CMSC 858L: Quantum Complexity

Instructor: Daniel Gottesman

Spring 2023

26 Multiprover Interactive Proofs

26.1 References

We are working towards a discussion of the result that MIP* = RE, which is from Ji, Nataranjan, Vidick, Wright, and Yuen, "MIP* = RE," arXiv:2001.04383. There are some nice blog posts by Yuen (https://quantumfrontiers.com/2020/03/01/the-shape-of-mip-re/#fn-14879-6) and Vidick (https://mycqstate.wordpress.com/2020/01/14/a-masters-project/) that talk about the background and history leading up to the result.

26.2 Classical Multiprover Interactive Proofs

We have discussed interactive proofs with a single prover, but you can also consider interactive proofs with multiple provers. The definition for a classical multiprover interactive proof is the same as for a single-prover proof, but with 2 or more provers. (It turns out that any language that has a r-prover interactive proof has a 2-prover interactive proof, so we only need to consider the 2-prover case.) Critically, the provers cannot communicate during the protocol (otherwise you might as well consider the team of the provers to be a single prover), but they can pre-arrange a strategy for answering the verifier's questions (which could involve sharing random numbers unknown to the verifier). When the provers decide on their strategy, they don't know exactly what instance they wil be asked about, so they need to pick something that will work for any instance. Again, the provers are considered to be infinitely powerful (computationally speaking) and the verifier is polynomially bounded. This gives the complexity class MIP.

Multiprover interactive proofs have extra power relative to single-prover interactive proofs because the verifier can cross-check answers from the different provers, who don't know what the other prover has been asked. This is like police questioning two suspects in different rooms to try to catch incosistencies in their stories.

For a concrete example of how it can help, consider the problem of k-coloring a very large bounded-degree graph, specified by a small circuit. (I.e., the circuit takes as input a node label and returns a list of labels for the adjacent nodes.) The verifier wishes to know if the graph can be colored such that each node has one of k colors, but no two adjacent nodes have the same color. Obviously, when the graph is polynomial in size, this problem is in NP: The witness is a valid coloring. When the graph is exponential size, though, it is not clear how even a single-prover interactive proof can prove there is a valid coloring. The verifier can ask about the colorings of some nodes, but for any limited set of nodes, it might be that there is a valid coloring of the subgraph even though globally there is some obstruction to the coloring. (As a very simple example of a global obstruction, imagine a graph which is a large circle with an odd number of nodes; this graph cannot be 2-colored, but any proper subgraph can be.) That is, there are some graphs where the most obvious interactive proofs have 0 soundness: The prover can always (100%) cheat.

But there is a straightforward protocol to get at least a small amount of soundness with two provers: The verifier can ask one prover for the colors of two random adjacent nodes (a, b), and ask the second prover for the colors of either (a, c) or (b, d) (with a 50% chance of each), where a and c are adjacent and b and d

are adjacent. (It is possible that $c = b$.) The verifier rejects if either adjacent pair has the same color or if the two provers give different colors for the node that both are asked.

Claim 1. *If the graph has N nodes and cannot be k-colored, there is a chance at least $\Omega(1/N)$ that the verifier rejects.*

Proof. Let us suppose that we have a particular run of the protocol; we can fix the random bits used by the provers to determine their strategy. Thus, the prover strategy can be simply described by two large functions $f_i(a, b)$ that determine what pair of colors prover i will say when asked about nodes a and b. If on query (a, b), prover 1 returns color g for node a, then if prover 2 does not also return color g for a for every query (a, c), there is at least a chance $\Omega(1/N)$ that the verifier will find an inconsistent color for node a and reject. Thus, we may assume that prover 2 always returns color g for a in any query. Similarly, prover 1 must return color g for a for any query. Thus, the prover strategies are actually a single function $g(a)$ that returns, for any node, the same color $g(a)$ from either prover, regardless of the other node queried to that prover.

That is, the prover strategy is a particular coloring of the graph. But by hypthosis, there is no valid k-coloring, so there must be at least one edge of the graph where the colors are the same. The chance that the verifier picks this edge to ask about is $\Omega(1/N)$, in which case the verifier will reject. □

Now, obviously the soundness of this particular protocol is quite low, particularly when N is exponentially large in the size of the circuit specifying the graph. However, note that we could not get any soundness at all for the single-prover interactive proof. Also, note that we only made a single query to each prover with a constant-size answer (although the query has to specify some nodes, and specifying even a single node takes polynomial space) and are still getting some results about an exponential-size graph.

There are much smarter protocols for this problem and related ones. It turns out that MIP = NEXP. NEXP is like NP, but with exponential-size witnesses and checking circuits. This is believed to be strictly stronger than EXP (exponential-time algorithms), which in turn is believed to be strictly stronger than PSPACE. We can actually prove that NEXP $\neq$ NP.

26.3 Quantum Multiprover Interactive Proofs

Now let us talk about quantum mutlprover interactive proofs. We immediately have to make some choices. We have two apparently separate orthogonal ways to make it quantum. We could allow the verifier and the communication with the provers to be quantum, and we could allow the provers to have pre-shared entanglement as part of their pre-arranged strategy. This gives us potentially three different quantum multiprover interactive proof complexity classes, one where the verifier is quantum and the provers don't share entanglement, one where the verifier is classical but the provers share entanglement, and one where the verifier is quantum *and* the provers share entanglement. However, it turns out the latter two are the same, which usually is considered to have a classical verifier and thus goes by the name MIP*.

The case where the provers don't share entanglement but the verifier is quantum hasn't received too much study, as far as I know, except in the case of *non-interactive proofs* with two unentangled provers (which is known as the class QMA(2)). Partly this is because in the interactive case, the verifier can provide the provers with some entanglement to use, but we now know that the full power of MIP* requires the provers to share more than a polynomial amount of entanglement, which is all that the verifier could provide. Thus, this would definitely be a different complexity class than MIP*.

However, MIP* is a much more natural class to consider. If the provers are quantum, then why wouldn't they be able to share entanglement when they plan their strategy? Entanglement is the quantum version of shared randomness, and the provers can share the entanglement without knowing what the instance will be.

As with the classical case, it turns out that 2 provers is enough for MIP*. MIP* also has a parallel repetition result, so we don't have to be too concerned about the precise cutoff probabilities for completeness and soundness.

It's not immediately clear whether MIP* is bigger or smaller than MIP (Or the same or incomparable.) It could be smaller: The shared entanglement between the provers can let them generate correlations stronger

than any classical systems, and this might be useful in cheating in some protocols. On the other hand, it can allow them to correlate their answers to harder problems and potentially use new protocols to prove more langauges than could be proven in MIP.

This was an open question for many years. One challenge is that it is not clear how much entanglement the provers will need to use for their optimal protocols. If it is a limited amount of entanglement, then you can describe the system with a semidefinite program and prove some concrete upper bounds on the complexity. But the provers have infinite computational power, so they might want to perform very very large unitaries, which would then potentially require large entangled states.

26.4 Nonlocal Games

MIP* has a close connection with *non-local games*, such as the CHSH and Bell inequalities. A non-local game can be thought of as a two-round multiprover interactive proof with 2 or more provers interacting with a single verifier. For instance, in the "non-local box" game, the two provers receive single bits a and b and must return x and y such that $x \oplus y = a$ AND b; that is $x = y$ unless $a = b = 1$, in which case $x \neq y$. Classical provers can only succeed in dong this with a maximum probability of $3/4$, whereas if the provers share entanglement, they can succeed with a probability of about 85%.

Non-local games relate to the foundations of quantum mechanics, as any non-local game where the provers can be do better with entanglement is a proof that there is no local hidden variable theory for quantum mechanics. In a local hidden variable theory, the underlying theory is really classical (with no or very limited long-range communication. Problem 2 on problem set 2 relates to this question as well.

One question of interest is what kinds of correlations can actually be generated by quantum systems and which cannot. In particular, when there is an *infinite* amount of entanglement, there is some ambiguity as to how the define this. One natural approach is to say that we look at entanglement in finite-dimensional Hilbert spaces, and then take the limit as the dimension goes to infinity. This gives us one possible set of allowed correlations. We should also probably include any set of correlations that can't be achieved exactly but can be achieved to arbitrary precision. In this case, the Hilbert spaces under consideration are tensor products of finite-dimensional Hilbert spaces held by the two provers.

However, there is an alternative way to think about infinite-dimensional entanglement. In quantum field theory, the notion of locality is a little bit fuzzy because the vacuum is entangled and operators spread over some area. Instead of tensor-product operators, people often consider something local if operators in the different "locations" (in this case, held by different provers) commute. *Tsirelson's problem* asks if these two definitions give the same set of correlations.

Tsirelson's problem turns out to be connected to an important problem in the mathematics of operator algebras. Mathematicians have a tendency to consider finite-dimensional Hilbert spaces to be boring, because there is only one Hilbert space of dimension D (up to isomorphism). They are usually much more interested in infinite-dimensional Hilbert spaces, of which there are many kinds. Mathematicians would like to classify them. The *Connes embedding problem* relates to this goal, and turns out to be equivalent to Tsirelson's problem: Resolving one would resolve the other as well.

26.5 MIP* in RE

Now back to MIP*. In the absence of any bound on the entanglement used in a quantum multi-prover interactive proof, it is hard to prove an upper bound on the complexity class MIP*. One bound we can prove easily is that MIP* $\subseteq$ RE.

Definition 1. *RE is the class of languages L for which there exists a Turing machine that will list all the elements of L.*

"RE" stands for *recursively enumerable*. Note that one disinguishing feature of this class compared to all the ones we have previously talked about is that there is no resource bound at all. It doesn't matter how long it takes the Turing machine to list an element, so long as it appears somewhere on the list. In particular, RE contains the halting problem: Consider a Turing machine that shares time through all programs by first

running program 0 for t steps, then program 1 for $t-1$ steps, then program 2 for $t-2$ steps, all the way to program $t-1$ for 1 step. Then the machine increments t by 1 and repeats, outputting the number of the program for any simulated program that halts. If program s halts after T steps, then once $t = s + T$, program s will be running T or more steps and will halt and therefore be output.

The halting problem is, in fact, RE-complete: Given a language $L \in$ RE, there exists a Turing machine to list its "yes" instances. We can modify this Turing machine to create one that takes input x and then lists all "yes" instances for L, halting when it sees x. This modified Turing machine halts iff $x \in L$, so a halting oracle would imply solving any $L \in$ RE.

Theorem 1. *MIP* $\subseteq$ RE*

Proof. We wish to construct a Turing machine that will output all "yes" instances for a language L in MIP*. Given any potential multi-prover interactive protocol for L, we can guess a bound on the amount of entanglement used by the provers, run through all possible prover strategies using that amount of entanglement and simulate that on a Turing machine to a very high degree of accuracy. If the result shows a strategy for the provers that exceeds the maximum possible value for a "no" instance, then we know that it must be a "yes" instance and we can output the instance. If not, we can try more entanglement, higher accuracy, or a different protocol. All of these things are unbounded in this case case, but the Turing machine is unbounded as well. For any "yes" instance, we will eventually get to a protocol, a strategy for the provers, and an accuracy that will make us certain that we have a "yes" instance. Again, we have no guarantee how long this will take, but that dosen't matter for this case. □

RE is a huge class, so we would expect that MIP* would be much smaller. Suprisingly, however, it is not. In fact, MIP* = RE! Somehow, the provers are able to prove even uncomputable things to a polynomially-bounded prover.

The proof that MIP* = RE also applies to Tsirelson's problem, since there is the possibility of arbitrary amounts of entanglement. It turns out that for provers in the commuting operators model instead of the tensor product model, there is an infinite sequence of semidefinite programs that upper bounds the probability of succeeding at a non-local game. That means that the class of quantum multiprover interactive proofs with provers who have commuting operators is included in co-RE. Since co-RE $\neq$ RE, this implies that there is some non-local game such that the commuting operators model of correlations succeeds with a different probability than the tensor product model. This resolves Tsirelson's problem, which in turn means that the Connes Embedding conjecture is false.