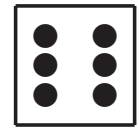

Quantum Computing Today

At the 25th Solvay Conference on Physics in 2011, John Preskill asked a question about quantum computing for which we still have no answer:

> Is controlling large-scale quantum systems merely **really, really hard**, or is it **ridiculously hard**?[1]

Preskill, who is the Richard P. Feynman Professor of Theoretical Physics at the California Institute of Technology, was asking if building ever larger quantum computers of the kind we envisioned in the last chapter is merely a matter of better engineering, or if there are fundamental limits about the nature of physics, computation, and reality itself that will get in the way. That is, are we likely to have working quantum computers "going beyond what can be achieved with ordinary digital computers" – what Preskill called "quantum supremacy" – after "a few decades of very hard work"? Or are we likely to come up short after even centuries of effort?

Preskill didn't have an answer, but he was enthusiastic about the quest: even if efforts to build a working large-scale quantum computer failed, humanity would still learn important fundamental truths about the fundamental nature of the universe.

[1] Preskill, "Quantum Computing and The Entanglement Frontier" (2012), emphasis in the original.

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

In the last chapter we discussed the first three great applications that have been envisioned for quantum computers: simulating quantum mechanical systems (Feynman), factoring large numbers (Shor), and speeding the search for solutions to any mathematical problem for which it is possible to construct a quantum oracle (Grover). All of these applications were developed by theoreticians working with nothing more than the metaphorical pencil and paper, and the ability to discuss ideas with their collaborators. Actually realizing these applications requires something more: a large-scale, reliable quantum computer.

Companies and research labs are racing to answer Preskill's question. Some are large, established technology powerhouses, like Google, IBM, and Microsoft. Others are well-funded emerging players, such as ColdQuanta, D-Wave and Rigetti. Most are building actual physics packages, with super-cooled superconductors and parts that are literally gold-plated. In most but not all cases, the results of these quantum computers can be reliably simulated using clusters of conventional computers. However, in a few cases, machines have been constructed that can solve problems beyond the capacity of today's digital computers – even when millions of those computers are networked together.

"I proposed the term 'quantum supremacy' to describe the point where quantum computers can do things that classical computers can't, regardless of whether those tasks are useful," Preskill wrote in 2019.[2] "With that new term, I wanted to emphasize that this is a privileged time in the history of our planet, when information technologies based on principles of quantum physics are ascendant."

After gaining traction, Preskill's term *quantum supremacy* has been somewhat supplanted by the term *quantum advantage*. Some researchers prefer this term, because it rightfully implies that quantum computers will be working alongside classical computers to literally confer advantage, just as a modern computer might offload some computations to a graphics processing unit (GPU).

Quantum computers have not scaled up at the same rate as their electronic computing predecessors. We have yet to experience a quantum form of Moore's Law (see Section 3.5, p. 98), in part because quantum engineers have not found a suitable quantum mechanism equivalent to the digital discipline that allows creating ever-larger

[2]Preskill, "Why I Called It 'Quantum Supremacy'" (2019).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

digital circuits without ever-increasing amounts of systemic error (see Section 3.3 (p. 84)). Although quantum error correction schemes exist, it is unclear if they can scale to allow for meaningfully complex computations, because these schemes themselves require higher quality qubits operational for longer timescales than are currently possible. Without resolving this issue, we will still likely be able to create analog *quantum simulators* for solving questions in physics, chemistry, and biology, but the goal of using quantum computers to crack codes may remain forever out of reach. Nevertheless, researchers at both Google and the University of Science and Technology of China created quantum computing systems that clearly meet Preskill's requirement for quantum supremacy.

In this first section of this chapter we will describe in abstract the basics of how the current generation of quantum computers work. Next, in Section 6.2.2 (p. 237) we discuss the hardware efforts of today and the near future. We discuss what will need to be overcome in Section 6.3 (p. 242). Finally we conclude this chapter with Section 6.4 (p. 253).

6.1 How to Build a Quantum Computer

In Chapter 4 we introduced the basic idea of the Fredkin and Toffoli gates, and in Chapter 5 we discussed the two quantum algorithms that started serious money flowing into the creation of actual quantum computers. In this chapter we'll briefly look at a simple quantum circuit and discuss the barriers to creating quantum circuits of the size necessary to accomplish the computational goals set out in the previous chapter.

In a now classic article, David P. DiVincenzo, then at the IBM T.J. Watson Research Center, formulated five requirements for quantum computing:[3]

1. There needed to be something that could "hold data and perform computation." For simplicity, scientists have focused systems that have two precise states, which we call qubits. Whereas a classical bit can only have two values, `0` and `1`, quantum bits are a superposition of these two states. This superposition is typically written using Paul Dirac's Bra-ket notation as $a|0\rangle + b|1\rangle$, where a and b are taken to be complex numbers such that $|a|^2 + |b|^2 = 1$ during the course of the computation,

[3] DiVincenzo, "Topics in Quantum Computers" (1997).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

but which become either `0` or `1` when they are measured at the end of the computation.[4] This measurement corresponds to "opening the box" in Schrödinger's famous thought experiment (see p. 523).[5]

2. The ability to initialize the qubits to a known "fiducial starting quantum state." This requirement is akin to resetting all of the bits in a classical computer to `0`. In his 1997 article, DiVincenzo wrote "I do not think that this 'initial state preparation' requirement will be the most difficult one to achieve for quantum computation." Three years later in his follow-up article, DiVincenzo was less sanguine: "The problem of continuous initialization does not have to be solved very soon; still, experimentalists should be aware that the speed with which a qubit can be zeroed will eventually be a very important issue."[6]

3. The ability to interact with each other using some form of *quantum gate*. This is where the Feynman and Toffoli gates from Section 4.5 (p. 151) become relevant. Each gate mixes the quantum state of two, three or more qubits together to perform some sort of simple computation. The physical construction of the quantum computer determines which qubits can be connected together. Ideally, the quantum gates are *universal*, so that they can be used to describe any computation (provided that you have sufficient qubits and time).

   As we will see in Chapter 3, this design makes the construction and programming of quantum computers fundamentally different from the way we have built classical computers. In classical computers the bits represented by the presence or absence of an electric charge move through the electronic circuits, which are fixed at the time the computer is manufactured. In a quantum computer, it is the qubits that are fixed when the computer is manufactured, and the system is programmed by playing a sequence of circuits through the qubits to perform

[4] With two qubits, the systems state is described by a four-dimensional vector: $a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle$.

[5] Qubits must be physically isolated from the universe such that there is no external energy that would bias the qubit towards being `0` or `1` on measurement. This is why qubits do not need to be isolated from gravity: both the $|0\rangle$ and the $|1\rangle$ states have the same mass.

[6] DiVincenzo, "The Physical Implementation of Quantum Computation" (2000).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

the desired computation. Thus, the computing speed of the quantum computer fundamentally depends on the number of qubits that it has and the speed at which the circuits can be constructed; this speed is exactly analogous to the clock speed of a modern microprocessor.[7]

4. The ability to keep the qubits in their *coherent*, *entangled* state for an extended period of time. This period of time is not measured in seconds, but in terms of how many gates can be played through the qubits. In his article, DiVincenzo suggested that it would be necessary to execute between a thousand and ten thousand gates in order to be able to perform meaningful computations with sufficient quantum error correction.[8]

   An added complication is how error propagates as the quantum computer begins to lose its coherency: if errors are correlated rather than randomly scattered through the system, it may adversely impact the ability to perform meaningful quantum error correction.

5. The ability to measure each qubit at the end of the computation.

We show what this looks like in Figures 6.1 through 6.3. This adder, which would be a small part of a much larger quantum circuit, takes two numbers between `0` and `15` and adds them together. The key difference between this adder and the 4-bit adder that you might find in a classical computer (such as Figure 3.5) is that this adder is reversible. The adder in Figure 6.3 uses 13 qubits and requires 30 gates. The design in Figure 6.3 also requires 30 cycles to operate because none of the gates execute at the same time. However, this algorithm can be optimized (Figure 6.4) by having many of the gates acting simultaneously. This optimized algorithm can run in just 7 cycles.

By *reversible*, we mean that this adder needs to be able to run in reverse. That is, it needs to be able to take the result of the addition, a single number between `0` and `15`, and provide the two specific input numbers that were used to create it. This may seem like a magic trick! If we told you that the number `9` is the sum of

[7] In his 1997 and 2000 articles, the requirement of "a 'universal' set of quantum gates" is presented as the fourth DiVincenzo criterion.

[8] Long decoherence time was originally presented as the third DiVincenzo criterion.

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

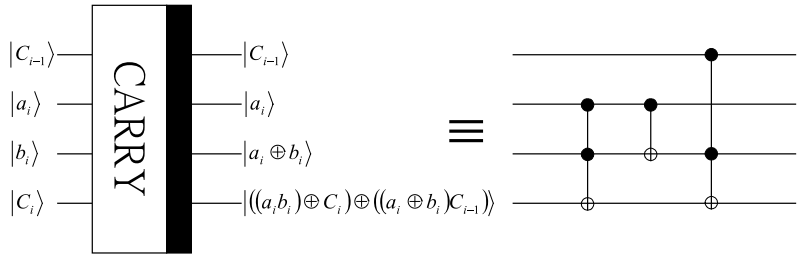

Figure 6.1. A 2-bit quantum carry gate, from Cheng and Tseng, "Quantum Plain and Carry Look-Ahead Adders" (2002), used with permission. The gate reversibly determines whether adding two bits produces a carry operation.

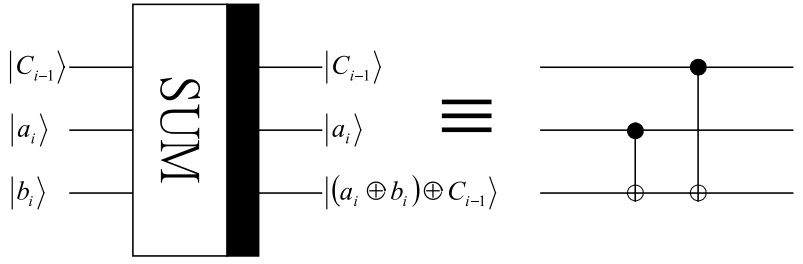

Figure 6.2. A 2-bit quantum sum gate, from Cheng and Tseng, "Quantum Plain and Carry Look-Ahead Adders" (2002), used with permission. The gate reversibly determines whether adding two bits produces a sum.

two numbers and asked you what they were, you would be unable to tell us: the answer might be `0` and `9`, or `1` and `8`, or `2` and `7`, and so on. As a result, the quantum 4-bit adder needs more than 4 bits of output: besides the 4-bit sum, it also preserves half of the input bits. The adder also has an additional input bit called z and an output bit that combines z with the carry bit. Such additional qubits are sometimes called an *ancillary* or *ancilla qubits*; designing efficient quantum circuits that use a minimum number of ancilla qubits is one of the current challenges of quantum computer programming, due to the small number of qubits and the short decoherence times. Programming quantum computers at the circuit level in this manner is exactly analogous to the way that computing's pioneers in the 1940s and 1950s modified the hardware of their computers to add new instructions and programmed the machines using machine code.

In summary, in order to compute at the quantum level, one must

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

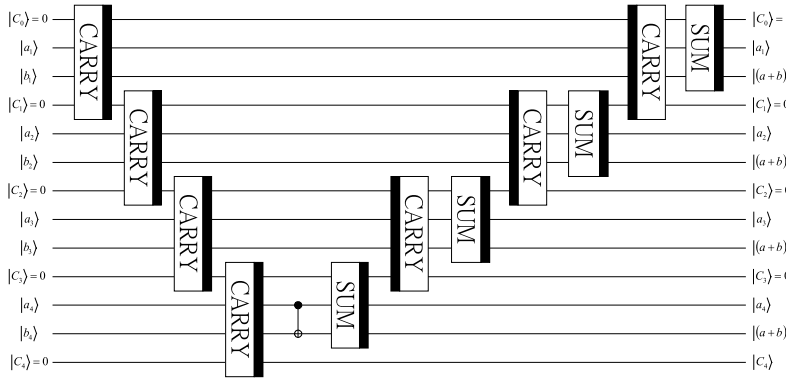

Figure 6.3. A 4-bit quantum adder circuit, from Cheng and Tseng, "Quantum Plain and Carry Look-Ahead Adders" (2002), used with permission. The inputs on the left are the nibbles $a_4a_3a_2a_1$ and $b_4b_3b_2b_1$ and the carry bit C_0. The output bits on the right are the sum $(a+b)_4(a+b)_3(a+b)_2(a+b)_1$, the input value $a_4a_3a_2a_1$, and the carry bit C_4.Time flows from left to right. Compare this with Figure 3.5, the 4-bit classical adder.

be able to generate, maintain, manipulate, and measure quantum states. Thus, quantum sensors are a precursor technology for quantum computing, and this is why this book presented quantum sensing first. In many ways, today's quantum computers are really just large-scale quantum sensor arrays.

6.2 The Quantum Computer Landscape

Preskill's 2019 article argues that the question he posed in 2012 is all but answered, and that we have moved from the era of quantum computing's first steps and into the era of noisy intermediate-scale quantum devices – NISQ – another term that he coined.

Unlike classical computers, which are nearly all based on silicon semiconductors, today's NISQ computers are not dominated by a single physical substrate. Instead, we are in a period of experimentation – one that might stretch out for decades. Today's quantum innovators are experimenting with different approaches to creating and managing the quantum states necessary for computation. To date, no one has realized the scale required for solving meaningful problems outside the world of experimental physics. The different media are promising in different ways, with some offering longer coherence times and greater interconnection, while others lack the need for specialized cooling or have engineering characteristics that might

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

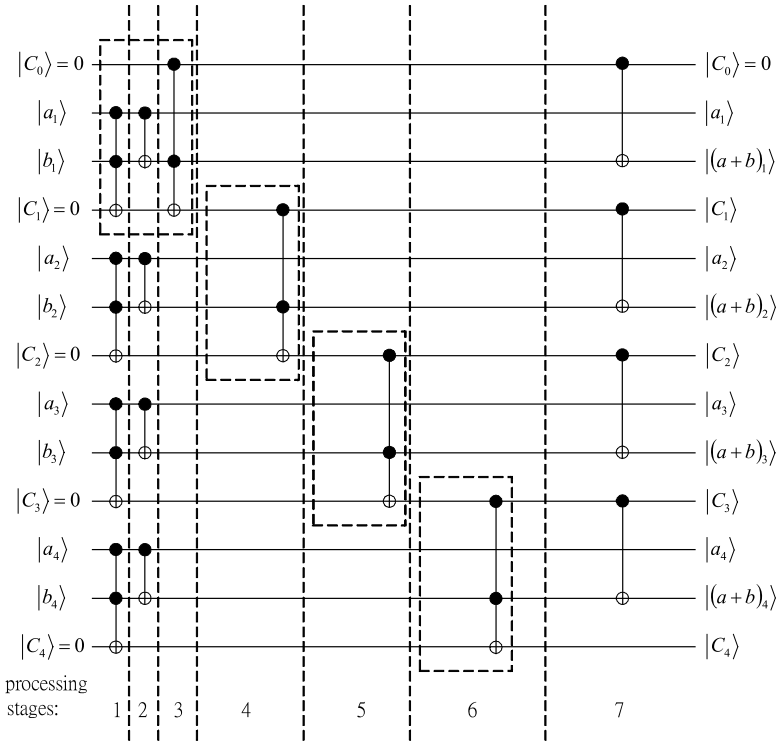

Figure 6.4. The 4-bit quantum adder from Figure 6.3, optimized to execute in fewer cycles. From Cheng and Tseng, "Quantum Plain and Carry Look-Ahead Adders" (2002), used with permission.

make a large-scale computer possible. We don't know which will be the winner.

6.2.1 Comparing Quantum Media

Understanding the quantum computing landscape is challenging because virtually every device that's been produced has different characteristics and capabilities. Some competitors claim to have relatively large-scale qubit devices, yet these may not be as interconnected as smaller devices, and large devices' size and architecture may be noisier and less stable than smaller devices. One cannot evaluate today's quantum computers simply by comparing the number of qubits they possess.

Adding to the difficulty, companies claims' on quantum computers may be strategically shaped to capture para-hardware markets,

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

such as software and services. Companies have created vocabularies and software frameworks that are explicitly helpful to them and their business model. Even when claimed to be neutral and universal, these vocabularies and frameworks cannot help but seek to establish a software ecosystem that is favorable to their creators.

Competitors in the field all seek the *logical qubit*, a qubit that can overcome the problems of gate errors, environmental noise, and decoherence long enough to perform quantum operations. Understandably, competitors have chosen different paths for the construction of a stable quantum computer. The paths chosen reflect a deeper design approach philosophy where some innovators are focused on small devices with high levels of interconnectivity and stability, while others are focused on building the largest device possible. The philosophy of the large devices is that with many *physical qubits*, the device can manage its own error.[9]

We've seen this behavior before repeatedly over the 70-year history of computing. Computer engineers in the 1950s experimented with a variety of computing and storage media before settling on silicon for switching, core memory for short-term storage, and a combination of hard drives, magnetic tape and punch cards for long-term storage. Similar technology competitions and selections took place in the world of high-performance supercomputers in the 1970s and 1980s. This fight played out once again during the emergence of cloud computing in the 2000s, with the surprising (to some) discovery that vast computing clouds built from commodity hardware could outperform specialized high-performance systems on a wide variety of tasks, once companies like Amazon and Google developed approaches for overcoming the challenges with scale.

6.2.2 Five Kinds of Quantum Computers

The word "quantum" is attached to a range of devices, and terminology in the field sometimes takes a functional approach. That is, the category of the device is cast by its use rather than its underlying architecture and capabilities. The lines between different categories of quantum computers blur. When it comes to computing, the word *quantum* can describe:

[9] Doug Finke, the publisher of the Quantum Computing Report, maintains the most comprehensive and up-to-date summary and categorization of hardware and software approaches by competitors. Finke's site carefully tracks claims of device size, quality, and construction (Finke, "Qubit Count" (2021)).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

- **Simulations of quantum computers.** On the most basic level, classical computers can be optimized to simulate quantum effects. The fundamental problem with using classical computers to simulate quantum systems it that today's algorithms require exponentially more steps to simulate a quantum system as the number of quantum particles increases; quantum computers do not have this problem (see **Section 5.1.2, "Modeling Chemical Reactions"**). However, we do not know if this exponential scaling is fundamental or not; an answer to that question would likely also result in an answer to the question of whether or not $P = NP$.

- **Quantum annealers.** Quantum annealers achieve quantum effects in specially prepared materials. D-Wave System's quantum annealer is the most well-known device in this category. A quantum annealer uses a metal material that exhibits quantum properties as it is cooled to temperatures close to absolute zero. Unlike a general purpose quantum computer, which uses gates to process qubits, the annealer is analog. The annealing process directly manipulates qubits.

  Quantum annealers are limited in function. Although D-Wave's machines have literally thousands of qubits,[10] the numbers cannot be compared with other kinds of quantum computers because the D-Wave qubits are not universal: they can only be used to solve a limited range of quantum problems. Specifically, the D-Wave can only solve problems phrased as quadratic unconstrained binary optimization (QUBO) calculations. When it comes to QUBO problems, D-Wave can solve problems that are significantly larger than almost all private companies in the field. D-Wave also hopes that its ability to solve optimization problems will make the system commercially attractive today to companies not interested in learning about quantum computing, but interested in actually using quantum computing to solve other problems. At this point, however, there is no clear

[10] D-Wave Systems scaled its annealer from 128 qubits, the D-Wave "One" released in 2011, to the D-Wave 2,000Q, a 2000-qubit annealer, in 2017. The 2,000Q has been commercially available since 2017; popular reporting suggests a $15m price tag (Temperton, "Got a Spare $15 Million? Why Not Buy Your Very Own D-Wave Quantum Computer" (2017)). The D-Wave advantage (2020) has 5000 qubits.

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

evidence that D-Wave's systems are more cost effective at optimizing than existing commercial optimizers such as CPLEX and Gurobi, run on traditional electronic computers.

- **Quantum simulators.** The Feynman vision that quantum computers would simulate quantum interactions is being pursued in the form of quantum simulators, devices that use, "entanglement and other many-particle quantum phenomena to explore and solve hard scientific, engineering, and computational problems," as described by a report signed by 37 attendees of a 2019 workshop organized by the National Science Foundation. According to the workshop report, there are now more than 300 quantum simulators operating around the world based on a wide variety of underlying platforms. Those working in the field are pursuing a two-phase strategy: in the first phase, early prototypes are built that are research curiosities in themselves. These early devices are intended to bridge to a second phase where a broader set of researchers can employ quantum simulation, with a goal of moving second-generation devices out of quantum computing applied research laboratories and into other fields such as botany, chemistry, materials science, astronomy, and in the creation of other quantum devices, including quantum internet technologies (discussed in Chapter 7). That is, the goal is to stop doing research on quantum simulators, and to start doing research *with* quantum simulators.

  Quantum simulators are similar in design to quantum computers, but as with quantum annealers, quantum simulators are not universal: simulators are constructed with a single goal of simulating quantum mechanical systems, and often on a single scientific problem, such as understanding photosynthesis. By taking the complexities involved in the pursuit of universality off the table, some see quantum physics simulators as the most compelling near-term strategy for quantum computing. The NSF group predicted: "Scaling existing bottom-up quantum simulators to hundreds or even thousands of interacting, entangled, and well-controlled quantum elements is realistically within reach."[11]

[11] Altman et al., "Quantum Simulators: Architectures and Opportunities" (2019).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

- **Noisy Intermediate-Scale Quantum Devices (NISQ).** NISQs represent the state-of-the-science in programmable digital quantum computing. Universities, research labs, and private companies are pouring untold sums of money into developing an "intermediate-scale" device that could lend insights into the building of larger devices. That is, a mid-scale quantum computer with 50–100 qubits might reveal characteristics of materials or engineering that make creation of a 500-qubit device possible, and so on.

  NISQs are being built with several technology substrates, all familiar to readers of **Chapter 2, "Quantum Sensing and Metrology"**. Several large companies such as Google and IBM are betting on the superconducting circuit approach, where Josephson junctions form the basis of the architecture. This is the same underlying approach as superconducting quantum interference devices discussed in Section 2.2 (p. 40).

  Others, such as Honeywell, are experimenting with ion trap approaches (see Figure 6.5), where charged electronic particles are held in position with lasers, magnetic fields, or even in a physical substrate, such as the nitrogen-vacancy approach discussed in Section 2.2 (p. 41). Ion traps do not require supercooling and enjoy long coherence times, but to date have been very limited in their number of qubits.[12]

  Photons are another option for NISQs. Photonic approaches also avoid supercooling and have good stability, and can be implemented using existing materials, like silicon and optical devices from commercial providers such as ThorLabs. As of this writing, the largest quantum computer is a photonic interferometer in China, but the device is limited to a single scientific application (see Figure 6.6).

  Microsoft is pursuing a cutting-edge approach known as "topological qubits," which involves splitting an electron in order to store information redundantly and thus manage noise problems

[12] In June 2020, Honeywell announced that it had created "the world's highest performing quantum computer," bench-marking it with IBM's notion of a "quantum volume" of 64 (Honeywell, "The World's Highest Performing Quantum Computer Is Here" (2020)). The computer had only six qubits, yet its interconnection and low noise led the company to make dramatic performance claims (Crane, "Honeywell Claims It Has Built The Most Powerful Quantum Computer Ever" (2020)).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

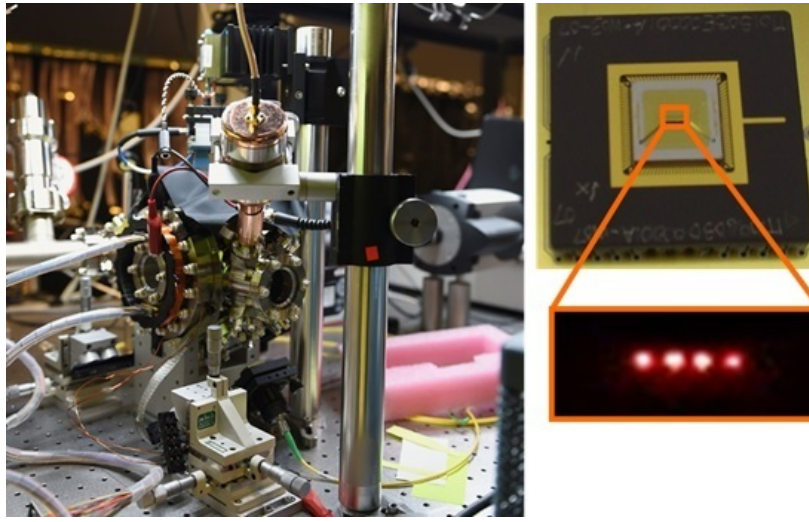

Figure 6.5. The device on the left is a vacuum chamber that houses four trapped ytterbium ions (on right) from Sandia National Laboratory. These ions can be measured using single-photon-sensitive media and are hoped to be a substrate for quantum computing and quantum memory. Photo courtesy US Air Force.

that cause decoherence. This approach is promising, but it is not nearly as developed as other approaches.

Despite their cutting-edge engineering, The National Academies of Sciences (NAS) characterizes NISQs as having "primitive" gate operations and as being plagued by error and decoherence. NAS' 2019 report concluded that today's NISQs will never scale to become the large-scale, general purpose quantum machines so desired.[13]

- **Large-scale quantum computers.** For many of the above-described efforts, the goal is to create a large, stable, universal digital quantum computer with millions of error-corrected qubits. Such a device would be similar to a modern high-performance computer. Stored in its creator's cloud warehouse, its universal functionality could be leased out to users to solve all manner of interesting problems. The question is now to realize that goal.

[13] Grumbling and Horowitz, *Quantum Computing: Progress and Prospects* (2019).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

One path is through fundamental discoveries in materials science, chemistry, or physics that can be applied to manage qubits. Indeed, while cryptanalysis grabs the news headlines, companies in quantum computing identify chemistry and materials science as their research focus. This is because with a mid-scale quantum computer, one might discover fundamental insights in materials design and in chemistry that elucidate strategies to build a larger quantum computer. Thus, like classical computers before it, quantum computer strategy is to trigger a virtuous cycle of growth. This insight also foreshadows an innovation policy issue: groups that can make those fundamental observations are likely to pull ahead of the pack, building ever-larger computers with teams that were trained over decades, using discoveries that competitors cannot obtain. In this large-scale scenario, quantum computing could be a *winner-take-all technology*, suggesting that the first innovator might well become the most successful one.

Alternatively, the path to the large-scale quantum computer may be just a matter of scaling up existing approaches. This appears to be the strategy of several reputable companies in the quantum computing field that are creating ever-larger devices based on superconducting circuits. Perhaps the manufacture of densely produced, well connected and controlled Josephson junctions will yield room-sized quantum computers with millions of qubits.

When will a large-scale quantum device be built? Even scientists at companies known to enthusiastically promote their technologies say that it will take a decade. Some say several decades. Others say this task is impossible. The next section turns to the reasons why building a quantum computer is so difficult.

6.3 Skeptics Present Quantum Computing's Challenges

Almost 20 years ago, physicists Jonathan P. Dowling and Gerard J. Milburn wrote that humankind had entered a new stage of quantum information science: the second quantum revolution. In the first quantum revolution, scientists used quantum mechanics to better understand our reality. Truly a scientific revolution, the first period of QIS started with theory and expanded over the century as more insights were gained (see Appendix A and Appendix B). The second

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

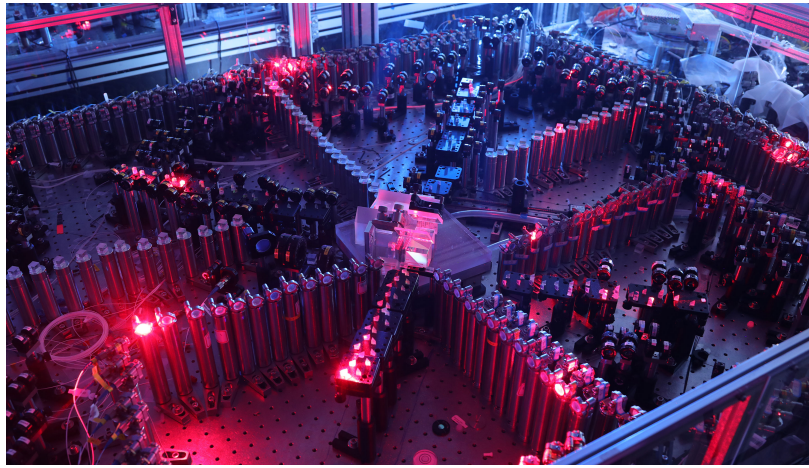

Figure 6.6. In 2020, Jian-Wei Pan and Chao-Yang Lu at the University of Science and Technology of China built a large-scale interferometer to solve the "boson sampling" problem, a task insoluble with classical computers. With 25 laser sources and 100 single-photon sensors, the Jiuzhang Device demonstrates the link between quantum sensing and computing. Image courtesy of Jian-Wei Pan.

quantum revolution is a technological one, where scientists actively employ "quantum mechanics to alter the quantum face of our physical world."

Dowling and Milburn canvassed the exciting state-of-the-science developments of this second revolution. Finally they warned that, "A solid-state quantum computer is probably the most daunting quantum technological challenge of all and will require huge advances in almost all the areas of quantum technology we have discussed."[14]

Significant progress has been made since then. Nevertheless, quantum computing still depends on realizing a number of technical feats. Until now we've presented the challenges as significant but surmountable. However, a significant number of well-credentialed experts maintain that general purpose quantum computing is simply not achievable with physics as we understand it today. This section details those challenges.

6.3.1 Scientific Challenges

A 2019 National Academies of Sciences (NAS) report characterized quantum computing as consisting of creating small, proof-of-concept,

[14]Dowling and Milburn, "Quantum Technology: The Second Quantum Revolution" (2003).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

demonstration devices.[15] This is because quantum computing requires a mastery of quantum superposition and entanglement, development of software and control systems, and management of costly, difficult physical conditions. But more than that, breakthroughs in quantum computing may also require fundamental breakthroughs in basic physics – or at very least, transitioning phenomena that have only been observed in a laboratory setting (and only in the last decade) into engineering prototypes.

To get an idea of the gap between theoretical advance and engineering realization, consider that Microsoft's approach, the "topological qubit,"[16] is based on a 1937 theoretical prediction that single electrons can be split into subparticles.[17] Now Microsoft hopes to use the phenomena to create a working quantum computer. But it took 75 years between the theory's discovery to produce evidence that the subparticles exist.[18] Microsoft collaborated with the Delft University of Technology (TU Delft), the oldest and largest Dutch public technological university in the Netherlands to substantiate the existence of the particles. In 2018, Microsoft published a paper with more evidence but the paper was retracted in 2021.[19]

Some argue that quantum computing will never be achieved; indeed, some claim that modern quantum computing research efforts are reaching the end of what they can accomplish. Physicist Mikhail Dyakonov wrote a short book about the challenges and reprinted a warning that Rolf Landauer urged scientists to include in their papers and talks: "This scheme, like all other schemes for quantum computation, relies on speculative technology, does not in its current form take into account all possible sources of noise, unreliability and manufacturing error, and probably will not work."[20]

A chorus of other commentators have downplayed quantum computing as an overhyped phenomenon. In 2015, a US Air Force advisory board found that technology advocates "herald[ed]" imminent

[15] Grumbling and Horowitz, *Quantum Computing: Progress and Prospects* (2019).
[16] Microsoft Corp., "Developing a Topological Qubit" (2018).
[17] Majorana and Maiani, "A Symmetric Theory of Electrons and Positrons" (2006).
[18] Mourik et al., "Signatures of Majorana Fermions in Hybrid Superconductor-Semiconductor Nanowire Devices" (2012).
[19] H. Zhang et al., "Quantized Majorana Conductance" (2018).
[20] Dyakonov, *Will We Ever Have a Quantum Computer?* (2020); Dyakonov, "When Will Useful Quantum Computers Be Constructed? Not in The Foreseeable Future, This Physicist Argues. Here's Why: The Case against: Quantum Computing" (2019).

breakthroughs but nevertheless, "no compelling evidence exists that quantum computers can be usefully applied to computing problems of interest to the Air Force."[21]

The most specific critique comes from the 2019 NAS report of the field that made both economic and technological assessments. On the economic front, the NAS group observed that there are essentially no economically advantaged uses for quantum computers for the foreseeable future (and obviously no consumer ones either).[22] This is directly different from the history of computing, in which spending money on computing was advantageous from the very first dollar spent. From the beginning, spending money on computing – be it mechanical, electromechanical, or electronic – made it possible to do something that wasn't otherwise possible, or to do it faster, or for less money overall. Although quantum computing might one day make it possible to train large-scale artificial intelligence machine learning models faster and with far less electricity than is currently the case, this does not seem to be a breakthrough that is plainly visible on the short-term horizon.

6.3.2 Engineering Challenges

Without uses that produce big savings or profits in the near term, funding for quantum computing is likely to be limited to governments and the largest technology companies. As such, quantum computing lacks the "virtuous cycle," like what was enjoyed with classical computers, with increasing commercial and consumer utility driving demands and willingness to pay for fantastic technological innovations.

The NAS survey's core technological critique relates to the difficulty of scaling up today's quantum systems into larger systems that can be used to solve meaningful problems. As a result of these challenges, the survey found it too uncertain to predict when a scalable quantum computer would be invented and that existing devices could never scale into general-purpose machines.

Quantum computers are characterized by the integration of multiple qubits. Thus, for a quantum computer to work, one needs to be able to encode, entangle, manipulate, and maintain an array of qubits, raising the challenges visited in Chapter 2. The challenges inherent in quantum computing are thus different from the obstacles

[21] US Air Force Scientific Advisory Board, *Utility of Quantum Systems for The Air Force Study Abstract* (2016).

[22] Grumbling and Horowitz, *Quantum Computing: Progress and Prospects* (2019).

encountered by engineers building and then scaling digital computers. Classical computers went through an evolution of mechanical, to relay, to tube, and to discrete transistors, and finally to integrated circuits. Each improvement produced systems that were smaller, faster, and required less overall energy to perform a computation. Semiconductors enjoyed their own virtuous cycle, providing chip makers with tools for designing and manufacturing computers that were ever more complex yet less expensive. Quantum computing has not realized a scaling breakthrough on the level of the transistor. Perhaps more to the point, there is no such breakthrough lurking in the future of any realistic technology road map. In many ways this is similar to the days of mechanical, electromechanical and tube-based computing, when larger computers might be faster than smaller ones, but they were also dramatically more expensive and less reliable.

Different technologies can be used to create qubits, but for each, quantum scientists must be able to master and control events at quantum scales (see Appendix A). Mastery and control require substantial technical expertise, reflected in the multidisciplinary nature of quantum computing teams (engineers, physicists, mathematicians, computer scientists, chemists, materials science). This is also a difference from the last 70 years of computing, which generally required mastery of fewer technical domains, and where modularization and isolation between technical domains meant less interdisciplinary work.

Quantum computers require that their qubits be entangled, cohered into a group that can be operated upon. But at the same time, quantum computers must be shielded from the universe, lest noise in the environment cause those qubits to decohere. This makes the quantum computer challenge fundamentally different from the classical computer. The transistor allowed scale with intricately managed stability. However, with quantum computers, scale requires the management of additional, exquisitely fragile quantum states.

When qubits decohere, they lose information. Thus, quantum algorithms have to be crafted to be efficient enough to execute before coherence is lost. As of this writing, some state-of-the-science devices have coherence in the hundreds of *micro*seconds, a time too short for the quantum gates of today to process significant numbers of qubits. This is a time period so short that human physical experience has no analogue for it. A blink of the eye takes about 100 000 microseconds.

The longer quantum computers run, the more performance de-

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

grades. In classical computing, extra bits are used to correct ordinary errors that occur in processing. This approach works because of all the engineering performed in classical computers to avoid quantum effects like tunneling. In quantum computing, many of the qubits employed are dedicated to error correction, so many that it creates significant overhead and degrades computing performance. Current thinking is that to emerge from the era of NISQ machines, as many as 90 percent of a quantum computer's qubits might have to be dedicated to error correction.[23] Initially, one might suggest just adding more qubits to achieve reliability, but as more qubits are added, system complexity increases, and quantum devices become more prone to both random environmental interference and to noise from the computer's own control system.

Quantum computers are not fault tolerant. In addition to temperature, vibration and electromagnetic interference can easily destabilize quantum computers. Conventional electronic computers rely on the digital discipline to smooth out errors so that they effectively do not matter.[24] In quantum devices, by contrast, errors are not rounded out, but instead compound until the conclusion of the computation.

To shield quantum computers from environmental noise that triggers decoherence, many quantum computer architectures require supercooling. This cooling is *super* because it is colder than even the background temperature of the universe. Extreme frigidity is needed both to elicit quantum properties from materials (for instance, in analog quantum annealers) but also because heat increases the chances that random energy collisions will generate noise that will interfere with quantum states or cause decoherence.

Keeping quantum devices at 15 millikelvin (−273 °C, −459 °F) means that quantum computer scientists need liquid helium, an increasingly rare and valuable element, of which there is a finite supply on Earth. There are currently no limits on the usage of Earth's helium supply.[25] Unlike quantum computing, many other quantum

[23] Möller and Vuik, "On The Impact of Quantum Computing Technology on Future Developments in High-Performance Scientific Computing" (2017).

[24] In classical computing, bits of data are either a `0` or `1`. In that environment, error appears as a decimal value such as `0.1` or `0.9` that can be easily rounded to `0` or `1`. For more information, see p. 84.

[25] Some hope that early quantum computers will solve fundamental challenges in fusion. If that happens, we could create helium via hydrogen fusion.

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

technologies do not require supercooling. This means that some sensing and communications technologies can be miniaturized, commercialized, and deployed in many more challenging contexts (in outer space, underwater, in missiles) than quantum computers.

6.3.3 Validation Challenges

It will be necessary to validate quantum computers to make sure that the answers they produce are correct. Ironically (and annoyingly), validation is easy for many of the hard, long-term applications for quantum computing, but likely to be harder for the more probable, near-term applications.

For the algorithms like factoring with Shor's algorithm and search with Grover's, validation is easy: just try the answer provided by the quantum computer and see if it works. That is, if the quantum computer says that the `2227` are `131` and `17`, one need merely multiply 131×17 to determine if the factorization is correct or not. The same logic applies to using Grover's algorithm to crack an AES-128 key: just try to decrypt the encrypted message: if the message decrypts, the AES-128 key is correct.

On the other hand, approaches for both error correction and validation are less developed for analog quantum simulators. One approach suggested in the 2019 NSF report is to run simulations forward and backwards (theoretically possible, since the computations should be reversible) to see if the simulator retraces its steps. Another approach is to see if different systems that should have equivalent outcomes do indeed have similar outcomes.

6.3.4 Ecosystem Challenges

A final challenge is not technical, but organizational. Significant work still needs to be done to create a rich ecosystem of quantum software. Beyond basic programming languages and compilers, which exist today, there is need for documentation for people at multiple levels of expertise, programming courses, systems on which to run those programs, and finally organizations willing to pay for training and to hire quantum programmers.

On the software front, many teams are developing languages to make interaction with quantum computers more routine and standardized. As of 2021, a growing "zoo" of quantum algorithms in-

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

cluded 430 papers.[26] But the overwhelming number of these algorithms are expressed as *papers* in *scientific journals* or on *preprint servers*; they are not code on sites like GitHub that can be downloaded, incorporated into other, larger quantum programs, and run. Recall that Ed Fredkin got himself hired without a college degree to write programs for BBN's first computer in 1956 (and which he convinced BBN to purchase – see Section 4.4.1 (p. 146)). We have not yet reached the point where it is possible to teach yourself quantum programming and get a job at a company that needs someone to write quantum algorithms to run on their quantum computer.

6.3.5 Quantum Supremacy and Quantum Advantage

Quantum Supremacy is an awkward term. As Preskill defined it in 2012, the goal is to perform a computation – any computation – that cannot be performed with a classical computer. But the term is misleading, because quantum engineers in China and the US have clearly achieved "supremacy" as defined by Preskill, but quantum computers are not supreme: for the vast majority of computations performed on planet Earth, you would not be able to use one of today's quantum computers. And even if reliable, large-scale quantum computers are available in the future, it is hard to imagine that these machines will be used for more than a tiny fraction of the world's computing problems. And even in these applications, quantum computers are likely to be co-processors that depend on classical computers for many functions. For these reasons, we prefer the term "quantum advantage" to describe the achievement of solving a problem with a quantum device that cannot be solved with a classical computer.

In December 2020, Jian-Wei Pan and Chao-Yang Lu made the most compelling claim of quantum advantage to date.[27] Their team built a large-scale interferometer to compute a specific problem, Gaussian Boson Sampling (GBS). The team named their device Jiuzhang, for the ancient Chinese manuscript focused upon applied mathematics, *Nine Chapters on the Mathematical Art*. But as exciting as the Jiuzhang development is, the device can perform just one computation. However, it's really fast!

Previously, Google researchers announced in October 2019 that they had achieved quantum supremacy using their 54-qubit Syca-

[26] Montanaro, "Quantum Algorithms: an Overview" (2016); S. P. Jordan, "Quantum Algorithm Zoo" (2021).

[27] Zhong et al., "Quantum Computational Advantage Using Photons" (2020).

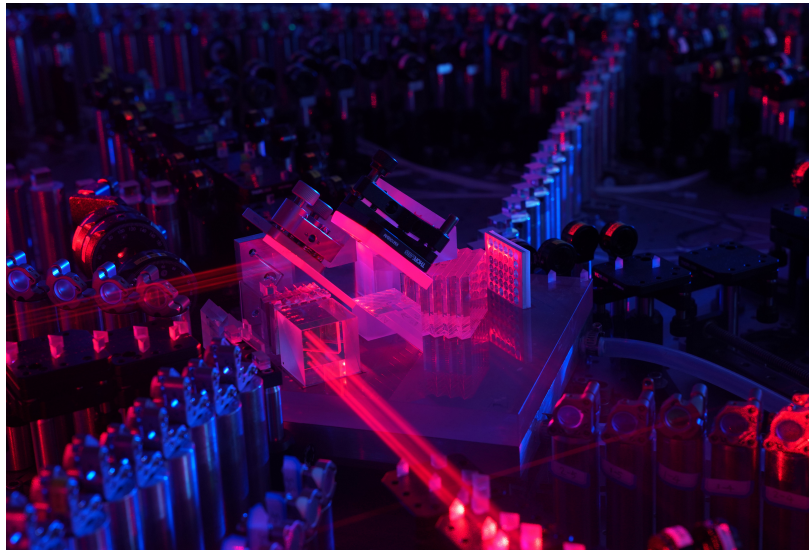

Figure 6.7. Computing a specific distribution of photons that would have taken 600 million years to solve on the fastest existing classical supercomputer in 2020 was done in 200 seconds with a reported 99 percent fidelity by Jian-Wei Pan and Chao-Yang Lu at the Hefei National Laboratory, University of Science and Technology of China. However, turning the device into a "fault-tolerant universal quantum computer, is a very long-term goal and requires many more challenges to tackle, including ultra-high-efficiency quantum light sources and detectors, and ultra-fast and ultra-low-loss optical switch," Lu told us. Image courtesy of Jian-Wei Pan.

more superconducting approach.[28] Google's researchers programmed their computer to create and then evaluate random quantum circuits. IBM, a chief rival to Google, quickly disputed the supremacy claim, arguing on its research blog that "ideal simulation of the same task can be performed on a classical system in 2.5 days and with far greater fidelity."[29] In March 2021, two Chinese scientists claimed that they replicated the Google approach with higher fidelity using classical GPUs.[30] The scientists concluded with a humble brag that their "proposed algorithm can be used straightforwardly for simulating and verifying existing and near-future NISQ quantum circuits" and helpfully posted their approach on GitHub. These quick retorts

[28] Arute et al., "Quantum Supremacy Using a Programmable Superconducting Processor" (2019).

[29] Pednault et al., "On 'Quantum Supremacy'" (2019).

[30] Pan and P. Zhang, "Simulating The Sycamore Quantum Supremacy Circuits" (2021).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

The Helium Challenge

Helium's stability, non-reactivity, and phase as a fluid at near absolute zero makes it useful for cooling both quantum computers and the magnets in Magnetic Resonance Imaging machines. And while helium is abundant in the universe, on Earth it is a non-renewable resource. The small amount of helium that our planet has is the result of underground radioactive decay. Helium is rendered along with natural gas; if it is released and not captured, it is no longer financially viable to collect from the air.

The US and Qatar are the largest producers of helium, with the US supply provided by a storage and enrichment facility in Amarillo, Texas, run by the US Bureau of Land Management. Russia's Gazprom and China are building plants in order to reduce their reliance on US sources. Because of helium's many uses, limited availability, and strategic relevance, conservationists have called for an international helium agency to preserve supply and prevent a crisis in availability, and to expand extraction of helium from existing natural gas plants.[a] But don't feel guilty about helium balloons: such consumption is inconsequential compared to industrial and medical uses.

Today the biggest consumers of helium are MRI machines and devices used at border crossings to detect dirty bombs and other nuclear devices. Quantum computers use less helium, and modern cryogenics equipment attempts to conserve and recycle it. D-Wave explicitly markets its annealer as recycling helium to avoid the need to continuously resupply the machine's local store of helium.

Some quantum computers require light helium, Helium-3. This is extracted from nuclear reactors, and is somewhat controlled. IBM's plans for a 1000-qubit superconducting device caused the company to develop a custom dilution refrigerator. Others are building supercooling capacities that do not use a cryogen like helium or liquid nitrogen. These non-cryogen coolers have a major disadvantage: they require much more electricity for cooling. However, as nations signal an interest in decoupling their technology stacks, nations without access to helium sales may simply turn to electric cooling.

[a] Nuttall, Clarke, and Glowacki, "Stop Squandering Helium" (2012).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

to Google's claim demonstrate how scientists value their quantum computing bragging rights, even if the bragging is only about the ability to solve otherwise meaningless random quantum puzzles.

The Jiuzhang device is a clear demonstration of quantum advantage, but the device has no practical application. Whereas Google's claim of advantage stands on contested ground, its Sycamore device can be programmed to solve problems other than random puzzles, so it is probably more important from a commercial point of view.

For computer scientists, achieving quantum advantage was long seen as a kind of Rubicon. But for most organizations, the real quantum computing Rubicon will be the moment that quantum computing can perform some useful commercial, defense, or intelligence application. Competitors strive to make the case that they have some advantage to sell from quantum computing. Perhaps the most promising in the near term are proposals that use quantum computers to solve part of a problem or those that apply "low-depth algorithms" that promise some quantum speedup with practical payoff. For instance, Goldman Sachs proclaimed that by optimizing algorithms, there will be a quantum advantage in derivatives pricing from even small quantum computers by 2025.[31] If they are correct – or even if other financial services firms believe that Goldman Sachs is correct – the development could create a gold rush in quantum computing.

How can one make sense of quantum computers' power when they rely on different physical media (ranging from photonics to trapped ions to annealing) and when innovators claim to have more qubits than competing devices? Quantum computers cannot be evaluated simply by the number of qubits they have, otherwise D-Wave's 2000-qubit system would be leagues ahead of teams at IBM, Google, and Microsoft – even when those systems can clearly perform computations that the quantum annealer can't. To evaluate quantum devices, IBM created its own metric called *quantum volume*.[32] A computer's quantum volume is "the largest random circuit of equal width and depth that the computer successfully implements." Thus, quantum volumes are necessarily perfect squares: 2, 4, 9, 16, and so on. Unfortunately, the largest quantum volume that IBM measured was 16,

[31] Giurgica-Tiron et al., "Low Depth Algorithms for Quantum Amplitude Estimation" (2020).

[32] Cross et al., "Validating Quantum Computers Using Randomized Model Circuits" (2019).

on a machine with 4 qubits running a circuit with a depth of four gates. "We conjecture that systems with higher connectivity will have higher quantum volume given otherwise similar performance parameters," the authors state.

Despite all these challenges, governments and large technology companies (e.g. Fujitsu, Google, IBM, Microsoft, Toshiba) have devoted major resources to quantum computing, and several startups (e.g. IonQ, Rigetti, Xanadu) are betting the company on it. Competition has produced wonderful resources to learn about and even experiment with quantum computing. For instance, IBM and others have made instructional videos, extensive, carefully curated explanatory material, and even made rudimentary quantum computers available through the Web at quantum-computing.ibm.com for anyone who wants to try their hand at programming the machines.

Quantum computing efforts are either basic or applied research. Basic research projects, like the Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN), can be huge impressive projects that reveal fundamental truths about the nature of the universe: at a cost of approximately $9 billion, the LHC is one of the most expensive scientific instruments ever built, and it is responsible for the "discovery" of the Higgs boson, but it is hard to draw a line from the LHC to improvements in day-to-day life of anyone except for several thousand construction workers, physicists, and science journalists. On the other hand, nuclear fission was discovered in December 1938 by physicists Lise Meitner and Otto Frisch,[33] which led to the creation of a working nuclear bomb within just seven years and the first nuclear power plants in 1954. Such is the unpredictability of research.

6.4 The Outlook for Quantum Computing

The long-term outlook for quantum computing may be hazy, but the near-term outlook for quantum computing companies appears to be quite bright.

As we saw in the last chapter, although it was the potential for quantum computers to crack codes that led to the initial burst of enthusiasm, interest in quantum computing is likely being sustained by the promise of using quantum technology as an advanced scientific instrument for learning more about quantum physics and quantum

[33] Tretkoff, "This Month in Physics History: December 1938: Discovery of Nuclear Fission" (2007).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

chemistry. The payoffs may be directly in these fields, or they may simply be the development of superior quantum sensors that are usable throughout the military industrial complex.

As such, there are many practical regulatory implications at least in the short term:

1. Because of their expense and complexity, only large firms and governments are likely to be able to afford quantum computers for some time. This means that governments have a relatively small number of players to police in quantum computing, and that the technologies may be easier to monitor and control. This period of large-organization exclusivity may continue for decades. Consider that classical computers were the domain of universities, governments, and large companies until the personal computer revolution of the 1970s.

2. Because of their complexity, quantum computers require teams of multidisciplinary experts. This means that one cannot simply sell a quantum computer and expect a user to make sense of it. Sellers will be on the premises of buyers and will probably know about the buyers' intended uses of the devices. The business model may be selling services as much as selling the device itself.

3. Because of their sensitivity to interference of all types, quantum computers are likely to be placed in low-noise environments. For instance, the D-Wave system occupies a $10 \times 10 \times 10$ foot housing plus three auxiliary cabinets for control systems. The cabinet is part of a system to produce quantum effects in D-Wave's annealer, where the chip is the size of a thumbnail. This requires a vacuum environment, a low-vibration floor, shielding to 50 000 times less than the Earth's magnetic field, and cooling to 0.0012 K.[34] Such devices are unlikely to be installed in jets for forward-deployed use, although they might be deployable in a suitably outfitted ship.

4. Finally, large firms that build the first quantum computers are likely to offer services through the cloud until the engineering becomes easier and medium-sized enterprises can purchase their own devices. Until then, quantum computing is likely to

[34] R. Copeland, "The International Quantum Race" (2017).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

be offered as an enhanced service, one optimized for specific problems.[35,36]

Taken together, these limits will shape the trajectory and offerings of quantum computers.

Despite the lack of a practical demonstration, many scientists believe that sufficiently large quantum computers will be much more powerful than classical computers for solving certain kinds of problems. We lack *proof* that quantum computers will be innately more powerful for the same reason that we lack proof that factoring is fundamentally more difficult than primality testing, or that mixed integer linear programming is fundamentally harder than linear programming. That is, we don't have a proof that $P \neq NP$.

[35] Ibid.

[36] Gibney, "The Quantum Gold Rush" (2019).

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press

https://doi.org/10.1017/9781108883719.010 Published online by Cambridge University Press