CMSC 858L: Quantum Complexity

Instructor: Daniel Gottesman

Spring 2023

24 3-Round Quantum Interactive Proofs

24.1 References

Vidick and Watrous, "Quantum Proofs," arXiv:1610.01664 is a good reference for quantum interactive proofs (and also QMA).

24.2 Project

The project will have student presentations in class on May 9 and 11, and the written component of the project is due Friday, May 12. As a reminder, the written part should be in the range 10–30 pages; most of the groups are small, so I am hoping for the lower end of that. If you are in a group of greater than 1, you should include a section saying what each person in the group contributed to the project. My expectation is that grades for the project will in most cases be the same among everyone in the group, but in case one part is substantially better or worse than the rest, I could possibly give a different grade to members of the same group.

There are 14 projects, so the presentations will have to be at most 10 minutes, including questions and changeover time. I will put up a Google Docs sheet next week for groups to sign up for time slots.

24.3 Completeness of Quantum Interactive Proofs

The general definition of quantum interactive proofs, as we defined it last time, is:

Definition 1. *Let QIP be the set of languages L for which there exists an interactive quantum protocol Π such that for any instance x of L,*

- *If $x \in L$, then there exists a verifier Alice such that Bob outputs "accept" with probability $\geq 2/3$.*
- *If $x \notin L$, then for all actions by the verifier Alice, Bob outputs "reject" with probability $\geq 2/3$.*

Of course, the precise cutoff $2/3$ is not important to the definition of the class. The cutoff can be made closer to 1 using *sequential repetition*, where we run the protocol, and then run it again and again, only accepting if each iteration is accepted. Note that *parallel repetition*, in which we run many times all at once, also works, but requires additional proof that the prover cannot cause the verifier's different runs of the protocol to give correlated accept/reject results. In the case of interactive proofs, it turns out that parallel repetition gives independent accept/reject probabilities in the best case for the prover, but the analogous statement is not true for all different kinds of interactive protocols, so you do need to be careful about this sort of thing.

While this lets the acceptance probability get very close to 1, it turns out that there is a different way to modify a quantum interactive proof so that for a "yes" instance, the verifier accepts with probability exactly 1, and in the "no" instance, Bob rejects with probability at least $1/9$. This means that the protocol has *perfect completeness.*

Recall that we can formalize interactive protocols as a sequence of unitaries U_i^A by prover Alice alternating with unitaries V_i^B by verifier Bob. U_i^A are computationally unbounded and act on Hilbert space $H_A \otimes H_M$ and V_i^B are efficiently implementable unitaries acting on Hilbert space $H_M \otimes H_B$. Which one goes first depends on the specific protocol. We can assume Alice's and Bob's operations are unitaries because we can purify the protocol by turning all CPTP maps into unitaries on a larger Hilbert space and simply not reusing any qubits that are supposed to be discarded.

I am going to let each message be a round, so either the U_i^A are only defined for even i or for odd i, with V_i^B defined for the other parity of i.

Note that Bob is always assumed to be following the protocol as specified. In the "yes" instance, Alice should perform the specific V_i^B given in the protocol description since they lead to a high acceptance probability by Bob, but in the "no" instance, we need to take into account the possibility that Alice will perform other unitaries.

The purification property is critical to the modification we need to do to show perfect completeness. Suppose we have a protocol and run it with some particular prover actions and verifier actions, stopping right before Bob makes his final measurement. The final state at this point is $\sqrt{1-p}|0\rangle|\psi_0\rangle+\sqrt{p}|1\rangle|\psi_1\rangle$, where p is acceptance probability by Bob, the first qubit is the qubit Bob will measure to determine acceptance and the states $|\psi_b\rangle$ are states of all the remaining qubits in $H_A \otimes H_M \otimes H_B$.

We then modify the protocol so that Bob introduces an extra ancilla qubit in the state $|0\rangle$ and does a CNOT from his measurement qubit to the ancilla. The state is now

$$\sqrt{1-p}|0\rangle_{anc}|0\rangle|\psi_0\rangle + \sqrt{p}|1\rangle_{anc}|1\rangle|\psi_1\rangle. \tag{1}$$

Bob then sends all of his other qubits (except for the new ancilla) to Alice. Alice then performs a unitary that takes $|0\rangle|\psi_0\rangle \mapsto |0\rangle|00\ldots0\rangle$ and $|1\rangle|\psi_1\rangle \mapsto |0\rangle|00\ldots0\rangle$. This is definitely unitary since the two input states and the two output states are orthogonal. It might be a hard unitary to implement, but Alice is computationally unbounded. Note also that Bob always follows the protocol, so Alice knows what $|\psi_0\rangle$ and $|\psi_1\rangle$ are.

Then Alice sends the first qubit along with the value of p back to Bob, who now holds the state $\sqrt{1-p}|00\rangle + \sqrt{p}|11\rangle$. Provided the amplitudes $\sqrt{p}$ and $\sqrt{1-p}$ can be implemented using the gate set available to Bob, he can make a measurement in a basis including this state. He accepts is this is the state he has (provided $p \geq 2/3$) and rejects otherwise. If p in the original protocol is not a number exactly constructible by Bob, Alice can slightly lower p until it is by adding a small amplitude to states that will always cause Bob to reject.

This gives a modified interactive protocol with two extra rounds compared to the original. If we have a "yes" instance of the original protocol, and Alice follows the protocol, she then has a strategy that will cause Bob to accept with 100% probability. But what if it is a "no" instance? Now $p \leq 1/3$, and whatever unitary Alice performs, the state Bob has is of the form

$$|\phi_t\rangle = \sqrt{1-p}|0\rangle|\phi_0\rangle + \sqrt{p}|1\rangle|\phi_1\rangle. \tag{2}$$

Alice is claiming the acceptance probability $p' \geq 2/3$, which would imply Bob should have the state

$$|\phi_f\rangle = \sqrt{1-p'}|00\rangle + \sqrt{p'}|11\rangle, \tag{3}$$

but

$$\begin{aligned} |\langle\phi_f|\phi_t\rangle|^2 &= |\sqrt{(1-p)(1-p')}\langle 0|\phi_0\rangle + \sqrt{pp'}\langle 1|\phi_1\rangle|^2 && (4)\\ &\leq |\sqrt{(1-p)(1-p')} + \sqrt{pp'}|^2 && (5)\\ &= 1-(p+p') + 2pp' + 2\sqrt{pp'(1-p)(1-p')} && (6)\\ &\leq 8/9, && (7)\end{aligned}$$

where the last line uses the fact that the previous line is maximized for p taking its maximum value 1/3 and p' taking its minimum value 2/3. Thus, Bob accepts with probability at most 8/9 for a "no" instance.

Bob's chance of rejecting no instances can be increased by repeating the protocol, either sequentially or in parallel, as discussed above.

Note also that this is inherently a quantum modification of the protocol, making intrinsic use of purification and entangled states.

24.4 3-Round Protocols

Given a protocol like the above that satisfies perfect completeness, suppose the original protocol runs for $T = 2t$ or $T = 2t - 1$ rounds. Then we can create a modified protocol that reduces the number of rounds to $t + 1$ or $t + 2$, to a minimum of 3, as follows:

1. Alice runs a simulation of Bob and performs the protocol with the simulation up to round t, poducing the state $|\psi_t\rangle$ of $H_A \otimes H_M \otimes H_B$ that would be produced under the original protocol at this point. She then sends to Bob $H_M \otimes H_B$.

2. Bob chooses random bit $b = 0$ or $b = 1$.

    (a) If $b = 0$, Alice and Bob perform the remainder of the protocol starting from round t with the real Bob now. Bob accepts according to the usual measurement at the end of the protocol.

    (b) If $b = 1$, Alice and Bob now perform the protocol in *reverse* starting from round t back to the initial state, each performing $(U_i^A)^\dagger$ or $(V_i^B)^\dagger$ at the appropriate time. Bob accepts if his state at the end of this procedure is $|00\ldots0\rangle$, the assumed initial state of the original protocol.

Note that the last round of communication must always have Alice sending to Bob so that Bob can make the final measurement. Thus, if the original protocol has a number of rounds that is 0 mod 4, then at round t, going either forward or backward, Bob is supposed to send to Alice, so he can combine sending b with that transmission, and the number of rounds in the modified protocol is $t + 1$ (step 1, and then the t steps of the half protocol chosen). If the original number of rounds is 2 mod 4, then at round t, going either direction means Alice has to send next, so the total number of rounds is $t + 2$ (step 1, Bob sending b, and then the t rounds of the half protocol). If the original number of rounds is odd, Alice sends first, so if we have $b = 1$, we don't need to reverse through Alice's initial unitary, only through Bob's first unitary. This means that for either 1 mod 4 or 3 mod 4 rounds in the original protocol, we can have at most $t + 1$ rounds

In particular, if the original protocol had 2 rounds, we have actually increased the number of rounds to 3. If it was originally 3, then the new protocol has 3 rounds again. If it originally had 4 rounds, the new protocol has 3 rounds, and if it originally had 5 rounds, the new protocol has 4 rounds. For greater than 5 rounds, the number of rounds is definitely decreased, by close to a factor of 2 for a large number of rounds.

Now, for a "yes" instance, Alice can just follow the protocol, and since the original protocol has perfect completeness, the modified protocol is also always accepted by Bob: If $b = 0$, they just end up running the original protocol and Bob accepts 100% of the time for that, and if $b = 1$, they reverse perfectly (since in step 1, Alice has accurately simulated Bob's actions, so Bob can perfectly reverse them) and Bob always accepts.

For a "no" instance, however, Alice could produce any state of her choice to give to Bob after step 1. However, if it is far from the state that would normally be produced by the protocol at step t, then Bob is likely to reject it when $b = 1$. If it is close to the state that would normally be produced by the protocol at step t, then Bob will reject with high probability when $b = 0$ (since it is a "no" instance).

In particular, suppose in step 1, Alice prepares the state $|\phi_t\rangle$. Let $|\psi_t\rangle$ be the state with highest fidelity to $|\phi_t\rangle$ among all states that could be produced by Alice and Bob at round t running the protocol from the correct initial state with Bob performing the correct unitaries V_i^B but Alice using any unitaries she wishes. Suppose $|\langle\phi_t|\psi_t\rangle| = v$.

When $b = 0$, Alice and Bob perform the rest of the protocol forward using the unitaries U_i^A for Alice (of her choice, not necessarily specified by the protocol) and V_i^B (as specified by the protocol). If they were to perform those unitaries on $|\psi_t\rangle$, they would get $|\psi_T\rangle$ at the end of the protocol, and because it is a "no" instance, the probability that Bob accepts with this state is at most $c < 1$, indicating an inner product of at most $\sqrt{c}$ with a perfect accepting state. The actual final state $|\phi_T\rangle$ has inner product v with $|\psi_T\rangle$, so if $v > \sqrt{c}$, there is a non-zero chance that the state is rejected by Bob.

When $b = 1$, Alice and Bob run the protocol in reverse from step t. Suppose while doing so, Alice performs the unitaries $(U_i^A)^\dagger$; and of course, Bob performs $(V_i^B)^\dagger$. Starting with the initial state $|00\ldots0\rangle$, if they had run U_i^A and V_i^B, they would have gotten a state $|\tilde{\psi}_i\rangle$, and $|\langle\phi_t|\tilde{\psi}_t\rangle| \leq v$ by the definition of $|\psi_t\rangle$. But

$$|\langle 00\ldots0| \prod (U_i^A)^\dagger (V_i^B)^\dagger |\phi_t\rangle| = |\langle\tilde{\psi}_t|\phi_t\rangle| \leq v, \tag{8}$$

so the probability of Bob accepting when $b = 1$ is at most v^2.

Therefore, the actual acceptance probability is the average of the value for $b = 0$ and $b = 1$, and at least one of those is strictly less than 1, so whatever strategy Alice picks, the acceptance probability for a "no" instance is also less than 1.

We can repeat this procedure $O(\log T)$ times to get the number of rounds down to the minimum 3. Each time we do so, the probability of rejecting in the "no" instance gets closer to 0, so by the time we reach 3 rounds, the probability of rejection might be $O(1/\text{poly} n)$. At this point, we can increase the probability of rejection again by repetition, but of course, if we want the number of rounds to stay 3, we should use parallel repetition. We have the following theorem:

Theorem 1. *If there is a quantum interactive proof for a language L, there is a quantum interactive proof using only 3 rounds of communication.*

This is in stark contrast to the classical case, for which it is known that if there is a classical interactive proof with a constant number of rounds for every language in PSPACE, then the polynomial hierarchy collapses. Again, our argument is explicitly quantum, although it may be less obvious how it is. The ability to reverse the protocol when $b = 1$ relies on the purification; otherwise it would not be possible to do this in a way that still contains the full randomness of Bob's choices.