Lecture 1

Introduction

1.1 Presentation of the course

Our goal in this course is to build towards a mostly self-contained presentation of two recent papers in quantum computing:

(a) *Classical verification of quantum computations*, by Mahadev [Mah18]. This paper addresses the question of *verification of quantum computation*: given *classical* data that is obtained from a quantum device that claims to have the ability to execute arbitrary quantum circuits (of polynomial size), how can a classical "verifier" ensure that the reported data indicates the correct outcome of the computation? For problems that have a natural classical certificate of correctness this is a simple task. For example, if the problem is to determine, given as input an integer n, if n has a prime factor larger than (say) $n^{1/4}$, then a positive answer can be certified by providing such a factor when it exists (and a negative answer can be certified by providing a complete prime decomposition of n). Such a factor, or more generally the prime decomposition of n, can be determined in quantum polynomial time using Shor's factoring algorithm; verifying it can be done in classical polynomial time.

However, not all problems that can be solved in quantum polynomial time are believed to lie in the class NP, i.e. not all quantum computations have outcomes that can be certified using an easily verifiable classical "witness". The well-known complexity-theoretic inclusion $\mathrm{BQP} \subseteq \mathrm{IP}$, that we discuss later in the course, implies that all problems in BQP have a classical randomized polynomial-time *interactive* verification procedure; however, in this procedure the prover may be asked to perform computations that are *harder* than BQP.

What Mahadev shows in her breakthrough result is that every polynomial time quantum computation can nevertheless be verified in classical randomized polynomial time by interacting with a quantum polynomial time device *as long as* one can ascertain (or believes) that the quantum device does not have the power to break a natural "post-quantum" cryptographic assumption (namely, the "Learning with Errors" (LWE) assumption).

(b) $\mathrm{MIP}^* = \mathrm{RE}$, by Ji, Natarajan, Vidick, Wright and Yuen [JNV$^+$20]. The equality that gives this paper its title is an equality between complexity classes, i.e. computational problems that have a similar level of "difficulty" in a given model of computation (e.g. BQP and NP are complexity classes). Here, MIP^* designates all those computational problems that can be decided efficiently in classical randomized polynomial time (similar to the verification in the previous item) by asking (classical) questions to *two infinitely powerful quantum provers sharing entanglement*. Here the provers have unbounded

computational capabilities and can hence solve any computational problem they like, including the one that the verifier is concerned with. However, it is assumed that the provers are not trusted: they will always try to convince the verifier that the answer to the problem in question is "yes" (e.g. "yes, this graph does have a valid 3-coloring"). The verifier has to employ certain tricks, or "tests", to detect any malicious behavior by the provers, so as never to make the wrong decision. The class RE denotes all problems for which there is an algorithm, running in any amount of time, that always eventually halts with the answer "yes" when this is the case (the algorithm does not need to halt in other cases). RE is a mind-boglingly large class of problems; it contains any decidable problem and even some undecidable ones such as the halting problem, which is complete for the class.

The equality $\text{MIP}^* = \text{RE}$ is surprising in that it shows that access to untrusted quantum provers grants exceedingly large verification power to the polynomial-time verifier. In contrast, it is known that the same class without entanglement between the provers, denoted MIP, is much smaller, $\text{MIP} = \text{NEXP}$. Other motivations for the result tie it to questions in the foundations of quantum non-locality ("Tsirelson's problem") and the theory of operator algebras ("Connes' embedding problem") which we will discuss in due time.

There is a deep connection between the two papers mentioned above (as well as the many works that led to them—although the course is not meant to be a comprehensive survey, we will review the most relevant references in due course). Indeed, at their heart both works identify means by which a classical procedure is able to certify an appropriate "quantum computation workspace" within one (or two) quantum devices, using only a classical interaction with it. In this sense, and to borrow the title of one of the important earlier papers in this area [RUV13], both works provide techniques to tie a "classical leash" around a quantum system. In order to achieve this both works ultimately have to tackle the same fundamental problem: what are classical signatures of quantum processes that can be leveraged to certify an entire computation? Very informally, we will see that in the case of (a) this signature is provided by the *uncertainty principle*: the fact that certain measurements in quantum mechanics are intrinsically incompatible; in the case of (b) it will be provided by *quantum non-locality*: the fact that entanglement allows distant parties to generate correlations that have no classical equivalent.

We have motivated our choice of topics by arguing that they tackle a fundamental problem, that of classically testing a quantum system. This is a problem of practical relevance (given an experimental quantum device, does it really do anything quantum?) as well as one that reaches deep to the power and limitations of the scientific method (for a discussion from an epistemological point of view, I recommend the great presentation by Aharonov and Vazirani [AV12]). It is also a problem that has turned out to stimulate many recent advances in quantum cryptography and complexity. For example, early works in delegated computation encouraged the development of quantum authentication codes, such as the Clifford code, that have found wide uses in cryptography; the study of quantum entangled-prover interactive proof systems has brought many discoveries in the foundations of quantum non-locality, such as dimension witnesses. The topics are connected through the common framework of interactive proof systems and share many techniques. Throughout the course we will make an effort to highlight the many open questions that stimulate research in this area and hope that our choice of results will provide a compelling entry point to it.

Having set the stage, let us discuss the structure of the course. We will start by tackling arguably the most fundamental question in this line of work: *what is a qubit*? How to define it mathematically, and how does one "certify" its existence? What does it mean to "test a qubit," and how can it be done? This question will occupy us for the four lectures. Exploring it will give us the opportunity to lay common foundations for the discussion of results (a) and (b) above. The remaining 6 lectures will be equally divided in two sets of 3 lectures each. The first set will examine the problem of delegation of polynomial-time

quantum computations in general, and then focus on a presentation of Mahadev's result. The second set will discuss the theory of multi-prover interactive proof systems and build towards an understanding of the main ingredients in the result by Ji et al.

We will start the course slowly, so that the first few lectures are accessible to as broad a public as possible. The content will be mathematically precise and as self-contained as can be, so that it is possible to follow starting with only an elementary background in quantum computing at the level of the Nielsen & Chuang textbook [NC02]. Later lectures will be more involved and may require a somewhat more advanced knowledge of certain topics in cryptography or in complexity theory; I will aim to keep the presentation as self-contained as possible. My goal is that by the end of the course an assiduous participant may dive into either paper with a solid understanding of the main steps and key techniques in mind.

1.2 What is a qubit?

Let's start with the basics; as we will see shortly the simplest questions are not always the least interesting.

According to Wikipedia, "The qubit is the basic unit of information for quantum computers." The associated entry goes on to declare that "A qubit is a two-state (or two-level) quantum-mechanical system," the latter being defined as "a quantum system that can exist in any quantum superposition of two independent (physically distinguishable) quantum states." So a qubit is some kind of "system" whose state space has two fundamental states, and moreover such that the system can be in any "superposition" of these two states. Note that the definition makes an important distinction between the system (qubit) itself and its state. The latter is easier to define, and indeed it is through its state space that the "qubit" is generally introduced in quantum computing courses: the *state space of a qubit* is the set of unit vectors in the complex vector space $\mathbb{C}^2$; we denote this space as $S(\mathbb{C}^2)$.[1] Thus the *state of a qubit* can be represented by a unit vector in $\mathbb{C}^2$; for example,

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad \text{and} \qquad |+\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

are both valid states for a qubit.[2]

So we all know how to recognize a valid state for a qubit. But what about *the qubit itself*? We would like a precise definition that captures the intuition given in the Wikipedia entry: a qubit is a "two-state system" that can be in any "independent superposition" of its two states. In particular, we would like our definition to clearly distinguish this notion from that of a classical "probabilistic bit", which could also be considered a two-level system that can be in any "superposition" of its two states, $p\,\overline{0} + (1-p)\,\overline{1}$ for any $p \in [0,1]$. What distinguishes the real, "1-dimensional" degree of freedom $p \in [0,1]$ of the probabilistic bit from the complex, "2-dimensional" degree of freedom of the qubit?

As we will soon see, for our purposes it turns out to be most meaningful to attempt a definition in terms of the "Heisenberg representation" of quantum mechanics, that places *observable* quantities at a forefront. Informally, we will distinguish a quantum degree of freedom from a classical one by requiring that the quantum degree of freedom can be measured (observed) in two *mutually incompatible ways*. In order to make this precise we make a small detour to introduce the formalism associated with observable quantities in quantum mechanics.

[1] To be consistent with standard mathematical terminology we'd call this the complex projective space and denote it $P^2(\mathbb{C})$. We'll follow standard quantum computing conventions instead.

[2] For the entirety of this course (except the last lecture) we take the traditional perspective from quantum computer science: unless specified otherwise we consider that Hilbert spaces are always finite dimensional and computation is always performed on qubits in $\mathbb{C}^2$ or possibly higher-dimensional qudits in $\mathbb{C}^d$ for $d \geq 1$ a (finite) integer.

1.2.1 Observables

While the state space of a *qubit* is generally taken to be $\mathbb{C}^2$, more generally a quantum mechanical state lies in an arbitrary separable Hilbert space $\mathcal{H}$.[3] The term "observable" is used to denote any quantity that can in principle be obtained as the result of a measurement: for example, position, momentum, spin, energy (with respect to a certain Hamiltonian), are all observables. In full generality an observable is specified by a Hermitian operator O on $\mathcal{H}$.[4] The interpretation of this is that each eigenvalue of O represents a possible outcome under a measurement of the observable, and the associated eigenvectors denote states under which the observable deterministically yields that outcome. Thus if $O = \sum_i \lambda_i \Pi_i$ is a spectral decomposition of O, then any state $|\psi\rangle$ such that $\Pi_i|\psi\rangle = |\psi\rangle$ will deterministically yield the outcome λ_i when measured according to O. In particular, we see that the important part of an observable is its eigenprojections rather than the associated eigenvalues: the latter are real numbers that are associated to the different possible experimental outcomes. These numbers are generally associated a physical meaning (such as position, momentum, etc.) but they can easily be changed by post-processing.

A collection of projection operators $\{\Pi_i\}$ that sum to identity is called a *Projector-Valued Measure* (PVM). When the Π_i are no longer required to be projections, it is called a *Positive Operator-Valued Measure* (POVM). This is the most general kind of measurement that is allowed in quantum mechanics. Given a POVM $\{\Pi_i\}$ the *Born rule* specifies that a measurement of an arbitrary state $|\psi\rangle$ under it will yield measurement outcome i with probability $\|\sqrt{\Pi_i}|\psi\rangle\|^2 = \langle\psi|\Pi_i|\psi\rangle$. Since the Π_i form a resolution of the identity ($\sum_i \Pi_i = \mathrm{Id}$) and $\||\psi\rangle\|^2 = 1$ we see that these probabilities always sum to 1, as they should. When the outcome i is obtained, the state evolves to a *post-measurement state* $|\psi'\rangle = \sqrt{\Pi_i}|\psi\rangle / \|\sqrt{\Pi_i}|\psi\rangle\|$. (Here $\sqrt{\Pi_i}$ is generally taken to be the positive square root of Π_i. However, other square roots can be used as well: since different square roots differ only by a unitary degree of freedom, choosing the one over the other is analogous to imposing an additional reversible evolution on the post-measurement state, which can be considered to be part of the measurement itself.)

In case the $\{\Pi_i\}$ are obtained from some observable O, we may use the associated eigenvalues to associate a real value to each measurement outcome i. In this case the *expectation* of the outcome of the measurement is

$$\sum_i \lambda_i \|\Pi_i|\psi\rangle\|^2 = \sum_i \lambda_i \langle\psi|\Pi_i|\psi\rangle = \langle\psi|O|\psi\rangle ,$$

a quantity that is sometimes referred to as the "overlap" of $|\psi\rangle$ on O. We will use this formula often.

An observable such that $O^2 = \mathrm{Id}$ has at most two eigenvalues, which by convention we take to be -1 and $+1$. Such an observable is called a *binary* observable; it is the most frequent kind of observable that we will encounter.

1.2.2 First definition of a qubit

With a precise mathematical definition of an observable quantity we are ready to make precise our informal definition of a "qubit" as "a system that can be observed in two mutually incompatible ways".

[3] "Separable" means that $\mathcal{H}$ has a countable basis. In fact quantum states can also live in non-separable Hilbert spaces; we make the restriction for convenience. In fact for the purposes of these notes you may as well think of $\mathcal{H}$ as being finite-dimensional, i.e. $\mathcal{H}$ is isomorphic to the complex vector space $\mathbb{C}^d$, for some integer $d \geq 1$. Allowing infinite-dimensional spaces gives us a little more generality.

[4] For us "Hermitian" means that O is self-adjoint, $O = O^\dagger$, where $O^\dagger$ is the operator such that $\langle O^\dagger u|v\rangle = \langle u|Ov\rangle$ for all $|v\rangle$ in the domain of O. If $\mathcal{H}$ is finite-dimensional then a matrix representation of $O^\dagger$ is obtained by taking the conjugate-transpose of a matrix representation of O.

Definition 1 (Qubit, Take 1)**.** A *qubit* is a triple $(\mathcal{H}, X, Z)$ consisting of a separable Hilbert space $\mathcal{H}$ and a pair of Hermitian operators X, Z acting on a $\mathcal{H}$ such that $X^2 = Z^2 = \mathrm{Id}$ and $\{X, Z\} = XZ + ZX = 0$.

Let's see why Definition 1 captures the intuitive notion of "mutually incompatible" observables. Let the 'computational basis' be an eigenbasis of Z, and the 'Hadamard basis' an eigenbasis of X. Then we claim that the anticommutation relation $XZ + ZX = 0$ ensures that any vector in the former makes a 45° angle with any vector in the latter. To see this, let $|\psi\rangle$ be an eigenvector of X with associated eigenvalue $\varepsilon \in \{\pm 1\}$. Then $\langle\psi|XZ + ZX|\psi\rangle = 0$ immediately implies $2\varepsilon\langle\psi|Z|\psi\rangle = 0$. Given that Z only has -1 and $+1$ as eigenvalues, this relation implies that the projections of $|\psi\rangle$ on the two eigenspaces of Z have equal length; in other words, $|\psi\rangle$ lies exactly between the $+1$ and -1 eigenspaces of Z. (Yet another way of saying this is that all principal angles between an eigenspace of X and one of Z are $\frac{\pi}{4}$; we will make this precise soon.) In this sense any X and Z satisfying the conditions of the definition are "maximally incompatible": any definite state for the one is entirely undetermined (i.e. yields uniformly random outcomes when measured) under the other.

There is a problem with this definition: by allowing the underlying Hilbert space $\mathcal{H}$ to be arbitrary we seem to have all but dropped the earlier requirement that a qubit is a system whose state space is "two-level" and thus identifiable with the projective space $\mathrm{S}(\mathbb{C}^2)$. Luckily, the following lemma allows us to make the connection with this requirement.

Lemma 2. *Let $(\mathcal{H}, X, Z)$ be a qubit. Then there is a Hilbert space $\mathcal{H}'$ and an isomorphism $\mathcal{H} \simeq \mathbb{C}^2 \otimes \mathcal{H}'$ such that under the same isomorphism, $X \simeq \sigma_X \otimes \mathrm{Id}$ and $Z \simeq \sigma_Z \otimes \mathrm{Id}$.*[5] *Here, σ_X and σ_Z are the usual Pauli observables on $\mathbb{C}^2$: in matrix form,*

$$\sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \textit{and} \qquad \sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Note that a consequence of the lemma is that qubits, as defined in Defintion 1, only exist in spaces of even (or infinite) dimension! In particular, qubits don't exist in dimension 1; indeed, in dimension 1 all operators commute. This is satisfactory: intuitively, a situation in which all possible observables commute ought to be considered "classical" (for instance, because there is a complete set of simultaneous eigenvectors for all observables).

It will be essential for our later goals that Definition 1 does not *a priori* require $\mathcal{H}$ to be a two-dimensional space. Indeed, how would one test such a claim? One does not "see" the dimension of the state space; while it is possible to probe parts of it it can never be excluded that the state space is larger than what is accessible to the experimentalist's setup. In this sense Definition 1 has a nice "operational" flavor to it: it refers to *observables* of the system and their properties. Although much more work is needed before we are able to make any of these statement formal, we see the definition as a good step towards giving us the ability to "test" that a system "is a qubit". In addition, the definition clearly has meaningful consequences; in particular it implies that qubits do not have a "classical explanation", so that a "test for a qubit" can serve as a "test for quantumness", i.e. a test that distinguishes quantum from classical behavior.

The proof of the lemma makes use of an elementary but fundamental tool in the analysis of many quantum information protocols, the CS (for "Cosine-Sine") decomposition. This decomposition is also known as "Jordan's lemma," after Camille Jordan's *Traité des substitutions et des équations algébriques* from 1870 (see this EHESS PhD thesis [Bre06] for a masterful 730-page account of the history behind the

[5]The reader might wonder what happened to σ_Y... Don't we need it to define our qubit? Here we are taking the "operator algebraists"' perspective, which is that if the system supports X and Z observables then it also supports $Y = iXZ$. Because Y is determined by X and Z, we do not include it in the definition.

use of Jordan's name alongside this theorem). Given that we will use the lemma frequently we give it a self-contained treatment in the next section.

1.2.3 Jordan's lemma

The discussion of Jordan's lemma in this section is mostly borrowed from `https://cims.nyu.edu/~regev/teaching/quantum_fall_2005/ln/qma.pdf`. Let P, Q be two orthogonal projections on a separable Hilbert space $\mathcal{H}$, and consider their sum $R = P + Q$. Then R is Hermitian so it has an orthonormal set of eigenvectors that is also a basis for $\mathcal{H}$.[6] Let $|\varphi\rangle$ be any eigenvector of R with associated eigenvalue λ. Consider two cases for the vector $P|\varphi\rangle$. The first case is that $P|\varphi\rangle$ is parallel to $|\varphi\rangle$. In this case, since

$$Q|\varphi\rangle = R|\varphi\rangle - P|\varphi\rangle = \lambda|\varphi\rangle - P|\varphi\rangle \tag{1.1}$$

it follows that $Q|\varphi\rangle$ is also parallel to $|\varphi\rangle$, so $|\varphi\rangle$ is a common eigenvector of P and Q. The second case is that $P|\varphi\rangle$ is linearly independent from $|\varphi\rangle$. In this case the two-dimensional subspace $\mathcal{S}$ spanned by $|\varphi\rangle$ and $P|\varphi\rangle$ is invariant under P, because being a projection P satisfies $P^2 = P$. Moreover, by (1.1) we see that $Q|\varphi\rangle$ lies in $\mathcal{S}$, and

$$QP|\varphi\rangle \;=\; Q(\lambda|\varphi\rangle - Q|\varphi\rangle) = (\lambda - 1)Q|\varphi\rangle$$

is also in that subspace. Using that P, Q are projections and that P is neither 0 nor the identity on $\mathcal{S}$ (otherwise we would have been in the first case) it follows that there is a basis of $\mathcal{S}$ such that in that basis,

$$P = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \qquad \text{and} \qquad Q = \begin{pmatrix} c^2 & cs \\ cs & s^2 \end{pmatrix}, \tag{1.2}$$

where $c = \cos\theta$ and $s = \sin\theta$ for some $\theta \in [0, \pi/2)$. (Other values of θ reduce to this case because the sign of cs can be flipped by negating the second basis vector for the chosen basis of $\mathcal{S}$.) Note finally that since $\mathcal{S}$ is stable by both P and Q it is stable by $R = P + Q$, so it has a basis made of eigenvectors of R, the vector $|\varphi\rangle$ that we started from and its orthogonal in R. Proceeding in this way inductively this lets us identify an eigenbasis of R such that its vectors are either isolated (stable by both P and Q) or in pairs (spanning a 2-dimensional subspace that is stable by both P and Q).

The following lemma summarizes the discussion so far.

Lemma 3. *Let P, Q be projections on a separable Hilbert space $\mathcal{H}$. Then there exists an orthogonal decomposition $\mathcal{H} = \oplus_i \mathcal{S}_i$ such that each $\mathcal{S}_i$ is a* 1- *or* 2*-dimensional subspace that is stable by P and Q. Furthermore, whenever $\mathcal{S}_i$ is* 2*-dimensional there is a basis for it in which P and Q take the form* (1.2)*, for some c_i and s_i that may depend on $\mathcal{S}_i$.*[7]

This very useful lemma informally says that, when only two projections are concerned, we can reduce the analysis to a 2-dimensional problem. Let's apply it to show Lemma 2.

Proof of Lemma 2. Let $\mathcal{H}$, X, Z be as in the statement of the lemma. Let $P = \frac{1}{2}(Z + \mathrm{Id})$ and $Q = \frac{1}{2}(X + \mathrm{Id})$. Then P, Q are projections on $\mathcal{H}$ so we can decompose them according to Lemma 3. Let $(|e_i\rangle, |f_i\rangle)$ be a basis for the i-th space $\mathcal{S}_i$ in which the matrices for P and Q have the form (1.2). Using $XZ + ZX = 0$ it follows that (i) there cannot be any 1-dimensional blocks, because these necessarily

[6] Here we are using that $\mathcal{H}$ is separable.

[7] For the case of 1-dimensional subspaces, since P and Q are projections they are each either identically 0 or identity in those subspaces.

commute, and (ii) in any two-dimensional block, the angle θ must equal $\pi/4$, as this is the only value in $[0, \pi/2)$ that leads to anti-commuting operators. Thus in each space $\mathcal{S}_i$, Z acts exactly as σ_Z and X as σ_X. Let $\mathcal{H}'$ have canonical basis $\{|i\rangle\}$, where i ranges over the block indices in the decomposition of P and Q. The required isomorphism is obtained by e.g. mapping $|e_i\rangle \in \mathcal{H}$ to $|0\rangle \otimes |i\rangle \in \mathbb{C}^2 \otimes \mathcal{H}'$ and $|f_i\rangle \in \mathcal{H}$ to $|1\rangle \otimes |i\rangle \in \mathbb{C}^2 \otimes \mathcal{H}'$. □

1.2.4 n qubits

Now that we have a working definition of a qubit, how about two, three, or even n qubits? What we mean when we say that a system "has n qubits" is that (i) it should have n copies of one qubit, so there should be $(X_1, Z_1), (X_2, Z_2), \ldots, (X_n, Z_n)$ on $\mathcal{H}$ such that each pair satisfies the definition of a qubit, and moreover (ii) the qubits should be "independent": indeed, we wouldn't want something like $X_1 = X_2$ to happen. How do we prevent this? Intuitively "independence" of the qubits should be reflected in the fact that they can be observed "independently", in any order, such that if e.g. we " observe" qubits 1 and 2 and then discard the outcome associated with qubit 1, we should obtain an outcome that is identically distributed as if we had only observed qubit 2 in the first place. These considerations suggest the following definition.

Definition 4 (n qubits, Take 1)**.** A system of n *qubits* is a tuple $(\mathcal{H}, X_1, Z_1, \ldots, X_n, Z_n)$ consisting of a separable Hilbert space $\mathcal{H}$ and n pairs of Hermitian operators (X_i, Z_i) for $i \in \{1, \ldots, n\}$ acting on $\mathcal{H}$ such that

(i) For each $i \in \{1, \ldots, n\}$, $(\mathcal{H}, X_i, Z_i)$ is a qubit;

(ii) For each $i \neq j \in \{1, \ldots, n\}$, qubits i and j are independent:

$$[X_i, X_j] = [X_i, Z_j] = [Z_i, X_j] = [Z_i, Z_j] = 0\,,$$

where for arbitrary operators A, B on $\mathcal{H}$, $[A, B] = AB - BA$ denotes the algebra commutator.

The commutation condition (ii) indeed implies that measurements on different qubits can be performed "independently". For example, the expectation of a measurement of qubit 2 in the Hadamard basis *after* qubit 1 has been measured in the computational basis is given by (where we write $X_1 = X_1^0 - X_1^1$ for the spectral decomposition of the observable X_1)

$$\begin{aligned}\frac{1}{2}\langle\psi|X_1^0 Z_2 X_1^0|\psi\rangle + \frac{1}{2}\langle\psi|X_1^1 Z_2 X_1^1|\psi\rangle &= \langle\psi|X_1 Z_2 X_1|\psi\rangle \\ &= \langle\psi|Z_2 X_1^2|\psi\rangle \\ &= \langle\psi|Z_2|\psi\rangle\,,\end{aligned}$$

as desired. (Here, for the first equality we used that $\langle\psi|X_1^0 Z X_1^1|\psi\rangle = \langle\psi|X_1^1 Z X_1^0|\psi\rangle$, for the second we used the commutation condition $[X_1, Z_2] = 0$, and for the last we used $X_1^2 = \mathrm{Id}$.) This calculation can be done for any pair of qubits, or even any sequence of measurements of qubits, and it shows that item (ii) in Definition 4 indeed captures the idea that each of the n qubits can be measured independently.

Note however that this "independence" is not necessarily quite the same as there truly being n qubits. In particular, if our definition of a qubit only required the use of a single observable X, as a classical bit would, then taking $X_1 = \cdots = X_n$ would satisfy both (i) and (ii), since an operator always commutes with itself. Indeed, just as we're used to a qubit being defined through its state space $\mathrm{S}(\mathbb{C}^2)$, we're used to n qubits being defined through their state space $\mathrm{S}(\mathbb{C}^2) \otimes \cdots \otimes \mathrm{S}(\mathbb{C}^2)$. Where is the tensor product in Definition 4, isn't it missing? The following lemma shows that an n-fold tensor product is in fact implicit in the definition.

Lemma 5. *Let $(\mathcal{H}, X_1, Z_1, \ldots, X_n, Z_n)$ be a system of n qubits. Then there exists a Hilbert space $\mathcal{H}'$ and an isormorphism $\mathcal{H} \simeq (\mathbb{C}^2)^{\otimes n} \otimes \mathcal{H}'$ such that under the same isomorphism, for every $i \in \{1, \ldots, n\}$ and $W \in \{X, Z\}$, $W_i \simeq \sigma_{W,i} \otimes \mathrm{Id}_{\mathcal{H}'}$, where here $\sigma_{W,i}$ denotes the Pauli W operator acting on the i-th copy of $\mathbb{C}^2$.*

Proof. We show the lemma by induction on $n \geq 1$. The case $n = 1$ is provided by Lemma 2. Suppose the lemma shown for some $n \geq 1$, show it for $(n+1)$. Let $(\mathcal{H}, X_1, Z_1, \ldots, X_{n+1}, Z_{n+1})$ be a system of $(n+1)$ qubits. Since $(\mathcal{H}, X_1, Z_1, \ldots, X_n, Z_n)$ is a system n qubits we can apply the induction hypothesis to it. Let $\mathcal{H}'$ and π' be the promised space and isomorphism. The key step is provided by the following claim.

Claim 6. *Let W be an Hermitian operator on $\mathcal{H}$ such that $[W, X_i] = [W, Z_i] = 0$ for all $i \in \{1, \ldots, n\}$. Then there exists W' Hermitian acting on $\mathcal{H}'$ such that under π', $W \simeq \mathrm{Id}_{(\mathbb{C}^2)^{\otimes n}} \otimes W'$.*

The claim immediately gives us the induction step: by applying it to X_{n+1} and Z_{n+1} we find X'_{n+1} and Z'_{n+1} on $\mathcal{H}'$ such that $(\mathcal{H}', X'_{n+1}, Z'_{n+1})$ is a qubit. Applying Lemma 2 to this qubit and composing the isomorphism obtained with π' completes the induction step. Therefore, it only remains to prove the claim.

Proof of Claim 6. Clearly it suffices to prove the statement "under π", i.e. for the case where $X_i = \sigma_{X,i}$ and $Z_i = \sigma_{Z,i}$. We introduce the following notation: for $a, b \in \{0,1\}^n$,

$$\sigma_X(a) = \sigma_{X,1}^{a_1} \otimes \cdots \otimes \sigma_{X,n}^{a_n} \qquad \text{and} \qquad \sigma_Z(b) = \sigma_{Z,1}^{b_1} \otimes \cdots \otimes \sigma_{Z,n}^{b_n} .$$

Using that the four 1-qubit Pauli matrixes Id, σ_X, σ_Z and $\sigma_X\sigma_Z$ form a basis for the complex vector space of linear operators on $\mathbb{C}^2$, W has a decomposition

$$W = \sum_{a,b} \sigma_X(a)\sigma_Z(b) \otimes W_{a,b} ,$$

where $W_{a,b}$ are arbitrary operators on $\mathcal{H}'$ (they are not necessarily Hermitian). Let's write out the left and right products of W with $\sigma_X(c)\sigma_Z(d)$, for some $c, d \in \{0,1\}^n$:

$$\begin{aligned} \sigma_X(c)\sigma_Z(d)W &= \sum_{a,b} \sigma_X(c)\sigma_Z(d)\sigma_X(a)\sigma_Z(b) \otimes W_{a,b} \\ &= \sum_{a,b} (-1)^{a\cdot d} \sigma_X(a+c)\sigma_Z(b+d) \otimes W_{a,b} , \end{aligned} \tag{1.3}$$

where in the second line we used the anti-commutation relation $\sigma_Z(d)\sigma_X(a) = (-1)^{a\cdot d}\sigma_X(a)\sigma_Z(d)$, as well as the "additivity" relations $\sigma_X(a+c) = \sigma_X(a)\sigma_X(c)$, and similarly for σ_Z. Similarly,

$$\begin{aligned} W\sigma_X(c)\sigma_Z(d) &= \sum_{a,b} \sigma_X(a)\sigma_Z(b)\sigma_X(c)\sigma_Z(d) \otimes W_{a,b} \\ &= \sum_{a,b} (-1)^{b\cdot c} \sigma_X(a+c)\sigma_Z(b+d) \otimes W_{a,b} . \end{aligned} \tag{1.4}$$

Using that the $\sigma_X(a)\sigma_Z(b)$ are linearly independent we can identify terms in (1.3) and (1.4); it follows that for any a, b, c, d, $(-1)^{b\cdot c} W_{a,b} = (-1)^{a\cdot d} W_{a,b}$. For any (a, b) unless $a = b = 0$ we can find strings c, d such that the two terms in $W_{a,b}$ are given opposite signs. Thus $W_{a,b} = 0$ whenever $(a, b) \neq (0, 0)$, and $W = \mathrm{Id} \otimes W_{0,0}$. Since W is Hermitian, $W_{0,0}$ is also Hermitian, proving the claim. □

□

Remark 7. The statement of Lemma 5 can be reformulated in the language of group representation theory, and this reformulation will be useful later on. The "n qubit Weyl-Heisenberg group" is the $2 \cdot 4^n$-element group G_n that is generated by the n-qubit σ_X and σ_Z matrices; its elements are $(-1)^c \sigma_X(a)\sigma_Z(b)$ for $a, b \in \{0,1\}^n$ and $c \in \{0,1\}$. From any system of n qubits $(\mathcal{H}, X_1, Z_1, \ldots, X_n, Z_n)$ it is straightforward to specify a representation ϕ of G_n by setting $\phi((-1)^c \sigma_X(a)\sigma_Z(b)) = (-1)^c \prod_i X_i^{a_i} \prod_i Z_i^{b_i}$. The lemma can be adapted to show that any representation of G_n that in addition sends -1 to -1, as ϕ does, must be a direct sum of copies of the representation by Pauli matrices.

1.2.5 Approximate qubits

An important theme of this course will be that the objects that we observe and manipulate in our "protocols" or "experiments" generally cannot be assumed to be perfectly "clean" or "noise-free". In this respect, the following exercise is a simple test that we ought to make on our definition; furthermore, it is a good exercise to practice the use of the CS decomposition.

Exercise 1.1. Suppose that X and Z are binary observables on $\mathcal{H}$ such that $\|\{X, Z\}\| \leq \varepsilon$ for some $\varepsilon \geq 0$. Show that there exists a qubit $(\mathcal{H}, X', Z')$ such that $\max\{\|X - X'\|, \|Z - Z'\|\} \leq \delta(\varepsilon)$. State the best dependence δ that you can get.

The exercise can be extended to consider n approximate qubits, but the proof is more delicate as some work is needed to keep the errors under control. The following is shown in [CRSV17].

Theorem 8. *Let $X_1, Z_1, \ldots, X_n, Z_n$ be binary observables on $\mathcal{H}$ and $\varepsilon \geq 0$ such that $\varepsilon/(1-\varepsilon)^2 \leq 1/(64n)$ and $\|\{X_i, Z_i\}\| \leq \varepsilon$ and $\|[S_i, T_j]\| \leq \varepsilon$ for all $i \neq j \in \{1, \ldots, n\}$ and $S, T \in \{X, Z\}$. Then there exists binary observables $X'_1, Z'_1, \ldots, X'_n, Z'_n$ on $\mathcal{H}$ such that $\{X'_i, Z'_i\} = 0$, $[S'_i, T'_j] = 0$ and moreover $\|S'_j - S_j\| \leq 4n\varepsilon/(1-\varepsilon)^2 + \varepsilon$ for all $i \neq j \in \{1, \ldots, n\}$ and $S, T \in \{X, Z\}$.*

The theorem shows that "n approximate qubits" are close to "n exact qubits" according to our definitions. Note that there is a dependence of the error on n, but not on the dimension of $\mathcal{H}$. In [CRSV17] it is shown by an explicit example that a linear dependence on n is necessary.

1.2.6 An operational definition?

Earlier we qualified our definition of a qubit as being "operational". This term is generally used to refer to a definition that can be "observed", i.e. the definition should imply a natural "test" that can be performed experimentally and that "certifies" that a certain object satisfies the definition in some way or another. Definition 1, or even the approximate version of it suggested in the previous section, lacks severely in that matter: who gets to observe "operators"? How can the anti-commutator be "witnessed"? In the next lecture we will refine our definition so that we are able to answer these questions.