Lecture 7

Verification for n qubit Hamiltonians in $XX - ZZ$ form

Moving beyond the case of a single qubit, our goal in this lecture is to generalize the results from Section 6.2 to a procedure for extracting n qubits from a prover, together with a statement that allows us to relate measurements in the standard or Hadamard basis of the extracted qubits to measurements performed by the prover in the protocol. Once this has been put in place we will not be far from a delegation protocol along the lines of the Fitzsimons-Morimae protocol from Section 5.2, but, crucially, with classical verifier and communication.

7.1 Setup

To set the stage we first give a straightforward generalization of the single-qubit verification protocol from Figure 6.1 to the case of n qubits. Recall from Section 5.2.2 that for our purposes it suffices to consider Hamiltonians that take the form (5.3). This form allows us to restrict our attention to a collection of measurement outcomes on an n-qubit state such that all qubits are measured in the same basis, computational or Hadamard. In particular we do not need to consider "mixed" measurements, with some qubits measured in the standard basis and other qubits in the Hadamard basis, because (5.3) does not have mixed terms such as $\sigma_{X,i}\sigma_{Z,j}$. (See Remark 7.1 regarding extensions to the mixed case.) It is then natural to request that the honest prover behaves exactly as in the single-qubit verification protocol, except that each action should be repeated independently for each of the n qubits: in the first phase the verifier sends the information for n functions $f_{pk_1}, \ldots, f_{pk_n}$, the prover executes the encoding procedure from the proof of Lemma 6.8 independently for each of the n qubits of its claimed low-energy eigenstate $|\varphi\rangle$ of $H_{\mathcal{C}}$ and reports the n images $y_1, \ldots, y_n$ obtained; in the second phase the verifier sends a single-bit challenge $c \in \{0,1\}$ and the prover measures all its qubits in the computational or Hadamard basis and returns the outcomes $x_1, \ldots, x_n$ or $d_1, \ldots, d_n$ respectively. Note that the only part that is not repeated is the challenge, which is identical for each of the n concurrent repetitions. The reason that we can restrict ourselves to such challenges is due to the form of $H_{\mathcal{C}}$ from (5.3) and, as we will see, greatly simplifies the analysis. The complete protocol is given in Figure 7.1.

Before proceeding to the analysis of the protocol we examine the question, "Where are the qubits?" For the single-qubit verification protocol our initial intuition came from the qubit computational test from lecture 4, for which we were able to argue that the prover indeed has a qubit $(|\psi\rangle, Z, X)$. For the verification protocol seen in the last lecture we saw that in order to allow verification of other states than the $|+\rangle$ state

Let $\mathcal{F}$ be a 2-to-1 trapdoor claw-free function family and $\lambda \in \mathbb{N}$ a security parameter. Let $\varepsilon, \delta > 0$ be accuracy parameters. Let $\gamma = 0$ and $N = \frac{C}{\delta^2}\binom{n}{2}^2 \ln(1/\varepsilon)$. The verifier and prover repeat the following interaction N times.

1. The verifier selects a pair $i \neq j \in \{1, \ldots, n\}$ and $W \in \{X, Z\}$ uniformly at random.
2. For $\ell = 1, \ldots, n$ the verifier generates $(pk_\ell, td_\ell) \leftarrow \textsc{Gen}(1^\lambda)$. It sends $(pk_1, \ldots, pk_n)$ to the prover.
3. The prover returns $y_1, \ldots, y_n \in \{0,1\}^m$, where $m = m(\lambda)$.
4. The verifier selects a uniformly random challenge $c \leftarrow_R \{0,1\}$ and sends c to the prover.
5. (a) *(Computational basis, $c = 0$:)* In case $c = 0$ the prover is expected to return $x_1, \ldots, x_n \in \{0,1\}^m$. If $f_{pk_\ell}(x_\ell) \neq 0$ for any ℓ then the verifier immediately aborts. The verifier sets

$$\gamma \leftarrow \gamma - J_{ij}\,(-1)^{b(x_i)}(-1)^{b(x_j)}\,.$$

   (b) *(Hadamard basis, $c = 1$:)* In case $c = 1$ the prover is expected to return $d_1, \ldots, d_n \in \{0,1\}^m$. The verifier uses td_i and td_j to determine the preimages $(x_{i,0}, x_{i,1})$ of y_i by f_{pk_i} and $(x_{j,0}, x_{j,1})$ of y_j by f_{pk_j} respectively. She sets

$$\gamma \leftarrow \gamma - J_{ij}\,(-1)^{d_i \cdot (x_{i,0} + x_{i,1})}(-1)^{d_j \cdot (x_{j,0} + x_{j,1})}\,.$$

If the verifier has not aborted at any of the steps $c = 0$, she returns the real number $o = \frac{1}{N}\binom{n}{2}\gamma$.

Figure 7.1: Verification protocol $\mathfrak{V}_n$ for an n-qubit Hamiltonian $H_{\mathcal{C}} = -\sum_{i,j} \frac{J_{ij}}{2}(\sigma_{X,i}\sigma_{X,j} + \sigma_{Z,i}\sigma_{Z,j})$.

we had to remove some of the tests done by the verifier (specifically, the equation check) and that due to this we were no longer able to guarantee a qubit in the sense of Definition 2.2. Nevertheless we were able to get around this by defining an abstract *extracted qubit* that did not directly correspond to the prover's observables but was still such that measurement outcomes on the extracted qubit could be shown to have a distribution that is negligibly close to outcomes obtained from the prover in the actual protocol (Lemma 6.6).

For the case of demonstrating n qubits a priori one would have to show that the prover has a state $|\psi\rangle$ and two families of observables $\{X(a) : a \in \{0,1\}^n\}$ and $\{Z(b) : b \in \{0,1\}^n\}$ that satisfy the Pauli commutation and anti-commutation relations when they act on $|\psi\rangle$. Indeed, a straightforward generalization of Lemma 2.3 then guarantees the existence of a suitable isometry with the space of n actual qubits. Showing this is challenging; luckily, for our purposes it is also not necessary.[1] Indeed, just as in the single-qubit case it is worth emphasizing that in the context of verification we do not need to guarantee that the prover has a certain quantum state, nor that it is able to perform certain measurements on it. The only real requirement is that a state $|\varphi\rangle$ *exists* such that $\langle\varphi|H_C|\varphi\rangle \leq a$. Thus, as we did in the analysis of the single-qubit verification protocol we will first introduce a abstract *extracted n qubit* defined from the prover's state and actions in the protocol but that also include additional ingredients that make it at first unclear how they relate to the prover itself. The definition of the extracted qubits is given in Section 7.2. Once this has been defined we will perform the second, crucial step, which is to relate the distribution of measurement outcomes on the extracted qubits to quantities that are directly observable in the protocol. This is done in Section 7.3. Finally in Section 7.4 we put everything together and show the completeness and soundness properties of the verification protocol given in Figure 7.1. In addition in Section 7.5 we will sketch a construction of a function family based on the Learning With Errors (LWE) problem that (approximately) satisfies all required assumptions and can thus be used to instantiate the protocol.

7.2 The n extracted qubits

7.2.1 Modeling the prover

We start by introducing notation that allows us to model an arbitrary prover in the protocol. Similarly to how we modeled the prover for the analysis of the computational qubit test in Section 4.3, a prover in the n-qubit verification protocol from Figure 7.1 can be represented using the following objects:

1. A state $|\psi\rangle$, that may depend on $pk_1, \ldots, pk_n$ and $y_1, \ldots, y_n$, such that $|\psi\rangle \in \mathcal{H}_{\mathsf{X}_1} \otimes \cdots \mathcal{H}_{\mathsf{X}_n} \otimes \mathcal{H}_{\mathsf{P}}$ with each space $\mathcal{H}_{\mathsf{X}_i}$ isomorphic to $(\mathbb{C}^2)^{\otimes m}$. The state $|\psi\rangle$ represents the state of the prover and the message registers at the end of step 3 in the protocol.

2. For the case $c = 0$, the prover directly measures all the X registers in the standard basis to obtain $x_1, \ldots, x_n$ that it returns to the verifier. For a string $a \in \{0,1\}^n$ we let

$$Z(a) = \sum_{x_1,\ldots,x_n} (-1)^{a_1 \cdot b_1(x_1)} \cdots (-1)^{a_n \cdot b_n(x_n)} |x_1\rangle\langle x_1| \otimes \cdots |x_n\rangle\langle x_n| \,, \tag{7.1}$$

   where the functions b_i are not necessarily all equal since they may depend on pk_i. This is analogous to (4.3).

[1] In the last three lectures of the course we will see how in the context of spatial assumptions it is known how to test n qubits in this sense; for computational assumptions we do not yet know how to do it.

3. For the case $c = 1$, the prover applies an arbitrary unitary U followed by a measurement of the qubits in X in the Hadamard basis to obtain $d_1, \ldots, d_n$. For a string $b \in \{0,1\}^n$ we let

$$X(b) = \sum_{d_1,\ldots,d_n} (-1)^{b_1(d_1\cdot(x_{1,0}+x_{1,1}))} \cdots (-1)^{b_n(d_n\cdot(x_{n,0}+x_{n,1}))} \cdot U^\dagger \big(H_{\mathsf{X}}^{\otimes nm} \otimes \mathrm{Id}_{\mathsf{P}}\big)^\dagger \big(|d_1,\ldots,d_n\rangle\langle d_1,\ldots,d_n|_{\mathsf{X}} \otimes \mathrm{Id}_{\mathsf{P}}\big)\big(H_{\mathsf{X}}^{\otimes nm} \otimes \mathrm{Id}_{\mathsf{P}}\big) U \,. \quad (7.2)$$

Remark 7.1. Note that the fact that the protocol only has two different challenges, $c = 0$ and $c = 1$, allows us to have a simple description for all $Z(a)$ and all $X(b)$ observables that involves only one "adversarial" unitary U. If we had to design a protocol that allows more general Hamiltonians with mixed terms of the form $\sigma_{X,i}\sigma_{Z,j}$ we would need to consider more challenges, and this would require a more complex analysis. This is done in [Mah18].

7.2.2 The isometry V

Next we define the n-qubit isometry V, and the extracted qubits.

Claim 7.2. *Let $|\psi\rangle \in \mathcal{H}$ and for every $a, b \in \{0,1\}^n$, $X(a)$ and $Z(b)$ observables on $\mathcal{H}$ such that all $X(a)$ (resp. all $Z(b)$) mutually commute and moreover $X(a)X(a') = X(a+a')$ for any $a, a' \in \{0,1\}^n$. Let $V : \mathcal{H} \to \mathcal{H}_{\mathsf{Q}} \otimes \mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}'$ where each of $\mathcal{H}_{\mathsf{Q}}$ and $\mathcal{H}_{\mathsf{A}}$ is $(\mathbb{C}^2)^{\otimes n}$ and $\mathcal{H}' \simeq \mathcal{H}$ be defined for all $|\varphi\rangle \in \mathcal{H}$ as*

$$V|\varphi\rangle \;=\; \Big(\frac{1}{2^n}\sum_{a,b} \mathrm{Id} \otimes \sigma_X(a)\sigma_Z(b) \otimes X(a)Z(b)\Big)|\phi^+\rangle^{\otimes n}|\varphi\rangle \;, \quad (7.3)$$

where each EPR pair $|\phi^+\rangle$ has one qubit in register Q and the other in register A and the σ_X and σ_Z operators act on register A. Then V is an isometry.

The proof of the claim is immediate and only uses that the family of states

$$\big\{(\sigma_X(a)\sigma_Z(b) \otimes \mathrm{Id})|\phi^+\rangle^{\otimes n} : \, a, b \in \{0,1\}^n\big\}$$

is orthonormal. Similarly to Definition 6.5 we can now define the n extracted qubits.

Definition 7.3 (Extracted qubits)**.** Let P be a prover in the verification protocol $\mathfrak{V}_n$ described in Figure 7.1. Let $|\psi\rangle$ be the state of P after having sent $y_1, \ldots, y_n$ at step 3 of the t-th iteration, for some $t \in \{1, \ldots, N\}$. Let V be defined in (7.3). Then we call the reduced density of $V|\psi\rangle$ on register Q the *extracted qubits* (implicitly, at iteration t) and denote them by $\rho_{\mathsf{Q}_1\cdots\mathsf{Q}_n}$.

7.3 Measurements on the extracted qubits

We start with the following analogue to Claim 6.4, which gives an explicit formula for the distribution of measurements in the standard or Hadamard basis on the n extracted qubits as a function of the prover's state and observables.

Claim 7.4. *The following hold for any prover, with ρ the n extracted qubits at any iteration (Definition 7.3):*

$$\forall b \in \{0,1\}^n \,, \qquad \mathrm{Tr}\big(\sigma_Z(b)\rho\big) = \langle\psi|Z(b)|\psi\rangle \,, \quad (7.4)$$

$$\forall a \in \{0,1\}^n \,, \qquad Tr\big(\sigma_X(a)\rho\big) = \frac{1}{2^n}\sum_b (-1)^{a\cdot b}\langle\psi|Z(b)X(a)Z(b)|\psi\rangle \,. \quad (7.5)$$

The claim can be illustrated using the following generalization of (6.5)

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{V} & \mathbb{C}^2 \otimes \mathcal{H}' \\ {\scriptstyle Z(b)} \Big\downarrow & & \Big\downarrow {\scriptstyle \sigma_Z(b)\otimes \mathrm{Id}} \\ {\scriptstyle \mathrm{E}_{b\in\{0,1\}}(-1)^{b\cdot a} Z(b)X(a)Z(b)} \Big\downarrow & & \Big\downarrow {\scriptstyle \sigma_X(a)\otimes \mathrm{Id}} \\ \mathcal{H} & \xrightarrow{V} & \mathbb{C}^2 \otimes \mathcal{H}' \end{array} \tag{7.6}$$

Proof. Eq. (7.4) is immediate using that $X(a)$ are observables and $\langle \phi^+|^{\otimes n} \sigma_X(a')\sigma_Z(b') \otimes \sigma_Z(b) |\phi^+\rangle^{\otimes n}$ is zero unless $a' = 0$ and $b = b'$. Eq. (7.5) is shown similarly by direct calculation, using $X(a')X(a'') = X(a' + a'')$ and $\sigma_Z(b)\sigma_X(a)\sigma_Z(b) = (-1)^{a\cdot b}\sigma_X(a)$. □

The next lemma is the key lemma. It argues that for computationally bounded provers, the quantity on the right-hand side of (7.5) is close to the simpler quantity $\langle \psi | X(a) | \psi \rangle$, that in particular can be inferred in the protocol from the prover's outcomes y_i and d_i (for those i such that $a_i = 1$). Before we can state the lemma we need to introduce one last assumption on the function family $\mathcal{F}$. Intuitively, this assumption is a natural quantum analogue of the classical property of collision resistance, but is stronger than it.

(**F.5**) Consider the following abstract game between an arbitrary "adversary" (think prover) and a trusted (quantum) "challenger" (think verifier). First, the adversary is provided a label pk (generated at random by the challenger) and required to prepare an arbitrary state of the form $|\phi\rangle = \sum_x \alpha_x |x\rangle$, where x ranges over the domain of f_{pk}. (In general the adversary may keep an additional register entangled with this state. For ease of notation we do not consider such entanglement in this description.) The adversary hands the state $|\phi\rangle$ over to the challenger, who evaluates f_{pk} in superposition on $|\phi\rangle$ and measures the image register, obtaining a y in the range of f_{pk} and the (suitably re-normalized) post-measurement state $|\phi'\rangle = \sum_{x:f_{pk}(x)=y} \alpha_x |x\rangle$. The challenger then returns to the adversary the string c together with *either* the state $|\phi'\rangle$ *or* the probabilistic mixture $\sum_{x:f(x)=c} |\alpha_x|^2 |x\rangle\langle x|$ obtained by measuring the same state $|\phi'\rangle$ in the computational basis (and throwing away the outcome). The adversary wins if it correctly guesses which is the case. Assumption (**F.5**) on the function family $\mathcal{F}$ states that for any QPT adversary $\mathcal{A}$ there is a negligible function μ such that for any λ, $\mathcal{A}$ succeeds in this game with probability that deviates from $\frac{1}{2}$ by at most $\mu(\lambda)$.

Remark 7.5. Assumption (**F.5**) is referred to as the "collapsing" property for the function family $\mathcal{F}$. This property was introduced by Unruh as a strengthening of the classical property of collision resistance required for his work on the security of commitment protocols that are computationally binding against quantum adversaries [Unr16]. The reason that this assumption implies collision resistance is that, if the function were not collision resistant, the adversary could identify a colliding pair (x_0, x_1) and submit $|\phi\rangle = \frac{1}{\sqrt{2}}(|x_0\rangle + |x_1\rangle)$ to the challenger. It could then measure the challenger's response in a basis containing the two states $\frac{1}{\sqrt{2}}(|x_0\rangle \pm |x_1\rangle)$ and guess that, in case the "$-$" outcome is obtained, the challenger must have measured; in the other case, the adversary guesses at random.

Note that assumption (**F.2**) *also* trivially implies collision resistance, since the ability to identify a claw allows one to generate arbitrary equations in it. It is possible to show that both (**F.2**) and (**F.5**) are strictly stronger than collision resistance. It is likely that the two assumptions are incomparable, but I have not tried to show this explicitly.

We can now state and prove the key lemma.

Lemma 7.6. *Let P be a prover that succeeds with probability 1 in protocol $\mathfrak{V}_n$. Let ρ be the n extracted qubits, as defined in Definition 7.3 (for any iteration). Then the following hold for any $i \neq j \in \{1,\ldots,n\}$:*

- *(Z-measurement:) The outcome of measuring qubits i and j of ρ in the computational basis is identically distributed to the bits $b(x_i)$ and $b(x_j)$ obtained from the prover in case $c = 0$.*
- *(X-measurement:) Under assumptions (**F.2**) and (**F.5**) the outcome of measuring qubits i and j of ρ in the Hadamard basis is* computationally indistinguishable *from the pair of bits $d_i \cdot (x_{i,0} + x_{i,1})$ and $d_j \cdot (x_{j,0} + x_{j,1})$ where d_i and d_j are obtained from the prover in case $c = 1$.*

As already noted in the previous lecture, for distributions on two bits the notions of computational and statistical indistinguishability are essentially equivalent. The lemma generalizes to the joint distribution of any number of bits, and in this case it is only the weaker computational indistiguishability that is obtained. For simplicity we restrict ourselves to proving the lemma for the setting of two bis only.

Proof. For the case of a measurement in the computational basis the lemma follows directly from (7.4) in Claim 7.4 and the definition of the extracted qubits. For the case of a measurement in the Hadamard basis we proceed in two steps.

In the first step we show that for any $b \in \{0,1\}^n$ the states $|\psi\rangle$ and $Z(b)|\psi\rangle$ are computationally indistinguishable. We show this by performing a reduction to an adversary that breaks assumption (**F.5**). Fix any $i \in \{1,\ldots n\}$ and suppose for contradiction that there exists an efficient observable R such that

$$\left|\langle\psi|R|\psi\rangle - \langle\psi|Z(e_i)RZ(e_i)|\psi\rangle\right| \;>\; \frac{1}{q(\lambda)}\,,$$

for some polynomial q and where the left-hand side should be understood on expectation over the creation of a state $|\psi\rangle$ according to the first three steps of protocol $\mathfrak{V}_n$. Let i be a position in which $b_i \neq 0$. Our goal is to reach a contradiction with (**F.5**). Towards this we construct an adversary $\mathcal{A}$ to the collapsing game that underlies assumption (**F.5**). Upon input pk, $\mathcal{A}$ creates the state $|\psi\rangle$ and returns to the challenger only the m-qubit register X_i. Note that the first part of the challenger's actions in the game does not change $|\psi\rangle$, since the prover has already collapsed it to a pair of preimages. The two cases correspond to the challenger returning either the mixed state $\sum_{b\in\{0,1\}} Z_{b,i}|\psi\rangle\langle\psi|Z_{b,i}$ or $|\psi\rangle\langle\psi|$, where $Z_{b,i} = (\mathrm{Id} + (-1)^b Z(e_i))/2$. The adversary $\mathcal{A}$ measures R and returns the outcome. The advantage of $\mathcal{A}$ in distinguishing the two cases is

$$\left|\langle\psi|R|\psi\rangle - \sum_b \langle\psi|Z_{b,i}RZ_{b,i}|\psi\rangle\right| \;=\; \frac{1}{2}\left|\langle\psi|R|\psi\rangle - \langle\psi|Z(e_i)RZ(e_i)|\psi\rangle\right|\,,$$

where the equality follows by definition of $Z_{b,i}$. Since by (**F.5**) this advantage should be negligible, we deduce that for every $i \in \{1,\ldots,n\}$ and every efficient observable R it must be that

$$\langle\psi|R|\psi\rangle \approx \langle\psi|Z(e_i)RZ(e_i)|\psi\rangle\,. \tag{7.7}$$

Since for any b the observable $Z(b)RZ(b)$ itself is efficient, applying (7.7) n times (with different choices of R) we deduce that for any b and efficient R,

$$\begin{aligned}\langle\psi|R|\psi\rangle &\approx \langle\psi|Z(b_1e_1)RZ(b_1e_1)|\psi\rangle\\ &\approx \langle\psi|Z(b_1e_1+b_2e_2)RZ(b_1e_1+b_2e_2)|\psi\rangle\\ &\approx \cdots\\ &\approx \langle\psi|Z(b)RZ(b)|\psi\rangle\,.\end{aligned}$$

We now extend the preceding reasoning to show that for any a of the form $a = a_i e_i + a_j e_j$ with e_i, e_j the canonical basis vectors and $a_i, a_j \in \{0,1\}$,

$$\Big| \frac{1}{2^n} \sum_b (-1)^{a\cdot b} \langle \psi | Z(b) X(a) Z(b) | \psi \rangle \\ - \frac{1}{4} \sum_{b_i, b_j \in \{0,1\}} (-1)^{a_i b_i + a_j b_j} \langle \psi | Z(b_i e_i + b_j e_j) X(a) Z(b_i e_i + b_j e_j) | \psi \rangle \Big| \leq \mu(\lambda) \,, \tag{7.8}$$

for some negligible function ν. Supposing this were not the case, by the triangle inequality and an averaging argument there must exist a b such that

$$\big| \langle \psi | Z(b) X(a) Z(b) | \psi \rangle - \langle \psi | Z(b_i e_i + b_j e_j) X(a) Z(b_i e_i + b_j e_j) | \psi \rangle \big| > \frac{1}{q(\lambda)} \,,$$

for some polynomial q. This leads to a contradiction with (**F.5**) using the same reasoning as before, because from the point of view of the statement of (**F.5**) for qubits not in positions i and j, the observable $X(a)$ is efficient, as its computation only requires trapdoors td_i and td_j.

To obtain the second part of the claim it remains to handle the $Z(e_i)$ and $Z(e_j)$ operators. For the positions i and j the associated trapdoor information *is* used in the computation of $X(a)$, so the preceding reasoning cannot be applied. Instead, we proceed similarly to the proof of Lemma 6.6, by reduction to the adaptive hardcore bit property, assumption (**F.2**). Note that if $a_i = 0$ or $a_j = 0$ then our task is exactly the task handled in Lemma 6.6. So assume $a_i = a_j = 1$. We perform a reduction to Lemma 6.6 via a simple hybrid argument. Suppose for the sake of contradiction that

$$\Big| \langle \psi | X(a) | \psi \rangle - \frac{1}{4} \sum_{b_i, b_j \in \{0,1\}} (-1)^{a_i b_i} \langle \psi | Z(b_i e_i + b_j e_j) X(a) Z(b_i e_i + b_j e_j) | \psi \rangle \Big| > \frac{1}{q(\lambda)} \,.$$

Then by the triangle inequality and averaging it must be that either

$$\Big| \langle \psi | X(a) | \psi \rangle - \frac{1}{2} \sum_{b_i \in \{0,1\}} (-1)^{a_i b_i + a_j b_j} \langle \psi | Z(b_i e_i) X(a) Z(b_i e_i) | \psi \rangle \Big| > \frac{1}{q(\lambda)} \,,$$

or

$$\Big| \langle \psi | Z(b_j) X(a) Z(b_j) | \psi \rangle - \frac{1}{2} \sum_{b_i \in \{0,1\}} (-1)^{a_i b_i} \langle \psi | Z(b_i e_i + e_j) X(a) Z(b_i e_i + e_j) | \psi \rangle \Big| > \frac{1}{q(\lambda)} \,.$$

The first case is ruled out directly by Lemma 6.6. The second case is ruled out by the same lemma, simply considering a prover that creates the state $Z(b_j)|\psi\rangle$ instead of $|\psi\rangle$ at the 3rd step (i.e. the step where $|\psi\rangle$ is defined). □

7.4 An n-qubit verification protocol

The following theorem is the main result of the past four lectures. It generalizes Proposition 6.9 to the case of an n-qubit Hamiltonian.

Theorem 7.7. *Let $\mathcal{F}$ be a function family satisfying (**F.1**), (**F.2**), (**F.3**), (**F.4'**) and (**F.5**). Let $H_{\mathcal{C}}$ be an n-qubit Hamiltonian of the form* (5.3) *and $\delta, \varepsilon > 0$ accuracy parameters. Then the verification protocol from Figure 7.1 has the following properties:*

1. *(Completeness:) For any n-qubit state $|\varphi\rangle$, there is a QPT prover that is accepted with probability* 1 *in the protocol and such that the value o returned by the verifier at the end of the protocol satisfies* $\mathrm{E}[o] = \langle\varphi|H|\varphi\rangle$.

2. *(Soundness:) For any QPT prover that is accepted with probability* 1 *in the protocol, there is an n-qubit state ρ such that the value o returned by the verifier at the end of the protocol satisfies* $\Pr(|o - \mathrm{Tr}(H\rho)| > \delta) \leq \varepsilon$.

Remark 7.8. The protocol in Figure 7.1, as the one in Figure 6.1, involves N repetitions of an elementary 4-message procedure. It is possible to parallelize the protocol to a single repetition in which the prover is asked to perform measurements on all N qubits of a ground state of $H_{\mathcal{C}}$. This however requires more work, because in the parallelized protocol the verifier needs to request "mixed" measurements from the prover; see Remark 7.1.

Remark 7.9. We pause to insist on how amazing Theorem 7.7 is. Due to Kitaev's circuit-to-Hamiltonian construction (Section 5.2.1) it is known that, under the widely believed assumption that QMA $\neq$ QCMA (where QCMA is the class of languages that admit classical proofs verifiable by QPT verifiers), there exist families of Hamilitonians of the form $H_{\mathcal{C}}$ such that any sufficiently low-energy eigenstate of $H_{\mathcal{C}}$ cannot have a simple classical description; in particular, there is no small quantum circuit to prepare such eigenstates, they must have high entanglement, etc. Yet Theorem 7.7 states that through an efficient classical interaction with a device that has the ability to prepare such states it is possible to *efficiently* verify their *existence*. There are two ways in which one might aim to strengthen that statement. First, in the spirit of "proofs of knowledge" we might aim to show that the prover *has* such a state, and not only that it *exists*. Showing this requires a formalization of the notion of the prover "having" a certain quantum state, but it can be done without any modification to the protocol itself; see [VZ20]. Second, in the spirit of our "test for a qubit" we might aim to show that the prover *has n qubits*. This we do not know how to show in the computational setting: it is an open question. (See the full notes at http://users.cms.caltech.edu/ vidick/teaching/fsmp/fsmp.pdf for how to achieve this broader goal under a different assumption, that of spatial assumption between two provers.)

Remark 7.10. The assumption that the prover succeeds with probability 1 that is made in the soundness statement is not difficult to relax; see Remark 6.10.

Proof. The completeness statement is entirely analogous to the same statement for Proposition 6.9. In slightly more detail, at each of the N iterations the honest prover prepares a fresh copy of the state $|\varphi\rangle$ and then applies the procedure described in the proof of Lemma 6.8 independently to each of the n qubits of $|\varphi\rangle$, using the key pk_i for the i-th qubit and obtaining an outcome y_i. For each qubit the post-measurement state is in an m-qubit register $\mathcal{X}_i$ that the prover measures in the standard basis in case of challenge $c = 0$, and Hadamard basis in case $c = 1$. It can then be verified by direct calculation that in case $c = 0$ for any pair $i \neq j$ the parity $(-1)^{b(x_i)+b(x_j)}$ is distributed as a measurement of $\sigma_Z(e_i + e_j)$ on $|\varphi\rangle$, and similarly in case $c = 1$ for any pair $i \neq j$ the parity $(-1)^{d_i\cdot(x_{i,0}+x_{i,1})+d_j\cdot(x_{j,0}+x_{j,1})}$ is distributed as a measurement of $\sigma_X(e_i + e_j)$ on $|\varphi\rangle$.

For soundness we use Lemma 7.6. The lemma shows that for any iteration $t = 1, \ldots, N$ in the protocol we can define a state ρ_t such that averaging over the verifier's choice of qubits i and j it holds that, whenever $c = 0$ then

$$\mathrm{E}\left[J_{ij}\,(-1)^{b(x_i)}(-1)^{b(x_j)}\right] \;=\; J_{ij}\,\mathrm{Tr}\big(\sigma_{Z,i}\sigma_{Z,j}\rho_t\big)\;.$$

and whenever $c = 1$ then

$$\mathrm{E}\left[J_{ij}\,(-1)^{d_i\cdot(x_{i,0}+x_{i,1})}(-1)^{d_j\cdot(x_{j,0}+x_{j,1})}\right] \;\approx\; J_{ij}\,\mathrm{Tr}\big(\sigma_{X,i}\sigma_{X,j}\rho_t\big)\;,$$

where the approximation is up to some negligible quantity in λ. Averaging these two quantities we see that on average over all the rounds,

$$\begin{aligned}\mathrm{E}[o] &\approx \frac{1}{N}\sum_{t=1}^{N}\sum_{i\neq j}\Big(-\frac{1}{2}J_{ij}\,\mathrm{Tr}\big(\sigma_{Z,i}\sigma_{Z,j}\rho_t\big)-\frac{1}{2}J_{ij}\,\mathrm{Tr}\big(\sigma_{X,i}\sigma_{X,j}\rho_t\big)\Big)\\ &= \frac{1}{N}\sum_{t=1}^{N}\mathrm{Tr}\big(H_{\mathcal{C}}\rho_t\big)\\ &= \mathrm{Tr}\big(H_{\mathcal{C}}\rho\big)\ ,\end{aligned}$$

where we defined $\rho = \frac{1}{N}\sum_t \rho_t$. The more quantitative statement given in the soundness part of the theorem follows directly by using a martingale concentration argument, provided the constant C in the definition of N is chosen large enough. □

7.5 Construction of a claw-free function family $\mathcal{F}$

The presentation of this section is adapted from [Vid20].

In Section 4.3 we have identified four assumptions (we added a fifth one in Section 7.3) on a family of functions $\{f_{pk(\lambda)} : \{0,1\}^{m(\lambda)} \to \{0,1\}^{m(\lambda)}\}_{\lambda\in\mathbb{N}}$, such that the five assumptions together are sufficient for the resulting delegated computation protocol to be sound. Can the five assumptions be simultaneously satisfied? Strictly speaking, we do not know the answer. In this section we sketch a construction that *nearly* satisfies the assumptions. The construction appears in [BCM+18], and a mild modification of it is used in Mahadev's protocol. Even though the desired assumptions will not all be strictly satisfied by the construction,[2] it is possible to verify that the protocol itself remains sound.

7.5.1 The LWE problem

Our starting point is the *Learning with Errors* problem, introduced by Regev [Reg09]. The hardness of this problem has become a widely used computational assumption in cryptography, for at least three reasons. The first is that it is very versatile, allowing the implementation of advanced primitives such as fully homomorphic encryption [Gen09, BV14], attribute-based encryption [GVW15], program obfuscation [WZ17, GKW17], traitor tracing [GKW18], and many others. The second is that the assumption can be reduced to the hardness of *worst-case* computational problems on lattices: an efficient procedure that breaks the LWE assumption *on average* can be used to solve the closest vector problem in (almost) any lattice. The third reason, that is most relevant to the use of the LWE assumption made here, is that in contrast to the RSA assumption on the hardness of factoring or the discrete logarithm problem so far it is believed that the LWE problem may be hard for quantum computers, so that cryptographic schemes based on it remain (to the best of published knowledge) secure against quantum attacks.

The LWE assumption comes in multiple flavors, all roughly equivalent. Here we formulate the *decisional LWE* assumption on the difficulty of distinguishing samples from two distributions. To state the problem, fix a size parameter $n \geq 1$, an integer modulus $q \geq 2$, a number of equations $m \geq n\log q$, and an

[2] In particular, we construct functions from $\mathbb{Z}_q^m$ to Z_q^m for some q that is required to be large and may not necessarily be chosen even. The definition of assumption (**F.2**) considers equations modulo 2, and this is naturally tailored to the capabilities of a quantum prover, for whom it is possible to generate such equations by measuring in the Hadamard basis. The family of functions constructed in this section can be shown to possess the hardcore bit property over $\mathbb{Z}_q$, but proving it over $\mathbb{Z}_2$ requires more work.

error distribution χ over $\mathbb{Z}_q$.[3] Given χ, write χ^m for the distribution over $\mathbb{Z}_q^m$ that is obtained by sampling each entry of a vector independently according to χ. The decisional LWE assumption is the following.

> *(Decisional LWE, informal)* Let A be a uniformly random matrix in $\mathbb{Z}_q^{m\times n}$, s a uniformly random vector in $\{0,1\}^n$, e a random vector in $\mathbb{Z}_q^m$ drawn from χ^m, and r a uniformly random vector in $\mathbb{Z}_q^m$. Then no classical or quantum probabilistic polynomial-time procedure can distinguish $(A, As+e)$ from (A, r).

Note that the distribution of $(A, As+e)$ and the distribution of (A, r) are in general very far from each other: provided m is sufficiently larger than n a random vector r will not lie in the column span of A, nor even be close to it. What the (decisional) LWE assumption asserts is that, even though in principle these distributions are far from each other, it is computationally difficult, given a sample from the one or the other, to tell which is the case. Note that without the error vector e the task would be easy: given (A, y), solve for $As = y$ and check whether the solution has coefficients in $\{0,1\}$. The LWE assumption is that the inclusion of e makes the task substantially more arduous. In particular, it is well-known that Gaussian elimination is very sensitive to errors, which rules out the most natural approach.

The definition we gave is informal because we have not specified how the parameters n, m and q should be chosen as a function of the security parameter λ, and we have not specified the distribution χ. In general one can make the decisional LWE assumption for any choice of these parameters—but for some choices the assumption will be invalidated by existing algorithms. We comment on come choices of parameters that are made in cryptography. The integer n should generally be thought of as commensurate with the security parameter λ, i.e. $n = \Theta(\lambda)$. The modulus q should be at least polynomial in n, but can be as large as exponential; this will be the case in our construction. The error distribution χ can be chosen in multiple ways. A common choice is to set χ a discretized centered Gaussian distribution with variance αq, for some small parameter α (typically chosen as an inverse polynomial function of n); this is generally denoted $D_{\mathbb{Z}_q,\alpha q}$. For more details on LWE and its applications, we refer to the survey [P+16].

7.5.2 Construction

To specify the function family $\mathcal{F}$ we first describe how public and private parameters for the function are chosen. Let λ be the security parameter (i.e. the number 2^λ is thought of as an estimate of the time required to break assumptions such as (**F.2**)).

First, integers n, m and a modulus q are chosen such that $n = \Omega(\lambda)$, $q \geq 2$ is a prime, and $m = \Omega(n \log q)$. Then, a matrix $A \in \mathbb{Z}_q^{m\times n}$ is sampled at random, together with a "trapdoor" in the form of a matrix $R \in \mathbb{Z}_q^{\ell\times m}$, where $n \leq \ell \leq m$ is a parameter. The sampling procedure has the property that the distribution of A is statistically close to uniform, and R is such that $G = RA \in \mathbb{Z}_q^{\ell\times n}$ is a "nice" matrix, in the sense that given $b = Gs + e$, for any $s \in \mathbb{Z}_q^n$ and e small enough, it is computationally easy to recover s.[4] That such a sampling procedure would exist and be efficiently implementable is non-trivial, and relies on the underlying lattice structure given by the columns of A; see [MP12]. Finally, a uniformly random $s \in \{0,1\}^n$, and a random $e \in \mathbb{Z}_q^m$ distributed according to $D_{\mathbb{Z}_q,\alpha q}$ with α of order $1/(\sqrt{mn\log q})$,[5]

[3] The use of the parameters n, m and q is local to this section. In particular, the m that specifies the domain and range of the function f_{pk} is not identical to the m here; see below.

[4] One can think of G as a matrix whose rows are almost orthonormal, so that Gaussian elimination on G induces only small propagation of the errors.

[5] The precise choice of α is delicate, and the parameters given here should only be treated as indicative; we refer to [BCM+18, Section 8] for the right setting of parameters.

are sampled. The public information is $pk = (A, z = As + e)$. The trapdoor information is the pair $td = (R, s)$. Note that pk is not uniformly distributed, but pairs (pk, td) can be sampled in randomized polynomial time in λ.

Next we discuss how the function $f = f_{pk}$ can be evaluated, given the public parameters $pk = (A, z)$. We define two functions f_0, f_1 that should be understood as $f(0\|\cdot)$ and $f(1\|\cdot)$ respectively. Each function goes from $\mathbb{Z}_2^{wn}$ to $\mathbb{Z}_2^{wm}$ for $w = \lceil \log q \rceil$. For $b \in \{0,1\}$ the function f_b takes as input an $x \in \mathbb{Z}_q^n$ (that can be seen as an element of $\mathbb{Z}_2^{wn}$ through its binary representation) and returns $Ax + e' + bz$, which is an element of $\mathbb{Z}_q^m \subseteq \mathbb{Z}_2^{wm}$. Here, e' is a vector sampled at random from a distribution $D_{\mathbb{Z}_q,\alpha' q}$ such that α' is "much larger" than α. The inclusion of e' makes f a "randomized" function, which is the main way in which the construction differs from the requirements expressed in Section 4.3. A formal way around this is to think of f_b as the function that returns not $Ax + e' + bz$, but the *distribution* of $Ax + e' + bz$, when $e' \sim D_{\mathbb{Z}_q,\alpha' q}$ and all other variables are fixed. In practice, the evaluation of f on a quantum computer (as required of the honest prover in the verification protocol) involves preparing a weighted superposition over all error vectors, and computing the function in superposition.

We would, of course, rather do away with this complication. Why is the error vector necessary? It is there to satisfy the important requirement that the functions f_0 and f_1 are injective with overlapping ranges, so that f itself is 2-to-1. Injectivity follows from the existence of the trapdoor for A and an appropriate setting of the standard deviation of the error distribution, which guarantee that (given the trapdoor) x can be recovered from $Ax + e' + bz$ (with high probability over the choice of e'). To make the function ranges overlap, we need the distribution of $Ax + e'$ to be statistically close to the distribution of $Ax' + e' + z = A(x' + s) + (e' + e)$. The first distribution considers an arbitrary vector in the column span of A, shifted by e; the second considers the same, except that the shift is by $(e' + e)$. For the two distributions to (almost) match, we need the distribution of e' to (almost) match the distribution of $e + e'$. This is possible as long as the standard deviation $\sigma' = \alpha' q$ is substantially larger than the standard deviation $\sigma = \alpha q$; provided this holds it is an exercise to compute the statistical distance between the two Gaussian and verify that it can be made very close to 1.

With this important caveat in place, we have specified the function f and verified property (**F.1**). Property (**F.3**) follows from the existence of the secret information $td = (R, s)$. Given a $b \in \{0,1\}$ and an element $y = Ax + e' + bz = A(x + bs) + (e' + be)$ in the range of f_b it is possible to use the trapdoor matrix R to recover $x + bs$ and subtract bs to deduce the preimage x of z under f_b. Property (**F.4**) holds trivially from the construction. Note that the function f has domain and range that are different. In particular, here the domain is larger than the range, and in case q is not a power of 2 f is only defined on a subset of its natural domain $\mathbb{Z}_2^{wn}$. These points are not very important and can be ignored at the level of our discussion.

Showing the hardcore bit property (**F.2**) and the collapsing condition (**F.5**) require more work, and we refer to [BCM+18] for a detailed exposition.[6] Similar "hardcore bit" properties to (**F.2**) have been shown for many LWE-based cryptographic schemes (see e.g. [AGV09]). Usually the property states that "for any vector $d \in \mathbb{Z}_q^n \backslash \{0\}$, the value $d \cdot s \in \mathbb{Z}_q$ is indistinguishable from uniform, even given a sample $(A, As + e)$". Our property (**F.2**) is subtly stronger, in that the adversary may choose the vector d itself, possibly as a function of the sample $(A, As + e)$. An additional difficulty stems from the specific equation that the adversary is asked to return. In the definition of Assumption (**F.2**) this is a d such that $d \cdot (x_0 + x_1) = 0$, where x_0, x_1 are the *binary representation* of the two preimages in $\mathbb{Z}_q^n$ of the prover's first message string $y \in \mathbb{Z}_q^m$. (The use of the binary representation comes from the requirements on the honest prover, that is asked to perform a measurement in the Hadamard basis, yielding a binary string of outcomes.) So here

[6] The collapsing condition is not shown in [BCM+18]. It is implicitly shown in [Mah18], where it can be seen to follow from property 2 in Definition 4.4 of an extended trapdoor claw-free family. (The connection is made explicit in [GV19].)

$x_0 = (0, r_0)$ and $x_1 = (1, r_1)$ such that r_0, r_1 are binary representations for two elements $x'_0, x'_1 \in \mathbb{Z}_q^n$ such that $x'_1 = x'_0 - s$ over $\mathbb{Z}_q$. Since the binary representation is not linear the equation obtained is not directly a linear equation in the secret s. Completing the argument showing that a procedure that returns the information asked for in Assumption (**F.2**), i.e. the pair $(x = (b, r_b), d)$, can be turned into a procedure that breaks the decisional LWE assumption, requires a little more work; this is where we need to assume that the secret vector s is a binary vector.

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.

[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.

[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.

[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.

[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.

[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.

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

[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.

[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.

[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.

[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.

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

[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.

[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.