Entangled Provers and Nonlocal Games: An Introduction to $\mathsf{MIP}^* = \mathsf{RE}$

Nathan Sheffield [1]

[1]MIT

Introduction

In 2020, Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright and Henry Yuen published a breakthrough result in the field of Quantum Complexity Theory. Building on years of prior work, and drawing on sophisticated methods from classical complexity, they showed that the class of languages decidable by proof systems with multiple entangled quantum provers is equivalent to the class of all recursively enumerable languages – implying, in particular, that they are capable of deciding the halting problem. This result has profound implications for complexity theory, quantum mechanics, and operator theory. This poster aims to provide the background needed to understand the result, and then give an intuitive explanation of what it is, where it comes from, and why it's important. We'll give intuitive sketches of the very most high-level mechanics of the proof, ignoring all technical details but hopefully distilling an appreciation for why it's true.

Classical Interactive Proofs

To understand what **MIP*** $=$ **RE** is saying, we first need to understand what the "IP" in **MIP*** means. As Agi discussed in one of the tutorials, **IP** corresponds to the idea of an "interactive proof system". Suppose you have a computational problem you're trying to solve – for example, imagine you have a pair of graphs, and you're trying to tell if they're isomorphic (that is, whether its possible to move the vertices around to make them look like the exact same graph).

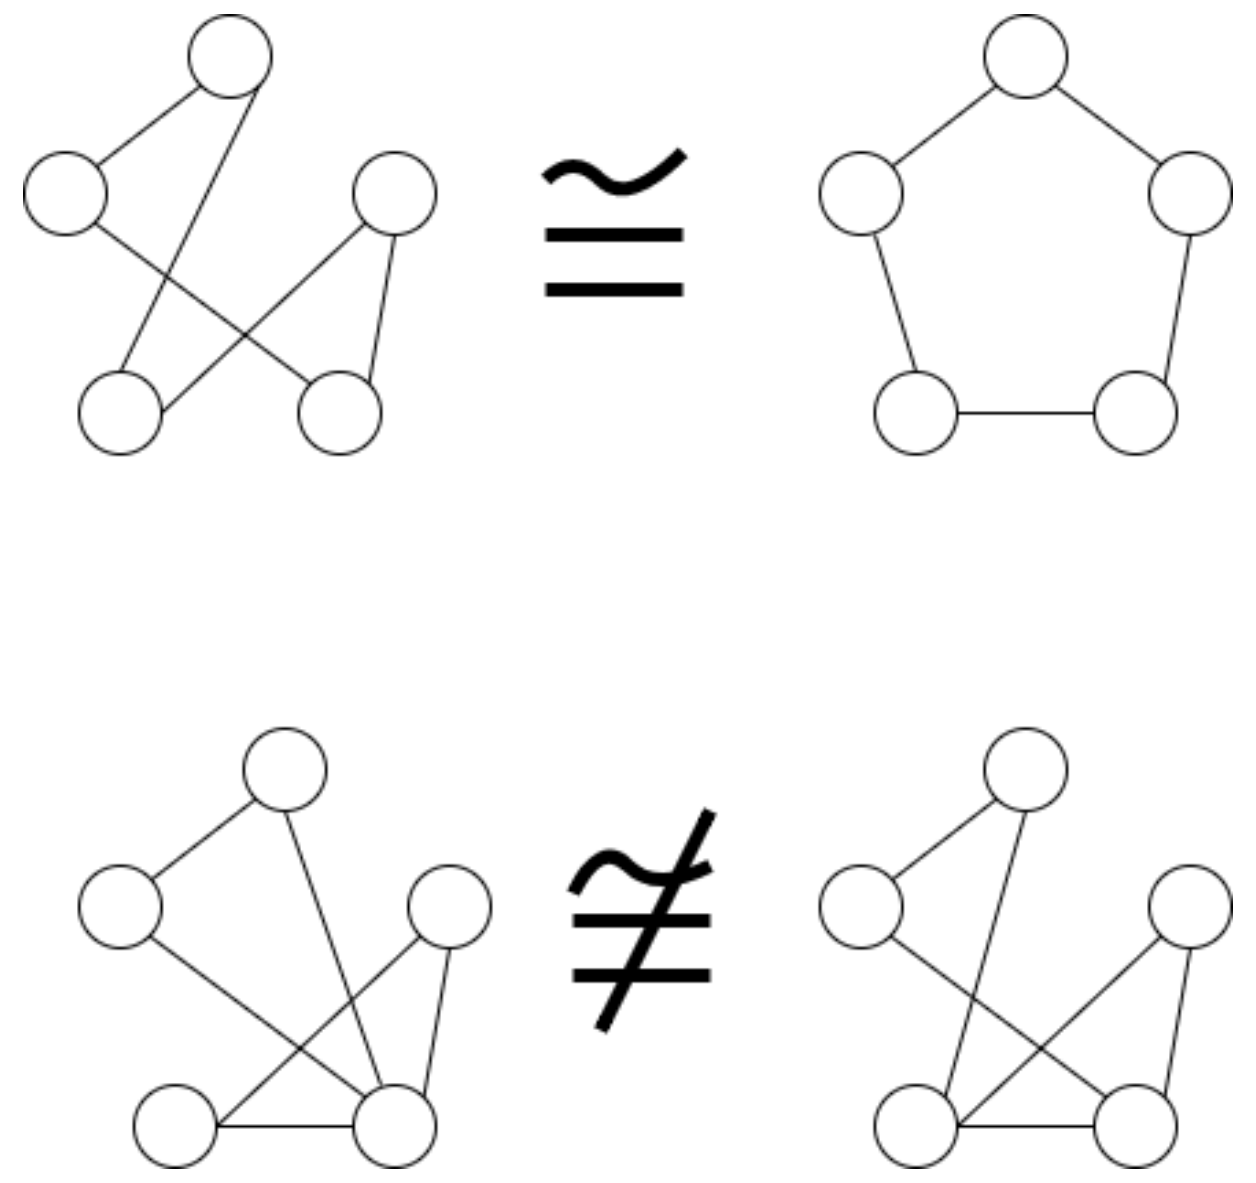

Figure 1. Top, a pair of isomorphic graphs. Bottom, two non-isomorphic graphs.

No efficient algorithm is known to solve this problem – it seems very hard to tell whether the two graphs are the same unless you try scrambling them in all the possible ways. However, imagine you had a very smart friend who could perform any calculation instantly. If you knew this friend was trustworthy, the problem would be very easy – you could just ask her the answer! The tricky bit is that your friend is known to be a bit of a prankster. Sometimes, she'll try to convince you of something false just to mess with you. You would like to have a series of questions you can ask her such that if she's being honest you'll almost always believe her, but if she's lying to you you'll almost always catch her. This defines the following class of problems:

> **Def (IP):** A problem is in **IP** if you (the "verifier") have some question-asking strategy such that
> - whenever the answer is `YES`, there's some question-answering strategy your friend (the "prover") can use to convince you to say "`YES`" at least 99% of the time
> - whenever the answer is `NO`, no matter what question-answering strategy the prover uses you'll say "`NO`" at least 99% of the time

Returning to the case of the pair of graphs, if you were trying to determine if two graphs were nonisomorphic, you might employ the following strategy:

1. Ask your friend if the graphs are isomorphic. If she says they are, say "`NO`".
2. If she says the graphs are different, choose one of the graphs at random, scramble it randomly, and then send it to her.
3. Now, ask her to tell you which of the original graphs the one you just sent her came from. If the graphs truly were nonisomorphic, she should be able to tell which this one is isomorphic to. If they actually *were* isomorphic though, and she was lying to us earlier, she won't be able to tell which one this came from, and so has no choice but to guess.
4. If you repeat this process many times, you can guarantee that the chance that your friend was lying to you and managed to guess right every time is less than 1%.

Proof systems with multiple provers

Now that we've defined the "IP" part of "**MIP*** $=$ **RE**, the next step is the "M". Here, "M" stands for "Multiple" – now, instead of having just *one* all-powerful friend, you have *two* (named Alice and Bob). To understand intuitively why this might be helpful for us, consider the classic analogy of a witness interrogation. If you're interrogating a single witness, it's possible for them to change their story to adapt to your line of questioning, so you have to be careful if you want to force them into a trap. If you're interrogating two witnesses that can't communicate with each other, though, then you can force them to give you a consistent story:

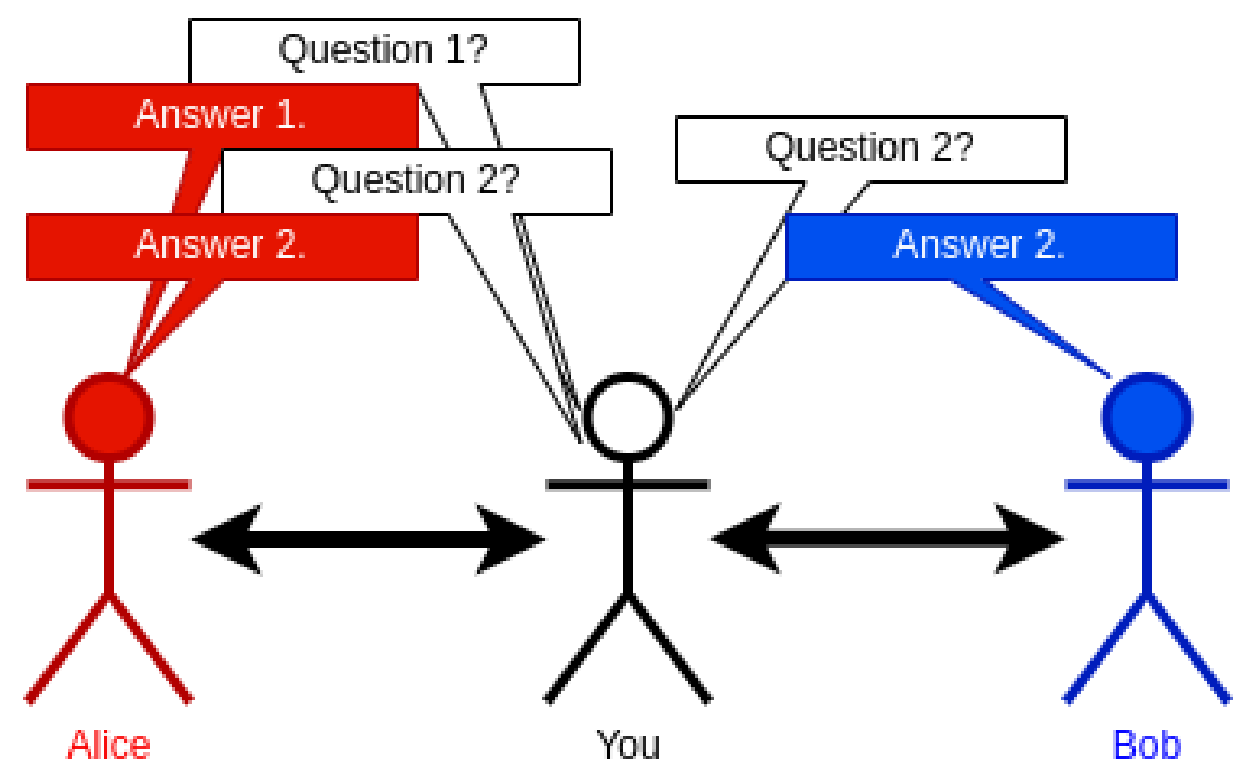

Figure 2. Example of the power of a **MIP** system. You can force Alice and Bob to decide on a fixed story in advance by having the conversation as normal with just Alice, but choosing a random question you asked her and giving it to Bob with no context. If Alice was adapting her responses based on the other questions you'd asked previously, then Bob wouldn't be able to give the same answer, because he didn't see the previous questions.

This intuitive helpfulness does indeed translate to an increase in strength: using a system with multiple interactive provers, it's possible to solve anything in **NEXP**, a very large class of problems.

Quantum Analog: MIP*

Now, finally, we see where quantum computation comes in. In the classical case, we think of the two provers as being **isolated**, but also **classically correlated**. That is to say, before we start interrogating them, Alice and Bob are allowed to share information between each other and come up with a shared strategy. Then, once the interrogation starts, they're no longer able to communicate. However, we know from quantum mechanics that not all physical correlations are explainable by local hidden variables! This leads to a natural generalization of **MIP** known as **MIP***, in which the provers are still isolated during the interrogation, but now are allowed to be **quantum correlated**. In other words, instead of their coordinated strategy consisting just of the shared classical information of what answers they plan to give to what questions, it can also consist of shared *quantum* information, in the form of entangled qubits.

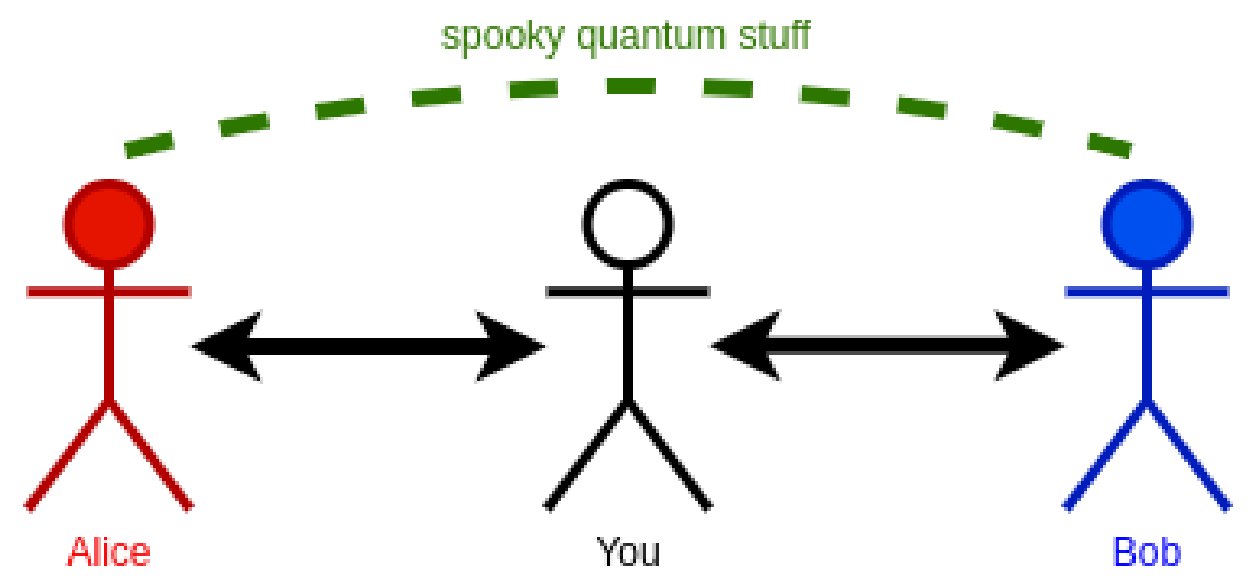

Figure 3. Although Alice and Bob can't communicate, their quantum states may be entangled with each other, allowing them to coordinate in ways that wouldn't be possible classically

How powerful is MIP*?

Looking at the definition of **MIP***, it seems intuitively like we've weakened the power of **MIP**. Before, we were restricting Alice and Bob's degree of information sharing, only allowing them to exchange classical bits beforehand. Now, even though they can't communicate any new information, allowing Alice and Bob to be entangled seems like it should make them more slippery in our interrogation. For instance, suppose Alice and Bob each held one half of an EPR state:

$$\frac{|0\rangle_A \otimes |0\rangle_B + |1\rangle_A \otimes |1\rangle_B}{2}$$

Now, if Alice chooses a response based on a measurement of her qubit of this state, it will collapse Bob's qubit to the same value, ensuring that they can give the same answer despite not having decided ahead of time what that answer would be. Of course, it is not immediately clear how this particular idea could help Alice and Bob – the value Alice gets from her measurement is totally random, so it seems like she could have just agreed with Bob ahead of time on which of the two random options was better. Is it possible that their shared quantum information can make Alice and Bob more likely to trick us than two classical provers could be? In order to discuss this question, we'll present a reframing of the **MIP*** system in terms of **nonlocal games**.

Nonlocal Games

Suppose, instead of a series of questions, we ask each of Alice and Bob only a *single* question, and compare their answers. This may sound like a weakening of the multiple interactive proof system, but with some investigation it can actually be shown to be equivalent. This allows us to recast the story of provers conspiring to trick us with a story of Alice and Bob playing a cooperative, nonlocal game. A nonlocal game works as follows:

1. Alice and Bob develop a shared strategy
2. They are isolated from each other (e.g. Alice stays on Earth and Bob is sent to Alpha Centauri, so that the speed of light restricts them from communicating over the timescale of the game)
3. A random pair of inputs (x, y) is generated by a probabilistic sampling algorithm $\mathcal{S}$
4. Alice receives x and must choose some output value a to respond with. Similarly, Bob receives y and must choose some output value b to respond with.
5. A deterministic algorithm $\mathcal{D}(x, y, a, b)$ determines whether or not Alice and Bob win the game.

Given that Alice and Bob know $\mathcal{S}$ and $\mathcal{D}$ beforehand, what is the maximum success probability they can achieve? We saw nonlocal games in lecture in the form of the CHSH game; here we'll present a different nonlocal game called the "**Magic Square Game**", which gives a clean illustration of the advantage of quantum strategies.

Magic Square Game

We think of the magic square game as being played on a tic-tac-toe board. The game works as follows:

- $\mathbb{S}$ chooses 2 numbers i and j, each uniformly at random from $\{1, 2, 3\}$, sending i to Alice and j to Bob
- $\mathbb{D}$ accepts a pair of responses a and b from Alice and Bob respectively if and only if
  - Both a and b are length three sequences of `X`'s and `O`'s
  - a has an even number of `X`'s
  - b has an odd number of `X`'s
  - The jth entry of a is equal to the ith entry of b

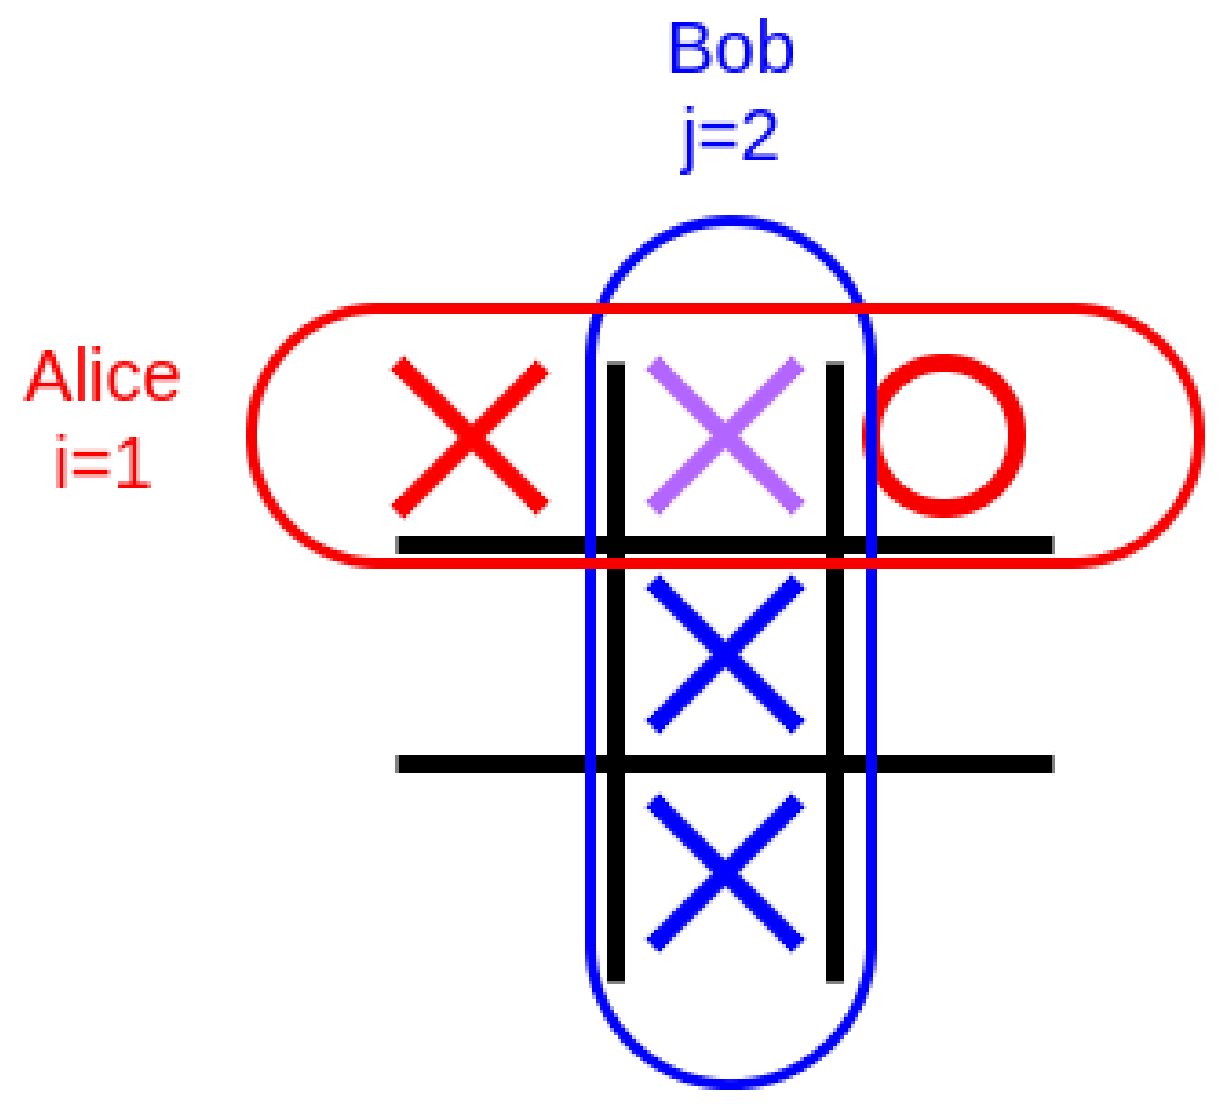

Figure 4. We think of Alice as returning the ith row of a tic-tac-toe board, and Bob as returning the jth column, with the restriction that all rows have an even number of `X`'s and all columns have an odd number of `X`'s. Here, Alice's response of `XXO` and Bob response of `XXX` causes them to win, since they do indeed agree on the overlap.

Consider a classical strategy for this game. Such a strategy must consist of a fixed response for each player for each number; that is, each player has a filled-in tic-tac-toe board in mind, and when queried they provide the corresponding row/column of their board. But note that since Alice's board has an even number of `X`'s in each row, and Bob's has an odd number of `X`'s in each column, Alice's board has an even number of `X`'s overall, while Bob's has an odd number. So, they must differ in at least one place, meaning that with probability $\frac{1}{9}$ the random values chosen by $\mathcal{S}$ will choose a row and column where Alice and Bob differ on the overlap. So, no classical strategy gets win probability greater than $\frac{8}{9}$. On the other hand, we can describe the following quantum strategy:

> In the quantum strategy for the Magic Square Game, Alice and Bob initially share two pairs maximally entangled qubits. That is to say, they start with 4 qubits between them in the state
>
> $$\frac{|00\rangle_A|00\rangle_B + |01\rangle_A|01\rangle_B + |10\rangle_A|10\rangle_B + |11\rangle_A|11\rangle_B}{2}$$
>
> Then, upon receiving i, Alice performs measurements on her qubits corresponding to the ith row of the following matrix, assigning an `X` to squares corresponding to measurements of $|1\rangle$, and `O` to squares corresponding to measurements of $|0\rangle$:
>
> 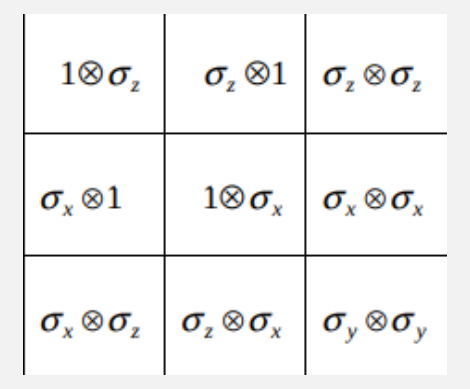
>
> | | | |
> |---|---|---|
> | $1 \otimes \sigma_z$ | $\sigma_z \otimes 1$ | $\sigma_z \otimes \sigma_z$ |
> | $\sigma_x \otimes 1$ | $1 \otimes \sigma_x$ | $\sigma_x \otimes \sigma_x$ |
> | $\sigma_x \otimes \sigma_z$ | $\sigma_z \otimes \sigma_x$ | $\sigma_y \otimes \sigma_y$ |
>
> Figure 5. Winning strategy for magic square problem [2]
>
> Similarly, Bob determines his response by performing the measurements in the jth column of the above matrix. By multiplying out the Pauli gates, we discover that this will result in winning assignments for all i, j.

Although any classical strategy for this game fails at least $\frac{1}{9}$ of the time, there exists a quantum strategy that succeeds every time. Recalling the corespondence between multiprover proof systems and nonlocal games (given a problem, Alice and Bob win by convincing the verifier of the wrong answer), this seems to justify our intuition that **MIP*** is a more restrictive class than **MIP**, because it seems as though Alice and Bob have more power to decieve us. However, as we will see, this intuition turns out to be entirely false! By carefully designing protocols around entanglement, we can use quantum multiprover systems to decide a shockingly large class of problems.

$\mathsf{MIP}^* = \mathsf{RE}$ Statement

As it turns out, the class **MIP*** of problems decidable by a polynomial-time verifier with access to multiple all-powerful entangled provers is actually equal to **RE**. **RE** is the class of **recursively enumerable** problems; that is, all problems of difficulty less than or equal to the halting problem. In particular, this class is **larger than the class of all problems decidable by any finite-time algorithm**! This is very surprising, since this is substantially more powerful than its classical counterpart **MIP**. In addition to being a profound result about the computational power of quantum information, the fact that these systems can decide undecidable problems has implications for both quantum physics and the mathematical theory of von Neumann algebras.

$\mathsf{MIP}^* \subseteq \mathsf{RE}$

The first part of this theorem is the upper bound – namely, that **MIP*** $\subseteq$ **RE**. There are classes of undecidable problems even larger than **RE**, so it could be possible for a proof system to actually decide an even larger language – but the proof of this upper bound is not very difficult. The basic idea is as follows: a quantum strategy for the provers consists of some number of qubits, and some quantum circuit that acts on those qubits to produce first an initial state before the game, and then on both those qubits and the inputs once the game begins. So, we consider an algorithm that starts by trying all possible circuits on 1 qubit, then 2 qubits, then 3, etc., checking to see whether if Alice and Bob used that strategy they'd be able to convince you with a 99% probability. If there's some strategy that works for Alice and Bob, this algorithm will find it in constant time and print "`YES`", otherwise it will run forever – so, the halting problem (a problem representative of the difficulty of **RE**) is at least as hard as any problem in **MIP***.

$\mathsf{RE} \subseteq \mathsf{MIP}^*$ sketch

The proof of the reverse inclusion is the subject of [4] – this proof is very involved, spanning around 200 pages and drawing on a lot of ideas from classical complexity theory. The basic idea, though, is that by asking Alice and Bob questions from an appropriate distribution, we can force them to have shared quantum information, and make whatever measurements we want them to make on it. This gives us substantially more control over what information Alice and Bob are allowed to see, allowing us to give them more power while carefully controlling it to prevent them from lying to us. The two main ideas we'll make use of to exploit this are **introspection** and **hiding**.

Introspection

The basic idea for how the proof happens is that we want to offload the work that we have to do to Alice and Bob. Concretely, suppose we have some verification procedure, but it requires the verifier to do a lot of work. That is, $\mathcal{S}$ takes a long time to because it has to generate and communicate long strings of random bits, and $\mathcal{D}$ takes a long time because it has to read through long responses to decide whether they win. We would like, ideally, for Alice and Bob to take over as much of this work as possible from us, because they have unbounded computational power and can do it all easily. However, this doesn't seem possible – after all, the whole point of this random checking and then verification of responses is because we can't trust Alice and Bob to be honest. If we ask Alice and Bob to do the random number generation for us, there's nothing stopping them from just choosing convenient values and pretending they're random!

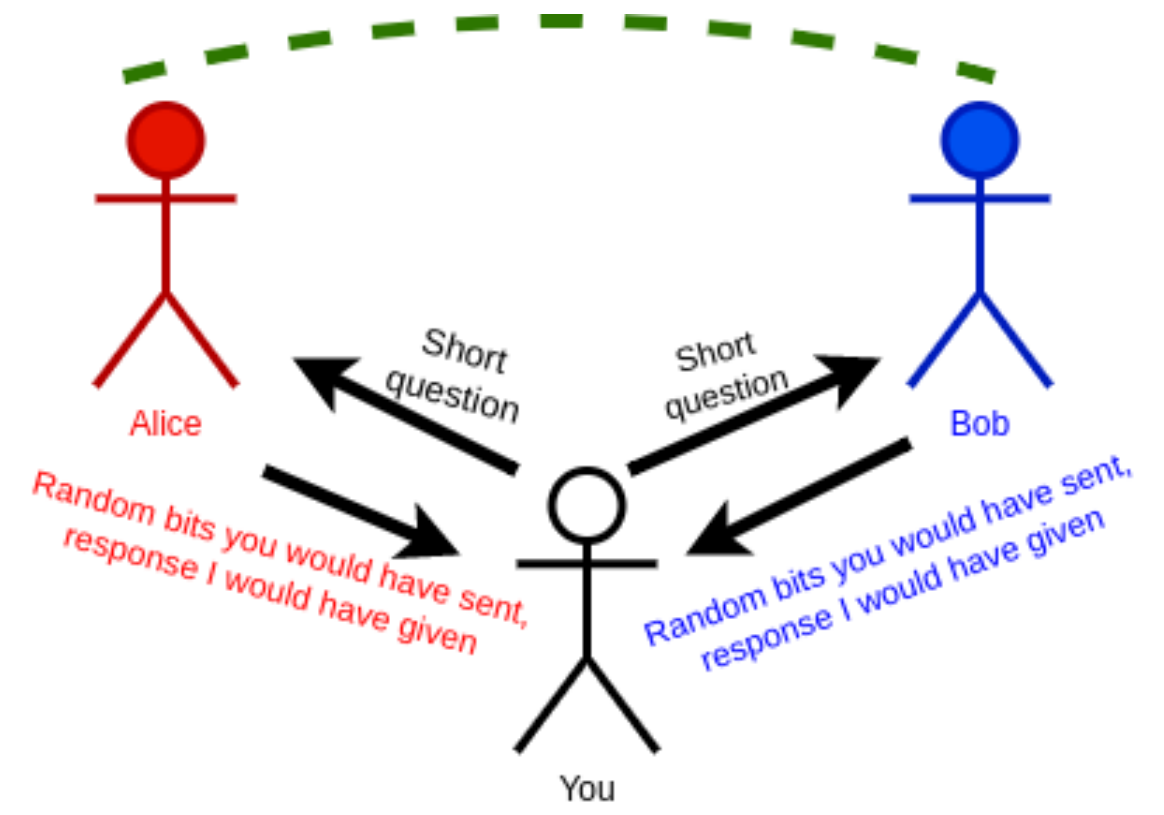

Alice
Short question
Short question
Bob
Random bits you would have sent, response I would have given
Random bits you would have sent, response I would have given
You

Figure 6. We would like it if Alice and Bob could come up with the questions we *would* have asked them, instead of us having to come up with them ourselves. This would save a lot of work, but seems hard to do without letting them trick us.

Also, in the original protocol, the random questions that we asked Alice and Bob were *related* to each other – in an interactive proof protocol, we want them similar things to allow them to provide a check on each others' answers. However, if Alice and Bob generate their own questions, they won't get to see the one generated for the other person. Both of these issues mean that in a classical **MIP** system there is no simplification we can do – however, these are both issues that are addressed by quantum mechanics! Quantum information has two properties that will be useful to us:

- **Randomness can be verified**. By performing small random checks on Alice and Bob, we can ensure that the values they're sending us are actually drawn from a truly random distribution, as opposed to letting them cheat. This works because we can force them to draw their randomness from **shared entangled qubits**, meaning that if one is not actually responding according to that random source, we can find out by asking the other prover to measure their copies of those qubits and see if the answers agree.
- **Randomness can be correlated**. If Alice could see all of the random bits Bob was using, she could cheat and the whole system would be no more powerful than a single prover. But, we want the bits she generates to depend on Bob's bits. Once again, the power of **shared entangled qubits** allows this to happen – by allowing Alice to measure only a subset of the shared qubits (the mechanism for ensuring this is explained in the next part), she can generate random questions for herself that depend on Bob's observations, but do not reveal too much information about them.

If we can develop a procedure that allows us to do only a communicate a small amount of information with Alice and Bob, and force them to generate their own questions honestly, then using a result called the **PCP theorem** from classical complexity we can modify the procedure such that their responses are short too. This will allow us to get an exponential reduction in the amount of work we have to do, since now a large fraction of the work of both generating questions and checking answers is delegated to the two provers.

Hiding

The key tool we'll use for ensuring that Alice and Bob see the right amount of information (e.g. enough to accurately simulate the original question-asker, but not enough to cheat) is the idea of **hiding**. We'll illustrate this with a example with particular importance to the full proof:

> **Classical low-degree testing**
>
> Suppose Alice and Bob have a function, and we're trying to verify that this function is a **low-degree polynomial**. To do so, we can send Alice a randomly chosen point and ask her for the value of the function at that point. Then we send Bob a randomly chosen line passing through that point, and ask him for a low-degree polynomial representing the function restricted to that line. Bob knows that Alice's point is somewhere on the line, but he doesn't know where it is, so if he tries to fit a low-degree polynomial to the original function it is likely to disagree with Alice's value unless the function really was low-degree.

To translate this protocol to the introspected prover setting, we'll need to ensure that Bob generates a random line passing through Alice's point, without allowing him to know which point along the line is Alice's point. To implement this, we can force Alice and Bob to start out with a set of 6 pairs of qubits, entangled uniformly over all tuples $|x_{1x}x_{1y}x_{2x}x_{2y}x_{3x}x_{3y}\rangle$, such that the points (x_{1x}, x_{1y}), (x_{2x}, x_{2y}), and (x_{3x}, x_{3y}) are colinear. Now, if Alice measures x_{1x} and x_{1y}, she'll get a point uniformly at random, and if Bob measures the remaining qubits he'll get a uniform random line passing through (x_{1x}, x_{1y}), without learning which point is Alice's.

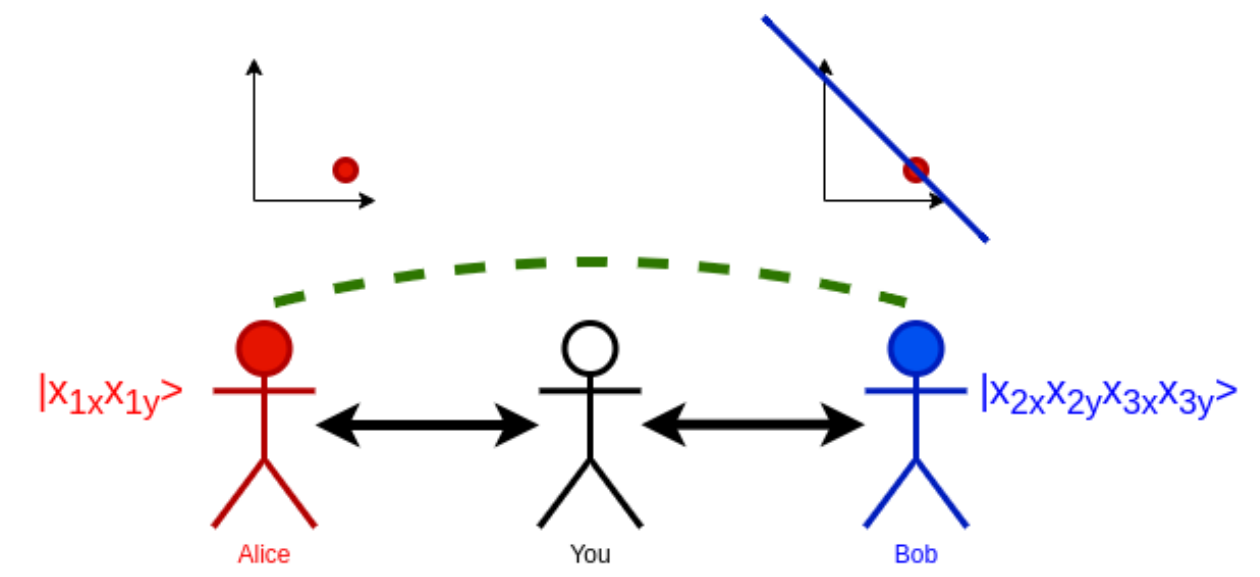

Figure 7. Alice and Bob each have only the information that would have been able to see in the classical low-degree test.

How do we ensure that, for instance, Bob doesn't sneakily measure $|x_{1x}x_{1y}\rangle$ to learn what Alice's point was? The trick is that we force him to send us, in addition to low-degree polynomial from the original protocol, **the result of σ_x measurements on** $|x_{1x}x_{1y}\rangle$! Half the time, we'll perform the protocol as usual, asking Alice to measure $|x_{1x}x_{1y}\rangle$ in the σ_z basis. But, the other half of the time, we'll actually ask her to perform a σ_x measurement instead, and make sure Bob's supposed σ_x measurement coincides. This forces Bob to be honestly performing σ_x measurements on the first two qubits, otherwise we'll catch him – but because of quantum mechanics this means that he cannot also read off Alice's values from the σ_z basis!

Overall Strategy

These ideas form the general primitives used in the **MIP*** $=$ **RE** proof. Using introspection techniques, we can convert a **MIP*** protocol into another **MIP*** protocol that solves the same problem while requiring exponentially less work for the verifier, by forcing the provers to make more use of entanglement. Informally, taking the "limit" of this procedure repeated many times allows us to speed up our computation so much that we can determine whether or not any given Turing Machine will halt, showing that **RE** $\subseteq$ **MIP***.

Conclusion

The proof of **MIP*** $=$ **RE** was revolutionary for the field of quantum complexity, and people are still trying to make sense of how it fits into our understanding. In addition to being an important statement in theory of computation, this problem was motivated by a long-unsolved mathematical problem in von Neumann algebra, known as **Connes Embedding Problem**, which concerns families of linear approximations to infinite-dimensional operators. It turns out that the fact that **MIP*** contains undecidable problems implies a negative result to this conjecture, giving (among other things) an unexpected difference in the predictions of the two major theories of quantum nonlocality! The outline here is only the simplest possible distillation of the original ideas of the paper; if you're interested, the original is very much worth reading.

References

[1] Simons institute quantum colloquium talks, April 2020.
[2] P. K. Aravind. A simple demonstration of bell's theorem involving two observers and no probabilities or inequalities, 2002.
[3] Sanjeev Arora and Boaz Barak. Computational complexity: A modern approach, 2016.
[4] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Mip*=re, 2020.