# CMSC 657: Introduction to Quantum Information Processing
# Lecture 15

Instructor: Daniel Gottesman

Fall 2024

## 1 Implementations of Quantum Computers

### 1.1 DiVincenzo Criteria

What properties does a system need to make it a good candidate for building a quantum computer? The main ones were laid out by David DiVincenzo and are widely known as the "DiVincenzo Criteria." They are not actually either sufficient nor are all of them strictly necessary for building a quantum computer, but they do provide a good guide, and most of the leading candidates for building quantum computers satisfy them:

1. A scalable physical system with well-characterized qubits.
2. The ability to initialize the state of the qubits to a simple fiducial state, such as $|000\ldots\rangle$.
3. Long relevant decoherence times, much longer than the gate operation time.
4. A universal set of quantum gates.
5. A qubit-specific measurement capability.
6. The ability to interconvert between stationary and flying qubits.
7. The ability to faithfully transmit flying qubits between specified locations.

The last two are only needed if you want to do quantum communications and interface the quantum computer with the communications system. The first five are for building a computational device.

### 1.2 Ion Traps

The first type of implementation we will discuss uses ions — single atoms missing an electron each — as qubits. It is the approach used by the companies IonQ and Quantinuum. IonQ currently has a 36-qubit quantum computer as their largest public device and Quantinuum has a 56-qubit device.

#### 1.2.1 Ions

In ion traps, one qubit = one atom, so first, let us talk about atoms. Atoms have a nucleus made of protons and neutrons which is positively charged, and lighter electrons which are negatively charged. Protons and electrons have the same magnitude of charge, so an atom is neutral (no electric charge) if the number of electrons and protons match. Electrons are a type of particle known as a *fermion*, which satisfy the *Pauli exclusion principle*. This means that no two electrons can be in exactly the same state. It is a vaguely reasonable approximation to say that the electrons act independently exact for the Pauli exclusion principle and that therefore, we can write down a Hamiltonian for a single electron bound by the nucleus. It has a

discrete set of energy levels (but infinite; they get closer together as the energy goes up) — eigenstates of the Hamiltonian — and at low temperatures, the electrons bound to an atom fill up the energy levels from the lowest on up.

The qubit in an ion trap quantum computer is an atom, generally a group II atom (from the second column of the periodic table). Each row of periodic table corresponds to a significant leap in energy between the eigenstates being filled and atoms are more stable when these shells of electrons are completely full. (This heuristic explains a great deal of chemistry just by itself.) A group II atom has two electrons in its outer shell, so it has a tendency to lose electrons. For quantum computing purposes, it is useful to have atoms which are ionized once (meaning they lost one electron), giving them charge +1 and leaving one electron in the outer shell. Because the ion is charged, it interacts with electric and magnetic fields and can be trapped in this way. (There are different trap designs.)

We can basically ignore the other electrons. They will fill the lower-lying shells and remain there, as there is a big energy cost for any of them to leave. However, the remaining outer shell electron can be more easily excited into different energy levels. We often draw pictures representing the different possible levels, e.g.:

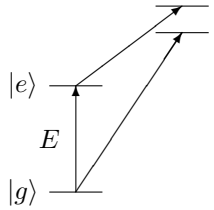


The horizontal axis in this kind of drawing doesn't always have a definite meaning. It just helps to separate out different kinds of levels so we can see better. Sometimes it might represent some other property (such as angular momentum) that distinguishes different levels. The specific energies of the energy levels depend on the detailed physics of the atom (no longer making simplifying approximations) and energy levels have other properties associated with them. One important factor is the energy difference between levels, and there may be other constraints ("selection rules") that limit when the electron can move between levels. The arrows represent interesting transitions between levels.

In the absence of any other consideration, electrons tend to drop to the lowest open energy level. This is because room temperature is still "cold" from the perspective of atomic energy levels, so the amount of energy in the atom in a thermal state is not enough to excite the electron to higher energy levels. Energy is conserved, so to drop down to a lower energy level, the electron needs to emit a photon, a particle of light. However, because of the selection rules, it might not be possible for an electron to go from a particular excited state to the ground state directly or might require more than one photon. Excited states that have this property can therefore be highly stable, lasting for minutes or even longer, which is enough time to do $10^5$ or more quantum gates.

The qubit is then the state of the outer shell electron. Two particularly stable states are chosen to be the computational states, $|0\rangle = |g\rangle$ the ground state and some excited state $|1\rangle = |e\rangle$. This gives us the first DiVincenzo criterion — multiple ions means multiple qubits and the ground and excited states are well-defined with specific known energies. Other excited states will also be relevant for some operations.

### 1.2.2 One-qubit gates in an ion trap

We control the qubit via lasers shining on the atom. Quantum particles of light are called photons, and photons, in contrast to electrons, are *bosons*. This means that unlike fermions, they *prefer* to be in the same state. A laser is a Bose condensate of light, with many photons all in the same state. When you shine a laser on an atom, the atom can absorb one of the laser's many photons to excite the outer shell electron

to a higher energy level, or the Bose statistics can induce the atom to emit an photon into the laser mode, joining the many photons already there.

While discussing the ion trap, we will consider the laser to be a classical object and just write down a Hamiltonian for the atom. Let us first consider the case of a laser whose frequency $\omega = E/\hbar$ is equal to the energy difference $E$ between the ground and excited states and consider only the two computational states. It is important that the energy difference between the $|0\rangle$ and $|1\rangle$ states be different from the difference between either of those states and any other electronic state so that other states are not accidentally excited by this laser. The Hamiltonian is then approximately

$$H = H_{\text{atom}} + H_{\text{laser}}, \tag{1}$$

with

$$H_{\text{atom}} = E|e\rangle\langle e| \tag{2}$$

$$H_{\text{laser}} = J(e^{i\omega t}|g\rangle\langle e| + e^{-i\omega t}|e\rangle\langle g|). \tag{3}$$

Recall that $\omega = E/\hbar$ is the frequency of the laser. The phase factors in $H_{\text{laser}}$ appear because the laser is absorbing or emitting energy into the atom, so the phase of the laser changes according to the amount of energy. We can assume that the ground state energy is 0.

This coupling induces oscillations over time between ground and excited, which are eigenstates of $H_{\text{atom}}$ but not of $H$. $H$ is just a $2 \times 2$ matrix

$$H = \begin{pmatrix} 0 & Je^{i\omega t} \\ Je^{-i\omega t} & E \end{pmatrix}. \tag{4}$$

This is a little more complicated than the Hamiltonians we have discussed previously because of the explicit time-dependence. The Schrödinger equation gives us for a state $|\psi(t)\rangle = \alpha(t)|g\rangle + \beta(t)|e\rangle$,

$$i\hbar(\alpha'(t)|g\rangle + \beta'(t)|e\rangle) = Je^{i\omega t}\beta(t)|g\rangle + (Je^{-i\omega t}\alpha(t) + E\beta(t))|e\rangle. \tag{5}$$

A solution for $|\psi(0)\rangle = |g\rangle$ is

$$|\psi(t)\rangle = \cos(Jt/\hbar)|g\rangle - i\sin(Jt/\hbar)e^{-iEt/\hbar}|e\rangle, \tag{6}$$

giving

$$i\hbar\alpha'(t) = -iJ\sin(Jt/\hbar) = Je^{i\omega t}\beta(t) \tag{7}$$

$$i\hbar\beta'(t) = J\cos(Jt/\hbar)e^{-iEt/\hbar} - iE\sin(Jt/\hbar)e^{-iEt/\hbar} \tag{8}$$

$$= Je^{i\omega t}\alpha(t) + E\beta(t). \tag{9}$$

In particular, if the laser is on for a time $t$, the probability of finding the state $|g\rangle$ afterwards is $\cos^2(Jt/\hbar)$. This is known as a Rabi oscillation.

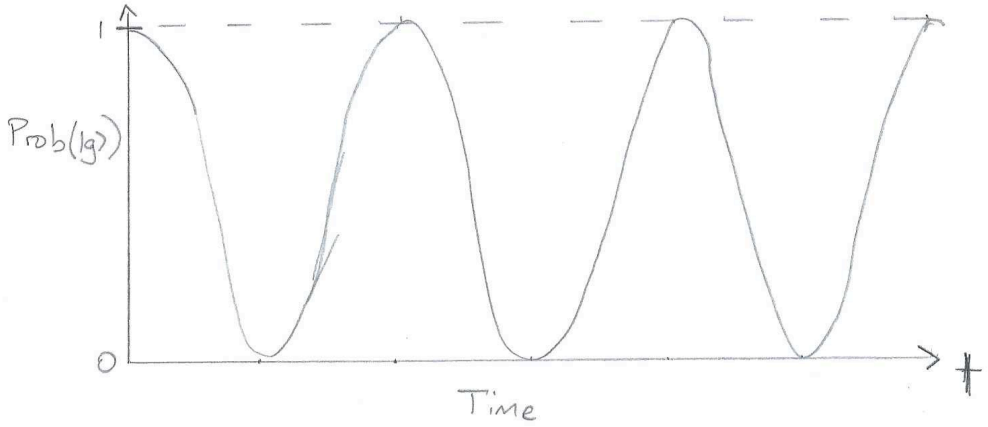


However, decoherence will cause the fringes to decay: Suppose the environment measures the state at time $t$ which is not one of the peaks. This is a version of an error source known as *dephasing* and results in the loss of relative phase information between the $|0\rangle$ and $|1\rangle$ states. We get a mixture of g vs. e, then the oscillation continues from there, leading to an oscillation with lower peaks. Really this is a continuous process, so the peaks gradually get smaller with time. How long the decay takes is a measure of the decoherence time.

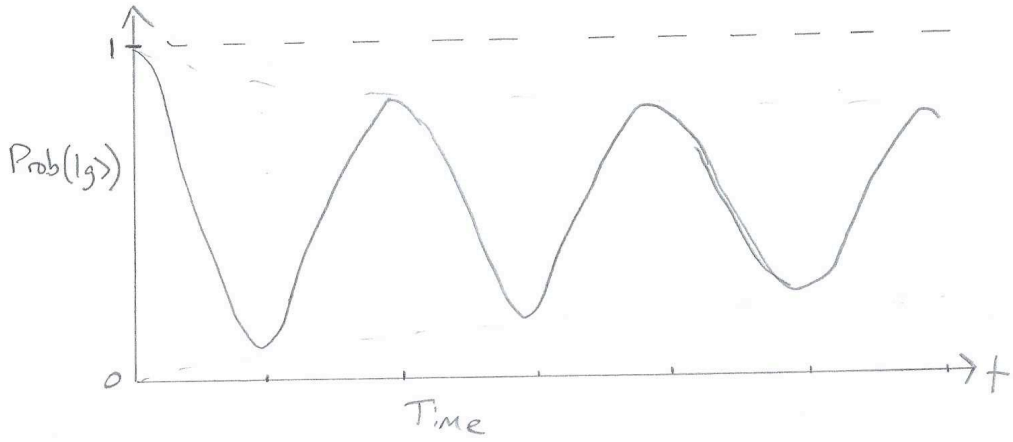


In ion traps, the atoms are very well insulated from the environment, so decoherence is low. The same technology used for ion traps is also used to make atomic clocks. We still get dephasing due to interactions with stray atoms nearby, but that is rare. There is also dephasing due to field fluctuations in the magnetic trap holding the atoms, which changes the energy levels slightly, leading to unwanted phase shifts. Spontaneous emission from excited to ground causes a different type of error known as *amplitude damping*, but that is much rarer in ion traps. Because of the low decoherence rate, you get $10^4$ or more periods before seeing significant decay. This is why ion traps are a good candidate for quantum computing, and in particular why they satisfy condition 3 of the DiVincenzo criteria.

The Rabi oscillations also give us an opportunity to do single-qubit gates. By turning the laser on for an appropriate period of time, we can do a bit flip or a gate creating superpositions such as the Hadamard gate. The duration of a pulse can be named by the angle it goes through. Thus, a bit flip is a $\pi/2$ pulse and a Hadamard is $\pi/4$ pulse. Note that a $\pi$ pulse restores us to the initial state, but with global phase $-1$. (This will be important later.)

Diagonal phase gates are in a sense even easier. If the laser is off, time evolution is just $e^{-iEt/\hbar|e\rangle\langle e|}$. That is, $|e\rangle$ acquires over time a relative phase compared to $|g\rangle$. By keeping track of time, we can do whichever phase gate we want. Note that this is happening whether we want it to or not, so we always need to keep track of time with every gate to make sure we know the correct relative phase of $|g\rangle$ and $|e\rangle$.

(Actually, in a multiple-qubit ion trap computer, this won't work for diagonal gates unless you want to apply the same phase gate to every qubit. But by fiddling with the relative phase of the lasers on different qubits, you can accomplish the same thing.)

Rotations about two axes let us get any single-qubit rotation using Euler angles.

### 1.2.3 Measurement and preparation

The usual way measurement is done is to use an additional excited state $|f\rangle$. Unlike $|e\rangle$, $|f\rangle$ has a short lifetime and rapidly decays via spontaneous emission to $|g\rangle$. Suppose we shine a bright laser on the atom, with a frequency resonant to the $gf$ energy difference.

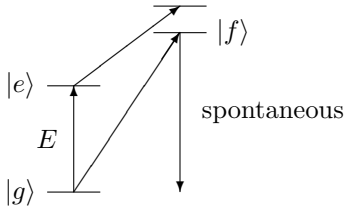


If the state of the atom is $|g\rangle$ initially, it will now rapidly cycle between $|g\rangle$ and $|f\rangle$. Sometimes, while in $|f\rangle$, it will spontaneously decay back to $|g\rangle$ rather than the stimulated decay that is part of the Rabi osciallation between $|g\rangle$ and $|f\rangle$. This emits light in a different direction, which can be picked up by a photodetector. On the other hand, if the state of the atom was $|e\rangle$, the laser doesn't interact with the atom, so no light is emitted away from the laser. Consequently, we have a measurement in the $|g\rangle, |e\rangle$ basis. This measurement is extremely accurate (condition 5).

Once we know whether we have $|g\rangle$ or $|e\rangle$, we can easily reset the qubit to $|g\rangle$ if it is not there already, using a resonant laser pulse (condition 2).

### 1.2.4 Two-qubit ion traps and gates

We can put multiple ions in the same linear ion trap, which strongly confines the ions except in one direction, and weakly confines the ions in that direction. Because they are all positively charged, they repel each other and form a line. If one of them moves, it pushes the others. When atoms get too far from each other, the trap pushes them back together. Therefore, the natural motion is a combination of different kinds of oscillation, broken down into *normal modes* that repesent different kinds of motion that can be analyzed separately. Because these are single atoms, the oscillation is quantized and can be described via *phonon* modes. A phonon is a quantum unit of vibration. In particular, the same phonon mode is shared by all the atoms.

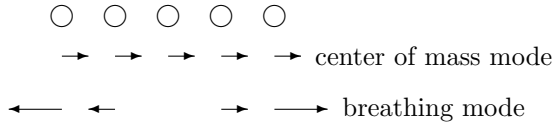


Let us a take a short digression to talk about harmonic oscillators. They show up in many places in physics (and will show up again when talking about other implementations). One way they can arise is from

a Hamiltonian $H = \hat{p}^2/2m + \frac{1}{2}m\nu^2\hat{x}^2$. The classical version of this Hamiltonian is a mass on a spring. As the spring stretches or shrinks from its default position ($x = 0$), there is a restoring force which is linear in the distance from $x = 0$, resulting in a quadratic energy. It shows up many places because if you have a potential $V(x)$, you can write it as a Taylor series $V(x) = V(0) + V'(0)x + \frac{1}{2}V''(0)x^2 + \ldots$. When a particle is at a local energy minimum, $V'(0) = 0$, so the first non-constant term is the $x^2$ term, and when it is close to the minimum, the higher order terms are smaller.