Quantum Complexity Theory
Sevag Gharibian
Summer 2019, University of Paderborn

Lecture 5: Quantum Merlin Arthur (QMA) and strong error reduction

"I have had my results for a long time, but I do not yet know how to arrive at them."
— Carl F. Gauss.

"If only I had the theorems! Then I should find the proofs easily enough."
— Georg B. Riemann.

Contents

Introduction. We have thus far defined BQP, studied the task of "solving" linear systems, and shown that matrix inversion is BQP-complete. We now wish to define a quantum analogue of NP. This is unfortunately a somewhat delicate issue; indeed, there are almost as many known quantum generalizations of NP as Snow White had dwarves — there's QMA, QMA$_1$, QMA(2), QCMA, StoqMA, and NQP. With this said, there is a *de facto* definition of "quantum NP" used by the community: Quantum Merlin Arthur (QMA).

In this lecture, we begin by defining Merlin Arthur (MA) and Quantum Merlin Arthur (QMA). We then study the surprising *strong error reduction* property of QMA. Finally, we close by discussing the relationship of QMA to known complexity classes. As suggested by the opening quotes of this lecture, a key theme will be the power of proofs (particularly quantum proofs); as with NP, these proofs will in general be hard to produce, but easy to verify.

1 Quantum Merlin Arthur (QMA)

Just as PromiseBPP was the correct class to generalize to BQP, to define QMA we begin with the promise-problem probabilistic generalization of NP, PromiseMA.

Definition 1 (PromiseMA). *A promise problem $\mathbb{A} = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ is in* PromiseMA *if there exists a (deterministic) TM M and fixed polynomials $p, s, r : \mathbb{N} \mapsto \mathbb{R}^+$, such that for any input $x \in \{0,1\}^n$, M takes in "proof" $y \in \{0,1\}^{p(n)}$ and string $z \in \{0,1\}^{s(n)}$, halts in at most $O(r(n))$ steps, and:*

- *(Completeness/YES case) If $x \in A_{\text{yes}}$, there exists a proof $y \in \{0,1\}^{p(n)}$, such that for at least $2/3$ of the choices of $z \in \{0,1\}^{s(n)}$, M accepts.*
- *(Soundness/NO case) If $x \in A_{\text{no}}$, then for all proofs $y \in \{0,1\}^{p(n)}$, at most $1/3$ of the choices of $z \in \{0,1\}^{s(n)}$ cause M to accept.*
- *(Invalid case) If $x \in A_{\text{inv}}$, then M may accept or reject arbitrarily.*

Exercise. How might we define Merlin-Arthur (MA), instead of PromiseMA (i.e. how does the definition above change if we drop the promise)?

Exercise. Show that the completeness and soundness parameters of $2/3$ and $1/3$, respectively, can be amplified without loss of generality to $1-2^{-n}$ and 2^{-n}, respectively. How many copies of the proof y suffice for this amplification?

The quantum analogue of PromiseMA which we focus on in this lecture is Quantum Merlin Arthur (QMA) (which, again, is really PromiseQMA, just as BQP is really PromiseBQP).

Definition 2 (Quantum Merlin Arthur (QMA))**.** *A promise problem $\mathbb{A} = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ is in QMA if there exists a P-uniform quantum circuit family $\{Q_n\}$ and polynomials $p, q : \mathbb{N} \mapsto \mathbb{N}$ satisfying the following properties. For any input $x \in \{0,1\}^n$, Q_n takes in $n + p(n) + q(n)$ qubits as input, consisting of the input x on register A, $p(n)$ qubits initialized to a "quantum proof" $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$ on register B, and $q(n)$ ancilla qubits initialized to $|0\rangle$ on register C. The first qubit of register C, denoted C_1, is the designated output qubit, a measurement of which in the standard basis after applying Q_n yields the following:*

- *(Completeness/YES case) If $x \in A_{\text{yes}}$, there exists proof $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, such that Q_n accepts with probability at least* $2/3$.
- *(Soundness/NO case) If $x \in A_{\text{no}}$, then for all proofs $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, Q_n accepts with probability at most* $1/3$.
- *(Invalid case) If $x \in A_{\text{inv}}$, Q_n may accept or reject arbitrarily.*

Exercise. If we replace the quantum proof $|\psi\rangle$ with a classical proof $y \in \{0,1\}^{p(n)}$ in the definition of QMA, do we recover PromiseMA?

A few comments on QMA are in order:

1. **Weak error reduction.** Similar to PromiseMA, parallel repetition suffices to amplify the QMA completeness and soundness parameters to $1-2^{-n}$ and 2^{-n}, respectively. However, the proof of this fact is not entirely trivial.

   Exercise. Suppose the QMA prover sends k copies of its proof, $|\psi\rangle$, instead of a single copy. On the jth copy of the proof, the verifier runs the verification circuit Q_n. Finally, the verifier measures the output qubits of all runs of Q_n, takes a majority vote of the resulting bits, and accepts if and only if the majority function yields 1. Prove that this procedure indeed amplifies the completeness and soundness parameters for QMA. (Hint: In the NO case, a cheating prover is *not* obligated to send k copies of some state $|\psi\rangle$ in tensor product, but rather can cheat by sending a large entangled state $|\phi\rangle \in (\mathbb{C}^2)^{\otimes k \cdot p(n)}$ across all k proof registers. Why does entanglement across proofs not help the prover in the NO case?)

   Observe we have denoted the use of parallel repetition above as *weak* error reduction. This is because the amplification step blows up the size of the proof register. Naively, one may expect this blowup to be necessary, since *a priori* it seems we cannot "reuse" the quantum proof $|\psi\rangle$ — indeed, the QMA verifier's measurement of its output qubit disturbs its quantum state, and the no-cloning theorem says the verifier cannot sidestep this by simply copying its input proof $|\psi\rangle$ before verifying it. Nevertheless, it turns out that amplification without a blowup in proof size *is* possible — this is called *strong* error reduction (Section 2), a simple and elegant application of which is to show that QMA $\subseteq$ PP (Section 3).

2. **Pure versus mixed proofs.** We have assumed the proof $|\psi\rangle$ in QMA to be a *pure* state, as opposed to a mixed state. Let us now reformulate the optimal acceptance probability of the quantum verifier Q_n as an eigenvalue problem; along the way, we will not only see that the "pure state proof" assumption is without loss of generality, but the reformulation we derive will prove crucial in our analysis of Section 2.

   Let Q_n be the circuit from Definition 2, acting on $n + p(n) + q(n)$ qubits. Recall that A, B, C denote the input, proof, and ancilla registers, respectively, and C_1 the designated output qubit of Q_n. Then, the probability that Q_n accepts proof $|\psi\rangle$ is

$$\begin{aligned}\Pr[\text{accept}] &= \left\| \, |1\rangle\langle 1|_{C_1} Q_n |x\rangle_A \otimes |\psi\rangle_B \otimes |0\cdots 0\rangle_C \, \right\|_2^2 \\ &= \langle x|_A \otimes \langle\psi|_B \otimes \langle 0\cdots 0|_C Q_n^\dagger |1\rangle\langle 1|_{C_1} Q_n |x\rangle_A \otimes |\psi\rangle_B \otimes |0\cdots 0\rangle_C \\ &= \mathrm{Tr}\left[\left(\langle x|_A \otimes I_B \otimes \langle 0\cdots 0|_C Q_n^\dagger |1\rangle\langle 1|_{C_1} Q_n |x\rangle_A \otimes I_B \otimes |0\cdots 0\rangle_C \right) |\psi\rangle\langle\psi|_B \right] \\ &= \mathrm{Tr}(P_x |\psi\rangle\langle\psi|), \end{aligned} \qquad (1)$$

   where the third statement follows by cyclicity of the trace, the fourth via identity $\mathrm{Tr}(M_{A,B} N_A \otimes I_B) = \mathrm{Tr}(\mathrm{Tr}_B(M_{A,B}) N_A)$, and the fifth by defining for convenience

$$P_x := \langle x|_A \otimes I_B \otimes \langle 0\cdots 0|_C Q_n^\dagger |1\rangle\langle 1|_{C_1} Q_n |x\rangle_A \otimes I_B \otimes |0\cdots 0\rangle_C.$$

   Henceforth, we shall abuse terminology by refering to P_x as the *POVM*[1] *for verifier* Q_n.

   Exercise. What space does P_x act on?

   Exercise. Prove $P_x \succeq 0$. (Hint: Prove first that if $A \succeq 0$, then $BAB^\dagger \succeq 0$ for any matrix $B \succeq 0$; the proof will be easier if you choose the "right" definition of positive semi-definiteness to work with. Second, recall the partial trace preserves positivity[2].)

   Now we are ready to address the question: *What happens if we consider a mixed proof ρ in place of a pure state $|\psi\rangle\langle\psi|$?*

   Exercise. Prove that for any density operator $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, there exists an i such that $\mathrm{Tr}(P_x\rho) \leq \mathrm{Tr}(P_x|\psi_i\rangle\langle\psi_i|)$. Conclude that, without loss of generality, quantum proofs in Definition 2 can be pure states.

   Finally, we stated earlier that the optimal acceptance probability can be reformulated as an eigenvalue problem. Indeed, the optimal acceptance probability over all proofs $|\psi\rangle$ is now

$$\max_{\text{unit vectors } |\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}} \langle\psi|P_x|\psi\rangle = \lambda_{\max}(P_x),$$

   attained by any eigenvector $|\psi\rangle$ of P_x with eigenvalue $\lambda_{\max}(P_x)$.

[1] Formally, POVM stands for Positive-Operator Valued Measure, and it normally denotes an alternate approach for modelling measurements. Specifically, a POVM P acting on n qubits is a set of operators $P = \{P_1, \ldots, P_k\}$ for some $k > 0$, such that $P_i \succeq 0$ and $\sum_i P_i = I$. As with projective measurements, each i denotes a distinct outcome of the measurement encoded by P, and the probability of outcome i is $\mathrm{Tr}(P_i\rho)$ when measuring state ρ. Unlike projective measurements, we do *not* require that $P_iP_j = 0$ for $i \neq j$; in this sense, POVMs generalize projective measurements. (Note that also unlike projective measurements, the postmeasurement state upon obtaining outcome i is no longer specified by $P_i\rho P_i/\mathrm{Tr}(\rho P_i)$.) In the context of a QMA verifier Q_n, we may view the application of Q_n and subsequent measurement of the output qubit of Q_n in the standard basis as a POVM consisting of two elements: $P = \{I - P_x, P_x\}$ (since the measurement has only two outputs, $|0\rangle$ or $|1\rangle$, respectively). Since this two-outcome POVM P is fully specified by P_x, for simplicity we choose to abuse terminology and refer to P by P_x.

[2] The partial trace is an example of a *completely positive* map from operators to operators; the fancy name for maps from operators to operators is a *superoperator*. We have not yet developed the formal framework for discussing such "superoperators", but for now it suffices to note that the partial trace correctly maps density operators to density operators.

Exercise. Prove the equality above using the Courant-Fischer variational characterization of eigenvalues. The latter states: Let $A \in \text{Herm}(\mathbb{C}^N)$ have eigenvalues $\lambda_1 \leq \cdots \leq \lambda_N$. Then,

$$\lambda_k = \min_{\text{subspaces } S \subseteq \mathbb{C}^N \text{ of dimension } k} \quad \max_{\text{unit vectors } |\psi\rangle \in S} \quad \langle\psi|A|\psi\rangle.$$

2 Strong error reduction for QMA

In the setting of PromiseMA, a previous exercise essentially asked you to show that at given a single copy of proof y, the completeness and soundness parameters of the PromiseMA verifier could be amplified to exponentially close to 1 and 0, respectively. Quantumly, one might naively expect an analogous statement to be false — a quantum verifier measures and hence disturbs its state, so how can it "reuse" its proof, $|\psi\rangle$? A simple solution would be for the quantum verifier to create multiple copies of $|\psi\rangle$ before beginning its verification; unfortunately, the quantum no-cloning theorem rules this out. The following theorem hence comes as a surprise.

Theorem 3 (Strong error reduction for QMA)**.** *Let Q_n be a QMA verifier for promise problem $A = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$, where we assume the terminology of Definition 2. Then, for any polynomial $r : \mathbb{N} \mapsto \mathbb{N}$, there exists a polynomial-time deterministic TM mapping Q_n to a (polynomial-size) quantum circuit R_n with the following properties for any input $x \in \{0,1\}^n$:*

- *(Completeness/YES case) If $x \in A_{\text{yes}}$, there exists proof $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, such that R_n accepts with probability at least $1 - 2^{-r(n)}$.*
- *(Soundness/NO case) If $x \in A_{\text{no}}$, then for all proofs $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, R_n accepts with probability at most $2^{-r(n)}$.*
- *(Invalid case) If $x \in A_{\text{inv}}$, R_n may accept or reject arbitrarily.*

It is crucial to note that both Q_n and R_n take in the same number of proof qubits, $p(n)$.

Remark. As with weak error reduction, Theorem 3 holds even if the completeness parameter $c(n)$ and soundness parameter $s(n)$ for Q_n satisfy $c(n) - s(n) \geq t(n)$ for some fixed polynomial $t : \mathbb{N} \mapsto \mathbb{N}$.

2.1 Intuition: A spinning top

To most easily see the intuition behind the proof of Theorem 3, assume Q_n has completeness 1, i.e. in the YES case, there exists a proof $|\psi\rangle$ accepted by Q_n with certainty. Assume first that $x \in A_{\text{yes}}$. There are two high-level steps to the amplification procedure:

1. *Run the verification.* Apply Q_n to obtain

$$|\phi\rangle = Q_n|x\rangle_A|\psi\rangle_B|0\cdots 0\rangle_C \in (\mathbb{C}^2)^{\otimes n+p(n)+q(n)}. \tag{2}$$

   Since $x \in A_{\text{yes}}$ and Q_n has perfect completeness, we know

$$|\phi\rangle = |1\rangle_{C_1}|\phi'\rangle \text{ for some unit vector } |\phi'\rangle \in (\mathbb{C}^2)^{\otimes n+p(n)+q(n)-1}.$$

   Thus, if we measure the output qubit, C_1, in the standard basis via projectors $\Pi^{\text{accept}} = |1\rangle\langle 1|, \Pi^{\text{reject}} = |0\rangle\langle 0| \in \mathcal{L}(\mathbb{C}^2)$, we not only obtain outcome Π^{accept} with certainty, but the postmeasurement state is $\Pi^{\text{accept}}_{C_1}|\phi\rangle = |\phi\rangle$. In other words, the measurement *does not disturb* the output of Q_n.

2. *Run the verification in reverse.* Since measuring $|\phi\rangle$ did not disturb it, applying Q_n in reverse now trivially reverts us to our initial state:

$$Q_n^\dagger\left(\Pi^{\text{accept}}_{C_1}|\phi\rangle\right) = Q_n^\dagger|\phi\rangle = Q_n^\dagger\left(Q_n|x\rangle_A|\psi\rangle_B|0\cdots 0\rangle_C\right) = |x\rangle_A|\psi\rangle_B|0\cdots 0\rangle_C.$$

If we now measure projectors[3] $\Pi^{\text{reset}} = |x\rangle\langle x|_A \otimes |0\cdots0\rangle\langle 0\cdots 0|_{A,C}, \Pi^{\text{new}} = I - \Pi^{\text{reset}} \in \mathcal{L}((\mathbb{C}^2)^{\otimes n+q(n)})$, we again leave the state invariant, i.e. with probability 1 we obtain outcome Π^{reset}, and the postmeasurement state satisfies

$$\Pi^{\text{reset}}|x\rangle_A|\psi\rangle_B|0\cdots0\rangle_C = |x\rangle_A|\psi\rangle_B|0\cdots0\rangle_C.$$

Note that we may repeat this procedure as many times as we like, and the outcomes will always be the same. This is akin to a perfectly spinning top — if we think of the each application of the amplification procedure as giving the top a supplemental twirl, the top will continue to spin blissfully along in a "steady state".

The interesting part is now the NO case. Here, in Step 1 of the amplification procedure above, since proof $|\psi\rangle$ is accepted with probability at most 1/3, we know $|\phi\rangle$ has form

$$|\phi\rangle = \alpha_0|0\rangle_{C_1}|(\phi')^{\perp}\rangle + \alpha_1|1\rangle_{C_1}|\phi'\rangle$$

for some orthonormal unit vectors $|\phi'\rangle, |(\phi')^{\perp}\rangle \in (\mathbb{C}^2)^{\otimes n+p(n)+q(n)-1}$, $|\alpha_0|^2 + |\alpha_1|^2 = 1$, and $|\alpha_1|^2 \leq 1/3$. The last of these properties guarantees that if we are lucky enough to measure $|1\rangle$ in C_1, the postmeasurement collapse will disturb $|\phi\rangle$ greatly (since most of the weight of $|\phi\rangle$ is on $|0\rangle_{C_1}$). This, in turn, suggests that when we now run Step 2 by inverting Q_n and measuring $\{\Pi^{\text{reset}}, \Pi^{\text{new}}\}$, we will obtain outcome Π^{new} with non-trivial probability, and again disturb our state greatly. And applying these two steps repeatedly will presumably amplify the disturbances further. This is analogous to saying that if we start with a top spinning with a slight wobble, each twirl we perform will further amplify the wobble until the top spins out of control.

2.2 Proof of strong error reduction

While Section 2.1 gave intuition as to why the amplification procedure might work, a formal analysis reveals the "motion of our top" can be tracked in a very elegant and precise fashion, even if we drop the assumption of perfect completeness.

Proof of Theorem 3. We begin by following Section 2.1. Let Q_n be a QMA verifier and $x \in \{0,1\}^n$ an input string. For brevity, we henceforth simply write Q for Q_n. To ease the analysis, we also rename our projectors

$$S_0 := \Pi^{\text{new}}, \qquad S_1 := \Pi^{\text{reset}}, \qquad E_0 := \Pi^{\text{reject}}, \qquad E_1 := \Pi^{\text{accept}}, \tag{3}$$

where S in S_i stands for "start" (since this measurement is on the start state) and E_i stands for "end" (since this measurement is on the end state).

The new verification procedure. The new circuit R_n (henceforth R for brevity) acts as follows.

1. Set $i = t = 0$.
2. While $i \leq N$:
   (a) (Run the verification) Apply Q and measure output qubit C_1 with respect to $\{E_0, E_1\}$. If the outcome is E_j, set $y_i = j \in \{0,1\}$, and increment i.
   (b) (Run the verification in reverse) Apply $Q^\dagger$ and measure input and ancilla registers A and C with respect to $\{S_0, S_1\}$. If the outcome is S_j, set $y_i = j \in \{0,1\}$, and increment i.
3. (Postprocessing) If the number of indices $i \in \{0, \dots, N-1\}$ such that $y_i = y_{i+1}$ is at least $N/2$, accept. Otherwise, reject.

It suffices to set $N = 8r(n)/9$. Note the mapping from Q to V takes time polynomial in n.

Exercise. How many times does the while loop above run with respect to N?

[3] For clarity, the superscript for Π^{reset} means the A and C registers are *reset* to their original states $|x\rangle$ and $|0\cdots0\rangle$, respectively, and for Π^{new} means the registers are set to a "new" start state.

Correctness. If we are in a YES case with perfect completeness, it is clear in Step 3 above that $y_i = y_{i+1}$ for all $i \in \{0, \ldots, N-1\}$. The aim is thus to show a similar statement for general YES (resp. NO) cases; that we are *more likely* to maintain (resp. flip) the value of y_i in setting y_{i+1} in the YES (resp. NO) case, thus leading to the correct answer with high probability in Step 3.

The starting point for the formal analysis is Equation (1), which said the probability that Q accepts proof $|\psi\rangle$ is $\mathrm{Tr}(P_x|\psi\rangle\langle\psi|)$, for positive-semidefinite operator

$$P_x := \langle x|_A \otimes I_B \otimes \langle 0\cdots 0|_C Q^\dagger E_1 Q |x\rangle_A \otimes I_B \otimes |0\cdots 0\rangle_C.$$

It turns out that if we restrict our attention to eigenvectors $|\psi\rangle$ of P_x, we obtain a clean closed form solution.

Closed form solution when $|\psi\rangle$ is an eigenvector of P_x. In the case when $|\psi\rangle$ is an eigenvector of P_x, we can *exactly* write down the acceptance probability of $|\psi\rangle$ by R, and this will turn out to suffice for the entire correctness analysis.

Lemma 4. *Suppose $|\psi\rangle$ is an eigenvector of P_x accepted by Q with probability p. Then, for any $i \in \{0, \ldots, N-1\}$, $\Pr[y_i = y_{i+1}] = p$. (Thus, $\Pr[y_i \neq y_{i+1}] = 1 - p$.)*

The magic of Lemma 4 is that even though *a priori* the action of R on $|\psi\rangle$ seems difficult to predict, when $|\psi\rangle$ is an eigenvector of P_x, each step 2(a) and 2(b) of R is just a *Bernoulli trial*[4]: Independently of all previous measurement outcomes, with probability p we don't flip our bit, and with probability $1-p$ we do flip our bit. This means we can later apply powerful tail bounds like the Chernoff bound to analyze the acceptance probability of R on eigenvectors of $|\psi\rangle$.

Proof of Lemma 4. Assume $P_x|\psi\rangle = p|\psi\rangle$ for $0 < p < 1$.

Exercise. Prove the claim in the setting $p = 0$ and $p = 1$.

Recall from Equation (3) that S_1 and E_1 denote successful projections at the start and end of verification (i.e. onto the original input x and all-zeroes ancilla, and onto the accepting output qubit, respectively), and $S_0 = I - S_1$ and $E_0 = I - E_1$. Define for brevity

$$|\phi\rangle := |x\rangle_A|\psi\rangle_B|0\cdots 0\rangle_C \qquad \text{and} \qquad \Gamma := S_1 Q^\dagger E_1 Q S_1.$$

A key identity is now the following.

Exercise. Prove that

$$\Gamma|\phi\rangle = S_1 Q^\dagger E_1 Q S_1 |\phi\rangle = p|\phi\rangle. \tag{4}$$

Why must we include projectors S_1 in the definition of Γ to make this a well-defined equality?

To show the claim, we trace through the first iteration of the while loop of R.

- The first run of Step 2(a) applies Q to $|\psi\rangle$. Since $E_0 + E_1 = I$, this step hence performs mapping

  $$|\phi\rangle \to Q|\phi\rangle = E_0 Q|\phi\rangle + E_1 Q|\phi\rangle.$$

  If we now measure $\{E_0, E_1\}$, we collapse to state

  $$|e_1\rangle := \frac{E_1 Q|\phi\rangle}{\parallel E_1 Q|\phi\rangle \parallel_2} \text{ with probability } \parallel E_1 Q|\phi\rangle \parallel_2^2 = \langle\phi|Q^\dagger E_1 Q|\phi\rangle = p.$$

[4] Recall from probability theory that *Bernoulli trials* refer to independently repeating a two-outcome sampling experiment.

Exercise. Why is $\langle\phi|Q^\dagger E_1 Q|\phi\rangle = p$? (Hint: What is $S_1|\phi\rangle$?)

Using the identity $E_0 = I - E_1$, we analogously collapse to

$$|e_0\rangle := \frac{E_0 Q|\phi\rangle}{\parallel E_0 Q|\phi\rangle \parallel_2} \text{ with probability } \parallel E_0 Q|\phi\rangle \parallel_2^2 = \langle\phi|Q^\dagger E_0 Q|\phi\rangle = 1 - \langle\phi|Q^\dagger E_1 Q|\phi\rangle = 1-p.$$

Note that together, these statements imply

$$Q|\phi\rangle = \sqrt{1-p}|e_0\rangle + \sqrt{p}|e_1\rangle. \tag{5}$$

- After Step 2(a), we have either $|e_0\rangle$ or $|e_1\rangle$. Step 2(b) now applies $Q^\dagger$, yielding one of two possible transitions:

$$\begin{aligned} Q^\dagger|e_1\rangle &= S_0 Q^\dagger|e_1\rangle + S_1 Q^\dagger|e_1\rangle && (6)\\ Q^\dagger|e_0\rangle &= S_0 Q^\dagger|e_0\rangle + S_1 Q^\dagger|e_0\rangle && (7)\end{aligned}$$

  We may simplify each term on the right hand side as:

$$\begin{aligned} S_1 Q^\dagger|e_1\rangle &= S_1 Q^\dagger \frac{E_1 Q|\phi\rangle}{\sqrt{p}} = \frac{1}{\sqrt{p}}\Gamma|\phi\rangle = \sqrt{p}|\phi\rangle && (8)\\ S_0 Q^\dagger|e_1\rangle &= S_0 Q^\dagger \frac{E_1 Q|\phi\rangle}{\sqrt{p}} = \frac{1}{\sqrt{p}} S_0 Q^\dagger E_1 Q|\phi\rangle && (9)\\ S_1 Q^\dagger|e_0\rangle &= S_1 Q^\dagger \frac{E_0 Q|\phi\rangle}{\sqrt{1-p}} = \frac{1}{\sqrt{1-p}}\left(|\phi\rangle - \Gamma|\phi\rangle\right) = \sqrt{1-p}|\phi\rangle && (10)\\ S_0 Q^\dagger|e_0\rangle &= S_0 Q^\dagger \frac{E_0 Q|\phi\rangle}{\sqrt{1-p}} = \frac{1}{\sqrt{1-p}}\left(S_0|\phi\rangle - S_0 Q^\dagger E_1 Q|\phi\rangle\right) = -\frac{1}{\sqrt{1-p}} S_0 Q^\dagger E_1 Q|\phi\rangle. && (11)\end{aligned}$$

Exercise. Prove each of the four statements above. (Hint: Use the fact that $E_0 + E_1 = I$. Also, why is $S_0|\phi\rangle = 0$?)

Exercise. What is $\left\| S_0 Q^\dagger E_1 Q|\phi\rangle \right\|_2$?

Exercise. Prove that after we measure the right hand side of Equation (6) with $\{S_0, S_1\}$, we obtain S_1 with probability p and S_0 with probability $1-p$. Similarly, measuring the right hand side of Equation (7) with S_0, S_1 yields S_0 with probability p and S_1 with probability $1-p$.

Exercise. Conclude from the last exercise that after the first iteration of the while loop, $y_1 = y_2$ with probability p. Accordingly, $y_1 \neq y_2$ with probability $1-p$.

This is precisely the behavior we are seeking. In sum, the analysis of Step 2(b) yields the following.

Exercise. Define $|s_0\rangle := \frac{S_0 Q^\dagger E_1 Q|\phi\rangle}{\parallel S_0 Q^\dagger E_1 Q|\phi\rangle \parallel_2}$ and $|s_1\rangle := |\phi\rangle$. Prove the following:

$$\begin{aligned} Q^\dagger|e_0\rangle &= -\sqrt{p}|s_0\rangle + \sqrt{1-p}|s_1\rangle\\ Q^\dagger|e_1\rangle &= \sqrt{1-p}|s_0\rangle + \sqrt{p}|s_1\rangle. \end{aligned}$$

Exercise. Prove that $Q|s_0\rangle = -\sqrt{p}|e_0\rangle + \sqrt{1-p}|e_1\rangle \in S_v$. (Hint: Use Equation (4).)

It will now be fruitful to step back and see the bigger picture which is emerging. Define $S_v := \text{Span}(|e_0\rangle, |e_1\rangle)$ and $S_w := \text{Span}(|s_0\rangle, |s_1\rangle)$. Then, our analysis above showed that Q maps S_v into S_w. Conversely, $Q^\dagger$ maps S_w back into S_v. In fact, $S_v = S_w$, as you will now show.

Exercise. Prove that $\{|e_0\rangle, |e_1\rangle\}$ is an orthonormal set.

Exercise. Prove that $\{|s_0\rangle, |s_1\rangle\}$ is an orthonormal set.

Exercise. Prove that $S_v = S_w$.

In other words, the evolution of R is *entirely confined* in a two-dimensional space: $S := S_v = S_w$. In space S, our analysis reveals the entire action of Q:

$$\begin{aligned}
Q|s_0\rangle &= -\sqrt{p}|e_0\rangle + \sqrt{1-p}|e_1\rangle \\
Q|s_1\rangle &= \sqrt{1-p}|e_0\rangle + \sqrt{p}|e_1\rangle \\
Q^\dagger|e_0\rangle &= -\sqrt{p}|s_0\rangle + \sqrt{1-p}|s_1\rangle \\
Q^\dagger|e_1\rangle &= \sqrt{1-p}|s_0\rangle + \sqrt{p}|s_1\rangle.
\end{aligned}$$

We are now ready to finish the proof of Lemma 4.

Exercise. Observe that $S_i|s_i\rangle \propto |s_i\rangle$, and $S_i|s_{i\oplus 1}\rangle = 0$. Similarly, $E_i|e_i\rangle \propto |e_i\rangle$, and $E_i|e_{i\oplus 1}\rangle = 0$. Why can we now conclude the analysis for a single loop iteration suffices to prove all of Lemma 4? □

With Lemma 4 in hand, we have a clean characterization of how R behaves on any proof $|\psi\rangle$ which is an eigenvector of P_x. We are now ready to complete the proof of Theorem 3.

Correctness proof for YES case. In the YES case, from Section 1 we know that the optimal proof for Q is an eigenvector $|\psi\rangle$ of P_x accepted with probability $p \geq 2/3$. Thus, by Lemma 4, for each $i \in \{0, \ldots, N-1\}$, $y_i = y_{i+1}$ with probability at least $2/3$. By the Chernoff bound (which we may apply since Lemma 4 reduces us to Bernoulli trials), the claim now follows.

Correctness proof for NO case. In the NO case, the optimal proof for Q is an eigenvector $|\psi\rangle$ of P_x accepted with probability $p \leq 1/3$. Unfortunately, here we cannot proceed as in the YES case by assuming a cheating prover sends an eigenvector of P_x as a proof. Luckily, it turns out that by applying a similar, but slightly more general, analysis to that above, one can explicitly show the desired bound for the NO case as well. We omit this additional analysis.

□

3 Relationship to other classes

3.1 The many cousins of QMA

There are a number of variants of QMA, the most prominent of which are arguably the following (in no particular order).

- **One-Sided Error Quantum Merlin Arthur (QMA$_1$).** QMA with perfect completeness, i.e. in the YES case there exists a proof accepted with probability 1.

- **Quantum Classical Merlin Arthur (QCMA).** QMA, but with a classical proof $y \in \{0,1\}^{p(n)}$.
- **Stoquastic Merlin Arthur (StoqMA).** QMA, except (1) the ancilla qubits are allowed to be initialized (independently) to either $|0\rangle$ or $|+\rangle$, (2) the verification circuit consists solely of reversible classical gates, and (3) the final single-qubit measurement is in the X basis (i.e. $\{|+\rangle, |-\rangle\}$.
- **Quantum Merlin Arthur with Two Proofs (QMA(2))** QMA, except the proof is promised to be a tensor product of two proofs, i.e. $|\psi\rangle = |\phi_1\rangle \otimes |\phi_2\rangle$ for some $|\phi_1\rangle, |\phi_2\rangle$.

Again, despite the nomenclature, each of these is a class of promise problems, not a class of languages. Here is what is known about the relationships between these:

- $\text{BQP} \subseteq \text{QCMA} \subseteq \text{QMA}_1 \subseteq \text{QMA} \subseteq \text{QMA}(2) \subseteq \text{NEXP}$.

  Exercise. Which of these inclusions are trivial? Why?

  Exercise. Why is it not clear that QMA = QMA(2)?

  Two remarks: (1) The inclusion $\text{QCMA} \subseteq \text{QMA}_1$ follows because one can show $\text{QCMA} = \text{QCMA}_1$, i.e. without loss of generality we may assume QCMA has perfect completeness. The analogous statement is *not* known for QMA. (2) The containment $\text{QMA}(2) \subseteq \text{NEXP}$ is, remarkably and sadly, the best trivial upper bound on QMA(2). This leaves quite a chasm between QMA and QMA(2), with the former contained in[5] PP. The only known *non-trivial*[6] upper bound on QMA(2) is

  $$\text{QMA}(2) \subseteq \text{Q}\Sigma_3 \subseteq \text{NEXP},$$

  where $\text{Q}\Sigma_3$ is a quantum analogue[7] of Σ_3^{p}, the third level of PH. It is not yet clear if this should be construed as strong evidence that $\text{QMA}(2) \neq \text{NEXP}$; not much is known about $\text{Q}\Sigma_3$, and it is entirely possible that $\text{QMA}(2) = \text{Q}\Sigma_3 = \text{NEXP}$. On the other hand, if the study of the classical analogue of $\text{Q}\Sigma_3$ is any guide, it would suggest $\text{QMA}(2) \neq \text{Q}\Sigma_3$ (and hence $\text{QMA}(2) \neq \text{NEXP}$), as classically alternating quantifiers are strongly believed to add power to a proof system (otherwise, PH collapses).

  Exercise. Why would QMA(2) = NEXP imply that alternating quantifiers do not strictly increase the power of a QMA(2) proof system?

- As for StoqMA, it is a rather strange fellow — in terms of lower bounds, it is the only "quantum" cousin of QMA which is *not* believed to contain BQP. Indeed, $\text{StoqMA} \subseteq \text{PH}$ (more precisely, it is in Arthur-Merlin (AM)), whereas it is believed BQP is not contained in PH. In terms of upper bounds, it is not clear whether $\text{StoqMA} \subseteq \text{QCMA}$; this because the former has a quantum proof but "classical" verification, whereas the latter has a classical proof but quantum verification. In fact, it is not even known whether weak error reduction holds for StoqMA, since the final X-basis measurement appears to prevent the standard "parallel repetition plus majority-vote" technique.

Upper bounds on QMA. The most "mainstream" upper bound on QMA is $\text{QMA} \subseteq \text{PP}$. However, there are two strictly stronger known bounds (assuming standard complexity theoretic conjectures):

[5] Recall $\text{PP} \subseteq \text{PSPACE} \subseteq \text{EXP} \subseteq \text{NEXP}$.

[6] That $\text{QMA}(2) \subseteq \text{Q}\Sigma_3$ is trivial; it is the containment $\text{Q}\Sigma_3 \subseteq \text{NEXP}$ which is non-trivial.

[7] Roughly, in a YES instance for $\text{Q}\Sigma_3$, there is a proof ρ_1, such that for all proofs ρ_2, there exists a ρ_3 leading the quantum verifier to accept (ρ_1, ρ_2, ρ_3) with probability at least $2/3$. Analogously for a NO instance, for all proofs ρ_1, there is a proof ρ_2, such that for all ρ_3 the quantum verifier accepts with probability at most $1/3$. All proofs are polynomial-size and allowed to be mixed. In strong contrast to QMA, it is *not* clear whether the proofs can be assumed pure without loss of generality.

1. QMA $\subseteq$ A$_0$PP $\subseteq$ PP. Rather than defining A$_0$PP, we will define the *quantum* class SBQP = A$_0$PP. SBQP is the class of decision problems for which there exists a quantum polynomial time algorithm which, on input $x \in \{0,1\}^*$, accepts in the YES case with probability at least $2 \cdot 2^{-p(|x|)}$, and accepts in the NO case with probability at most $2^{-p(|x|)}$, for some polynomial p. Note that it is strongly believed that SBQP = A$_0$PP $\neq$ PP, since equality would collapse PH.

2. QMA $\subseteq$ P$^{\text{QMA[log]}}$ $\subseteq$ PP. Here, P$^{\text{QMA[log]}}$ is the set of decision problems solved by a P machine which can make at most $O(\log n)$ (adaptive) queries to a QMA oracle. Again, this strictly separates QMA from PP, in the sense that it is unlikely that QMA = P$^{\text{QMA[log]}}$. This is because the latter contains both QMA and co-QMA (the complement of QMA), and so QMA = P$^{\text{QMA[log]}}$ would have the unlikely implication that QMA $\supseteq$ co-QMA.

   Exercise. Why does co-QMA $\subseteq$ P$^{\text{QMA[log]}}$ hold?

While these two upper bounds on QMA are likely stronger than PP, we now close the lecture by showing the weaker bound QMA $\subseteq$ PP; this latter containment follows via a simple application of strong error reduction.

3.2 Using strong error reduction to show QMA $\subseteq$ PP

Theorem 5. QMA $\subseteq$ PP.

Proof idea. The proof idea is most easily grasped by using it to show NP $\subseteq$ PP. For this, suppose we have a 3-SAT input formula $\phi : \{0,1\}^n \mapsto \{0,1\}$. To put NP in PP, our goal is to show that there exists a probabilistic polynomial-time algorithm A which, given ϕ, accepts with probability strictly larger than $1/2$ if ϕ is satisfiable, and accepts with probability at most $1/2$ otherwise. The approach for doing so is simple — if and only if ϕ is satisfiable, it has a satisfying assignment x; so, A randomly picks an assignment $y \in \{0,1\}^n$, and outputs $\phi(y)$.

Exercise. Prove that if ϕ is satisfiable, A accepts with probability at least $1/2^n$. On the other hand, if ϕ is unsatisfiable, A accepts with probability 0. Why is this enough to imply 3-SAT $\in$ PP?

Proof of Theorem 5. We shall show QMA $\subseteq$ PQP, for PQP defined essentially identically to PP except with a P-uniform quantum circuit family in place of a Turing machine. It is known that PQP = PP, whose proof we omit here.

Let $\mathbb{A} = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ be a QMA promise problem, and $x \in \{0,1\}^n$ an input. The overall proof idea is the same as in the classical case — the PQP machine A simply "guesses" a quantum proof $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, feeds it into verifier Q_n, and outputs Q_n's answer. Formally, to model a "random proof" $|\psi\rangle$, A instead feeds Q_n the maximally mixed state $I/2^{p(n)} \in \mathcal{L}((\mathbb{C}^2)^{\otimes p(n)})$.

Exercise. Why does $\frac{1}{2^{p(n)}}I$ correctly model a random pure state $|\psi\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$?

Recall we may now write the acceptance probability of Q_n on proof $I/2^{p(n)}$ as (for POVM P_x defined as in Equation (1))

$$\Pr[\text{accept}] = \operatorname{Tr}\left(P_x \cdot \frac{I}{2^{p(n)}}\right) = \frac{1}{2^{p(n)}}\operatorname{Tr}(P_x).$$

Exercise. Recall that $P_x \succeq 0$, and that $\operatorname{Tr}(P_x)$ is the sum of all eigenvalues of P_x. Why does the Pr[accept] above not suffice to separate YES from NO cases of $\mathbb{A}$?

As the exercise above shows, this naive idea alone does not work. Rather, we must first use strong error reduction to amplify the completeness and soundness parameters of Q_n. Specifically, recall that Q_n takes in $p(n)$ proof qubits, for some polynomial p. For any polynomial r, Theorem 3 says we may map Q_n to a new circuit R_n which still takes in $p(n)$ proof qubits, but has completeness and soundness parameters $1 - 2^{-r(n)}$ and $2^{-r(n)}$, respectively. This now suffices to complete the proof.

Exercise. Prove that for sufficiently large fixed r, feeding $I/2^n$ into R_n and outputting its answer suffices to decide in PQP whether $x \in A_{\text{yes}}$ or $x \in A_{\text{no}}$. More formally, let P_x^R denote the POVM for verifier R_n (c.f. Equation (1)). Prove that:

- If $x \in A_{\text{yes}}$, then $\frac{1}{2^{p(n)}}\text{Tr}(P_x^R) \geq \frac{1}{2^{p(n)}} - \frac{1}{2^{p(n)+r(n)}}$.
- If $x \in A_{\text{no}}$, then $\frac{1}{2^{p(n)}}\text{Tr}(P_x^R) \leq \frac{1}{2^{r(n)}}$.

What choice of r hence suffices to distinguish YES from NO cases in PQP? (Hint: You do not need to use the precise structure of P_x^R; the relationship between the optimal probability of acceptance of R_n and the eigenvalues of P_x^R suffices.)

Exercise. Would the approach above work if we used weak error reduction instead of strong error reduction? Why or why not? □