Technical Roadmap for Fault-Tolerant Quantum Computing

October 2016

Amir Fruchtman and Iris Choi, University of Oxford

NQIT Networked Quantum Information Technologies

UK NATIONAL QUANTUM TECHNOLOGIES PROGRAMME

Networked Quantum Information Technologies (NQIT) is funded by the Engineering and Physical Sciences Research Council and is a part of the UK National Quantum Technologies Programme

Authors: Amir Fruchtman and Iris Choi, University of Oxford

To contact the authors, please email:
iris.choi@physics.ox.ac.uk

Design & layout: Hannah Rowlands, based on a design by Hunts

Acknowledgements: The authors would like to thank the following individuals for their helpful technical discussions and contributions to this document:

- Mete Atature
- Simon Benjamin
- Niel de Beaudrap
- Matthew Hoban
- Peter Leek
- David Lucas
- Gavin Morley
- Joshua Nunn
- Joe O'Gorman
- Andrea Rocchetto
- Jason Smith
- Nathan Walk

Table of Contents

Executive Summary

As the quantum computing field is gaining momentum, a small quantum computer with 10 - 200 qubits is on the horizon. Industrialists expressed a demand for a technical roadmap which explains the complex concepts of fault-tolerant quantum computing for a broad audience, and to identify the potential applications for a small quantum computer. Applications from quantum chemistry, quantum assisted computing, secret sharing and machine learning are described in this technical roadmap. Where possible, the authors have indicated the number of qubits needed for small quantum computer applications.

It is our intention to provide an impartial and accurate presentation of the fault-tolerant quantum computing technology, its developments and the potential applications for a small quantum computer. We hope that this technical report will be helpful to those who want to understand, engage, develop, manufacture or invest in this technology.

1. Introduction & Scope

With the promise of performing previously impossible computing tasks, there has been a substantial increase in momentum for quantum computing research, leading to a race to realise the world's first universal quantum computing machine. This progress has opened up many commercial opportunities, which created a substantial amount of interest from industry. Through the interactions with the Networked Quantum Information Technologies (NQIT) and the National UK Quantum Technologies Programme industrial networks, many industrialists expressed that the subject of quantum computing is complex, therefore, there is need for a detailed roadmap in order to clarify its technical development and potential applications. Such a roadmap will help them to understand the status of this technology and make relevant business decisions.

This technical roadmap is a direct response to these industrial requests. Our report focuses specifically on the subject of fault-tolerant quantum computing, therefore, does not cover other applications of quantum technologies such as quantum communications, cryptography, enhanced sensors, random-number generators, and so on.

This report aims to show the technical steps needed to build a fully functional quantum computer. We give an overview of the subject, and review leading technologies to realise such a computer. We include an estimate of the resources needed for real world problems, which address the common concerns. We also discuss the possible applications that would become available during the process towards building a fully universal quantum computer, i.e. with using only a “small” quantum computer. These applications apply to fields such as physics and chemistry simulations, encryption, and optimisation.

2. Two Kinds of Quantum Computers

There are two broad approaches to building a quantum computer: the **circuit model approach**[a], which can execute a sequence of operations, much like a CPU, and the **adiabatic approach**, which is oriented towards solving specific types of optimisation problems. As we explain below, the two kinds are theoretically equivalent, meaning there is a mapping between one and another, but in practice it is not easy to use this mapping.

Both kinds are made of quantum bits - **qubits**. A qubit has two distinct features that differentiate it from a regular bit:

1) **Superposition** – In contrast to a regular bit, which can be either 0 or 1, a qubit can be in be in a *superposition* of both 1 and 0, at the same time. It might be more in 0 or more in 1, or have any ratio between them. A qubit is best thought of as an arrow that can point to any direction in 3D space: when it points up, the qubit is said to be in the 0 state, down is the 1 state, and any other direction is a combination of both. There are two quantities that define the state of the qubit then: the angle of rotation in the up-down direction, and the rotation in the left-right direction. The up-down angle defines the probability of finding the qubit in the 0 or 1 direction, and the left-right direction is a purely quantum property which is called the *phase* of the qubit. Figure 1 shows an illustration of the state of a single qubit.

2) **Entanglement** – Two (or more) different qubits can exhibit correlations, which means that the state of each qubit is not independent from the rest. A typical example is the *Bell state* of two qubits; this is a state where there is an equal probability for finding both qubits in the 0 state, and finding both in the 1 state, but there is no chance of finding one of them in the 1 state and one in the 0. When two qubits are correlated like they are said to be *entangled*.

These correlations are thought to be the key to the superiority of quantum computers over their classical counterpart. In classical computing, a state of n bits can be described using n numbers (zero or ones), while a state of n qubits can only be described using 2^n-1 complex numbers, i.e. exponentially more information. This means that an exponential number of classical bits would be needed to store the state of a quantum computer, even approximately.

Qubit A qubit, or quantum bit, is a unit of quantum information, similar to a 'bit' in classical computing. However, unlike a bit, which can either be 0 or 1, a qubit can be 0 and 1 at the same time - a quantum superposition of both states. When multiple qubits are combined, they can store vastly complex data.

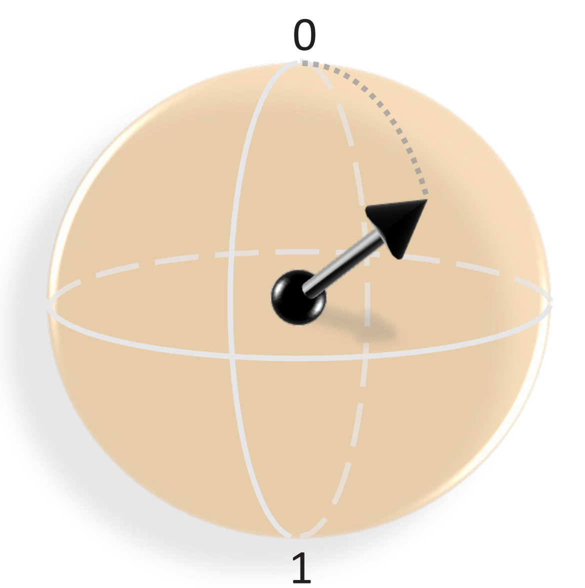

Figure 1: An illustration of a state of a single qubit. This 3D diagram is also known as the Bloch Sphere. The qubit's state is given by a direction in 3D space, or by a point on the sphere

[a] There are other models of quantum computing, such as measurement based quantum computing. However, as these models are very closely related to the circuit model, we have made this broad distinction for brevity's sake.

2.1 Errors, Decoherence, and the Threshold Theorem

A physical qubit does not hold its state indefinitely. It undergoes random bit-flips and loses its phase over time. This is called **decoherence**. To overcome this, physicists have come up with clever tricks for **error correction**, where the state of a logical qubit is encoded within several physical qubits. It is then possible to detect and correct errors that occur to the physical qubits, without changing the state of the logical one. This way the logical qubit is protected against errors, and holds information for longer than any underlying individual qubit can. The smallest number of physical qubits that can encode and correct arbitrary errors on a logical qubit is five[1,2].

The protocol for error-correction requires the physical qubits to be constantly monitored and manipulated. The **threshold theorem** [3] states that error-correction can only be achieved if the required manipulations can be performed with a very low error, below a certain threshold. In other words, if qubits can be manipulated with high precision, above a certain threshold, the errors can be corrected. If not - more errors would be introduced than fixed. Below this threshold, the performance of the logical qubit does not benefit from the error correction scheme, and therefore cannot be fault tolerant. Above the threshold, the more physical qubits used to encode a logical one, the greater the suppression of errors that can be achieved.

Surface Code The most prominent method for error-correction in quantum computing is called the surface code.

The threshold depends on the method used to encode and error-correct the logical qubit. The most prominent method for error-correction is called the **surface code** [4], for which the threshold sits at **99%** [5, 6]. The minimum number of qubits that can implement the surface code is nine [7]. Only very recently, as discussed below, advances in science and engineering enable us to create qubit operations that are above the threshold, meaning errors can be corrected. This signifies a major step towards building a large quantum computer.

The circuit model can be thought of as a direct analogue to a conventional computer, where instead of bits to represent data, we have qubits (for an explanation, see the beginning of this section). A circuit model quantum computer with a universal set of gates can perform any quantum algorithm. The most famous ones include Shor's algorithm for factorising a number into its prime factors, which is exponentially faster than the best classical algorithm [10], and Grover's search algorithm for searching an unstructured database, which gives a quadratic speedup over any possible classical algorithm [11].

Currently, the circuit-based model approach is the only known way to achieve fault-tolerant quantum computing.

Currently, the circuit-based model approach is the only known way to achieve fault-tolerant quantum computing. A circuit-based quantum computer is a machine that has qubits, has an ability to initialise them, perform *gates* on them, and *measure* them, i.e. check whether they are in the 0 or 1 state. In order to realise the full potential of quantum computation, the computer needs to be able to perform a "universal" set of quantum gates on the qubits. These gates are the basic building blocks of the algorithm. They are the *single-qubit gates*, which are also called *rotations* from imagining the qubit as pointing to a direction in space, and *two-qubit gates*, which creates entanglement between two qubits.

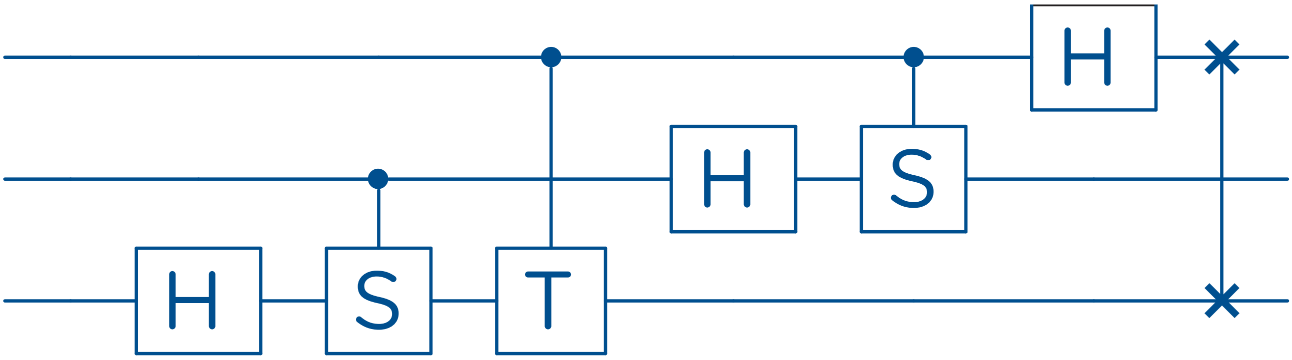

Figure 2: A subroutine of a quantum Fourier transform algorithm, to be read from left to right. Each row is a single qubit. Squares with H represent a single-qubit rotations called Hadamard. Vertical lines represent two-qubit gates.

2.2 Fault-tolerant Quantum Computation

When the qubits are encoded in an error-correcting protocol such as the surface-code, using many physical qubits to represent one logical qubit, errors that occur in the system can be supressed. However, the Eastin Knill theorem [12] states that some types of error-correcting schemes, including the aforementioned surface-code, cannot reliably achieve a universal set of gates. Gates that cannot be reliably implemented are referred to as non-robust gates. This means that additional resources are needed for a universal set of gates. Some error-correcting schemes get around this restriction [13], but it seems that these require higher-quality qubits to realise, i.e. produce a more stringent threshold for error correction [14].

For a fully functional quantum computation, using the "surface code" error-correction method, one needs access to special states, called **magic states**. These are prepared using independent qubits, and then injected into the system whenever a non-robust gate needs to be performed. These magic states are needed in vast quantities, and their preparation requires a significant portion of a device to operate as a dedicated **magic state factory**.

Magic States Magic states are the additional resources needed by the surface code quantum error correction protocol.

As an example, to perform Shor's algorithm to factorise a 1000 bit number, more than **10 billion** magic states are needed [8]. However, the performance of the qubits used for providing the magic states is not as stringent as is required for the qubits that encode the logical qubits.

Thus the overhead for a **fault-tolerant universal** quantum computer is substantial, and depending on the quality of qubits, it is estimated we need at least **a few hundreds of millions** of them, most constantly creating magic states, in order to realise a real-world quantum algorithm that outperforms conventional computers. On the bright side, the qubits needed for the magic state factory don't have to have very high quality as the ones encoding the logical qubits.

SECTION SUMMARY

A universal and fault-tolerant quantum computer can out-perform conventional computers, implementing any quantum algorithm. The main barrier today is that **hundreds of millions** [8] of qubits are needed to be accurately controlled and manipulated for solving real-world problems such as factorisation of large numbers. Currently, the highest number of qubits with the required characteristics realised is nine [9].

2.3 Modular Architecture

A large-scale universal quantum computer needs to link together millions or even billions of qubits, maintaining large-scale quantum phenomena. There are two different approaches to fabricating such a device.

The first approach is to have a single homogenous architecture, where each qubit is directly connected to its neighbours, meaning that two-qubit gates can be applied directly between the neighbours. Whilst this approach is appealing, it might not scale very well. Depending on the qubit technology, it may require many control lasers or microwaves pointing to the same physical region, or a very large vacuum chamber holding all of the qubits, or similarly a large cryogenic refrigerator.

An alternative approach is the network architecture [15], where the quantum computer is formed from many small *modules*, each consists of only a (relatively) small number of qubits, according to what is permissible by the technology used. These modules need to be linked together to perform inter-module gates between remote qubits. This requires either physically moving qubits from one module close to the qubits of another (for trapped ion qubits this is called shuttling [16]), or by using different qubits whose role is to mediate the quantum links between the qubits of the computer.

One way this can be done is by using photons to create entanglement between two 'communication qubits', one in each module. These communication qubits are qubits that are not used to encode the state of the quantum computer, but rather, to enable performing inter-module gates. Once two communication qubits are entangled, it is possible to perform any quantum gate between the two modules by performing local gates within each module [17].

For fault-tolerant quantum computation, the link between the modules, i.e. the entangling of remote communication qubits, can be done with lower quality than is needed for the gates that operate on the qubits that encode the state of the computer [18]. This is mainly thanks to a procedure called 'entanglement purification' [19]. This procedure turns several low-quality entanglement operations into a single high-quality one. Remote entanglement has been shown experimentally between two qubits of all major qubit technologies, including ion traps [20], superconducting qubits [21], and solid-state qubits [22], where for the latter case, the entanglement was between two electrons separated by 1.3 kilometres.

2.4 Adiabatic Quantum Computing

In contrast to the circuit-based quantum computer, an adiabatic quantum computation, takes a different approach to quantum computation. It performs a very specific kind of computation, but an extremely useful one: **optimisation**. An adiabatic quantum computer is used to find a set of variables that minimises some multi-variable function, and it is believed to be able to do so faster than a conventional computer [23], although the mathematics behind it is not completely clear at the moment. It has been proven [24] that any standard quantum computation can be done, in principle, with the adiabatic approach, and vice versa, but unfortunately the proof does not show a *simple* way to go from one system to the other.

It is possible to efficiently simulate the evolution of an adiabatic quantum computer using a circuit-based quantum computer, by slicing the adiabatic time-evolution into many small steps, where each step is then translated into a circuit [25]. The other way, i.e. simulating a circuit-based quantum computer on an adiabatic one, is also possible; though both ways generally result in an overhead. That said, several efficient protocols for solving specific problems other than optimisation on an adiabatic quantum computer are available, including the famous Grover's search problem [26].

The main concept behind the physics of adiabatic quantum computing is as follows: the optimisation problem is mapped onto a physical system of interacting qubits. Finding the best solution to the problem is equivalent to finding a qubit configuration that minimises the energy of the physical system, which is known as the ground state of the system. Physics shows that if any system is initiated in its ground state, or lowest energy state, and the parameters are *slowly* changed in the system, the system always remains in the ground state (of the new system). In practice, the user can initiate a very simple system (no interaction at all) where its ground-state is known, and slowly change the interaction to reach to the system of interest. Then, the user can probe the system to find out what value each qubit has, hence gaining a solution to the original optimisation problem.

Currently there are two interesting research challenges in the adiabatic paradigm:

a) **Error Correction** - in contrast to the circuit-based quantum computer, it is not known how to do error correction in the adiabatic machine, or how errors affect the computations.

b) **The gap problem** - The main concept behind the physics of adiabatic quantum computing is that when the system's parameters are *slowly* changed, the system remains in the ground state of the new system. The quantification of 'how slow' depends inversely on the *"energy gap between the ground and first excited state"* [23]. The value of this energy gap is not well-understood theoretically, and depends on the specific problem. It is known that this gap becomes smaller as the size of the system grows. This means that the computation time gets longer. Thus the adiabatic quantum computer is limited by the size of the problems it can solve.

An ideal adiabatic quantum computer should operate at an extremely low temperature, such that thermal fluctuations cannot bridge the energy gap mentioned above. A machine which operates at a finite temperature (i.e. above what is required for the adiabatic model) is known as a **quantum annealing machine** and is commercially available **today**. It provides a platform to understand these limitations and provide scientific breakthroughs.

Although an adiabatic quantum computer is limited by the gap problem, it might very well be able to outperform classical computers for the next decade.

3. Towards a Fault-tolerant Quantum Computer

The main effort from academia and some private entities such as Google, IBM, and Intel, is to build a gate-model quantum computer, and in 2016 the European Commission announced a €1 billion investment for research over the next 10 years.

The flagship goal of the NQIT Programme is to reach step four: a single, fault-tolerant logical qubit.

THE STEPS FOR BUILDING A MODULAR UNIVERSAL FAULT-TOLERANT QUANTUM COMPUTER, THAT IMPLEMENTS THE SURFACE CODE FOR ERROR-CORRECTION, ARE:

1. Build a module containing several high-quality qubits (isolated from environmental noise) ✓
2. Develop a precise method of performing single and two qubit logic gates, and measurements, with errors below the threshold (< 1% error). ✓
3. Build a second module, and connect the two modules in a quantum manner, with high quality.
4. Build several of these modules, combine them together to have a **single**, **fault-tolerant logical qubit**. This qubit can live for a very long time without any errors.
5. Build a **magic state factory** that enables a universal set of gates for the logical qubit.
6. Build **N** logical qubits + magic-state factory bundles, and couple them all together, to have an **N** qubit, universal, fault-tolerant, quantum computer.

Right now most of the practical research is focused on steps **1 and 2**, which have both been achieved in principle. However, many groups are trying to make higher-quality qubits, and better controlling methods. This is because the lower the error-rate, the number of physical qubits needed to realise one logical qubit decreases. Improvements in these aspects are extremely important and accelerate the realisation of a fully universal quantum computer.

The flagship goal of the NQIT Programme is to reach step four: a single, fault-tolerant logical qubit.

3.1 Popular Questions about Quantum Computers

HOW MANY LOGICAL QUBITS ARE NEEDED FOR OUTPERFORMING CONVENTIONAL COMPUTERS?

We can *simulate* a quantum computer on a conventional one, where the number of (logical) qubits we can simulate is limited by the memory of the computer, rather than its computing power. The largest super-computer today has ~10 Petabytes of memory, which lets it simulate ~ 48 qubits [27]. The largest simulation so far was made on a smaller machine which simulated 42 qubits [28]. So any quantum computer with more than roughly **50 qubits cannot be simulated on a computer**. This is called **quantum supremacy**.

Note that these qubits don't necessarily need to be error-corrected. It might be possible to use physical qubits that are good enough to undergo thousands of operations before an error occurs. However, such a machine is not scalable, and could only perform short algorithms. Nevertheless, it could not be simulated by a conventional computer, thus it could show quantum supremacy, and it might be useful for some applications.

HOW MANY PHYSICAL QUBITS ARE NEEDED FOR REAL-WORLD PROBLEMS?

As an example, to factorise a 1000-bit number, using Shor's algorithm, if the gate error is ~0.1%, we need **~166 million physical qubits**, but if the gate error is ~0.01%, then only **~5.5 million physical qubits** are needed. It will take the computer **6.6 weeks** to do if it runs at 1 MHz speed, or **11 hours** at 100 MHz [8] (this number does not depend on the gate error, just the speed). For comparison, the largest number to have been factorised on a conventional computer was a 768-bit number, and it took **more than two years** of many hundreds of CPUs to do so [29]. Factorising a 1000-bit number is roughly a 1000 times harder than that.

An Illustration of a super-computer. The largest one today can simulate 48 qubits.

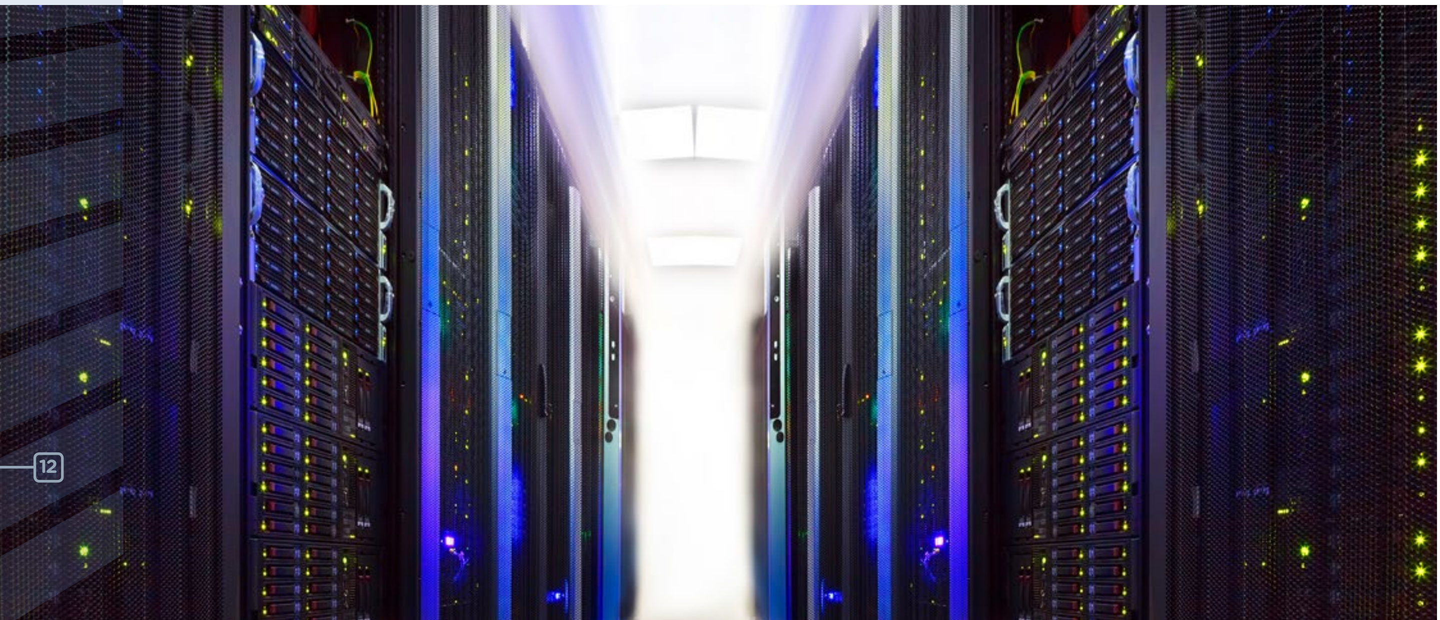

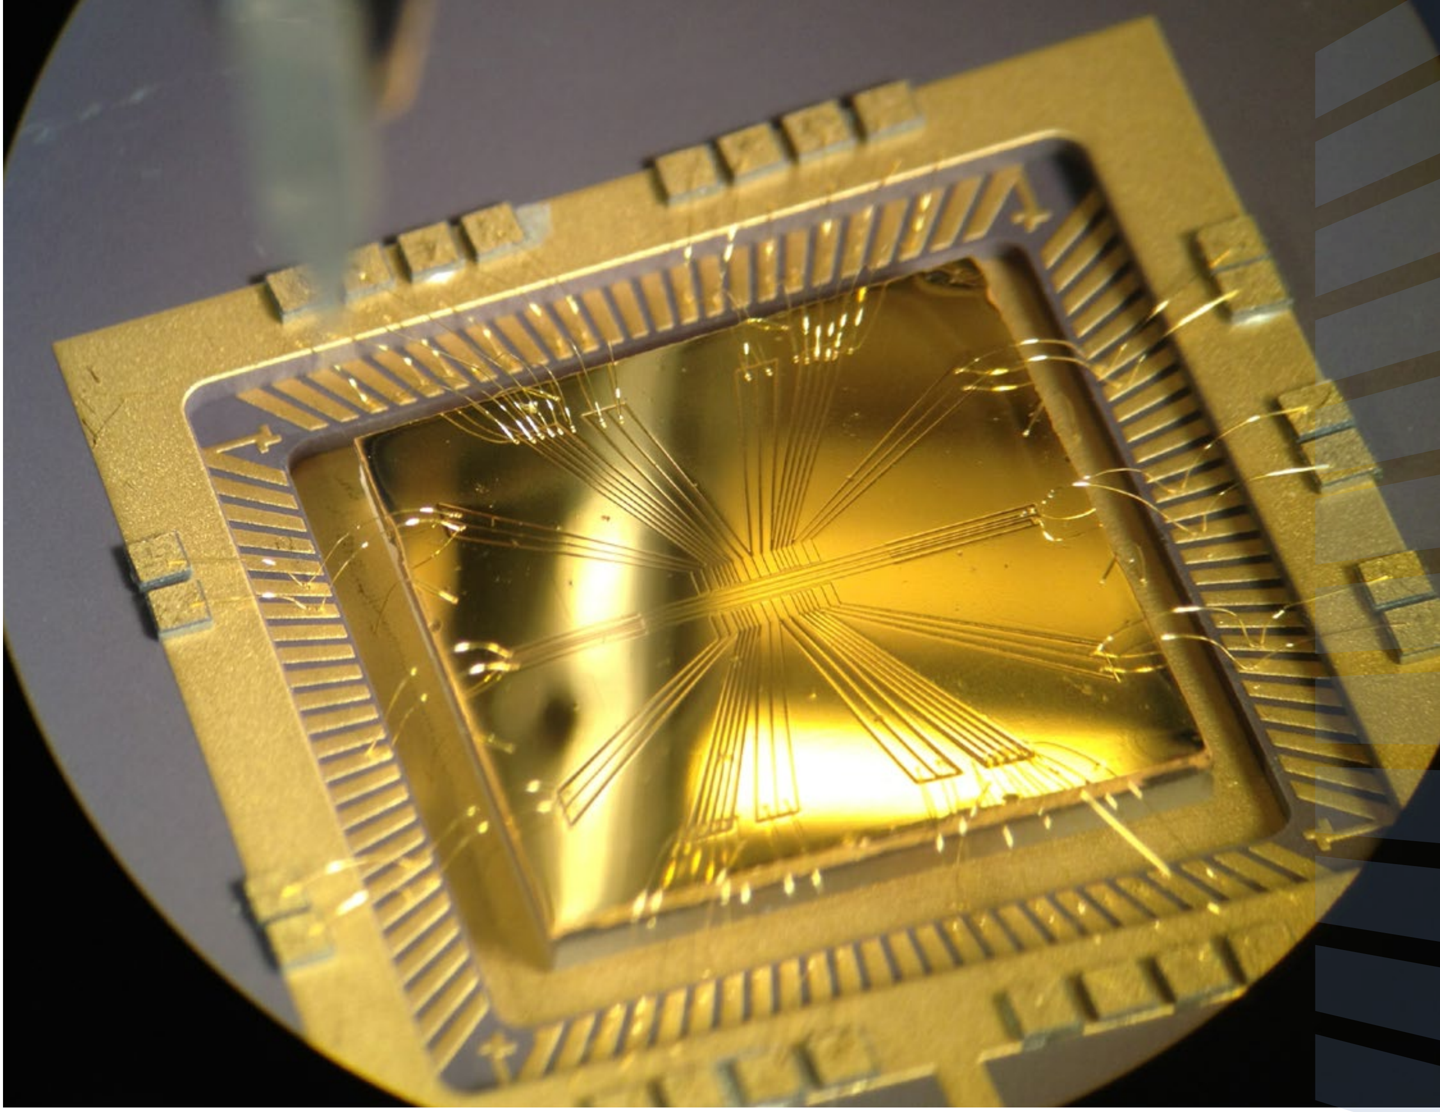

An ion trap from the Oxford University Ion Trap Group. The ions are held in vacuum just above the surface using electric fields / *D. Aude-Craik & D. Allcock*

3.2 Leading Technologies

There are several research approaches to realising the basic building block of a quantum computer - a qubit. The most popular approaches include: Ion trap qubits, superconducting qubits, nitrogen-vacancy centre qubits, photonic qubits, and silicon qubits.

Figure 3 shows the progress of the **ion-trap** and **superconducting qubits** technologies since 2002. The fidelity of a 2-qubit gate, a key performance parameter for quantum computing, improves from below the fault-tolerant threshold in 2002 (ion-traps), to beating the threshold in 2008 and to being an order of magnitude higher than the threshold in 2016. These were major scientific breakthroughs, as the ability to implement fault tolerant qubits enables scalability and real world applications. Both ion-trap qubits and superconducting qubits are now demonstrated to have fidelities above the threshold. Higher fidelities lead to a less stringent requirement in the total number of qubits for fault tolerant calculations. Therefore, for efficient performance it is essential to operate with fidelities as high above the threshold as possible.

Figure 3 also shows the various technologies and approaches to implementing ion-traps and superconducting qubits. These include using different ion species, different ions isotopes, multiple qubit demonstrations, near field and far-field microwave control methods.

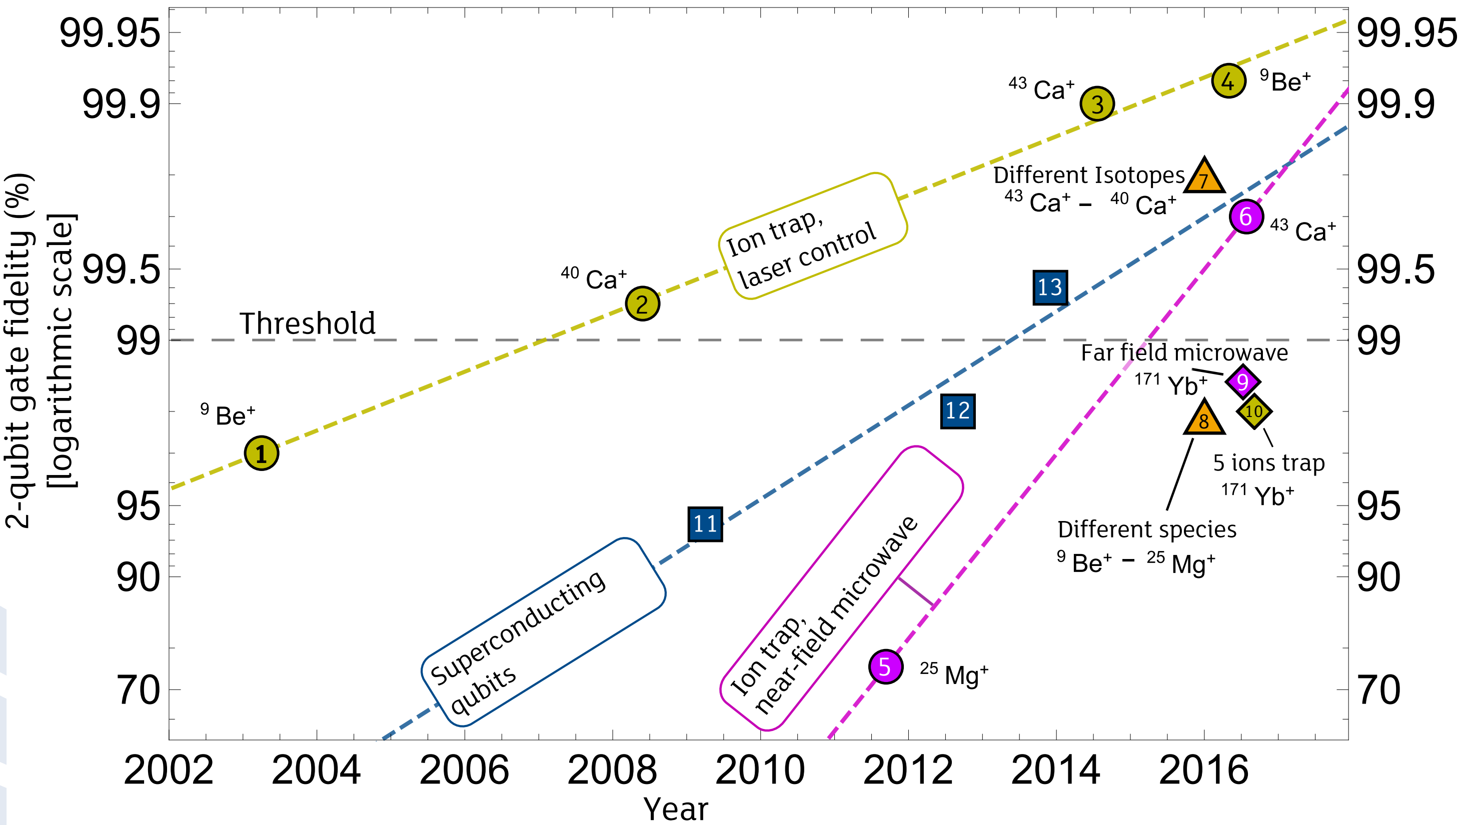

Figure 3: Progress of ion-qubit and superconducting qubit technology over time. y-axis: 2-qubit gate-fidelity (precision), in logarithmic scale. Above the threshold of 99% it is possible to reduce errors using the surface code error-correcting scheme.

References for each data point in Figure 3 are listed below and can be viewed in the References Section:

Technology	Data Point	Reference
Ion-trap, laser controlled	1 $^{9}Be^{+}$	Reference 30
	2 $^{40}Ca^{+}$	Reference 31
	3 $^{43}Ca^{+}$	Reference 32
	4 $^{9}Be^{+}$	Reference 33
Ion-trap, microwave control	5 $^{25}Mg^{+}$	Reference 34
	6 $^{43}Ca^{+}$	Reference 35
Ion-trap, different isotopes/species	7 $^{43}Ca^{+}$ - $^{40}Ca^{+}$	Reference 36
	8 $^{9}Be^{+}$ - $^{25}Mg^{+}$	Reference 37
Far field microwave	9 $^{171}Yb^{+}$	Reference 38
5-ions in a trap	10 $^{171}Yb^{+}$	Reference 39
Superconducting qubits	11	Reference 40
	12	Reference 41
	13	Reference 42

3.2.1 Various Qubit Technologies

Here we compare some of the different qubit technologies available. Table 1 below summarises the various parameters for the **Ion-traps qubits** and **Superconducting-qubits** technologies. Table 2 summarises parameters for the **Nitrogen-vacancy centres in diamond** and **Photonic qubits.**

3.2.1.1 Ion Traps

Trapped atomic ions are the most mature technology for implementing qubits and quantum logic gates. Simple algorithms and quantum simulations have been performed on up to around 20 trapped-ion qubits. They were the first system in which deterministic (i.e. "on demand") two-qubit logic gates were demonstrated, about 20 years ago, and the precision with which multi-qubit quantum logic gates can be performed has increased steadily since then, doubling every few years to reach the 99.9% level in the last couple of years. This represents a major milestone, as it is significantly beyond the minimum precision required to implement some quantum error correction codes - without which it is not possible to build a general-purpose quantum computer. Indeed, the performance of all elementary qubit operations (memory, state input, state readout, logic gates) for trapped-ion qubits is presently unrivalled by any other technology. The gate speed is, however, significantly slower than in solid state platforms, being typically in the 10-100 microsecond range, but importantly this is not a fundamental limitation and several research groups are investigating methods of implementing faster gates. Trapped ions involve the overheads of laser and high-vacuum technology, but the laser systems can be relatively simple, and the highest performances so far demonstrated have been in room-temperature setups.

3.2.1.2 Superconducting Qubits

Superconducting circuits started out as an unlikely candidate for quantum computing - the first demonstrations around 2000 had nanosecond coherence lifetimes. Understanding of coherence has however improved quite dramatically over the last decade, and superconducting circuits are now arguably at a comparable development level to trapped ions, with operation of 5-10 qubit circuits demonstrated in several research groups (albeit with individual gate fidelities still lower than with ions). The advantages of superconducting circuits as an architecture lie in their ability to be very flexibly designed and micro fabricated, and in the very strong non-linearities that can be engineered within them (making two-qubit gates and entanglement easy to engineer provided good coherence is maintained). This latter feature makes it possible to engineer fast gate times (in the range of 10 ns), and hence fast quantum calculations. It should in principle be possible to micro-fabricate large scale superconducting quantum circuits using techniques similar to those well established for the semiconductor industry. The dependence on superconductivity and the relatively low microwave frequency of the qubits does however mean that the circuits must be operated at very low (~10 mK) temperatures, requiring them to be housed in expensive dilution refrigerators. This places a limit on the size, convenience, or cost of a superconducting circuit based quantum computer, but it does not limit their potential to operate at the 100-1000 qubit level, and likely beyond.

Table 1: Comparison of ion traps and superconducting qubits technologies

	Ion Traps	Superconducting Qubits
Description	Ionised atoms trapped in empty space by oscillating electric fields, where qubit is encoded in the direction of spin of the atom's electrons with respect to an external magnetic field	Circuits of a superconducting material, kept in cryogenic temperature, where the qubit is encoded in the current that flows through the system
# of qubits realised with ability to perform single and two-qubit gates	5 [b] (Ref: 39)	9 (Ref: 9)
Lifetime of a qubit (T2 time)	~50 seconds	~50 micro-seconds
Best gate precision (2-qubit gate fidelity)	~ 99.9% [c] (Ref: 33 & 43)	~99.4% [d] (Ref: 42)
Time to perform a 2-qubit gate	~50 micro-seconds (Laser) ~3 millisecond (Microwave)	~50 nano-seconds
Physical size of a qubit	~1 micrometre	~100 micrometres
Pros	- Can operate at room temperature (but vacuum) - Long-lived qubit memory	- The chip can be manufactured using standard techniques - Fast gate time
Cons	- High power to perform gates	- Needs to be kept in cryogenic temperatures (~10mK) - Poor qubit T2 lifetime
Examples of research groups (in alphabetical order)	Blatt Group, University of Innsbruck, Austria (R. Blatt) Ion Storage Group, NIST, US (D. Wineland) Ion Trap Quantum Computing, University of Oxford, UK (D. Lucas, A. Steane) Trapped Ion Quantum Information, University of Maryland, US (C. Monroe)	DiCarlo Lab, Delft, NL (L. DiCarlo) IBM Quantum Computing group (M. Steffen, J. Chow) Martinis Group, UCSB/Google, US (J. Martinis) QuDev Lab, ETH Zurich, CH (A. Wallraff) Schoelkopf Lab & QuLab, Yale University, US (R. Schoelkopf, M. Devoret)

[b] Other groups have shown more qubits but without the ability to entangle arbitrary pair

[c] Best gate precision (1-qubit gate fidelity) for ion trap qubit: 99.9999% (Ref: 44)

[d] Best gate precision (1-qubit gate fidelity) for superconducting qubit: 99.92% (Ref: 42)

3.2.1.3 Nitrogen-Vacancy Centres

Nitrogen vacancy centres (NVC) in diamond benefit from electron and nuclear spins with coherence times of over 1 millisecond and 1 second respectively, at room temperature and atmospheric pressure. In this solid state system, single spins are polarized and read out by exciting with green light and detecting red fluorescence; microwave and RF pulses are used to provide coherent spin control. By sending the fluorescence through optical fibres, two NVC have been entangled despite being separated by over 1km, setting the benchmark for distributed entanglement in matter qubit system.

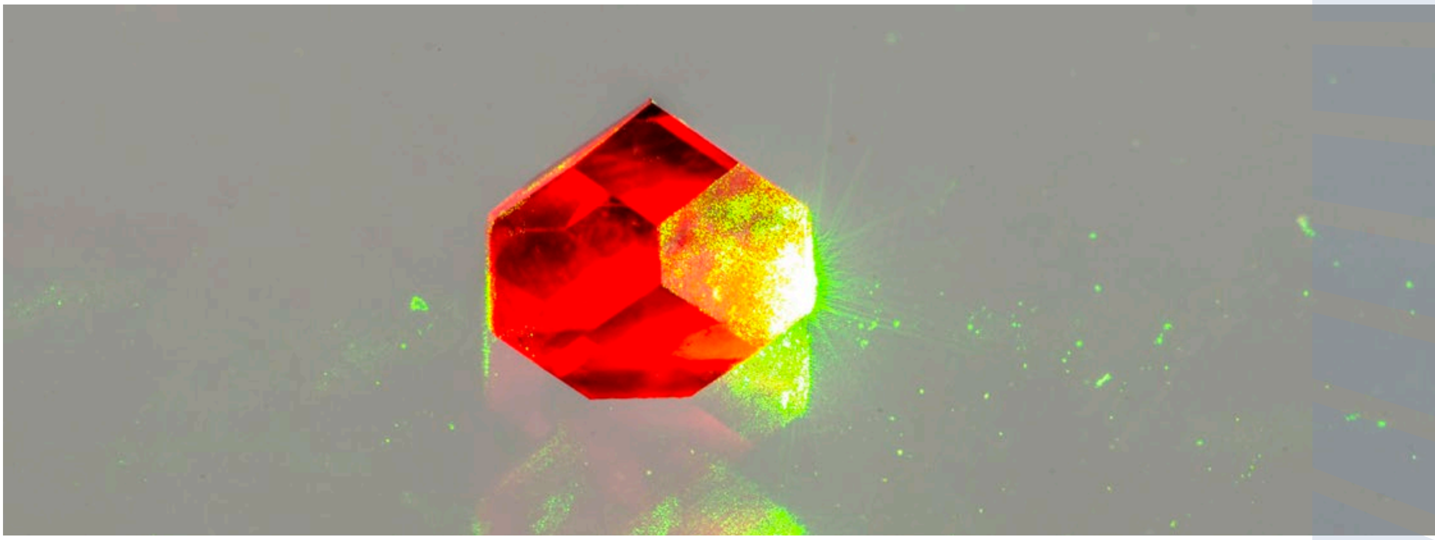

Diamond containing nitrogen vacancies fluorescing due to illumination with green light / *Jon Newland*

3.2.1.4 Photonics

Another way to implement quantum information processing is to use light, with qubits encoded in the occupation of on-chip waveguides by photons. Optical signals have carrier frequencies >100 THz, giving them an effective temperature ~10,000 Celsius - this means at room temperature there is no thermal noise and there is no need to cool down a photonic system to access its quantum features. Photons have no charge and do not require magnetic or electric shielding; no vacuum systems are needed and they can be sent down fibres or through free space to form either short-range or global-scale quantum links. Their high carrier frequency means they support extremely large data bandwidths, limited only by the speed of electronic demodulation (currently ~10 GHz in conventional telecommunication). However, as photons do not interact directly, two-qubit quantum gates are induced by interference followed by measurement, and are therefore probabilistic. This has limited the scale of quantum photonic processors to just a few operations. To achieve a scalable architecture, multiplexing strategies utilising optical switching, frequency modulation and light storage are being developed, to enable successful entangling operations to be actively concatenated. But these enabling technologies are yet to be incorporated into a large-scale demonstration.

Table 2: Comparison of Nitrogen-vacancy centres in diamonds and Photonics qubits technologies

	Nitrogen-vacancy centres in diamond	Photonics
Description	Qubits are encoded in the state of a nitrogen atom and isotopes of carbon atoms that are embedded in a diamond.	Single photons propagating in waveguides, where qubits are encoded in the occupation of a pair of waveguides.
# of qubits realised with ability to perform single and two-qubit gates	6 [e] (Ref: 45)	10 (Ref: 46)
Lifetime of a qubit (T2 time)	~10 seconds	~150 micro-seconds in a fibre
Best gate precision (2-qubit gate fidelity)	~88% (Ref: 47)	~98% (Ref: 48)
Time to perform a 2-qubit gate	~100 micro-seconds	~1 nano-second
Physical size of a qubit	~0.5 micrometres	~0.1-5 micrometres
Pros	No vacuum needed Can operate at room temperature [f]	Operates at room temperature Inherently suitable for the modular architecture
Cons	Current gate fidelity is under the threshold	Primitive elements are probabilistic. Multiplexing is needed to overcome this, resulting in very large overheads. Current gate fidelity is under the threshold.
Examples of research groups (in alphabetical order)	Diamond Research Group, University of Warwick, UK (M. Newton, G. Morley) Hanson Lab, TU Delft, Netherlands (R. Hanson) Lukin Group, Harvard University, US (M. Lukin) Photonic Nanomaterials Group, University of Oxford, UK (J. Smith) Quantum Information and Nanoscale Metrology Group, University of Cambridge, UK (M. Atature) Wrachtrup Group, University of Stuttgart, Germany (J. Wrachtrup)	QT Lab, University of Queensland, Australia (A. White) Quantum Information Lab, Sapienza University, Rome, Italy (F. Sciarrino) Quantum Photonics Group, University of Bristol, UK (J. O'Brien) Quantum Photonics Laboratory, MIT, US (D. Englund) Ultrafast quantum optics and optical metrology Group, University of Oxford, UK (I. Walmsley) Walther Group, Institute for Quantum Optics and Quantum Information, Vienna, Austria (P. Walther)

[e] This is one electron spin and 5 nuclear spins

[f] With reduced qubit lifetime (T2 time) of ~1 second

3.2.1.5 Spin Qubits in Silicon

Recently **silicon quantum computing** has attracted substantial research interests and commercial investments. Here, a qubit is encoded in the nuclear or electronic spin state of a donor atom in a silicon chip. These atoms are either naturally or intentionally placed into a silicon base, and controlled using microwave pulses. Like superconducting qubits, the silicon qubits must be cooled down to a fraction of a degree above absolute zero to operate. Qubit lifetimes of as long as 30 seconds have been reported [49]. With good single-qubit gate fidelity, and a recently demonstrated 2-qubit gate [50], the technology seems promising.

Silicon is an attractive base for a universal quantum computer because it is potentially compatible with the microelectronics of existing computers, using industry-standard silicon CMOS devices.

Examples of research groups working on a silicon quantum computer (in alphabetical order):

- Centre for Quantum Computation, University of New South Wales, Australia (M. Simmons, A. Morello, A. Dzurak)
- Eriksson Group, University of Wisconsin-Madison, US (M. Eriksson)
- Quantum Spin Dynamics, University College London, UK (J. Morton)
- Quantum Transport Group, TU Delft, Netherlands (L. Kouwenhoven, L. Vandersypen)

4. Applications for A Small Quantum Computer

SECTION SUMMARY

Even a small quantum computer consisting of less than 200 qubits can outperform standard computing and can be useful for simulations, optimisations, and encryption.

Before the arrival of a full quantum computer, we would only have a small quantum computer consisting of tens of error-corrected qubits, or several hundred non error-corrected ones.

As we mention in the Introduction and Section 3.1 Popular Questions about Quantum Computers, a classical computer today can simulate roughly 40 qubits. Still there are some scenarios where even a quantum computer consisting of less than this number of qubits would be more suitable than a classical one.

3D illustration of an ion trap quantum computer / *Chip Nyman*

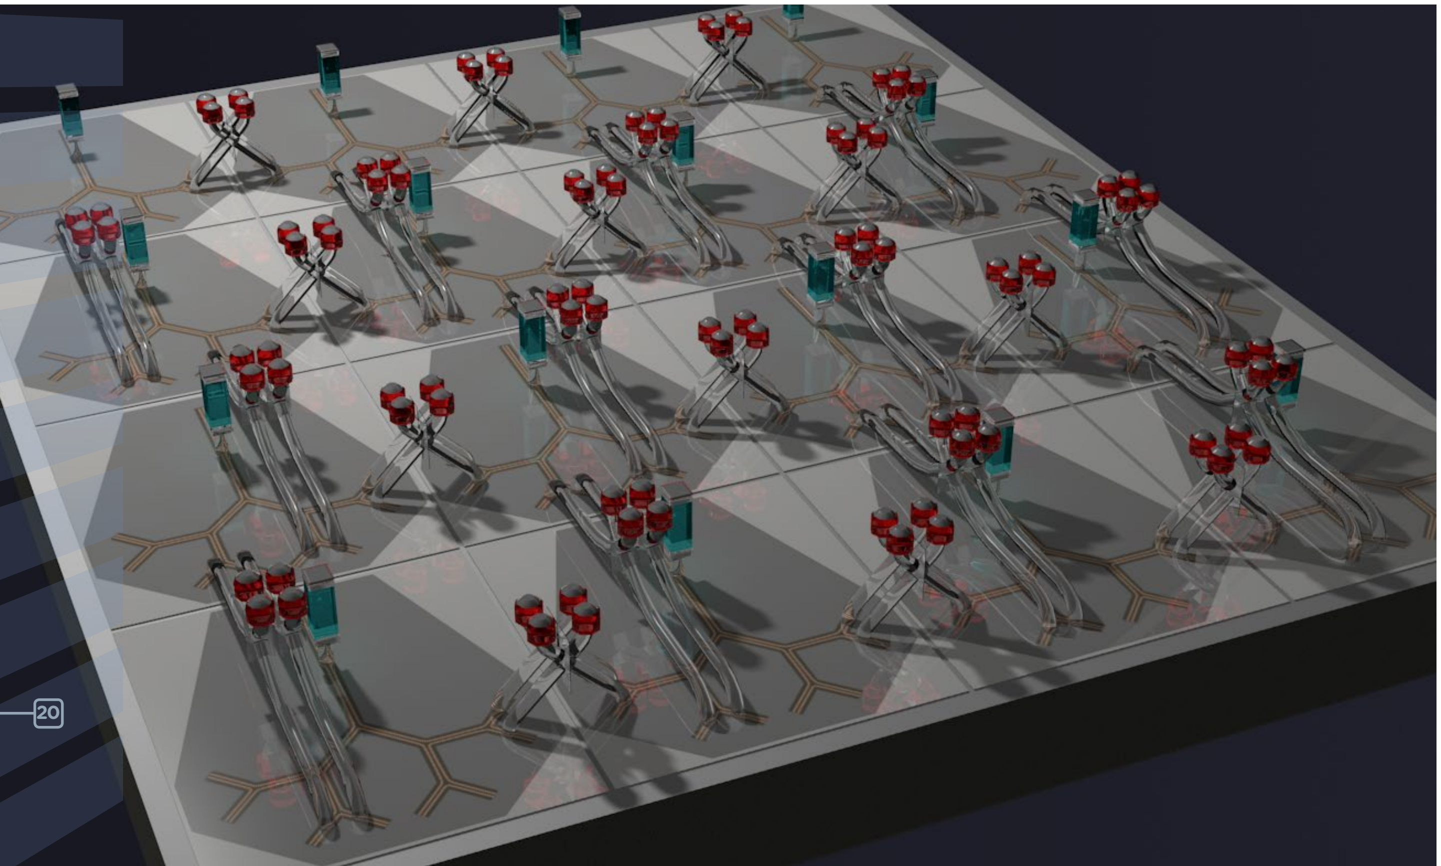

Even simulating a small quantum computer (tens of qubits) requires a large amount of computational resource.

First, the number of possible simulated qubits on a classical computer is memory-limited. The number 40 quoted above, taken from Reference 27, is for a double-precision accuracy for the real and imaginary parts of the complex numbers which represent the state of the quantum computer (16 bytes per complex number). This is for a system with a total of 32 Terabytes of memory shared across 1000 nodes. If a better accuracy is needed, fewer qubits could be simulated. For example, quad precision will result in a maximum of 39 qubits on the same machine. **A true quantum computer does not hold this limitation, and its accuracy is only limited by the precision of the instruments themselves.**

Second, the time it takes to simulate a quantum gate is limited by the network communications bandwidth between the nodes in the simulation, or more precisely on the time it takes to transfer all the memory stored in one compute node to another. Returning to the above example, simulating a single 2-qubit gate, can take more than 3 seconds with network bandwidth of 5.5 Gb/s and 32 Gigabytes per node. This is compared to a true quantum computer which, depending on the physical realisation, can perform the same gate in less than a millisecond (Ion-traps) or even less than a microsecond (Superconducting qubits) (see Table 1).

Thus there are reasons to believe that classical computers cannot efficiently simulate even a small quantum computer. This means realising a small quantum computer will open up an entirely new platform for applications that are currently impossible using conventional computers. Indeed, many researchers are looking for the “killer application” of a small quantum computer.

We note that there is a possibility to realise a small quantum computer using very good qubits, even without error correction. Such a machine will be limited in the number of operations it can perform, i.e. its operation time would be limited. Nevertheless, even a quantum computer with more than 50 qubits that can perform several hundreds of gates without errors can still outperform a conventional computer, and might even have useful applications.

Below we review some candidate applications of a small quantum computer.

4.1 Quantum Chemistry

Perhaps the most promising application is quantum chemistry. Even after all the recent progress in classical algorithms for chemistry, some molecules exhibit what is known as “strong correlation”, which in practice means that conventional methods fail to yield reliable solutions for them in reasonable time. This is because available approximations normally rely on having very weak correlations.

It is estimated that a small quantum computer of only 150 (error corrected) qubits would be able to calculate the exact ground-state of some of such molecules, without relying on the classical approximation. Moreover, the bigger the quantum computer gets, the larger and more complex molecules it can tackle. The quantum algorithms for these kind of problems have a statistical aspect. It consists of repeatedly measuring the state of the system to find some average measured values, then reset it to a new state which depends on these values. This means that the qubits of the quantum computer are only needed to be isolated from any noise for a very limited time, and not for the entire computation. Moreover, because of the averaging, some recent simulations suggest that these algorithms may give accurate results even when some noise or errors occur during the computation. This makes many researchers optimistic as to the possibility of performing meaningful chemistry calculations on a quantum computer in the foreseeable future.

Some recent proof-of-concept experiments have already been shown on a five qubit machine [51]. Reference 52 gives a review about quantum chemistry.

4.2 Quantum Co-Processor

Another way that a small quantum computer can boost computation is by combining it with a classical computer, in a way where some of the classical computation, or some steps in it, are off-loaded to the quantum computer. This is called a quantum co-processor.

Currently this idea can be used for the following:

- For simulating a larger quantum computer - we can use the resources of a small quantum computer with n qubits, to simulate (on a classical computer) a quantum computer with n+k qubits, with classical resources that only scale exponentially with k (rather than exponentially with n+k if no quantum computer was present)[53].
- For simulating physics - The Hubbard model, a widely used model of electrons in a solid, can be studied using a method called dynamical mean-field theory (DMFT). It maps the original Hubbard model of a spin lattice onto a simpler impurity problem. It was recently suggested that this method can be studied using a quantum co-processor. In DMFT, a very complex function is needed to be evaluated many times over the course of the computation, and this evaluation can be done efficiently on a quantum computer with as low as four qubits. Each evaluation is fed into the classical computer for the rest of the computation. Having more qubits simply makes the evaluation more accurate, thus for accelerating this simulation any size of small quantum computer would be useful [54, 55].

WHAT IS THE DIFFERENCE BETWEEN A LOGICAL QUBIT AND A PHYSICAL QUBIT?

A **physical** qubit is a real physical system, that can be in one of two (quantum) states, or in any superposition of both. A quantum computer executes algorithms which rely on abstracted **logical** qubits, which simply represents an idealised qubit. In practice, to achieve fault-tolerant quantum computation, a single **logical** qubit is represented by many **physical** qubits. In this case, error-correcting schemes can be employed to preserve the state of the **logical** qubit even when errors occur.

4.3 Secret Sharing

An interesting application of having several good error-corrected qubits that can hold quantum information for a long time, without computing, is secret sharing.

Secret sharing is a task where a dealer sends a secret S to n (possibly, dishonest [9]) players in a way such that the cooperation of a minimum of $k \leq n$ players is required to decode the secret; i.e., k-1 players should be unable to decode it even if they collaborated, but any k is fine. Protocols that accomplish this task are known as (k,n)-threshold schemes. The need for such a task appears naturally in a variety of situations, from children's games and online chats, to banking, industry, and military security: the secret message cannot be entrusted to any individual, but coordinated action is required for it to be decrypted in order to prevent wrongdoings.

Using N error-corrected logical qubits, it is possible to implement a secret-sharing protocol between N parties, each holds one qubit [56]. A secret is encoded in a way which is secure even against eavesdropping, and can be reconstructed only if k parties collaborate. The secret can be encoded for the lifetime of the qubit, which can be arbitrarily long for good error-corrected qubits.

4.4 Machine Learning

Machine learning has recently attracted a growing interest in the quantum community. Learning algorithms have indeed a series of features that are believed to be suited for a quantum implementation: they need to manipulate high dimensional vectors and use probabilistic processes that require sources of randomness. Quantum computers can accelerate these processes by carefully exploiting the probabilistic nature of the measurement process and the ability to perform certain linear algebra operations efficiently (e.g. finding eigenvalues with quantum phase estimation).

Several classical algorithms have been translated into the quantum realm, showing theoretical speedup over their classical counterparts. These include support vector machines, k-means clustering, and deep learning [57]. Most of these algorithms obtain their speedup using either Grover's search or the Harrow, Hassidim, Lloyd algorithm to solve linear systems of equations [58]. The latter approach is the only one that guarantees an exponential speedup but it requires very specific conditions on the structure of the dataset. Whether the attained speedup can remain on a practical problem remains an open question.

4.5 A Full Scale Quantum Computer

After an intermediate stage of only having a small quantum computer, we expect to have a full scale computer including error correction and a magic state factory. The applications for such a machine are vast, and they vary from doing linear algebra, simulating physics or chemistry faster than classical computers, to machine learning applications, fast database searches, financial analysis, and more [59]. For these applications hundreds of millions of good-quality physical qubits that can be controlled very accurately would be needed.

[9] If instead of qubits we have modes, which are in a sense a generalisation of a qubit, then the players can be dishonest as well.

5. Roadmap for Quantum Computing Applications

Figure 4 shows a roadmap for Quantum Computing applications, expressed in increments of available qubits. Potential applications for each qubit increment are shown (please refer to Section 4 for details). Table 3 summarises the minimum requirements for each application.

Table 3: applications of a small quantum computer

Application	Minimum Requirement
Quantum Chemistry	~ 150 error-corrected qubits
Quantum Co-Processor	Any number of qubits
Quantum Supremacy	More qubits than a classical computer can simulate. Currently 50
Secret Sharing	Very long-lived qubits
Machine Learning	A few hundred qubits

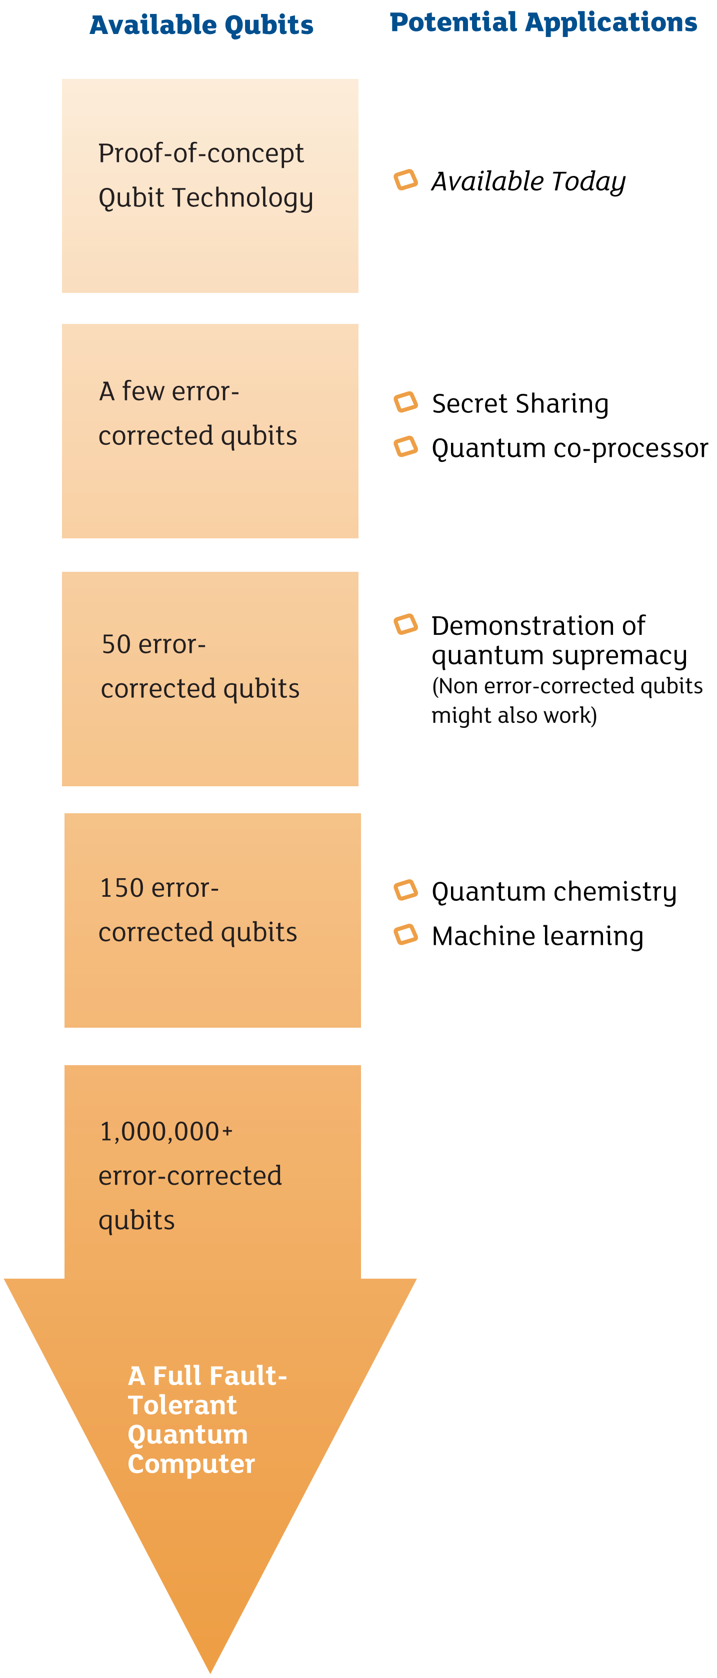

Figure 4: Applications of a small quantum computer, towards a full quantum computer

6. Summary

This technical roadmap has explained several complex concepts and performance indicators used in Quantum Computing for non-experts. Clear steps for technology developments are listed to explain to readers the necessary progress towards a full scale quantum computer. These steps serve as indicators to the length of time it will take for a full scale quantum computer to be realised.

This technical report also answered some popular questions on quantum computing, such as *'how many qubits are needed to out-perform conventional computers?'*. To deepen the understanding of this technology, a detailed discussion on the universal fault-tolerant quantum computing is given.

After the theoretical concepts are presented, several promising hardware approaches for basic building blocks of a quantum computer are listed. The performance of ion-trap qubits, superconducting, solid states and photonic qubits are discussed and compared. In addition, the increase in precision, as function of time, of the ion-trap and superconducting qubits technologies are plotted in Figure 3. As precision is one of the key performance parameters, Figure 3 captures the scientific and technological development over the last 14 years.

As the quantum computing field is gaining momentum, a small quantum computer with 10- 100 qubits is in the horizon. It will be useful for readers to identify the potential applications for a small quantum computer. Applications from quantum chemistry, quantum assisted computing, secret sharing and machine learning are described. Where possible, the authors have indicated how many qubits are needed for the small quantum computer applications.

Finally, from the historical developments and the state-of-the-art, a roadmap towards a full scale quantum computer is given in Section 5. This roadmap (Figure 4) is shown in increment of qubits. Potential applications for each increment are also presented.

Quantum computing is one of the most promising emerging technologies today. With the promise of performing previously impossible computing tasks, the arrival of quantum computing is expected to bring revolutionary changes to general computing. It is our intention to provide an impartial and accurate presentation of quantum computing technology, its developments and potential applications for a small quantum computer. We hope that this technical report will be helpful to those who want to understand, engage, develop, manufacture or invest in this technology.

References

1 R. Laflamme, C. Miquel, J. P. Paz, and W. H. Zurek, “Perfect Quantum Error Correcting Code,” *Phys. Rev. Lett.*, vol. 77, no. 1, pp. 198–201, 1996.

2 C. H. Bennett, D. P. DiVincenzo, J. a. Smolin, and W. K. Wootters, “Mixed-state entanglement and quantum error correction,” *Phys. Rev. A*, vol. 54, no. 5, p. 3824, 1996.

3 E. Knill, R. Laflamme, and W. H. Zurek, “Resilient Quantum Computation,” *Science* (80-.)., vol. 279, no. 5349, pp. 342–345, 1998.

4 A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Phys. Rev. A*, vol. 86, no. 3, 2012.

5 A. G. Fowler, A. C. Whiteside, and L. C. L. Hollenberg, “Towards practical classical processing for the surface code: Timing analysis,” *Phys. Rev. A*, vol. 86, no. 4, pp. 4–7, 2012.

6 D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, “Surface code quantum computing with error rates over 1%,” *Phys. Rev. A*, vol. 83, no. 2, pp. 2–5, 2011.

7 C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, “Surface code quantum computing by lattice surgery,” *New J. Phys.*, vol. 14, no. 12, p. 123011, Dec. 2012.

8 J. O’Gorman and E. T. Campbell, “Quantum computation with realistic magic state factories,” *arXiv*:1605.07197, May 2016.

9 J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, “State preservation by repetitive error detection in a superconducting quantum circuit,” *Nature*, vol. 519, no. 7541, pp. 66–69, 2015.

10 P. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” *SIAM J. Comput.*, vol. 26, no. 5, pp. 1484–1509, 1997.

11 L. K. Grover, “A fast quantum mechanical algorithm for database search,” *Proc. twenty-eighth Annu. ACM Symp. Theory Comput. - STOC ’96*, pp. 212–219, 1996.

12 B. Eastin and E. Knill, “Restrictions on transversal encoded quantum gate sets,” *Phys. Rev. Lett.*, vol. 102, no. 11, pp. 1–4, 2009.

13 A. Paetznick and B. W. Reichardt, “Universal fault-tolerant quantum computation with only transversal gates and error correction,” *Phys. Rev. Lett.*, vol. 111, no. 9, pp. 1–5, 2013.

14 B. J. Brown, N. H. Nickerson, and D. E. Browne, “Fault-tolerant error correction with the gauge color code,” *Nat. Commun.*, vol. 7, p. 12302, Jul. 2016.

15 C. R. Monroe, R. J. Schoelkopf, and M. D. Lukin, “Quantum Connections,” *Sci. Am.*, vol. 314, no. 5, pp. 50–57, Apr. 2016.

16 D. Kielpinski, C. Monroe, and D. J. Wineland, “Architecture for a large-scale ion-trap quantum computer,” *Nature*, vol. 417, no. 6890, pp. 709–711, 2002.

17 R. Nigmatullin, C. J. Ballance, N. de Beaudrap, and S. C. Benjamin, “Minimally complex ion traps as modules for quantum communication and computing,” *arXiv*:1605.00111, Apr. 2016.

18 N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, “Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links,” *Phys. Rev. X*, vol. 4, no. 4, pp. 1–17, 2014.

19 C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. a. Smolin, and W. K. Wootters, “Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels,” *Phys. Rev. Lett.*, vol. 76, no. 5, pp. 722–725, Jan. 1996.

20 C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L. M. Duan, and J. Kim, “Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects,” *Phys. Rev. A*, vol. 89, no. 2, pp. 1–16, 2014.

21 A. Narla, S. Shankar, M. Hatridge, Z. Leghtas, K. M. Sliwa, E. Zalys-Geller, S. O. Mundhada, W. Pfaff, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, “Robust Concurrent Remote Entanglement Between Two Superconducting Qubits,” *Phys. Rev. X*, vol. 6, no. 3, p. 31036, 2016.

22 B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, “Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres,” *Nature*, vol. 526, no. 7575, pp. 682–686, 2015.

23 E. Cohen and B. Tamir, “Quantum annealing – foundations and frontiers,” *Eur. Phys. J. Spec. Top.*, vol. 224, no. 1, pp. 89–110, 2015.

24 D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, “Adiabatic Quantum Computation Is Equivalent to Standard Quantum Computation,” *Siam Rev.*, vol. 50, no. 4, pp. 755–787, 2008.

25 E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, “Quantum Computation by Adiabatic Evolution,” *arXiv*:0001106, p. 24, Jan. 2000.

26 I. Hen, “Continuous-time quantum algorithms for unstructured problems,” *J. Phys. A Math. Theor.*, vol. 47, no. 4, p. 45305, 2014.

27 M. Smelyanskiy, N. P. D. Sawaya, and A. Aspuru-Guzik, “qHiPSTER: The Quantum High Performance Software Testing Environment,” *arXiv*:1601.07195, 2016.

28 S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, J. M. Martinis, and H. Neven, “Characterizing quantum supremacy in near-term devices,” *arXiv*:1608.00263, 2016.

29 T. Kleinjung, K. Aoki, J. Franke, A. K. Lenstra, E. Thomé, J. W. Bos, P. Gaudry, A. Kruppa, P. L. Montgomery, D. A. Osvik, H. Riele, A. Timofeev, and P. Zimmermann, “Factorization of a 768-Bit RSA Modulus,” in *Proceedings of CRYPTO, 2010*, pp. 333–350.

30 D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovi , C. Langer, T. Rosenband, and D. J. Wineland, “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate.,” *Nature*, vol. 422, no. 6930, pp. 412–5, 2003.

31 J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt, “Towards fault-tolerant quantum computing with trapped ions,” *Nat. Phys.*, vol. 4, no. 6, pp. 463–466, 2008.

32 C. Ballance, T. Harty, N. Linke, and D. Lucas, “High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions,” arXiv:1406.5473, 2014.

33 J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, “High-Fidelity Universal Gate Set for 9Be+ Ion Qubits,” *Phys. Rev. Lett.*, vol. 117, no. 6, p. 60505, Aug. 2016.

34 C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland, "Microwave quantum logic gates for trapped ions," *Nature*, vol. 476, pp.181-184 , 2011.

35 T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas, *Phys. Rev. Lett.* 117 140501 (Sep.2016)

36 C. J. Ballance, V. M. Schäfer, J. P. Home, D. J. Szwer, S. C. Webster, D. T. C. Allcock, N. M. Linke, T. P. Harty, D. P. L. Aude Craik, D. N. Stacey, A. M. Steane, and D. M. Lucas, “Hybrid quantum logic and a test of Bell’s inequality using two different atomic isotopes,” *Nature*, vol. 528, no. 7582, pp. 384–386, Dec. 2015.

37 T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland, “Multi-element logic gates for trapped-ion qubits,” *Nature*, vol. 528, no. 7582, pp. 380–383, 2015.

38 S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger, “Trapped-ion quantum logic with global radiation fields,” *arXiv*:1603.03384, Mar. 2016.

39 S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, “Demonstration of a small programmable quantum computer with atomic qubits.,” *Nature*, vol. 536, no. 7614, pp. 63–6, Aug. 2016.

40 L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, “Demonstration of two-qubit algorithms with a superconducting quantum processor.,” *Nature*, vol. 460, no. 7252, pp. 240–244, 2009.

41 J. M. Chow, J. M. Gambetta, A. D. Córcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen, M. B. Ketchen, and M. Steffen, “Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits,” *Phys. Rev. Lett.*, vol. 109, no. 6, pp. 1–5, 2012.

42 R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, a N. Korotkov, a N. Cleland, and J. M. Martinis, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature*, vol. 508, no. 7497, pp. 500–503, Apr. 2014.

43 C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, “High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits,” *Phys. Rev. Lett.*, vol. 117, no. 6, p. 60504, Aug. 2016.

44 T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, *Phys. Rev. Lett.* 113, 220501 (Nov. 2014)

45 A. Reiserer, N. Kalb, M. S. Blok, K. J. M. M. Van Bemmelen, T. H. Taminiau, R. Hanson, D. J. Twitchen, and M. Markham, “Robust Quantum-Network Memory Using Decoherence-Protected Subspaces of Nuclear Spins,” *Phys. Rev. X*, vol. 6, no. 2, p. 21040, 2016.

46 X. Wang, L.-K. Chen, W. Li, H. Huang, C. Liu, C. Chen, Y. Luo, Z.-E. Su, D. Wu, Z.-D. Li, H. Lu, Y. Hu, X. Jiang, C.-Z. Peng, L. Li, N.-L. Liu, Y.-A. Chen, C.-Y. Lu, and J.-W. Pan, “Experimental ten-photon entanglement,” *Arxiv*:1605.08547, May 2016.

47 G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-Herbrüggen, H. Abe, T. Ohshima, J. Isoya, J. F. Du, P. Neumann, and J. Wrachtrup, “Quantum error correction in a solid-state hybrid spin register.,” *Nature*, vol. 506, no. 7487, pp. 204–7, 2014.

48 B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O’Brien, A. Gilchrist, and A. G. White, “Simplifying quantum logic using higher-dimensional Hilbert spaces,” *Nat. Phys.*, vol. 5, no. 2, pp. 134–140, 2009.

49 J. T. Muhonen, J. P. Dehollain, A. Laucht, F. E. Hudson, R. Kalra, T. Sekiguchi, K. M. Itoh, D. N. Jamieson, J. C. McCallum, A. S. Dzurak, and A. Morello, “Storing quantum information for 30 seconds in a nanoelectronic device,” *Nat. Nanotechnol.*, vol. 9, no. 12, pp. 986–991, Oct. 2014.

50 M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, “A two-qubit logic gate in silicon,” *Nature*, vol. 526, no. 7573, pp. 410–414, Oct. 2015.

51 P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, A. Megrant, J. Y. Mutus, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, “Scalable Quantum Simulation of Molecular Energies,” *Phys. Rev. X,* vol. 31007, p. 1512.06860, Dec. 2016.

52 L. Mueck, “Quantum reform,” *Nat. Chem.*, vol. 7, no. 5, pp. 361–363, 2015.

53 S. Bravyi, G. Smith, and J. A. Smolin, “Trading Classical and Quantum Computational Resources,” *Phys. Rev. X*, vol. 6, no. 2, p. 21043, Jun. 2016.

54 J. M. Kreula, L. García, – Álvarez, L. Lamata, S. R. Clark, E. Solano, and D. Jaksch, “Few-qubit quantum-classical simulation of strongly correlated lattice fermions,” *EPJ Quantum Technol.*, pp. 1–19, 2016.

55 J. M. Kreula, S. R. Clark, and D. Jaksch, “Non-linear quantum-classical scheme to simulate non-equilibrium strongly correlated fermionic many-body dynamics,” *Sci. Rep.*, vol. 6, p. 32940, Sep. 2016.

56 R. Cleve, D. Gottesman, and H. -K. Lo, "How to Share a Quantum Secret," *Phys. Rev. Lett*. vol. 83, no. 3, pp. 648-651, Jul. 1999.

57 M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemp. Phys.*, vol. 56, no. 2, pp. 172–185, Apr. 2015.

58 A. W. Harrow, A. Hassidim, and S. Lloyd, “Quantum Algorithm for Linear Systems of Equations,” *Phys. Rev. Lett.*, vol. 103, no. 15, p. 150502, Oct. 2009.

59 A. Montanaro, “Quantum algorithms: an overview,” *npj Quantum Inf.*, vol. 2, p. 15023, 2016.

For more information about NQIT,
please visit our website:

http://www.nqit.ox.ac.uk

Or send us an email:

Or get in touch

email: contact@nqit.org

@NQIT_QTHub

NQIT Networked Quantum Information Technologies

UK NATIONAL QUANTUM TECHNOLOGIES PROGRAMME