1

Introduction

The internet is one of the key technological achievements of the twentieth century, an enabling factor in every aspect of our everyday use of modern technology. Whereas digital computing was the definitive technology of the twentieth century, quantum technologies will be for the 21st [127, 23].

Perhaps the most exciting prospect in the quantum age is the development of quantum computers. Richard Feynman [65] was the first to ask the question *'If quantum systems are so exponentially complex that we are unable to simulate them on our classical computers, can those same quantum systems be exploited in a controlled way to exponentially outperform our classical computers?'* Subsequently, the Deutsch-Jozsa algorithm [52] demonstrated for the first time that algorithms can run on a quantum computer, exponentially outperforming any classical algorithm. Since then, an enormous amount of research has been dedicated to finding new quantum algorithms, and the search has indeed been a very fruitful one,[1] with many important applications having been found, including, amongst many others:

- Searching unstructured databases:
  - Grover's algorithm [83].
  - Quadratic speedup.
- Satisfiability and optimisation problems:[2]
  - Grover's algorithm.
  - Quadratic speedup.
  - Includes solving **NP**-complete problems and brute-force cracking of private encryption keys.

[1] See the Quantum Algorithm Zoo for a comprehensive summary of the current state of knowledge on quantum algorithms.

[2] A satisfiability problem is one in which we search a function's input space for a solution(s) satisfying a given output constraint. The hardest such problems, like the archetypal 3-SAT problem, are **NP**-complete.

https://doi.org/10.1017/9781108868815.003 Published online by Cambridge University Press

  - Many optimisation problems are **NP**-complete or can be approximated in **NP**-complete.
- Period finding and integer factorisation:
  - Shor's algorithm [165].
  - Exponential speedup.
  - This compromises both Rivest, Shamir and Adleman (RSA) and elliptic-curve public-key cryptography [141], the most widely used cryptographic protocols on the internet today.
  - This problem is believed to be **NP**-intermediate – an **NP** problem that lies outside **P** (and is therefore classically hard) but that is not **NP**-complete (the 'hardest' of the **NP** problems).
- Simulation of quantum systems:
  - Lloyd's algorithm [107].
  - Exponential speedup.
  - This includes simulation of molecular and atomic interactions in the study of quantum chemistry or nuclear physics; interactions between drug molecules and organic molecules for drug design; genetic interactions for the study of genetics and genetic medicine; nanoscale semiconductor physics for integrated circuit design; and much more.
- Simulation of quantum field theories:
  - Jordan-Lee-Preskill algorithm [94, 34].
  - Exponential speedup.
  - A key area of fundamental physics research.
- Topological data analysis:
  - Lloyd's algorithm [108].
  - Exponential speedup.
  - Broad applications including social media network analysis; consumer behaviour; behavioural dynamics; neuroscience; and higher-dimensional signal and image processing.
- Solving linear systems of equations:
  - Algorithms by [84, 26].
  - Exponential speedup.
  - Widespread applications in linear algebra and calculus.
- Quantum machine learning:
  - Lloyd's algorithm [109].
  - This includes putting an end to humanity.

https://doi.org/10.1017/9781108868815.003 Published online by Cambridge University Press

Some of these are discussed in more detail in Chapter 28.

It is likely we have not yet begun to fully recognise the capabilities of quantum computers and the full plethora of applications they may have in the future. We stand at the beginning of the emergence of an entirely new type of technology.

In addition to many practical applications, the onset of quantum computing carries with it deep philosophical implications, specifically, the extended Church-Turing (ECT) thesis hypothesises that any physically realisable system can be *efficiently*[3] simulated by a universal Turing machine (i.e., classical computer). The believed exponential complexity of quantum systems inclines quantum computer scientists to believe that the ECT thesis is therefore false [50].[4] The demonstration of large-scale quantum computers, though unable to prove or disprove the ECT thesis,[5] could at least provide some convincing evidence against the ECT conjecture.

From a computational complexity theorist's perspective, it is strongly believed that the complexity classes of problems efficiently solvable on classical computers (**P** and **BPP**) and quantum computers (**BQP**) are distinct. Specifically, it is believed that $\mathbf{BPP} \subset \mathbf{BQP}$. If this conjecture is correct, it implies the existence of quantum algorithms superpolynomially faster than the best classical ones and that the ECT thesis is not correct. More specifically, Figure 1.1 illustrates the believed relationships between some of the most important complexity classes relevant to quantum computing.

In addition to quantum computing, quantum cryptography holds the promise of uncrackable cryptographic protocols, guaranteed not by the assumed complexity of solving certain mathematical problems like integer factorisation or brute-force searching but by the laws of quantum mechanics. That is, provided that our understanding of quantum mechanics is correct, quantum cryptographic protocols exist that cannot be cracked, irrespective of the computational resources of an adversary.

Already we are beginning to see elementary realisations of essential quantum technologies such as quantum computing, cryptography and metrology. As these technologies become increasingly viable and more ubiquitous, the demand for networking them and sharing quantum resources between them will become a pressing issue. Most notable, quantum cryptography and *cloud quantum computing* will be pivotal in the proliferation of quantum technology, which necessarily requires reliable quantum communications channels.

[3] The term 'efficient' is one coined by the computer scientist to mean that a problem can be solved in time at most polynomial in the size of the problem.

[4] We have discovered a truly marvellous proof of this, which this footnote is too narrow to contain.

[5] When we talk about 'scalability' or the 'ECT thesis' we are talking about asymptotic relationships. Clearly no finite-sized experiment can prove asymptotic scaling with certainty. But with a sufficiently large quantum computer at our disposal, demonstrating exponentially more computational power than its classical sibling, we might be reasonably satisfied in convincing ourselves about the nature of the scaling of different computational models.

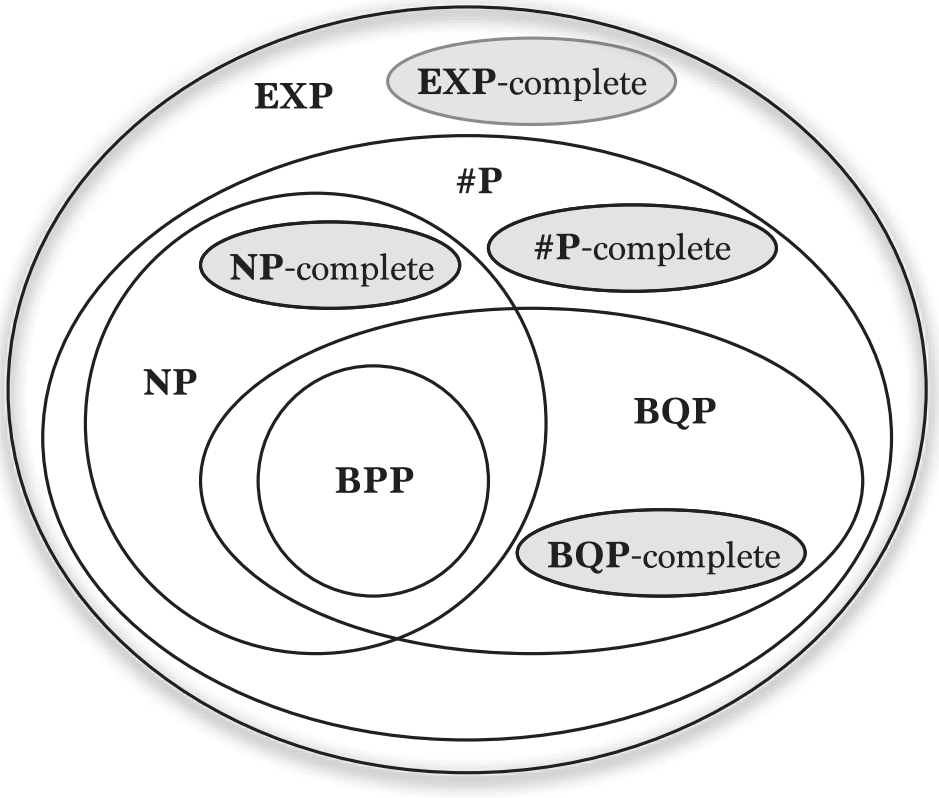

Figure 1.1 Believed relationships between the complexity classes most relevant to quantum computing. **BPP** is the class of polynomial-time probabilistic classical algorithms. **NP** is the class of problems verifiable in polynomial time using classical algorithms. **NP**-complete are the subset of **NP** problems polynomial-time reducible to any other problem in **NP**, similarly for other 'complete' problems. **BQP** is the class of probabilistic algorithms solvable in polynomial time on universal quantum computers. **#P** is the set of counting problems that count satisfying solutions to **P** problems (**P** is the same as **BPP** but deterministic rather than probabilisitic). **EXP** is the class of all algorithms that require exponential time. Note that it is actually unproven whether $\mathbf{P} = \mathbf{BPP}$ or $\mathbf{P} \subset \mathbf{BPP}$. There are examples where the best known **BPP** algorithms outperform the best known **P** algorithms, which could arise because the two classes are inequivalent or because we simply have not tried hard enough to find the best deterministic algorithms. Furthermore, though it is known that $\mathbf{P} \subseteq \mathbf{NP}$, it is not known whether $\mathbf{BPP} \subseteq \mathbf{NP}$. For the sake of illustration in our Venn diagram we have taken the view that it is. **BPP** is regarded as the class of problems efficiently solvable on universal Turing machines (i.e., classical computers), whereas **BQP** is the class efficiently solvable on universal quantum computers. The computational superiority of quantum computers is based on the (strongly believed, yet unproven) assumption that $\mathbf{BPP} \subset \mathbf{BQP}$.

https://doi.org/10.1017/9781108868815.003 Published online by Cambridge University Press

The first demonstrations of digital computer networks were nothing more than simple two-party, point-to-point (P2P) communication. However, the internet we have today extends far beyond this, allowing essentially arbitrary worldwide networking across completely ad hoc networks comprising many different mediums, with any number of parties, in an entirely plug-and-play and decentralised fashion. Similarly, elementary demonstrations of quantum communication have been performed across a small number of parties, and much work has been done on analysing quantum channel capacities in this context. But, as with digital computing, demand for a future *quantum internet* is foreseeable, enabling the arbitrary communication of quantum resources, between any number of parties, over ad hoc networks.

The digital internet may be considered a technology stack, such as TCP/IP (Transmission Control Protocol/Internet Protocol), comprising different levels of abstraction of digital information [177]. At the lowest level we have raw digital data we wish to communicate across a physical medium. Above this, we decompose the data into packets. The packets are transmitted over a network, and TCP is responsible for routing the packets to their destination and guaranteeing data integrity and Quality of Service (QoS). Finally, the packets received by the recipient are combined and the raw data are reconstructed.

The TCP layer remains largely transparent to the end-user, enabling virtual software interfaces to remote digital assets that behave as though they were local. This allows high-level services such as the File Transfer Protocol (FTP), the worldwide web, video and audio streaming and outsourced computation on supercomputers, as though everything were taking place locally, with the end-user oblivious to the underlying networking protocols, which have been abstracted away. To the user, YouTube videos or Spotify tracks behave as though they were held as local copies. And FTP or DropBox allows storage on a distant data centre to be mounted as though it were a local volume. We foresee a demand for these same criteria in the quantum era.

In the context of a quantum internet, packets of data will instead be quantum states, and the transmission control protocol is responsible for guiding them to their destination and ensuring quality control.

Our treatment of quantum networks will be optics heavy, based on the reasonable assumption that communications channels will almost certainly be optical, albeit with many possible choices of optical states and mediums. However, this does not preclude nonoptical systems from representing quantum information that is not in transit, and we consider such 'hybrid' architectures in detail, as well as the interfacing between optical and nonoptical systems. Indeed, it is almost certain that future large-scale quantum computers will not be all-optical, necessitating interfacing different physical architectures.

Shared quantum entanglement is a primitive resource with direct applications in countless protocols. This warrants special treatment of quantum networks that do not implement a full network stack but instead specialise in just this one task – entanglement distribution. We will see that such a specialised network will already be immensely useful for a broad range of applications, and its simplicity brings with it many inherent advantages.

The quantum internet will enable advances in the large-scale deployment of quantum technologies. Most notable, in the context of quantum computing it will allow initially very expensive technology to be economically viable and broadly accessible via the outsourcing of computations for both consumers who cannot afford quantum computers and to well-resourced hosts who can – *cloud quantum computing*.

With the addition of recent advances in homomorphic encryption and blind quantum computing, such cloud quantum computing can be performed securely, guaranteeing privacy of both data and algorithms, secure even against the host performing the computation. This opens up entirely new economic models and applications for the licensing of compute time on future quantum computers in the cloud.

The unique behaviour of quantum computing, in terms of the superclassical scaling in its computational power, brings with it many important economic and strategic considerations that are extremely important to give attention to in the postclassical world.

But quantum technologies extend far beyond computation. Many other exciting applications for controlled quantum systems exist, with new ones frequently emerging. Thus, the quantum internet will find utility beyond cloud quantum computing, enabling the global exchange of quantum resources and assets. This could include the networking of elementary quantum resources such as state preparation, entanglement sharing and teleportation and quantum measurements or scale all the way up to massively distributed quantum computation or a global quantum cryptography network.

It is hard to foresee the future trajectory of quantum technology, much as no one foresaw the advances digital technology has made over the last half century. But it is certain that as the internet transformed digital technology, the quantum internet will define the future of quantum technologies.