Lecture 1 – QISE – Bits and Qubits, An Introduction

Intro to AQET NRT

In this second quantum revolution, dubbed Quantum 2.0, society will leverage the quantum-mechanical properties of light and matter to enable new technologies in computation, communication, and sensing.

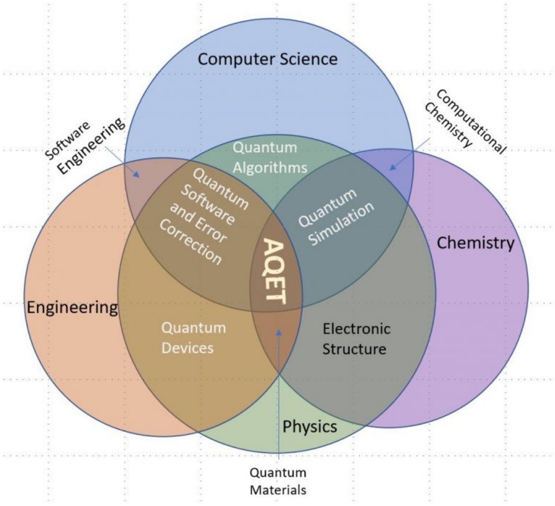

Accelerating Quantum-Enabled Technologies (AQET) is a National Science Foundation Research Traineeship program (NRT) at the University of Washington that seeks to train the next generation of scientists and engineers to enable this revolution.

AQET expands on UW faculty research in quantum information science and engineering (QISE) by establishing a unique curriculum that bridges the gap from physics to chemistry, computer science & engineering, electrical & computer engineering, and materials science & engineering. Quantum 2.0 will require teams of expertise from across disciplines, and AQET is one of the first programs to bring hardware and software scientists and engineers together at the trainee level.

AQET Courses (From Proposal)

Course 1: Introduction to Quantum Information/Quantum Computing addresses the lack of QIST training currently offered in curricula outside physics. These courses will be offered during Autumn quarter. Three different course options will be offered, given the different domain-specific backgrounds of the students and the types of different research problems.

- **Introduction To Quantum Information Science and Engineering For Chemists And Materials Scientists (CHEM 561):** This course targets chemistry and materials students, using familiar language and framework to explore core concepts of QIST, including computation, logic, and quantum physics. This course utilizes three hardware platforms to help students understand the core scientific concepts that govern their operation and how materials play a role: trapped ions, defects in crystals, and superconducting qubits. An introduction to sensing, communication, and computation applications is provided at the end of the course through student presentations.
- **Introduction to Quantum Computing (CSE 599Z):** This course is primarily targeted at CSE and computationally-oriented ECE students, providing an introduction to quantum computing and quantum algorithms. No prior knowledge of quantum mechanics is assumed. This course covers topics like the circuit model of quantum computing, query complexity, and quantum complexity theory.
- **Quantum Information (PHYS 521A):** Offered in physics, this course targets physics students and interested CSE and ECE majors with appropriate background in quantum mechanics. This course provides students with the background needed to understand

modern quantum information hardware on both the physical and the mathematical levels by focusing on both implementations of quantum computing and quantum information theory.

Course 2: Implementations in Quantum Information (CHEM 560) brings together all students in an innovative, project-based course that highlights the challenges in implementing quantum information systems. The course combines the different skills sets to implement and characterize quantum information processing performance on IBM quantum computers using the Qiskit platform. Topics include quantum tomography, entanglement witnesses, randomized benchmarking, and quantum control. This course will be offered in Winter quarter.

Course 3: Advanced Topics in Quantum Information the third phase of the program, encompasses a range of domain-specific courses in advanced topics. Many different courses can satisfy this requirement.

What is QISE?

Quantum information science and engineering encompasses a lot of different areas including computing, sensing, communication, cryptography, etc. At the heart of all of these is the qubit or quantum bit. To understand what a qubit is, it helps to first establish what a bit is. A bit (short for binary digit) is the smallest unit of data in computer a computer – it represents a logical state and has a single binary value, either 0 or 1 (true/false, on/off, +/-). The correspondence between these values and the physical state of the underlying system is dependent on user/programmer encoding.

We can draw an analogy or contrast between the classical bit and a simple light switch – it is either on or off and that state is not dependent on me being in the room to observe it, just by the position of the switch.

So what about a quantum bit? Let’s consider our light switch again. A quantum light switch could be represented by a sphere with the north pole being OFF and the south pole being ON. Imagine I put this special switch in a room separate from the light and control the switch by putting my finger on the sphere. If I put my finger on the north pole, the light is definitely off and in the south pole, the light is definitely on. You can go to the room and check – this is always the case. If I move my finger anywhere else on the sphere the light might be on or off when you check. If you don’t check, the light is in an indeterminate state…. It’s not dimmed, it just exists with some probability of being on or off when seen. As soon as I walk into the room, the light is again on or off – the act of observation forced it to be so. This is a 2-state quantum system.

A quantum bit or qubit is another such 2-state quantum system and it is the basic information unit in quantum computing. With a qubit we replace the on/off (0/1) with $|1>$ and $|0>$. The important point about a qubit vs a classical bit is that it can be in a state other than 0 or 1 – it is possible to form linear combinations of states, which we call superpositions:

$$|\Psi> = \alpha|0> + \beta|1>$$

Remarkably, we can’t examine a qubit to determine its quantum state – we can’t determine a and b. Instead, quantum mechanics tells us we can only learn a much more restricted piece of information about the qubit. When we measure it we get either the result 0 with probability $|\alpha|^2$ or 1 with probability $|\beta|^2$ (and $|\alpha|^2 + |\beta|^2 = 1$ since the probabilities must sum to 1).

This dichotomy between the unobservable state of the qubit and the observations we make is at the heart of quantum computation and quantum information. It makes it hard to intuit the behavior of the quantum system, but there is an indirect correspondence because qubit states can be manipulated and transformed in ways that lead to measurement outcomes that depend on the different properties of the state. So these quantum states have real, experimentally observable consequences.

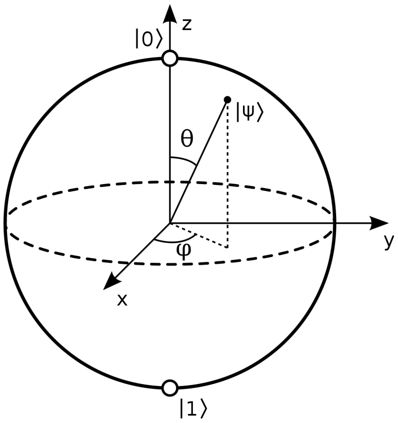

Because $|\alpha|^2 + |\beta|^2 = 1$, we can define $|\Psi>$ using a 3D sphere, just like we represented our quantum light switch. Here

$$|\Psi> = \cos(\theta/2)|0> + e^{i\varphi}\sin(\theta/2)|1>$$

This is called a Bloch sphere and our position on the sphere is defined by two angles, θ and φ.

We know a classical bit stores two values only. We can combine them into large units to store a greater range of values. How much information can a qubit store? There are an infinite number of points on our unit sphere, so in principle we could store the entirety of Romeo & Juliet with this single qubit… but this is misleading because a qubit will give only either 0 or 1 when observed. And measurement changes the state of a qubit, collapsing it from the superposition of |0> and |1> to the specific state consistent with the measurement. So from a single measurement of a qubit, we only get a single bit of information about the state of the qubit. But how much information is represented by a qubit if we do not measure it? This is a trick question because how much information can be represented by a qubit if we don’t measure it? 0… Nevertheless, there is something really important here because when Nature evolves a closed quantum system of qubits, she keeps track of all the continuous variables describing the state, like a and b. So in the state of a qubit, Nature conceals a great deal of hidden information and as we will see this extra information grows exponentially with the number of qubits. Understanding this hidden quantum information is at the heart of what makes quantum mechanics a powerful tool for

# of qubits	# bits / # loops	RAM
1	2	2 bits
2	4	4 bits
3	8	1 byte
4	16	2 bytes
5	32	4 bytes
6	64	8 bytes
7	128	16 bytes
8	256	32 bytes
9	512	64 bytes
10	1024	128 bytes
11	2048	256 bytes
12	4096	512 bytes
13	8192	1 kB

# of qubits	# bits / # loops	RAM	Time
13	8192	1 kB	$2.73\text{x}10^{-6}$ s
20	1048576	128 kB	$3.5\text{x}10^{-4}$s
23	8388608	1 MB	$2.8\text{x}10^{-3}$s
33	8589934592	1 GB	2.9 s
43	$8.8\text{x}10^{12}$	1 TB	49 mins
53	$9.0\text{x}10^{15}$	1 PB	35 hours
63	$9.2\text{x}10^{18}$	1 EB	97.5 years
1000	$1.1\text{x}10^{301}$	$1.3\text{x}10^{282}$ EB	$1.1\text{x}10^{284}$ years

information processing. (If we look at two bits, they can take the following values: 0, 0 // 0, 1 // 1, 0 // 1, 1. Two qubits take all those values at once. One qubit can take the value of two bits. Two qubits can take the values of four bits. In general, n qubits can take the values of 2^n bits.)

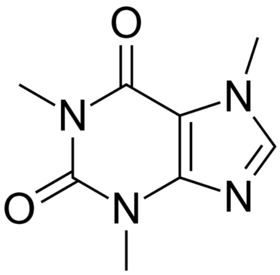

So what? Let's take chemistry as our example – if we could do chemistry in a computer instead of a beaker, we would need it to happen with full fidelity. The atoms and molecules modeled by the computer should behave exactly like they do in the laboratory. If we could build such a full fidelity model, we could more efficiently solve many global challenges – cure diseases, figure out how to feed our planet, discover new materials, understand how proteins function, etc. Is this plausible? Let's consider a simple molecule that we all know and love to see… caffeine (1,3,7-trimethylxanthine).
An 8 oz cup of coffee has 95 mg caffeine = 2.95 x 10^{20} molecules. That's a big number. Scientists estimate that there are 10^{49}-10^{50} atoms on our planet. [A gigabyte (1 billion bytes) is 10^9 bytes (1 byte = 8 bits).] But we're not asking to model all the molecules in a cup of coffee, just 1 at a single instant. Caffeine is small – it is made up of 24 atoms and can exist in 10^{48} distinct quantum states (configurations of these atoms). That means for a classical computer to perfectly represent caffeine would require 10^{48} bits… this is comparable to 1-10% of the number of atoms on earth! And this is just one small molecule! Somehow Nature manages to deal with all this information… How? We don't know! What we know is that we have no hope of providing enough traditional computer storage to hold this much information… Nature isn't classical (Feynman). But, 160 qubits could hold 2^{160} = 1.46 x 10^{48} bits while they were involved in computation. I'm not saying how we would get all the data into those qubits and I'm not saying how many more we would need to do something interesting with that information. But it does give us hope!

So, with this brief intro, where are we going in this class?
Outline of syllabus and learning objectives.

Week 1, September 28/30: [Brandi] What is Quantum Information Science and Engineering? What is a quantum computer? Why does this matter?

- Friday Reading Discussion: Sutor Chapter 1
- National Strategic Overview for Quantum Information Science (September 2018)
  - https://www.quantum.gov/wp-content/uploads/2020/10/2018_NSTC_National_Strategic_Overview_QIS.pdf
- NSF QISE Research Page
  - https://www.nsf.gov/mps/quantum/quantum_research_at_nsf.jsp

Week 2, October 3/5/7: [Peter] Applications and Challenges: Quantum sensing, quantum communication, quantum computing
Friday Reading Discussion:

Week 3, October 10/12/14: [Peter / Sutor Ch. 2-6] Math Bootcamp: Imaginary Numbers and Linear algebra
Friday Reading Discussion:

Week 4, October 17/19/21: [Peter / N&C Ch. 2] Intro to Quantum mechanics
Friday Reading Discussion:

Week 5, October 24/26/28: [Brandi / Sutor Ch. 7] One Qubit
Friday Reading Discussion:

Week 6, October 31/November 2/4: [Brandi / Sutor Ch. 11 + supplements] What Does it Take to be a Qubit? The relationship between material properties and quantum memory/quantum coherence
Friday Reading Discussion:

Week 7, November 7/9/11: [Brandi / Sutor Ch. 8, 9, 11] Entanglement and multi-qubit gates --> quantum circuits and error correction
Friday Reading Discussion:

Week 8-11: My favorite qubit

- Week 8, November 14/16/18: [Brandi] Spins and Molecular Systems
  - Guest Lecture Stefan Stoll (UW Chemistry), Wednesday 11/16
- Week 9/10, November 21/28/30/December 2: [Peter] Defect Centers, Superconductors, Trapped Atoms/Ions
  - Guest Lecture Kai-Mei Fu (UW Physics), Wednesday 11/30
- Week 11, December 5/7/9: [Brandi] Photonic Systems
  - Guest Lecture Arka Majumdar (UW ECE), Wednesday 12/7

My favorite qubit presentations:

- trapped atoms/ions
- superconducting qubits
- spin qubits
  - classic solid state/epitaxial qubits
  - designer defects (diamond, silicon, etc.)
  - molecular qubits
- optical (single photon) qubits and cavity QED
- topological quantum materials

Assessment – weekly reading quizzes (30%), homework graded for completion (30%), my favorite qubit project (working in pairs, 30%), participation and discussion (10%)
Monday, Wednesday = lecture (week 1-7) or student presentations (week 8-10)
Friday = reading discussion, working problems, and/or guest lecture

Canvas Site – https://canvas.uw.edu/courses/1604983

Friday Discussion
Expectations: In CHEM/MSE 561 we seek to build community, share knowledge, and create a foundation of support. We will work together to support and learn from our peers. We commit to:

- Be prepared and engaged - come to class having reviewed and reflected on the reading. Actively contribute to the discussions and engage with guest speakers.
- Ask for help - from each other and the larger UW QISE community - we do not expect to be the source of all information. We will proactively help students find the resources they need.
- Accessibility - accommodate students' needs and ensure format of class is given in a way that caters to all learning types, recognizing that we all come from different backgrounds.
- Be patient - Grappling with the topics in this course can be challenging. Remember to give one another the space and time to think and reflect. Silence is OK.

Quantum smart workforce: Education at early stage including K-12; Fundamental academic research

NISQ: In the noisy intermediate-scale quantum era the leading quantum processors contain about 50 to a few hundred qubits, but are not advanced enough to reach fault-tolerance nor large enough to profit sustainably from quantum supremacy. The term was coined by John Preskill in 2018. It is used to describe the current state of the art in the fabrication of quantum processors.

The term 'noisy' refers to the fact that quantum processors are very sensitive to the environment and may lose their quantum state due to quantum decoherence. In the NISQ era, the quantum processors are not sophisticated enough to continuously implement quantum error correction. The term 'intermediate-scale' refers to the quantum volume related to the not-so-large number of qubits and moderate gate fidelity.

Quantum essential supporting technologies: Quantum states are fragile → requires new classes of high-performance classical components

Sneak Peek at Qubit Platforms

Superconducting Qubits
To maintain coherence, it is essential to keep electron-electron interactions, and also interactions between electrons and other degrees of freedom (such as phonons in the solid), under control. Superconductors have the advantage in this regard because the electrons condense into Cooper pairs that form a single superfluid. This superfluid is able to move through the metal lattice without any resistance (i.e., without interactions) because it takes a certain amount of energy, known as the energy gap, to break up the Cooper pairs.

The behavior of the electron superfluid is completely determined by a single quantum wavefunction. The amplitude of this wavefunction determines the number of Cooper pairs, while the value of the phase is related to the supercurrent and any magnetic field that is present. The amplitude and phase of the wavefunction are conjugate variables – that is, they are related by an uncertainty principle that means we cannot measure both of their values with arbitrary precision at the same time. The two primary types of superconducting qubit, the charge qubit and the flux qubit, are directly related to these two variables: charge qubits are associated with the amplitude, while flux qubits are related to the phase.

A key component in most superconducting qubits is a device called a Josephson junction. This consists of a thin layer of aluminum oxide, which is an insulator, sandwiched between two superconducting layers of aluminum, which becomes a superconductor when cooled below 1.2 K. The insulating layer is so thin (a few nanometers) that Cooper pairs can tunnel through it and couple the superconducting wavefunctions on either side of the barrier. Most of the circuits for superconducting qubits built so far consist of Josephson junctions and other components like capacitors connected by superconducting leads made of aluminum.

Trapped Ion Qubits

For trapped ions, internal electronic states of the ion are used for the qubit states |0> and |1>. Trapped-ion qubits can generally be considered as one of four types: hyperfine qubits, where the qubit states are hyperfine states of the ion separated by an energy splitting of order gigahertz; Zeeman qubits, where the qubit states are magnetic sublevels split by an applied field and typically have tens of megahertz frequencies; fine structure qubits, where the qubit states reside in the fine structure levels and are separated by typically tens of terahertz; and optical qubits, where the qubit states are separated by an optical transition (typically hundreds of terahertz).

Initialization and readout in trapped ions are both performed by laser manipulation of the ion internal and motional states. Initialization is performed via optical pumping into the |1> state, often accompanied by cooling of the ions' quantized motion to the trap harmonic oscillator ground state. State readout is likewise very simple: a resonant laser couples the |1> state to a cycling transition which scatters many photons that can be collected by a detector, while no similar transition exists for the |0> state which therefore remains dark. High fidelity state preparation and readout have both been performed in less than 1 ms.

Advantages: Coherence times can be exceptionally long for all four types of qubits. Hyperfine qubit coherence times as high as 50 s have been achieved without using spin-echo or other dynamical decoupling techniques and, such coherence times were extended up to 600 s with the aid of dynamical decoupling. These coherence times are effectively T2 times, limited by technical sources of dephasing rather than by the fundamental state lifetime.

Trapped ions also benefit from the fact that all ions of a given species and isotope are fundamentally identical.

While any ion contains additional internal states, the number of additional levels that must be accounted for in performing quantum operations is small when compared with the continuum of additional states that exist in solid-state qubits.

An ion can be trapped for many hours, or in some cases up to months for heavier ion species in deep traps, without being lost.

Disadvantages: While it is in principle easy to trap larger and larger numbers of ions in linear chains or two dimensional arrays, in practice the scaling to larger numbers of trapped ions has been slow. Arrays of up to thousands of superconducting qubits—such as the DWave 2000Q machine—have been fabricated with elementary control over each qubit, although these large arrays have limited connectivity, typically very short coherence times, and have not been used to demonstrate entanglement even between two qubits. While clouds of many thousands of ions can easily be trapped in deep macroscopic RF traps, such large clouds typically afford little meaningful control over individual ions and lack ion-specific readout. The largest systems of trapped ions with meaningful control and readout include 300ion crystals in Penning traps and linear chains of ∼100 ions in RF traps; neither of these systems has yet demonstrated entanglement between arbitrary ions in the system. The difficulties of implementing the necessary optical and electronic control have slowed progress towards larger numbers of trapped ions as compared with other technologies

Spin Qubits

Spin qubits have been most widely developed for solid-state semiconductor quantum dots. They are formed when electrons or holes are confined in a static potential well in a semiconductor, giving them a quantized energy spectrum. The simplest spin qubit is a single electron spin located in a quantum dot, but many additional varieties have been developed, some containing multiple spins in multiple quantum dots, each of which has different benefits and drawbacks. Although these spins act as simple quantum systems in many ways, they also experience complex effects due to their semiconductor environment. They can be controlled by both magnetic and electric fields depending on their configuration and are therefore dephased by magnetic and electric field noise, with different types of spin qubits having different control mechanisms and noise susceptibilities. Initial experiments were primarily performed in gallium arsenide–based materials, but silicon qubits have developed substantially and research on qubits in silicon metal-oxide-semiconductor, silicon/silicon germanium heterostructures, and donors in silicon is also being pursued.

The solid-state quantum dot is described by a central spin of a single charge (electron or hole) coupled to $N \approx 10^3 - 10^5$ nuclear spins via hyperfine interaction. The lack of translational motion combined with the mismatch in electron and nuclear spin energies suppresses relaxation, providing long spin lifetimes required for spin qubits. However, thorough understanding of spin relaxation is complicated by the multitude and complexity of the residual environment couplings, which include electron–phonon interactions, quadrupolar coupling of nuclear spins to strain, nuclear spin diffusion, and electron cotunneling arising from proximity of the Fermi reservoir. Moreover, impurity charge traps adjacent to QDs degrade spin qubit lifetimes.

NVs: In 1997 Jörg Wrachtrup and colleagues at the University of Technology Chemnitz in Germany showed that a single NV defect could be manipulated and provide an optical output at room temperature (Science 276 2012). This discovery sparked the field of diamond quantum technology (see Nature 505 472 for a more detailed history). The process is called optically detected magnetic resonance (ODMR) and, with the NV defect, it is observed when measuring a change in fluorescence after shining green light on a single NV defect, or on an ensemble of them, while scanning an applied microwave field. When the field hits resonance with the spin quantum numbers (ms) causing a transition from ms = 0 to ms = ±1, a decrease in fluorescence is observed. So, by measuring the intensity of the fluorescence, you can read out the spin state of the defect.

Photonic Qubits

In optical systems for quantum information processing, the unit of light in a given mode—or photon—is used to represent a qubit. Superpositions of quantum states can be easily represented, encrypted, transmitted, and detected using photons. Quantum states of light are durable. Photons seldom interact with each other, which means we can easily avoid uncontrolled interactions that destroy their quantum state. However, to build a universal quantum computer, we need multi-qubit gates, which means that photons must be made to communicate with each other somehow! We can make photons affect each other by using a material as a mediator. These materials need to be perfect single photon sources.

Xanadu System: Infrared laser pulses fired into the chip are coupled together with microscopic resonators to generate so-called “squeezed states” consisting of superpositions of multiple photons. The light next flows to a series of beam splitters and phase shifters that perform the desired computation. The photons then flow out the chip to superconducting detectors that count the photon numbers to extract the answer to the quantum computation.

Topological Qubits

Topology is a branch of mathematics describing structures that experience physical changes such as being bent, twisted, compacted, or stretched, yet still maintain the properties of the original form. When applied to quantum computing, topological properties create a level of protection that helps a qubit retain information despite what’s happening in the environment. The topological qubit achieves this extra protection in two different ways:

Electron fractionalization. By splitting the electron, quantum information is stored in both halves, behaving similarly to data redundancy. If one half of the electron runs into interference, there is still enough information stored in the other half to allow the computation to continue.

Ground state degeneracy. Topological qubits are engineered to have two ground states—known as ground state degeneracy—making them much more resistant to environmental noise. Normally, achieving this protection isn’t feasible because there’s no way to discriminate between the two ground states. However, topological systems can use braiding or measurement to distinguish the difference, allowing them to achieve this additional protection.

Building on two decades of scientific research and recent investments in simulation and fabrication, the Azure Quantum team has engineered devices that allow them to induce a topological phase of matter bookended by a pair of Majorana zero modes. These quantum excitations don’t normally exist in nature and must be coaxed into appearing under incredibly precise conditions. The ability to create and sustain a quantum phase with Majorana zero modes and a measurable topological gap removes the biggest obstacle to producing a unique type of qubit, which Microsoft’s quantum machine will use to store and compute information, called a topological qubit. It’s the foundation for Microsoft’s approach to building a quantum computer that is expected to be more stable than machines built with other types of known qubits, and therefore scale like no other.