}$$

Recall that $\varepsilon < \eta/20 < 1/20$, so the quantity above is greater than $\eta \cdot 2^{-3}d/4n > \eta/(20n) > \varepsilon/n$, as needed (using $d > 1$). Thus in either of our two cases we conclude $\mathbb{E}[V_{i,j} | \tau_i, \tau_j] \geq \varepsilon/n$, so $\mathbb{E}[V_{i,j}] \geq \varepsilon/n$ unconditioned as well. Summing over all $i < j$, we find $\mathbb{E}[V] \geq \varepsilon \binom{m'}{2}/n = \Omega((m')^2/n)$.

Next we upper-bound $\mathbb{E}[V^2] = \sum_{i < j, k < l} \mathbb{E}[V_{i,j}V_{k,l}]$. There are $\binom{m'}{2}$ terms for which $(i, j) = (k, l)$. For each such term $\mathbb{E}[V_{i,j}^2] = \mathbb{E}[V_{i,j}]$. Condition on the vertex v_i outputted by D_i . Fixing any such choice of v_i , the probability that $V_{i,j} = 1$ is of course upper-bounded by the probability that v_j is equal or adjacent to v_i in G . Since v_j is uniform on S_j and v_i is of degree d , this probability is at most $(d + 1)/|S_i| \leq (d + 1)/((1 - \varepsilon)n)$, so $\mathbb{E}[V_{i,j} | v_i] \leq (d + 1)/((1 - \varepsilon)n)$. As v_i was an arbitrary conditioning, we conclude $\mathbb{E}[V_{i,j}] \leq (d + 1)/((1 - \varepsilon)n)$. Thus the contribution to $\mathbb{E}[V^2]$ from terms where $(i, j) = (k, l)$ is at most $\binom{m'}{2}(d + 1)/((1 - \varepsilon)n) = O((m')^2/n)$.

If $(i, j), (k, l)$ consists of three distinct indices, assume that $j = l$, the other cases being handled similarly. Condition on any choice of v_j . Then $V_{i,j}V_{j,k} = 1$ can only occur if v_j and v_k are each either incident on or equal to v_i . These two events are independent after conditioning on v_j since D_i, D_j, D_k are independent. Thus $\mathbb{E}[V_{i,j}V_{j,k}] \leq [(d + 1)/((1 - \varepsilon)n)]^2$.

For any three distinct indices $a < b < c \leq m'$, there are six tuples $(i < j), (k < l)$ for which $\{i, j, k, l\} = \{a, b, c\}$. Thus the contribution to $\mathbb{E}[V^2]$ from these ‘‘triplet’’ terms is at most $6 \binom{m'}{3} \cdot [(d + 1)/((1 - \varepsilon)n)]^2 = O((m')^3/n^2)$.

If $(i, j), (k, l)$ are four distinct elements of $[m']$, then the pair $V_{i,j}, V_{j,k}$ depend on disjoint sets of independent random variables, so that $\mathbb{E}[V_{i,j}V_{k,l}] = \mathbb{E}[V_{i,j}]\mathbb{E}[V_{k,l}]$. Thus the contribution to $\mathbb{E}[V^2]$ from these terms is upper-bounded by $\sum_{i<j,k<l} \mathbb{E}[V_{i,j}]\mathbb{E}[V_{k,l}] = \mathbb{E}[V]^2$.

Putting things together,

$$\mathbb{E}[V^2] < O\left(\frac{(m')^2}{n} + \frac{(m')^3}{n^2}\right) + \mathbb{E}[V]^2.$$

With this bound in hand, we apply Chebyshev's inequality:

$$\begin{aligned} \Pr[V = 0] &\leq \Pr[|V - \mathbb{E}[V]| \geq \mathbb{E}[V]] \\ &\leq \frac{\mathbb{E}[V^2] - \mathbb{E}[V]^2}{\mathbb{E}[V]^2} \\ &\leq O\left(\frac{n^2}{(m')^4} \left(\frac{(m')^2}{n} + \frac{(m')^3}{n^2}\right)\right) \\ &= O\left(\frac{n}{(m')^2} + \frac{1}{m'}\right), \end{aligned}$$

which is at most .01 if we take m' to be a suitably large value in $O(\sqrt{n})$. This proves the lemma. \square

5 Acknowledgements

We thank Scott Aaronson for helpful comments and suggestions.

References

- [ABD⁺09] S. Aaronson, S. Beigi, A. Drucker, B. Fefferman, and P. Shor. The power of unentanglement. *Theory of Computing*, 5(1):1–42, 2009. Earlier version in Complexity'08.
- [BBPS96] C. Bennett, H. Bernstein, S. Popescu, and B. Schumacher. Concentrating partial entanglement by local operations. *Phys. Rev. A*, 53(4):2046–2052, Apr 1996.
- [BCWdW01] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, Sep 2001. arXiv:quant-ph/0102001v1.
- [BCY10] F. Brandao, M. Christandl, and J. Yard. Faithful squashed entanglement, 2010. arXiv:1010.1750.
- [Bei10] S. Beigi. NP vs $\text{QMA}_{\log}(2)$. *Quantum Information and Computation*, 54(1&2):0141–0151, 2010. arXiv:0810.5109.
- [Bra08] F. Brandao. *Entanglement Theory and the Quantum Simulation of Many-Body Physics*. PhD thesis, Imperial College, London, 2008. arXiv:0810.0026.
- [BT09] H. Blier and A. Tapp. All languages in NP have very short quantum proofs. In *Proc. of 3rd International Conference on Quantum, Nano and Micro Technologies*, 2009. arXiv:0709.0738.

- [Din07] I. Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12, 2007. Earlier version in STOC'06.
- [HM10] A. Harrow and A. Montanaro. An efficient test for product states, with applications to quantum Merlin-Arthur games, 2010. arXiv:1001.0017v3. To appear in FOCS'10.
- [KMY03] H. Kobayashi, K. Matsumoto, and T. Yamakami. Quantum Merlin-Arthur proof systems: Are multiple Merlins more helpful to Arthur? In Toshihide Ibaraki, Naoki Katoh, and Hirotaka Ono, editors, *ISAAC*, volume 2906 of *Lecture Notes in Computer Science*, pages 189–198. Springer, 2003.
- [Nie99] M. Nielsen. Conditions for a class of entanglement transformations. *Phys. Rev. Lett.*, 83(2):436–439, Jul 1999.