STANDARDS Australia 100
CELEBRATING 100 YEARS 1922-2022

RESEARCH REPORT

At the Intersection Between Scalable Quantum Computing and Standardisation

by Nathan K Langford and Simon J Devitt

ABOUT STANDARDS AUSTRALIA

Standards Australia is an independent, non-government, not for profit organisation. We are the nation's peak non-government standards development organisation.

The work of Standards Australia and our staff, stakeholders, members and contributors enhances the nation's economic efficiency, international competitiveness and contributes to a safe and sustainable environment for all Australians.

Standards Australia's vision is to be a global leader in trusted solutions that improve life – today and tomorrow.

Authors

Nathan Langford

A/Prof. Nathan Langford is a lecturer in physics at UTS, a former ARC Future Fellow, and Director of the newly formed Advanced Prototype Packaging Facility (UTS/USYD). He leads the Millikelvin Quantum Science laboratory and the new Hardware Team in the UTS Centre for Quantum Software and Information. Nathan has over 20 years' experience in advanced experimental quantum science research, working at top labs in Australia, Europe and the US. His CirQuS (Circuit Quantum Science) research group focuses on performance characterisation and optimisation in circuit-QED based Digital Quantum Simulations, as well as Quantum Chaos, Quantum Control and quantum device benchmarking techniques including quantum tomography. He is an active member of the Sydney Quantum Academy, having served on its Technical Advisory Committee for the 4 years since it was established, and through his recent infrastructure development activities, Nathan is also engaging widely with local NSW quantum and advanced manufacturing industry. Since 2022, Nathan has been a member of the ISO/IEC JTC 1 Working Group 14 on Quantum Information Technology. He is playing an influential role in both WG 14's international standards development work, and the strategic development of Standards Australia's national quantum standardisation activities.

Simon Devitt

A/Prof. Simon Devitt is Research Director at the Centre for Quantum Software & Information at UTS. His research focuses on Quantum Error Correction and Fault-tolerance, the design of large-scale quantum computing and communications systems and the compilation and resource optimisation of quantum algorithms. Simon is the co-founder and managing director of h-bar quantum consultants and Co-founder of a new quantum start-up, Eigensystems Pty Ltd. Having spent over a decade overseas working in the quantum computing programs of the UK and Japan, Simon has worked with numerous corporations, startups and VC firms on their expansion into the quantum technology space and advised multiple government agencies on multimillion-dollar R&D initiatives.

Contents

Acknowledgements

Standards Australia would like to acknowledge the support of the Australian Department of Industry, Science and Resources in the development of this white paper.

The authors would like to acknowledge Jared Cole for reading and providing feedback on this report.

Executive Summary

From Frontier Science to New Industry: Quantum computing has emerged as, arguably, a defining technology of the 21st century, at the heart of the so-called second quantum revolution (Dowling & Milburn, 2002). Slowly developed as an outgrowth of the physics and computer science community in the late 1960s and 1970s, quantum computing and quantum technology, in general, have experienced several expansion phases, including:

- the quantum optics era that dominated the 1980s and 1990s;
- the race for new quantum bits and gates that took off in the late 1990s, following breakthroughs discoveries of the first algorithms (Shor's algorithm (Shor, 1994), Grover's algorithm (Grover, 1996), universal quantum simulations (Lloyd, 1996), and quantum error correction (Shor, 1995) and fault-tolerance (Shor, 1996)), and led to first small-scale gate-based quantum processors (Gershenfeld & Chuang, 1997);
- and the birth of a new deep-tech quantum industry which started to take off in the mid-to-late 2010s, several years after the first major big-tech players (like IBM, Microsoft, Google, and Intel) entered the quantum computing space seriously.

Quantum computing has now definitively crossed from solely academic research into the commercial world as a nascent new market, a paradigmatic example of deep-tech industry: On one hand, the new quantum industry aims to engineer large-scale quantum computers and new applications to tackle problems of commercial interest for economic and societal benefit. On the other, this is still built on very new, fundamental science, and there is still a lot that is unknown about what quantum computers will ultimately be good for.

Towards a Commercial Future: Investment in quantum hardware dominates over investment in quantum software, with 73% of investments since 2018 going towards quantum hardware (World Economic Forum, 2022). We now have numerous corporates and major research institutions building and demonstrating routine control over quantum processors at "early" intermediate scales that are pushing into "beyond-classical" performance regimes, beyond the few-qubit devices that were state-of-the-art a decade ago. Yet there remains still a vast chasm to be crossed in both scale and performance before cutting-edge quantum computers realise a size where they can be expected to solve problems of known utility that are currently out of reach of conventional computing. And it is likely that numerous disruptive innovations will still be required to reach this target.

Commercial quantum software activity is now also growing, with many major industrial quantum efforts having significant software and algorithms teams. Yet the majority of this effort is arguably still narrowly focussed on hardware-related tasks of hardware control and algorithms for the Noisy Intermediate-Scale Quantum (NISQ) era (Preskill, 2018). Current investment dynamics in NISQ algorithms are likely driven at least partly by a mismatch between natural short-term commercial incentives and necessary realism about the momentous scale of the task of building a quantum computer of genuine utility. Some consensus is now perhaps emerging that little-to-no genuine utility can be expected from small-to-medium-scale NISQ quantum processors (Brandhofer, et al., 2021), and there is a growing realisation that more dedicated attention needs to be focussed on the main long-term goal of scalable quantum computing. Despite the longer time horizon, this remains the scenario where the expected very significant benefits from quantum computing will emerge. Significant government investment (G. Brennen, 2012) will likely need to continue to help bridge the gap between the natural short-termism of risk tolerance in the private sector and the very significant, yet ultimately long-term pay-offs that quantum computers are likely to bring. However, a targeted, strategic focus on standardisation in contexts carefully chosen for appropriateness and broad benefit, could also play a key role in smoothing over some of the bumps in the quantum computing development path ahead.

What is this report about? We start this report by introducing the current context of quantum software and hardware development. We then argue the importance of concentrating more focus on the longer-term aim of scalable quantum computing and explore the challenges that lie on the path to achieving that goal. To do this, we discuss how the scale and type of challenges will vary across a range of potential end-game scenarios for quantum computing technologies, from smartphone-embedded quantum “accelerators” all the way up to a Large Hadron Collider scale Quantum Computing Collaboration. We discuss the algorithms and hence applications that will be relevant at each scale, and the end-user markets that these will unlock. We conclude that potentially useful applications exist at all scales, and that the most feasible scenarios at the scale of cloud quantum computing data centres or corporate quantum high-performance computing clusters should already provide benefits across many sectors.

To discuss the role that standardisation processes can play in the development of quantum computing, we look at five different spheres of stakeholder engagement for the quantum computing community, explore the needs and opportunities for standardisation within these spheres. We consider what we can learn from the existing or near-term standardisation landscape.

Key needs and opportunities for standardisation in scalable quantum computing

1. **Quantum computing is heterogeneous: Engaging within the community**

- The existence of widely varying hardware and software platforms creates significant challenges for standardisation processes, but also opportunities.
- In both hardware and software, standardisation could improve communication and compatibility between different technology platforms and layers.
- Focussed initiatives driven by governments, broad industry consortia or key players, are proving most agile at starting to develop standards.

2. **Quantum computing has extreme technology requirements: Engaging with suppliers**

- Quantum computing technologies often demand key aspects of performance from supporting technologies far surpassing what they were originally developed for.
- Developing Requirements Specifications for supporting equipment is one of the most important near-term opportunities for formal standards development.
- The growing quantum market offers new and lucrative economic opportunities for conventional technology manufacturers willing to engage with new requirements.

3. **Quantum computing will have interdisciplinary impact: Engaging with end-users**

- End-user applications for quantum computing impact across diverse sectors; developing them requires highly interdisciplinary collaboration with end users.
- New algorithms and use-cases could strongly influence both its reach across sectors, and the technology scales it could be usefully deployed at.
- Appropriate standardisation could aid communication between quantum and end-user domain experts, and help establish reliable resource benchmarking for both classical and quantum algorithms, that will be crucial for algorithm development.

4. **Quantum computing is rapidly developing deep tech: Educating a quantum workforce**

- Due to its rapid growth and reliance on cutting-edge science, the new quantum industry is currently facing global pipeline shortages in relevant expertise areas.
- Currently, the best opportunity for standardisation is to establish Terminology and Requirements standards to help train expert technicians in supporting areas.
- Innovation is still required, to identify if standardisation can play a further substantive role in addressing issues in the quantum workforce pipeline.

5. **Quantum is deep tech: Engaging with the wider community**

- As a deep-tech industry, rapid growth quantum computing risks out-pacing public licence and regulatory confidence, as other fields like AI have done.
- Developing robust, objective, widely adopted (and developed), comprehensive and accessible performance benchmarking standards will be difficult, but vital for due diligence and regulatory compliance for end-users, investors and government.
- Good terminology and informative standards could facilitate better communication with external stakeholders, to help build consumer/investor confidence.

We also consider what we can learn from the existing or near-term standardisation landscape, e.g., in relation to how to maximise agility and pace of standards development while maintaining quality control. We argue that, for maximum impact, initial standardisation efforts should focus either 1) on areas which provide commercial benefit without encroaching on commercial IP and competitiveness; or 2) on areas where there is significant public and community interest in defining rigorous performance standards, such as for comparative benchmarking and quality assurance processes.

We then conclude our report by discussing how the role and opportunities for standardisation might vary under each of the different quantum computing end-game scenarios, guided again by considering the different potential spheres of engagement that standards development processes have impact on. For example, standards are likely to be open but internally developed for an LHC-type Quantum Computing Collaboration, but installing quantum computing accelerator chips in smartphones would require full, open, international, formal standardisation. And while hardware would have to be fully miniaturised, monolithic and highly standardised for smartphone use, an LHC-scale quantum computer could tolerate much more complex, highly customised and hybrid-compatible designs.

1. Quantum computing: A second generation quantum technology

Technologies such as Quantum Computing, Quantum Sensors and Quantum Communications systems, that arise out of quantum information science, are sometimes referred to as second-generation quantum technologies. First-generation quantum technologies are often characterised in terms of three quintessential examples:

- The transistor (1956 Nobel prize in physics) (Schockley, et al., 1956)
- Nuclear Magnetic Resonance (1944 Physics Nobel, 1952 Physics Nobel) (Rabi, 1944) (Bloch & Purcell, 1952)
- The Maser-Laser (1964 Physics Nobel) (Townes, et al., 1964)

These three technologies, which are fundamentally enabled by quantum mechanical effects, have each ushered in their own technological revolution: the transister underpins the entire digital information revolution of the late 20th century; nuclear magnetic resonance introduced a completely new method for non-invasive medical imaging; and the laser has transformed many technologies, including driving a fibre optics telecommunications revolution which has helped the internet thrive.

The distinction between first and second-generation quantum technologies is how the quantum mechanical effects of matter and light are exploited. In first-generation technologies, these underlying effects do not manifest directly at the operational level, which can still generally be explained using the principles of classical physics. For example, the operation of a transistor can be entirely described in terms of distinct, macroscopic voltage levels that behave and can be manipulated as a classical bit ("off" or "on"). In these contexts, any direct manifestation of the underlying quantum effects generally acts as a noise process that interrupts the desired operational behaviour, and significant effort is sometimes expended to mitigate these effects. Indeed, the emergence of uncontrollable quantum effects with increasing transistor miniaturization is one of the main reasons why the decades long Moore's Law growth in classical computing power has now all but broken down.

Second-generation quantum technologies, by contrast, focus on the direct manipulation and exploitation of a given quantum system's quantum mechanical wave function, in a way which cannot be described classically. This could be through manipulating fundamental particles such as electrons, photons or nucleons or through macroscopic quantum wave functions such as the superconducting current of an electrical circuit or the collective wave function of a Bose-Einstein condensate.

We have discovered that, using direct manipulation of the wave function of a quantum mechanical system (2012 Physics Nobel) (Haroche & Wineland, 2012), we can perform specific information processing tasks that are not possible when the underlying physical system is not directly exploiting the quantum physics, i.e. it is behaving classically. These new information processing possibilities have direct relevance to numerous areas of human life, from information security and cryptography to ultra-high resolution sensing to high-performance computation. Combined with the fact that humanity now has the engineering skills to manipulate and control quantum systems directly, these have driven the expansion of quantum information science into one of the largest fields of physics.

2. The current state of quantum computing

In this section of the report, we give a brief overview of where the field of quantum computing is now and how we got here, in terms of both hardware and software. We then consider what the path to scalable quantum computing currently looks like, and why it is important.

Quantum Hardware

As the quantum industry has evolved over the past 20 years, we have seen a plethora of different technologies proposed, and in some cases demonstrated, for scalable quantum computing. In the early to mid-2000s there were literally dozens of hardware proposals published and these have been slowly whittled down to approximately nine primary modalities that have received significant investment from governments, corporations or the private equity community. Australia has been at the forefront of many of the hardware technologies that are currently being developed for large-scale quantum computing and still has significant capacity (either experimentally or theoretically) in many of them, as illustrated in Figure 1 on the next page.

Fig. 1: Australian strengths in quantum computing hardware

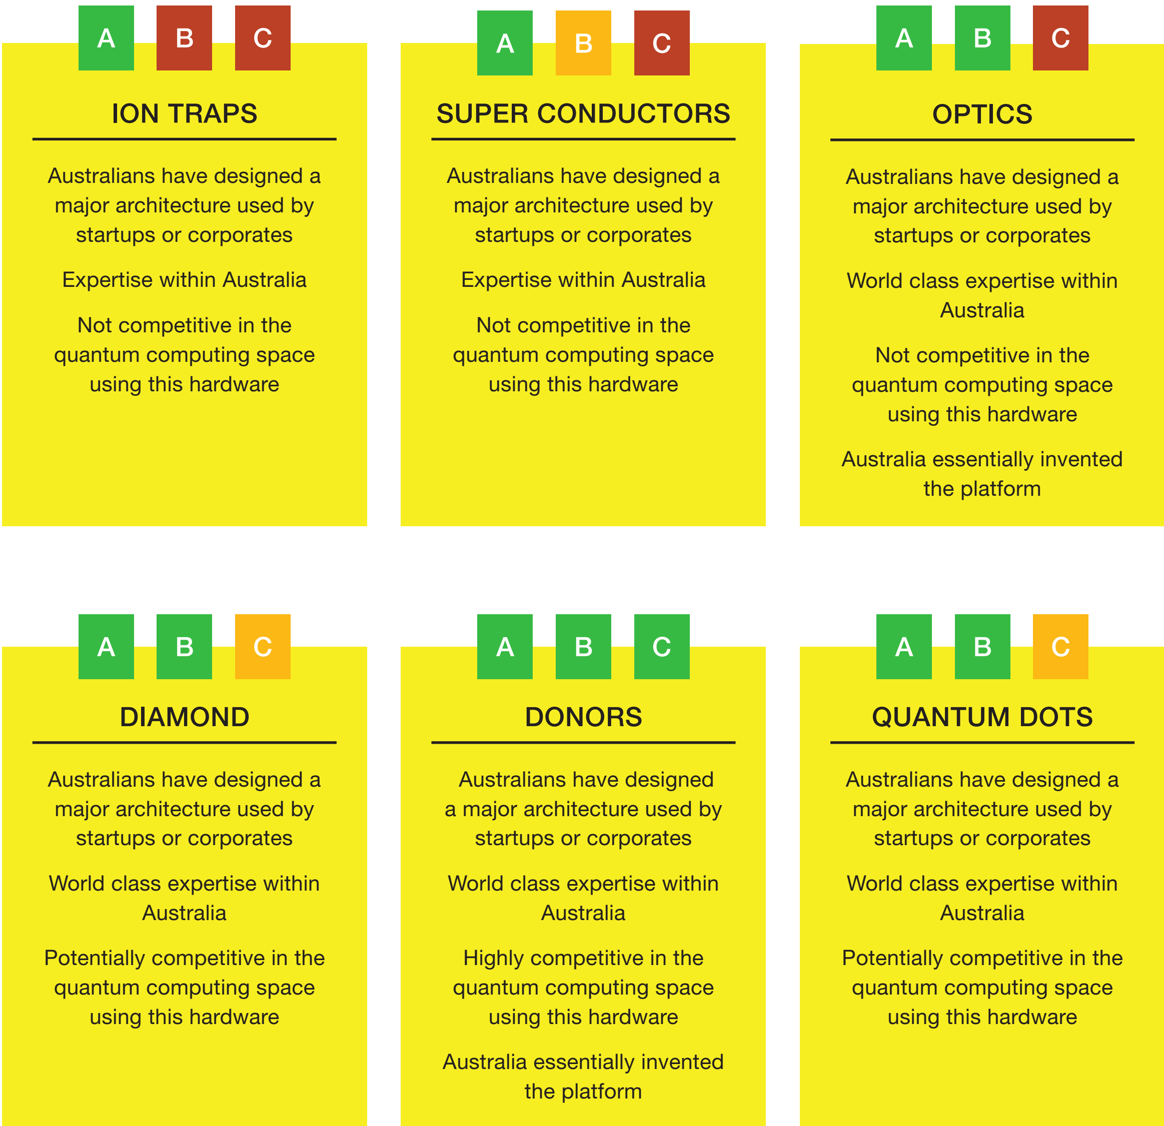

A — Have Australians designed the architectures and blueprints that are actively being built?

B — Do we have local talent to be highly competitive if resources are available?

C — Are we competitive in this space now?

While the specific timeline and developments for quantum hardware would fill multiple books, approximately nine primary hardware modalities for quantum computing have currently taken shape. These systems demonstrated the ability to manufacture quantum bits in a semi-reliable manner and has significant funding from either universities or national programs. These nine major systems are:

- Ion Traps (Zoller & Cirac, 1995).
- Superconductors (DiVincenzo, 1997).
- Diamond and other colour centres (Shahriar, et al., 2002)
- Photonics, discrete variable, single photon (Knill, et al., 2001) and continuous variable coherent laser pulses (Braunstein & Lloyd, 1999).
- Quantum Dots (Loss & DiVincenzo, 1998)
- Donor Based system (Kane, 1998)
- Neutral Atoms (Brennen, et al., 1999) (Jaksch, et al., 2000)
- Topological states of Matter (Kitaev, 2003)

These nine systems remain the dominant systems under development for quantum computing today, even though during the 2000s, they were part of a much more extensive list of proposed hardware modalities.

By the early to mid-2010s, each of these systems (except for topological states of matter) could claim to be able to routinely fabricate and control physical qubits, perform universal gate operations at moderate to high fidelity and even run small-scale test protocols such as quantum algorithms, and error-correction codes or communications protocols. This was when quantum computers began to move out of the laboratory and into the commercial world. The figure below shows a representative, but non-exhaustive evolution of qubit processor sizes from 1998 until 2022, illustrative of the general evolution over the past 25 years.

Fig. 2: Historical size of gate-based quantum computing chips

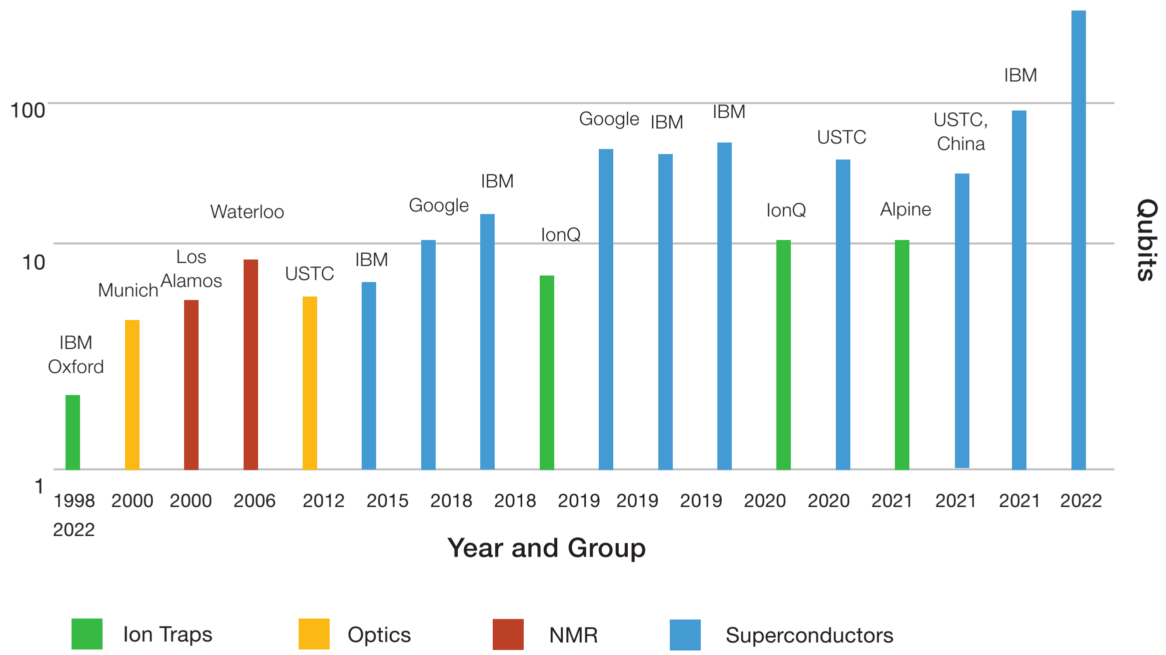

While chipset integration is getting better, especially with the increased investment from the private sector, quantum computers are still constrained by the error rates associated with the qubits and gates themselves.

When various qubit systems were first experimentally demonstrated starting in the late 1990s, 20-30% error rates were not uncommon (Monroe, et al., 1995) (Nakamura, et al., 1999). Since then, research and development across essentially all major platforms have decreased these physical error rates to the 1% and sub-1% level, depending on the type of quantum operation and the type of hardware system. It is now common across several qubit platforms to reliably fabricate and control qubits with high fidelities, in some cases above 99% (Barends, et al., 2014) (Wang, et al., 2020).

Due to the limited functionality of current quantum computing systems, researchers have spent a significant amount of time developing core theory surrounding the concept of how to demonstrate beyond classical performance (also termed quantum supremacy (Preskill, 2012)). This area of research is to deliberately design a quantum algorithm/circuit that is difficult for a classical computer to simulate/emulate. It was shown by a variety of authors (Boixo, et al., 2018) that a classical computer cannot efficiently emulate this sampling procedure as the number of qubits is increased (Bremner, et al., 2010) (Bremner, et al., 2016). These sampling algorithms/ circuits were explicitly designed to be the smallest possible quantum algorithm/circuit that we could prove could not be effectively simulated or emulated on classical computers, not to produce an algorithm of any particular scientific or commercial utility.

This challenge of demonstrating beyond-classical performance was taken up by the Google Quantum AI team, who, in 2019, published a paper that claimed to have demonstrated random circuit sampling in a chip-set of 53 superconducting qubits (Arute, et al., 2019). Whether this result really is “beyond classical” is somewhat controversial, as the Google experiment sits right on the threshold of what is potentially simulable with a classical machine, rather than clearly beyond the classical regime. However, the work demonstrated many highly beneficial aspects of their technology, including suppressing complex error channels and fabricating, testing and calibrating an extremely complex quantum chip. In 2021, researchers from USTC, China, unveiled the Zuchongzhi 2 superconducting chip, which realised random circuit sampling over 56 qubits. This is a more explicit demonstration of quantum random circuit sampling beyond the classical limit than the 2019 Google result (Wu, et al., 2021).

From a quantum hardware perspective, the advances described above are remarkable and represent a tour de force of deep-tech development combining cutting-edge scientific and engineering innovations. From a quantum software perspective, however, state-of-the-art quantum computing technology is still effectively limited to “toy” devices and proof-of-concept systems. In contradiction to the marketing efforts of a large part of the quantum computing sector, there is no application of scientific or commercial utility that can be implemented on a quantum computer that cannot be effectively solved on current classical systems. There is a significant effort within the theoretical quantum community and across many domain experts to both identify new domain problems that require enhanced computational power, benchmark the utility of quantum computing for these applications and provide rigorous estimates for the size of a quantum computer needed to run these new algorithms (DARPA, n.d.).

The major bottleneck is related to the additional physical resources required to effectively error-correct quantum chipsets. Quantum algorithms are extremely sensitive to errors during computation and Quantum Error Correction (QEC) protocols, needed to reduce the errors in the chipsets, require resources of their own.

For a quantum computer that is large enough to be of utility, Quantum Error Correction constitutes the vast majority of the computation performed by the system. That is, the principle computation performed by a large-scale quantum computer is to correct its own errors. For example, factoring a large, composite number, using Shor’s algorithm is one of the most impactful applications of a quantum computer, due to its utility to compromise RSA public-key cryptosystems. Without error correction, a quantum computer consisting of approximately 5000 physical qubits would be sufficient to break RSA-2048 (2048-bit public keys are on the higher end of what is currently used across the classical internet) (Brandhofer, et al., 2021). This is not possible using our current technology as the error rate associated with these qubits would need to be one part in one quadrillion (< $1x10^{-15}$). To perform enough error correction to allow this algorithm to be successfully run, we actually require a machine containing 20 million physical qubits (Gidney & Ekera, 2021) – 4,000 times more than the algorithm actually requires. This overhead is solely required to reduce the error in the system from 0.1% at the physical level to the 10^{-13}% needed to successfully implement the algorithm.

Fig. 3: Cheap Qubits - Why qubit cost cannot be tomorrow's problem

Factoring

SHOR 2048-FACTORING Application: Cryptoanalysis	20M PHYSICAL QUBITS
Total cost, $1000/qubit	$20 Billion
Total cost, $100/qubit	$2 Billion
Total cost, $10/qubit	$200 Million
Total cost, $1/qubit	$20 Million

C Gidney, M Ekera Quantum 5, 433 (2020)

Quantum simulation

Nitrogen Fixation Application: Agriculture/ Fertilizer	200.000 Logical Qubits
Total cost, $1000/qubit	$200 Billion
Total cost, $100/qubit	$20 Billion
Total cost, $10/qubit	$2 Billion
Total cost, $1/qubit	$200 Million

Microsoft QuARC group, PNAS, 114 (29) 7555-7560 (2017)

What quantum computing will ultimately be when we reach the scale needed for significant commercial or scientific utility is still unclear. Depending on how the systems scale, quantum computers may end up being ubiquitous technology, like the classical computer or they may be such a capital-intensive device, governments will have to partner just to build a single one. Much of this depends on the ultimate cost of fabricating qubits. Figure 3 above uses potential applications – factoring large numbers (Gidney & Ekera, 2021) and simulating the quantum behaviour of a biological molecule useful for creating fertiliser (Reiher, et al., 2017) – and imagines the cost of a machine capable of executing these applications as a function of the individual cost of a qubit. It's clear that unless qubits (and their supporting control and infrastructure technology) ultimately reach the sub-$1 level, a quantum computer may be like other large and complex scientific machines – such as the Large Hadron Collider in Switzerland – we may only be able to build one.

Quantum Software

Contrary to what one might imagine from looking at what gets the most coverage in media (and arguably funding), quantum computing is even now not a technology that is driven by advances in quantum hardware. In fact, many of the most important research results that ultimately drive investment in quantum computing have come, not even out of the quantum physics community, but out of theoretical computer science. For example, arguably the most important result in quantum computing was Peter Shor's paper describing an efficient quantum algorithm for factoring large numbers (Shor, 1994). Similarly influential has been Richard Feynman's paper arguing that a quantum computer is the best tool to model the behaviour of complex physical quantum systems, a task known as quantum simulation (Feynman, 1982). Suffice it to say, despite perhaps receiving comparatively little of the glory, quantum theory, algorithms and software have played and will continue to play every bit as important a role as quantum hardware in reaching the ultimate goal of realising economic, commercial and societal benefit from real-world quantum computers.

Currently, despite their immense individual promise and ubiquity of potential impact, there are relatively few tasks known for which quantum computers will be able to outperform conventional "classical" computers. Broadly speaking, these tasks include (Montanaro, 2016):

- Cryptanalysis (exponential speed-up): Utilising subroutines like factoring and discrete log (Shor, 1994), a quantum computer can break most commonly used cryptographic protocols like RSA and Blockchain (Aggarwal, et al., 2018), by efficiently solving the complex mathematical tasks that underpin them.
- Quantum simulations (exponential speed-up): A quantum computer can efficiently model the behaviour (including dynamics) of complex physical quantum systems (Feynman, 1982) (Lloyd, 1996)
- Linear systems solving (exponential speed-up with caveats): Under certain problem conditions, and provided the solution can be used in an efficient way, the HHL algorithm enables quantum computers to efficiently "solve" a system of linear equations (Harrow, et al., 2008). While it provides an explicitly "quantum" solution which may not always be convenient to access, linear systems are so ubiquitous across so many areas of science, applied mathematics and computer science, that this is still of great potential interest.
- Search and optimisation (mainly quadratic speed-ups): Grover's search algorithm (Grover, 1996) and amplitude amplification (Brassard, et al., 1998) allow quantum computers to solve, respectively, unstructured and heuristic search problems (slightly) more efficiently than the best known classical algorithms. Quantum walks (Venegas-Andraca, 2012) can be used to provide sometimes exponential, but mainly quadratic speed-ups over classical optimisation tasks by using a quantum analogue of stochastic numerical algorithms.
- Quantum machine learning (unknown speed-ups) (Peral Garcia, et al., 2022): Currently, it is not known whether quantum computers can provide any concrete speed-up over classical algorithms for machine learning and artificial intelligence, but given the broad reach of these application contexts, this is the subject of a great deal of ongoing research. Part of the challenge here is in identifying what is the best-known classical algorithm. For example, proposals for quantum machine learning have been known to inspire new, and much more efficient, classical algorithms, in turn negating a previously identified quantum speed-up.

One of the great challenges to finding quantum algorithms with large speed-ups over classical computing is that, by definition, quantum computers can only provide a meaningful advantage for tasks which cannot be solved efficiently by classical computers. It is important to remember that classical computing is already incredibly advanced, a robust and powerful technology that is extraordinarily well suited to the tasks it does well. So any task that can already be solved on a classical computer, should be solved that way, because it will inevitably be cheaper and more reliable.

Fig. 4: Quantum software - The four major levels of the stack

Quantum software contains four broad categories of packages, all being developed at different rates.

ALGORITHMIC

Development of quantum algorithms appropriate to commercial and/or scientific problems with provable speedup over best known classical techniques.

LEVEL 1

CIRCUIT

Creating quantum circuits for high level algorithms and optimising for total number of algorithmic or logical qubits, total number of timesteps etc ...

LEVEL 2

ERROR CORRECTION

Embed and optimise complex error correction protocols.

LEVEL 3

HARDWARE CONTROL

What signals do we actually send into the computer to perform a gate? How can we “tune“ these pulses to reduce error rates as much as possible to reduce error correction requirements.

LEVEL 4

Quantum software covers a broad spectrum of research and technology development, from algorithm design to hardware control optimisation. There are four primary levels to the quantum software stack. Developing quantum algorithms like those above sits at the highest level (Montanaro, 2016). Beneath this is the software required to compile abstract quantum algorithms into the machine-level instructions needed to execute the algorithm. This is commonly referred to as the quantum circuit description and would be akin to individual transistor instructions in a classical microprocessor (Heim, 2020).

The level beneath this covers quantum algorithms that require error-correction protocols. This takes the quantum circuit specification and generates another, different set of machine-level instructions that embeds error-correction code into the computation. The majority of the work of a quantum computer for these algorithms is actually the error-correction part of the computation, with estimates for some algorithms indicating that over 90% (Gidney & Ekera, 2021) of the circuitry needed to implement the algorithm is dedicated to the computer correcting its own errors.

Finally, we have hardware control or firmware software (Ball, et al., 2021). This is the software that takes as input a gate call – i.e. do this quantum gate on this particular qubit – and converts this into the physical physical “thing” that manipulates the qubit. This could be a laser pulse in the case of ion traps or it could be an electrical signal in the case of superconducting circuits. Embedded within this software level are numerous techniques to increase the accuracy of the quantum gate. Commonly referred to as error mitigation techniques (Cai, et al., 2022), quantum control software does its best to create the most accurate quantum gates possible at the physical layer, as this reduces the burden on the active quantum error-correction layer that carries a physical resource overhead.

There are now numerous quantum startups and corporations building out software solutions at each layer of the quantum software stack, as illustrated in Figure 5 below. Typically, these efforts are all highly individual, producing tools that are often designed to augment a specific hardware vendor and written in different languages. As such, these software solutions do not generally talk to each other or allow for a common framework to translate them to various different potential hardware modalities. Another issue is that these tools are generally designed and marketed by following analogies to classical software and programming languages. Quantum computing design and hardware are so nascent, however, and their conceptual underpinning so different from classical computing, that the quantum/classical analogies that motivate their design are often inaccurate.

Fig. 5: Quantum Programming Languages

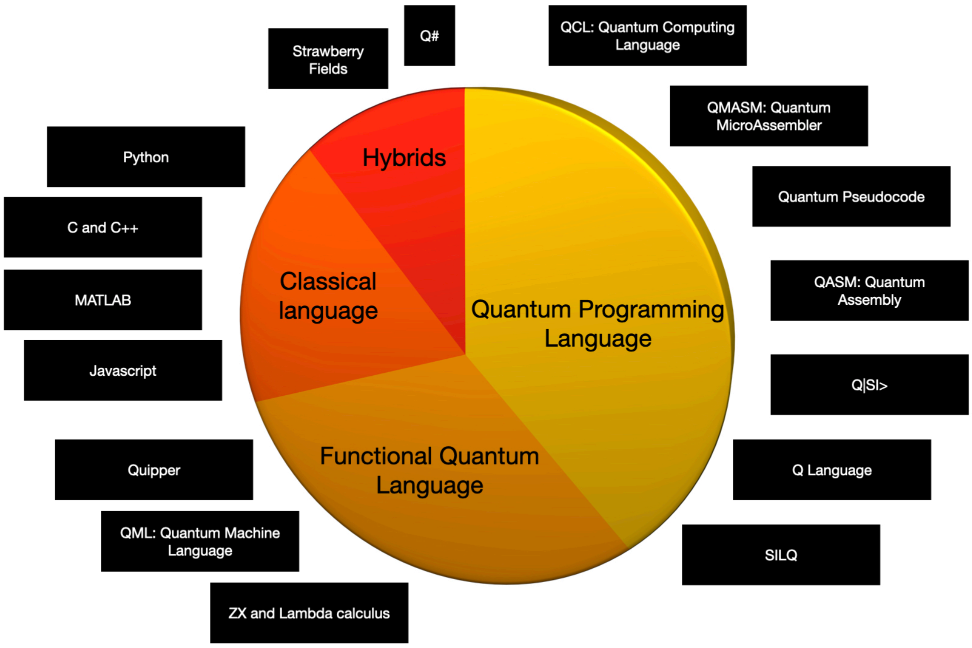

The classical programming frameworks of languages, compilers, transpilers, etc, have been built up over the nearly 75-year history of classical computing to provide higher and higher levels of programming abstraction away from the details of the operation of a classical computer. When a program is run, however, the coding language a programmer uses, usually designed to be easily readable and debuggable, undergoes an extensive and complex process of translation into individual transistor instructions that control the microprocessor, RAM and Cache memory, and input and output control, including the monitor, keyboard, mouse, etc. This part of the software control and compilation for a classical computing system is invisible to the vast majority of users and even professional programmers of classical computing systems.

In the case of quantum computing, we are currently programming literally at the level of individual transistor instructions. The only layer of abstraction that is somewhat commonly available today in the quantum computing ecosystem is how a particular quantum gate call is converted into the particular laser pulse or electronic signal that is ultimately sent to the physical qubit (the firmware control layer). At this level of abstraction, however, given the large number of competing hardware modalities being used, these translations are very specific to the hardware and platform, and even to the calibration of an individual processor, and as such must be tailored from system to system.

The need for Scalable machines.

With quantum processors starting to reach sizes of 10s-100s of physical qubits, mostly operating at high error rates, quantum computers have entered the so called (N)oisy, (I)ntermediate (S)cale (Q)uantum regime, a term coined in 2018 by Caltech professor John Preskill (Preskill, 2018). While the definition of NISQ quantum computing can be at times very nebulous, a useful working definition is (Brandhofer, et al., 2021):

> ***A quantum computation small enough that it can be faithfully executed on a noisy, small to medium-scale quantum chip, without the need for resource-costly quantum error correction protocols.***

In recent years, NISQ quantum computing has been a major commercial focus for many quantum hardware and software companies, and it has been hoped that algorithms could be found that could provide commercial utility even for NISQ-scale quantum computers. As of 2023, however, the hope of finding commercially valuable quantum algorithms in the NISQ era is starting to diminish. Many hardware companies are now developing roadmaps for large-scale quantum computing architectures (IBM Roadmap, n.d.) (Google Roadmap, n.d.) (Rigetti Roadmap, n.d.) (AWS Roadmap, n.d.) and focusing on how to quantify the resources required for commercially viable quantum computing and find new techniques at the algorithmic and error-correction level to reduce these resource costs as much as possible.

These roadmaps generally fall into three categories:

1. Build a large processor that does not allow scalable quantum computing, and try to work out whether it can do something useful that surpasses classical capabilities. In this category: DWave and companies targeting analogue quantum simulations (D-Wave Roadmap, n.d.).
2. Build a small-scale processor that works as a general purpose quantum computer at small scales, and then try to solve the platform's scaling problems as you make bigger processors. This includes most superconducting qubit processors (IBM, Google, Rigetti, Quantum Circuits Inc., etc) and most ion-trap quantum computing (IonQ, Alpine Quantum, Quantinuum, etc).
3. Build a harder processor that tries to embed a solution to scalability intrinsically from the ground up, and then try to make it work at small scales. Examples include: Microsoft (topological quantum computing), PsiQuantum (one-way photonic quantum computing), donor and quantum-dot based silicon quantum computing (SQC, Diraq, Quantum Motion, and others).

The vast majority of current commercial quantum computing efforts are aiming to follow one of the latter two approaches.

The transition from small-scale quantum chipsets to progressively larger systems that can provide utility beyond quantum education or device testing will be a long process. This is not just because of the significant engineering challenges associated with the fabrication and control of large qubit chipsets, but also the development of the infrastructure that will eventually be needed to produce large quantum computers at a cost appropriate for their utility.

Examining the historical trend of classical transistors, the digital revolution of the 20th and 21st centuries has taken place because computational power has become ubiquitous and commonplace. A well-known quote from IBM's Thomas Watson in 1943, where he stated that "I think there is a world market for maybe five computers", in many ways encapsulates the potential dilemma now facing quantum. The digital revolution could not have been so profoundly influential without economies of scale allowing for massive investment in fabrication infrastructure, driving down costs to the point where individual transistors are one of the cheapest functional objects in existence. Will quantum technology similarly become as cheap and ubiquitous as we move towards the 22nd century, or will Thomas Watson prove to be correct after-all, but for the quantum age?

As we have noted, the current status of quantum computing algorithms development suggests that there is little utility in small to medium-scale quantum algorithms. The known algorithms are either not providing a computational advantage commensurate with their cost relative to state-of-the-art classical techniques, or, to reach the scale where they provide definitive benefit, require quantum computers larger in size than the estimated 20 million physical qubits necessary to compromise RSA-2048 public key cryptosystems (Gidney & Ekera, 2021). Indeed, a key current goal for algorithms researchers is to improve algorithmic and error-correction techniques sufficiently to even reduce the number of qubits needed for clear utility to “this side of Shor” (that is, below 20 million physical qubits). At this level, there are reasons for optimism, courtesy of research into Shor’s algorithm itself. In 2012, estimates (Devitt, et al., 2013) benchmarked out the cost of compromising RSA-2048 encryption keys at approximately 10 billion physical qubits. Between 2012 and 2019, this was reduced to the 20 million number, commonly accepted today (Gidney & Ekera, 2021). This reduction of nearly a factor of 1000x was done purely using better techniques in designing and implementing the algorithm and error correction protocols. Both results assumed the exact same performance of the quantum hardware itself.

Despite the rapid growth we have seen in quantum computing over the last decade, there are, unfortunately, still comparatively few experts, worldwide, in analysing and optimising quantum algorithms to allow commercial or scientific utility to be achieved using smaller quantum computers. The example above provides us strong evidence that once appropriate researcher-hours are dedicated to the problem, significant advances and innovation can be found. Currently, the only other algorithm to receive significant effort into optimisation is that of quantum simulations, but even there, there are still many more questions than answers. As other quantum algorithms become as well studied as Shor’s algorithm, there can be much optimism that significant new optimisations will be found. Will this allow us to find an extremely useful application that requires only 100,000 physical qubits or less? We do not know the answer.

To the best of our knowledge, quantum algorithms providing a computational advantage for any problem of use will require a machine containing at least 1 million physical qubits (Babbush, et al., 2018) and for many problems, a machine much larger than this (Reiher, et al., 2017) (Webber, et al., 2022). The developmental timeframe of such a large device will be dependent on how the investment landscape evolves, but even relatively optimistic estimates would usually put the timeframe at the middle of the 2030s at the earliest.

The most significant unknown here would be a scenario arising where a minor advance in computational heuristics can be achieved with a smaller-scale quantum machine. A quantum computer acting as a computational co-processor, even providing only a small percentage increase in efficiency for computations needed for design, manufacturing or logistics of economies of scale, could provide a significant advantage to the nation or entity that possesses that technology.

3. The quantum computing end game

As the new quantum computing industry develops apace, research and investment decisions are usually based on some kind of attempt to predict how the industry is going to develop over the next 5, 10 or 20 years, driven by the promise of significant commercial reward to be gained by those who will, in hindsight, be found to have made the best choice. In a field like quantum computing, however, which is growing and changing so rapidly, where vital progress is being made across many disciplines simultaneously, and where the variability between different platforms and approaches is so large, such predictions are adrift in uncertainty. This is only magnified by the immense challenge of the task being faced in building a scalable quantum computer.

But while the specific near-term path is very hard to predict, what about the eventual destination? After almost three decades of research since Shor published his revolutionary first application of large-scale quantum computing, there is now a broad consensus that there is no foundational reason to doubt that quantum computing is possible. It is now not a question of if, but when, we will see its first real-world applications demonstrated. In this report, we do not try to predict how quantum computing will develop in the near term, but instead imagine five different possible end-game scenarios for where quantum computing technologies could eventually land, and consider what we can learn from them about the path needed to get there. Which of these scenarios ends up being closest to reality will emerge out of how the tension between two key aspects is resolved: namely 1) how ubiquitous / cheap / noiseless qubits become at the hardware level, and 2) how far algorithmic resource requirements can be optimised at the software level.

To reach parameter regimes of interest, quantum computing circuits must be big. For example, Shor's algorithm was used to factor numbers up to four bits long without approximation on an Ion Trap system in 2016 (Monz, et al., 2016), but standard key lengths for RSA cryptosystems are upwards of 2048-bits. One particular algorithmic implementation of Shor's algorithm to break a key of that size would require 4100 qubits and 274 million gate steps (Beauregard, 2003). At this size, physical errors in a NISQ-style quantum computer would need to be absurdly low, well beyond anything capable by any experimental system in the foreseeable future. Consequently, extensive quantum error correction will be required to achieve algorithms of this size, which includes effectively all algorithms that have been argued to have scientific or commercial utility for quantum computing.

Currently, the best estimate of the number of physical qubits required to factor RSA-2048 with an error-corrected fault-tolerant computer is approximately 20 million qubits for a superconducting quantum computing architecture (Gidney & Ekera, 2021). However, this value can be further reduced through improvements in qubit performance, and there may also be significant room to improve qubit counts in the software level through algorithmic optimisation. These questions also depend very heavily on the precise platform underpinning such an eventual quantum computer.

1. The LHC scenario: This scenario represents the case that building a quantum computer at sufficient scale to realise commercially or societally beneficial applications requires a massive coordinated global collaboration involving academia and industry participants across many, many countries, as was the case for the Large Hadron Collider that discovered the Higgs boson, or forthcoming Square Kilometre Array telescopes. Such a scenario would likely need to be driven and funded by massive multilateral agreements at the government level with one or maybe a handful of such quantum computers being constructed. In an LHC-type quantum computer, one can imagine that each component can be highly optimised and highly specialised, likely the result of one or more dedicated PhD-level research projects. Industry involvement in such an enterprise may take the form more of a service model where companies may provide specialist parts on commission, rather than a standard commercial model. But like in the moon-shot space race, we could also expect many profound innovations and spin-off technologies to arise from the research required to realise the end goal.
2. The corporate provider/data centre scenario: In this scenario, the closest to the way companies currently provide access to cutting-edge quantum hardware, each quantum computer is built and maintained by a commercial party, with access (usually cloud-based) to computing time sold to end users. While such efforts may be government-subsidised to some degree initially, eventually these quantum computing data centres would be funded by the commercial manufacturer/provider, with cost being recouped via the chosen access model. One could imagine seeing up to some 10s of providers offering quantum computing data centres based on different technology platforms, potentially with different competitive benefits, costs and specialised focuses.
3. The corporate user/HPC cluster scenario: In this scenario, quantum computing hardware providers would manufacture and sell quantum computing servers to individual corporate users or consortia, or other large local entities like government departments. Similar to the way that high-performance computing clusters are sold today, the user would then maintain

and operate the quantum computer, with individual control over access. Depending on cost and other resources, such a scenario might see thousands to millions of computers worldwide, and providers offering differing technology platforms potentially even with different application specialities.

4. The home desktop/QPU scenario: If per qubit cost comes way down and environmental constraints are relaxed, the next scenario would see quantum processing units added to individual home desktop computers. Apps would then access the QPU as an accelerator component, a bit like an advanced GPU. Such QPUs would need to be highly compatible with the standard classical computing technology they would be embedded in. One could imagine some overall cooling might be required, but effectively the QPUs would need to operate at room temperature.
5. The smartphone/tablet scenario: For this scenario, quantum computing hardware would need to be manufactured at extremely small scales with frugal resource requirements. Fully miniaturised and monolithic platforms would likely be required, and would have to be fully room-temperature operable.

The last two scenarios would require extraordinary, not-easily-predicted technology breakthroughs. Most corporate quantum computing players currently seem to be anticipating quantum computers to end up somewhere in the regions of scenarios 2 or 3.

Scales and requirements

The table in Figure 6 summarises some of the practical constraints quantum computers would face in each scenario.

Fig. 6: Quantum computing requirements at different deployment scales

LHC	Multinational corporate provider/Data centre	Corporate user/ HPC cluster	Home desktop/ QPU scenario	Personal (portable) device
1 or few	10s	1000s-millions	100s millions – billions	Several billions
Major cryogenic infrastructure possible	Major cryogenic infrastructure possible	Some cryogenic infrastructure possible	Room temperature required	Room temperature required
Increased cost per qubit tolerated	Requires decreasing cost per qubit →			Cheapest cost per qubit required
~$20B @ $1000/ qubit for RSA-2048	~$200M @ $10/ qubit for RSA-2048	~20M @ $1/qubit for RSA-2048	$4000 @ $0.02 per 100 qubits for RSA-2048	$400 @ $0.02 per 1000 qubits for RSA-2048
Increased power consumption per qubit tolerated	Requires decreasing power consumption per qubit →			Lowest power consumption per qubit required
Large number of networked modules tolerated; hybrid platforms optimised to specific functionalities feasible	Requires more miniaturised, more monolithic → ← More tolerance for modularity and hybrid platforms			Fully miniaturised and monolithic required; hybrid platforms unlikely

Given the widely varying physical and resource requirements associated with different potential scales for quantum computing platforms, there will also obviously be different key challenges to overcome if their development is to be achieved. For example, an LHC-scale quantum computer could still be operated at cryogenic temperatures, whereas smartphone QPU accelerators would absolutely need to operate at room temperature. An LHC-scale quantum computer could readily adopt hybrid hardware platform designs, whereas smartphone QPUs would need to be monolithic and highly miniaturised. The table in Figure 7 summarises some of the key roadblocks across software and hardware that quantum computing will face across different deployment scales.

To provide some context on cost, consider just two operational aspects of a simple superconducting quantum computing experiment based on transmons controlled by off-the-shelf electronics, microwave and DC control. A typical transmon experiment could require up to 3 arbitrary fast voltage control channels and 1 high-precision microwave generator. Even without considering all the other experimental paraphernalia (amplifiers, readout controls, cryogenic operating expenses, nanofabrication costs, and much more), this would put the per-qubit cost of a transmon experiment in excess of $20k.

Fig. 7: Roadblocks to quantum computing at different deployment scales

Scale	Key roadblocks
LHC	Software: Control compatibility, algorithm optimisation, new algorithms; Hardware: Component compatibility, Cost
Corporate provider/Data centre	Software: Algorithm optimisation and validation, new algorithms, data security, software compatibility; Hardware: Cost, performance verification and benchmarking
Corporate user/ HPC cluster	Software: Algorithm optimisation, new algorithms, programmability, software compatibility and portability, software abstraction, verification and validation; Hardware: Cost, operating temperature, robustness and maintainability, due diligence/compliance/performance benchmarking
Home desktop	Software: Lack of relevant applications/algorithms, software compatibility and portability; Hardware: Operating temperature, power consumption, cost, miniaturisation
Personal (portable) device	Software: Lack of relevant applications/algorithms, software compatibility and portability; Hardware: Operating temperature, power consumption, cost, miniaturisation

Scales, applications and end users

Different quantum computing applications may be required or useful, depending on the technology scenario realised. In the table in Figure 8, we consider, we consider the main known quantum computing algorithms which offer beyond classical performance, and make some guesses at the different main quantum computing applications that would be most relevant at each scale.

Fig. 8: Relevant quantum computing applications at different deployment scales

likely relevant possibly relevant unlikely relevant

Application area	Quantum Algorithms	Speedups known?	Qubits	LHC	Corporate provider/Data centre	Corporate user/HPC cluster	Home desktop	Personal (portable) device
Cryptanalysis / Cryptography	Shor's factoring/ discrete log	Exponential	~20M	unlikely relevant	likely relevant	likely relevant	possibly relevant	possibly relevant
Quantum simulations	Trotterisation	Exponential	~200M	likely relevant	likely relevant	likely relevant	unlikely relevant	unlikely relevant
Search and Optimisation	Grover's, Amplitude amplification, Simulated annealing	Quadratic		unlikely relevant	likely relevant	likely relevant	possibly relevant	possibly relevant
Quantum machine learning, AI	QAOA, Variational algorithms	Polynomial?		possibly relevant	likely relevant	likely relevant	likely relevant	possibly relevant
Beyond classical benchmark problems	BosonSampling, IQP	Exponential	10s-100s	unlikely relevant	unlikely relevant	unlikely relevant	unlikely relevant	unlikely relevant
Maths: Solving linear equations	HHL	Exponential*	unknown	possibly relevant	likely relevant	likely relevant	likely relevant	likely relevant
Maths: Integration, Summation etc	Nayak and Wu	Quadratic		unlikely relevant	possibly relevant	possibly relevant	possibly relevant	possibly relevant
Maths: Graph theory	Quantum walks	Polynomial/ Exponential	??	possibly relevant	likely relevant	likely relevant	possibly relevant	possibly relevant

Different applications are also likely to be relevant to different categories of end users. For example, it is unlikely that pharmaceutical simulations would ever actually need to be run in a home computing context, and health providers would be unlikely to have the resources to maintain and operate their own QHPC facility for individual diagnostic capacity. In the table in Figure 9, we consider, we consider different classes of end users and anticipate which technology scale might be most relevant for some potential applications (some more speculative than others).

Fig. 9: Relevance of quantum computing to end users at different deployment scales

likely relevant — possibly relevant — unlikely relevant

End user	Relevant applications	LHC	Corporate provider/ Data centre	Corporate user/HPC cluster	Home desktop	Personal (portable) device
Global academic and government consortia	Major humanitarian problems	likely relevant	likely relevant	likely relevant	unlikely relevant	unlikely relevant
National defence / military / security	Cryptanalysis, Search and optimisation, Machine learning and AI	unlikely relevant	possibly relevant	likely relevant	likely relevant	unlikely relevant
Finance corporates	Portfolio optimisation	unlikely relevant	likely relevant	likely relevant	unlikely relevant	unlikely relevant
Pharmaceutical corporates	Drug and vaccine development, Studying disease pathways	unlikely relevant	likely relevant	likely relevant	unlikely relevant	unlikely relevant
Health / medical diagnostics	Disease diagnosis	unlikely relevant	possibly relevant	unlikely relevant	likely relevant	unlikely relevant
Advanced manufacturing / materials	Novel materials, logistics and optimisation	unlikely relevant	likely relevant	likely relevant	possibly relevant	unlikely relevant
Logistics / Transport / Utilities	Supply chain optimisation, scheduling optimisation	unlikely relevant	likely relevant	likely relevant	unlikely relevant	unlikely relevant
Personal / home users	Subroutine acceleration incl. linear systems and machine learning/AI	unlikely relevant	unlikely relevant	unlikely relevant	likely relevant	likely relevant

Figures 8 and 9 show that it can be instructive to think about the eventual product scale scenario to consider which industry verticals are most likely to be able to exploit future scalable quantum computing technology, and even how far along the technology development path we need to travel before we start hitting a point where advances and miniaturisation offer diminishing returns. One point that is very clear, the question of whether benefit is likely to be realised for home-scale quantum computing could depend very much on what new quantum algorithms and subroutines are discovered in the coming decades.

4. The role of standards in quantum computing

Standards have for centuries played a critical and positive role in commercial activities across all walks of life. Whether that be the centuries-old standardisation of ingredients in German beer, the awarding of licences to drive cars and practise medicine, the reduction of fuel emissions in vehicles, or the recent redefinition of the kilogram in terms of fundamental physical constants, many daily activities are built on layers upon layers of agreements on standardisation. From a commercial perspective, standards can remove barriers to global trade, improve supply chain resilience and reliability, and open up new markets for product distribution. From a consumer perspective, standards can improve consumer safety, facilitate accurate performance comparisons, and increase confidence in new technologies. As a young field that is rapidly growing, but also maturing in terms of breadth and critical mass of activity, quantum computing is now reaching the point where many of these considerations have great commercial relevance.

One of the main goals of this report is to discuss the role that standards and standardisation can play in the future development of scalable quantum computing. In the next section, we suggest some motivations and opportunities for developing standards in quantum computing, briefly summarising the current status of quantum computing standards. Finally, we consider how the standards landscape might change depending on what the final scale ends up being for practical scalable quantum computing.

Motivations for standardisation: Needs and opportunities

Formal standards development has emerged only relatively recently in quantum computing. Before the recent avalanche of new commercial players entering the quantum industry, adoption of quantum computing standards happened slowly and organically. Being largely publication driven, this did result in some kind of ad hoc consensus building, but in an environment where outcomes can be distorted for example by positions driven by influential players, it was sometimes an imperfect consensus outcome. While reasonably robust, this organic evolution has also been unable to keep up with the rapid pace of expansion we are currently witnessing within the quantum industry.

At the other end of the scale, we are now seeing formal quantum standards being developed by a range of recognised Standards Development Organisations (SDOs). The ISO/IEC JTC 1 has established a working group on Quantum Computing (WG14) with a variety of open projects. The IEC has published a white paper on Quantum Information Technology (IEC Market Strategy Board, 2021) which advocates that relevant SDOs need to maintain active relationships to better track parallel quantum computing standardisation efforts, and propose the development of a mechanism to assess standardisation readiness levels for emerging technologies to better inform ongoing quantum standardisation decisions. The IEC has also created the Standardization Evaluation Group (SEG) 14 on Quantum Technology aiming, among other things, to summarise use cases and propose a roadmap for standardisation in quantum technologies. The aim of formal international standards development processes is to produce robust and equitable consensus among the broadest possible cross-section of stakeholders, but a significant challenge is that these massively multilateral, committee driven processes also operate slowly and with considerable inertia by design. They consequently also struggle to keep pace with the rapid progress taking place in quantum computing. This is further complicated by the fact that there is currently limited overlap between the communities of experts in standardisation and quantum computing.

Currently, the most agile formal standards development is currently being driven more locally by national (or European) standards bodies (NSBs), professional bodies and industry-engaged consortia, arising especially out of flagship quantum strategies such as those of the US National Quantum Initiative and the EU Quantum Flagship. The IEEE, Quantum Economic Development Consortium (QED-C) and European CEN-CENELEC all have active standardisation activities covering a range of topics. One successful strategy these organisations are using to engage with large groups of quantum experts from industry and academia, is through focussed workshops aiming to bridge gaps between quantum and standards. These enable a rapid and agile inclusive position to be reached within an incomplete, but significant cohort of experts.

At its most fundamental level, standards are a tool to facilitate and streamline interactions between different parties to social or commercial engagement, particularly when they come from divergent backgrounds. We consider five different spheres of interaction and engagement relevant to quantum computing, and identify motivations and opportunities for quantum computing standards development within each of these spheres (see (Abdelkafi, et al., 2018), (ISO, n.d.) and (Standards Australia, n.d.) for useful general information about different types of formal standardisation activities):

1. Quantum computing is heterogeneous: Engaging within the community

KEY POINTS:

- The existence of widely varying hardware and software platforms creates significant challenges for standardisation processes, but also opportunities.
- In both hardware and software, standardisation could improve communication and compatibility between different technology platforms and layers.
- Focussed initiatives driven by governments, broad industry consortia or key players, are proving most agile at starting to develop standards.

The need: Currently, there are numerous hardware platforms under active development and at different degrees of maturity. They build on wildly disparate physical technologies, from optical-frequency travelling photons, to resonant microwave electronics, to magnetically trapped arrays of individual ions, and cross over fundamentally contrasting quantum computing paradigms, from gate-based quantum circuits, to measurement-driven consumption of massively many-body entangled resource states, to deformation-driven manipulation of topologically active quantum states. As such, even communicating between the platforms at either hardware or software levels is extremely challenging. Benchmarking meaningfully across platforms is a greater challenge still, especially since they are often following fundamentally diverging development pathways towards scalability.

The opportunities: In this environment, standardisation, while challenging, offers valuable opportunities to help bridge these communication gaps. Quantum computing start-ups and corporates commonly recruit staff from across physics, maths, computer science and engineering to work highly collaboratively. *Terminology standards* could help establish lingua francas to facilitate cross-disciplinary communication within the community, and accessible yet expert-driven, purely informative *Technical Reports* (as opposed to the more prescriptive "normative standards") could help disseminate broader understanding of development trajectories and key challenges. Add to that the fact that quantum computers may ultimately rely on a hybrid platform integrating multiple hardware paradigms (e.g., photonically mediated, local trapped-ion quantum computing nodes), and the benefits of improving cross-paradigm communication and understanding become even clearer. Simultaneously, Measurement and Test Method standards could both improve cross-platform benchmarking, and to some degree reduce the level of cross-disciplinary expertise required by creating standard interfacing touchpoints for communication. Yet it is worth noting that developing such benchmarking standards is still an active area of cutting-edge quantum computing research (Cross, et al., 2019) (Wack, et al., 2021) (Lubinski, et al., 2021).

Within the hardware space, formal normative *Specifications* and *System Architecture* standards could facilitate better interfacing between platforms. Early standards would likely take the form of *Practice Guidelines* outlining recommendations and possible capabilities. In a hypothetical far future of networked quantum computing nodes, however, such standards would likely need to evolve into strict *Requirements standards*. Formal interfacing standards will likely be even more critical within the software space, and arguably even sooner. In particular, Software *Specifications*, *Frameworks* for *System Architectures*, and formal *Reference Models* could help avoid proliferation and dispersal of software architectures and designs. The challenge is that many of the underlying frameworks of quantum computing (e.g., compiler layer, stack design) are as yet very undefined and topics of active, ongoing research, so it is important to avoid the creation of too-rigid standards that inhibit either adoption (due to irrelevance) or innovation. To avoid these scenarios, broad expert and industry engagement will be critical.

The short term: Within the realm of formal standards, the ISO/IEC JTC 1 WG14 is currently developing a *Terminology standard* for quantum computing which has recently progressed the Draft International Standard stage. The IEEE has established a dedicated Quantum Initiative for Standards (IEEE, n.d.) that now has several open projects including on Performance Metrics and Benchmarking (P7131) (IEEE, n.d.) and Quantum Computing Architecture (P3120) (IEEE, n.d.). The Quantum Economic Development Consortium (QED-C) has already produced an initial suite of Application-Oriented Performance Benchmarks for Quantum Computing (Lubinski, et al., 2021) and a report on “Practical Intermediate Representation for Quantum (PIRQ): Requirements and Near-Term Recommendations” (QED-C, n.d.) for software interfacing. CEN-CENELEC have in turn established a Focus Group on Quantum Technology (FGQT) with one main goal of creating a standardisation roadmap for quantum technologies (CEN-CENELEC, n.d.). An article recently published by delegates from FGQT proposes potential frameworks for defining the quantum hardware and software stacks, as well as a “supply-chain model” for identifying use-cases for quantum technology standards (van Deventer, et al., 2022).

In the earlier days of quantum computing research, standards were mostly informal and adopted organically, mainly arising in the domain of *Measurement and Test Methods* for benchmarking. Examples of informal consensus standards developed in this way included:

- Quantum state and process fidelities accessed either through quantum state and process tomography (James, et al., 2001) (Chuang & Nielsen, 1997) (Poyatos, et al., 1997), or through reduced, targeted measurement sets (O’Brien, et al., 2004) (Hofmann, 2005). With multiple mathematical definitions initially in use for fidelity, these too converged to a single, operationally meaningful definition. It is well understood that quantum state and process fidelity are primarily useful for characterising small systems and are inadequate for larger device or processor performance benchmarking.
- Gate fidelities accessed through various related processes which can be loosely grouped as randomized benchmarking (RB) protocols (Knill, et al., 2008) (Magesan, et al., 2011) (Magesan, et al., 2012). These measures largely superseded state and process fidelity metrics once small-scale quantum processors containing multiple qubits began to emerge. However, while RB protocols provide a somewhat rigorously defined performance metric, it is generally found that efficient device calibration and tuning requires more targeted and platform specific measurements and calibration protocols (Arute, et al., 2019).

There have also been instances where more targeted approaches have successfully produced de facto informal standards. For example, IBM researchers released OpenQASM, the Open Quantum Assembly Language, in 2017 (Cross, et al., 2017) to provide a generic quantum circuit description, and many quantum programming languages and hardware control systems maintain compatible interfaces with OpenQASM (see Case Study 2). National funding organisations have also leveraged targeted funding calls to stimulate focussed research activity into individual outcomes identified as critical, and have been used to drive the development of rigorous benchmarking standards. For example, the US Defense Advanced Research Projects Agency (DARPA) recently ran a funding round on Quantum Benchmarking for quantum algorithms

(DARPA, n.d.), and the US Department of Energy (DoE) have recently opened a call to fund research into quantum transpilation to develop methods for converting algorithms between different quantum computing paradigms (US Department of Energy, n.d.).

Hardware standardisation is currently non-existent beyond initial efforts described above to create *Terminology Standards*. Currently, quantum computing machines are not designed in any way to be interoperable, even within the context of devices that utilise the same underlying technology: superconducting qubit processors built by Google would not be easily interchangeable with the superconducting chipset designs from IBM, and serious attempts to standardise fabrication processes or control protocols, or define interfacing standards are not currently taking place. Indeed any attempt to do so at this stage should be approached with great caution, because it is still unclear how to effectively build a large-scale machine, capable of fault-tolerant, error-corrected computation, in any modality. The next generation of quantum chipsets will begin to inform how technological standardisation can proceed. The standards community should continue to work closely with the quantum ecosystem as forthcoming hardware generations evolve, and in the meantime, standards development in the hardware space should be focussing on establishing foundational and peripheral standards (such as *Terminology standards* and *Framework guidelines*) that prepare the way for and facilitate the development of normative *Requirements Specifications* when the time is ripe.

Standardisation across quantum software is arguably where the most significant short-term impacts from the standardisation community can currently be realised. A rapid proliferation of quantum programming languages, circuit synthesis tools, performance analytics solutions and hardware interface software is currently taking place as more and more companies enter the space. Unlike large-scale algorithmic benchmarking and standardisation, quantum software standardisation does not necessarily require the existence of scalable algorithms or scalable machines. The evolution of software standardisation for classical computing was a long and complex process, with de facto standards often built upon techniques that evolved as the technology became more sophisticated. It is as yet unclear if quantum will need to evolve in a similar manner or if a more structured and targeted approach could be beneficial. The many open questions about scalable quantum hardware architectures, the variety of algorithmic techniques and error-correction protocols, and the jockeying for position between many companies to position themselves as defacto standards, does create challenges for early standardisation. But provided the efforts are thoroughly expert-driven and industry-engaged, this area provides interesting opportunities to start creating networks and conversations around standardisation, and to find ways to promote early cooperation on technology development throughout the global quantum ecosystem.

2. Quantum computing has extreme technology requirements: Engaging with suppliers

KEY POINTS:

- Quantum computing technologies often demand key aspects of performance from supporting technologies far surpassing what they were originally developed for.
- Developing *Requirements Specifications* for supporting equipment is one of the most important near-term opportunities for formal standards development.
- The growing quantum market offers new and lucrative economic opportunities for conventional technology manufacturers willing to engage with new requirements.

The need: The technical challenge of building a quantum computer is so profound and fundamental, it has been compared with grand scale human endeavours such as the moon landing or building the large hadron collider. Such massive feats of (quantum) engineering rely crucially on layers upon layers of supporting (classical and quantum) infrastructure, but often present these underpinning technologies with extreme operational performance requirements. Examples include: unparalleled coherence times achieved in isotopically pure silicon-28 “semiconductor vacuums”, enhanced trapped-ion heating lifetimes in ultrahigh cryogenic vacuums, ultrahigh photon-loss sensitivity in integrated and fibre-based photonics. Here, quantum performance is otherwise limited by noise processes that are completely undetectable for related classical technologies. Without quantum-targeted specifications standards, auxiliary infrastructure development requires time-consuming industry engagement with classical technology developers whose attention is focussed on their established large-volume classical clientele.

The opportunities: This is one of the most immediately accessible areas for immediate, active standardisation in quantum computing. It is also the area that will benefit the most from and leverage existing technology standards. As above, *Terminology standards* will crucially help mediate the conversations between quantum users and enabling technology providers, and targeted informative *Technical Reports* could be developed to describe individual technology use cases or connected families of use cases within quantum computing. However, unlike standards about the quantum technologies themselves, standardisation on the supplier side could likely immediately begin targeting creation of detailed (normative) *Requirements Specifications* for quantum use cases, along with any *Measurement and Test Methods* for certifying and benchmarking the enabling technology’s performance. Such specifications would likely co-opt many aspects of existing standards, and simply extend those aspects particularly critical for quantum use cases. To date, successful scenarios where suppliers of supporting technologies have expanded into the quantum market (such as with cryogenic microwave components) have often arisen off the back of detailed bilateral engagement with one or more quantum users who were seeking a specific capability that was not otherwise available. A combination of standards like the above could enable new suppliers to explore early feasibility without having to find and carry out time-intensive engagement with quantum experts (or at least to shortcut the process). This could bring new suppliers into the quantum market, create valuable competition and drive innovation. Such standards would also streamline procurement processes for the quantum community, and provide added value to suppliers through benchmarking their products against others in the market.

The main risk to standardisation efforts with this scenario is if technology standards become rapidly obsolete because quantum platforms and techniques change too rapidly. However, while performance improvements in peripheral technologies have sometimes driven significant improvements in quantum device performance, in most cases, the nature of those enabling technologies can probably be expected to remain fairly constant over many years. For example, dilution refrigerators have become bigger and more powerful over time, but they have only undergone one major technology evolution in the past 25 years (the evolution from wet to dry fridges). Further, quantum hardware platforms are no longer coming and going over short time scales, and many of these enabling technologies are in any case shared between platforms.

The short term: As the quantum market grows, we are more and more seeing technology manufacturers offer products and services targeted specifically at the quantum community. Generally, these providers fall into one of three categories: 1) enabling technologies that have been developed primarily to serve quantum technologies (e.g., large-volume cryostats); 2) purely classical technologies that are keen to fine-tune and re-market their products to the growing quantum market (e.g., passive cryogenic microwave components); and 3) purely classical technologies that merit sufficient adaptation and specialisation to the quantum context to allow an opportunity for new start-up providers (often spun out by quantum phd graduates) to target that niche (e.g., arbitrary waveform control systems for qubit control).

There is an advantage to one part of the quantum community if a certain level of standardisation could be realised, namely the creators of infrastructure and support technology. Quantum computing systems need extensive support infrastructure, this could be complex from dilution refrigeration systems, designed to cool quantum chips to a few hundredths of a degree above absolute zero, to as simple as high quality coaxial cable for use in carrying the microwave signals used to control individual qubits. There are companies whose main value proposition is the construction and sale of this peripheral infrastructure, and defining standards could help them sell their products to the largest possible customer base from both researchers and quantum technology companies.

3. Quantum computing will have broadly interdisciplinary impact: Engaging with end-users

KEY POINTS:

- End-user applications for quantum computing impact across diverse sectors; developing them requires highly interdisciplinary collaboration with end users.
- New algorithms and use-cases could strongly influence both its reach across sectors, and the technology scales it could be usefully deployed at.
- Appropriate standardisation could aid communication between quantum and end-user domain experts, and help establish reliable resource benchmarking for both classical and quantum algorithms, that will be crucial for algorithm development.

The need: Quantum computing is predicted to have impact across many, many diverse sectors of society and the economy. And as the previous section of this report illustrates, specific algorithms and use cases can strongly influence the technology scales at which quantum computers could be usefully and profitably deployed. Yet identifying and optimising application opportunities within end-use verticals will require intimate collaboration between end-user domain experts and the quantum industry, people with even more diverse backgrounds than those already working within the quantum industry. Good standards can help mediate communication between disparate communities and help to facilitate better and more efficient collaborations, and could eventually even obviate the need for such collaborations since ultimately, practical deployment of quantum computing for end-users cannot rely on educating all end-users to be quantum experts themselves.

The opportunities: The first key opportunities in this sphere of engagement are again in establishing *Terminology standards* and targeted *Technical Reports* to facilitate communication between quantum experts and relevant end-user domain experts. It then becomes a bit more challenging here, because questions around identifying and optimising quantum algorithms for end-user applications are some of the most challenging and under-explored in the field today, and represent cutting-edge research areas, making standardisation inadvisable. Ultimately, however, the impact of quantum computing within these end-user domains will be greatly increased if quantum computing subroutines can be utilised by end-users who do not have significant quantum expertise. To achieve this effectively, it is necessary to design the user tools with as much abstraction as possible away from the hardware control layers. The software tools in this scenario could benefit greatly from detailed software and architecture *Specification standards* and *Frameworks*, as described in the quantum community context discussed above. Once quantum computing is established as a reliable *provider technology* that can be deployed in end-user applications verticals, detailed *Requirements Specifications* and *Measurement and Test Methods standards*, potentially as part of Quality Assurance standards, would be required for both hardware and software technology aspects. This would be crucial for end-user confidence and due diligence for procurement and use quantum computing technologies for commercially critical activities.

The short term: The question of at what scale quantum computers will provide computational utility for a problem – scientific or commercial – is difficult to answer beyond “the quantum computer will need to be pretty big”. Maybe the best encapsulation of the question – from the standpoint of scientific utility – is a quote from Prof. Rod Van Meter at Keio University in Japan, who asks:

> ***“When will the first Nature or Science paper be published where the results of the paper were computed by a quantum computer, but where the paper has nothing to do with the quantum computer that performed the calculation.”***

Essentially, when do quantum computers reach a scale and utility that puts them into the same category as any other High-Performance computing infrastructure for either commercial or scientific applications?

The NISQ era of computation was defined and studied extensively on the hopeful assumption that computational utility could be found for quantum computers without having to worry about incorporating resource-costly error correction protocols. However, it is now being tacitly acknowledged by more and more researchers, government agencies and most importantly quantum computing companies, that genuine computational utility will not be found with small-scale quantum chipsets and that we may need to go to devices containing millions of qubits before quantum computers can outperform state-of-the-art classical HPC simulation in any scientific or commercial context. Consequently, algorithmic benchmarking and by extension, standardisation, are only now receiving significant attention from companies and government agencies.

The standardisation of large-scale quantum algorithms is very tightly intertwined with algorithmic benchmarking. Without quantifying the utility of a quantum algorithm and what scale a machine needs to reach before that utility can be realised, robust standardised benchmarks will be impossible to construct.

The standardisation community is not yet part of the discussions with efforts within the quantum ecosystem to provide detailed benchmarks on large-scale quantum computation and this should be of priority for the standardisation community. While many efforts to understand the utility of quantum computing exist within the private quantum sector, the results of which are not necessarily being made public, there are also several government-level research programs that are open-source and are making significant inroads on these questions.

4. Quantum computing is deep tech and rapidly developing: Educating a quantum workforce

KEY POINTS:

- Due to its rapid growth and reliance on cutting-edge science, the new quantum industry is currently facing global pipeline shortages in relevant expertise areas.
- Currently, the best opportunity for standardisation is to establish *Terminology* and *Requirements standards* to help train expert technicians in supporting areas.
- Innovation is still required, to identify if standardisation can play a further substantive role in addressing issues in the quantum workforce pipeline.

The need: Quantum is rapidly developing. Progress is moving so fast that it can be difficult to keep up, even for researchers working in the thick of it. This is made worse by the fact that the quantum ecosystem is simultaneously trying to continue growing in size, despite the global shortage in expertise. The challenge is that the growing quantum ecosystem needs to rapidly onboard as many new members as possible with a range of depths of quantum expertise. The quantum community needs to use all available mechanisms possible to support this growth.

The opportunities: It is not immediately obvious where standardisation can play a role in solving this particular challenge, beyond the perhaps simplistic idea that appropriately used standards can facilitate better communication between groups from different backgrounds and training. Probably the most significant opportunity for standardisation in this area is to focus on the training of technical experts in supporting or enabling areas that are nevertheless vital for quantum innovation at serious scales in academia and industry. Examples include: cryogenic technicians, microwave measurement engineers, fabrication process engineers, EM modelling engineers and software engineers. In this direction, standardisation can have a positive impact in much the same way that it can on the engagement between the quantum community and suppliers of supporting technologies. *Terminology standards* and targeted *Technical Reports* could establish the relevant lingua franca, and detailed *Requirements Specifications* could be used to carefully specify the technical aspects defining the scope of the supporting technician's or engineer's work. It may then be possible to develop standard short courses or technical certifications to provide the required cross-skilling or up-skilling to bring that person into the quantum workforce.

Beyond these initial suggestions, to have an impact across the broader quantum community, some innovative thinking about the role of standardisation is clearly required.

The short term: We have already discussed that quantum computing start-ups and corporates need to recruit staff from across physics, maths, computer science and engineering to work in highly collaborative teams. Yet most of these people, to some degree, still need to be able to engage with and discuss relevant aspects of the quantum systems they are dealing with. Currently, companies spend significant resources in bringing their staff up to speed on the quantum side of their work. This drives the interest that companies have in focussed quantum education initiatives like the Sydney Quantum Academy (Sydney Quantum Academy, n.d.), since it relieves them of significant work. Education platforms are emerging to address the issue of trying to broaden out access to the "mysteries of quantum", such as "Qubit by Qubit" (Qubit x Qubit, n.d.) and the "Learn Quantum Computation using Qiskit" textbook from IBM (IBM, n.d.). As a first step, standards development organisations could consult with industry employers and existing quantum education organisations such as Sydney Quantum Academy or Qubit by Qubit to identify where, when and if to undertake activities in this area.

5. Quantum is deep tech: Engaging with the wider community

KEY POINTS:

- As a deep-tech industry, rapid growth quantum computing risks out-pacing public licence and regulatory confidence, as other fields like AI have done.
- Developing robust, objective, widely adopted (and developed), comprehensive and accessible performance benchmarking standards will be difficult, but vital for due diligence and regulatory compliance for end-users, investors and government.
- Good terminology and informative standards could facilitate better communication with external stakeholders, to help build consumer/investor confidence.

The need: Arthur C Clarke's so-called third law is that "any sufficiently advanced technology is indistinguishable from magic". As an archetypal field of deep-tech development, quantum computing falls into this category. Making advances in quantum computing requires people to operate simultaneously at the forefront of both science and engineering. The structures and principles of quantum computing are so new and different that it is hard to both imagine and plan for what it might be possible to do with it. In this sort of environment, it can be very hard to distinguish bad actors and misleading advertising from legitimate, cutting-edge players, even for researchers within the field, since it is not so easy to develop simple intuitions about how everything works. Without relevant standards, commercial parties are incentivised to report selective benchmarks that present their particular technology in the best light. By providing transparency and rigour to benchmarking, testbeds design, and certifications and compliance, standards can help establish confidence with governments, investors and the wider public, and will be an important tool for trying to solve this problem.

The second major need in this sphere of engagement relates to public and regulatory licence to operate. When new technologies are hard to understand and grow very rapidly, they can create uncertainty and backlash in the wider public and regulations. In quantum computing, we wish to avoid the problems AI is now facing because it was not developed in parallel with discussions about ethics, guiding principles and regulations, which are now scrambling to catch up. Quantum tech is potentially in a similar position and having clear standards could help avoid some of these pitfalls.

The opportunities: Standards development in quantum is unlikely to have much opportunity to engage directly with the wider community, at least for quite some time to come, and is more likely to engage through interactions with service providers, investors and government stakeholders. In the short term, the main opportunities for standards development to engage in this sphere is with investors and government stakeholders. Government bodies may be interested in quantum standards both from an investment due diligence perspective, and from a regulatory approvals perspective. In all of these cases, *Terminology standards* and targeted *Technical Reports* would be crucial for establishing a productive dialogue. After that, the most important contributions from standardisation in quantum to engender investment and regulatory confidence in either investment or government sectors would come from establishing expert-driven, widely adopted, detailed *Requirements Specifications*, *Measurement and Test Methods* (benchmarking), and Quality Assurance standards, for performance benchmarking, testbed development and certification purposes.

The short term: We describe the status of standardisation in quantum benchmarks in detail in Case Study 1 below.

Case Study 1: Quantum Computer Performance Benchmark

Quantum computing is still an extremely nascent technology and, as mentioned above, there are currently nine primary hardware systems being developed for quantum computers. This distinction between hardware platforms is not like the old Mac versus Windows or Intel versus PowerPC battles that occurred in the late 20th Century. After all, by that time, digital computer systems were all built on the same underlying Silicon-based transistor technology. These nine quantum platforms, however, rely on completely different operating principles and building blocks, from particles of light moving continuously at nearly 300,000km/sec, to precisely placed, individual Phosphorus atoms in an otherwise perfect crystal of Silicon. And such dramatic differences necessitate equally great variations in design, fabrication, engineering, infrastructure and control techniques.

So how do we compare the capacity or performance of systems that are so fundamentally different from each other as a photonics-based quantum computer is from a Silicon-based machine? Is such a comparison useful? Does it even make sense to try and compare them?

The quantum computing community has been grappling with these questions ever since the first qubit demonstrations emerged, and this is where much of the informal standards development in quantum computing has taken place to date. For example, for devices at the processor scale, metrics used to compare evolution and performance across different quantum computing machines include:

- Qubit count: How many physical qubits do I have in my quantum computer?
- Qubit fidelity: How good are my qubits, how accurately can they be controlled?
- Qubit connectivity and controllability: Can I interact large groups of qubits together, and do we have sufficient control to perform individual operations of each qubit?
- Gate speed: How fast are my quantum gates between qubits?
- Gate count: How many gates can my quantum computer run before the performance deteriorates below acceptable levels?
- Circuit execution: If we run a small-scale quantum protocol on our device, does it work or not?

Such metrics have generally been sufficient for the academic space. Experiments on ion-trap quantum computers, superconductors or photonics, when assessed by experts, could generally be compared and contrasted accurately and without too many complications. However, the community has not yet identified metrics that can capture all critical aspects of the quantum computing development pathway. So creating benchmarking standards is not just a matter of agreeing on some consensus: these are often cutting-edge research questions in their own right.

With quantum computing now moved into the commercial space, competing players are obviously incentivised to promote their own products. Since there are no agreed formal benchmarking standards, this leaves researchers and companies open to propose or adopt whichever metrics show their results in the best light. All platforms still face major technical scientific and engineering barriers ahead of them, and what may be a key roadblock in one platform, may not even show up in another.

With multiple commercial quantum computing systems from different vendors now being deployed, both in the cloud and as stand-alone systems, there have been attempts to provide some type of comparative metric. The most common was introduced by IBM in 2018, called Quantum Volume (QV) (Cross, et al., 2019).

Quantum volume is defined approximately in the following way.

1. Take a random quantum circuit consisting of N – qubits and N – layers of gates.
2. Calculate, using a classical computer, samples of the output expected from this circuit.
3. Run the circuit on the quantum computer and compare its output distributions against your classical simulation.
4. If the output samples match “well enough” (according to a defined statistical measure), then increase the size of N.
5. Repeat steps 1-4 until the quantum computer does not match the classical simulations.
6. The largest value of N that passes the test defines the Quantum Volume $QV = 2^N$ exponent.

This metric does provide a somewhat system-agnostic measure of the capacity of a quantum computer and is currently used quite broadly. Quantum computing vendors often report the current largest quantum volume they claim their system achieves. However, there are still a range of issues with this metric:

- Quantum Volume doesn’t take into account the speed of the quantum computer: Two systems with a quantum volume of $QV = 2^{10} = 1{,}024$ may execute, successfully, the same 10-qubit quantum circuit, but one may do it in 100 microseconds and another may do it in 1 second.
- Quantum Volume is exponential in qubit number: That is, passing the QV test with one extra qubit, corresponds to a doubling of quantum volume. Since this reflects the exponentially increasing difficulty of the corresponding classical simulation, this scaling could potentially be useful for tracking the accessible complexity of NISQ era devices. But if scalable quantum computers orders of magnitude larger are ultimately required to realise applications of utility and value, then the exponential scaling will make QV less intuitive for tracking real progress, by exaggerating the impact of incremental change.
- Vendors often report Quantum Volume without any outside certification. In recent results from Los Alamos National Labs (Pelofske, et al., 2022), quantum volumes were independently measured through the cloud interface some vendors provide. In every machine studied, the Quantum Volumes they could reproduce were lower than quoted.
- Quantum Volume requires a classical calculation to verify against. As quantum computers get larger, classical machines will not be able to perform this calculation. Quantum Volume would potentially have to be “bootstrapped”, with the hope that smaller quantum computers can be used to verify QV for larger machines.
- Quantum Volume is not able to characterise the scalability of a particular processor. How useful is it to track QV as a computer scales towards 1000 qubits, if there are major roadblocks preventing that machine scaling further?
- Quantum Volume does not characterise the ability of a quantum computer to successfully execute an algorithm of value, without assuming that random quantum circuits are representative of more general algorithm performance. But passing the QV test for a particular size does not guarantee that all quantum circuits of that size can be successfully executed with low enough error.

In 2021, IBM introduced a second metric attempting to fix some of these issues, most notably by building in a way to characterise the speed of their system. This new metric was called Circuit Layer Operations per Second (CLOPS) (Wack, et al., 2021). CLOPS uses the same circuits used to calculate Quantum Volume, but it also quantifies how many of these circuits can be run per unit of time, aiming to account for both computational time and any pre-compilation time required to design, specify and submit a quantum circuit to the actual quantum computer. However, many of the objections to Quantum Volume as a metric remain within this new CLOPS definition. For example, benchmarking quantum computing progress using a metric that ignores considerations of scalability tacitly favours short-term improvements in small-scale performance even at the expense of design choices that impact scalability. Companies choosing to design platforms that aim to embed scalability from the start would probably argue that worrying about Quantum Volume is just ignoring the big elephant in the room.

The release of new metrics, however, can be strongly impacted by motivations related to competitive dynamics. Metrics designed by individual organisations are often developed and deployed because they benefit those organisations. Consider the two well established metrics above: IBM first introduced Quantum Volume arguably as a counter narrative to Google's efforts between 2014 and 2019, which aimed to define a decisive milestone target in relation to beyond-classical performance (also termed quantum supremacy). Similarly, the release of CLOPS was designed to highlight the speed of a given quantum circuit computation. Around 2019-2020, ion-trap quantum computers began to overtake IBM in Quantum Volume, due to the fact that arrays of trapped ions, while slower to scale up in size, can often achieve significantly lower errors than superconducting qubits (Quantinuum, n.d.) (Ion-trap quantum computers use identical ions trapped in virtually perfect vacuum, with fewer potential noise channels than the more flexible superconducting electronic circuits fabricated on solid-state chips.). But ion-trap gates generally operate some 4 orders of magnitude slower than superconducting qubit gates, meaning that superconducting chipsets can outperform ion-trap systems in CLOPS, even while they have lower Quantum Volume.

Examples such as this highlight the need for independent and more holistic benchmarking standards for NISQ-era quantum computers. The above arguments may be valid, but it does not necessarily benefit the community as a whole for standards to be driven so strongly by individual commercial interests, and metrics can become entrenched in the absence of other options, if adopted broadly enough. Independent standardisation of new performance metrics for quantum computing is now being discussed by bodies such as ISO/IEC and IEEE, with input from Standards Australia and other NSBs, broad consortia like QED-C, and significant targeted research efforts are being driven by some funding agencies. Given the intrinsic differences between the nine major hardware modalities for quantum computing, this is not a trivial task, and it takes time. But even if NISQ-era quantum computing does not reach the scale necessary to provide scientific or commercial benefit in the computational space, being able to categorise and compare the performance of radically disparate types of machines will be critical as the quantum ecosystem continues to expand.

Case Study 2: Quantum Assembly Language

There is arguably only one current example of any type of agreed standard in the quantum software space. The Open Quantum Assembly Language OpenQASM was released in 2017 (Cross, et al., 2017), also by researchers at IBM, and is a gate-based description of a quantum circuit. An OpenQASM instruction set is essentially an ordered list of individual quantum gate calls that are executed sequentially by the quantum computer. The OpenQASM format is sometimes referred to as an intermediate representation. The idea is that higher-level, abstract quantum algorithms are first written as an OpenQASM list of gates, before quantum firmware software then translates these OpenQASM commands into specific controls that execute the gates on the quantum hardware. In the last 5 years, OpenQASM has been rather broadly adopted as a defacto standard with most quantum programming languages and software packages having the ability to read and write to the OpenQASM format, and most hardware vendors allowing users to run quantum circuits specified in OpenQASM format on small-scale hardware.

This is a context where "informal" standards development spearheaded by a single organisation has worked very effectively. At the current stage of commercial growth, quantum computing vendors across the board can derive clear benefits, in terms of opening up new markets, by providing users with a way to interface activity between different providers. As a result, it now attracts significant active engagement across the community, through both direct contributors and opportunities to submit feedback. And while it is a very formally defined language specification, it is not developed under the auspices of a collaborative standards body, and has therefore been able to build momentum rather quickly. Furthermore, this is a tool designed to connect activities from different providers, not differentiate them. IBM gains reputation and recognition for maintaining this project, without it significantly impacting on their core IP and commercial interests.

Other loose collaborations and efforts to provide some type of standardisation framework for quantum software include: the Quantum Intermediate Representation (QIR) alliance, initiated by Microsoft (Microsoft, n.d.); Mitiq, a cross-platform quantum compilation system developed by the community-based non-profit, the Unitary fund (Mitiq, n.d.); and the Practical intermediate representation for quantum (PIRQ) project of the QED-C (QED-C, n.d.).

These case studies highlight two important lessons for standardisation development in the quantum community moving forwards:

1. For better or worse, standards developed by individual organisations, or focussed consortia, can be more agile and adaptive, and move more rapidly to create de facto standards that are rather broadly adopted.
2. Standardisation processes will realise the most impact focussing on two scenario types:
   a. scenarios (like OpenQASM) with broad benefits for the commercial ecosystem but which do not impinge on areas of IP significance and commercial differentiation;
   b. and scenarios (like benchmarking) where there are compelling "public-interest" benefits for the wider community (e.g., for users and investors) for having rigorous, expert-driven specifications to facilitate objective market comparisons and quality assurance.

Standardisation at different scales

In the previous section, we have outlined some opportunities for standardisation to positively impact both the quantum computing industry and the wider community, and discussed relevant short-term considerations, each within 5 different spheres of engagement. Not all of these opportunities, however, will be relevant in different quantum computing development pathways, depending on where the technology ends up landing in terms of scale and complexity. In this section, we do not ask how (or how fast) quantum standardisation might develop moving forward, but instead skip forward to the quantum computing end game, and consider how the needs for standardisation might look under each of the 5 scale scenarios outlined in Section 3 above. These discussions are again then informed by the 5 spheres of engagement we identified for standardisation in quantum.

1. **The LHC scenario:** Collaborations at the scale of the Large Hadron Collider, or the Apollo moon mission, must necessarily make extensive use of standardisation. At the community level, transparent standardisation processes would assist in demonstrating the scrupulous governance required to back up that level of public investment. An LHC-like Quantum Computer Collaboration would require detailed *System Architectures* to be developed constructed at the hardware and software levels. Hardware and software systems would likely be highly modular, requiring comprehensive *Requirements Specifications* for performance benchmarking and interfacing. Hybrid QC platforms could be readily adopted in this scenario. Performance specifications would be especially critical to ensure compatibility with system components being provided by industry contributors. The infrastructure would be developed, maintained and operated by a combination of research-level academics and advanced technicians and engineers. Consequently, this would relax some constraints around software- and control-level abstractions, driven by demand and convenience to the expert operators. Access to such a facility would likely be highly competitive and controlled, and computational projects would be run by teams that would almost certainly include dedicated quantum expertise. Even at this, more research-driven scale, the QC Collaboration would involve experts from widely disparate domains, and would establish standard terminologies as required.

   However, while the QC Collaboration would almost certainly include industry partners, as a technology development project, it would likely have branched away from quantum computing's current highly commercially driven and fast-paced environment. Standards, while created extensively, would be managed through internal Collaboration processes, separate from the context of industry-focussed SDOs. Famous examples of standards and frameworks developed out of such efforts include the world wide web, invented at CERN, and the Technology Readiness Level framework developed by NASA to accurately track the maturity of new technologies being developed for inclusion in massively complex collaborative projects like space missions. And like the Apollo moon mission project, a QC Collaboration could be expected to spin off many advanced, commercially valuable, enabling technologies, driven by the innovation required to tackle such a large scale endeavour. At this project scale, quantum standardisation might not look the same as formal standards development, but it would still be crucial.

2. **The corporate provider/data centre scenario:** This scale would see the entire hardware and control structure housed internally within a single corporate entity (or less likely, an academic consortium). By retaining control over the entire hardware stack, this would reduce the need for formal standardisation at the hardware level. The provider would undoubtedly develop detailed internal hardware and software architectures and frameworks, but these would likely form part of the provider's competitive intellectual property portfolio, and the need for formal standards could be assessed case by case. *Requirements Specifications* could be developed, as required, to maintain supply chain compatibility. Providers may favour maximum achievable QC platform uniformity over hybrid approaches. A robust software and control framework maximising abstraction would be crucial to ensure accessibility

and usability by external customers. Particular emphasis would need to be paid to the development of detailed Software standards describing APIs and interfaces, and ideally also transpilers to allow the system to run code developed on competitor offerings.

With time on a QC data centre likely being extremely expensive, computing time would likely be purchased through rigorous procurement processes. In this environment of disparate commercial QC data centres, formal *Terminology standards* and *Technical Reports* would provide valuable context for interested customers. Customers would demand sufficient programming compatibility between systems to ensure that one-off purchases would not constrain future purchasing decisions. Benchmarking via formal, normative system-level Specifications and *Measurement and Test Methods*, and even full Quality Assurance standards, would provide significant commercial benefits to interested customers. These outward customer-focussed considerations would provide the main motivation for QC data centre providers' engagement with standardisation. Similar requirements would be necessary to ensure regulatory approval and investor confidence.

3. **The corporate user/HPC cluster scenario:** In this scenario, while individual providers may manufacture the entire product chain, the need for customer-side operation and maintenance of the quantum HPC (QHPC) would necessitate a stronger focus on conformity with formally recognised standards. These requirements may be similar to the QC data centre scenario from a program-/software-level user perspective, including in relation to benchmarking and quality assurance. Indeed, the inevitable movement of QHPC users between different QHPCs would create significant demand for a high degree of software user interface compatibility between different QHPC providers. But the requirements would likely be much more strict from a hardware- and control-level perspective. There may also be customer demand for interfacing compatibility between systems even from different providers. The need for experienced customer-side operations technicians and maintainers increase the need for broader workforce development.
4. **The home desktop/QPU scenario:** The inclusion and exploitation of quantum processor units (QPUs) in a home desktop device would not only require a step change in qubit cost and environmental constraints, but would also require a substantial step up in formal standards development and conformity. In addition to all the standardisation requirements discussed for scenarios 2 and 3 above, comprehensive Quality Assurance standards would likely be required for regulatory approval and consumer confidence in seeing the technology proliferate. QPUs would almost certainly require uniform (not hybrid) hardware platforms, and exploitation of QPU capabilities by commercial app developers would require further increases in user interface and software abstractions, and APIs, etc. Each increase of this sort would increase the role of standardisation. Yet it is perhaps questionable whether the degree of sufficient abstraction and knowledge level of commercial app developers at this level could ever be harmonised to the degree where QPUs in a home desktop could be adequately exploited to add value.
5. **The smartphone/tablet scenario:** In this final, most portable scenario, the role of and reliance on standardisation would likely be similar to the desktop/QPU scenario above, but with still more stringent requirements. Hardware platforms would not only need to be uniform, but likely also monolithic, and certainly fully miniaturised. In addition to almost unimaginably cheap qubits, this would also require a paradigm shift in hardware platform design, or fabrication processing technology, or both.

The table below summarises some of the requirements described above grouped into hardware and software domains.

Fig. 10: Quantum computing standardisation requirements at different deployment scales

Scale	Software standardisation	Hardware standardisation
LHC	• Extensive standardisation • Internally developed (not SDOs) • Open standards • Expert operators	• Extensive standardisation • Internally developed • Open standards • Expert operators • Hybrid designs possible
Corporate provider/Data centre	• High standardisation • Internal/external dev • Some open standards • Semi-expert operators • External compatibility	• Moderate standardisation • Internally developed • No operators • Platform uniformity favoured
Corporate user/HPC cluster	• High standardisation • Internal/external dev • Some open standards • Semi-expert operators • External compatibility	• High standardisation • Internal/external dev • Some open standards • Semi-expert operators • Interfaceable • Platform uniformity favoured
Home desktop	• Extensive standardisation • Detailed, open standards • Maximum abstraction, UX • Lay operators • Cross compatibility	• Extensive standardisation • Open standards • No operators can be required • Platform uniformity necessary • Interfaceable
Personal (portable) device	• Extensive standardisation • Detailed, open standards • Maximum abstraction, UX • Lay operators	• Extensive standardisation • Open standards • No operators can be required • Monolithic necessary • Miniaturised

Bibliography

Abdelkafi, N. et al., 2018. Understanding ICT Standardization: Principles and Practice. France: European Telecommunications Standards Institute.

Aggarwal, D. et al., 2018. Quantum Attacks on Bitcoin, and How to Protect Against Them. Ledger, Volume 3.

Arute, F., Arya, K., Babbush, R. & et al., 2019. Quantum supremacy using a programmable superconducting processor. Nature, Volume 574, pp. 505-510.

AWS Roadmap, n.d. Designing a fault-tolerant quantum computer based on Schrödinger-cat qubits. [Online]

Available at: https://aws.amazon.com/blogs/quantum-computing/designing-a-fault-tolerant-quantum-computer-with-cat-qubits/

Babbush, R., Gidney, C., Berry, D. & et al., 2018. Encoding electronic spectra in quantum circuits with linear T complexity. Phys. Rev. X., Volume 8, p. 041015.

Ball, H., Biercuk, M. & Hush, M., 2021. Quantum firmware and the quantum computing stack. Physics Today, Volume 74, p. 28.

Barends, R., Kelly, J., Megrant, A. & et al., 2014. Superconducting quantum circuits at the surface code threshold for fault tolerance. Nature, Volume 508, pp. 500-503.

Beauregard, S., 2003. Circuit for Shor's algorithm using 2n+3 qubits. Quantum Information and Computation, Volume 3, pp. 175-185.

Bloch, F. & Purcell, E., 1952. The Nobel Prize in Physics 1952. [Online]

Available at: https://www.nobelprize.org/prizes/physics/1952/summary/

Boixo, S., Isakov, S., Smelyanskiy, V. & et al., 2018. Characterizing quantum supremacy in near-term devices. Nature Physics, Volume 14, p. 595.

Brandhofer, S., Devitt, S., Wellens, T. & Polian, I., 2021. Noisy intermediate-scale quantum (NISQ) computers—how they work, how they fail, how to test them? 2021 IEEE 39th VLSI Test Symposium (VTS), San Diego, CA, USA, pp. 1-10.

Brassard, G., Hoyer, P. & Tapp, A., 1998. Quantum Counting. arXiv:quant-ph/9805082.

Braunstein, S. & Lloyd, S., 1999. Quantum Computation over Continuous Variables. Phys. Rev. Lett., Volume 82, pp. 1784-1787.

Bremner, M., Josza, R. & Shepherd, D., 2010. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. Proc. R. Soc. A, Volume 467, pp. 459-472.

Bremner, M., Montanaro, A. & Shepherd, D., 2016. Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations. Phys. Rev. Lett., Volume 117, p. 080501.

Brennen, G., Caves, C., Jessen, P. & Deutsch, I., 1999. Quantum Logic Gates in Optical Lattices. Phys. Rev. Lett., Volume 82, p. 1060.

Cai, Z., Babbush, R., Benjamin, S. & et al., 2022. Quantum Error Mitigation. arXiv:2210.00921.

CEN-CENELEC, n.d. Quantum Technologies. [Online]

Available at: https://www.cencenelec.eu/areas-of-work/cen-cenelec-topics/quantum-technologies/

Chuang, I. & Nielsen, M., 1997. Prescription for experimental determination of the dynamics of a quantum black box. Journal of Modern Optics, Volume 44, p. 2455.

Cross, A., Bishop, L., Sheldon, S. & et al., 2019. Validating quantum computers using randomized model circuits. Phys. Rev. A., Volume 100, p. 032328.

Cross, A., Bishop, L., Smolin, J. & Gambetta, J., 2017. Open Quantum Assembly Language. arXiv:1707.03429.

DARPA, n.d. Quantum Benchmarking. [Online]

Available at: https://www.darpa.mil/program/quantum-benchmarking

Devitt, S., Stephens, A., Munro, W. & Nemoto, K., 2013. Requirements for fault-tolerant factoring on an atom-optics quantum computer. Nature communications, Volume 4, p. 2524.

DiVincenzo, D., 1997. Topics in Quantum Computers. Mesoscopic Electron Transport, NATO, ASI., Volume 257.

Dowling, J. & Milburn, G., 2002. Quantum Technology: The Second Quantum Revolution. arXiv:quant-ph/0206091.

D-Wave Roadmap, n.d. The D-Wave Clarity Roadmap. [Online]

Available at: https://www.dwavesys.com/media/xvjpraig/clarity-roadmap_digital_v2.pdf

Feynman, R., 1982. Simulating physics with computers. Int J. Theor. Phys., Volume 21, pp. 467-488.

Gaebler, J. et al., 2012. Randomized Benchmarking of Multiqubit Gates. Physical Review Letters, Volume 108, p. 260503.

Gershenfeld, N. & Chuang, I., 1997. Bulk Spin-Resonance Quantum Computation. Science, pp. 275:350-356.

Gidney, C. & Ekera, M., 2021. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. Quantum, Volume 5, p. 433.

Google Roadmap, n.d. Our quantum computing journey. [Online]

Available at: https://quantumai.google/learn/map

Grover, L., 1996. A fast quantum mechanical algorithm for database search. Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, pp. 212-219.

Haroche, S. & Wineland, D., 2012. The Nobel Prize in Physics 2012. [Online]

Available at: https://www.nobelprize.org/prizes/physics/2012/summary/

Harrow, A., Hassidim, A. & Lloyd, S., 2008. Quantum algorithm for solving linear systems of equations. Phys. Rev. Lett., Volume 103, p. 150502.

Heim, B., 2020. Quantum programming languages. Nature physics, Volume 2, p. 709–722.

Hofmann, H., 2005. Complementary Classical Fidelities as an Efficient Criterion for the Evaluation of Experimentally Realized Quantum Operations. Physical Review Letters, Volume 94, p. 160504.

IBM Roadmap, n.d. The IBM Quantum Development Roadmap. [Online]

Available at: https://www.ibm.com/quantum/roadmap

IBM, n.d. Learning Quantum Computation using Qiskit. [Online]

Available at: https://qiskit.org/textbook/preface.html

IEC Market Strategy Board, 2021. White Paper on Quantum information technology. [Online]

Available at: https://www.iec.ch/basecamp/quantum-information-technology

IEEE, n.d. IEEE Quantum Initiative Support for Standards. [Online]

Available at: https://quantum.ieee.org/standards

IEEE, n.d. Standard for Quantum Computing Architecture. [Online]

Available at: https://standards.ieee.org/ieee/3120/10751/

IEEE, n.d. Standard for Quantum Computing Performance Metrics & Performance Benchmarking. [Online]

Available at: https://standards.ieee.org/ieee/7131/10681/

ISO, n.d. Deliverables. [Online]

Available at: https://www.iso.org/deliverables-all.html

Jaksch, D. et al., 2000. Fast Quantum Gates for Neutral Atoms. Phys. Rev. Lett., Volume 85, p. 2208.

James, D., Kwiat, P., Munro, W. & White, A., 2001. Measurement of qubits. Physical Review A, Volume 64, p. 052312.

Kane, B., 1998. A silicon-based nuclear spin quantum computer.. Nature, Volume 393, pp. 133-137.

Kitaev, A., 2003. Fault-tolerant quantum computation by anyons. Annals of Physics, Volume 303, pp. 2-30.

Knill, E., Laflamme, R. & Milburn, G., 2001. A scheme for efficient quantum computation with linear optics. Nature, Volume 409, pp. 46-52.

Knill, E., Leibfried, D., Reichle, R. & et al., 2008. Randomized benchmarking of quantum gates. Physical Review A, Volume 77, p. 012307.

Lloyd, S., 1996. Universal Quantum Simulators. Science, pp. 273:1073-1078.

Loss, D. & DiVincenzo, D., 1998. Quantum computation with quantum dots. Phys. Rev. A., Volume 57, p. 120.

Lubinski, T. et al., 2021. Application-Oriented Performance Benchmarks for Quantum Computing. arXiv:2110.03137.

Magesan, E., Gambetta, J. & Emerson, J., 2011. Scalable and Robust Randomized Benchmarking of Quantum Processes. Physical Review Letters, Volume 106, p. 180504.

Magesan, E., Gambetta, J. & Emerson, J., 2012. Characterizing quantum gates via randomized benchmarking. Physical Review A, Volume 85, p. 042311.

Microsoft, n.d. Introducing Quantum Intermediate Representation (QIR). [Online]

Available at: https://devblogs.microsoft.com/qsharp/introducing-quantum-intermediate-representation-qir/

Mitiq, n.d. [Online]

Available at: https://unitary.fund/mitiq.html

Monroe, C. et al., 1995. Demonstration of a Fundamental Quantum Logic Gate. Phys. Rev. Lett, Volume 75, p. 4714.

Montanaro, A., 2016. Quantum algorithms: an overview. npj Quantum Information, Volume 2, p. 15023.

Monz, T., Nigg, D., Martinez, E. & et al., 2016. Realization of a scalable Shor algorithm. Science, Volume 351, pp. 1068-1070.

Nakamura, Y., Pashkin, Y. A. & Tsai, J.-S., 1999. Coherent control of macroscopic quantum states in a single-Cooper-pair box. Nature, Volume 398, pp. 786-788.

O'Brien, J. et al., 2004. Quantum Process Tomography of a Controlled-NOT Gate. Physical Review Letters, Volume 93, p. 080502.

Pelofske, E., Bärtschi, A. & Eidenbenz, S., 2022. Quantum Volume in Practice: What Users Can Expect from NISQ Devices. arXiv:2203.038164.

Peral Garcia, D., Cruz-Benito, J. & Jose Garcia-Penalvo, F., 2022. Systematic Literature Review: Quantum Machine Learning and its applications. arXiv:2201.04093.

Poyatos, J., Cirac, J. & Zoller, P., 1997. Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate. Physical Review Letters, Volume 78, p. 390.

Preskill, J., 2012. Quantum computing and the entanglement frontier. arXiv:1203.5813.

Preskill, J., 2018. Quantum Computing in the NISQ era and beyond. Quantum, Volume 2, p. 79.

QED-C, n.d. Practical intermediate representation for quantum (PIRQ): Requirements and near-term recommendations. [Online]

Available at: https://quantumconsortium.org/pirq22/

Quantinuum, n.d. Quantinuum Sets New Record with Highest Ever Quantum Volume. [Online]

Available at: https://www.quantinuum.com/news/quantinuum-sets-new-record-with-highest-ever-quantum-volume

Qubit x Qubit, n.d. Training the Future Diverse Quantum Workforce. [Online]

Available at: https://www.qubitbyqubit.org/

Rabi, I., 1944. The Nobel Prize in Physics 1944. [Online]

Available at: https://www.nobelprize.org/prizes/physics/1944/summary/

Reiher, M., Wiebe, N., Svore, K. & et al., 2017. Elucidating reaction mechanisms on quantum computers. Proceedings of the National Academy of Sciences, Volume 114, pp. 7555-7560.

Rigetti Roadmap, n.d. Updated Strategic Plan. [Online]

Available at: https://investors.rigetti.com/news-releases/news-release-details/rigetti-computing-announces-updated-strategic-plan-prioritize

Schockley, W., Bardeen, J. & Brattain, W., 1956. The Nobel Prize in Physics 1956. [Online]

Available at: https://www.nobelprize.org/prizes/physics/1956/summary/

Shahriar, M. et al., 2002. Solid-state quantum computing using spectral holes. Phys. Rev. A., Volume 66, p. 032301.

Shor, P., 1994. Algorithms for quantum computation: discrete logarithms and factoring. Proceedings 35th Annual Symposium on Foundations of Computer Science. IEEE Comput. Soc. Press, pp. 124-134.

Shor, P., 1995. Scheme for reducing decoherence in quantum computer memory. Physical review A., p. 52: R2493–R2496.

Shor, P., 1996. Fault-tolerant quantum computation. Proceedings of 37th Conference on Foundations of Computer Science. Burlington, VT, USA: IEEE Comput. Soc. Press, pp. 56-65.

Standards Australia, n.d. Standardisation Guide 003: Standards and Other Publications. [Online]

Available at: https://www.standards.org.au/getmedia/d9da035d-2fbc-4417-98c1-aa9e85ef625d/SG-003-Standards-and-Other-Publications.pdf.aspx

Sydney Quantum Academy, n.d. [Online]

Available at: https://sydneyquantum.org/

Townes, C., Basov, N. & Prokhorov, A., 1964. The Nobel Prize in Physics 1964. [Online]

Available at: https://www.nobelprize.org/prizes/physics/1964/summary/

US Department of Energy, n.d. Exploratory Research for Extreme-Scale Science (EXPRESS) - Quantum Algorithms across Models. [Online]

Available at: https://science.osti.gov/-/media/grants/pdf/foas/2023/SC_FOA_0002950.pdf#page9

van Deventer, O., Spethmann, N., Loeffler, M. & et al., 2022. Towards European standards for quantum technologies. EPJ Quantum Technology, Volume 9, p. 33.

Venegas-Andraca, S., 2012. Quantum walks: a comprehensive review. Quantum Information Processing, Volume 11, pp. 1015-1106.

Wack, A., Paik, A., Javadi-Abhari, A. & et al., 2021. Scale, Quality, and Speed: three key attributes to measure the performance of near-term quantum computers. arXiv:2110.14108.

Wang, Y., Crain, S., Fang, C. & et al., 2020. High-Fidelity Two-Qubit Gates Using a Microelectromechanical-System-Based Beam Steering System for Individual Qubit Addressing. Phys. Rev. Lett., Volume 125, p. 150505.

Webber, M., Elfving, V., Weidt, S. & Hensinger, W., 2022. The impact of hardware specifications on reaching quantum advantage in the fault tolerant regime. AVS Quantum Sci, Volume 4, p. 013801.

World Economic Forum, 2022. State of Quantum Computing: Building a Quantum Economy. [Online]

Available at: https://www.weforum.org/reports/state-of-quantum-computing-building-a-quantum-economy/

Wu, Y., Bao, W., Cao, S. & et al., 2021. Strong Quantum Computational Advantage Using a Superconducting Quantum Processor. Phys. Rev. Lett., Volume 127, p. 180501.

Zoller, P. & Cirac, J., 1995. Quantum Computations with Cold Trapped Ions. Phys. Rev. Lett., Volume 74, p. 4091.