41076: Methods in Quantum Computing

`Quantum computing stack: hardware`

Dr. Mária Kieferová based on the materials from Dr. Min-Hsiu Hsieh
Centre for Quantum Software & Information, Faculty of Engineering and Information Technology, University of Technology Sydney

Abstract

Contents to be covered in this lecture are

1. Quantum computing stack
2. DiVincenzo's criteria
3. Decoherence in a quantum system
4. Tomography
5. Selected physical architectures - ion traps

1 The quantum computing stack

Now when we have the fundamental tools to discuss quantum information, we will turn our attention to the quantum software stack starting with its lowest level.

This figure does not include supporting quantum technologies such as refrigeration and supporting electronics necessary for making qubits stable and controllable. In many cases, individual levels can be mostly isolated so that progress is being simultaneously in many places.

2 Experimental requirements

The necessary properties of a quantum system that can be used for quantum computation are summarized by DiVincenzo's criteria:

1. A scalable physical system with well-characterized qubit
2. The ability to initialize the state of the qubits to a simple fiducial state
3. Long relevant decoherence times
4. A "universal" set of quantum gates
5. A qubit-specific measurement capability

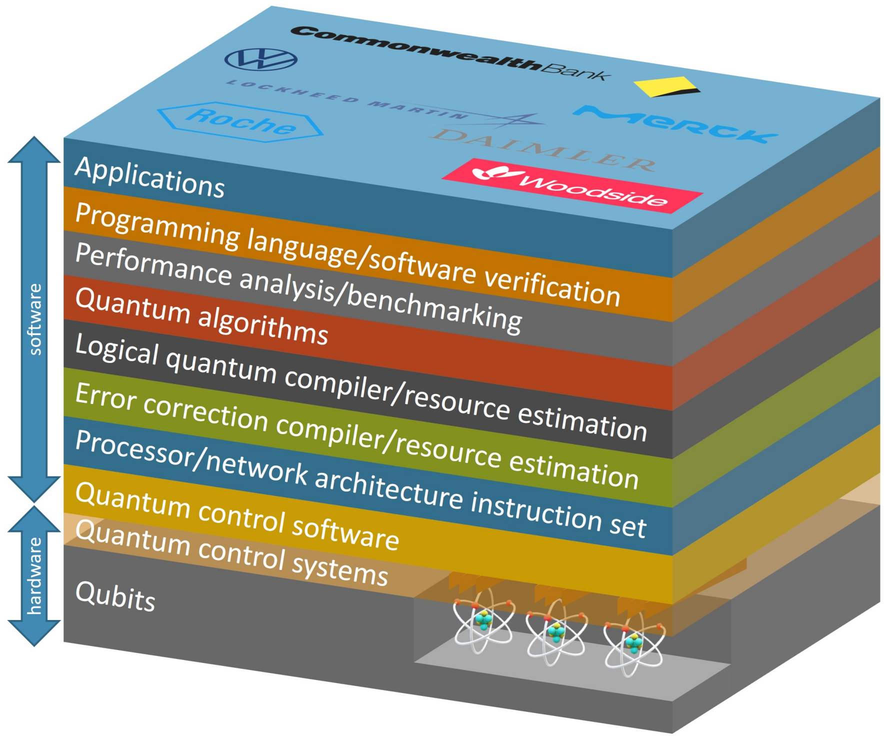

Figure 1: Quantum computing stack. Credit: Michael Bremner

Note that all the requirements need to be satisfied simultaneously.

Let us look into each requirement in more detail. First, we need to ensure the ability to define a qubit, an isolated 2-dimensional quantum state. More often than not, a quantum system will have more than two levels. For instance, say we would like to use the energy states of an atom as a qubit, however, they have many more levels than two and many of them are unstable.

We require the ability to be always able to increase the number of qubits in our computer. For example, one could define a nuclear spin and an electron spin in an atom as a system of two qubits, however, without the ability to interact with additional qubits, this system is not scalable. Another example is an early approach to quantum computing, nuclear magnetic resonance, however, it is now widely believed that this approach doesn't scale above a dozen qubits.

Besides fundamental limits to scalability, there are also practical constraints, such as the physical size of a chip that would contain 1000s of qubits. What needs to be kept in mind is that we still need to have access to control and measure individual qubits, which leads to additional constraints to scalability.

The second requirement is to be able to initialize the initial state of a quantum computer. A typical state for use in many calculations is the state $|0\dots0\rangle$, which is a pure state. However, contact with the environment leads to decoherence, i.e. noise. A difficulty in some systems is initializing all the qubits very close to the state $|0\dots0\rangle$ without restoring to a measurement. In other systems, the sources that create states are probabilistic, and creating a state with multiple qubits is in practice difficult (for example photonics).

The third requirement is that our qubits need to be stable for the duration of computation. We will learn about decoherence later in this class. I will just note that the decoherence times need to be compared to gate times in the same architecture rather than taking them as absolute numbers.

Apart from keeping our qubits stable, we need to be able to apply gates to them, specifically a set of gates that allows us to perform arbitrary quantum computation. Gates $\{H, T, CNOT\}$ and $\{H, Toff\}$ both create universal sets. This criterion often makes satisfying criterion 3 difficult - creating qubits that do not interact with anything makes it difficult to perform the required interactions to perform gates.

Lastly, we need to measure our qubits. It is sufficient to perform single-qubit measurements at the end of the computation but sometimes different measurements can be beneficiary. Measuring individual qubits with sufficient precision can be challenging and sometimes give additional engineering constraints to a quantum computer. For example, qubits in nano-diamonds or photonics can operate at high temperatures but their measurement uses superconductors that need temperatures of the order of Kelvins to function.

3 Quantum hardware

We described quantum computation as unitary operations applied on qubits followed by measurement. However, in practice, our qubits are not 2-dimensional and they don't hold quantum states for an extended time, we don't apply the unitaries we wish to but instead different channels, and the system gets measured by the environment instead of us.

One of the key enemies of quantum computation is decoherence - an unwanted interaction between our quantum state and the environment. Examples of decoherence are dephasing and depolarizing channels that we learned about last time.

Decoherence rapidly (exponentially) deteriorates the state of qubits which can be characterized by two times:

- T1 measures how fast a qubit loses energy. Often, the state $|0\rangle$ is encoded to the ground state and $|1\rangle$ into an excited state (i.e. state with higher energy). T1 measures the exponential decay time for a qubit to relax from $|1\rangle$ to $|0\rangle$.
- T2 measures the stability of a phase of a qubit. Starting from a particular state on the "equator" of a Bloch sphere, for a time $t \geq T2$ the phase disappears and the mixed state will be along the z (vertical) axis.

It should be stressed that while we have a description of noise such as depolarizing and dephasing channels, actual physical noise is more complex these channels are only simplified models.

After an interaction with an (unknown) noisy channel, our quantum state will end up in some unknown quantum state. A process of characterizing outputs of quantum computers is *quantum state tomography.*

Recall that if we measure a quantum state, we only sample from the possible outcomes. To learn a fully quantum-mechanical description of a quantum state, one needs to measure many copies (i.e. thousands for a single qubit) of a quantum state in different basis sets. If a particular measurement allows us to fully reconstruct a quantum state, we say that the measurement is tomographically complete.

Exercise 1. *A single qubit is fully characterized by a vector $\vec{r}, |r| \leq 1$ such that*

$$\rho = \frac{1}{2}I + r_0\sigma_x + r_1\sigma_y + r_2\sigma_z. \tag{1}$$

Take a set of operators

$$M = \left\{ \frac{I+X}{6}, \frac{I-X}{6}, \frac{I+Y}{6}, \frac{I-Y}{6}, , \frac{I+Z}{6}, \frac{I-Z}{6} \right\}. \tag{2}$$

Show that

1. *M is a POVM (operators are positive and sum to identity).*
2. *M is tomographically complete, i.e. measuring enough times will allow us to learn the vector r.*

Decoherence can be in practice minimized by building qubits that are very well isolated from the environment, often using high vacuum and dilution refrigerators. We need to have access to the qubits to perform desired interactions, i.e. quantum gates. We do not expect to ever reach noise levels that will be low enough to perform significantly long quantum computation - there will be a need to minimize errors through algorithms known as error correction.

Quantum gates are also never perfect and their characterization goes by the name *process tomography.* The formal definition of error rate is beyond the scope of this course but it quantifies how much an actual gate can deviate from an ideal one for the worst possible input. For example, take the unitaries Z and I. If we apply them to states $|0\rangle$ and $|1\rangle$, they appear to act the same. However, on the X basis, we would discover that they are in fact very different operations. Thus, comparing gates depends on the input.

In practice, experimentalists estimate how good their gate is by stating gate fidelity, i.e. their average performance. In 2022, 99% fidelity for 2-qubit gates and 99.9% for single qubits gates are considered to be very good numbers.

Exercise 2. *Suppose you have a 99% of percent of success when performing an operation. How many operations in sequence can you perform before the chance of successfully performing the sequence gets below 50%? You can assume that the errors are independent.*

Unfortunately, errors on quantum gates are not completely independent - performing two gates in parallel on nearby qubits can lead to cross talk.

Formally, the average fidelity of a channel is defined with respect to the identity channel

$$F(\mathcal{E}) = \int d\psi \langle\psi|\mathcal{E}(\psi)|\psi\rangle \tag{3}$$

as an average over all state fidelities. To obtain the average, we must integrate over all the quantum states in a given Hilbert space with equal weightings and satisfy $\int d\psi = 1$. This is known as integration over Haar measure.

Exercise 3. *Compute the fidelity of a qubit depolarizing channel* $\mathcal{E}(\rho) = (1-p)|\psi\rangle\langle\psi| + p\frac{I}{d}$.

We can then use the definition for any channel by seen as a perfect channel followed by a noise on the identity channel. Computing average gate fidelities can be further simplified using Nielsen's formula [2]. In a special case when the channel is unitary, we can compute its fidelity (with respect to the identity channel) as

$$F(U) = \frac{d + |Tr(U)|^2}{d + d^2}. \tag{4}$$

Exercise 4. *1. Verify that $F(I)=1$ in (4).*

2. The hottest quantum startup promises to do quantum computing by implementing Hadamard and Toffoli gates. However, they have a minor issue: their Toffoli gates are not working and they are simply doing nothing (i.e. identity gates). What is the fidelity of their "Toffoli" gate?

3. What if they replace all m-controlled-NOT gates with the identity?

Besides error rate and average fidelity, there are other properties of quantum channels such as channel capacity that are beyond the scope of this lecture.

4 Physical architectures

Here we will do a very brief overview of physical architectures that are popular in 2022. Quite frankly, we wouldn't do justice to any single architecture by trying to explain it within a lecture. Instead, we will only give a brief review of some of the strengths and weaknesses of different approaches without resulting in a full-on graduate-level physics class.

There are two broad approaches to making qubits - either one can use a physical object such as an atom or a photon and define states $|0\rangle$ and $|1\rangle$ on them or one can build their own qubits. Building qubits (sometimes referred to as "artificial atoms" but this is not a very good analogy) requires an extra step in building a quantum computer but it allows us to engineer desirable properties of qubits.

For each architecture, we need to ask what is its potential for satisfying DiVincenzo's criteria as well as consider the current level of progress. Currently, superconducting qubits have the highest number of highly controllable qubits followed by trapped ions (we are not including quantum annealers, analog computers, and non-universal systems here). At the same time, silicon qubits are some other architectures that are still working only on pairs of qubits. This is a large gap but it might not last forever. To illustrate a failure of an early advantage, in the early 2000s, NMR quantum computers were by far the most advanced systems but they pretty much disappeared in the last 10 years. However, having these early prototypes gave us a much better understanding of working with quantum systems and taught us how to build better quantum computers. Randomized benchmarking and Hamiltonian learning are two programs that emerged from early NMR computers but became standard tools across quantum computing.

4.1 Trapped ions

One approach to building a quantum computer uses two energy levels in a charged particle (ion) as a qubit. The ions are trapped in a dynamic electromagnetic field.

A qubit is encoded into two energy levels (either in the hyperfine structure of an ion or in a ground state and an excited state). Single qubit gates are performed by applying an external electromagnetic field to the atoms. The ion chain itself acts as linear harmonic oscilator and can allow us to couple the ions.

2-qubit gates can be implemented using auxiliary states in the ions and utilizing the Coulomb interaction between ions. Another approach is to instead implement the Mølmer–Sørensen gate that makes use of an interaction with an external field and can be extended into a many-qubit gate.

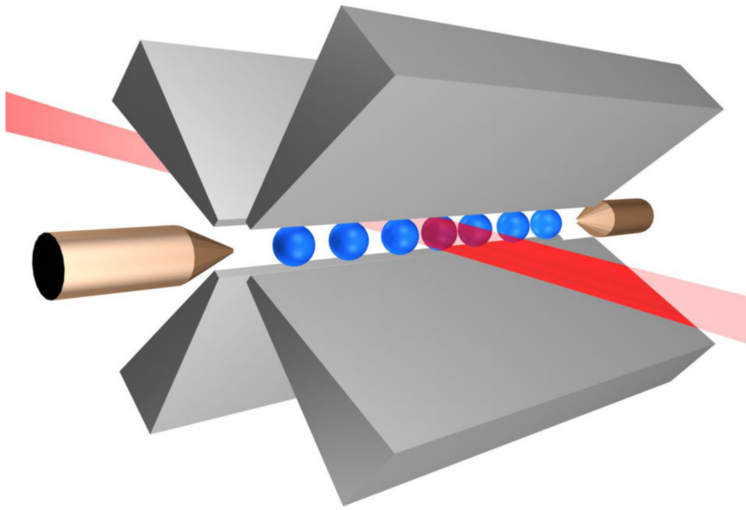

Figure 2: Linear ion trap. Source: Institute of Theoretical Physics, Innsbruck

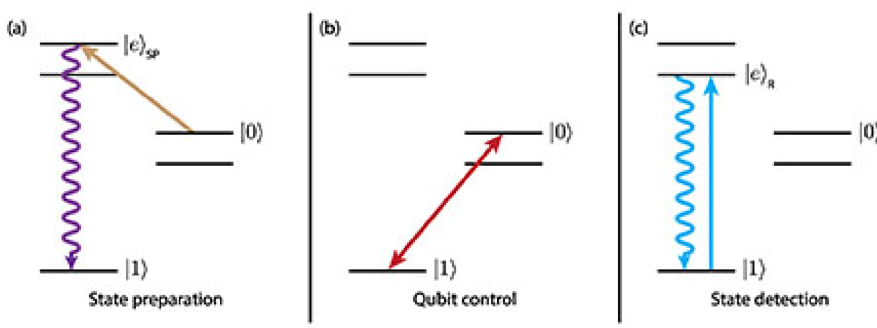

Figure 3: Using levels of an atomic structure for quantum computation. Source: MIT

Ion trap quantum computers are currently pursued by a number of academic groups and companies IonQ and Quantinuum.

+ ions create identical qubits (but the control is not uniform)
+ lowest gate errors out of all approaches
– very slow gates
– About 50 ions is the maximum for a trap without using individual control. Coupling different traps has not been very successful so far.

The issue with ion straps is their scalability -as the size of a trap grows, the qubits became unstable and hard to control. There are a number of proposal to address the problem, namely entangling traps, physically moving ions between them or redesign the traps completely. There is a significant effort to advance ion trap quantum computing but at the time of writing it is not clear if scalability can be significantly improved.

For more information about trapped ions on an accessible level see `https://pennylane.ai/qml/demos/tutorial_trapped_ions.html`

Further Reading

The good people from Xanadu(photonic quantum company) made a series of very accessible tutorials on quantum architectures at `https://pennylane.ai/qml/demos_quantum-computing.html`.

A seminal text on ion trap architectures [1]

References

[1] David Kielpinski, Chris Monroe, and David J Wineland, *Architecture for a large-scale ion-trap quantum computer*, Nature **417** (2002), no. 6890, 709–711.

[2] Michael A Nielsen, *A simple formula for the average gate fidelity of a quantum dynamical operation*, Physics Letters A **303** (2002), no. 4, 249–252.