Lecture 5

Delegating Quantum Computations

So far we have been entirely focused on the problem of certifying a "qubit", i.e. that a certain device to which the experimentalist, or verifier, has access to and is willing to make simple assumptions about (the device has two spatially isolated components/the device is computationally bounded) is, at some point in its execution, making a pair of anti-commuting measurements.

Our goal in the next three lectures is to go beyond the certification of a single qubit, to the verification that the device implements an entire quantum computation of the verifier's choice. In this lecture we define this problem of *delegating quantum computations* to an untrusted party and give an overview of existing approaches. Towards the end of the lecture we describe a protocol due to Fitzsimons and Morimae [MF16] that involves quantum communication from the prover to the verifier. In subsequent lectures we combine that protocol with the computational qubit test from the previous lecture and a few additional ideas to obtain a purely classical protocol due to Mahadev [Mah18].

5.1 Problem statement

5.1.1 Quantum circuits and the class BQP

For us, a quantum circuit $\mathcal{C}$ is specified by an integer n and an ordered sequence of elements of the form (G, i, j) where $G \in \{H, CNOT, T\}$ and $i, j \in \{1, \ldots, n\}$. Letting

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \qquad T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}, \tag{5.1}$$

the circuit $\mathcal{C} = ((G_1, i_1, j_1), \ldots, (G_T, i_k, j_k))$ represents the unitary U on $(\mathbb{C}^2)^{\otimes n}$ obtained as $U = U_T \cdots U_1$ where for all $t \in \{1, \ldots, T\}$, U_t acts as the unitary associated with G_t in (5.1) on qubits i_t and j_t and as identity on the other qubits. (In case $G_t \in \{H, T\}$ it is required that $i_t = j_t$.) The Solovay-Kitaev theorem shows that any n-qubit unitary can be arbitrary well-approximated, in operator norm, by the unitary derived from a circuit; however, the size of the circuit may (in fact, must) in general grow exponentially fast with n. Those unitaries that can be represented by small circuits are called "efficient".

[0] Some of the material for this lecture is taken from an overview of Mahadev's result written for a mathematical audience and published in the Bulletin of the AMS [Vid20]. Some of it is reproduced from lecture notes prepared for a winter school at UCSD: `http://cseweb.ucsd.edu/~slovett/workshops/quantum-computation-2018/`.

Given a quantum circuit $\mathcal{C}$ acting on n qubits and $x \in \{0,1\}^m$ for some $m \leq n$ we say that "$\mathcal{C}$ accepts input x with probability p" if the probability of obtaining the outcome 1 after a measurement in the computational basis of the first qubit of the n-qubit state obtained by applying the unitary $\mathcal{C}$ to the input state $|x\rangle|0^{n-m}\rangle$ is p.

Definition 5.1. We say that a promise language $L = (L_{yes}, L_{no})$ is in BQP if there exists a family of polynomial-time generated quantum circuits $\{\mathcal{C}_n\}_{n\in\mathbb{N}}$ such that for all integer n and $x \in \{0,1\}^n$,[1]

- (Completeness:) If $x \in L_{yes}$ then $\mathcal{C}_n$ accepts x with probability at least $\frac{2}{3}$;
- (Soundness:) If $x \in L_{no}$ then $\mathcal{C}_n$ accepts x with probability at most $\frac{1}{3}$.

Note the requirement that the family $\{\mathcal{C}_n\}$ is polynomial-time generated. This means that there exists a classical Turing Machine that on input 1^n runs in time $\mathrm{poly}(n)$ and returns a description of $\mathcal{C}_n$ as a sequence of gates taken from a fixed universal set—here we use (5.1), but the specific choice will not matter for us.

The definition of BQP sets arbitrary values $2/3$ and $1/3$ for the completeness and soundness parameters. Error amplification works just as for the case of BPP, by repeating the circuit sequentially. This requires intermediate measurements, but it is not hard to show that these can be postponed till the end of the computation by the use of ancilla qubits and CNOT gates. As a result, any choice of a, b such that $a - b > \mathrm{poly}^{-1}(n)$ gives the same definition: for any such a, b, and for any fixed polynomial q, $\mathrm{BQP}(a,b) = \mathrm{BQP} = \mathrm{BQP}(1-2^{-q}, 2^{-q})$.

Exercise 5.1. Show that BQP is included in PP, the class of languages for which there exists a probabilistic Turing machine that accepts YES inputs with probability $> 1/2$, and rejects NO inputs with probability $> 1/2$. (Hint: first show inclusion in PSPACE by giving space-efficient implementations of basic linear algebra operations. Inclusion in PP follows from similar arguments, but is a bit more delicate.)

The class PP lies outside of the polynomial hierarchy. The most commonly-held belief is that the intersection of BQP and PH is non-trivial: it is neither BPP, nor PH itself. Recently Raz and Tal [RT19] showed that an oracle problem introduced by Aaronson [Aar10] is in BQP but not in PH.

Recall the notion of interactive proof system that we introduced informally in Section 2.2. We end this section by defining a family of complexity classes associated with interactive proof systems.

Definition 5.2 (Adapted from [AG17])**.** Given complexity classes $\mathcal{P}$ and $\mathcal{Q}$, $\mathrm{IP}[\mathcal{P}, \mathcal{Q}]$ is the class of (promise) languages L such that there is a polynomial-time Turing machine M that on input 1^n returns the description of classical circuits for the verifier V_n in an interactive protocol with a prover P such that

- (Completeness:) There is a family of provers $\{P_n\}_{n\in\mathbb{N}}$ that lie in the class $\mathcal{P}$ such that for all $x \in L_{yes}$ the interaction of $V_{|x|}$ and $P_{|x|}$ on common input x accepts with probability at least $\frac{2}{3}$.
- (Soundness:) For any family of provers $\{P_n\}_{n\in\mathbb{N}}$ that lie in the class $\mathcal{Q}$, for all $x \in L_{no}$ the interaction of $V_{|x|}$ and $P_{|x|}$ on common input x accepts with probability at most $\frac{1}{3}$.

When the classes $\mathcal{P}$ and $\mathcal{Q}$ coincide we simply write $\mathrm{IP}[\mathcal{P}]$ for $\mathrm{IP}[\mathcal{P}, \mathcal{P}]$. We use the standard notation $\mathrm{IP} = \mathrm{IP}[\mathrm{BPP}, \mathrm{ALL}]$ with ALL the class of all languages (i.e. soundness is proves without any restriction on the prover).

The definition is slightly informal, because for some classes $\mathcal{P}$ it may not be clear what it means for the prover to lie in $\mathcal{P}$. For us the meaning will always be clear from context, as $\mathcal{P}$ and $\mathcal{Q}$ will always be either BPP, BQP or ALL.

[1] Note that in general, $\mathcal{C}_n$ may act on $\mathrm{poly}(n)$ qubits, the first n of which are by convention destined to receive the input x and the first of which also serves as output qubit.

5.1.2 Delegating quantum computations

The fact that BQP is not (believed to be) in NP implies that in general we do not expect there to exist classically verifiable proofs for the correctness of an arbitrary quantum computation. This poses a challenge: as we see quantum computers emerging, how will we test their predictions? This is a practical problem — will anyone trust the "quantum cloud" — but also a philosophical one — is quantum mechanics a testable theory? (For more on this, see [AV13].)

Not all is lost. What we *do* know is that BQP is included in PSPACE, the class of languages that can be decided using polynomial space (and arbitrary time); in fact Exercise 5.1 asked you to show a stronger statement. And even though it is not a trivial result, it is known that PSPACE = IP. So all languages in BQP have *classical* interactive proofs, with an efficient classical verifier! Unfortunately there is a major caveat to this observation. The proof that PSPACE is in IP is based on the classical SUM-CHECK protocol, which in general requires the server to execute PSPACE-complete computations (essentially, the server has to compute exponentially large sums in order to determine answers that will satisfy the client). (For an exposition of the proof we refer to the book [AB09].)

So, even though a protocol exists, it is unknown if there is such a protocol in which a honest server is only required to have the power of BQP. Today this is a major open question:

Open Question 5.3. Is $\mathrm{BQP} \subseteq \mathrm{IP}[\mathrm{BQP}, \mathrm{ALL}]$? In words, do all languages in BQP have single-server interactive proofs in which the client has the power of BPP and for which completeness holds with a BQP server and soundness holds against any server?

There are some partial impossibility results [ACGK17] on this question, as well as possibility results where completeness holds for provers that require more power than BQP but not necessarily the entire power of PSPACE; see e.g. [AG17]. If, however, one allows slightly more power to the verifier then there are scenario in which the question is known to have a positive answer:

1. The client has access to a limited quantum computer, such as the ability to prepare single qubits in arbitrary states and send them to the server, or receive single qubits from the server and make simple measurements on them;

2. The client is allowed to interact with multiple quantum servers sharing entanglement.

The question as formulated above asks for *verifiable* delegation: given a quantum circuit (deciding some BQP language L), is there a protocol that allows a classical client to extract the outcome of the circuit from a BQP server, in a way that any cheating server, attempting to convince the client of the wrong outcome, will be detected? A second desirable property of a delegation protocol is *blindness*: while the client would like to learn the valid outcome of her circuit, she might not want to disclose the particular circuit or input she is interested in to the server. This is a distinct property from verifiability; in particular, one may ask for blindness in the "honest-but-curious" model, where verifiability is trivial. The following definition introduces these properties slightly more formally.

Definition 5.4 (Delegated computation)**.** In the task of delegated computation, a client (sometimes called the *verifier*) has an input $(x, \mathcal{C})$, where x is a classical string and $\mathcal{C}$ the classical description of a quantum circuit. The client has a multiple-round interaction with a quantum server (sometimes also called *server*). At the end of the interaction, Alice either returns a classical output y, or she aborts. A protocol for delegated computation is called:

- *Correct* if whenever both the client and the server follow the protocol, with high probability Alice accepts (she does not abort) and $y = \mathcal{C}(x)$. (This property is sometimes called *completeness*.)

- *Verifiable* if for any server deviating from the protocol, the client either aborts or returns $y = \mathcal{C}(x)$. (This property is analogous to what we have been calling *soudness*.)
- *Blind* if for any server deviating from the protocol, at the end of the protocol the server has no information at all about the client's input $(x, \mathcal{C})$.

The definition remains rather informal. For example, how should we formalize the "information" that the server has at the end of the computation? This can be rather delicate, especially once one starts taking into account a small chance ε of deviation from the perfect properties. A precise definition satisfying all the desired properties (universal composability in particular) would take us too far. Such a definition was given using the framework of *abstract cryptography* in [DFPR14].

The informal definition will be sufficient for our purposes. Note that in spite of being rather similar neither of the properties of verifiability or blindness is known to directly implies the other. In practice verifiability often follows from blindness by arguing, using "traps", that if a protocol is already blind then the server's trustworthiness can be tested by making it run "dummy" computations for which Alice already knows the output, without the server being able to distinguish whether it is asked to do a real or dummy computation. We will see an example of this technique later on.

Open Question 5.5. Is there a general transformation from any protocol satisfying blindness, to a protocol satisfying both verifiability and blindness? See [Mor18] for how to achieve this by making use of post-hoc verification (cf. Section 5.2), and [KMW17] for another approach.

Remark 5.6. The problem of delegating computation is interesting even for classical computation. In this case the client herself could directly execute the classical circuit $\mathcal{C}$. But it makes sense to be even more demanding, and seek protocols where the client is super-efficient: the best we could hope for is a client that runs in time *linear* in the input length, and independent of the size of the circuit. In addition, we would like the overhead for the server to be as small as possible, so that the honest behavior requires a server effort of the same order as the size of the circuit, $|\mathcal{C}|$. This kind of interactive proofs are called *doubly efficient* interactive proofs [GKR08]. The paper [RRR16] shows how to achieve such proofs with client runtime that is linear in the input length, polynomial in the space required by $\mathcal{C}$, and polylogarithmic in $|\mathcal{C}|$. If one is willing to make computational assumptions (essentially, subexponential LWE) then even more efficient delegation is possible [KRR14], with client runtime that is linear in the input size and poly-logarithmic in $|\mathcal{C}|$.

These results usually do not put emphasis on the requirement of blindness: they focus on verifiability alone. One reason for this is that blindness is "trivially solved" by employing homomorphic encryption [Gen09]. This, however, requires computational assumptions, and induces significant computational overhead.

5.1.3 Approaches to delegating quantum computation

There are three main types of approaches: *prepare-and-send* (the client has ability to prepare single-qubit states and send them to the server), *receive-and-measure* (the client has the ability to receive single-qubit states from the server and measure them), and *two-server* (the client interacts classically with two spatially isolated servers. A great recent survey describing these approaches in detail is [GKK19]. Here we focus on a specific protocol of receive-and-measure type, that we will later build on to obtain a protocol with a classical verifier, under computational assumptions on the prover.

5.2 The Fitzsimons-Morimae protocol

We describe the receive-and-measure protocol from [MF16], as it will form the basis for the Mahadev protocol.

5.2.1 The circuit-to-Hamiltonian reduction

The Cook-Levin theorem showing NP-completeness of the 3SAT problem is based on what could be called a "circuit-to-formula" reduction: given a classical circuit, the computation performed by the circuit on some input is represented as a "tableau" such that the property of being a valid tableau can be encoded in a formula whose variables represent the state of any given wire in the circuit and whose constraints enforce correct propagation of the gates of the circuit.

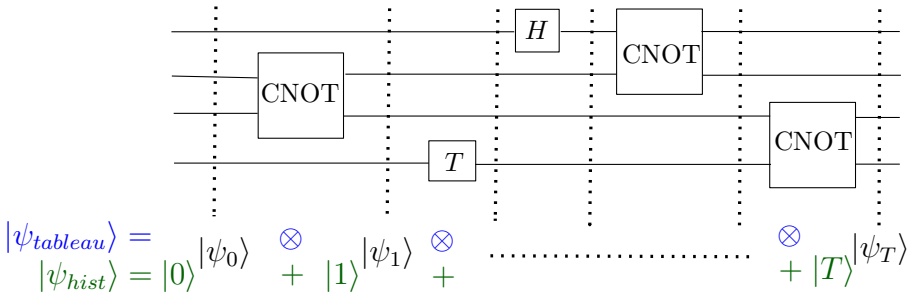

Figure 5.1: Two different ways to create a tableau from a quantum circuit. The state $|\psi_{tableau}\rangle$ is the tensor product of the state of the circuit at each time step. The state $|\psi_{hist}\rangle$ is their superposition, indexed by a clock register that goes from $|0\rangle$ to $|T\rangle$.

For quantum circuits the idea of a tableau of the computation is less straightforward. The most direct analogue is to consider the juxtaposition of the quantum state of a T-gate circuit at each step of the computation, i.e. the tensor product $|\psi_0\rangle \otimes \cdots \otimes |\psi_T\rangle$ of the states $|\psi_i\rangle$ obtained by executing the circuit from scratch and stopping after i gates have been applied. While this is a well-defined $n(T+1)$-qubit quantum state (see Figure 5.1) the property of being a valid "quantum tableau" cannot be enforced using *local* constraints! The reason is subtle, and has to do with the possible presence of entanglement at intermediate steps of the computation. Indeed, there are quantum states that are very different, in the sense that they are perfectly distinguishable by some *global* observable, yet cannot be distinguished at all by any *local* observable, that would act on at most, say, half the qubits. An example is given by the two n-qubit "cat" (named after the homonymous animal) states

$$|\psi_\pm\rangle = \frac{1}{\sqrt{2}}\big(|0\cdots 0\rangle \pm |1\cdots 1\rangle\big) .$$

The two states $|\psi_+\rangle$ and $|\psi_-\rangle$ are easily seen to be orthogonal, so that they can be perfectly distinguished by a measurement. But it is an exercise to verify that for any observable that acts on at most $(n-1)$ of the n qubits, both states give exactly the same expectation value. (Informally, this is because any measurement on a strict subset of the qubits of the state necessarily destroys the coherence; the only relevant information, the $\pm$ sign, is encoded "globally" and cannot be accessed locally.) Note that this is a uniquely quantum phenomenon: if two classical strings of bits have each of their bits equal, one pair at a time, then the strings are "globally" identical. Not so for quantum states.

So naïve tableaus will not do. In the late 1990s Alexei Kitaev introduced a very powerful idea that provides a solution. Kitaev's idea is to replace the juxtaposition of snapshot states by their *superposition* (see Figure 5.1). A special ancilla system, called the "clock", is introduced to index different elements of the superposition. Thus, instead of defining a tableau as $|\psi_0\rangle \cdots |\psi_T\rangle$, Kitaev considers the state

$$|\psi_{hist}\rangle \;=\; \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} |t\rangle|\psi_t\rangle \;. \tag{5.2}$$

Note that this takes less qubits to store, but this is not the important point. Kitaev showed that, assuming the clock register is encoded in unary, it is possible to check the correct propagation of every step of the circuit directly on this superposition by only applying local observables: there is a set of observables H_{in} that checks that $|\psi_0\rangle$ has the right format; a set of observables H_{prop} that checks propagation of the circuit, and an observable H_{out} that checks that the output qubit of the circuit is in the right state. (In addition, there is a term H_{clock} that checks that the clock register is well-formed, i.e. contains the representation of an integer in unary. This can be done locally by penalizing configurations of the form "$\cdots 10 \cdots$".) The key point that makes this possible is that, while equality of quantum states cannot be decided locally when the states are juxtaposed, it becomes possible when they are given in superposition. As an exercise, we can verify that a measurement of the first qubit of the state

$$|\psi_{SWAP}\rangle \;=\; \frac{1}{\sqrt{2}}\big(|0\rangle|\psi_0\rangle + |1\rangle|\psi_1\rangle\big)$$

in the Hadamard basis $\{|+\rangle, |-\rangle\}$ returns the first outcome with probability exactly $\frac{1}{2}(1 + |\langle\psi_0|\psi_1\rangle|^2)$. With more work, replacing the use of gadgets in the classical Cook-Levin reduction by techniques from perturbation theory, it is possible to write the resulting observables as a linear combination of local terms that all take a particularly simple form. The result is the following theorem from [CM16].

Theorem 5.7. *For any integer $n \geq 1$ there are $n' = \mathrm{poly}(n)$, $a = a(n)$ and $\delta \geq 1/\mathrm{poly}(n)$ such that the following holds. Given a T-gate quantum circuit $\mathcal{C} = ((G_1, i_1, j_1), \ldots, (G_T, i_T, j_T))$ acting on n qubits, such that $T = \mathrm{poly}(n)$, and an input x for the circuit, there exist efficiently computable real weights $\{J_{ij},\ i, j \in \{1, \ldots, n'\}\}$ such that $|J_{ij}| \leq 1$ for all i, j and moreover if*

$$H_{\mathcal{C}} \;=\; -\sum_{i,j} \frac{J_{ij}}{2} \left(\sigma_{\mathrm{X},i}\sigma_{\mathrm{X},j} + \sigma_{\mathrm{Z},i}\sigma_{\mathrm{Z},j}\right) , \tag{5.3}$$

where $\sigma_{\mathrm{X},i}$ and $\sigma_{\mathrm{Z},j}$ denote single-qubit Pauli X and Z operators acting on the i-th and j-th qubit respectively, then:

- *(Completeness) If the circuit $\mathcal{C}$ accepts its input x with probability at least $2/3$, then the smallest eigenvalue of $H_{\mathcal{C}}$ is at most a;*
- *(Soundness) If the circuit $\mathcal{C}$ accepts its input x with probability at most $1/3$, then the smallest eigenvalue of $H_{\mathcal{C}}$ is at least $a + \delta$.*

Remark 5.8. It is possible to modify Theorem 5.7 so that the completeness and soundness statements specify that "if there exists a state $|\phi\rangle$ such that $\mathcal{C}$ accepts on input $(x, |\phi\rangle)$ with probability at least $2/3$..." and "if there does not exist a state $|\phi\rangle$ such that $\mathcal{C}$ accepts on input $(x, |\phi\rangle)$ with probability greater than $1/3$..." respectively. Thus, Theorem 5.7 can be adapted to show that the problem of estimating the minimal energy of a Hamiltonian of the form (5.3) is a QMA-complete problem.

Theorem 5.7 provides us with a roadmap for the verification of quantum circuits: it is sufficient to verify the *existence* of a quantum state that yields certain statistics, when some of its qubits are measured in the computational (σ_Z observable) or Hadamard (σ_X observable) basis. The reason this can be considered progress is that we no longer need to check the time evolution of a quantum state under a quantum cicuit; it is sufficient to collect measurement statistics and estimate the "energy" $\langle\psi|H|\psi\rangle$. In particular, the theorem readily leads to a verification protocol in a model where the prover has a full quantum computer, and the verifier only has a limited quantum device — namely, a one-qubit memory, together with the ability to measure the qubit using either the σ_X or σ_Z observables.

5.2.2 The protocol

Such a verification protocol was introduced by Fitzsimons and Morimae and refined in a paper with Hadjušek. The protocol is summarized in Figure 5.2. In the protocol, the prover is required to prepare a smallest eigenstate of the Hamiltonian $H_\mathcal{C}$ given in (5.3). While it may not be immediately obvious at the level of our description, it is possible to prepare such a "history state" (5.2) by executing a quantum circuit that is only mildly more complex than the original circuit $\mathcal{C}$.

Let $\mathcal{C}$ be a quantum circuit provided as input, and $H_\mathcal{C}$ the n-qubit Hamiltonian obtained from $\mathcal{C}$ as in (5.3).

1. The verifier initializes a counter γ to 0. She executes the following interaction with the prover independently $N = \frac{C}{\delta^2}\binom{n'}{2}^2 \ln(1/\varepsilon)$ times, where C is a large enough universal constant:
   (a) The prover creates an eigenstate $|\psi\rangle$ of H with smallest eigenvalue.
   (b) The prover sends the qubits of $|\psi\rangle$ one by one to the verifier.
   (c) The verifier selects a measurement $W \in \{X, Z\}$ uniformly at random, and measures each qubit in the associated basis upon reception. Let $b_{W,i} \in \{-1, 1\}$ be the outcome for the i-th qubit.
   (d) The verifier selects $i \neq j \in \{1, \ldots, n'\}$ uniformly at random. She updates her counter $\gamma \leftarrow \gamma - J_{ij} b_{W,i} b_{W,j}$.
2. If $\frac{\gamma}{N}\binom{n'}{2} \leq a + \delta/2$ the verifier accepts the interaction. Otherwise, she rejects.

Figure 5.2: The Fitzsimons-Hadjuček-Morimae verification protocol, parametrized by a quantum circuit $\mathcal{C}$ and an accuracy parameter $\varepsilon > 0$.

We note that in the protocol, the verifier measures the qubits in a randomly chosen basis, and then selects a single pair (i, j) such that $J_{ij} \neq 0$ uniformly at random to update her counter. One could imagine small optimizations where e.g. a maximum matching of such pairs is measured at each step. Such optimizations only bring marginal improvements in efficiency of the protocol; moreover they complicate the extension to a classical verifier that we will see later. For this reason, we prefer to keep the simplest expression possible for the protocol.

Theorem 5.9. *Let $\mathcal{C}$ be a quantum circuit and $H_\mathcal{C}$ the Hamiltonian associated to it as in* (5.3)*. Let x be an input to the circuit $\mathcal{C}$ and $\varepsilon > 0$ a parameter for the protocol. Then the following hold:*

- *(Completeness:) If $\mathcal{C}$ accepts x with probability at least $2/3$, then there is a QPT prover that is accepted with probability at least $1 - \varepsilon$*

- *(Soundness:) If $\mathcal{C}$ accepts x with probability at most $1/3$, then any prover is accepted with probability at most ε.*

Note that in the theorem, the soundness statement does not place any computational assumption on the prover.

Proof. The key calculation that underlies the proof is the following.

Claim 5.10. *Let ρ be the density matrix that represents the mixture over the N n'-qubit states sent by the prover in the protocol (in general these states may be entangled). Then the expectation of γ/N is exactly*

$$\mathrm{E}\Big[\frac{\gamma}{N}\Big] \;=\; -\frac{1}{\binom{n'}{2}} \sum_{i\neq j} \frac{J_{ij}}{2} \operatorname{Tr}\big((\sigma_X^i\sigma_X^j + \sigma_Z^i\sigma_Z^j)\rho\big) \;=\; \frac{1}{\binom{n'}{2}} \operatorname{Tr}\big(H\rho\big)\,. \tag{5.4}$$

Moreover, for N chosen as in the protocol for a large enough choice of the constant C it holds that

$$\Pr\left(\Big|\frac{\gamma}{N}\binom{n'}{2} - \operatorname{Tr}\big(H\rho\big)\Big| > \frac{\delta}{2}\right) \;\leq\; \varepsilon\,. \tag{5.5}$$

Proof. For $t \in \{1,\ldots,n\}$ let G_t denote the product of the two outcomes $b_{W,i}$ and $b_{W,j}$ obtained by the verifier at step (c) of the protocol, where W, i and j are as sampled at step (d). Then the random variables $G_t \in \{-1,1\}$ are i.i.d. such that for each t, $\mathrm{E}[G_t] = \operatorname{Tr}(\sigma_W^i\sigma_W^j\rho)$, with W, i and j are the values sampled in step t. Since $\gamma = -\sum_t J_{ij}G_t$, averaging over those choices gives (5.4). Using $|J_{ij}| \leq 1$, by Hoeffding's inequality for any $s > 0$

$$\Pr\big(\big|\gamma - \mathrm{E}[\gamma]\big| > s\big) \;\leq\; e^{-\frac{2s^2}{4N}}\,.$$

By choosing N sufficiently large with respect to $\binom{n'}{2}^2\delta^{-2}\ln(1/\varepsilon)$ we get (5.5). □

Based on Claim 5.10 the proof of Theorem 5.9 follows rather directly. For the completeness, we take $\rho = |\psi\rangle\langle\psi|$ such that $\langle\psi|H|\psi\rangle \leq a$, whose existence is guaranteed by the completeness case of Theorem 5.7. As noted above, this ρ can be prepared efficiently by a QPT prover. Using (5.5) it follows that this prover is accepted with probability at least $1-\varepsilon$. For the soundness, ρ is arbitrary. Using the soundness case of Theorem 5.7 it must be that $\operatorname{Tr}(H\rho) \geq a+\delta$, so that the conclusion follows again from (5.5). □

Even though the verifier's "quantumness" in this protocol is limited — she only needs to hold one qubit at a time — this capability is crucial for the analysis, as it is used to guarantee the "existence" of the state that is being measured: it allows us to meaningfully talk about "the state ρ whose first qubit is the first qubit received by the verifier; whose second qubit is the second qubit received by the verifier; etc.". These qubits are distinct, because the verifier has seen and then discarded them (it would be a different matter if they were returned to the prover). In particular, the fact that a one-qubit computer can be trivially simulated on a classical piece of paper is immaterial to the argument.

With a classical verifier things become substantially more delicate. How can we verify the existence of an n-qubit state with certain properties, while having only access to classical data about the state, data that, for all we know a priori, could have been generated by a simple — classical — laptop? To achieve this we need to find a way for the verifier to establish that the prover holds an n-qubit state, without ever having the ability to directly probe even a single qubit of that state. In the previous lecture we saw a means to achieve this for a single qubit based on the computational hardness of certain functions called "claw-free". In the next lecture we extend that method to introduce a protocol by which the prover can certify the existence

of any single-qubit state that is a low-energy eigenstate of a single-qubit Hamiltonian. In the following lecture we combine this extension with the Fizsimons-Morimae protocol to obtain a protocol for delegating quantum computations with a classical client.

Bibliography

[Aar10] Scott Aaronson. Bqp and the polynomial hierarchy. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 141–150. ACM, 2010.

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

[ACGK17] Scott Aaronson, Alexandru Cojocaru, Alexandru Gheorghiu, and Elham Kashefi. On the implausibility of classical client blind quantum computing. *arXiv preprint arXiv:1704.08482*, 2017.

[AG17] Dorit Aharonov and Ayal Green. A quantum inspired proof of $\mathrm{P}^{\setminus \#\mathrm{p}} \subseteq$ IP. *arXiv preprint arXiv:1710.09078*, 2017.

[AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *Theory of Cryptography Conference*, pages 474–495. Springer, 2009.

[AV12] Dorit Aharonov and Umesh Vazirani. Is quantum mechanics falsifiable? A computational perspective on the foundations of quantum mechanics. *arXiv preprint arXiv:1206.3686*, 2012.

[AV13] Dorit Aharonov and Umesh Vazirani. *Is quantum mechanics falsifiable? A computational perspective on the foundations of quantum mechanics*. Computability: Turing, Gödel, Church, and Beyond. MIT Press, 2013.

[BCM+18] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331. IEEE, 2018.

[Bel64] John S Bell. On the einstein podolsky rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.

[BOGKW19] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pages 373–410. 2019.

[Bre06] Frédéric Brechenmacher. *Histoire du théorème de Jordan de la décomposition matricielle (1870-1930). Formes de représentation et méthodes de décomposition.* PhD thesis, 2006.

[BV14] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on Computing*, 43(2):831–871, 2014.

[Cap15] Valerio Capraro. *Connes' Embedding Conjecture*, pages 73–107. Springer International Publishing, Cham, 2015.

[CCKW19] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. Qfactory: classically-instructed remote secret qubits preparation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 615–645. Springer, 2019.

[CLS17] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1):012202, 2017.

[CM14] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *International Colloquium on Automata, Languages, and Programming*, pages 320–331. Springer, 2014.

[CM16] Toby Cubitt and Ashley Montanaro. Complexity classification of local hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016.

[Con76] Alain Connes. Classification of injective factors cases II_1, II_∞, III_λ, $\lambda \neq 1$. *Annals of Mathematics*, pages 73–115, 1976.

[CR20] Rui Chao and Ben W Reichardt. Quantum dimension test using the uncertainty principle. *arXiv preprint arXiv:2002.12432*, 2020.

[CRSV17] Rui Chao, Ben W Reichardt, Chris Sutherland, and Thomas Vidick. Overlapping qubits. *arXiv preprint arXiv:1701.01062*, 2017.

[CRSV18] Rui Chao, Ben W Reichardt, Chris Sutherland, and Thomas Vidick. Test for a large amount of entanglement, using few measurements. *Quantum*, 2:92, 2018.

[CS17] Andrea Coladangelo and Jalex Stark. Robust self-testing for linear constraint system games. *arXiv preprint arXiv:1709.09267*, 2017.

[CS18] Andrea Coladangelo and Jalex Stark. Unconditional separation of finite and infinite-dimensional quantum correlations. *arXiv preprint arXiv:1804.05116*, 2018.

[DFPR14] Vedran Dunjko, Joseph F Fitzsimons, Christopher Portmann, and Renato Renner. Composable security of delegated quantum computation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 406–425. Springer, 2014.

[EPR35] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical review*, 47(10):777, 1935.

[Fri12] Tobias Fritz. Tsirelson's problem and Kirchberg's conjecture. *Reviews in Mathematical Physics*, 24(05):1250012, 2012.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, volume 9, pages 169–178, 2009.

[GKK19] Alexandru Gheorghiu, Theodoros Kapourniotis, and Elham Kashefi. Verification of quantum computation: An overview of existing approaches. *Theory of computing systems*, 63(4):715–808, 2019.

[GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N Rothblum. Delegating computation: interactive proofs for muggles. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 113–122. ACM, 2008.

[GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 612–621. IEEE, 2017.

[GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 660–670. ACM, 2018.

[GMR85] Shafi Goldwasser, Silvio Micali, and Ronald Rivest. A "paradoxical" solution to the signature problem. In *Advances in Cryptology*, pages 467–467. Springer, 1985.

[GV19] Alexandru Gheorghiu and Thomas Vidick. Computationally-secure and composable remote state preparation. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1024–1033. IEEE, 2019.

[GVW01] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. In *International Colloquium on Automata, Languages, and Programming*, pages 334–345. Springer, 2001.

[GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. *Journal of the ACM (JACM)*, 62(6):45, 2015.

[Ito10] Tsuyoshi Ito. Polynomial-space approximation of no-signaling provers. In *International Colloquium on Automata, Languages, and Programming*, pages 140–151. Springer, 2010.

[Ji13] Zhengfeng Ji. Binary constraint system games and locally commutative reductions. *arXiv preprint arXiv:1310.3794*, 2013.

[JNP+11] Marius Junge, Miguel Navascues, Carlos Palazuelos, David Perez-Garcia, Volkher B Scholz, and Reinhard F Werner. Connes' embedding problem and Tsirelson's problem. *Journal of Mathematical Physics*, 52(1):012102, 2011.

[JNV+20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*= RE. *arXiv preprint arXiv:2001.04383*, 2020.

[JP11] Marius Junge and Carlos Palazuelos. Large violation of bell inequalities with low entanglement. *Communications in Mathematical Physics*, 306(3):695, 2011.

[Kir93] Eberhard Kirchberg. On non-semisplit extensions, tensor products and exactness of group C^*-algebras. *Inventiones mathematicae*, 112(1):449–489, 1993.

[KKMV09] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. Using entanglement in quantum multi-prover interactive proofs. *Computational Complexity*, 18(2):273–307, 2009.

[KMW17] Elham Kashefi, Luka Music, and Petros Wallden. The quantum cut-and-choose technique and quantum two-party computation. *arXiv preprint arXiv:1703.03754*, 2017.

[KRR14] Yael Tauman Kalai, Ran Raz, and Ron D Rothblum. How to delegate computations: the power of no-signaling proofs. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 485–494. ACM, 2014.

[Mah18] Urmila Mahadev. Classical verification of quantum computations. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–267. IEEE, 2018.

[Mer90] N David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical review letters*, 65(27):3373, 1990.

[Mer93] N David Mermin. Hidden variables and the two theorems of john bell. *Reviews of Modern Physics*, 65(3):803, 1993.

[MF16] Tomoyuki Morimae and Joseph F Fitzsimons. Post hoc verification with a single prover. *arXiv preprint arXiv:1603.06046*, 2016.

[Mor18] Tomoyuki Morimae. Blind quantum computing can always be made verifiable. *arXiv preprint arXiv:1803.06624*, 2018.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 700–718. Springer, 2012.

[MYS12] Matthew McKague, Tzyh Haur Yang, and Valerio Scarani. Robust self-testing of the singlet. *Journal of Physics A: Mathematical and Theoretical*, 45(45):455304, 2012.

[NC02] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10(7):073013, 2008.

[NW19] Anand Natarajan and John Wright. Neexp is contained in mip. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518. IEEE, 2019.

[Oza13a] Narutaka Ozawa. About the Connes embedding conjecture. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.

[Oza13b] Narutaka Ozawa. About the connes embedding conjecture. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.

[P$^+$16] Chris Peikert et al. A decade of lattice cryptography. *Foundations and Trends® in Theoretical Computer Science*, 10(4):283–424, 2016.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):34, 2009.

[RRR16] Omer Reingold, Guy N Rothblum, and Ron D Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 49–62. ACM, 2016.

[RT19] Ran Raz and Avishay Tal. Oracle separation of BQP and PH. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 13–23, 2019.

[RUV13] Ben W Reichardt, Falk Unger, and Umesh Vazirani. Classical command of quantum systems. *Nature*, 496(7446):456–460, 2013.

[Slo19] William Slofstra. The set of quantum correlations is not closed. In *Forum of Mathematics, Pi*, volume 7. Cambridge University Press, 2019.

[SW87] Stephen J Summers and Reinhard Werner. Maximal violation of bell's inequalities is generic in quantum field theory. *Communications in Mathematical Physics*, 110(2):247–259, 1987.

[Tsi93] Boris S Tsirelson. Some results and problems on quantum Bell-type inequalities. *Hadronic Journal Supplement*, 8(4):329–345, 1993.

[Unr16] Dominique Unruh. Computationally binding quantum commitments. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 497–527. Springer, 2016.

[Vid19] Thomas Vidick. From operator algebras to complexity theory and back. *Notices of the American Mathematical Society*, 66(10), 2019.

[Vid20] Thomas Vidick. Verifying quantum computations at scale: A cryptographic leash on quantum devices. *Bulletin of the American Mathematical Society*, 57(1):39–76, 2020.

[VN32] J Von Neumann. Mathematische grundlagen der quantenmechanik. 1932.

[VW16] Thomas Vidick and John Watrous. Quantum proofs. *Foundations and Trends® in Theoretical Computer Science*, 11(1-2):1–215, 2016.

[VZ20] Thomas Vidick and Tina Zhang. Classical proofs of quantum knowledge. *arXiv preprint arXiv:2005.01691*, 2020.

[WBMS16] Xingyao Wu, Jean-Daniel Bancal, Matthew McKague, and Valerio Scarani. Device-independent parallel self-testing of two singlets. *Physical Review A*, 93(6):062121, 2016.

[WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 600–611. IEEE, 2017.