In the previous class, we have defined the class QMA, which can be seen as the quantum analogue of NP. We have also shown how to amplify its acceptance probabilities to be exponentially close to 0 and 1 (the proof can be found in [3]). The basic idea was to apply the original verifier several times and then take the majority vote among the answers. Since the (quantum) witness is potentially ruined after a measurement, we had to supply the new verifier with several copies of the original witness.

In this lecture we will show an alternative approach to amplifying QMA, due to Marriott and Watrous [4]. The advantage of this approach is that the witness remains the same. In other words, only one copy of the original witness is enough. In addition to being an interesting result by its own right, such 'witness-preserving amplification' has several applications, such as the one we'll see at the end of this lecture. Moreover, a similar approach was recently used in a result of Watrous on zero knowledge against quantum attacks [8].

Let us recall the definition of the class QMA. We start with the well-known classical class NP. We say that a family indexed by $\{0,1\}^*$ is *uniformly generated*, if there exists an algorithm that given a bit string x, generates the corresponding element of the family in time polynomial in the length of x. Then the class NP can be defined as follows.

DEFINITION 1 *The class* NP *consists of all languages $L \subseteq \{0,1\}^*$ for which there exists a uniformly generated family of classical, deterministic, poly-size circuits $\{V_x : x \in \{0,1\}^*\}$ and a polynomial m, such that:*

1. *For all $x \in L$ there exists an $m(|x|)$-bit witness w such that $V_x(w) = 1$;*
2. *For all $x \notin L$ and for all $m(|x|)$-bit witness w, $V_x(w) = 0$.*

By allowing randomization in the verification process, we obtain a class known as MA.

DEFINITION 2 *The class* MA *consists of all languages $L \subseteq \{0,1\}^*$ for which there exists a uniformly generated family of classical, randomized, poly-size circuits $\{V_x : x \in \{0,1\}^*\}$ and a polynomial m, such that:*

1. *For all $x \in L$, there exists an $m(|x|)$-bit witness w such that $\Pr(V_x(w) = 1) \geq 2/3$;*
2. *For all $x \notin L$ and for all $m(|x|)$-bit witness w, $\Pr(V_x(w) = 1) \leq 1/3$.*

Finally, by allowing a quantum witness and a quantum verifier, we obtain the class QMA (which is sometimes called BQNP). See Figure 1.

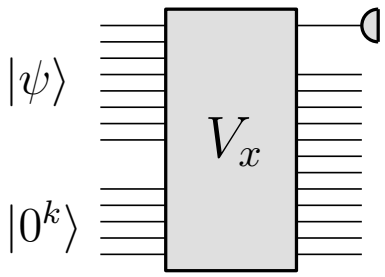

Figure 1: A QMA verifier

DEFINITION 3 *The class* QMA *consists of all languages $L \subseteq \{0,1\}^*$ for which there exists a uniformly generated family of quantum poly-size circuits $\{V_x : x \in \{0,1\}^*\}$ and polynomials m, k, where each V_x has $m(|x|)$ input qubits, $k(|x|)$ auxiliary qubits and its output is given by the first output qubit, such that:*

1. *For all $x \in L$ there exists an $m(|x|)$-qubit witness $|\psi\rangle$ such that $\Pr(V_x$ accepts $|\psi\rangle) \geq 2/3$;*

2. *For all $x \notin L$ and for all $m(|x|)$-qubit witness $|\psi\rangle$, $\Pr(V_x$ accepts $|\psi\rangle) \leq 1/3$.*

Two comments about this definition are in place. First, let us denote by $\mathsf{QMA}(a, b)$ the class of languages as in the definition above with $2/3$ and $1/3$ replaced by a and b respectively. Then, by the QMA amplification we saw in the last lecture, $\mathsf{QMA}(2/3, 1/3) = \mathsf{QMA}(1 - 2^{-\mathrm{poly}(|x|)}, 2^{-\mathrm{poly}(|x|)})$ and so we can say that QMA is robust with respects to these parameters. Second, we consider only *pure* states $|\psi\rangle$. Allowing mixed states would result in the same class, since the maximum acceptance probability of a verifier is always obtained by a pure state.

The choice of QMA as *the* quantum analogue of NP is not entirely obvious. Let us mention two other possible definitions:

- In the definition of QMA we take both the witness and the verifier to be quantum. If instead we take the verifier to be quantum but keep the witness classical, we obtain a class known as QCMA. Not much is known about this class beyond the obvious containments $\mathsf{NP} \subseteq \mathsf{MA} \subseteq \mathsf{QCMA} \subseteq \mathsf{QMA}$.

- An alternative way to define NP is as the class of languages L for which there exists a family of randomized, uniformly generated, circuits $\{V_x : x \in \{0,1\}^*\}$ such that $x \in L$ iff the probability that V_x outputs 1 is nonzero. Now, by considering quantum circuits instead of randomized ones, we obtain another quantum analogue of NP known as NQP. It turns out that NQP is equal to a known classical complexity class called $\mathsf{coC_{=}P}$, which is very powerful.

There are only a few problems known to be in QMA that are not known to be in NP. One notable example is the Group Non-Membership problem. Another important example are problems that are *complete* for QMA, such as the Local Hamiltonian problem. We will discuss these problems in a later class.

1 A Tale of Two Subspaces

As we all know, for any two lines in $\mathbb{R}^n$ that go through the origin (i.e., one-dimensional subspaces), one can define the *angle* between them. If we take a line and a plane, we can again define the angle between them in a natural way. But what happens if we take two planes? Here our three-dimensional intuition is no longer good enough. Indeed, in three-dimensions, two two-dimensional subspaces always intersect in a line, and orthogonal to that line we find the angle between the two subspaces. This is no longer true in higher dimensions: Starting from four dimensions, two two-dimensional subspaces generally have a trivial intersection, and instead of forming an angle, they form *two* angles!

In more generality, the question we consider in this section is how two subspaces interact. This question turns out to have a very elegant answer, as we shall soon see. This answer, which was first given in a remarkable paper of C. Jordan in 1875, was since rediscovered many times by mathematicians, statisticians, physicists, and computer scientists. In addition to being a crucial component in witness-preserving amplification of QMA, this question also plays an important role in many recent results in quantum computation, often in an implicit way (see, e.g., [4, 8, 5, 7, 1, 6]). This topic is also covered in Chapter VII of the book by Bhatia [2].

In our analysis below, we describe a subspace by the *projector* on the subspace. This is very convenient since a projector is independent of the choice of basis of the subspace. To recall, a projector is a Hermitian matrix satisfying $\Pi^2 = \Pi$, i.e., all its eigenvalues are either 0 or 1. The eigenspace corresponding to the eigenvalue 1 is the subspace on which Π projects.

So the question we are trying to answer is: given two projectors Π_1 and Π_2, what can we say about how they interact? We will show that one can find an orthogonal decomposition of space into one-dimensional

and two-dimensional subspaces such that both Π_1 and Π_2 are block-diagonal in this decomposition. An example of this is shown in Figure 2. There, two random 5-dimensional subspaces were chosen in 11-dimensional space. Based on these two subspaces, a decomposition of the 11-dimensional space into five two-dimensional subspaces and one one-dimensional subspace was computed. Then, a basis was chosen inside each two-dimensional subspace so that Π_1 is diagonal. The figure shows the two projectors in this basis. We remark that the upper matrix has a nicer form just because of the way we chose the basis inside each two-dimensional subspace.

$$\begin{pmatrix}
1. & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1. & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1. & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1. & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1. & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{pmatrix}$$

$$\begin{pmatrix}
0.957739 & 0.201185 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.201185 & 0.0422613 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0.782673 & 0.412427 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0.412427 & 0.217327 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0.31525 & -0.464615 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -0.464615 & 0.68475 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0.0819297 & -0.274258 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -0.274258 & 0.91807 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.00684302 & -0.082439 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.082439 & 0.993157 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{pmatrix}$$

Figure 2: Two randomly chosen rank-5 projectors in 11-dimensional space

LEMMA 1 *For any two projectors Π_1, Π_2 there exists an orthogonal decomposition of the Hilbert space into one-dimensional and two-dimensional subspaces that are invariant under both Π_1 and Π_2. Moreover, inside each two-dimensional subspace, Π_1 and Π_2 are rank-one projectors (in other words, inside each two-dimensional subspace there are two unit vectors $|v\rangle$ and $|w\rangle$ such that Π_1 projects on $|v\rangle$ and Π_2 projects on $|w\rangle$).*

PROOF: Consider the matrix $\Pi_1 + \Pi_2$. Clearly, this is a Hermitian matrix, and as such it has a complete set of eigenvectors. We show that this set of vectors can be partitioned into sets of size either one or two, and that each set spans an invariant subspace as in the lemma.

Let $|\varphi\rangle$ be an eigenvector of $\Pi_1 + \Pi_2$, normalized to be of unit length, and let λ be the corresponding eigenvalue. Then by definition,

$$\Pi_1|\varphi\rangle + \Pi_2|\varphi\rangle = \lambda|\varphi\rangle. \tag{1}$$

Assume first that $\Pi_1|\varphi\rangle$ is in $\mathrm{span}(|\varphi\rangle)$. By (1), we also have that $\Pi_2|\varphi\rangle$ is in $\mathrm{span}(|\varphi\rangle)$. Hence, we obtain that $\mathrm{span}(|\varphi\rangle)$ is a one-dimensional invariant subspace of both Π_1 and Π_2. We remark that since both Π_1 and Π_2 are projectors, we in fact have that $\Pi_1|\varphi\rangle$ is either 0 or $|\varphi\rangle$ and similarly $\Pi_2|\varphi\rangle$ is 0 or $|\varphi\rangle$.

So assume now that $\Pi_1|\varphi\rangle$ is not in $\mathrm{span}(|\varphi\rangle)$ and consider the two-dimensional space S spanned by $|\varphi\rangle$ and $\Pi_1|\varphi\rangle$. Then this space is invariant by Π_1 since for any α and β,

$$\Pi_1(\alpha|\varphi\rangle + \beta\Pi_1|\varphi\rangle) = (\alpha+\beta)\Pi_1|\varphi\rangle \in S.$$

It is also invariant by Π_2 since

$$\Pi_2|\varphi\rangle = \lambda|\varphi\rangle - \Pi_1|\varphi\rangle \in S$$

by (1) and

$$\Pi_2\Pi_1|\varphi\rangle = \Pi_2(\lambda|\varphi\rangle - \Pi_2|\varphi\rangle) = (\lambda - 1)\Pi_2|\varphi\rangle \in S$$

where the containment holds by the previous step. Being invariant by both Π_1 and Π_2, S is clearly also invariant by $\Pi_1 + \Pi_2$. Hence, the vector orthogonal to $|\varphi\rangle$ in S is another eigenvector of $\Pi_1 + \Pi_2$, and so S is spanned by two eigenvectors of $\Pi_1 + \Pi_2$, as promised. Finally, it follows easily from (1) that inside S, Π_1 and Π_2 are rank-one projectors. □

This lemma gives us a very good picture of the interaction between Π_1 and Π_2. Let us consider the two-dimensional subspaces $S_1, S_2, \ldots$ (the one-dimensional subspaces are in some sense trivial and won't matter for our applications). Inside each subspace S_i, Π_1 is the projector $|v_i\rangle\langle v_i|$ for some unit vector $|v_i\rangle$ and similarly Π_2 is the projector $|w_i\rangle\langle w_i|$ for some unit vector $|w_i\rangle$ (see Figure 3). The angles $\theta_i = \arccos(|\langle v_i|w_i\rangle|) \in [0, \pi/2]$ are essentially what is known as the *principal angles* between the two subspaces. For convenience, we set the phase of $|w_i\rangle$ so that $\langle v_i|w_i\rangle$ is a non-negative real. For instance, this implies that $\Pi_2|v_i\rangle = \cos\theta_i|w_i\rangle$ and $\Pi_1|w_i\rangle = \cos\theta_i|v_i\rangle$.

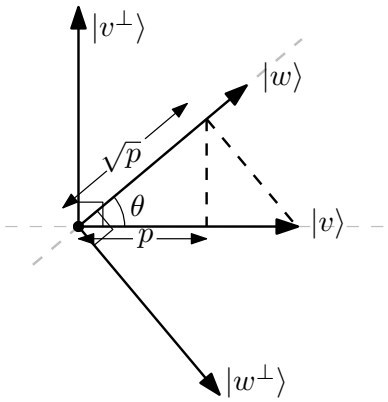

Figure 3: Π_1 and Π_2 inside a two-dimensional invariant space

Let us introduce some extra notation that will be useful in the next section. We define $p_i = \cos^2\theta_i = |\langle v_i|w_i\rangle|^2$. We also let $|v_i^\perp\rangle$ and $|w_i^\perp\rangle$ be unit vectors in S_i orthogonal to $|v_i\rangle$ and $|w_i\rangle$ respectively, with phases chosen so that we have

$$|w_i\rangle = \sqrt{p_i}|v_i\rangle + \sqrt{1-p_i}|{v_i}^\perp\rangle$$
$$|v_i\rangle = \sqrt{p_i}|w_i\rangle + \sqrt{1-p_i}|{w_i}^\perp\rangle.$$

This immediately implies that

$$|{v_i}^\perp\rangle = \sqrt{1-p_i}|w_i\rangle - \sqrt{p_i}|w_i^\perp\rangle$$
$$|{w_i}^\perp\rangle = \sqrt{1-p_i}|v_i\rangle - \sqrt{p_i}|v_i^\perp\rangle.$$

The lemma allows us to easily answer questions regarding matrices obtained from combinations of Π_1 and Π_2. For instance, what can we say about $\Pi_1\Pi_2\Pi_1$? It is clearly block-diagonal in the S_i decomposition, and inside each S_i, it is

$$|v_i\rangle\langle v_i||w_i\rangle\langle w_i||v_i\rangle\langle v_i| = |\langle v_i|w_i\rangle|^2 \cdot |v_i\rangle\langle v_i| = p_i \cdot |v_i\rangle\langle v_i|,$$

i.e., it is a projector on $|v_i\rangle$ multiplied by the squared cosine of the angle between $|v_i\rangle$ and $|w_i\rangle$.

2 Witness-preserving Amplification of QMA

We now prove the main result of this lecture.

THEOREM 2 *For any witness size $m = m(|x|) \in \mathrm{poly}$, any $a, b : \mathbb{N} \to [0,1]$ such that $a(n) - b(n) \geq 1/\mathrm{poly}(n)$ and any $r = \mathrm{poly}(n)$,*

$$\mathsf{QMA}_m(a,b) \subseteq \mathsf{QMA}_m(1-2^{-r}, 2^{-r})$$

where the subscript indicates the witness size.

PROOF: Let $L \in \mathsf{QMA}_m(a,b)$ and let $\{V_x : x \in \{0,1\}^*\}$ be the corresponding family of quantum verifiers. Without loss of generality, we can assume that for any witness the acceptance probability of V_x is strictly smaller than 1 and strictly bigger than 0. This follows by a simple modification to the circuit: for instance, start by tossing a fair three-sided dice; depending on the result, either accept, reject, or run the original verifier.

To prove the theorem, we shall construct a new verifier V'_x that expects the same witness as V_x and whose acceptance probabilities are exponentially close to 0 and 1. The verifier V'_x repeats the following two steps N times (see Figure 4):[1]

1. Apply V_x, measure if the output qubit is 1, and then apply $V_x^\dagger$;
2. Measure if all auxiliary qubits are 0.

It then writes down all $2N$ measurement results as $a_1, a_2, \ldots, a_{2N} \in \{\mathsf{Y}, \mathsf{N}\}$ and counts the number of times that $a_i = a_{i+1}$. For example, if it observes the results YYNYNNYNNN, the count is 4. If this number is at least $\frac{a+b}{2}(2N-1)$, it accepts; otherwise it rejects. We remark that there are several alternative ways to construct an amplified verifier V'_x, and an example of such a construction appears in the homework.

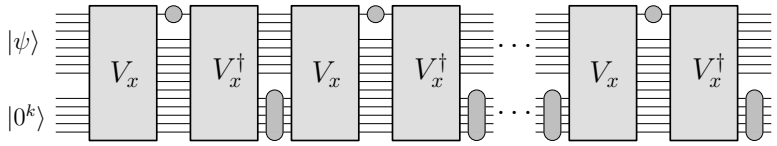

Figure 4: The new QMA verifier

We now analyze the verifier. Define Π_1 as the projector on the auxiliary qubits being zero, and Π_2 as $V_x^\dagger \Delta V_x$ where Δ is the projector on the output qubit being 1. We first notice that the maximum acceptance probability of the original verifier V_x is given by the largest eigenvalue of $\Pi_1\Pi_2\Pi_1$.[2] To see that, notice first that the subspace of all states with auxiliary qubits orthogonal to $|0^k\rangle$ is an eigenspace with eigenvalue 0 (because of Π_1). In the orthogonal space (spanned by states whose auxiliary qubits are $|0^k\rangle$),

$$\langle\psi|\Pi_1\Pi_2\Pi_1|\psi\rangle = \langle\psi|V_x^\dagger \Delta V_x|\psi\rangle = \|\Delta V_x|\psi\rangle\|^2$$

which is the acceptance probability of the witness $|\psi\rangle$. Hence the maximum acceptance probability is the maximum eigenvalue of $\Pi_1\Pi_2\Pi_1$ and the corresponding eigenvector is the witness that achieves it.

We now apply the analysis of the previous section to Π_1 and Π_2, and consider the resulting decomposition. Using the notation of the previous section, we have two-dimensional subspaces $S_1, S_2, \ldots$ with two unit vectors $|v_i\rangle$, $|w_i\rangle$ inside each subspace, and $p_i = |\langle v_i|w_i\rangle|^2$.

[1] This verifier can also be implemented without any intermediate measurements in a standard way.

[2] This observation is already used in the analysis of plain QMA amplification.

First we notice that Π_1 (or, to be precise, the subspace on which it projects) is spanned by $|v_1\rangle, |v_2\rangle, \ldots$. To see that, notice that in each one-dimensional subspace, Π_1 must be zero, since otherwise we obtain a vector in Π_1 (i.e., a legal witness) whose acceptance probability is either 0 or 1 (depending on whether Π_2 is 0 or 1 in that subspace), contradicting our assumption on V_x. So from now on we only need to consider the two-dimensional subspaces.

As we have seen in the previous section, the matrix $\Pi_1\Pi_2\Pi_1$ can be written as

$$\Pi_1\Pi_2\Pi_1 = \sum_{i=1,2,\ldots} p_i \cdot |v_i\rangle\langle v_i|.$$

In particular, we see that the maximum eigenvalue of $\Pi_1\Pi_2\Pi_1$, which is the maximum acceptance probability of V_x, is $\max_i p_i$.

We now go back to analyzing the new verifier V'_x. We can think of it as performing a sequence of projective measurements, alternating between the measurement $\{\Pi_2, I-\Pi_2\}$ and the measurement $\{\Pi_1, I-\Pi_1\}$. Let us analyze its behavior when given one of the vectors $|v_i\rangle$ as input (see Figure 5). In the first measurement, with probability $|\langle v_i|w_i\rangle|^2 = p_i$ we obtain Y and otherwise we obtain N. The resulting state is either $|w_i\rangle$ in case of Y or $|w_i^\perp\rangle$ in case of N. If the result was Y, then in the second measurement we obtain Y with probability $|\langle w_i|v_i\rangle|^2 = p_i$ and $|v_i^\perp\rangle$ otherwise. Similarly, if the result was $|w_i^\perp\rangle$, then in the second measurement we obtain Y with probability $|\langle w_i^\perp|v_i\rangle|^2 = 1-p_i$ and N otherwise. To summarize, we see that the probability of a transition from Y to Y or from N to N is exactly p_i.

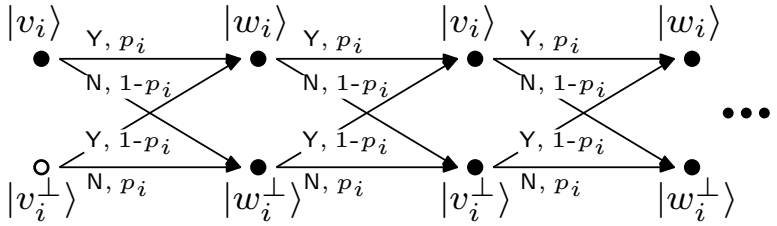

Figure 5: Transition probabilities in V'_x

The most important point to notice here is that we never go out of the two-dimensional space S_i. This is the key idea that makes the construction work. In a sense, even though we perform measurements, the original state $|v_i\rangle$ is not lost. For instance, whenever we measure Π_1, we're back to the original state $|v_i\rangle$.

From the above we can easily prove the correctness of V'_x in the case that $x \in L$. In this case, the largest eigenvalue of $\Pi_1\Pi_2\Pi_1$ is at least a. Hence there exists an i such that $p_i \geq a$. If we give $|v_i\rangle$ as a witness to V'_x, then by the above analysis, we see that V'_x essentially performs $2N-1$ independent coin flips with bias p_i. By a standard Chernoff bound, the probability that the fraction of heads is at least $\frac{a+b}{2}$ can be made greater than $1-2^{-r}$ by choosing N to be a large enough polynomial.

It remains to handle the case that $x \notin L$. Here we know that the largest eigenvalue of $\Pi_1\Pi_2\Pi_1$ is at most b and hence for all i, $p_i \leq b$. Our goal is to show that for any possible witness $|\psi\rangle \in \Pi_1$, V'_x accepts with probability at most 2^{-r}. In case $|\psi\rangle$ happens to be one of the eigenvectors $|v_i\rangle$, the analysis is essentially the same as in the previous case: each measurement is an independent coin toss with bias at most b. Hence by choosing N to be a large enough polynomial, we can make sure that for all $|v_i\rangle$, the probability that V'_x accepts is at most 2^{-r}.

For the general case, one can simply write an arbitrary witness $|\psi\rangle$ as a linear combination of the $|v_i\rangle$'s and perform an analysis similar to the one above on such combinations. It can be shown that at each step the probability of a transition N $\rightarrow$ N or Y $\rightarrow$ Y is at most b, no matter what the current state is. A more elegant way to argue this is the following. Consider the quantum operation that measures the name i of the

block S_i and traces i out. This operation can be described by the super-operator

$$\rho \to \sum_i \Pi_{S_i} \rho \Pi_{S_i}.$$

Since all the measurements performed by V'_x are block-diagonal in the S_i decomposition, this super-operator commutes with all of them. Hence, we can apply it before we start V'_x without changing the behavior of the verifier. But now the input to V'_x is a mixture of states $|v_i\rangle$, and we already know that V'_x accepts any such witness with probability at most 2^{-r}. □

3 One Application of Witness-preserving Amplification

In this section we show that witnesses of logarithmic size do not help. Let's start with the classical case, where the goal is to show that $\mathsf{MA}_{\log} = \mathsf{BPP}$. The idea is to simply guess a random witness. Since there is only a polynomial number of them, we have a good chance of hitting the correct one. To make sure we don't have too many false positives, we first amplify the probabilities.

THEOREM 3 *For any logarithmic* $m = m(|x|)$, $\mathsf{MA}_m(\frac{2}{3}, \frac{1}{3}) = \mathsf{BPP}$.

PROOF: The containment $\mathsf{BPP} \subseteq \mathsf{MA}_m(\frac{2}{3}, \frac{1}{3})$ follows from the definitions. For the other direction, let L be a language in $\mathsf{MA}_m(\frac{2}{3}, \frac{1}{3})$. By (classical) amplification, we have that $L \in \mathsf{MA}_m(\frac{3}{4}, \frac{1}{4}2^{-m})$. Let $\{V_x\}$ be the corresponding family of verifiers. Consider the BPP machine that on input x applies V_x to a witness chosen uniformly from all 2^m possible witnesses, and accepts if and only if V_x accepts.

If $x \in L$ then there is a witness w that V_x accepts with probability at least $\frac{3}{4}$. Thus, our BPP machine accepts with probability at least $\frac{3}{4}2^{-m}$. If $x \notin L$, then for any witness, V_x accepts with probability at most $\frac{1}{4}2^{-m}$. Hence, the probability that our BPP machine accepts is also at most $\frac{1}{4}2^{-m}$. The difference between $\frac{1}{4}2^{-m}$ and $\frac{3}{4}2^{-m}$ is inverse polynomial, and hence this is indeed a BPP machine. □

Notice that in the above proof it is crucial that the amplification process does not increase the size of the witness (or at least not too much). This will also be crucial in the proof below, where we show that $\mathsf{QMA}_{\log} = \mathsf{BQP}$. Unlike the classical case, it is not immediately clear how to choose a random witness. First of all, there is an infinite number of them. So instead we can try to take a finite set of witnesses that well-approximates all possible witnesses. This is indeed possible; however, since the witness is on a logarithmic number of qubits, it lives in a space of polynomial dimension, which implies that such a finite set must be of exponential size! Luckily, as we will see below, there is a very easy solution: use as a witness the completely mixed state.

THEOREM 4 *For any logarithmic* $m = m(|x|)$, $\mathsf{QMA}_m(\frac{2}{3}, \frac{1}{3}) = \mathsf{BQP}$.

PROOF: The containment $\mathsf{BQP} \subseteq \mathsf{QMA}_m(\frac{2}{3}, \frac{1}{3})$ follows from the definitions. For the other direction, let L be a language in $\mathsf{QMA}_m(\frac{2}{3}, \frac{1}{3})$. According to witness-preserving amplification, we have that $L \in \mathsf{QMA}_m(\frac{3}{4}, \frac{1}{4}2^{-m})$, and let $\{V_x\}$ be the corresponding family of quantum verifiers. Consider the BQP machine that applies V_x with the witness being a completely mixed state on m qubits (i.e., having density matrix $I/2^m$), and accepts if and only if V_x accepts. An equivalent description is that we choose a random m-bit (classical) string, and use it as a witness to V_x.

For the analysis, define

$$Q_x = (I_m \otimes \langle 0|^k) V_x^\dagger \Delta V_x (I_m \otimes |0\rangle^k)$$

where Δ is the projector on the output qubit being 1. Then the probability that V_x accepts an m-qubit witness $|\varphi\rangle$ can be written as $\langle\varphi|Q_x|\varphi\rangle$. The acceptance probability of the BQP machine is thus

$$\frac{1}{2^m}\sum_{i\in\{0,1\}^m}\langle i|Q_x|i\rangle = \frac{1}{2^m}\mathrm{tr}(Q_x).$$

Another way to see that is to notice that

$$\frac{1}{2^m}\mathrm{tr}(Q_x) = \mathrm{tr}(\Delta\cdot V_x(2^{-m}I_m\otimes|0^k\rangle\langle 0^k|)V_x^\dagger)$$

and the right hand side is the exactly the probability to measure 1 in the output qubit of the state obtained by applying V_x to $2^{-m}I_m\otimes|0^k\rangle\langle 0^k|$.

If $x\in L$ then there is a witness $|\psi\rangle$ that is accepted by V_x with probability at least $\frac{3}{4}$ and so we have $\langle\psi|Q_x|\psi\rangle\geq\frac{3}{4}$. In particular, this implies that $\mathrm{tr}(Q_x)\geq\frac{3}{4}$ and hence the acceptance probability of the BQP machine is at least $\frac{3}{4}\cdot 2^{-m}$. On the other hand, if $x\notin L$ then V_x accepts any witness with probability at most $\frac{1}{4}2^{-m}$. This implies that $\mathrm{tr}(Q_x)\leq\frac{1}{4}$ and the acceptance probability of the BQP machine is at most $\frac{1}{4}\cdot 2^{-m}$. The difference between these two probabilities is inverse polynomial and hence this is indeed a BQP machine. □

References

[1] A. Ambainis, J. Kempe, and A. Rivosh. Coins make quantum walks faster. In *Proc. 16th ACM-SIAM SODA*, pages 1099–1108, 2005. quant-ph/0402107.

[2] R. Bhatia. *Matrix Analysis*. Number 169 in Graduate Texts in Mathematics. Springer-Verlag, New York, 1997.

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

[4] C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005.

[5] L. Masanes. Extremal quantum correlations for N parties with two dichotomic observables per site. quant-ph/0512100, 2005.

[6] M. Szegedy. Spectra of quantized walks and a $\sqrt{\delta\varepsilon}$ rule. quant-ph/0401053, 2004.

[7] B. Toner and F. Verstraete. Manuscript, 2006.

[8] J. Watrous. Zero-knowledge against quantum attacks. In *Proc. 38th ACM Symp. on Theory of Computing (STOC)*, 2006.