$\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}$: De-Merlinizing Quantum Protocols

Scott Aaronson*
University of Waterloo

Abstract

This paper introduces a new technique for removing existential quantifiers over quantum states. Using this technique, we show that there is no way to pack an exponential number of bits into a polynomial-size quantum state, in such a way that the value of any one of those bits can later be proven with the help of a polynomial-size quantum witness. We also show that any problem in QMA *with polynomial-size quantum advice, is also in* PSPACE *with polynomial-size classical advice. This builds on our earlier result that* $\mathsf{BQP/qpoly} \subseteq \mathsf{PP/poly}$*, and offers an intriguing counterpoint to the recent discovery of Raz that* $\mathsf{QIP/qpoly} = \mathsf{ALL}$*. Finally, we show that* $\mathsf{QCMA/qpoly} \subseteq \mathsf{PP/poly}$ *and that* $\mathsf{QMA/rpoly} = \mathsf{QMA/poly}$.

1. Introduction

Let Bob be a graduate student, and let x be an n-bit string representing his thesis problem. Bob's goal is to learn $f(x)$, where $f : \{0,1\}^n \to \{0,1\}$ is a function that maps every thesis problem to its binary answer ("yes" or "no"). Bob knows x (his problem), but is completely ignorant of f (how to solve the problem). So to evaluate $f(x)$, he's going to need help from his thesis advisor, Alice. Like most advisors, Alice is infinitely powerful, wise, and benevolent. But also like most advisors, she's too busy to find out what problems her students are working on. Instead, she just doles out the same advice s to all of them, which she hopes will let them evaluate $f(x)$ for any x they might encounter. The question is, how long does s have to be, for Bob to be able to evaluate $f(x)$ for any x?

Clearly, the answer is that s has to be 2^n bits long—since otherwise s will underdetermine the truth table of f. Indeed, let $g(x,s)$ be Bob's best guess as to $f(x)$, given x and s. Then even if Alice can choose s probabilistically, and we only require that $g(x,s) = f(x)$ with probability at least $2/3$ for every x, still one can show that s needs to be $\Omega(2^n)$ bits long.

*Email: scott@scottaaronson.com. Supported by ARDA, CIAR, and IQC. Part of this work was done at Caltech.

But what if Alice is a quantum advisor, who can send Bob a quantum state $|\psi_f\rangle$? Even in that case, Ambainis et al. [4] showed that Alice has to send $\Omega(2^n/n)$ qubits for Bob to succeed with probability at least $2/3$ on every x. Subsequently Nayak [11] improved this to $\Omega(2^n)$, meaning that there is no quantum improvement over the classical bound. Since 2^n qubits is too many for Alice to communicate during her weekly meetings with Bob, it seems Bob is out of luck.

So in desperation, Bob turns for help to Merlin, the star student in his department. Merlin knows f as well as x, and can thus evaluate $f(x)$. The trouble is that Merlin would prefer to take credit for evaluating $f(x)$ himself, so he might deliberately mislead Bob. Furthermore, Merlin (whose brilliance is surpassed only by his ego) insists that all communication with lesser students be one-way: Bob is to listen in silence while Merlin lectures him. On the other hand, Merlin has no time to give an exponentially long lecture, any more than Alice does.

With "helpers" like these, Bob might ask, who needs adversaries? And yet, is it possible that Bob could play Alice and Merlin against each other—cross-checking Merlin's specific but unreliable assertions against Alice's vague but reliable advice? In other words, does there exist a randomized protocol satisfying the following properties?

(i) Alice and Merlin both send Bob $\mathrm{poly}(n)$ bits.

(ii) If Merlin tells Bob the truth about $f(x)$, then there exists a message from Merlin that causes Bob to accept with probability at least $2/3$.

(iii) If Merlin lies about $f(x)$ (i.e., claims that $f(x) = 1$ when $f(x) = 0$ or vice versa), then no message from Merlin causes Bob to accept with probability greater than $1/3$.

It is relatively easy to show that the answer is no: if Alice sends a bits to Bob and Merlin sends w bits, then for Bob to succeed we must have $a(w+1) = \Omega(2^n)$. Indeed, this is basically tight: for all $w \geq 1$, there exists a protocol in which Merlin sends w bits and Alice sends $O\left(\frac{2^n}{w} + n\right)$ bits. Of course, even if Merlin didn't send anything, it would suffice for Alice to send 2^n bits. At the other extreme, if Merlin sends 2^n bits, then it suffices for Alice to send an $\Theta(n)$-bit "fingerprint" to authenticate Merlin's message. But in any event, either Alice or Merlin will have to send an exponentially-long message.

On the other hand, what if Alice and Merlin can both send *quantum* messages? Our main result will show that, even in this most general scenario, *Bob is still out of luck*. Indeed, if Alice sends a qubits to Bob, and Merlin sends w qubits, then Bob cannot succeed unless $a(w+1) = \Omega(2^n/n^2)$. Apart from the n^2 factor (which we conjecture can be removed), this implies that no quantum protocol is asymptotically better than the classical one. It follows, then, that Bob ought to drop out of grad school and send his resume to Google.

1.1. Banishing Merlin

But why should anyone care about this result, apart from Alice, Bob, Merlin, and the Google recruiters? One reason is that the proof introduces a new technique for removing existential quantifiers over quantum states, which might be useful in other contexts. The basic idea is for Bob to loop over all possible messages that Merlin could have sent, and accept if and only if there exists a message that would cause him to accept. The problem is that in the quantum case, the number of possible messages from Merlin is doubly-exponential. So to loop over all of them, it seems we'd first need to amplify Alice's message an exponential number of times. But surprisingly, we show that this intuition is wrong: to account for any possible quantum message from Merlin, it suffices to loop over all possible *classical* messages from Merlin! For, loosely speaking, any quantum state can eventually be detected by the "shadows" it casts on computational basis states. However, turning this insight into a "de-Merlinization" procedure requires some work: we need to amplify Alice's and Merlin's messages in a subtle way, and then deal with the degradation of Alice's message that occurs regardless.

1.2. QMA With Quantum Advice

In any case, the main motivation for our result is that it implies a new containment in quantum complexity theory: namely that

$$\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}.$$

Here QMA is the quantum version of MA, and $\mathsf{/qpoly}$ means "with polynomial-size quantum advice." Previously, it was not even known whether $\mathsf{QMA/qpoly} = \mathsf{ALL}$, where ALL is the class of all languages! Nevertheless, some context might be helpful for understanding why our new containment is of more than zoological interest.

Aaronson [1] showed that $\mathsf{BQP/qpoly} \subseteq \mathsf{PP/poly}$, where $\mathsf{BQP/qpoly}$ is the class of problems solvable in BQP with polynomial-size quantum advice. He also gave an oracle relative to which $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$. Together, these results seemed to place strong limits on the power of quantum advice.

However, recently Raz [14] reopened the subject, by showing that in some cases quantum advice can be extraordinarily powerful. In particular, Raz showed that $\mathsf{QIP}(2)/\mathsf{qpoly} = \mathsf{ALL}$, where $\mathsf{QIP}(2)$ is the class of problems that admit two-round quantum interactive proof systems. Raz's result was actually foreshadowed by an observation in [1], that $\mathsf{PostBQP/qpoly} = \mathsf{ALL}$. Here $\mathsf{PostBQP}$ is the class of problems solvable in quantum polynomial time, if at any time we can measure the computer's state and then "postselect" on a particular outcome occurring.[1]

These results should make any complexity theorist a little queasy, and not only because jumping from $\mathsf{QIP}(2)$ or $\mathsf{PostBQP}$ to ALL is like jumping from a hilltop to the edge of the universe. A more serious problem is that these results fail to "commute" with standard complexity inclusions. For example, even though $\mathsf{PostBQP}$ is strictly contained in $\mathsf{BQEXPEXP}$, notice that $\mathsf{BQEXPEXP/qpoly}$ is (very) strictly contained in $\mathsf{PostBQP/qpoly}$!

1.3. The Quantum Advice Hypothesis

On the other hand, the same pathologies would occur with classical randomized advice. For neither the result of Raz [14], nor that of Aaronson [1], makes any essential use of quantum mechanics. That is, instead of saying that

$$\mathsf{QIP}(2)/\mathsf{qpoly} = \mathsf{PostBQP/qpoly} = \mathsf{ALL},$$

we could equally well have said that

$$\mathsf{IP}(2)/\mathsf{rpoly} = \mathsf{PostBPP/rpoly} = \mathsf{ALL},$$

where $\mathsf{IP}(2)$ and $\mathsf{PostBPP}$ are the classical analogues of $\mathsf{QIP}(2)$ and $\mathsf{PostBQP}$ respectively, and $\mathsf{/rpoly}$ means "with polynomial-size randomized advice."

Inspired by this observation, here we propose a general hypothesis: that *whenever quantum advice behaves like exponentially-long classical advice, the reason has nothing to do with quantum mechanics*. More concretely:

- **The Quantum Advice Hypothesis:** For any "natural" complexity class $\mathcal{C}$, if $\mathcal{C}/\mathsf{qpoly} = \mathsf{ALL}$, then $\mathcal{C}/\mathsf{rpoly} = \mathsf{ALL}$ as well.

[1] Here is the proof: given a Boolean function $f : \{0,1\}^n \to \{0,1\}$, take

$$|\psi_n\rangle = \frac{1}{2^{n/2}} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle$$

as the advice. Then to evaluate $f(x)$ on any x, simply measure $|\psi_n\rangle$ in the standard basis, and then postselect on observing $|x\rangle$ in the first register.

The evidence for this hypothesis is simply that we have not been able to refute it. In particular, in Appendix 7 we will show that $\mathsf{QMA/rpoly} = \mathsf{QMA/poly}$. So if $\mathsf{QMA/qpoly}$ contained all languages—which (at least to us) seemed entirely possible *a priori*—then we would have a clear counterexample to the hypothesis. In our view, then, the significance of the $\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}$ result is that it confirms the quantum advice hypothesis in the most nontrivial case considered so far.

To summarize, the quantum advice hypothesis has been confirmed for at least four complexity classes: BQP, QMA, $\mathsf{PostBQP}$, and $\mathsf{QIP}(2)$. It remains open for other classes, such as $\mathsf{QMA}(2)$ (QMA with two unentangled yes-provers) and QS_2^p (QMA with competing yes-prover and no-prover).

1.4. Outline of Paper

- Section 2 surveys the complexity classes, communication complexity measures, and quantum information notions used in this paper.
- Section 3 states our "De-Merlinization Theorem," and then proves three of its implications: (i) a lower bound on the QMA communication complexity of random access coding, (ii) a general lower bound on QMA communication complexity, and (iii) the inclusion $\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}$.
- Section 4 proves the De-Merlinization Theorem itself.
- Section 5 concludes with some open problems.
- Appendix 7 proves a few other complexity results, including $\mathsf{QCMA/qpoly} \subseteq \mathsf{PP/poly}$ and $\mathsf{QMA/rpoly} = \mathsf{QMA/poly}$.

2. Preliminaries

2.1. Complexity Classes

We assume familiarity with standard complexity classes like BQP, $\mathsf{P/poly}$, and MA. The class QMA (Quantum Merlin-Arthur) consists of all languages for which a 'yes' answer can be verified in quantum polynomial time, given a polynomial-size quantum witness state $|\varphi\rangle$. The completeness and soundness errors are $1/3$. The class QCMA (Quantum Classical Merlin-Arthur) is the same as QMA, except that now the witness must be classical. It is not known whether $\mathsf{QMA} = \mathsf{QCMA}$. See the Complexity Zoo[2] for more information about these and other classes.

[2] http://qwiki.caltech.edu/wiki/Complexity_Zoo

Given a complexity class $\mathcal{C}$, we write $\mathcal{C}/\mathsf{poly}$, $\mathcal{C}/\mathsf{rpoly}$, and $\mathcal{C}/\mathsf{qpoly}$ to denote $\mathcal{C}$ with polynomial-size deterministic, randomized, and quantum advice respectively.[3] So for example, $\mathsf{BPP/rpoly}$ is the class of languages decidable by a BPP machine, given a sample from a distribution $\mathcal{D}_n$ over polynomial-size advice strings which depends only on the input length n. It is clear that $\mathsf{BPP/rpoly} = \mathsf{BPP/poly} = \mathsf{P/poly}$. However, in other cases the statement $\mathcal{C}/\mathsf{rpoly} = \mathcal{C}/\mathsf{poly}$ is harder to prove or is even false.

Admittedly, the $/\mathsf{rpoly}$ and $/\mathsf{qpoly}$ operators are not always well-defined: for example, $\mathsf{P/qpoly}$ is just silly, and $\mathsf{AM/rpoly}$ seems ambiguous (since who gets to sample from the advice distribution?). For interactive proof classes, the general rule we adopt is that *only the verifier gets to "measure" the advice*. In other words, the prover (or provers) knows the advice distribution $\mathcal{D}_n$ or advice state $|\psi_n\rangle$, but not the actual results of sampling from $\mathcal{D}_n$ or measuring $|\psi_n\rangle$. In the case of $/\mathsf{rpoly}$, the justification for this rule is that, if the prover knew the sample from $\mathcal{D}_n$, then we would immediately get $\mathcal{C}/\mathsf{rpoly} = \mathcal{C}/\mathsf{poly}$ for all interactive proof classes $\mathcal{C}$, which is too boring. In the case of $/\mathsf{qpoly}$, the justification is that the verifier should be allowed to measure $|\psi_n\rangle$ at any time and in any basis it likes, and it seems perverse to require the results of such measurements to be relayed instantly to the prover.

In a private-coin protocol, the verifier might choose to reveal some or all of the measurement results to the prover, but in a public-coin protocol, the verifier must send a uniform random message that is uncorrelated with the advice. Indeed, this explains how it can be true that $\mathsf{IP}(2)/\mathsf{rpoly} \neq \mathsf{AM/rpoly}$ (the former equals ALL, while the latter equals $\mathsf{NP/poly}$), even though Goldwasser and Sipser [5] famously showed that $\mathsf{IP}(2) = \mathsf{AM}$ in the uniform setting.

For the complexity classes $\mathcal{C}$ that appear in this paper, it should generally be obvious what we mean by $\mathcal{C}/\mathsf{rpoly}$ or $\mathcal{C}/\mathsf{qpoly}$. But to fix ideas, let us now formally define $\mathsf{QMA/qpoly}$.

Definition 1 $\mathsf{QMA/qpoly}$ *is the class of languages* $L \subseteq \{0,1\}^*$ *for which there exists a polynomial-time quantum verifier* $\mathcal{Q}$*, together with quantum advice states* $\{|\psi_n\rangle\}_{n\geq 1}$*, such that for all* $x \in \{0,1\}^n$*:*

(i) If $x \in L$*, then there exists a quantum witness* $|\varphi\rangle$ *such that* $\mathcal{Q}$ *accepts with probability at least* $2/3$ *given* $|x\rangle|\psi_n\rangle|\varphi\rangle$ *as input.*

(ii) If $x \notin L$*, then for all pure states*[4] $|\varphi\rangle$ *of the witness register,* $\mathcal{Q}$ *accepts with probability at most* $1/3$ *given* $|x\rangle|\psi_n\rangle|\varphi\rangle$ *as input.*

[3] We can also write $\mathcal{C}/\mathsf{rlog}$ (for $\mathcal{C}$ with logarithmic-size randomized advice), $\mathcal{C}/\mathsf{qlog}$, and so on.

[4] By linearity, this is equivalent to quantifying over all mixed states of the witness register.

Here $|\psi_n\rangle$ and $|\varphi\rangle$ both consist of $p(n)$ qubits for some fixed polynomial p. Also, $\mathcal{Q}$ can accept with arbitrary probability if given a state other than $|\psi_n\rangle$ in the advice register.

One other complexity class we will need is PostBQP, or BQP with postselection.

Definition 2 PostBQP *is the class of languages $L \subseteq \{0,1\}^*$ for which there exists a polynomial-time quantum algorithm such that for all $x \in \{0,1\}^n$, when the algorithm terminates:*

(i) The first qubit is $|1\rangle$ with nonzero probability.

(ii) If $x \in L$, then conditioned on the first qubit being $|1\rangle$, the second qubit is $|1\rangle$ with probability at least $2/3$.

(iii) If $x \notin L$, then conditioned on the first qubit being $|1\rangle$, the second qubit is $|1\rangle$ with probability at most $1/3$.

One can similarly define PostBQPSPACE, PostBQEXP, and so on. We will use a result of Aaronson [2], which characterizes PostBQP as simply the classical complexity class PP.

2.2. Communication Complexity

Let $f : \{0,1\}^N \times \{0,1\}^M \rightarrow \{0,1\}$ be a Boolean function. Suppose Alice has an N-bit string X and Bob has an M-bit string Y. Then $\mathrm{D}^1(f)$ is the deterministic one-way communication complexity of f: that is, the minimum number of bits that Alice must send to Bob, for Bob to be able to output $f(X,Y)$ with certainty for any (X,Y) pair. If we let Alice's messages be randomized, and only require Bob to be correct with probability $2/3$, then we obtain $\mathrm{R}^1(f)$, the bounded-error randomized one-way communication complexity of f. Finally, if we let Alice's messages be quantum, then we obtain $\mathrm{Q}^1(f)$, the bounded-error quantum one-way communication complexity of f.[5] Clearly $\mathrm{Q}^1(f) \leq \mathrm{R}^1(f) \leq \mathrm{D}^1(f)$ for all f. See Klauck [7] for more detailed definitions of these measures.

Now suppose that, in addition to a quantum message $|\psi_X\rangle$ from Alice, Bob also receives a quantum witness $|\varphi\rangle$ from Merlin, whose goal is to convince Bob that $f(X,Y) = 1$.[6] We say Alice and Bob *succeed* if for all X, Y,

(i) If $f(X,Y) = 1$, then there exists a $|\varphi\rangle$ such that Bob accepts $|Y\rangle|\psi_X\rangle|\varphi\rangle$ with probability at least $2/3$.

(ii) If $f(X,Y) = 0$, then for all $|\varphi\rangle$, Bob accepts $|Y\rangle|\psi_X\rangle|\varphi\rangle$ with probability at most $1/3$.

Call a protocol "(a,w)" if Alice's message consists of a qubits and Merlin's consists of w qubits. Then for all integers $w \geq 0$, we let $\mathrm{QMA}_w^1(f)$ denote the "QMA_w one-way communication complexity" of f: that is, the minimum a for which there exists an (a,w) protocol such that Alice and Bob succeed. Clearly $\mathrm{QMA}_w^1(f) \leq \mathrm{Q}^1(f)$, with equality when $w = 0$.

2.3. Quantum Information

Here we review some basic facts about mixed states. Further details can be found in Nielsen and Chuang [13] for example.

Given two mixed states ρ and σ, the *fidelity* $F(\rho,\sigma)$ is the maximum possible value of $\langle\psi|\varphi\rangle$, where $|\psi\rangle$ and $|\varphi\rangle$ are purifications of ρ and σ respectively. Also, given a measurement M, let $\mathcal{D}_M(\rho)$ be the probability distribution over measurement outcomes if M is applied to ρ. Then the *trace distance* $\|\rho - \sigma\|_{\mathrm{tr}}$ equals the maximum, over all possible measurements M, of $\|\mathcal{D}_M(\rho) - \mathcal{D}_M(\sigma)\|$, where

$$\|\mathcal{D} - \mathcal{D}'\| = \frac{1}{2}\sum_{i=1}^{N} |p_i - p_i'|$$

is the variation distance between $\mathcal{D} = (p_1, \ldots, p_N)$ and $\mathcal{D}' = (p_1', \ldots, p_N')$. For all ρ and σ, we have the following relation between fidelity and trace distance:

$$\|\rho - \sigma\|_{\mathrm{tr}} \leq \sqrt{1 - F(\rho,\sigma)^2}.$$

Throughout this paper, we use $\mathcal{H}_N$ to denote N-dimensional Hilbert space. One fact we will invoke repeatedly is that, if I is the maximally mixed state in $\mathcal{H}_N$, then

$$I = \frac{1}{N}\sum_{j=1}^{N} |\psi_j\rangle\langle\psi_j|$$

where $\{|\psi_1\rangle, \ldots, |\psi_N\rangle\}$ is *any* orthonormal basis for $\mathcal{H}_N$.

3. De-Merlinization and Its Applications

Our main result, the "De-Merlinization Theorem," allows us to lower-bound $\mathrm{QMA}_w^1(f)$ in terms of the ordinary quantum communication complexity $\mathrm{Q}^1(f)$. In this section we state the theorem and derive its implications for random access coding (in Section 3.1), one-way communication complexity (in Section 3.2), and complexity theory (in Section 3.3). The theorem itself will be proved in Section 4.

[5] We assume no shared randomness or entanglement. Also, we assume for simplicity that Alice can only send pure states; note that this increases the message length by at most a multiplicative factor of 2 (or an additive factor of $\log N$, if we use Newman's Theorem [12]).

[6] For convenience, from now on we assume that Merlin only needs to prove statements of the form $f(X,Y) = 1$, not $f(X,Y) = 0$. For our actual results, it will make no difference whether we adopt this assumption (corresponding to the class QMA), or the assumption in Section 1 (corresponding to QMA $\cap$ coQMA).

Theorem 3 (De-Merlinization Theorem) *For all Boolean functions f (partial or total) and all $w \geq 2$,*

$$\mathrm{Q}^1(f) = O\left(\mathrm{QMA}^1_w(f) \cdot w \log^2 w\right).$$

Furthermore, given an algorithm for the QMA^1_w protocol, Bob can efficiently generate an algorithm for the Q^1 protocol. If the former uses C gates and S qubits of memory, then the latter uses $C \cdot S^{O(S)}$ gates and $O\left(S^2 \log^2 S\right)$ qubits of memory.

3.1. Application I: Random Access Coding

Following Ambainis et al. [4], let us define the *random access coding* (or RAC) problem as follows. Alice has an N-bit string $X = x_1 \ldots x_N$ and Bob has an index $i \in \{1, \ldots, N\}$. The players' goal is for Bob to learn x_i.

In our setting, Bob receives not only an a-bit message from Alice, but also a w-bit message from Merlin. If $x_i = 1$, then there should exist a message from Merlin that causes Bob to accept with probability at least $2/3$; while if $x_i = 0$, then no message from Merlin should cause Bob to accept with probability greater than $1/3$. We are interested in the minimum a, w for which Alice and Bob can succeed.

For completeness, before stating our results for the quantum case, let us first pin down the classical case—that is, the case in which Alice and Merlin both send classical messages, and Alice's message can be randomized. Obviously, if Merlin sends 0 bits, then Alice needs to send $\Theta(N)$ bits; this is just the ordinary RAC problem studied by Ambainis et al. [4]. At the other extreme, if Merlin sends the N-bit message X, then it suffices for Alice to send an $O(\log N)$-bit fingerprint of X. For intermediate message lengths, we can interpolate between these two extremes.

Theorem 4 *For all a, w such that $aw \geq N$, there exists a randomized $(a + O(\log N), w)$ protocol for RAC—that is, a protocol in which Alice sends $a + O(\log N)$ bits and Merlin sends w bits.*

Proof. The protocol is as follows: first Alice divides her string $X = x_1 \ldots x_N$ into a substrings $Y_1, \ldots, Y_a$, each at most w bits long. She then maps each Y_j to an encoded substring $Y'_j = g(Y_j)$, where $g : \{0,1\}^w \rightarrow \{0,1\}^W$ is a constant-rate error-correcting code satisfying $W = O(w)$. Next she chooses $k \in \{1, \ldots, W\}$ uniformly at random. Finally, she sends Bob k (which requires $O(\log N)$ bits of communication), together with the k^{th} bit of Y'_j for every $j \in \{1, \ldots, a\}$.

Now if Merlin is honest, then he sends Bob the substring $Y_j \in \{0,1\}^w$ of X containing the x_i that Bob is interested in. This allows Bob to learn x_i. Furthermore, if Merlin cheats by sending some $Y \neq Y_j$, then Bob can detect this with constant probability, by cross-checking the k^{th} bit of $g(Y)$ against the k^{th} bit of Y'_j as sent by Alice. ■

Using a straightforward amplification trick, we can show that the protocol of Theorem 4 is essentially optimal.

Theorem 5 *If there exists a randomized (a, w) protocol for RAC, then $a(w+1) = \Omega(N)$ and $a = \Omega(\log N)$.*

Proof. We first show that $a(w+1) = \Omega(N)$. First Alice amplifies her message to Bob by sending $W = O(w+1)$ independent copies of it. For any fixed message of Merlin, this reduces Bob's error probability to at most (say) $2^{-2(w+1)}$. So now Bob can ignore Merlin, and loop over all 2^w messages $z \in \{0,1\}^w$ that Merlin *could* have sent, accepting if and only if there exists a z that would cause him to accept. This yields an ordinary protocol for the RAC problem in which Alice sends aW bits to Bob. But Ambainis et al. [4] showed that any such protocol requires $\Omega(N)$ bits; hence $a(w+1) = \Omega(N)$.

That Alice needs to send $\Omega(\log N)$ bits follows by a simple counting argument: let $\mathcal{D}_X$ be Alice's message distribution given an input X. Then $\mathcal{D}_X$ and $\mathcal{D}_Y$ must have constant variation distance for all $X \neq Y$, if Bob is to distinguish X from Y with constant bias. ■

Together, Theorems 4 and 5 provide the complete story for the classical case, up to a constant factor. In the quantum case, the situation is no longer so simple, but we can give a bound that is tight up to a polylog factor.

Theorem 6 *If there exists a quantum (a, w) protocol for RAC, then*

$$a(w+1) = \Omega\left(\frac{N}{\log^2 N}\right).$$

Proof. If $w = 0$ or $w = 1$ then clearly $a = \Omega(N)$, so assume $w \geq 2$. By Theorem 3,

$$\begin{aligned} \mathrm{Q}^1(\mathrm{RAC}) &= O\left(\mathrm{QMA}^1_w(\mathrm{RAC}) \cdot w \log^2 w\right) \\ &= O\left(aw \cdot \log^2 N\right). \end{aligned}$$

But Nayak [11] showed that $\mathrm{Q}^1(\mathrm{RAC}) = \Omega(N)$, and hence $aw = \Omega\left(N / \log^2 N\right)$. ■

Clearly Theorem 6 can be improved when w is very small or very large. For when $w = 0$, we have $a = \Omega(N)$; while for any w, a simple counting argument (as in the classical case) yields $a = \Omega(\log N)$. We believe that Theorem 6 can be improved for intermediate w as well, since we do not know of any quantum protocol that beats the classical protocol of Theorem 4.

3.2. Application II: One-Way Communication

Theorem 3 yields lower bounds on QMA communication complexity, not only for the random access coding problem, but for other problems as well. For Aaronson [1] showed the following general relationship between $\mathrm{D}^1(f)$ and $\mathrm{Q}_2^1(f)$:

Theorem 7 ([1]) *For all Boolean functions $f : \{0,1\}^N \times \{0,1\}^M \to \{0,1\}$ (partial or total),*

$$\mathrm{D}^1(f) = O\left(M\,\mathrm{Q}_2^1(f)\log \mathrm{Q}_2^1(f)\right).$$

Combining Theorem 7 with Theorem 3, we obtain the following relationship between $\mathrm{D}^1(f)$ and $\mathrm{QMA}_w^1(f)$. For all $f : \{0,1\}^N \times \{0,1\}^M \to \{0,1\}$ (partial or total) and all $w \geq 2$,

$$\mathrm{D}^1(f) = O\left(M \cdot w\log^3 w \cdot \mathrm{QMA}_w^1(f)\log \mathrm{QMA}_w^1(f)\right).$$

3.3. Application III: Upper-Bounding QMA/qpoly

We now explain why the containment $\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}$ follows from the De-Merlinization Theorem. The first step is to observe a weaker result that follows from that theorem:

Lemma 8 $\mathsf{QMA/qpoly} \subseteq \mathsf{BQPSPACE/qpoly}$.

Proof. Given a language $L \in \mathsf{QMA/qpoly}$, let $L_n : \{0,1\}^n \to \{0,1\}$ be the Boolean function defined by $L_n(x) = 1$ if $x \in L$ and $L_n(x) = 0$ otherwise. Then if we interpret Alice's input as the truth table of L_n, Bob's input as x, and S as the number of qubits used by the $\mathsf{QMA/qpoly}$ machine, the lemma follows immediately from Theorem 3. ■

Naïvely, Lemma 8 might seem obvious, since it is well-known that $\mathsf{QMA} \subseteq \mathsf{PSPACE}$. But remember that even if $\mathcal{C} \subseteq \mathcal{D}$, it need not follow that $\mathcal{C}/\mathsf{qpoly} \subseteq \mathcal{D}/\mathsf{qpoly}$.

The next step is to replace the quantum advice by classical advice.

Lemma 9 $\mathsf{BQPSPACE/qpoly} \subseteq \mathsf{PostBQPSPACE/poly}$.

Proof. Follows from the same argument used by Aaronson [1] to show that $\mathsf{BQP/qpoly} \subseteq \mathsf{PostBQP/poly}$. All we need to do is replace polynomial time by polynomial space. ■

Finally, we observe a simple generalization of Watrous's theorem [15] that $\mathsf{BQPSPACE} = \mathsf{PSPACE}$.

Lemma 10 $\mathsf{PostBQPSPACE} = \mathsf{PSPACE}$.

Proof Sketch. Ladner [9] showed that $\mathsf{PPSPACE} = \mathsf{PSPACE}$. Intuitively, given the computation graph of a $\mathsf{PPSPACE}$ machine, we want to decide in PSPACE whether the number of accepting paths exceeds the number of rejecting paths. To do so we use divide-and-conquer, as in the proof of Savitch's theorem that $\mathsf{NPSPACE} = \mathsf{PSPACE}$. An obvious difficulty is that the numbers of paths could be *doubly* exponential, and therefore take exponentially many bits to store. But we can deal with that by computing each bit of the numbers separately. Here we use the fact that there exist NC^1 circuits for addition, and hence addition of 2^n-bit integers is "locally" in PSPACE.

If each path is weighted by a complex amplitude, then it is easy to see that the same idea lets us sum the amplitudes over all paths. We can thereby simulate $\mathsf{BQPSPACE}$ and $\mathsf{PostBQPSPACE}$ in PSPACE as well. ■

In particular, Lemma 10 implies that $\mathsf{PostBQPSPACE/poly} = \mathsf{PSPACE/poly}$. (For note that unlike randomized and quantum advice, deterministic advice commutes with standard complexity class inclusions.)

Putting it all together, we obtain:

Theorem 11 $\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}$.

As a final remark, let QAM be the quantum analogue of AM, in which Arthur sends a public random string to Merlin, and then Merlin responds with a quantum state. Marriott and Watrous [10] observed that $\mathsf{QAM} = \mathsf{BP}{\cdot}\mathsf{QMA}$. So

$$\mathsf{QAM/qpoly} = \mathsf{BP}{\cdot}\mathsf{QMA/qpoly} = \mathsf{QMA/qpoly},$$

since we can hardwire the random string into the quantum advice. Hence $\mathsf{QAM/qpoly} \subseteq \mathsf{PSPACE/poly}$ as well. This offers an interesting contrast with the result of Raz [14] that $\mathsf{QIP}(2)/\mathsf{qpoly} = \mathsf{ALL}$.

4. Proof of The De-Merlinization Theorem

We now proceed to the proof of Theorem 3. In Section 4.1 we prove several lemmas about damage to quantum states, and in particular, the effect of the damage caused by earlier measurements of a state on the outcomes of later measurements. Section 4.2 then gives our procedure for amplifying Bob's error probability, after explaining why the more obvious procedures fail. Finally, Section 4.3 puts together the pieces.

4.1. Quantum Information Lemmas

In this section we prove several lemmas that will be needed for the main result. The first lemma is a simple variant of Lemma 2.2 from [1]; we include a proof for completeness.

Lemma 12 (Almost As Good As New Lemma) *Suppose a 2-outcome POVM measurement of a mixed state ρ yields outcome 1 with probability ε. Then after the measurement, and assuming outcome 0 is observed, we obtain a new state ρ_0 such that $\|\rho - \rho_0\|_{\mathrm{tr}} \leq \sqrt{\varepsilon}$.*

Proof. Let $|\psi\rangle$ be a purification of ρ. Then we can write $|\psi\rangle$ as $\sqrt{1-\varepsilon}\,|\psi_0\rangle + \sqrt{\varepsilon}\,|\psi_1\rangle$, where $|\psi_0\rangle$ is a purification of ρ_0 and $\langle\psi_0|\psi_1\rangle = 0$. So the fidelity between ρ and ρ_0 is

$$F(\rho,\rho_0) \geq \langle\psi|\psi_0\rangle = \sqrt{1-\varepsilon}.$$

Therefore

$$\|\rho-\rho_0\|_{\mathrm{tr}} \leq \sqrt{1-F(\rho,\rho_0)^2} \leq \sqrt{\varepsilon}.$$

■

The next lemma, which we call the "quantum union bound," abstracts one of the main ideas from [4].

Lemma 13 (Quantum Union Bound) *Let ρ be a mixed state, and let $\{\Lambda_1,\ldots,\Lambda_T\}$ be a set of 2-outcome POVM measurements. Suppose each Λ_t yields outcome 1 with probability at most ε when applied to ρ. Then if we apply $\Lambda_1,\ldots,\Lambda_T$ in sequence to ρ, the probability that at least one of these measurements yields outcome 1 is at most $T\sqrt{\varepsilon}$.*

Proof. Follows from a hybrid argument, almost identical to Claim 4.1 of Ambainis et al. [4]. More explicitly, by the principle of deferred measurement, we can replace each measurement Λ_t by a unitary U_t that CNOT's the measurement outcome into an ancilla qubit. Let $\rho_0 = \rho\otimes|0\cdots0\rangle\langle0\cdots0|$ be the initial state of the system plus T ancilla qubits. Then by the same idea as in Lemma 12, for all t we have

$$\left\|U_t\rho_0U_t^{-1}-\rho_0\right\|_{\mathrm{tr}} \leq \sqrt{\varepsilon}.$$

So letting

$$\rho_t := U_T\cdots U_{T-t+1}\rho_0U_{T-t+1}^{-1}\cdots U_T^{-1},$$

by unitarity we also have

$$\begin{aligned}\|\rho_{t+1}-\rho_t\|_{\mathrm{tr}} &= \left\|\begin{array}{c} U_T\cdots U_{T-t}\rho_0U_{T-t}^{-1}\cdots U_T^{-1} \\ -U_T\cdots U_{T-t+1}\rho_0U_{T-t+1}^{-1}\cdots U_T^{-1}\end{array}\right\|_{\mathrm{tr}} \\ &= \left\|U_{T-t}\rho_0U_{T-t}^{-1}-\rho_0\right\|_{\mathrm{tr}} \\ &\leq \sqrt{\varepsilon},\end{aligned}$$

and hence $\|\rho_T-\rho_0\|_{\mathrm{tr}} \leq T\sqrt{\varepsilon}$ by the triangle inequality.

Now let M be a measurement that returns the logical OR of the T ancilla qubits, and let $\mathcal{D}(\rho)$ be the distribution over the outcomes (0 and 1) when M is applied to ρ. Suppose M yields outcome 1 with probability p when applied to ρ_T. Then since M yields outcome 1 with probability 0 when applied to ρ_0, the variation distance $\|\mathcal{D}(\rho_T)-\mathcal{D}(\rho_0)\|$ is equal to p. So by the definition of trace distance,

$$p \leq \|\rho_T-\rho_0\|_{\mathrm{tr}} \leq T\sqrt{\varepsilon}.$$

■

Finally, we give a lemma that is key to our result. This lemma, which we call the "quantum OR bound," is a sort of converse to the quantum union bound. It says that, for all quantum circuits Λ and advice states $|\psi\rangle$, if there exists a witness state $|\varphi\rangle$ such that Λ accepts $|\psi\rangle|\varphi\rangle$ with high probability, then we can also cause Λ to accept with high probability by repeatedly running Λ on $|\psi\rangle|j\rangle$, where $|j\rangle$ is a random basis state of the witness register, and then taking the logical OR of the outcomes. One might worry that, as we run Λ with various $|j\rangle$'s, the state of the advice register might become corrupted to something far from $|\psi\rangle$. However, we show that if this happens, then it can only be because one of the measurements has already accepted with high probability.

Lemma 14 (Quantum OR Bound) *Let Λ be a 2-outcome POVM measurement on a bipartite Hilbert space $\mathcal{H}_A\otimes\mathcal{H}_B$. Also, let $\{|1\rangle,\ldots,|N\rangle\}$ be any orthonormal basis for $\mathcal{H}_B$, and for all $j\in\{1,\ldots,N\}$, let Λ_j be the POVM on $\mathcal{H}_A$ induced by applying Λ to $\mathcal{H}_A\otimes|j\rangle$. Suppose there exists a product state $\rho\otimes\sigma$ in $\mathcal{H}_A\otimes\mathcal{H}_B$ such that Λ yields outcome 1 with probability at least $\eta>0$ when applied to $\rho\otimes\sigma$. Then if we apply $\Lambda_{j_1},\ldots,\Lambda_{j_T}$ in sequence to ρ, where $j_1,\ldots,j_T$ are drawn uniformly and independently from $\{1,\ldots,N\}$ and $T\geq N/\eta^2$, the probability that at least one of these measurements yields outcome 1 is at least $\left(\eta-\sqrt{N/T}\right)^2$.*

Proof. Let E_t denote the event that one of the first t measurements of ρ yields outcome 1. Also, let $\alpha := \left(\eta-\sqrt{N/T}\right)^2$. Then our goal is to show that $\Pr[E_t]\geq\alpha$ for some t, where the probability is over the choice of $j_1,\ldots,j_T$ as well as the measurement outcomes. Suppose $\Pr[E_t]<\alpha$ for all t; we will derive a contradiction.

Let ρ_t be the state in $\mathcal{H}_A$ after the first t measurements, averaged over all choices of $j_1,\ldots,j_t$ and assuming E_t does not occur. Suppose $\|\rho_t-\rho\|_{\mathrm{tr}}>\sqrt{\alpha}$ for some t. Then interpreting the first t measurements as a single measurement, and taking the contrapositive of Lemma 12, we find that $\Pr[E_t]>\alpha$, and we are done. So we can assume without loss of generality that $\|\rho_t-\rho\|_{\mathrm{tr}}\leq\sqrt{\alpha}$ for all t.

For all mixed states ς in $\mathcal{H}_A\otimes\mathcal{H}_B$, let $P_\Lambda(\varsigma)$ be the probability that Λ yields outcome 1 when applied to ς. By the definition of trace distance, we have

$$P_\Lambda(\varsigma')\geq P_\Lambda(\varsigma)-\|\varsigma-\varsigma'\|_{\mathrm{tr}}$$

for all ς,ς'. Therefore

$$\begin{aligned}P_\Lambda(\rho_t\otimes\sigma) &\geq P_\Lambda(\rho\otimes\sigma)-\|\rho_t\otimes\sigma-\rho\otimes\sigma\|_{\mathrm{tr}} \\ &= P_\Lambda(\rho\otimes\sigma)-\|\rho_t-\rho\|_{\mathrm{tr}} \\ &\geq \eta-\sqrt{\alpha}.\end{aligned}$$

Hence

$$P_\Lambda(\rho_t\otimes I)\geq\frac{\eta-\sqrt{\alpha}}{N},$$

where

$$I = \frac{1}{N} \sum_{j=1}^{N} |j\rangle \langle j|$$

is the maximally mixed state in $\mathcal{H}_B$. It follows that for all t,

$$\underset{j \in \{1,\ldots,N\}}{\mathrm{EX}} \left[P_\Lambda \left(\rho_t \otimes |j\rangle \langle j| \right) \right] \geq \frac{\eta - \sqrt{\alpha}}{N}.$$

Now notice that

$$\Pr\left[E_t | \neg E_{t-1}\right] = \underset{j_t \in \{1,\ldots,N\}}{\mathrm{EX}} \left[P_\Lambda \left(\rho_{t-1} \otimes |j_t\rangle \langle j_t| \right) \right]$$

for all t. Furthermore, since $E_{t-1} \Rightarrow E_t$, the events $\neg E_{t-1} \wedge E_t$ are disjoint. Therefore

$$\begin{aligned}
\Pr\left[E_T\right] &= \sum_{t=1}^{T} \Pr\left[\neg E_{t-1} \wedge E_t\right] \\
&= \sum_{t=1}^{T} \Pr\left[\neg E_{t-1}\right] \Pr\left[E_t | \neg E_{t-1}\right] \\
&\geq \sum_{t=1}^{T} (1-\alpha) \cdot \underset{j_t \in \{1,\ldots,N\}}{\mathrm{EX}} \left[P_\Lambda \left(\rho_{t-1} \otimes |j_t\rangle \langle j_t| \right) \right] \\
&\geq (1-\alpha)\, T \left(\frac{\eta - \sqrt{\alpha}}{N} \right) \\
&\geq \left(\eta - \sqrt{\alpha} \right)^2 \frac{T}{N} \\
&= 1,
\end{aligned}$$

which is certainly greater than $\alpha = \left(\eta - \sqrt{N/T} \right)^2$. Here we are using the fact that $T \geq N/\eta^2$, and hence $\alpha \leq 1$. ■

4.2. Amplification

Before proceeding further, we need to decrease Bob's soundness error (that is, the probability that he accepts a dishonest claim from Merlin). The simplest approach would be to have Alice and Merlin both send ℓ copies of their messages for some ℓ, and then have Bob run his verification algorithm ℓ times in parallel and output the majority answer. However, this approach fails, since the decrease in error probability is more than cancelled out by the *increase* in Merlin's message length (recall that we will have to loop over all possible classical messages from Merlin). So then why not use the "in-place amplification" technique of Marriott and Watrous [10]? Because unfortunately, that technique only works for Merlin's message; we do not know whether it can be generalized to handle Alice's message as well.[7] Happily, there is a "custom" amplification procedure with the properties we want:

[7] In any such generalization, certainly Alice will still have to send multiple copies of her message. The question is whether Merlin will also have to send multiple copies of *his* message.

Lemma 15 *Suppose Bob receives an a-qubit message $|\psi\rangle$ from Alice and a w-qubit message $|\varphi\rangle$ from Merlin, where $w \geq 2$. Let $A = O\left(aw \log^2 w\right)$ and $W = O(w \log w)$. Then by using A qubits from Alice and W qubits from Merlin, Bob can amplify his soundness error to 5^{-W} while keeping his completeness error $1/3$.*

Proof. We will actually use two layers of amplification. In the "inner" layer, we replace Alice's message $|\psi\rangle$ by the $a\ell$-qubit message $|\psi\rangle^{\otimes \ell}$, where $\ell = O(\log w)$. We also replace Merlin's message $|\varphi\rangle$ by the $w\ell$-qubit message $|\varphi\rangle^{\otimes \ell}$. We then run Bob's algorithm ℓ times in parallel and output the majority answer. By a Chernoff bound, together with the same observations used by Kitaev and Watrous [6] to show amplification for QMA, this reduces both the completeness and the soundness errors to $\varepsilon = \frac{1}{1000w^3}$, for suitable $\ell = O(\log w)$.

In the "outer" layer, we replace Alice's message by $|\psi\rangle^{\otimes \ell u}$, where $u = O(W)$. We then run the inner layer u times, once for each copy of $|\psi\rangle^{\otimes \ell}$, but reusing the same register for Merlin's message each time. (Also, after each invocation of the inner layer, we uncompute everything except the final answer.) Finally, we output the majority answer among these u invocations.

Call Bob's original algorithm $\mathcal{Q}$, and call the amplified algorithm $\mathcal{Q}_*$. Then our first claim is that if $\mathcal{Q}$ accepts all w-qubit messages from Merlin with probability at most $1/3$, then $\mathcal{Q}_*$ accepts all W-qubit messages with probability at most 5^{-W}, for suitable $u = O(W)$. This follows from a Chernoff bound—since even if we condition on the first through t^{th} invocations of the inner layer, the $(t+1)^{st}$ invocation will still receive a "fresh" copy of $|\psi\rangle^{\otimes \ell}$, and will therefore accept with probability at most $\varepsilon \leq 1/3$. The state of Merlin's message register before the $(t+1)^{st}$ invocation is irrelevant.

Our second claim is that, if $\mathcal{Q}$ accepts some $|\varphi\rangle$ with probability at least $2/3$, then $\mathcal{Q}_*$ accepts $|\varphi\rangle^{\otimes \ell}$ with probability at least $2/3$. For recall that a single invocation of the inner layer rejects $|\varphi\rangle^{\otimes \ell}$ with probability at most ε. So by Lemma 13, even if we invoke the inner layer u times in sequence, the probability that one or more invocations reject is at most $u\sqrt{\varepsilon}$, which is less than $1/3$ for suitable $u = O(W)$. ■

4.3. Main Result

We are now ready to prove Theorem 3: that for all Boolean functions f and all $w \geq 2$,

$$\mathrm{Q}^1(f) = O\left(\mathrm{QMA}_w^1(f) \cdot w \log^2 w\right).$$

Furthermore, if Bob uses C gates and S qubits in the QMA_w^1 protocol, then he uses $C \cdot S^{O(S)}$ gates and $O\left(S^2 \log^2 S\right)$ qubits in the Q^1 protocol.

Proof of Theorem 3. Let $\mathcal{Q}$ be Bob's algorithm. Also, suppose Alice's message has a qubits and Merlin's message has w qubits. The first step is to replace $\mathcal{Q}$ by the amplified algorithm $\mathcal{Q}_*$ from Lemma 15, which takes an A-qubit advice state $|\Psi\rangle$ from Alice and a W-qubit witness state from Merlin, where $A = O\left(aw\log^2 w\right)$ and $W = O\left(w \log w\right)$. From now on, we use $\mathcal{Q}_*\left(|\Phi\rangle\right)$ as a shorthand for $\mathcal{Q}_*$ run with witness $|\Phi\rangle$, together with an advice register that originally contains Alice's message $|\Psi\rangle$ (but that might become corrupted as Bob uses it). Then Bob's goal is to decide whether there exists a $|\Phi\rangle$ such that $\mathcal{Q}_*\left(|\Phi\rangle\right)$ accepts with high probability.

To do so, Bob uses the following procedure $\mathcal{M}$. Given Alice's message $|\Psi\rangle$, this procedure runs $\mathcal{Q}_*\left(|z\rangle\right)$ for $9\left(2^W\right)$ computational basis states $|z\rangle$ of the witness register chosen uniformly at random. Finally it returns the logical OR of the measurement outcomes.

let $|c\rangle$ be a counter initialized to $|0\rangle$
for $t := 1$ to $9\left(2^W\right)$
  choose $z \in \{0,1\}^W$ uniformly at random
  run $\mathcal{Q}_*\left(|z\rangle\right)$, and let b be $\mathcal{Q}_*$'s output
    // 1 *for accept,* 0 *for reject*
  set $|c\rangle := |c+b\rangle$
  run $\mathcal{Q}_*^{-1}\left(|z\rangle\right)$ to uncompute garbage
next t
if $c=0$ then return $f\left(x,y\right)=0$;
  otherwise return $f\left(x,y\right)=1$

Let us first show that $\mathcal{M}$ is correct. First suppose that $f\left(x,y\right)=0$. By Lemma 15, we know that $\mathcal{Q}_*\left(|\Phi\rangle\right)$ accepts with probability at most 5^{-W} for all states $|\Phi\rangle$ of the witness register. So in particular, $\mathcal{Q}_*\left(|z\rangle\right)$ accepts with probability at most 5^{-W} for all basis states $|z\rangle$. By Lemma 13, it follows that when $\mathcal{M}$ is finished, the counter c will have been incremented at least once (and hence $\mathcal{M}$ itself will have accepted) with probability at most

$$\frac{9\left(2^W\right)}{\sqrt{5^W}} \ll \frac{1}{9}.$$

Next suppose that $f\left(x,y\right)=1$. By assumption, there exists a $|\Phi\rangle$ such that $\mathcal{Q}_*\left(|\Phi\rangle\right)$ accepts with probability at least $2/3$. So setting $\eta = 2/3$, $N = 2^W$, and $T = 9\left(2^W\right)$, Lemma 14 implies that $\mathcal{M}$ will accept with probability at least

$$\left(\eta - \sqrt{\frac{N}{T}}\right)^2 = \left(\frac{2}{3} - \sqrt{\frac{1}{9}}\right)^2 = \frac{1}{9}.$$

It remains only to upper-bound $\mathcal{M}$'s complexity. If Bob's original algorithm $\mathcal{Q}$ used C gates and S qubits, then clearly the amplified algorithm $\mathcal{Q}_*$ uses $O\left(C \cdot w\log^2 w\right)$ gates and $O\left(S \cdot w \log^2 w\right)$ qubits. Hence $\mathcal{M}$ uses

$$O\left(C \cdot w\log^2 w \cdot 2^W\right) = C \cdot S^{O(S)}$$

gates and $O\left(S^2 \log^2 S\right)$ qubits, where we have used the fact that $w \leq S$. This completes the proof. ■

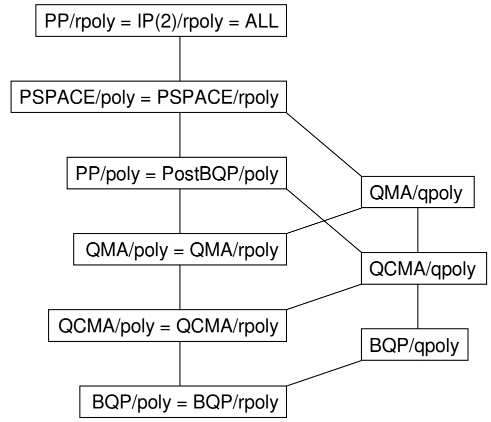

Figure 1. Known containments among classical and quantum advice classes.

5. Conclusions and Open Problems

Figure 1 shows the known relationships among deterministic, randomized, and quantum advice classes, in light of this paper's results. We still know remarkably little about quantum advice, compared to other computational resources. But our results provide new evidence for a general hypothesis: that if you're strong enough to squeeze an exponential amount of advice out of a quantum state, then you're also strong enough to squeeze an exponential amount of advice out of a probability distribution.

We end with some open problems.

- Can we find a counterexample to the quantum advice hypothesis? What about $\mathsf{QMA}\left(2\right)$, or $\mathsf{QMA}\left(k\right)$ for $k > 2$, or QS_2^p? Currently, we do not even know whether $\mathsf{QMA}\left(2\right)/\mathsf{rpoly} = \mathsf{ALL}$; this seems related to the difficult open question of amplification for $\mathsf{QMA}\left(2\right)$ (see Kobayashi et al. [8]).
- Is there a class $\mathcal{C}$ such that $\mathcal{C}/\mathsf{rpoly} \neq \mathcal{C}/\mathsf{poly}$ but $\mathcal{C}/\mathsf{rpoly} \neq \mathsf{ALL}$?
- Can we tighten the $\Omega\left(N/\log^2 N\right)$ lower bound of Theorem 6 to $\Omega\left(N\right)$? One approach would be to tighten Lemma 15, by generalizing the in-place QMA amplification of Marriott and Watrous [10].
- Can we improve the containment $\mathsf{QMA}/\mathsf{qpoly} \subseteq \mathsf{PSPACE}/\mathsf{poly}$ to $\mathsf{QMA}/\mathsf{qpoly} \subseteq \mathsf{PP}/\mathsf{poly}$? Alterna-

tively, can we construct an oracle (possibly a 'quantum oracle' [3]) relative to which $\mathsf{QMA/qpoly} \not\subset \mathsf{PostBQP/poly}$? This would indicate that the upper bound of $\mathsf{PSPACE/poly}$ might be difficult to improve.

6. Acknowledgments

Greg Kuperberg collaborated in the research project of which this paper was an offshoot, and I am grateful to him for comments and advice, as well as for several observations including Proposition 16. I also thank Oded Regev for first suggesting to me the problem of proving an upper bound on $\mathsf{QMA/qpoly}$; Ashwin Nayak and Hirotada Kobayashi for pointing out errors in an earlier version of Section 4.1; Harumichi Nishimura for helpful discussions; Jon Yard for pointing out a gap in an earlier version of Section 3.3; and Ronald de Wolf and the anonymous reviewers for comments on the manuscript.

References

[1] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. quant-ph/0402095.

[2] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. Roy. Soc. London*, A461(2063):3473–3482, 2005. quant-ph/0412187.

[3] S. Aaronson and G. Kuperberg. Quantum versus classical proofs and advice. In preparation, 2006.

[4] A. Ambainis, A. Nayak, A. Ta-Shma, and U. V. Vazirani. Quantum dense coding and quantum finite automata. *J. ACM*, 49:496–511, 2002. Earlier version in ACM STOC 1999, pp. 376-383. quant-ph/9804043.

[5] S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. In *Randomness and Computation*, volume 5 of *Advances in Computing Research*. JAI Press, 1989.

[6] A. Kitaev and J. Watrous. Parallelization, amplification, and exponential-time simulation of quantum interactive proof systems. In *Proc. ACM STOC*, pages 608–617, 2000.

[7] H. Klauck. Quantum communication complexity. In *Proc. Intl. Colloquium on Automata, Languages, and Programming (ICALP)*, pages 241–252, 2000. quant-ph/0005032.

[8] H. Kobayashi, K. Matsumoto, and T. Yamakami. Quantum Merlin-Arthur proof systems: are multiple Merlins more helpful to Arthur? In *ISAAC*, pages 189–198, 2003. quant-ph/0306051.

[9] R. E. Ladner. Polynomial space counting problems. *SIAM J. Comput.*, 18:1087–1097, 1989.

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

[11] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proc. IEEE FOCS*, pages 369–377, 1999. quant-ph/9904093.

[12] I. Newman. Private vs. common random bits in communication complexity. *Inform. Proc. Lett.*, 39:67–71, 1991.

[13] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[14] R. Raz. Quantum information and the PCP theorem. In *Proc. IEEE FOCS*, 2005. quant-ph/0504075.

[15] J. Watrous. Space-bounded quantum complexity. *J. Comput. Sys. Sci.*, 59(2):281–326, 1999.

7. Appendix: Other Complexity Results

The purpose of this appendix is to show that, in upper-bounding $\mathsf{QMA/qpoly}$, the computational difficulty really does arise from the need to handle quantum advice and quantum witnesses simultaneously: if either or both are "dequantized," then the upper bound of $\mathsf{PSPACE/poly}$ can be improved. In particular, and in increasing order of nontriviality, Theorem 17 will show that $\mathsf{MA/rpoly} = \mathsf{MA/poly}$ (and likewise that $\mathsf{QCMA/rpoly} = \mathsf{QCMA/poly}$), Theorem 19 will show that $\mathsf{QMA/rpoly} = \mathsf{QMA/poly}$, and Theorem 20 will show that $\mathsf{QCMA/qpoly} \subseteq \mathsf{PP/poly}$.

First, however, let us make a cautionary observation, which illustrates why such upper bounds cannot be blithely assumed. Recall that $\mathsf{MA_{EXP}}$ is the exponential-time analogue of MA.

Proposition 16 $\mathsf{MA_{EXP}/rpoly} = \mathsf{ALL}$.

Proof. Given an arbitrary Boolean function $f : \{0,1\}^n \to \{0,1\}$, an honest Merlin's message will consist of the truth table of f, while the randomized advice will consist of an $O(n)$-bit fingerprint of the truth table. ■

We can also "scale down" Proposition 16 by an exponential, to obtain $\mathsf{MA/poly} \subseteq \mathsf{MA/rlog}$. More explicitly, in the $\mathsf{MA/rlog}$ simulation, an honest Merlin's message will contain the advice s to the $\mathsf{MA/poly}$ machine, while the rlog advice will consist of an $O(\log n)$-bit fingerprint of s.

We next show that $\mathsf{MA/rpoly} = \mathsf{MA/poly}$. Combined with the above observation, this result has the surprising implication that

$$\mathsf{MA/rlog} = \mathsf{NP/poly} = \mathsf{MA/poly} = \mathsf{MA/rpoly}.$$

In other words, for an MA machine, $\mathrm{poly}(n)$ bits of randomized advice are no more powerful than $\log(n)$ bits.

Theorem 17 $\mathsf{MA/rpoly} = \mathsf{MA/poly}$.

Proof. Let L be a language in $\mathsf{MA/rpoly}$, and let $\mathcal{A}(x,r,z)$ be Arthur's verification algorithm run on input x, advice string r, and witness $z \in \{0,1\}^{w(n)}$, for some polynomial w. (We assume without loss of generality that Arthur is deterministic, since the randomized advice can provide his coins.) Also, let $\mathcal{D}$ be the distribution from which r is drawn. Then for all $x \in L$, there exists a z such that

$$\Pr_{r\in\mathcal{D}}[\mathcal{A}(x,r,z) \text{ accepts}] \geq \frac{2}{3},$$

whereas for all $x \notin L$ and all z,

$$\Pr_{r \in \mathcal{D}} [\mathcal{A}(x,r,z) \text{ accepts}] \leq \frac{1}{3}.$$

Let $R = (r_1, \ldots, r_{p(n)})$ be a $p(n)$-tuple of independent samples from $\mathcal{D}$, for some $p(n) = \Theta(n + w(n))$. Then there exists a boosted verifier $\mathcal{A}_*$ such that for all $x \in L$, there exists a z such that

$$\Pr_{R \in \mathcal{D}^{p(n)}} [\mathcal{A}_*(x,R,z) \text{ accepts}] \geq 1 - \frac{1}{2^n 2^{w(n)}},$$

whereas for all $x \notin L$ and all z,

$$\Pr_{R \in \mathcal{D}^{p(n)}} [\mathcal{A}_*(x,R,z) \text{ accepts}] \leq \frac{1}{2^n 2^{w(n)}}.$$

So by a simple counting argument, there exists a fixed advice string R such that for all $x \in L$, there exists a z such that Arthur accepts; whereas for all $x \notin L$ and all z, Arthur rejects. ■

Indeed, using the same techniques we can show that

$$\mathsf{QCMA/rlog} = \mathsf{QCMA/qlog} = \mathsf{QCMA/poly} = \mathsf{QCMA/rpoly}.$$

Next we want to show a somewhat harder result, that $\mathsf{QMA/rpoly} = \mathsf{QMA/poly}$. To do so we will need the following theorem of Marriott and Watrous.

Theorem 18 (Marriott and Watrous [10]) *The error probability in any* QMA *protocol can be made exponentially small without increasing the size of Merlin's quantum witness.*

We can now prove the analogue of Theorem 17 for QMA.

Theorem 19 $\mathsf{QMA/rpoly} = \mathsf{QMA/poly}$.

Proof. Given a language $L \in \mathsf{QMA/rpoly}$, let $\mathcal{D}$ be the distribution from which Arthur's advice is drawn, and let $\mathcal{Q}(x, r, |\varphi\rangle)$ be Arthur's verification algorithm run on input x, advice string r, and witness $|\varphi\rangle \in \mathcal{H}_2^{\otimes w(n)}$. Then for all $x \in L$, there exists a $|\varphi\rangle$ such that

$$\Pr_{r \in \mathcal{D}} [\mathcal{Q}(x,r,|\varphi\rangle) \text{ accepts}] \geq \frac{2}{3},$$

whereas for all $x \notin L$ and all $|\varphi\rangle$,

$$\Pr_{r \in \mathcal{D}} [\mathcal{Q}(x,r,|\varphi\rangle) \text{ accepts}] \leq \frac{1}{3}.$$

Here the probability is taken over $\mathcal{Q}$'s internal randomness as well as r.

By Theorem 18, we can make the error probability exponentially small without increasing the size of $|\psi\rangle$. So let $R = (r_1, \ldots, r_{p(n)})$ be a $p(n)$-tuple of independent samples from $\mathcal{D}$, for some $p(n) = \Theta(n + w(n))$. Then there exists a boosted verifier $\mathcal{Q}_*$ such that for all $x \in L$, there exists a $|\varphi\rangle$ such that

$$\Pr_{R \in \mathcal{D}^{p(n)}} [\mathcal{Q}_*(x,R,|\varphi\rangle) \text{ accepts}] \geq 1 - \frac{1}{2^n 2^{3w(n)}},$$

whereas for all $x \notin L$ and all $|\varphi\rangle$,

$$\Pr_{R \in \mathcal{D}^{p(n)}} [\mathcal{Q}_*(x,R,|\varphi\rangle) \text{ accepts}] \leq \frac{1}{2^n 2^{3w(n)}}.$$

So by a simple counting argument, there exists a fixed advice string R_1 such that for all $x \in L$, there exists a $|\varphi\rangle$ such that Arthur accepts with probability at least $1 - 2^{-3w(n)}$. However, we still need to handle the case $x \notin L$. Since the number of states $|\varphi\rangle \in \mathcal{H}_2^{\otimes w(n)}$ with small pairwise inner product is *doubly* exponential, a naïve counting argument no longer works. Instead, observe that there exists a fixed advice string R_0 such that for all $x \notin L$ and all computational basis states $|z\rangle$ with $z \in \{0,1\}^{w(n)}$,

$$\begin{aligned}\Pr[\mathcal{Q}_*(x,R_0,|z\rangle) \text{ accepts}] &\leq 2^n 2^{w(n)} \cdot \frac{1}{2^n 2^{3w(n)}} \\ &= \frac{1}{2^{2w(n)}}.\end{aligned}$$

Now suppose by contradiction that there exists a $|\varphi\rangle$ such that

$$\Pr[\mathcal{Q}_*(x,R_0,|\varphi\rangle) \text{ accepts}] > \frac{1}{3}.$$

Then

$$\Pr[\mathcal{Q}_*(x,R_0,I) \text{ accepts}] > \frac{1}{3} \cdot \frac{1}{2^{w(n)}},$$

where

$$I = \frac{1}{2^{w(n)}} \sum_{z \in \{0,1\}^{w(n)}} |z\rangle \langle z|$$

is the maximally mixed state on $w(n)$ qubits. But this implies that there exists a basis state $|z\rangle$ such that

$$\Pr[\mathcal{Q}_*(x,R_0,|z\rangle) \text{ accepts}] > \frac{1}{3} \cdot \frac{1}{2^{w(n)}},$$

which yields the desired contradiction. Finally, by a union bound, there exists a fixed advice string R that combines the properties of R_0 and R_1. ■

7.1. Upper-Bounding QCMA/qpoly

We now show that $\mathsf{QCMA/qpoly} \subseteq \mathsf{PP/poly}$. Conceptually, the proof is similar to the proof that $\mathsf{QMA/qpoly} \subseteq \mathsf{PSPACE/poly}$, but with three differences. First, since the witnesses are now classical, they can be provided to the simulating machine as part of the advice. Second, since the witnesses are provided, there is no longer any need to try exponentially many random witnesses. Indeed, this is what improves the upper bound from $\mathsf{PSPACE/poly}$ to $\mathsf{PP/poly}$. And third, we can no longer exploit the fact that $\mathsf{BQPSPACE/qpoly} = \mathsf{PSPACE/poly}$, in order to split the proof neatly into a "de-Merlinization" part (which is new) and an "advice" part (which follows from earlier work of Aaronson [1]). Instead, we need to generalize the machinery from [1] to the QCMA setting.

Theorem 20 $\mathsf{QCMA/qpoly} \subseteq \mathsf{PP/poly}$.

Proof. Let L be a language in $\mathsf{QCMA/qpoly}$, and let $L(x) = 1$ if $x \in L$ and $L(x) = 0$ otherwise. Also, let $\mathcal{Q}$ be a verifier for L, which takes a a-qubit quantum advice state $|\psi\rangle$ and w-bit classical witness z for some polynomials a and w (for convenience, we omit the dependence on n). Then the first step is to replace $\mathcal{Q}$ by an amplified verifier $\mathcal{Q}_*$, which takes an A-qubit advice state $|\Psi\rangle := |\psi\rangle^{\otimes \ell}$, where $A = a\ell$ and $\ell = O(\log a)$. As a result, $\mathcal{Q}_*$ has completeness and soundness errors $1/A^4$.

Let $\mathcal{Q}_*(x, \rho, z)$ be shorthand for $\mathcal{Q}_*$ run with input x, advice ρ, and witness z. Then given x, our goal is to simulate $\mathcal{Q}_*(x, |\Psi\rangle, z(x))$, where $z(x)$ is an optimal witness for x. We will do so using a $\mathsf{PP/poly}$ machine $\mathcal{M}$. The classical advice to $\mathcal{M}$ will consist of a "Darwinian training set" $(x_1, z_1), \ldots, (x_T, z_T)$ for $T = O(A)$, together with $L(x_t)$ for every $t \in \{1, \ldots, T\}$. Here each $x_t \in \{0,1\}^n$ is an input and each $z_t \in \{0,1\}^w$ is its corresponding witness. Given this advice, $\mathcal{M}$ runs the following procedure to compute $L(x)$.

let $\rho := I_A$ be the maximally mixed state on A qubits
for $t := 1$ to T
  let $|b\rangle$ be a qubit initialized to $|0\rangle$
  run $\mathcal{Q}_*(x_t, \rho, z_t)$, and CNOT the output into $|b\rangle$
  run $\mathcal{Q}_*^{-1}(x_t, \rho, z_t)$ to uncompute garbage
  measure $|b\rangle$, and postselect on observing $b = L(x_t)$
next t
for all $z \in \{0,1\}^w$,
  let λ_z be the probability that $\mathcal{Q}_*(x, \rho, z)$ accepts
if there exists a z such that $\lambda_z \geq 2/3$, then accept
otherwise, if $\lambda_z \leq 1/3$ for all z, then reject

Let us first see why $\mathcal{M}$ can be simulated in $\mathsf{PP/poly}$. The 'for' loop is just a postselected quantum computation, and can clearly be simulated by the result of Aaronson [2] that $\mathsf{PostBQP} = \mathsf{PP}$. The one nontrivial step is to decide whether there exists a z such that $\lambda_z \geq 2/3$, or whether $\lambda_z \leq 1/3$ for all z. We do this as follows. Let ρ_t be the state of the advice register after the first t postselection steps, conditioned on those steps succeeding. We first amplify by repeating the 'for' loop $J = O(w)$ times, using a different advice register each time. This yields J copies of ρ_T. We then replace $\mathcal{Q}_*(x, \rho_T, z)$ by the doubly-amplified verifier $\mathcal{Q}'_*\left(x, \rho_T^{\otimes J}, z\right)$, which runs $\mathcal{Q}_*(x, \rho_T, z)$ once for each of the J advice registers, and returns the majority outcome. Let λ'_z be the probability that $\mathcal{Q}'_*\left(x, \rho_T^{\otimes J}, z\right)$ accepts. Then by a Chernoff bound, and assuming the constant in $J = O(w)$ is sufficiently large, we have reduced the problem to deciding whether

(1) there exists a $z \in \{0,1\}^w$ such that $\lambda'_z \geq 1 - 2^{-2w}$, or

(2) $\lambda'_z \leq 2^{-2w}$ for all z.

Now let
$$S := \frac{1}{2^w} \sum_{z \in \{0,1\}^w} \lambda'_z.$$

Then $S \geq 2^{-w-1}$ in case (1), whereas $S \leq 2^{-2w}$ in case (2). So it suffices to give a $\mathsf{PP/poly}$ machine with $\alpha + \beta S$ accepting paths, for some positive constants α and β. Our machine will simply do the following:

- Choose z uniformly at random.
- Simulate a $\mathsf{PostBQP}$ computation that accepts with probability proportional to λ'_z.

The reason this works is that the probability of the T postselection steps in the 'for' loop all succeeding is independent of z.

It remains only to show $\mathcal{M}$'s correctness. Let p_t be the probability that the first t postselection steps in the 'for' loop all succeed. We choose the "training inputs" $x_1, \ldots, x_T$ and witnesses $z_1, \ldots, z_T$ in such a way that

(a) $p_{t+1} \leq \frac{2}{3} p_t$ for all $t \in \{0, \ldots, T-1\}$.

(b) z_t is a valid witness for x_t whenever $x_t \in L$, meaning that $\mathcal{Q}_*(x_t, |\Psi\rangle, z_t)$ accepts with probability at least $1 - 1/A^4$.

(c) There is no larger training set that satisfies (a) and (b).

Then it suffices to prove the following two claims:

(i) $T = O(A)$ for all training sets that satisfy (a) and (b).

(ii) $\mathcal{M}$ correctly decides every input x, if we train it on some $(x_1, z_1), \ldots, (x_T, z_T)$ that satisfies (a), (b), and (c).

For Claim (i), notice that we can write the maximally mixed state I as a mixture of 2^A orthonormal vectors

$$I = \frac{1}{2^A} \sum_{i=1}^{2^A} |\Psi_i\rangle \langle \Psi_i|,$$

where $|\Psi_1\rangle := |\Psi\rangle$ is the "true" advice state. We argue that the $|\Psi_1\rangle\langle\Psi_1|$ component must survive all T post-selection steps with high probability. For if $x_t \notin L$, then $\mathcal{Q}_*(x_t, |\Psi\rangle, z_t)$ accepts with probability at most $1/A^4$, while if $x_t \in L$, then $\mathcal{Q}_*(x_t, |\Psi\rangle, z_t)$ rejects with probability at most $1/A^4$ by assumption (b). So by Lemma 13, the probability of outputting the wrong answer on any of $(x_1, z_1), \ldots, (x_T, z_T)$, using $|\Psi\rangle$ as the advice, is at most $T\sqrt{1/A^4} = T/A^2$. Hence

$$p_T \geq \frac{1}{2^A}\left(1 - \frac{T}{A^2}\right).$$

On the other hand, $p_{t+1} \leq \frac{2}{3}p_t$ for all t by assumption (a), and hence $p_T \leq (2/3)^T$. Combining we obtain $T = O(A)$.

For Claim (ii), suppose by way of contradiction that $\mathcal{M}$ rejects some $x \in L$. Then $\mathcal{Q}_*(x, \rho_T, z)$ accepts with probability less than $2/3$ for all z. But this implies that if we trained $\mathcal{M}$ on the enlarged set $(x_1, z_1), \ldots, (x_T, z_T), (x, z)$ for any z, then we would get $p_{T+1} \leq \frac{2}{3}p_T$, thereby contradicting the maximality of T. Likewise, suppose $\mathcal{M}$ accepts some $x \notin L$. Then there exists a "false witness" $\widehat{z}$ such that $\mathcal{Q}_*(x, \rho_T, \widehat{z})$ accepts with probability greater than $1/3$. So if we trained $\mathcal{M}$ on the enlarged set $(x_1, z_1), \ldots, (x_T, z_T), (x, \widehat{z})$, we would again get $p_{T+1} \leq \frac{2}{3}p_T$, contradicting the maximality of T. ■