41076: Methods in Quantum Computing

`The quantum computing stack and 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. Physical architectures

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

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

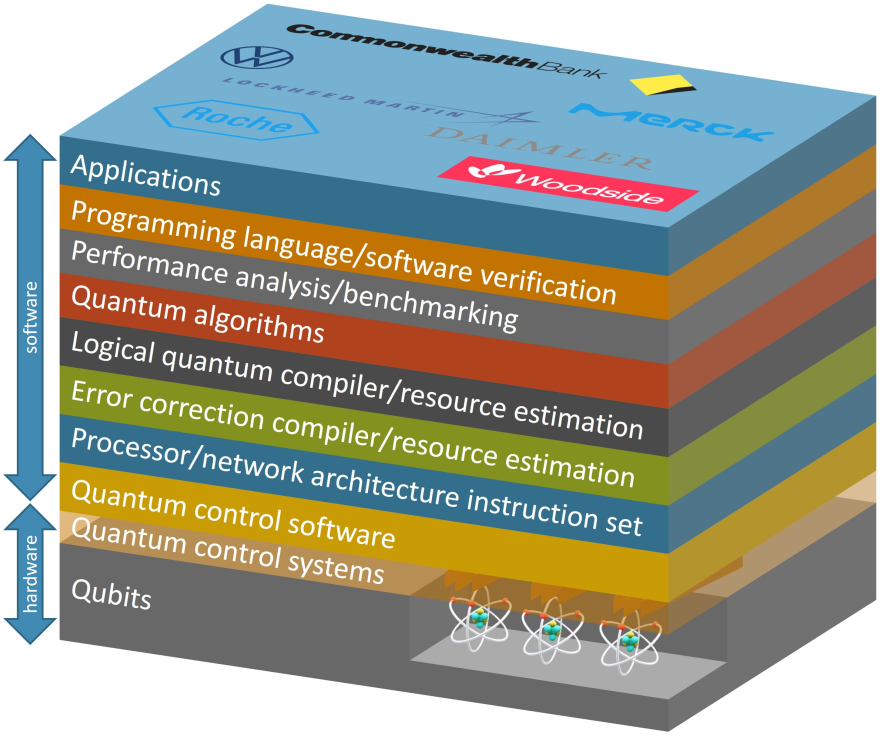

Figure 1: Quantum computing stack. Credit: Michael Bremner

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 $|0\ldots 0\rangle$, 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 $|0\ldots 0\rangle$ 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 Decoherence

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 which 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. a 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.

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 1. *Suppose you have a 99% 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.

3.1 NISQ

NISQ stands for Noisy Intermediate Scale Quantum which means quantum computers with dozens to hundred qubits without full error correction. These systems themselves can be too large to classically simulate (i.e. keep the wave function stored in a classical memory and apply gates by matrix multiplication or similar) and in some cases, can support computation that is not tractable by classical at that time. This is known as beyond-classical computation and was demonstrated by Google in 2019, although classical algorithms and computers improved since and can solve the same problem today.

Unfortunately, problems used for the beyond-classical experiments are not practical problems faced by industry or other scientists. Instead, they are artificial problems designed to maximize the gap between classical and quantum computers. A significant amount of research has focused on using a NISQ computer for problems of practical interest, specifically in areas of quantum chemistry, optimization, and machine learning. These approaches face two major drawbacks:

- The number of qubits is very limited, often allowing for input of very small instances. The connectivity of qubits is also often limited.
- All the gates have a chance of error and qubits themselves decohere. Thus, it is possible to execute only very short circuits until the computation becomes dominated by noise.

NISQ still lacks truly practical applications that go beyond demonstrations on small instances or simulating physical phenomena that can be solved classically (for example, braiding of anyons). While trying things on small problems first makes sense, there are many NISQ approaches that are not scalable.

3.2 Fault tolerant computation

Since quantum operations are never perfect, it seems that quantum computation would be inherently unstable and errors would eventually accumulate. Fortunately, it was shown that if the noise level is low enough, we can use redundancy and clever trick to detect and correct the errors. This process is known as error correction and quantum computers can are able to correct errors and run effectively error-free and are known as fault-tolerant.

4 Physical architectures

Here we will do a very brief overview of physical architectures that are popular in 2022. 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 Quantum physics

One of the observations that led to the discovery of quantum physics is that atoms have discrete spectra: there are only certain allowed energy levels that an electron can stay in. Electrons can also absorb particles of light (photons) and if the energy of a photon corresponds to the difference between allowed energy levels, the photon is absorbed by an electron and the electron can become excited to a higher energy state. Similarly, an electron in an excited state can emit a photon and fall to a state with a lower (allowed energy). The photon will have an energy corresponding to the difference between the energy levels and we can also compute its frequency f and wavelength λ as

$$f = \frac{E}{h}, \qquad \lambda = \frac{hc}{E} \tag{1}$$

where E stands for energy, h is the Planck constant, and c is the speed of light. We can see the energy levels and corresponding wavelengths of absorbed/emitted light of hydrogen

The energies of a system depend on the outside potential, interactions within a system, and kinetic energy of particles within a system. For an electron in hydrogen, this would correspond to a Coulomb potential caused by the nucleus and the electron's own potential energy. For helium, we would need to consider the Coulomb potential and kinetic energy of both electrons as well as the energy corresponding to the repulsive interaction between them. Mathematically, all the energy contributions are expressed by summing them to an operator known as Hamiltonian. The allowed

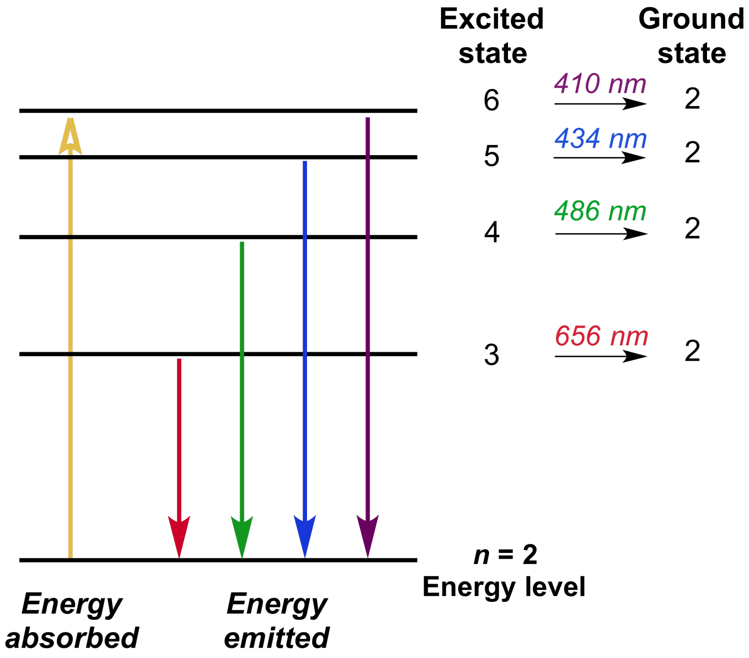

Figure 2: The simplest version of energy levels and spectrum of hydrogen.

energies are then the eigenvalues E_i of the Hamiltonian and allowed states are their corresponding eigenvectors $|\psi_i\rangle$

$$H|\psi_i\rangle = E_i|\psi_i\rangle. \tag{2}$$

Hamiltonian will also determine how the system evolves in time through the (time-dependent) Schrödinger equation

$$i\hbar\frac{d}{dt}|\psi(t)\rangle = H|\psi(t)\rangle. \tag{3}$$

The state with the lowest energy is known as the ground state.

The energy levels will be affected by an outside electric or magnetic field. It is *in principle* possible to define two of the levels as a qubit and apply a photon light beam to perform operations on it. Such a qubit would be incredibly unstable and hard to control but we will see that we can further refine this idea to propose more practical qubits.

The physics we described only applies perfectly at 0K. Low temperatures can approximate this behavior, but if the particle is at a higher temperature, it has more energy and can bounce between different levels which leads to thermal noise.

4.2 Spin qubits

Let us start with an electron spin, specifically an electron spin of a phosphorus atom surrounded by silicon. The spin acts as a small magnet and reacts to an outside magnetic field - it will line up with the outside magnetic field to its ground state, qubit state $|0\rangle$. The ground state will be stable if the temperature is very low, in practice people use several millikelvins, otherwise, the spin can spontaneously flip due to a thermal fluctuation. The opposite orientation will be our excited state, $|1\rangle$. We can then apply a pulse with a frequency corresponding to the energy difference between a ground state and the excited state. The length of the pulse is also important - during the duration of the pulse, our spin qubit will be rotating from $|0\rangle$ to $|1\rangle$. This allows us not only to apply the X gate but also to prepare superpositions. The excited state has higher energy than the surrounding electrons in silicon, which allows the excited electron to be absorbed by silicon, leading to a current that is detectable in experiments.

Instead of an electron spin, it is possible to use a nucleus spin as well. Nuclear spin is much weaker but also quite stable and it can be measured using an interaction between the nuclear spin and an electron spin and subsequently by observing a current.

There is a plethora of approaches to quantum computing based on spin. Sydney (UNSW) is one of the world leaders in silicon qubits with different groups pursuing either nuclear or electron spin and different manufacturing techniques. While building spin qubits was proposed in the 90s, manufacturing the first qubits and their interaction proved to be challenging. Two qubit fidelities above 99% were recently demonstrated by three different groups including UNSW.

+ some approaches have very good prospects for scalability

+ error rates below the threshold have been demonstrated

+ very fast gates (but perhaps too fast)

− building the first few qubits is incredibly challenging

4.3 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.

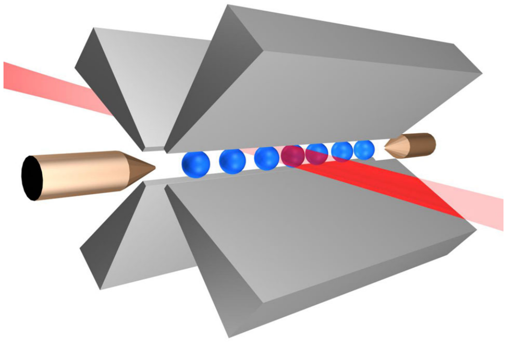

Figure 3: Linear ion trap. Source: Institute of Theoretical Physics, Innsbruck

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 a linear harmonic oscillator 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.

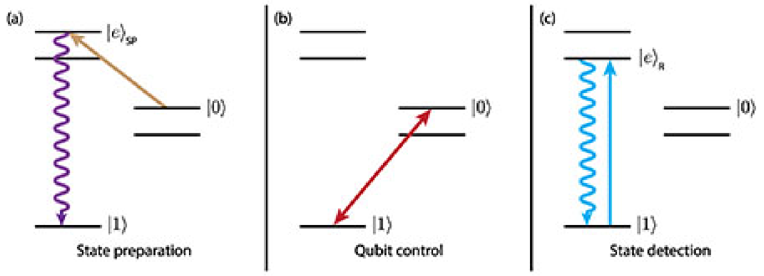

Figure 4: Using levels of an atomic structure for quantum computation. Source: MIT

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

- \+ ions create identical qubits (but the control is not uniform)
- \+ lowest gate errors out of all approaches
- – each qubit requires a significant physical space
- – 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 proposals to address the problem, namely entangling traps, physically moving ions between them, or redesigning 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`

4.4 Photonics

Photons are elementary particles that can be used as qubits. Different properties of a photon can be used for computation. One is to give the photon two possible paths it can travel in and call one of their state $|0\rangle$ and another [$|1\rangle$. These states are known as modes. Another approach is known as continuous variables quantum computation, which uses the continuous wave function of a photon

$$|\psi(x)\rangle = \int_x \psi(x)|x\rangle \tag{4}$$

expressed in the basis of position eigenstates. The same wave function can be also expressed in momentum eigenstates, however, it is impossible to determine both the position and the momentum

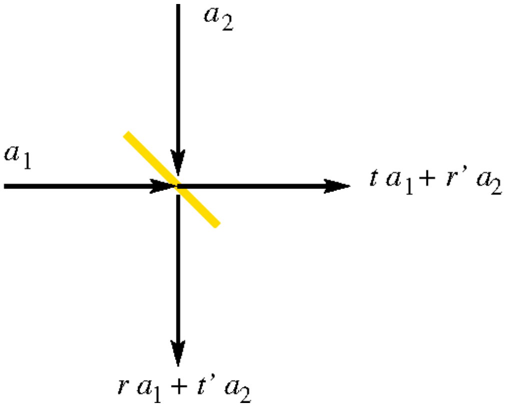

Figure 5: Beamsplitter can split a beam of light into two. If the light consists of a single photon, it will create a superposition across two different modes. Credit University of Potsdam

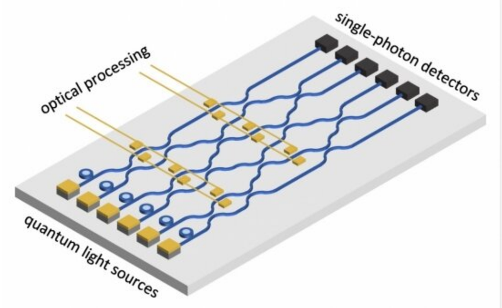

Figure 6: Quantum photonic on a chip. Source: Galan Moody

of a photon (or another particle for that matter) because of the Heisenberg uncertainty relation. This approach is pursued by Xanadu and differs a lot from the techniques discussed in this course but ultimately leads to an equivalent computational model.

Another approach is to take single atoms and use a beamsplitter to create two modes 5.

In the qubit formalism, a beamsplitter enacts a matrix operation

$$\begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \hat{c} \\ \hat{d} \end{pmatrix} \tag{5}$$

where a, b, c, d correspond to the 4 different paths. Other optical elements used are mirrors and phase shifters.

These elements are linear - a response to two photons is a linear combination of responses to each photon. Full quantum computation requires the addition of a nonlinearity. Nonlinear materials exist but they are very lossy creating decoherence. Another approach was proposed by Knill, Laflamme, and Milburn (KLM protocol). This approach shows how to perform universal quantum computation using only linear photonic, ancillae, and measurements. Based on the outcome of the measurement, further gates are applied adaptively.

Photonic quantum computation can be performed either in optical waveguides on q chip as in Fig. 6 or in free space. Using a chip is much more convenient for scalability but free space allows for sensing qubits over large distances for quantum communication.

In summary

- + photons make "perfect" qubits
- + qubits don't need to be cooled (but still need low temperatures for superconducting detectors)
- + low intrinsic decoherence
- – lack of single photon sources
- – many gates are probabilistic
- – since photos travel at a speed of light, gates need to be perfectly timed
- – difficult error-correction

4.5 Superconducting qubits

Superconducting qubits are one of the most successful qubits in 2022. Superconducting qubits are a type of "artificial atoms" created from small electric circuits with superconducting elements.

When a voltage U is applied to a regular conductor, a current starts I flowing through the circuit that is proportional to the voltage and inversely proportional to the resistance R of the circuit

$$I = \frac{U}{R}. \tag{6}$$

At very low temperatures (in the order of Kelvins, room temperature 300K), some materials become *superconductors* and have zero resistance. One of the effects that can be observed is that a current can flow through a superconductor without any voltage applied. Current is caused by a flow of electrons in a medium but in the regime of superconductivity, the electrons pair up into *Cooper pairs*. Cooper pairs observe different statistics that electrons and in the superconductive regime, they can be all in the same quantum state. This collective behavior leads to quantum mechanical effects that are observable on a macroscopic scale.

Different superconducting qubits are defined by their electric circuits that consist of standard electrical components (voltage, capacitors, inductors . . .) and also Josephson junctions. Josephson junction consists of two superconductors separated by a very thin barrier created by an insulator, see Fig. 7. The Cooper pairs use quantum tunneling to travel through the barrier resulting in a measurable current in the circuit.

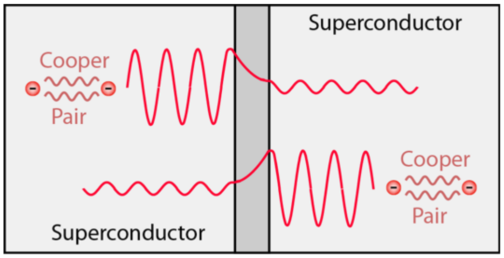

Figure 7: Josepson junction is an essential circuit element of superconducting qubits, source `http://hyperphysics.phy-astr.gsu.edu/`

Josephson tunnel junction behaves like a non-linear inductor and by adding a capacitor to the circuit, we can create an LC resonator and compute its quantum properties. Different circuits have different energy levels, see 8. It is important to have different spacing between individual levels, otherwise, we wouldn't be able to control our qubit.

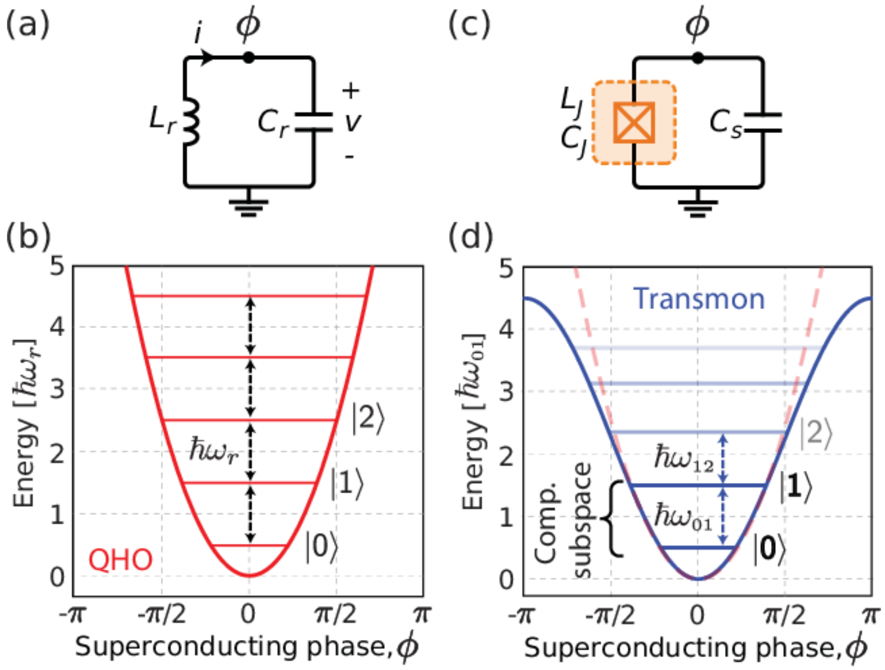

Figure 8: Quantum physics of electric circuits. Source [2]. Figures a) and b) create a quantum linear harmonic oscillator that is not suitable as a qubit because its energy levels are uniformly spaced. In c) and d) the inductor is replaced with Josephson junctions that add a nonlinearity which allows us to define a qubit in the lowest two levels. The box with a cross includes two Josephson junctions in parallel.

There are many types of superconducting qubits but in recent years transmons (and their variations) emerged as the leader for their controllability and robustness towards the noise. Both Google and IBM use transmons, in their architectures with 72 and 127 qubits respectively. The superconducting elements are made out of aluminum and kept at millikelvin temperatures and set on silicon wafers. The qubits can be controlled with electronics in GHz in the microwave regime.

While transmons are a leading quantum computing architecture with the highest number of qubits at the moment, they do have challenges. The strengths and weaknesses are summarized below:

+ Error rates are relatively low

− The qubits must be kept at mK temperatures. This is possible but requires a dilution refrigerator.

− The qubits are quite large, coupled with control electronic makes building chips above 1000s of qubits too large for standard dilution refrigerators

Further Reading

The good people from Xanadu 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]

A review of silicon quantum [4].

An essay about NISQ [3] is now slightly out-of-date but it informed a lot of the work in quantum in the last few years.

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] Philip Krantz, Morten Kjaergaard, Fei Yan, Terry P Orlando, Simon Gustavsson, and William D Oliver, *A quantum engineer's guide to superconducting qubits*, Applied Physics Reviews **6** (2019), no. 2, 021318.

[3] John Preskill, *Quantum computing in the nisq era and beyond*, Quantum **2** (2018), 79.

[4] Floris A Zwanenburg, Andrew S Dzurak, Andrea Morello, Michelle Y Simmons, Lloyd CL Hollenberg, Gerhard Klimeck, Sven Rogge, Susan N Coppersmith, and Mark A Eriksson, *Silicon quantum electronics*, Reviews of modern physics **85** (2013), no. 3, 961.