Quantum Information and the PCP Theorem

Ran Raz*
Weizmann Institute
ran.raz@weizmann.ac.il

Abstract

Our main result is that the membership $x \in SAT$ (for x of length n) can be proved by a logarithmic-size quantum state $|\Psi\rangle$, together with a polynomial-size classical proof consisting of blocks of length $polylog(n)$ bits each, such that after measuring the state $|\Psi\rangle$ the verifier only needs to read **one** block of the classical proof. This shows that if a short quantum witness is available then a (classical) PCP with only one query is possible.

Our second result is that the class $QIP/qpoly$ contains **all** languages. That is, for any language L (even non-recursive), the membership $x \in L$ (for x of length n) can be proved by a polynomial-size quantum interactive proof, where the verifier is a polynomial-size quantum circuit with working space initiated with some quantum state $|\Psi_{L,n}\rangle$ (depending only on L and n). Moreover, the interactive proof that we give is of only one round, and the messages communicated are classical. The advice $|\Psi_{L,n}\rangle$ given to the verifier can also be replaced by a classical probabilistic advice, as long as this advice is kept as a secret from the prover. Our result can hence be interpreted as: the class $IP/rpoly$ contains all languages.

For the proof of the second result, we introduce the *quantum low-degree-extension* of a string of bits. The main result requires an additional machinery of *quantum low-degree-test*.

1 Introduction

1.1 The Main Result

A probabilistic checkable proof (PCP) is a proof that can be (probabilistically) verified by reading only a small portion of it. The PCP theorem [BFL, FGLSS, AS1, ALMSS] states that for any $x \in SAT$ (where x is an input of length n), there is a PCP p for the membership $x \in SAT$, such that the proof p is of length $poly(n)$ bits and it can be (probabilistically)

*Research supported by Israel Science Foundation (ISF) grant.

verified by reading only a constant number of its bits. Moreover, there is a PCP p for the membership $x \in SAT$, such that the proof p is consisted of $poly(n)$ blocks of length $O(1)$ bits each and it can be (probabilistically) verified by reading only **two** of its blocks. A similar PCP that can be verified by reading only one of its blocks is obviously impossible, under standard hardness assumptions (even if we allow the length of each block to be almost linear).

In this paper, we show that the membership $x \in SAT$ (for x of length n) can be proved by a logarithmic-size quantum state $|\Psi\rangle$, together with a polynomial-size classical proof p consisting of blocks of length $polylog(n)$ bits each, such that after measuring the state $|\Psi\rangle$ the verifier only needs to read **one** block of the proof p.

More precisely, the verifier can be modelled by a polynomial-size quantum circuit. For any $x \in SAT$, there exists a logarithmic-size quantum state $|\Psi\rangle$ and an array p of $poly(n)$ blocks of length $polylog(n)$ bits each, that encode a proof for the membership $x \in SAT$ and can be verified as follows: The verifier applies on $|\Psi\rangle$ a carefully designed (probabilistic) unitary transformation U (that can be computed in quantum logarithmic time). The verifier measures some of the qubits of $U|\Psi\rangle$. Denote the collapsed state (after the measurement) by $|\Psi'\rangle$. Based on x and on the result of the measurement, the verifier composes a (classical) query q (of length $O(\log n)$ bits) and reads the qth block of p. Denote the value of that block by r. Based on x, q, r, the verifier applies a unitary transformation U' (that can be computed in quantum logarithmic time) on $|\Psi'\rangle$ and measures all bits of $U'|\Psi'\rangle$. Based on the result of the measurement, the verifier decides whether to *Accept* or *Reject*, where *Accept* is interpreted as $x \in SAT$ and *Reject* is interpreted as a declaration that the proof $(|\Psi\rangle, p)$ is not correct. We will have the following (standard) completeness and soundness properties (for any fixed constant $\epsilon > 0$):

1. For any $x \in SAT$, there exist $|\Psi\rangle$ and p that cause the verifier to accept with probability 1.

2. For any $x \notin SAT$, and any $|\Psi\rangle$ and p, the verifier rejects with probability $\geq 1 - \epsilon$.

1.2 The Information of a Quantum State

Our main result is only possible because the logarithmic length quantum state $|\Psi\rangle$ (from Subsection 1.1) contains information about all bits of the classical proof p. We hence proceed with a short discussion of the information of a quantum state.

A quantum state of n qubits contains an infinite amount of information. If the state is only given up to some fixed (say, constant) precision it still contains an exponential amount of information. On the other hand, a quantum measurement can only give n bits of information about the state. One way to formalize the last statement is given by Holevo's theorem [Hol].

A simplified version of Holevo's theorem can be stated as follows: If n (classical) bits $a_1, ..., a_n$ are encoded by a single quantum state $|\Psi\rangle = |\Psi(a_1, ..., a_n)\rangle$, such that the original values of the bits $a_1, ..., a_n$ can be retrieved from the state $|\Psi\rangle$, then $|\Psi\rangle$ is a state of at least

n qubits. In other words: Assume that Bob encodes n bits $a_1, ..., a_n$ by a quantum state $|\Psi\rangle$ and sends $|\Psi\rangle$ to Alice. Assume that Alice can retrieve the original values of $a_1, ..., a_n$ by measuring the state $|\Psi\rangle$. Then, $|\Psi\rangle$ is a state of at least n qubits. Moreover, if we only require that each a_i is retrieved correctly with probability $1 - \epsilon$, and allow an error to occur with probability ϵ, then $|\Psi\rangle$ is a state of at least $(1 - \mathrm{H}(\epsilon)) \cdot n$ qubits, where $\mathrm{H}(\epsilon)$ denotes the Shannon's entropy of the distribution $(\epsilon, 1 - \epsilon)$.

A strengthening of Holevo's theorem was suggested by Ambainis, Nayak, Ta-Shma, and Vazirani [ANTV] and was proved by Nayak [Nay]. A simplified version of Nayak's theorem can be stated as follows: Assume that Bob encodes n bits $a_1, ..., a_n$ by a quantum state $|\Psi\rangle$ and sends $|\Psi\rangle$ to Alice. Assume that for every index $i \in \{1, ..., n\}$ (of her choice), Alice can retrieve the original value of a_i by measuring the state $|\Psi\rangle$. Then, $|\Psi\rangle$ is a state of at least n qubits. Moreover, if we only require that Alice retrieves a_i correctly with probability $1 - \epsilon$, and allow an error to occur with probability ϵ, then $|\Psi\rangle$ is a state of at least $(1 - \mathrm{H}(\epsilon)) \cdot n$ qubits.

Note that the difference between Holevo's theorem and Nayak's theorem is that in Holevo's theorem we require that Alice can retrieve the values of **all** the original bits, whereas in Nayak's theorem we only require that Alice can retrieve the value of **one** bit of her choice. Note that by the uncertainty principle these two tasks are not necessarily equivalent. It was demonstrated in [ANTV] that the two tasks are indeed not equivalent.

In this paper, we suggest a protocol that works as follows: Bob will encode a large number of (classical) bits by a very short quantum state and will send that state to Alice. Alice will not be able to retrieve even one of the original bits by herself. Nevertheless, the value of each one of the original bits can be retrieved by an Arthur-Merlin protocol, with a third party, the infinitely powerful prover Merlin. In this protocol, Alice acts as the verifier Arthur. In other words, although Alice is not able to retrieve the value of the ith bit by herself, Merlin will tell her that value and will be able to convince her that this value is correct. Note that in this setting Bob is completely trustable and hence Alice can count that the quantum state given by Bob correctly encodes the original bits. Merlin, on the other hand, cannot be trusted and hence Alice needs to be convinced that his answer is correct.

Interestingly, the communication between Alice and Merlin in our protocol will be **classical**. They will not need to exchange quantum states. We can hence assume w.l.o.g. that Merlin is an infinitely powerful **classical** computer. Alice, on the other hand, will need to have the ability to measure the quantum state sent by Bob[1], but her computational power will be polynomially bounded (as required in an Arthur-Merlin protocol).

More precisely, we will construct a protocol that works as follows: Bob encodes 2^n (classical) bits $a_1, ..., a_{2^n}$ by a quantum state $|\Psi\rangle = |\Psi(a_1, ..., a_{2^n})\rangle$ of size $O(n)$ qubits, and sends $|\Psi\rangle$ to Alice. Alice measures the state $|\Psi\rangle$. Given an index $i \in \{1, ..., 2^n\}$ (of her choice), and based on the result of the measurement, Alice composes a (classical) question q of length $poly(n)$ bits and sends (i, q) to Merlin. After seeing (i, q), Merlin responds with a (classical)

[1] See however the discussion in Subsection 1.2.1

answer r of length $poly(n)$ bits. Based on i, q, r and the result of the measurement, Alice decides on a value $V \in \{0, 1, Err\}$, where 0 is interpreted as $a_i = 0$ and 1 is interpreted as $a_i = 1$, and Err is interpreted as a declaration that Merlin is cheating. We will have the following (standard) completeness and soundness properties for this protocol:

1. For any i, q, there is an answer r, such that $V = a_i$ (with probability 1).
2. For any i, q, r, we have that $V \in \{a_i, Err\}$ with probability $\geq 1 - 1/n^{\Omega(1)}$.

In other words, for any index i and question q, Merlin will be able to give an answer r that causes Alice to conclude the correct value of a_i, and on the other hand, no answer given by Merlin can cause Alice to conclude the incorrect value of a_i (with non-negligible probability).

Our results are in fact more general: We will be able to encode and retrieve $a_1, ..., a_{2^n}$ that can get $n^{O(1)}$ different values, rather than bits (i.e., each a_i can be a block of $O(\log n)$ bits). Moreover, we will be able to retrieve any constant number of values $a_{i_1}, ..., a_{i_k}$.

1.2.1 The Protocol is Actually Classical

Interestingly, the above protocol can be presented as a classical protocol[2]. In the above protocol, Bob sends a quantum state $|\Psi\rangle$ and Alice measures all qubits of that state. Instead, Bob can just send the result of the measurement. Since the result of the measurement is just a classical probabilistic string, Bob can produce a classical string chosen randomly according to the same distribution and send that string to Alice. Quantum power is hence not needed. We note, however, that the classical protocol works only if the classical probabilistic string supplied by Bob is kept as a secret from the prover Merlin. We note also that for our main PCP result we do need the quantum version of the protocol, as we will need to apply an additional machinery of *quantum low degree test*.

1.3 The Exceptional Power of $QIP/qpoly$

Interactive proofs were introduced by Goldwasser, Micali and Rackoff and by Babai and Moran [GMR, Bab, BM], and were extended to the quantum case by Watrous [Wat]. The simplest version of a quantum interactive proof is a **one-round** (i.e., two messages) quantum interactive proof, usually called a $QIP(2)$ proof.

In a $QIP(2)$ proof, the infinitely powerful prover Merlin tries to convince the verifier Arthur for a membership $x \in L$, (where L is some language and x is an input of length n, and both x and L are known to both parties). Both parties have quantum computers and they can communicate between them quantum states. Merlin's computational power is unlimited (but he must obey the laws of physics). Arthur's computational power, on the other hand,

[2]This was communicated to us by many colleagues.

is limited to (quantum) polynomial time. The proof has one round of communication, where the two parties exchange between them quantum states.

In this paper, we will not need the full power of $QIP(2)$ proofs. We will use a subclass of proofs that we call $QIP(2)^*$ proofs. In a $QIP(2)^*$ proof, Arthur and Merlin communicate between them **classical** messages, rather than quantum states. We can hence assume w.l.o.g. that Merlin is an infinitely powerful **classical** computer. Arthur, on the other hand, will need to have the ability to work with quantum states (in order to be able to work with the quantum advice discussed below).

A $QIP(2)^*$ proof has one round of communication: Based on x (and possibly on a random string), Arthur composes a classical question q of length $poly(n)$ bits and sends q to Merlin. After seeing (x, q), Merlin responds with a classical answer r of length $poly(n)$ bits. Based on x, q, r, Arthur decides on a value $V(x, q, r) \in \{Accept, Reject\}$, where $Accept$ is interpreted as $x \in L$ and $Reject$ is interpreted as a declaration that Merlin is cheating. The following completeness and soundness properties should be satisfied[3] (for some small constant ϵ):

1. For any $x \in L$ and any q, there is an answer r, such that $V(x, q, r) = Accept$, with probability $\geq 1 - \epsilon$.
2. For any $x \notin L$ and any q, r, we have that $V(x, q, r) = Reject$, with probability $\geq 1 - \epsilon$.

In other words, if $x \in L$ then for any question q Merlin will be able to give an answer r that causes Arthur to accept (with high probability), and on the other hand, if $x \notin L$ then for any question q no answer given by Merlin can cause Arthur to accept (with non-negligible probability).

In this paper, we are interested in the class $QIP(2)^*/qpoly$, that is, the class of languages that have polynomial-size $QIP(2)^*$ proofs with a **polynomial-size quantum advice**. A $QIP(2)^*/qpoly$ proof is the same as a $QIP(2)^*$ proof, except that the computational power of Arthur is quantum polynomial time with a polynomial-size quantum advice. In other words, Arthur is a quantum circuit in $BQP/qpoly$. We can think of a circuit in $BQP/qpoly$ as a polynomial-size quantum circuit with working space initiated with an arbitrary quantum state $|\Psi_{L,n}\rangle$ (depending only on L and n). We think of the state $|\Psi_{L,n}\rangle$ as a (polynomial-size) *quantum advice* (given to the verifier).

The notion of quantum advice was studied in several previous works [NY, Aar], as a quantum analog to the notion of classical advice (or classical non-uniformity). These works concentrated on the class $BQP/qpoly$ and proved some limitations of that class. In particular, Aaronson proved that the class $BQP/qpoly$ is contained in the classical class $PP/poly$ [Aar].

We show that the class $QIP(2)^*/qpoly$ contains **all** languages. That is, for any language L, there is a polynomial-size $QIP(2)^*/qpoly$ interactive proof for the membership $x \in L$. Since

[3] In the protocols constructed in this paper, we will actually have perfect completeness, i.e., $\epsilon = 0$ in the first item bellow.

any $QIP(2)^*/qpoly$ proof is also a $QIP(2)/qpoly$ proof, this obviously means that that the class $QIP(2)/qpoly$, and hence also $QIP/qpoly$, contain all languages.

We note that it is relatively easy to see that the class $PP/qpoly$ also contains all languages. This was first observed by Aaronson [Aar].

1.3.1 The Class $IP/rpoly$

As in Subsection 1.2, the result of this subsection can also be presented as a classical result. Arthur and Merlin do not really need to have quantum power and they only exchange between them classical bits. The quantum advice given to Arthur can be replaced by a classical probabilistic string. Equivalently, we can think of Arthur as a distribution over polynomial size classical circuits. Our result can hence be interpreted as: the class $IP/rpoly$ contains all languages (where IP stands for Interactive Proofs and $rpoly$ stands for a randomized polynomial size advice). We note, however, that this is true only if the advice given to Arthur is kept as a secret from Merlin. (If we think of Arthur as a distribution over polynomial size classical circuits then the exact circuit chosen from the distribution is kept as a secret from Merlin). The classical result is true only if the class $IP/rpoly$ is defined accordingly.

1.4 Methods

We combine methods previously used in the field of probabilistic checkable proofs and methods previously used in the field of quantum computations, together with some new ideas.

The quantum state $|\Psi(a_1, ..., a_{2^n})\rangle$, from Subsection 1.2, is a quantum representation of the so called, *low degree extension*, of $a_1, ..., a_{2^n}$. Low degree extensions were extensively used in the past in the study of randomness and derandomization and probabilistic checkable proofs. For the retrieval protocol of Subsection 1.2, we use the random self reducibility property and the locally decodability property of the low degree extension.

For the results discussed in Subsection 1.3, we will use as a quantum advice the quantum state $|\Psi(a_0, ..., a_{2^n-1})\rangle$, where $a_i = 1$ iff $i \in L$. The results of Subsection 1.3 will then follow immediately from the ones of Subsection 1.2.

For the results discussed in Subsection 1.1, we will use the quantum state $|\Psi(a_1, ..., a_m)\rangle$, where $(a_1, ..., a_m)$ is a (classical) PCP for the membership $x \in SAT$. Note, however, that in the setting of Subsection 1.1 the verifier cannot assume anything about the quantum state given to him, as it is given by the (un-trusted) prover. The verifier cannot even trust that the quantum state given to him is a correct representation of the low degree extension of some sequence of bits. A key step in our analysis will be a *quantum low degree test* that will ensure that the state is close to a quantum representation of some multivariate polynomial of low degree. Since this seems to be impossible for the verifier to do by himself, the test is done with the help of a classical PCP (or equivalently, with the help of a classical prover).

Note that in the setting of Subsection 1.1, the verifier cannot query the classical proof

more than once. Moreover, the verifier can measure the quantum state only once (as the state collapses after the measurement). Hence, the verifier cannot apply both the quantum low degree test and the retrieval protocol. We will hence need to integrate these two tasks. We will do that using ideas from [DFKRS]. A special attention is given to the probability of error, as we would like to keep it as small as possible (and in particular, sub-constant).

Most of the technical work in the paper is done in the proofs of the results discussed in Subsection 1.1 (including the proof for the correctness of the quantum low degree test).

1.5 Discussion

The PCP style results of Subsection 1.1 scale up to languages in $NEXP$. More precisely, for any language $L \in NEXP$, the membership $x \in L$ (for x of length n) can be proved by a polynomial-size quantum state $|\Psi\rangle$, together with an exponential-size classical proof p consisting of blocks of length $poly(n)$ bits each, such that after measuring the state $|\Psi\rangle$ the verifier only needs to read one block of the proof p.

There are several alternative ways to present the last result. One of them is the following: Consider a two-rounds interactive proofs model, where the prover has **quantum** power in the first round but only **classical** power in the second round (note that in the second round the prover still has an infinitely powerful classical computer, but he cannot access any quantum state). Then, for any language $L \in NEXP$, the membership $x \in L$ (for x of length n) can be proved by a polynomial-size interactive proof in this model.

Note that $IP = PSPACE$ [LFKN, Sha], and $QIP \subset EXP$ [KW]. Thus, if the prover has classical power in both rounds or quantum power in both rounds we are not likely to be able to prove memberships $x \in L$ even for languages $L \in NTIME(n^{\log n})$. In contrast, if the prover has quantum power in the first round and classical power in the second we are able to prove memberships $x \in L$ for any $L \in NEXP$.

One can ask why it is not possible to use the same protocol when the prover is quantum in both rounds. The reason is that if we do so, the answers given by the prover in the second round may depend on the results of a measurement of a quantum state that is entangled to the state supplied to the verifier in the first round. This forms a sophisticated version of the EPR paradox, in the spirit of [CHTW].

1.6 Preliminaries

We assume that the reader is familiar with the basic concepts and notations of quantum computation. For excellent surveys on the subject see [Aha, NC].

Let F be a field of size 2^a for some integer a (that will be a function of n and will be determined later on). Our basic quantum element will be a quantum register of a qubits, rather than a single qubit. Each such basic element represents a member of the Hilbert space $C^{|F|}$. We denote by $\{|e\rangle\}_{e \in F}$ the standard basis for that space.

The base for the logarithm in this paper is always 2. By $[m]$ we denote the set $\{1, ..., m\}$. We denote probabilities by Prob and expectations by Exp. We say that a multivariate polynomial is of total degree r if its total degree is **at most** r.

2 The Retrieval Protocol

In this section, we present[4] the results discussed in Subsection 1.2. We will encode 2^n (classical) bits $a_1, ..., a_{2^n}$ by a quantum state $|\Psi\rangle = |\Psi(a_1, ..., a_{2^n})\rangle$ of size $O(n)$ qubits. We will show how to retrieve the value of any of the original bits by a (polynomial-size) Arthur-Merlin protocol. Our protocol is in fact more general: We will be able to encode and retrieve $a_1, ..., a_{2^n}$ that can get $n^{O(1)}$ different values, rather than bits (i.e., each a_i can be a block of $O(\log n)$ bits).

As noted in Subsection 1.2, the protocol has a completely classical interpretation. We concentrate here on the quantum case as this is the version that will be needed for the rest of the paper.

2.1 Quantum Low Degree Extension

W.l.o.g., assume that $n > 4$ is an even power of 2 (otherwise, we just increase n to at most $4n$, by padding with zeros). Denote by F a field of size $2^a \doteq n^c$, where c is a large enough constant integer that will be determined later on (for the content of this section, $c = 2$ is enough). Let $H \subset F$ be any (efficiently enumerable) subset of size $\sqrt{n}$ (e.g., the lexicographically first elements in some representation of the field F). Denote $d = 2n/\log n$, and assume for simplicity of the presentation that d is integer. Note that $|H^d| = 2^n$. Denote by $\pi : H^d \to [2^n]$ any (efficiently computable) one-to-one function (e.g., the lexicographic order of H^d).

Let $a_1, ..., a_{2^n} \in F$. We can view $(a_1, ..., a_{2^n})$ as a function from H^d to F. More precisely, define $A : H^d \to F$ by $A(z) = a_{\pi(z)}$. A basic fact is that there exists a unique extension of A into a function $\tilde{A} : F^d \to F$, such that $\tilde{A}$ is a multivariate polynomial (in d variables) of degree at most $|H| - 1$ in each variable. The function $\tilde{A}$ is called, the *low degree extension* of $a_1, ..., a_{2^n}$. Note that the total degree of $\tilde{A}$ is lower than $2n^{1.5}/\log n < n^{1.5}$.

We define the *quantum low degree extension* of $a_1, ..., a_{2^n}$, by

$$|\Psi(a_1, ..., a_{2^n})\rangle = |F|^{-d/2} \cdot \sum_{z_1,..,z_d \in F} |z_1\rangle|z_2\rangle \cdots |z_d\rangle|\tilde{A}(z_1, ..., z_d)\rangle.$$

Note that $|\Psi(a_1, ..., a_{2^n})\rangle$ is a quantum state of $(d+1)c \log n = 2cn + c \log n = O(n)$ qubits.

[4] We present a simplified version of the argument, given by Adam Smith.

2.2 The Protocol

Assume now that Alice got the state $|\Psi\rangle = |\Psi(a_1, ..., a_{2^n})\rangle$ and she wants to retrieve the value of a_i for some $i \in [2^n]$, or more generally, the value of $\tilde{A}(w)$ for some $w \in F^d$. This can be done by the following interactive protocol with the infinitely powerful prover Merlin.

Alice Measures all qubits of $|\Psi\rangle$, and gets as a result a random $z \in F^d$ and the value $\tilde{A}(z)$. If $z \neq w$, Alice computes the line (i.e., affine subspace of dimension 1 in F^d) that contains both w and z. Formally, this line is the set

$$\ell = \{w + (z - w) \cdot t\}_{t \in F} \subset F^d$$

(where all operations are in the vector space F^d). Alice sends ℓ to Merlin[5]. Merlin is required to respond with the value of $\tilde{A}$ on all the points in ℓ. Denote by $g(t)$ the value given by Merlin for the point $w + (z - w) \cdot t$.

Roughly speaking, Alice will reject (i.e., conclude the value Err) if $g : F \to F$ is not a low degree polynomial in the variable t, or if $g(1)$ disagrees with the value $\tilde{A}(z)$ (which is the only value of $\tilde{A}$ that Alice knows).

Formally, denote by $\tilde{A}|_\ell : F \to F$ the restriction of $\tilde{A}$ to the line ℓ (parameterized by t). That is, $\tilde{A}|_\ell(t) = \tilde{A}(w + (z - w) \cdot t)$. Recall that the total degree of $\tilde{A}$ is $< n^{1.5}$. Hence, $\tilde{A}|_\ell$ is a polynomial (in the one variable t) of degree $< n^{1.5}$. If g is not a polynomial of degree $< n^{1.5}$ then Alice rejects automatically. Otherwise, Alice checks whether or not $g(1) = \tilde{A}(z)$. If $g(1) \neq \tilde{A}(z)$ Alice rejects (note that $g(1)$ is the value given by Merlin for the point z). Otherwise, Alice concludes the value $g(0)$ (i.e., the value given by Merlin for the point w).

2.3 Analysis of the Protocol

The analysis of the retrieval protocol of Subsection 2.2 is extremely simple.

Denote by r a strategy of Merlin in the protocol. Formally, r is just the set of all answers given by Merlin for all possible pairs (w, ℓ). W.l.o.g., we can assume that the strategy r is deterministic. Denote by $V_{R1}(|\Psi\rangle, w, r)$ the value concluded by Alice when applying the protocol on a quantum state $|\Psi\rangle$ and a point $w \in F^d$, when Merlin is applying the strategy r. Note that $V_{R1}(|\Psi\rangle, w, r)$ is a random variable. Recall that for $a_1, ..., a_{2^n} \in F$, we denote by $\tilde{A} : F^d \to F$ the low degree extension of $a_1, ..., a_{2^n}$ and by $|\Psi(a_1, ..., a_{2^n})\rangle$ the quantum low degree extension of $a_1, ..., a_{2^n}$, as defined in Subsection 2.1.

The completeness and soundness properties of the protocol are given by the following lemma.

Lemma 2.1 *For every $a_1, ..., a_{2^n} \in F$ and every $w \in F^d$,*

[5]This can be done by sending w and one additional point (say, the lexicographically first point) in ℓ. Note that Merlin doesn't know z. (We don't care if Merlin does know w, but note that we could also send ℓ by just sending two different points (say, the two lexicographically first points) in it).

1. $\exists r$, *s.t.* $V_{R1}(|\Psi(a_1,...,a_{2^n})\rangle, w, r) = \tilde{A}(w)$ *with probability 1.*

2. $\forall r$, $V_{R1}(|\Psi(a_1,...,a_{2^n})\rangle, w, r) \in \{\tilde{A}(w), Err\}$ *with probability* $\geq 1 - 1/n^{c-1.5}$.

Proof:
Obviously, if Merlin's answer on line ℓ is the polynomial $g = \tilde{A}|_\ell$ then Alice concludes the correct value $\tilde{A}(w)$ with probability 1. So, the first part is obvious.

For the second part, note that if Merlin's answer on a line ℓ is a polynomial g of degree less than $n^{1.5}$ then either g is the same polynomial as $\tilde{A}|_\ell$ or the two polynomials agree on less than $n^{1.5}$ points. In the first case, Alice concludes the correct value $\tilde{A}(w)$. In the second case, Alice will reject for every value $z \in \ell \setminus \{w\}$ on which the two polynomials disagree. (Recall that Merlin doesn't know z and only knows the description of the line ℓ). Thus, Alice rejects on a fraction of at least $1 - n^{1.5}/|F|$ of the points in $\ell \setminus \{w\}$. Summing over all lines, with probability of at least $1 - n^{1.5}/|F| = 1 - 1/n^{c-1.5}$ Alice will either conclude the correct value or reject. □

2.4 Retrieving More Values

Suppose now that we want Alice to be able to retrieve the values of $a_{i_1}, ..., a_{i_k}$, for $k > 1$. An obvious way to do that is by encoding $a_1, ..., a_{2^n}$ by the tensor product of $|\Psi(a_1, ..., a_{2^n})\rangle$ with itself k times, that is, by the state $|\Psi\rangle \otimes \cdots \otimes |\Psi\rangle$, where $|\Psi\rangle = |\Psi(a_1, ..., a_{2^n})\rangle$ is the quantum low degree extension of $a_1, ..., a_{2^n}$, (as before). Alice can now retrieve one value from each copy of the state $|\Psi\rangle$. Moreover, this can be done in parallel in one round[6].

In this paper, we will not use this method. We will need, however, a method to retrieve more than one value from only **one** copy of $|\Psi(a_1, ..., a_{2^n})\rangle$. This can be done by a generalization of the retrieval protocol of Subsection 2.2. For simplicity of the presentation, we will present here the retrieval of only two values. The same protocol generalizes to an arbitrary k. The complexity of the retrieval protocol, however, is exponential in k.

Assume that Alice got the state $|\Psi\rangle = |\Psi(a_1, ..., a_{2^n})\rangle$ and she wants to retrieve the values of $a_i, a_{i'}$ for some (different) $i, i' \in [2^n]$, or more generally, the values of $\tilde{A}(w), \tilde{A}(w')$ for some (different) $w, w' \in F^d$. This can be done by the following interactive protocol.

Alice Measures all qubits of $|\Psi\rangle$, and gets as a result a random $z \in F^d$ and the value $\tilde{A}(z)$. Alice computes the plane (i.e., affine subspace of dimension 2 in F^d) that contains all three points w, w', z. Formally, this plane[7] is the set

$$p = \{w + (z - w) \cdot t_1 + (w' - w) \cdot t_2\}_{t_1, t_2 \in F} \subset F^d.$$

[6]It is not hard to show that in this setting applying the protocol in parallel is practically equivalent to applying it sequentially. Issues of parallel repetition, such as the the ones in [Raz], are not a problem here.

[7]Note that if z, w, w' happen to be on the same line then p is a line rather than a plan. Nevertheless, we can proceed in the exact same way.

Alice sends p to Merlin, who is required to respond with the value of $\tilde{A}$ on all the points in p. Denote by $g(t_1, t_2)$ the value given by Merlin for the point $w + (z - w) \cdot t_1 + (w' - w) \cdot t_2$.

If g is not a polynomial of total degree $< n^{1.5}$ then Alice rejects automatically. If $g(1,0) \neq \tilde{A}(z)$ Alice rejects as well. Otherwise, Alice concludes the values $(g(0,0), g(0,1))$ (i.e., the values given by Merlin for the points w, w').

Denote by r a strategy of Merlin in the protocol. Denote by $V_{R2}(|\Psi\rangle, (w, w'), r)$ the values concluded by Alice when applying the protocol on a quantum state $|\Psi\rangle$ and points $w, w' \in F^d$, when Merlin is applying the strategy r. The completeness and soundness properties of the protocol are given by the following lemma.

Lemma 2.2 *For every $a_1, ..., a_{2^n} \in F$ and every $w, w' \in F^d$,*

1. *$\exists r$, s.t. $V_{R2}(|\Psi(a_1, ..., a_{2^n})\rangle, (w, w'), r) = (\tilde{A}(w), \tilde{A}(w'))$ with probability 1.*
2. *$\forall r$, $V_{R2}(|\Psi(a_1, ..., a_{2^n})\rangle, (w, w'), r) \in \{(\tilde{A}(w), \tilde{A}(w')), Err\}$ with probability $\geq 1 - 1/n^{c-1.5}$.*

Proof:
Same as the proof of Lemma 2.1 □

3 Interactive Proofs with Quantum Advice

In this section, we present the results discussed in Subsection 1.3. Quantum interactive proof systems were first introduced by Watrous (see [Wat] for the formal definition). In these proof systems, the verifier can be modelled by a polynomial-size quantum circuit. Quantum interactive proof systems with polynomial-size quantum advice are defined in the same way, except that the verifier is modelled by a polynomial-size quantum circuit with a polynomial-size quantum advice. That is, the verifier is a polynomial-size quantum circuit, with working space initiated with an arbitrary quantum state $|\Psi\rangle$. (The state $|\Psi\rangle$ is considered to be part of the description of the circuit and it cannot depend on the inputs to the circuit).

The class $QIP/qpoly$ is defined to be the class of all languages that have polynomial-size quantum interactive proofs with a polynomial-size quantum advice. We show that the class $QIP/qpoly$ contains all languages. For any language L, the membership $x \in L$ can be proved by a polynomial-size quantum interactive proof, with a polynomial-size quantum advice. Moreover, the interactive proofs that we construct for the membership $x \in L$ are of only one round, and all messages communicated are classical[8].

Theorem 3.1 *$QIP/qpoly$ contains all languages.*

[8]Note that formally in the standard definition of quantum interactive proofs the parties can only communicate between them quantum messages. Nevertheless, since a quantum states can encode classical messages, the model is equivalent to a model where the parties can communicate both quantum and classical messages. In the interactive proofs constructed here, the parties communicate only classical messages.

Proof:
Let L be any language. For a string i of length n bits, define $a_i = 1$ iff $i \in L$. We will use as a quantum advice for the verifier the quantum low degree extension $|\Psi(a_0, ..., a_{2^n-1})\rangle$ (see Subsection 2.1). The proof now follows by the retrieval protocol of Subsection 2.2. Given x of length n bits, the verifier uses the retrieval protocol to retrieve the value of a_x (by an interactive protocol with the prover). The verifier accepts iff the value concluded by the protocol is 1 (and rejects if the value concluded is 0 or Err). The completeness and soundness properties follow immediately by Lemma 2.1. More precisely, if $x \in L$ there is a strategy for the prover that causes the verifier to accept (with probability 1), and on the other hand, if $x \notin L$ then no strategy for the prover can cause the verifier to accept with non-negligible probability. □

4 Quantum Low Degree Testing

In the settings of Subsection 1.2 and Subsection 1.3, a verifier could assume that the quantum state given to him is a correct quantum low degree extension of some $a_1, ..., a_{2^n}$ (as defined in Subsection 2.1). In the setting of quantum proofs, and quantum versions of the PCP theorem, a verifier cannot assume anything about a quantum state given to him. A key step towards proving the results discussed in Subsection 1.1 is a *quantum low degree test*, developed in this section. Roughly speaking, a quantum low degree test intends to check whether a quantum state is close to a representation of a polynomial of small total degree.

4.1 Classical Low Degree Tests

Roughly speaking, a (classical) low degree test intends to check whether a multivariate function is close to a polynomial of small total degree. Low degree tests and their applications have been studied in numerous of works and have been central in the study of interactive proofs and probabilistic checkable proofs (see for example [BFL, FGLSS, AS1, ALMSS]). In this paper, we will need to use the "low-error" analysis of [RS2, AS2] for (versions of) the test presented in [RS1].

Let F be a field and let d be some integer. Let L be the set of all lines in F^d (i.e., the set of all affine subspaces of dimension 1). For every $\ell \in L$, let $g_\ell : \ell \to F$ be a polynomial[9] of degree r. Denote, $G = \{g_\ell\}_{\ell \in L}$.

For $f, f' : F^d \to F$ and for $G = \{g_\ell\}_{\ell \in L}$ as above, denote

$$\mathrm{Agr}[f, f'] = \mathrm{Prob}_{z \in F^d}[f(z) = f'(z)],$$

$$\mathrm{Agr}[f, g_\ell] = \mathrm{Prob}_{z \in \ell}[f(z) = g_\ell(z)],$$

[9]We assume here that ℓ is presented as $\ell = \{u + (v-u) \cdot t\}_{t \in F}$ for some $u, v \in \ell$, and hence we can think of g_ℓ as a polynomial in the variable t. Note that the degree of g_ℓ does not depend on the choice of u, v.

$$\mathrm{Agr}[f, G] = \mathrm{Exp}_{\ell \in L} \mathrm{Agr}[f, g_\ell],$$

where all probabilities and expectations are with respect to the uniform distribution.

The Rubinfeld-Sudan test [RS1] suggests that if $\mathrm{Agr}[f, G]$ is large then f is close to a polynomial of total degree r. The following lemma, that menages to work with quite small values of $\mathrm{Agr}[f, G]$, was proved in [AS2]. A similar lemma for planes, rather than lines, was proved in [RS2] (see also [DFKRS]). Here, we can use any of these tests. We note that the lemmas proved in [RS2, AS2] are in fact stronger in several ways. We present them here in a simpler form that will suffice for us.

Lemma 4.1 *(Arora-Sudan) Let $f : F^d \to F$ be any function, and let $G = \{g_\ell\}_{\ell \in L}$ be such that every $g_\ell : \ell \to F$ is a polynomial of degree r. Assume that*

$$\mathrm{Agr}[f, G] > \frac{c \cdot r}{|F|^\epsilon}$$

where c is a (large enough) universal constant and $\epsilon > 0$ is a (small enough) universal constant. Then, there exists $h : F^d \to F$ of total degree r, such that,

$$\mathrm{Agr}[h, f], \mathrm{Agr}[h, G] \geq (\mathrm{Agr}[f, G])^2/32.$$

Note that the lemma shows that if $\mathrm{Agr}[f, G]$ is large then **both** f and G are close to a polynomial h of total degree r. Interestingly, it will be easier for us to use the claim about G.

In this paper, we will need a slightly more general version of Lemma 4.1, where we allow the polynomials $g_\ell : \ell \to F$ to take multiple values. More generally, we allow each g_ℓ to be a random variable, distributed over polynomials of degree r. We update the above notations as follows.

For $f : F^d \to F$ and for $G = \{g_\ell\}_{\ell \in L}$ as above, denote

$$\mathrm{Agr}[f, g_\ell] = \mathrm{Exp}_{g_\ell} \mathrm{Prob}_{z \in \ell}[f(z) = g_\ell(z)],$$

$$\mathrm{Agr}[f, G] = \mathrm{Exp}_{\ell \in L} \mathrm{Agr}[f, g_\ell].$$

It is a folklore meta-theorem that all known low degree tests work as well when assignments can take multiple values. As before, if $\mathrm{Agr}[f, G]$ is large then both f and G are close to a polynomial h of total degree r.

Lemma 4.2 *Let $f : F^d \to F$ be any function, and let $G = \{g_\ell\}_{\ell \in L}$ be such that every $g_\ell : \ell \to F$ is a random variable, distributed over polynomials of degree r. Assume that*

$$\mathrm{Agr}[f, G] > \frac{c \cdot r}{|F|^\epsilon}$$

where c is a (large enough) universal constant and $\epsilon > 0$ is a (small enough) universal constant. Then, there exists $h : F^d \to F$ of total degree r, such that,

$$\mathrm{Agr}[h, f], \mathrm{Agr}[h, G] \geq (\mathrm{Agr}[f, G])^2/32.$$

The lemma follows by a reduction to Lemma 4.1, using well known methods (see for example [AS2, RS2, DFKRS]).

4.2 The Quantum Test

Let F be a field of size 2^a for some integer a, and let d be some integer. Recall that our basic quantum element is a quantum register of a qubits, rather than a single qubit. Each such basic element represents a member of the Hilbert space $C^{|F|}$. Denote by $\mathcal{H}_{d+1}$ and $\mathcal{H}_2$ the following Hilbert spaces

$$\mathcal{H}_{d+1} = C^{|F|^{d+1}},$$
$$\mathcal{H}_2 = C^{|F|^2}.$$

Let L be the set of all lines in F^d (as before). Our quantum low degree test intends to check whether a quantum state $|\Phi\rangle \in \mathcal{H}_{d+1}$ is close to a state of the form

$$|F|^{-d/2} \cdot \sum_{z_1,..,z_d \in F} |z_1\rangle|z_2\rangle \cdots |z_d\rangle|f(z_1,...,z_d)\rangle,$$

where $f : F^d \to F$ is some polynomial of total degree r. In addition to the state $|\Phi\rangle$, the test has access to a set of **(classical)** polynomials $G = \{g_\ell\}_{\ell \in L}$, where as before, for every $\ell \in L$ the polynomial $g_\ell : \ell \to F$ is of degree r (see footnote in Subsection 4.1). Each g_ℓ is supposed to be (in a correct proof) the restriction of f to the line ℓ. In our test, the verifier reads only one of the polynomials g_ℓ.

4.2.1 Step I

The verifier chooses a random regular (i.e., one to one) linear function $E : F^d \to F^d$. The function E defines a permutation U_E over the standard basis for $\mathcal{H}_{d+1}$, as follows: For every $z = (z_1, ..., z_d) \in F^d$ and every $y \in F$,

$$|z_1\rangle|z_2\rangle \cdots |z_d\rangle|y\rangle \longmapsto |E(z)_1\rangle|E(z)_2\rangle \cdots |E(z)_d\rangle|y\rangle.$$

Since U_E is a permutation over a basis for $\mathcal{H}_{d+1}$, it extends to a unitary transformation

$$U_E : \mathcal{H}_{d+1} \to \mathcal{H}_{d+1}.$$

The verifier computes the quantum state $U_E|\Phi\rangle$ and measures the first $d-1$ registers of that state (i.e., $|E(z)_1\rangle \cdots |E(z)_{d-1}\rangle$). Denote by $b_1, ..., b_{d-1} \in F$ the results of the measurement. The state $U_E|\Phi\rangle$ collapses into a state $|\Phi'\rangle \in \mathcal{H}_2$ (in the last two registers).

Note that the set of solutions for the set of linear equations

$$E(z)_1 = b_1, \ldots, E(z)_{d-1} = b_{d-1}$$

is a line $\ell \in L$. The line ℓ can be presented as $\ell = \{u + (v-u) \cdot t\}_{t \in F}$, where $u \in F^d$ is the unique solution for the set of linear equations $E(u) = (b_1, \ldots, b_{d-1}, 0)$, and $v \in F^d$ is the unique solution for the set of linear equations $E(v) = (b_1, \ldots, b_{d-1}, 1)$.

If the original state $|\Phi\rangle$ is indeed of the form

$$|F|^{-d/2} \cdot \sum_{z_1,..,z_d \in F} |z_1\rangle|z_2\rangle \cdots |z_d\rangle|f(z_1,...,z_d)\rangle,$$

then the collapsed state $|\Phi'\rangle \in \mathcal{H}_2$ will be

$$|\Phi'\rangle = |F|^{-1/2} \cdot \sum_{t \in F} |t\rangle|f_\ell(t)\rangle,$$

where $f_\ell : F \to F$ is the restriction of f to the line ℓ (parameterized by t), i.e., $f_\ell(t) = f(u + (v-u) \cdot t)$.

4.2.2 Step II

The verifier reads the polynomial g_ℓ. We can think of this polynomial as a polynomial $g_\ell : F \to F$, where the line ℓ is parameterized by the same t as above (i.e., the line ℓ is presented as $\ell = \{u + (v-u) \cdot t\}_{t \in F}$).

Denote by $|e_1\rangle \in \mathcal{H}_2$ the quantum state

$$|e_1\rangle = |F|^{-1/2} \cdot \sum_{t \in F} |t\rangle|g_\ell(t)\rangle.$$

The verifier wants to compare the states $|\Phi'\rangle$ and $|e_1\rangle$. This is done as follows. The verifier extends $|e_1\rangle$ into any orthonormal basis $\{|e_1\rangle, \ldots, |e_{|F|^2}\rangle\}$ for the space $\mathcal{H}_2$, and measures the state $|\Phi'\rangle$ according to this basis. The verifier accepts if the result of the measurement is 1 and rejects in any other case.

Note that if indeed

$$|\Phi'\rangle = |F|^{-1/2} \cdot \sum_{t \in F} |t\rangle|f_\ell(t)\rangle,$$

and $f_\ell = g_\ell$, then $|\Phi'\rangle = |e_1\rangle$ and the verifier accepts with probability 1. In general, the verifier accepts with probability

$$|\langle e_1|\Phi'\rangle|^2.$$

4.3 Complexity of the Verifier

The complexity of the verifier in the procedure of Subsection 4.2 is polynomial in $|F|$ and d. To see this, we need to check that both steps can be done in that complexity.

In the first step, the verifier needs to compute the quantum transformation

$$|z_1\rangle|z_2\rangle \cdots |z_d\rangle \longmapsto |E(z)_1\rangle|E(z)_2\rangle \cdots |E(z)_d\rangle.$$

It is enough to show that the classical transformation

$$(z_1, z_2, \ldots, z_d) \longmapsto (E(z)_1, E(z)_2, \ldots, E(z)_d)$$

has a reversible classical circuit of size $poly(|F|, d)$. This follows immediately by the fact that any such transformation can be expressed as a product of $poly(d)$ reversible operations on only two variables each. One way to do that is by the inverse of the Gauss elimination procedure, that shows how to diagonalize any $d \times d$ matrix E by a sequence of $poly(d)$ operations that work on only two rows each. Note that every operation that works on only two variables can be trivially translated into a quantum circuit of size $poly(|F|)$, as the dimension of the relevant Hilbert space, $\mathcal{H}_2$, is $|F|^2$.

In the second step, the verifier needs to measure the state $|\Phi'\rangle$ according to the basis $\{|e_1\rangle, \ldots, |e_{|F|^2}\rangle\}$. Note however that since the space $\mathcal{H}_2$ is of dimension $|F|^2$, this can trivially be done by a quantum circuit of size $poly(|F|)$.

4.4 Analysis of the Test

For a quantum state $|\Phi\rangle \in \mathcal{H}_{d+1}$ and for a set of polynomials $G = \{g_\ell\}_{\ell \in L}$ (where for every $\ell \in L$ the polynomial $g_\ell : \ell \to F$ is of degree r), denote by $V_{QLDT}(|\Phi\rangle, G)$ the probability that the quantum low degree test procedure of Subsection 4.2 accepts.

The completeness of the test is given by the following lemma. The lemma shows that if $|\Phi\rangle$ is indeed a correct representation of a polynomial $f : F^d \to F$ of total degree r, and each g_ℓ is the restriction of f to the line ℓ, then the test accepts with probability 1.

Lemma 4.3 *Assume that*

$$|\Phi\rangle = |F|^{-d/2} \cdot \sum_{z_1,..,z_d \in F} |z_1\rangle|z_2\rangle \cdots |z_d\rangle|f(z_1, ..., z_d)\rangle,$$

for some polynomial $f : F^d \to F$ of total degree r. Assume that $G = \{g_\ell\}_{\ell \in L}$, where every $g_\ell : \ell \to F$ is the restriction of f to the line ℓ. Then,

$$V_{QLDT}(|\Phi\rangle, G) = 1.$$

Proof:
The proof is straightforward. As mentioned above, after Step I we get the collapsed state $|\Phi'\rangle = |e_1\rangle$. Hence, the result of the measurement in Step II will always be 1. □

The soundness of the test is harder to prove and is given by the following lemma. The lemma shows that if $V_{QLDT}(|\Phi\rangle, G)$ is large then G is close to a polynomial h of total degree r. Recall that the original motivation of the test was to prove that $|\Phi\rangle$ is close to a representation of a polynomial h of low total degree. Nevertheless, it will be enough for us to have this property for G rather than $|\Phi\rangle$. For simplicity of the presentation, we state and prove the lemma only for G.

Lemma 4.4 *Let $G = \{g_\ell\}_{\ell \in L}$ be such that every $g_\ell : \ell \to F$ is a polynomial of degree r. Assume that for some quantum state $|\Phi\rangle \in \mathcal{H}_{d+1}$,*

$$V_{QLDT}(|\Phi\rangle, G) > \frac{c \cdot r}{|F|^\epsilon}$$

where c is a (large enough) universal constant and $\epsilon > 0$ is a (small enough) universal constant. Then, there exists $h : F^d \to F$ of total degree r, such that,

$$\mathrm{Agr}[h, G] \geq [V_{QLDT}(|\Phi\rangle, G)]^4/50.$$

The proof of the Lemma is given in Subsection 4.5.

As in the case of Lemma 4.1, we will need a slightly more general version of Lemma 4.4, where we allow the polynomials $g_\ell : \ell \to F$ to take multiple values. More generally, we allow each g_ℓ to be a random variable, distributed over polynomials of degree r. When reading g_ℓ, the verifier gets an evaluation of g_ℓ, that is, each polynomial of degree r is obtained with the probability that g_ℓ gets that value.

We denote by $V_{QLDT}(|\Phi\rangle, G)$ the probability that the quantum low degree test procedure accepts on a quantum state $|\Phi\rangle$ and on a set $G = \{g_\ell\}_{\ell \in L}$ as above.

Lemma 4.5 *Let $G = \{g_\ell\}_{\ell \in L}$ be such that every $g_\ell : \ell \to F$ is a random variable, distributed over polynomials of degree r. Assume that for some quantum state $|\Phi\rangle \in \mathcal{H}_{d+1}$,*

$$V_{QLDT}(|\Phi\rangle, G) > \frac{c \cdot r}{|F|^\epsilon}$$

where c is a (large enough) universal constant and $\epsilon > 0$ is a (small enough) universal constant. Then, there exists $h : F^d \to F$ of total degree r, such that,

$$\mathrm{Agr}[h, G] \geq [V_{QLDT}(|\Phi\rangle, G)]^4/50.$$

The proof of Lemma 4.5 is the same as the one of Lemma 4.4, using Lemma 4.2 rather than Lemma 4.1.

4.5 Proof of Lemma 4.4

4.5.1 Notations

First note that w.l.o.g. we can assume that $|\Phi\rangle$ is a **pure state**. For $z = (z_1, ..., z_d) \in F^d$ and $y \in F$, denote by $\phi_{z,y}$ the coefficient of $|z_1\rangle|z_2\rangle \cdots |z_d\rangle|y\rangle$ in $|\Phi\rangle$. That is,

$$|\Phi\rangle = \sum_{z \in F^d, y \in F} \phi_{z,y} |z\rangle|y\rangle.$$

For every $z \in F^d$, denote

$$\phi_z = \sqrt{\sum_{y \in F} |\phi_{z,y}|^2}$$

For every line $\ell \in L$, denote

$$\phi_\ell = \sqrt{\sum_{z \in \ell} |\phi_z|^2}$$

Denote

$$N = \frac{|F|^d - 1}{|F| - 1}$$

and note that N is the number of directions of lines ℓ in F^d. For every $z \in F^d$, denote by $L(z) \subset L$ the set of lines ℓ that contain z. Note that every $L(z)$ is a set of N lines.

We denote the total acceptance probability $V_{QLDT}(|\Phi\rangle, G)$ by γ, i.e.,

$$\gamma = V_{QLDT}(|\Phi\rangle, G).$$

4.5.2 The Acceptance Probability

Assume w.l.o.g. that every $\phi_{z,y}$ (and hence also every ϕ_z and every ϕ_ℓ) is none-zero. Otherwise, we just change the state $|\Phi\rangle$ to an extremely close state that satisfies that property. (This is done for the simplicity of the presentation, in order to avoid divisions by 0).

In Step I of the test, the linear function E determines the direction of the line ℓ that we obtain in that step. Since E is chosen with the uniform distribution, each direction is chosen with probability $1/N$. After the direction is chosen, each line ℓ in that direction is obtained with probability ϕ_ℓ^2. Altogether, every line ℓ is obtained with probability ϕ_ℓ^2/N.

If a line ℓ was obtained, the state $|\Phi\rangle$ collapses to the state

$$|\Phi'_\ell\rangle = \phi_\ell^{-1} \cdot \sum_{t \in F} \sum_{y \in F} \phi_{z(t),y} |t\rangle |y\rangle,$$

where $z(t) = u + (v - u) \cdot t$, and u, v are the ones defined in Subsection 4.2 (i.e., u, v are such that the line ℓ is presented as $\ell = \{u + (v - u) \cdot t\}_{t \in F}$, as described in Subsection 4.2).

Since $|e_1\rangle = |F|^{-1/2} \cdot \sum_{t \in F} |t\rangle |g_\ell(t)\rangle$, the acceptance probability $|\langle e_1 | \Phi'_\ell \rangle|^2$ is

$$|\langle e_1 | \Phi'_\ell \rangle|^2 = \left| \phi_\ell^{-1} \cdot |F|^{-1/2} \cdot \sum_{t \in F} \phi_{z(t),g_\ell(t)} \right|^2 = \left| \phi_\ell^{-1} \cdot |F|^{-1/2} \cdot \sum_{z \in \ell} \phi_{z,g_\ell(z)} \right|^2$$

where (for simplicity) we think of g_ℓ as a polynomial $g_\ell : F \to F$ when we write $g_\ell(t)$, and as a polynomial $g_\ell : \ell \to F$ when we write $g_\ell(z)$.

The total acceptance probability γ can now be expressed as

$$\gamma = \sum_{\ell \in L} \left(\frac{\phi_\ell^2}{N} \right) \cdot \left| \phi_\ell^{-1} \cdot |F|^{-1/2} \cdot \sum_{z \in \ell} \phi_{z,g_\ell(z)} \right|^2 = (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \left| \sum_{z \in \ell} \phi_{z,g_\ell(z)} \right|^2$$

We can see from this expression that w.l.o.g. we can assume that all the coefficients $\phi_{z,y}$ of the state $|\Phi\rangle$ are real and positive. (Otherwise, we change each $\phi_{z,y}$ to $|\phi_{z,y}|$ and we can only increase the total acceptance probability). Hence,

$$\gamma = (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \left(\sum_{z \in \ell} \phi_{z,g_\ell(z)} \right) \cdot \left(\sum_{z' \in \ell} \phi_{z',g_\ell(z')} \right)$$

$$= (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \sum_{z,z' \in \ell} \phi_{z,g_\ell(z)} \cdot \phi_{z',g_\ell(z')}$$

Since every $\phi_{z',g_\ell(z')}$ is at most $\phi_{z'}$, we can bound

$$\gamma \leq (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \sum_{z,z' \in \ell} \phi_{z,g_\ell(z)} \cdot \phi_{z'}$$

We will write the last expression as a sum of two expressions, according to whether or not $z = z'$. The first expression is the sum of all terms where $z = z'$. That expression is

$$(|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \sum_{z \in \ell} \phi_{z,g_\ell(z)} \cdot \phi_z \leq (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \sum_{z \in \ell} \phi_z^2$$

$$= (|F| \cdot N)^{-1} \cdot \sum_{z \in F^d} \sum_{\ell \in L(z)} \phi_z^2 = |F|^{-1} \cdot \sum_{z \in F^d} \phi_z^2 = |F|^{-1}$$

Hence,

$$\gamma - |F|^{-1} \leq (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \sum_{z \neq z' \in \ell} \phi_{z,g_\ell(z)} \cdot \phi_{z'}$$

$$= (|F| \cdot N)^{-1} \cdot \sum_{\ell \in L} \sum_{z \neq z' \in \ell} (\phi_z \cdot \phi_{z'}) \cdot (\phi_{z,g_\ell(z)}/\phi_z)$$

and hence by the Cauchy-Schwartz inequality,

$$\gamma - |F|^{-1} \leq (|F| \cdot N)^{-1} \cdot \sqrt{\sum_{\ell \in L} \sum_{z \neq z' \in \ell} (\phi_z \cdot \phi_{z'})^2} \cdot \sqrt{\sum_{\ell \in L} \sum_{z \neq z' \in \ell} (\phi_{z,g_\ell(z)}/\phi_z)^2}$$

In this formula, we can substitute

$$\sum_{\ell \in L} \sum_{z \neq z' \in \ell} (\phi_z \cdot \phi_{z'})^2 = \sum_{z \in F^d} \sum_{z' \neq z \in F^d} (\phi_z \cdot \phi_{z'})^2 \leq \sum_{z \in F^d} \sum_{z' \in F^d} (\phi_z \cdot \phi_{z'})^2 = \left(\sum_{z \in F^d} \phi_z^2 \right)^2 = 1$$

and

$$\sum_{\ell \in L} \sum_{z \neq z' \in \ell} (\phi_{z,g_\ell(z)}/\phi_z)^2 = \sum_{z \in F^d} \sum_{\ell \in L(z)} \sum_{z' \neq z \in \ell} (\phi_{z,g_\ell(z)}/\phi_z)^2 = (|F| - 1) \cdot \sum_{z \in F^d} \sum_{\ell \in L(z)} (\phi_{z,g_\ell(z)}/\phi_z)^2$$

and we get

$$\left(\gamma - |F|^{-1}\right)^2 \leq (|F| \cdot N)^{-2} \cdot (|F| - 1) \cdot \sum_{z \in F^d} \sum_{\ell \in L(z)} (\phi_{z,g_\ell(z)}/\phi_z)^2$$

$$\leq |F|^{-d} \cdot N^{-1} \cdot \sum_{z \in F^d} \sum_{\ell \in L(z)} (\phi_{z,g_\ell(z)}/\phi_z)^2. \tag{1}$$

4.5.3 Using the Classical Test

We will now define a probabilistic function $f : F^d \to F$. Formally, for every $z \in F^d$ we define $f(z)$ as a random variable in F. Alternatively, we can think of f as a distribution over functions from F^d to F. For every $z \in F^d$ and $y \in F$, we define

$$\text{Prob}[f(z) = y] = \left(\phi_{z,y}/\phi_z\right)^2.$$

Note that

$$\sum_{y \in F} (\phi_{z,y}/\phi_z)^2 = 1.$$

We extend the definition of $\text{Agr}[f, G]$ from Subsection 4.1 to probabilistic functions f. Formally, we define

$$\text{Agr}[f, g_\ell] = \text{Exp}_f \text{Prob}_{z \in \ell}[f(z) = g_\ell(z)],$$
$$\text{Agr}[f, G] = \text{Exp}_{\ell \in L} \text{Agr}[f, g_\ell].$$

Thus,

$$\text{Agr}[f, G] = \text{Exp}_{\ell \in L} \text{Exp}_f \text{Prob}_{z \in \ell}[f(z) = g_\ell(z)]$$

$$= \text{Exp}_{\ell \in L} \text{Exp}_{z \in \ell} \text{Prob}_f[f(z) = g_\ell(z)]$$

$$= \text{Exp}_{z \in F^d} \text{Exp}_{\ell \in L(z)} \text{Prob}_f[f(z) = g_\ell(z)]$$

$$= |F|^{-d} \cdot N^{-1} \cdot \sum_{z \in F^d} \sum_{\ell \in L(z)} \text{Prob}_f[f(z) = g_\ell(z)]$$

$$= |F|^{-d} \cdot N^{-1} \cdot \sum_{z \in F^d} \sum_{\ell \in L(z)} (\phi_{z, g_\ell(z)}/\phi_z)^2$$

$$\geq \left(\gamma - |F|^{-1}\right)^2$$

(by inequality 1).

Since we can think of f as a distribution over deterministic functions f', the agreement $\text{Agr}[f, G]$ is a convex combination of $\text{Agr}[f', G]$ for deterministic functions f'. Hence, there exists a (deterministic) function $f' : F^d \to F$ with

$$\text{Agr}[f', G] \geq \left(\gamma - |F|^{-1}\right)^2.$$

Hence, by Lemma 4.1, there exists $h : F^d \to F$ of total degree r, such that,

$$\text{Agr}[h, G] \geq \gamma^4/50$$

(under the assumption that the universal constant c is large enough and the universal constant ϵ is small enough). □

5 Quantum Information and the PCP Theorem

In this section, we present the results discussed in Subsection 1.1. Roughly speaking, we show that the membership $x \in SAT$ can be proved by a logarithmic-size quantum state, together with a polynomial-size classical proof of blocks of poly-logarithmic length, such that after measuring the quantum state the verifier only needs to read one of the blocks of the classical proof.

5.1 PCP with Short Quantum Witness

In all that comes below, a verifier is a polynomial-time machine that can process both quantum states and classical strings.

We define an (s_1, s_2)-verifier to be as follows: The verifier gets three inputs: $(x, |\Phi\rangle, p)$. The first input, x, is a classical string of length n bits. (We think of x as the input whose membership in a language L is to be verified). The second input, $|\Phi\rangle$, is a quantum state of length s_1 qubits. The third input, p, is a classical array of $poly(n)$ blocks of length s_2 bits each. (We think of $(|\Phi\rangle, p)$ as a possible proof for the membership $x \in L$). The verifier is allowed to query at most **one** of the blocks of the third input p.

We define the class QPCP$[s_1, s_2, \epsilon]$ as follows: A language L is in QPCP$[s_1, s_2, \epsilon]$ if there exists an $(O(s_1), O(s_2))$-verifier V, such that the following completeness and soundness properties are satisfied:

1. For any $x \in L$, there exist $|\Phi\rangle$ and p, such that

$$\text{Prob}[V(x, |\Phi\rangle, p) = accept] = 1.$$

2. For any $x \notin L$, and any $|\Phi\rangle$ and p,

$$\text{Prob}[V(x, |\Phi\rangle, p) = accept] \leq \epsilon.$$

The definition extends to promise problems, where we only consider inputs x that satisfy a certain *promise*.

Theorem 5.1 $SAT \in \text{QPCP}[log(n), polylog(n), o(1)]$

5.2 Proof of Theorem 5.1

5.2.1 Classical PCP

For the proof of Theorem 5.1, it is clearly enough to prove $L \in \text{QPCP}[log(n), polylog(n), o(1)]$ for any other NP-complete language or promise problem L. We will work with the following promise problem that we call $GAP(s, q, \epsilon)$:

An instance of the problem is $x = (\varphi_1, ..., \varphi_k)$, where $\varphi_1, ..., \varphi_k$ are predicates over a set of variables $\{Y_1, ..., Y_m\}$. Every variable Y_i can take 2^s different values (i.e., we can think of every Y_i as a block of s bits). Every predicate φ_i depends on at most q of the variables $Y_1, ..., Y_m$. The promise is that: either, there is an assignment to $Y_1, ..., Y_m$ that satisfies **all** predicates, or, any assignment to $Y_1, ..., Y_m$ satisfies at most ϵ fraction of the predicates. The goal is to accept iff the first possibility is correct (under the assumption that the promise is satisfied).

Different versions of the PCP theorem prove the NP-completeness of $GAP(s, q, \epsilon)$ for a large range of values of the parameters s, q, ϵ. Here, we will be interested in the following parameters: We require s to be at most $\log \log n$ (where n is the length of x). We require q to be constant. We would like ϵ to be as small as possible, preferably sub-constant. (The error of our verifier in the proof of Theorem 5.1 will be polynomially related to ϵ. Thus, a small constant ϵ is ok if we only want to achieve a small constant error).

The NP-completeness of $GAP(s, q, \epsilon)$, for some $s \leq \log \log n$, for some constant q, and for some $\epsilon \leq (\log n)^{-\Omega(1)}$ is known [RS2, AS2, DFKRS]. Moreover, if we only tried to achieve a small **constant** probability of error, we could have used many other versions of the PCP theorem. For example, we could have used the results in [Raz] and work with $q = 2$ and an arbitrarily small constant ϵ.

In all that comes below, we fix (s, q, ϵ) such that $GAP(s, q, \epsilon)$ is known to be NP-complete and such that: q is constant, s is at most $\log \log n$, and ϵ is sub-constant. We will show that the problem is in QPCP$[log(n), polylog(n), o(1)]$. The best probability of error that we are able to achieve is $(\log n)^{-\Omega(1)}$.

We will construct a verifier V such that on an instance $x = (\varphi_1, ..., \varphi_k)$ of $GAP(s, q, \epsilon)$, the following properties are satisfied:

1. If there exists an assignment to $Y_1, ..., Y_m$ that satisfies all predicates $\varphi_1, ..., \varphi_k$, then there exist $|\Phi\rangle$ and p, such that

$$\text{Prob}[V(x, |\Phi\rangle, p) = accept] = 1.$$

2. If any assignment to $Y_1, ..., Y_m$ satisfies at most ϵ fraction of the predicates $\varphi_1, ..., \varphi_k$, then for any $|\Phi\rangle$ and p,

$$\text{Prob}[V(x, |\Phi\rangle, p) = accept] \leq \epsilon'$$

(where $\epsilon' = o(1)$).

Every assignment to $Y_1, ..., Y_m$ that satisfies all predicates will translate into $(|\Phi\rangle, p)$, such that, $\text{Prob}[V(x, |\Phi\rangle, p) = accept] = 1$. We think of $(|\Phi\rangle, p)$ as a proof for the satisfiability of $\varphi_1, ..., \varphi_k$. We think of the verifier as a procedure that verifies that proof. We will first describe how to translate an assignment to $Y_1, ..., Y_m$ (that satisfies all predicates) into a (correct) proof $(|\Phi\rangle, p)$, and then describe the verification procedure.

5.2.2 Preliminaries

Let $x = (\varphi_1, ..., \varphi_k)$ be an instance of length n of $GAP(s, q, \epsilon)$. For simplicity of the notations, extend the set of variables $\{Y_1, ..., Y_m\}$ to $\{Y_1, ..., Y_n\}$, (by adding dummy variables). Denote by $t_1, ..., t_k$ the sets of variables that the predicates $\varphi_1, ..., \varphi_k$ depend on (respectively). Recall that each t_j is a set of size at most q. W.l.o.g., we can assume that $t_1, ..., t_k$ are all different. (This is only done for simplifying the notations). W.l.o.g., we can assume that every predicate in $\varphi_1, ..., \varphi_k$ has at least one satisfying assignment. (Otherwise, it is clear that there is no assignment that satisfies all predicates, and since s is at most $\log \log n$ the verifier can check that easily).

W.l.o.g., assume that $n = 2^{\hat{n}}$, such that $\hat{n} > 4$ is an even power of 2 (otherwise, we just increase n to at most n^4). The variable $\hat{n}$ will play the role of n in Subsection 2.1. As in Subsection 2.1, denote by F a field of size $2^a \doteq \hat{n}^c$, where c is a large enough constant integer that will be determined later on. As in Subsection 2.1, denote $d = 2\hat{n}/\log \hat{n}$ and assume for simplicity that d is integer. Assume that $d > q + 1$. Note that $|F^d|$ is polynomial in n.

Let $H \subset F$ be a subset of size $\sqrt{\hat{n}}$, as in Subsection 2.1, and let $\pi : H^d \to [2^{\hat{n}}]$ be a one-to-one function, as in Subsection 2.1. We will use here the inverse function $\pi^{-1} : [n] \to F^d$. This function maps the variables in $\{Y_1, ..., Y_n\}$ to F^d. Intuitively, we think of each variable Y_i as placed on the point $\pi^{-1}(i) \in F^d$. For every t_j, define $\tau_j = \pi^{-1}(t_j) \subset F^d$. W.l.o.g., assume that for every τ_j, the dimension of the smallest affine subspace of F^d that contains τ_j is exactly $q - 1$ (otherwise, we add arbitrary points to τ_j).

Let L be the set of all lines in F^d, as in Section 4. For any $\ell \in L$ and any τ_j, denote by $S(\ell, \tau_j)$ the smallest affine subspace of F^d that contains both ℓ and τ_j. This subspace will usually be of dimension $q + 1$, and will always be of dimension at most $q + 1$.

For an assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$, let $\tilde{A} : F^d \to F$ be the low degree extension of $a_1, ..., a_n$, as defined in Subsection 2.1. Recall that the total degree of $\tilde{A}$ is less than $\hat{n}^{1.5}$. For any affine subspace $S \subset F^d$, denote by $\tilde{A}|_S$ the restriction of $\tilde{A}$ to S. We think of $\tilde{A}|_S$ also as a function from $F^{d'}$ to F, where d' is the dimension of S, and where formally we assume that some (linear) parameterization of the affine space S is implicit.

5.2.3 The Correct Proof

An assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$, that satisfies all predicates, translates into $(|\Phi\rangle, p)$ that causes the verifier to accept with probability 1. We refer to that $(|\Phi\rangle, p)$ as the *correct proof*.

The state $|\Phi\rangle$ of the correct proof will be the quantum low degree extension of $a_1, ..., a_n$, i.e., $|\Psi(a_1, ..., a_n)\rangle$, as defined in Subsection 2.1. Note that this is a state of $O(\hat{n}) = O(\log n)$ qubits. The array p will have one block for every pair (τ_j, S), such that $S \subset F^d$ is a $q + 1$ dimensional affine subspace that contains τ_j. In the correct proof, this block will contain the restriction $\tilde{A}|_S$ (as in [DFKRS]). Note that the number of blocks is bounded by $k \cdot |F|^{2d}$ which is polynomial in n, and the size of each block (i.e., the number of bits it takes to describe each

$\tilde{A}|_S$) is bounded by $a \cdot (\hat{n}^{1.5})^{q+1}$, which is poly-logarithmic in n.

5.2.4 The Verification Procedure

Denote by $p(\tau_j, S)$ the content of the block indexed by (τ_j, S) of p. Recall that in the correct proof $p(\tau_j, S) = \tilde{A}|_S$ and recall that $\tilde{A}|_S$ is a polynomial of total degree at most $\hat{n}^{1.5}$. Note that from $\tilde{A}|_S$ one can induce the restriction of $\tilde{A}$ to τ_j, as well as the restriction of $\tilde{A}$ to any line ℓ contained in S. Note that the restriction of $\tilde{A}$ to τ_j gives the values of the assignment $a_1, ..., a_n$ to all the variables in t_j, by $a_i = \tilde{A}(\pi^{-1}(i))$ (for every $i \in t_j$). Given these values, one can check wether or not the predicate φ_j is satisfied. Note that in the correct proof φ_j must be satisfied.

In general, the verifier expects $p(\tau_j, S)$ to be a function from S to F (and we can think of this function also as a function from F^{q+1} to F). Whenever the verifier reads a block $p(\tau_j, S)$, the verifier can check that $p(\tau_j, S)$ is indeed a polynomial of total degree at most $\hat{n}^{1.5}$, and that the values induced from $p(\tau_j, S)$ to the set τ_j satisfy the predicate φ_j. Since the verifier rejects automatically whenever a block that doesn't pass these tests is read, we can assume w.l.o.g. that all the blocks pass these tests. That is, every $p(\tau_j, S)$ is a polynomial of total degree $\hat{n}^{1.5}$, and the values induced from it to the set τ_j satisfy the predicate φ_j.

The verification procedure goes as follows. The verifier performs Step I of the quantum low degree test, as described in Subsection 4.2, and proceeds to Step II. At the beginning of Step II, the verifier needs to read g_ℓ (for the line ℓ obtained in Step I). This is done as follows: The verifier chooses a random $j \in [k]$. If $S(\ell, \tau_j)$ is of dimension $q+1$, define $S = S(\ell, \tau_j)$. Otherwise, define S to be a random $q+1$ dimensional affine subspace that contains $S(\ell, \tau_j)$. The verifier reads $p(\tau_j, S)$ (and performs the above mentioned tests on $p(\tau_j, S)$), and define g_ℓ to be the polynomial induced from $p(\tau_j, S)$ (i.e., the restriction of $p(\tau_j, S)$ to ℓ). The verifier continues with Step II of the quantum low degree test, as described in Subsection 4.2.

5.2.5 Complexity of the Verifier

We only need the complexity of the verifier to be polynomial in n. Note, however, that if the verifier has random access to x, p, then all tasks in the verification procedure can be performed in time poly-logarithmic in n. While this is not important for the proof, it is essential for scaling up the proof to $NEXP$, and may be important for possible future applications.

5.2.6 Analysis of the Test

Denote by $V(x, |\Phi\rangle, p)$ the output of the verifier on inputs $(x, |\Phi\rangle, p)$. Note that $V(x, |\Phi\rangle, p)$ is a random variable.

The completeness of the test is straightforward. If there exists an assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$ that satisfies all predicates $\varphi_1, ..., \varphi_k$ then the correct proof $(|\Phi\rangle, p)$, as described in

Subsection 5.2.3, satisfies

$$\text{Prob}[V(x, |\Phi\rangle, p) = \textit{accept}] = 1.$$

The soundness of the test is given by the following lemma. The lemma shows that if for some $(|\Phi\rangle, p)$, $\text{Prob}[V(x, |\Phi\rangle, p) = \textit{accept}]$ is large then there is an assignment to $Y_1, ..., Y_n$ that satisfies many of the predicates $\varphi_1, ..., \varphi_k$.

Lemma 5.2 *Assume that for some $(|\Phi\rangle, p)$ and some γ,*

$$\text{Prob}[V(x, |\Phi\rangle, p) = \textit{accept}] \geq \gamma > \frac{c' \cdot \hat{n}^{1.5}}{|F|^{\epsilon}}$$

where c' is a (large enough) universal constant and $\epsilon > 0$ is a (small enough) universal constant (as in Lemma 4.5). Then, there exists an assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$ that satisfies at least $\gamma^4/100$ fraction of the predicates $\varphi_1, ..., \varphi_k$.

5.2.7 Proof of Lemma 5.2

Recall that w.l.o.g. we can assume that every $p(\tau_j, S) : S \to F$ is a polynomial of total degree at most $\hat{n}^{1.5}$, and that the values induced from it to the set τ_j satisfy the predicate φ_j. For every line $\ell \in L$ that is contained in S, denote by $p(\tau_j, S)|_\ell : \ell \to F$ the restriction of $p(\tau_j, S)$ to ℓ.

For every $\ell \in L$, and every $j \in [k]$, define $g_{\ell,j} : \ell \to F$ as follows: If $S(\ell, \tau_j)$ is of dimension $q + 1$, define $S = S(\ell, \tau_j)$. Otherwise, define S to be a random $q + 1$ dimensional affine subspace that contains $S(\ell, \tau_j)$. Define $g_{\ell,j}$ to be the restriction of $p(\tau_j, S)$ to ℓ, that is, $g_{\ell,j} = p(\tau_j, S)|_\ell$. Note that $g_{\ell,j}$ is the same as the polynomial g_ℓ defined by the verification procedure (see Subsection 5.2.4) in case that j is the random index in $[k]$ that was picked by the procedure. For every $\ell \in L$, define $g_\ell : \ell \to F$ as follows: Choose a random $j \in [k]$ and fix $g_\ell = g_{\ell,j}$. Note that g_ℓ is the same as the polynomial g_ℓ defined by the verification procedure (see Subsection 5.2.4).

Denote, $G = \{g_\ell\}_{\ell \in L}$. For every $j \in [k]$, denote $G_j = \{g_{\ell,j}\}_{\ell \in L}$. Recall the definition of $\text{Agr}[h, G]$ in Subsection 4.1. By Lemma 4.5 (under the assumption that c' is large enough and ϵ is small enough), there exists $h : F^d \to F$ of total degree $\hat{n}^{1.5}$, such that,

$$\text{Agr}[h, G] \geq \gamma^4/50.$$

By the definitions of G, G_j,

$$\text{Exp}_j \text{Agr}[h, G_j] = \text{Agr}[h, G].$$

Hence, for at least $\gamma^4/100$ fraction of the indices $j \in [k]$,

$$\text{Agr}[h, G_j] \geq \gamma^4/100. \tag{2}$$

Define the assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$ to be the assignment induced from h, that is, $\forall i,\ a_i = h(\pi^{-1}(i))$. We will show that for every $j \in [k]$ that satisfies inequality 2, the

assignment $a_1, ..., a_n$ satisfies the predicate φ_j. Hence, the assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$ satisfies at least $\gamma^4/100$ fraction of the predicates $\varphi_1, ..., \varphi_k$, and the lemma is proved.

Claim 5.3 *For every $j \in [k]$ that satisfies inequality 2, the assignment $a_1, ..., a_n$ to $Y_1, ..., Y_n$ satisfies the predicate φ_j.*

Proof:
Fix $j \in [k]$ that satisfies inequality 2. Denote by $L' \subset L$ the set of all lines ℓ, such that $S(\ell, \tau_j)$ is of dimension exactly $q+1$. Recall that for every $\ell \in L'$,

$$g_{\ell,j} = p(\tau_j, S(\ell, \tau_j))|_\ell$$

Since the dimension of the smallest affine subspace of F^d that contains τ_j is $q-1 < d-2$, most lines in L are also in L'. More precisely, the ratio $|L'|/|L|$ is larger than $1-|F|^{-1}$. Hence, by inequality 2,

$$\mathrm{Exp}_{\ell \in L'} \mathrm{Agr}[h, g_{\ell,j}] > \mathrm{Exp}_{\ell \in L} \mathrm{Agr}[h, g_{\ell,j}] - |F|^{-1}$$

$$= \mathrm{Agr}[h, G_j] - |F|^{-1} \geq \gamma^4/100 - |F|^{-1}. \quad (3)$$

Denote by $\mathcal{S}$ the set of all $q+1$ dimensional affine subspaces $S \subset F^d$ that contain τ_j. For every $S \in \mathcal{S}$, denote by L_S the set of all lines $\ell \in L$ that are contained in S, and denote by $L'_S \subset L_S$ the set of all lines $\ell \in L$, such that $S(\ell, \tau_j) = S$. In other words, $L'_S = L_S \cap L'$. Note that $\{L'_S\}_{S \in \mathcal{S}}$ is a partition of L'. Hence, by inequality 3,

$$\mathrm{Exp}_{S \in \mathcal{S}} \mathrm{Exp}_{\ell \in L'_S} \mathrm{Agr}[h, g_{\ell,j}] = \mathrm{Exp}_{\ell \in L'} \mathrm{Agr}[h, g_{\ell,j}]$$

$$\geq \gamma^4/100 - |F|^{-1}. \quad (4)$$

Note that for every $\ell \in L'_S$, we have $g_{\ell,j} = p(\tau_j, S(\ell, \tau_j))|_\ell = p(\tau_j, S)|_\ell$. Hence,

$$\mathrm{Exp}_{S \in \mathcal{S}} \mathrm{Exp}_{\ell \in L'_S} \mathrm{Agr}[h, g_{\ell,j}] =$$

$$\mathrm{Exp}_{S \in \mathcal{S}} \mathrm{Exp}_{\ell \in L'_S} \mathrm{Prob}_{z \in \ell}[h(z) = p(\tau_j, S)|_\ell(z)]. \quad (5)$$

For every $S \in \mathcal{S}$, the dimension of the smallest affine subspace of S that contains τ_j is $q-1 = (q+1)-2$. Hence, most lines in L_S are also in L'_S. More precisely, the ratio $|L'_S|/|L_S|$ is larger than $1 - 2|F|^{-1}$. Therefore, for every $S \in \mathcal{S}$,

$$\mathrm{Exp}_{\ell \in L_S} \mathrm{Prob}_{z \in \ell}[h(z) = p(\tau_j, S)|_\ell(z)] \geq$$

$$\mathrm{Exp}_{\ell \in L'_S} \mathrm{Prob}_{z \in \ell}[h(z) = p(\tau_j, S)|_\ell(z)] - 2|F|^{-1}$$

Hence, by inequality 4 and equality 5,

$$\mathrm{Exp}_{S\in\mathcal{S}}\mathrm{Exp}_{\ell\in L_S}\mathrm{Prob}_{z\in\ell}[h(z)=p(\tau_j,S)|_\ell(z)]\geq$$

$$\mathrm{Exp}_{S\in\mathcal{S}}\mathrm{Exp}_{\ell\in L'_S}\mathrm{Prob}_{z\in\ell}[h(z)=p(\tau_j,S)|_\ell(z)]-2|F|^{-1}$$

$$=\mathrm{Exp}_{S\in\mathcal{S}}\mathrm{Exp}_{\ell\in L'_S}\mathrm{Agr}[h,g_{\ell,j}]-2|F|^{-1}$$

$$\geq\gamma^4/100-3|F|^{-1}.$$

Hence, by a uniformity argument,

$$\mathrm{Exp}_{S\in\mathcal{S}}\mathrm{Prob}_{z\in S}[h(z)=p(\tau_j,S)(z)]=$$

$$\mathrm{Exp}_{S\in\mathcal{S}}\mathrm{Exp}_{\ell\in L_S}\mathrm{Prob}_{z\in\ell}[h(z)=p(\tau_j,S)|_\ell(z)]$$

$$\geq\gamma^4/100-3|F|^{-1}.$$

Hence, there exists (at least one) $S\in\mathcal{S}$, such that,

$$\mathrm{Prob}_{z\in S}[h(z)=p(\tau_j,S)(z)]\geq\gamma^4/100-3|F|^{-1}.$$

Recall that $h:F^d\rightarrow F$ and $p(\tau_j,S):S\rightarrow F$ are both polynomials of total degree at most $\hat{n}^{1.5}$. Thus, by Schwartz-Zippel's lemma, if they agree on a fraction larger than $\hat{n}^{1.5}/|F|$ of the points $z\in S$ they must agree on every point $z\in S$. Thus, under the assumption that the constant c (that determines the size of the field F) is large enough, h and $p(\tau_j,S)$ agree on every point $z\in S$. (Note that we have the freedom to fix c as large as we want).

Since we assumed that the values induced from $p(\tau_j,S)$ to the set τ_j satisfy the predicate φ_j, we conclude that the values induced from h to the set τ_j satisfy the predicate φ_j. That is, the assignment $a_1,...,a_n$ to $Y_1,...,Y_n$ satisfies the predicate φ_j.

This ends the proof of Claim 5.3. □

Since inequality 2 holds for at least $\gamma^4/100$ fraction of the indices $j\in[k]$, the assignment $a_1,...,a_n$ to $Y_1,...,Y_n$ satisfies at least $\gamma^4/100$ fraction of the predicates $\varphi_1,...,\varphi_k$.

This ends the proof of Lemma 5.2. □

5.2.8 Completing the Proof of Theorem 5.1

We have constructed an $(O(logn), polylog(n))$-verifier V, such that on an instance $x = (\varphi_1, ..., \varphi_k)$ of $GAP(s, q, \epsilon)$ the following properties are satisfied:

1. If there exists an assignment to $Y_1, ..., Y_m$ that satisfies all predicates $\varphi_1, ..., \varphi_k$, then there exist $|\Phi\rangle$ and p, such that

   $$\text{Prob}[V(x, |\Phi\rangle, p) = accept] = 1.$$

   (See Subsection 5.2.6).

2. If any assignment to $Y_1, ..., Y_m$ satisfies at most ϵ fraction of the predicates $\varphi_1, ..., \varphi_k$, then for any $|\Phi\rangle$ and p,

   $$\text{Prob}[V(x, |\Phi\rangle, p) = accept] \leq o(1).$$

   (By Lemma 5.2).

Hence, $GAP(s, q, \epsilon) \in \text{QPCP}[log(n), polylog(n), o(1)]$, and since $GAP(s, q, \epsilon)$ is NP-complete we conclude that $NP \subset \text{QPCP}[log(n), polylog(n), o(1)]$. □

Acknowledgment

I am grateful to Adam Smith for simplifying the retrieval protocol of Subsection 2.2 (and for allowing me to include here the simplified version), and to Amir Shpilka and Yael Tauman Kalai for very helpful conversations.

References

[Aar] Scott Aaronson: Limitations of Quantum Advice and One-Way Communication. Theory of Computing 1(1): 1-28 (2005)

[Aha] Dorit Aharonov: Quantum Computation- A Review. Annual Review of Computational Physics, World Scientific, volume VI, ed. Dietrich Stauffer (1998)

[ALMSS] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, Mario Szegedy: Proof Verification and the Hardness of Approximation Problems. J. ACM 45(3): 501-555 (1998)

[ANTV] Andris Ambainis, Ashwin Nayak, Amnon Ta-Shma, Umesh V. Vazirani: Dense Quantum Coding and Quantum Finite Automata. J. ACM 49(4): 496-511 (2002)

[AS1] Sanjeev Arora, Shmuel Safra: Probabilistic Checking of Proofs: A New Characterization of NP. J. ACM 45(1): 70-122 (1998)

[AS2] Sanjeev Arora, Madhu Sudan: Improved Low-Degree Testing and its Applications. Combinatorica 23(3): 365-426 (2003)

[Bab] Laszlo Babai: Trading Group Theory for Randomness. STOC 1985: 421-429

[BFL] Laszlo Babai, Lance Fortnow, Carsten Lund: Non-Deterministic Exponential Time has Two-Prover Interactive Protocols. Computational Complexity 1: 3-40 (1991)

[BM] Laszlo Babai, Shlomo Moran: Arthur-Merlin Games: A Randomized Proof System, and a Hierarchy of Complexity Classes. J. Comput. Syst. Sci. 36(2): 254-276 (1988)

[CHTW] Richard Cleve, Peter Hoyer, Benjamin Toner, John Watrous: Consequences and Limits of Nonlocal Strategies. IEEE Conference on Computational Complexity 2004: 236-249

[DFKRS] Irit Dinur, Eldar Fischer, Guy Kindler, Ran Raz, Shmuel Safra: PCP Characterizations of NP: Towards a Polynomially-Small Error-Probability. STOC 1999: 29-40 (full version in http://www.wisdom.weizmann.ac.il/∼ranraz/publications)

[FGLSS] Uriel Feige, Shafi Goldwasser, Laszlo Lovasz, Shmuel Safra, Mario Szegedy: Interactive Proofs and the Hardness of Approximating Cliques. J. ACM 43(2): 268-292 (1996)

[GMR] Shafi Goldwasser, Silvio Micali, Charles Rackoff: The Knowledge Complexity of Interactive Proof Systems. SIAM J. Comput. 18(1): 186-208 (1989)

[Hol] Alexander S. Holevo: Some Estimates for the Amount of Information Transmittable by a Quantum Communications Channel. Problemy Peredaci Informacii 9(3): 3-11 (1973). English translation: Problems of Information Transmission 9(3): 177-183 (1973)

[KW] Alexei Kitaev, John Watrous: Parallelization, Amplification, and Exponential Time Simulation of Quantum Interactive Proof Systems. STOC 2000: 608-617

[LFKN] Carsten Lund, Lance Fortnow, Howard J. Karloff, Noam Nisan: Algebraic Methods for Interactive Proof Systems. J. ACM 39(4): 859-868 (1992)

[Nay] Ashwin Nayak: Optimal Lower Bounds for Quantum Automata and Random Access Codes. FOCS 1999: 369-377

[NC] Michael A Nielsen, Isaac L Chuang. Quantum Computation and Quantum Information. Cambridge University Press, 2000.

[NY] Harumichi Nishimura, Tomoyuki Yamakami: Polynomial Time Quantum Computation with Advice. Electronic Colloquium on Computational Complexity (ECCC)(059): (2003)

[Raz] Ran Raz: A Parallel Repetition Theorem. SIAM J. Comput. 27(3): 763-803 (1998)

[RS1] Ronitt Rubinfeld, Madhu Sudan: Robust Characterizations of Polynomials with Applications to Program Testing. SIAM J. Comput. 25(2): 252-271 (1996)

[RS2] Ran Raz, Shmuel Safra: A Sub-Constant Error-Probability Low-Degree Test and a Sub-Constant Error-Probability PCP Characterization of NP. STOC 1997: 475-484

[Sha] Adi Shamir: IP = PSPACE. J. ACM 39(4): 869-877 (1992)

[Wat] John Watrous: PSPACE has Constant-Round Quantum Interactive Proof Systems. Theor. Comput. Sci. 292(3): 575-588 (2003)