Frontiers of Quantum complexity *Instructors: Anand Natarajan, Anurag Anshu*

Lecture 7: Friends of QMA

Sept 25, 2025 *Scribe: Anastasiia Struss*

1 Reminder on MA and friends

We start by recalling definitions of some classical complexity classes.

Definition 1.1. Language L is in NP if there exists a deterministic polynomial-time verifier V and a polynomial p such that for every input string x of length n

Completeness: $x \in L \Rightarrow \exists \pi \in \{0,1\}^{p(n)} : V(x,\pi) = 1$,

Soundness: $x \notin L \Rightarrow \forall \pi \in \{0,1\}^{p(n)} : V(x,\pi) = 0$.

If we additionally assume that in the completeness case, the proof is unique, we would obtain the class UP. Obviously, UP is contained in NP. Moreover, there is a randomized reduction from any problem in NP to a problem in UP [VV85].

Definition 1.2. Class MA is a natural generalization of NP. In an MA protocol, Merlin (prover) sends Arthur (verifier) a message, and then Arthur decides whether to accept or not in polynomial time using randomness. A language L is in MA if there exists a polynomial-time randomized verifier V and a polynomial p such that for every input string x of length n

Completeness: $x \in L \Rightarrow \exists \pi \in \{0,1\}^{p(n)} : \Pr[V(x,\pi)] \geq 2/3$,

Soundness: $x \notin L \Rightarrow \forall \pi \in \{0,1\}^{p(n)} : \Pr[V(x,\pi)] \leq 1/3$.

Using standard Chernoff-style amplification, we can reduce the soundness error of an MA verifier to $2^{-\mathrm{poly}(n)}$. Moreover, MA admits perfect completeness [ZF87]. However, we do not know whether MA $\subseteq$ NP, but under assumptions implying BPP $=$ **P**, we have MA $\subseteq$ NP and thus MA $=$ NP [AB09].

Now assume that a protocol starts with Arthur tossing some random coins and sending the outcome of all his coin tosses to Merlin. After that Merlin responds and Arthur deterministically verifies the proof. These types of protocols correspond to the class AM. It was shown that coins can be made private with only a small round overhead [GS86].

One can generalize all classes discussed above by allowing polynomially many rounds resulting in the class IP, which is equal to PSPACE [Sha92].

2 Quantum-Classical Merlin-Arthur (QCMA)

We already introduced class QMA, a quantum analog of MA. A classical proof is replaced with a quantum proof and a polynomial-time verifier is replaced with a polynomial-size quantum circuit. This is a compelling open question whether the proof in the YES case can be made unique.

If Merlin sends a classical proof to a quantum Arthur, we obtain the so-called Quantum-Classical Merlin-Arthur (QCMA) class.

Figure 1: Class Diagram

Definition 2.1. Formally, a language L is in QCMA if there exists a P-uniform quantum circuit family $\{V_n\}_n$ and polynomials p, q, such that V_n takes in the input x of length n, classical proof π of length $p(n)$ and $q(n)$ ancilla qubits initialized to $|0\rangle$. The first ancilla qubit is the output qubit. In the end of the protocol, after applying V_n, this qubit is measured in the standard basis such that

Completeness: $x \in L \Rightarrow \exists \pi \in \{0,1\}^{p(n)} : \Pr[V_n(x,\pi)] \geq 2/3$,

Soundness: $x \notin L \Rightarrow \forall \pi \in \{0,1\}^{p(n)} : \Pr[V_n(x,\pi)] \leq 1/3$.

Notice that QCMA verifier can simulate a classical algorithm or ignore the witness, implying that $\mathsf{MA}, \mathsf{BQP} \subseteq \mathsf{QCMA}$. On the other hand, we do not believe that all quantum computations are contained in MA. It is also not known whether QMA equals QCMA, although QCMA is contained in QMA. One non-triviality in this containment is that any quantum proof can be made classical by measuring it. It is useful to remember that we still do not know whether $\mathsf{P} = \mathsf{PSPACE}$. In particular, it means that there is no hope of immediately separating QCMA and QMA.

If $\mathsf{QCMA} = \mathsf{QMA}$, then the Local Hamiltonian problem would be contained in QCMA, which means that there must exist some efficiently quantumly checkable classical certificate that a Hamiltonian has low ground energy. One way this could happen is for all local Hamiltonians $\mathcal{H}$ there exists a polynomial-size circuit $\mathcal{C}$ such that $\mathcal{C}|0\rangle^{\otimes n} = |\psi\rangle$, where $|\psi\rangle$ is some low energy state of $\mathcal{H}$. This seems to be implausible. Next time we will discuss some complexity theoretic evidence that these classes are different.

While QCMA-hardness does not appear to capture the local Hamiltonian problem, this class has exciting properties that are not known to hold for QMA. For instance, $\mathsf{QCMA} = \mathsf{QCMA}_1$, where the subscript means perfect completeness [JKNN11]. The main proof idea is to convert a quantum circuit so that the output probabilities are always nice numbers by writing it with only Hadamard, Toffoli and X gates. Then one may assume that any acceptance probability of the verifier on a classical proof has $2^{\mathrm{poly}(n)}$ in its denominator. The honest prover will reveal the probability that the verifier is supposed to get. Second ingredient is a phase estimation trick which helps to boost the correctness to 1 if you know the probability you are supposed to accept with. In comparison, we do not know whether an analog statement for QMA holds, since the class is possibly gateset dependent.

Finally, we want to mention a QCMA-complete problem called the Ground State Connectivity problem (GSCON). The input of the problem are two ground states $|\psi_1\rangle$ and $|\psi_2\rangle$ of a local Hamiltonian. The goal is to decide whether there exist polynomially many 2-qubit gates $U_1, \ldots, U_m$ such that $U_i \cdots U_1 |\psi_1\rangle$ are low-energy states for all i and $U_m \cdots U_1 |\psi_1\rangle \approx |\psi_2\rangle$. To convince ourselves that GSCON is in QCMA, notice that the classical proof is just a description of gates $U_1, \ldots, U_m$. Similar to the proof of containment of the local Hamiltonian problem in QMA, the

first property can be checked using a quantum phase estimation algorithm [Kit95]. The second property can be checked using the SWAP test [BCWdW01] that we discussed in the problem sets.

3 Stoquastic Merlin-Arthur (StoqMA)

The next quantum version of MA is the class StoqMA. We start our discussion of StoqMA by looking at a specific type of local Hamiltonians.

Definition 3.1. Hamiltonian $\mathcal{H}$ is stoquastic if all off-diagonal entries are real and non-positive: for all $i \neq j$

$$\langle i| \mathcal{H} |j\rangle \leq 0,$$

where we represent $\mathcal{H}$ with respect to the computational basis.

Notice that if we change bases this condition need not still be true. One can show that the stoquastic local Hamiltonian problem is MA-hard and is contained in AM [BDOT06].

StoqMA is a complexity class associated with stoquastic Hamiltonians.

Definition 3.2. A language L is in StoqMA if there exists a P-uniform quantum circuit family $\{V_n\}_n$ and polynomials p, q, such that V_n takes in the input x of length n, quantum proof π of length $p(n)$ and $q(n)$ ancilla qubits initialized to $|0\rangle$. Let L_n be the number of layers of V_n. Then the first layer consists only of parallel Hadamard gates applied on ancilla qubits. In the final layer Hadamard gates may be applied to any qubits. All intermediate layers are classical, i.e. consist of X, CNOT and Toffoli gates. The first ancilla qubit is the output qubit. In the end of the protocol, after applying V_n, this qubit is measured in the standard basis such that

Completeness: $x \in L \Rightarrow \exists \pi \in (\mathbb{C}^2)^{\otimes p(n)} : \Pr[V_n(x, \pi)] \geq \alpha$,

Soundness: $x \notin L \Rightarrow \forall \pi \in (\mathbb{C}^2)^{\otimes p(n)} : \Pr[V_n(x, \pi)] \leq \beta$,

where $\beta - \alpha \geq 1/\operatorname{poly}(n)$.

Definition 3.3. One canonical StoqMA-complete language is the so-called transverse field Ising model (TIM) consisting of stoquastic local Hamiltonians of the form

$$\mathcal{H} = \sum_{1 \leq i < j \leq n} p_{ij} Z_i Z_j + \sum_{1 \leq k \leq n} q_k X_k + g_k Z_k,$$

where q_k, g_k and p_{ij} are real coefficients and $q_k \leq 0$ [BH17].

For the rest of the lecture we will focus on frustration-free stoquastic local Hamiltonians.

Definition 3.4. A local Hamiltonian $\mathcal{H} = \sum_i^m H_i$ is called frustration-free if its ground state is a simultaneous ground state of each local term.

The main result is the following theorem.

Theorem 3.5. *The frustration-free stoquastic local Hamiltonian problem is* MA*-complete [BT10].*

We will not prove the theorem in its entirety and only show the most interesting part, which is the containment in MA.

We start with some intuition. First notice that many physical Hamiltonians arise from kinetic and potential energy and have this form.

Moreover, recall that the propagation part of the Feynman–Kitaev Hamiltonian is a graph Laplacian on the clock connecting frustration-free stoquastic local Hamiltonian problem to classical random walks.

Next, consider the so-called no sign problem. It roughly says that if you want to compute the expectation value of some observable $\mathcal{O}$ in the Gibbs state, you can write is as a sum of non-negative quantities multiplied by some number that depends on the observable:

$$\mathrm{Tr}\left[\mathcal{O}e^{\beta H}\right] = \sum_x w(x)\cdot\mathcal{O}(x),$$

where $w(x)$ are non-negative. This means that sampling configurations proportional to $w(x)$ yields the expectation of $\mathcal{O}$ as the Monte-Carlo average of $\mathcal{O}(x)$.

Finally, if we additionally assume that $\mathcal{H} = \sum_{i=1}^m H_i$ with $||H_i|| \leq 1$, then $\mathcal{H}' = mI - \mathcal{H}$ has non-negative entries. The ground state of $\mathcal{H}$ is the maximum energy state of $\mathcal{H}'$ and the top eigenspace of $\mathcal{H}'$ is spanned by non-negative states [Per07], [FFF$^+$12]! Such states can be written as linear combination of some basis elements with non-negative coefficients:

$$|\psi\rangle = \sum \psi_x |x\rangle$$

where $\psi_x \geq 0$.

For the proof, think of the Hamiltonian as encoding a graph and the state $|\psi\rangle$ as encoding a distribution on the graph. We will set up a random walk, whose stationary state is this distribution.

First we need a little lemma that will let us set up the graph and transition probabilities of a random walk. We should think of a graph whose edges are given by a projector in the given Hamiltonian that couples the two basis states with an off-diagonal entry that is not zero.

Lemma 3.6. *Let Π be a projector with non-negative entries and let $|\psi\rangle$ be a state such that $\Pi|\psi\rangle = |\psi\rangle$. Let $S(\psi)$ be the support of ψ and $x \in S(\psi)$, then*

$$\langle x|\Pi|x\rangle > 0.$$

Moreover, for all y, $\langle y|\Pi|x\rangle > 0$ implies $y \in S(\psi)$ and

$$\frac{\langle y|\psi\rangle}{\langle x|\psi\rangle} = \sqrt{\frac{\langle y|\Pi|y\rangle}{\langle x|\Pi|x\rangle}}.$$

One can prove the lemma by reducing to the rank one case $\Pi = |\psi_a\rangle\langle\psi_a|$ and then performing a direct calculation.

Proof of Theorem 3.5 (containment in MA*).* We have to give a Merlin Arthur proof system with a randomized verifier, where Merlin provides a proof that causes the verifier to believe that a given Hamiltonian has low energy.

Assume that $\mathcal{H} = \sum_i^m H_i$, where $H_i = (I - \Pi_i)$, for some projectors Π_i. Further assume that in the YES case the minimal energy $\lambda_{\min}(\mathcal{H})$ vanishes and in the NO case $\lambda_{\min}(\mathcal{H}) \geq \varepsilon = 1/\operatorname{poly}(n)$.

In the beginning of the protocol Merlin sends $x_0 \in \{0,1\}^n$ to Arthur. Then Arthur performs the following algorithm.

- Pick $0 < \beta < 1$ and $L = \operatorname{poly}(n)$ so that $2^{n/2}(1 - \varepsilon\beta m^{-1})^L \leq 1/3$.
- For $j = 0, \ldots, L-1$
  - Check that $x_j \in S_{good}$:
    $$\forall i, \langle x_j | \Pi_i | x_j \rangle > 0.$$
    If not, reject.
  - Pick an $i \in [m]$ at random. Let
    $$G_i = I - \beta H_i,$$
    so that each G_i is positive and also has nonnegative entries everywhere.
  - Let the set of neighbors $N_i(x_j)$ be all strings y with $\langle y | \Pi_i | x_j \rangle > 0$.
  - Move to $y \in N_i(x_j)$ with probability
    $$P^i_{x_j \to y} = \underbrace{\sqrt{\frac{\langle y | \Pi_i | y \rangle}{\langle x_j | \Pi_i | x_j \rangle}}}_{=: r^i_{x_j \to y}} \cdot \langle y | G_i | x_j \rangle .$$
    Set $x_{j+1} = y$, $r_{j+1} = r^i_{x_j \to y}$.
- After L iterations: if we haven't rejected yet, and if
  $$\prod_{i=1}^{L} r_i \leq 1,$$
  then accept. Else reject.

Now we will analyze the algorithm. Assume that $\mathcal{H}$ has zero minimal energy. Then honest Merlin will send x_0 from the support of the ground state such that x_0 has the highest weight. Notice that using the lemma and the fact that $\mathcal{H}$ is frustration-free we can also confirm that $P^i_{x_i \to y}$ is indeed a probability:

$$\begin{aligned}
\sum_y P^i_{x_i \to y} &= \sum_y \frac{\langle y | \psi \rangle}{\langle x_i | \psi \rangle} \langle y | (I - \beta H_i) | x_i \rangle \\
&= \frac{1}{\langle x_i | \psi \rangle} \cdot \langle x_i | (I - \beta H_i) \sum_y |y\rangle \langle y | \psi \rangle \\
&= \frac{1}{\langle x_i | \psi \rangle} \cdot \langle x_i | \psi \rangle \\
&= 1.
\end{aligned}$$

Set $\pi(x_j) = \langle x_j|\psi\rangle^2$. We claim that π is stationary.

$$\begin{aligned}
\pi(x_{j+1}) &= \frac{1}{m}\sum_i \sum_{x_j} P^i_{x_j \to x_{j+1}} \pi(x_j) \\
&= \frac{1}{m}\sum_i \sum_{x_j} \frac{\langle x_{j+1}|\psi\rangle}{\langle x_j|\psi\rangle} \langle x_{j+1}| (I - \beta H_i) |x_j\rangle (\langle x_j|\psi\rangle)^2 \\
&= \langle x_{j+1}|\psi\rangle \frac{1}{m}\sum_i \sum_{x_j} \langle x_{j+1}| (I - \beta H_i) |x_j\rangle \langle x_j|\psi\rangle \\
&= \langle x_{j+1}|\psi\rangle^2.
\end{aligned}$$

The lemma will also ensure that x_0 and all further $x'_j s$ are good. Moreover, the fact that x_0 has the highest weight will ensure $\prod_{i=1}^L r_i = \frac{\langle x_L|\psi\rangle}{\langle x_0|\psi\rangle} \leq 1$ and Arthur will accept.

In the soundness case, we use the bound:

$$\begin{aligned}
p_{accept}(x_0) &= \frac{1}{m^L} \sum_{i_1,\ldots,i_L} \sum_{x_1,\ldots,x_L \in S_{good}} \left(\prod_{j=1}^L r_j \langle x_{j-1}| G_{i_j} |x_j\rangle \right) \cdot \mathbf{1}\left[\prod_{j=1}^L r_j \leq 1 \right] \\
&\leq \frac{1}{m^L} \sum_{i_1,\ldots,i_L} \sum_{x_1,\ldots,x_L \in S_{good}} \langle x_0| G_{i_1} |x_1\rangle \langle x_1| G_{i_2} |x_2\rangle \ldots \langle x_{L-1}| G_{i_L} |x_L\rangle \\
&\leq \frac{1}{m^L} \sum_{i_1,\ldots,i_L} \sum_{x_1,\ldots,x_L \in \{0,1\}^n} \langle x_0| G_{i_1} |x_1\rangle \langle x_1| G_{i_2} |x_2\rangle \ldots \langle x_{L-1}| G_{i_L} |x_L\rangle \\
&= \sum_{x_L \in \{0,1\}^n} \langle x_0| (\underbrace{G}_{I - \beta \frac{1}{m}\sum_i H_i})^L |x_L\rangle \\
&= \sqrt{2^n} \cdot \langle x_0| G^L |+\rangle^{\otimes n} \\
&\leq \sqrt{2^n}(1 - \varepsilon\beta/m)^L \leq 1/3.
\end{aligned}$$

□

We conclude by an intuition on where $\prod_{i=1}^L r_i \leq 1$ comes from. Notice that a stoquastic Hamiltonian can be frustrated for two reasons: the walk can take you outside of the good set, and the projectors can have the same good set but be "misaligned." Condition $\prod_{i=1}^L r_i \leq 1$ helps Arthur to detect these cases.

References

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.

[BCWdW01] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical review letters*, 87(16):167902, 2001.

[BDOT06] Sergey Bravyi, David P Divincenzo, Roberto I Oliveira, and Barbara M Terhal. The complexity of stoquastic local hamiltonian problems. *arXiv preprint quant-ph/0606140*, 2006.

[BH17] Sergey Bravyi and Matthew Hastings. On complexity of the quantum ising model. *Communications in Mathematical Physics*, 349(1):1–45, 2017.

[BT10] Sergey Bravyi and Barbara Terhal. Complexity of stoquastic frustration-free hamiltonians. *Siam journal on computing*, 39(4):1462–1485, 2010.

[FFF+12] Georg Frobenius, Ferdinand Georg Frobenius, Ferdinand Georg Frobenius, Ferdinand Georg Frobenius, and Germany Mathematician. Über matrizen aus nicht negativen elementen. 1912.

[GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 59–68, 1986.

[JKNN11] Stephen P Jordan, Hirotada Kobayashi, Daniel Nagaj, and Harumichi Nishimura. Achieving perfect completeness in classical-witness quantum merlin-arthur proof systems. *arXiv preprint arXiv:1111.5306*, 2011.

[Kit95] A Yu Kitaev. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995.

[Per07] Oskar Perron. Zur theorie der matrices. *Mathematische Annalen*, 64(2):248–263, 1907.

[Sha92] Adi Shamir. Ip= pspace. *Journal of the ACM (JACM)*, 39(4):869–877, 1992.

[VV85] Leslie G Valiant and Vijay V Vazirani. Np is as easy as detecting unique solutions. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 458–463, 1985.

[ZF87] Stathis Zachos and Martin Furer. Probabilistic quantifiers vs. distrustful adversaries. In *International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 443–455. Springer, 1987.