

MANNING



MEAP Edition Manning Early Access Program Learn Quantum Computing with Python and Q#

A Hands-on approach Version 8

Copyright 2020 Manning Publications

For more information on this and other Manning titles go to $\underline{\text{manning.com}}$

©Manning Publications Co. To comment go to liveBook

Thank you for joining the MEAP for Learn Quantum Computing with Python and Q#: A Hands-on approach we hope you'll enjoy getting started with quantum computing!

Quantum computing has recently been gaining more interest as the research field is maturing and large tech companies are investing heavily in it. With all this hype, it is both interesting and important to learn what we could do with these new devices as well as what will likely be out of scope. Our goal with this book is to both introduce you to this exciting new field, as well as Microsoft's domain-specific programming language for quantum computers: Q#.

When we were learning quantum computing ourselves, it was very exciting but also a bit scary and intimidating. Looking back, it doesn't have to be that way, as a lot of what makes topics like quantum computing confusing is the way in which they are presented, not the content itself. With this book, there is an opportunity both for us as authors and for you as readers to learn without this obstruction, using modern programming languages like Python and Q# to help along the way.

We've written the book to be accessible to developers, rather than the "textbook" style common to most other quantum computing books. If you have done some programming before and are familiar with matrices, vectors, and some basic operations involving matrices, such as what you might use in computer graphics or machine learning, you should be good to go! We are importantly *not* assuming any prior knowledge of quantum mechanics or physics, we will help you learn what you need along the way.

By the end of the book, you should be able to:

- Understand what a quantum computer is and why it is such a rapidly expanding field
- Write quantum programs in Q#, Microsoft's new quantum programming language
- Predict what kinds of problems might be suitable to solve with a quantum computer

• Program and use quantum simulators that can be used to run quantum algorithms on classical computers today

We are excited for you to start your own journey through quantum computing by joining this MEAP, and we look forward to hearing from you in the <u>liveBook Discussion Forum</u> about what you find along your way!

—Sarah Kaiser and Chris Granade

brief contents

PART 1: GETTING STARTED WITH QUANTUM

- 1 Introducing Quantum Computing
- 2 Qubits: The Building Blocks
- 3 Sharing Secrets With Quantum Key Distribution
- 4 Nonlocal Games: Working With Multiple Qubits
- 5 Teleportation and Entanglement: Moving Quantum Data Around

PART 2: PROGRAMMING QUANTUM ALGORITHMS IN Q#

- 6 Changing the odds: An introduction to Q#
- 7 What is a Quantum Algorithm?
- 8 Quantum sensing: It's not just a phase

PART 3: APPLIED QUANTUM COMPUTING

- 9 Solving Chemistry Problems with Quantum Computers
- 10 Searching with Quantum Computers
- 11 Arithmetic With Quantum Computers

APPENDIXES

A Installing Required Software

B Linear Algebra Refresher

Introducing Quantum Computing

This chapter covers:

- · Why people are excited about quantum computing,
- · What a quantum computer is,
- · What a quantum computer is and is not capable of, and
- How a quantum computer relates to classical programming.

Quantum computing has been an increasingly popular research field and source of hype over the last few years. There seem to be news articles daily discussing new breakthroughs and developments in quantum computing research, promising that we can solve any number of different problems faster and with lower energy costs. Quantum computing can make an impact across society, making it an exciting time to get involved and learn how to program quantum computers and apply quantum resources to solve problems that matter.

In all of the buzz about the advantages quantum computing offers, however, it is easy to lose sight of the real scope of the advantages. We have some interesting historical precedent for what can happen when promises about a technology outpace reality. In the 1970s, machine learning and artificial intelligence suffered from dramatically reduced funding, as the hype and excitement around AI outstripped its results; this would later be called the "AI winter." Similarly, Internet companies faced the same danger trying to overcome the dot-com bust.

One way forward is by critically understanding what the promise offered by quantum computing is, how quantum computers work, and what they can do, and what is not in scope for quantum computing. Also it's just really cool to learn about an entirely new computing model! To develop that understanding, as you read this book you'll learn

how quantum computers work by programming simulations that you can run on your laptop today. These simulations will show many of the essential elements of what we expect real commercial quantum programming to be like, while useful commercial hardware is coming online.

1.1 WHO THIS BOOK IS FOR

This book is intended for people who are interested in quantum computing and have had little to no experience with quantum mechanics, but some programming background. As we learn to write quantum simulators in Python and quantum programs in Q#, Microsoft's specialized language for quantum computing, we'll use traditional programming ideas and techniques to help us out. A general understanding of programming concepts like loops, functions, and variable assignments will be helpful. Similarly, we will be using some mathematical concepts from linear algebra such as vectors and matrices, to help us describe the quantum concepts — if you're familiar with computer graphics or machine learning, many of the concepts we need here will be similar. We'll use Python to review the most important mathematical concepts along the way, but familiarity with linear algebra will be helpful.

1.2 WHO THIS BOOK IS NOT FOR

Quantum computing is a wondrous and fascinating new field that offers new ways of thinking about computation, and new tools for solving difficult problems — this book can help you get your start in quantum computing, so that you can continue to explore and learn. That said, this book isn't a textbook, and isn't intended to prepare you for quantum computing research all on its own. As with classical algorithms, developing new quantum algorithms is a mathematical art as much as anything else; while we touch on the math in this book and use it to explain algorithms, there's a variety of different textbooks available that can help you build on the ideas that we cover here.

If you're already familiar with quantum computing, and want to go further into the physics or mathematics, we suggest one of the following resources:

1.2.1 Textbooks and other resources for learning further

- Quantum Computing: A Gentle Introduction by by Eleanor G. Rieffel and Wolfgang H. Polak (ISBN-13: 9780262526678)
- *Quantum Computing since Democritus* by Scott Aaronson (ISBN-13: 9780521199568)
- Quantum Computation and Quantum Information by Michael A. Nielsen and Isaac L. Chuang (ISBN-13: 9781107002173)
- *Quantum Processes Systems, and Information* by Benjamin Schumacher and Michael Westmoreland (ISBN-13: 978S0521875349)

1.3 HOW THIS BOOK IS ORGANIZED

The goal of this text is to enable you to start exploring and using the practical tools we

have now for quantum computing. The text of this book is broken up into three main parts that build on each other.

- Part I will gently introduce the concepts needed to describe *qubits*, the fundamental unit of a quantum computer. This Part will describe how to simulate qubits in Python, making it easy to write simple quantum programs.
- Part II will describe how to use the Quantum Development Kit and the Q# programming language to compose qubits, and to run quantum algorithms that perform differently than any known classical algorithms.
- In Part III, we will apply the tools and methods from the previous two Parts to learn how quantum computers might be applied to real-world problems such as simulating chemical properties.

DEEP DIVE: It's OK to snorkel!

Quantum computing is a richly interdisciplinary area of study, bringing together ideas from programming, physics, mathematics, engineering, and computer science. From time to time throughout the book, we'll take a moment to point to how quantum computing draws on ideas from these other fields to put the concepts we're learning about into that richer context.

While these asides are meant to spark curiosity and further exploration, they are by nature tangential. You'll get everything you need to enjoy quantum programming in Python and Q# from this book whether or not you plunge into these deep dives. Taking a deep dive can be fun and enlightening, but if deep dives aren't your thing, that's OK too; it's perfectly fine to snorkel.

1.4 WHY DOES QUANTUM COMPUTING MATTER?

Computing technology advances at a truly stunning pace. Three decades ago, the 80486 processor would allow users to execute 50 MIPS (million instructions per second), but today, small computers like the Raspberry Pi can reach 5,000 MIPS, while desktop processors can easily reach 50,000 to 300,000 MIPS. If you have an exceptionally difficult computational problem you'd like to solve, a very reasonable strategy is to simply wait for the next generation of processors to make your life easier, your videos stream faster, and your games more colorful.

For many problems that we care about, however, we're not so lucky. We might hope that getting a CPU that's twice as fast lets us solve problems twice as big, but just as with so much in life, more is different. Suppose we want to sort a list of 10 million numbers and find that it takes about 1 second. If we later want to sort a list of 1 billion numbers in one second, we'll need a CPU that's 130 times faster, not just 100 times. Some problems make this even worse: for some problems in graphics, going from 10 million to 1 billion points would take 13,000 times longer.

Problems as widely varied as routing traffic in a city and predicting chemical reactions get more difficult *much* more quickly still. If quantum computing were simply a computer that runs 1,000 times as fast, we would barely make a dent in the daunting computational challenges that we want to solve. Thankfully, quantum computers are much more interesting than that. In fact, we expect that quantum computers will likely be much *slower* than classical computers, but that the resources required to solve many

problems *scale* differently, such that if we look at the right kinds of problems we can break through "more is different." At the same time, quantum computers aren't a magic bullet, in that some problems will remain hard. For example, while it is likely that quantum computers can help us immensely with predicting chemical reactions, it is possible that they won't be of much help with other difficult problems.

Investigating exactly which problems we can obtain such an advantage in and developing quantum algorithms to do so has been a large focus of quantum computing research. Understanding where we can find advantages by using quantum computers has recently become a significant focus for quantum software development in industry as well. Software platforms such as the Quantum Development Kit make it easier to develop software for solving problems on quantum computers, and in turn to assess how easy different problems are to solve using quantum resources.

Up until this point, it has been very hard to assess quantum approaches in this way, as doing so required extensive mathematical skill to write out quantum algorithms and to understand all of the subtleties of quantum mechanics. Now, industry has started developing software packages and even new languages and frameworks to help connect developers to quantum computing. By leveraging the entire Quantum Development Kit, we can abstract away most of the mathematical complexities of quantum computing and help people get to actually *understanding* and *using* quantum computers. The tools and techniques taught in this book allow developers to explore and understand what writing programs for this new hardware platform will be like.

Put differently, quantum computing is not going away, so understanding what scale of what problems we can solve with them matters quite a lot indeed! There are already many governments and CEOs that are convinced that quantum computing will be the next big thing in computing. Some people care about quantum attacks on cryptography, some want their own quantum computer next year. Independent of the whether or not a quantum "revolution" happens, quantum computing has and will factor heavily into decisions about how to develop computing resources over the next several decades.

1.4.1 Decisions that are strongly impacted by quantum computing.

- What assumptions are reasonable in information security?
- What skills are useful in degree programs?
- How to evaluate the market for computing solutions?

For those of us working in tech or related fields, we increasingly must make or provide input into these decisions. We have a responsibility to understand what quantum computing is, and perhaps more importantly still, what it is not. That way, we will be best prepared to step up and contribute to these new efforts and decisions.

All that aside, another reason that quantum computing is such a fascinating topic is that it is both similar to and very different from classical computing. Understanding both the similarities and differences between classical and quantum computing helps us to understand what is fundamental about computing in general. Both classical and quantum computation arise from different descriptions of physical laws such that

understanding computation can help us understand the universe itself in a new way.

What's absolutely critical, though, is that there is no one right or even best reason to be interested in quantum computing. Whether you're reading this because you want to have a nice nine-to-five job programing quantum computers or because you want to develop the skills you need to contribute to quantum computing research, you'll learn something interesting to you along the way.

Further reading

If you are interested in exploring the more philosophical or foundational implications of quantum computing, much of this is covered in computational complexity.

Some good resources along these lines are:

- The Complexity Zoo (https://complexityzoo.uwaterloo.ca/Complexity Zoo),
- Complexity Theory: A modern approach (ISBN-13: 9780521424264), or
- Quantum Computing Since Democritus (ISBN-13: 9780521199568).

1.5 WHAT CAN QUANTUM COMPUTERS DO?

As quantum programmers we would like to know:

If we have a particular problem, how do we know it makes sense to solve it with a quantum computer?

We are still learning about the exact extent of what quantum computers are capable of, and thus we dont have any concrete rules to answer this question yet. So far, we have found some examples of problems where quantum computers offer significant advantages over the best known classical approaches. In each case, the quantum algorithms that have been found to solve these problems exploit quantum effects to achieve the advantages, sometimes referred to as a quantum advantage.

1.5.1 Some useful quantum algorithms

- Grover's algorithm (Chapter 10): searches a list of N items in \sqrt{N} steps.
- Shor's algorithm (Chapter 11): quickly factors large integers, such as those used by cryptography to protect private data.

Though we'll see several more in this book, both Grover's and Shor's are good examples of how quantum algorithms work: each uses quantum effects to separate correct answers to computational problems from invalid solutions. One way to realize a quantum advantage is to find ways of using quantum effects to separate correct and incorrect solutions to classical problems.

What are quantum advantages?

Grover's and Shor's algorithms illustrate two distinct kinds of quantum advantage. Factoring integers might be easier classically than we suspect, as we haven't been able to prove that factoring is difficult. The evidence that we use to conjecture that factoring is hard classically is largely derived from experience, in that many people have tried very hard to factor integers quickly, but haven't succeeded. On the other hand, Grover's algorithm is provably faster than any classical algorithm, but uses a fundamentally different kind of input.

©Manning Publications Co. To comment go to liveBook

Other quantum advantages might derive from problems in which we must simulate quantum dynamics. Quantum effects such as interference are quite useful in simulating other quantum effects.

Finding a *provable* advantage for a practical problem is an active area of research in quantum computing. That said, quantum computers can be a powerful resource for solving problems even if we can't necessarily prove that there will never be a better classical algorithm. After all, Shor's algorithm already challenges the assumptions underlying large swaths of classical information security — a mathematical proof is in that sense made necessary only by the fact that we haven't yet built a quantum computer in practice.

Quantum computers also offer significant benefits to simulating properties of quantum systems, opening up applications to quantum chemistry and materials science. For instance, quantum computers could make it much easier to learn about the ground state energies of chemical systems. These ground state energies then provide insight into reaction rates, electronic configurations, thermodynamic properties, and other properties of immense interest in chemistry.

Along the way to developing these applications, we have also seen significant advantages in spin-off technologies such as quantum key distribution and quantum metrology, as we will see in the next few chapters. In learning to control and understand quantum devices for the purpose of computing, we also have learned valuable techniques for imaging, parameter estimation, security, and more. While these are not applications for quantum computing in a strict sense, they go a long way to showing the values of *thinking* in terms of quantum computation.

Of course, new applications of quantum computers are much easier to discover when we have a concrete understanding of how quantum algorithms work and how to build new algorithms from basic principles. From that perspective, quantum programming is a great resource to learn how to discover entirely new applications.

1.6 WHAT IS A QUANTUM COMPUTER?

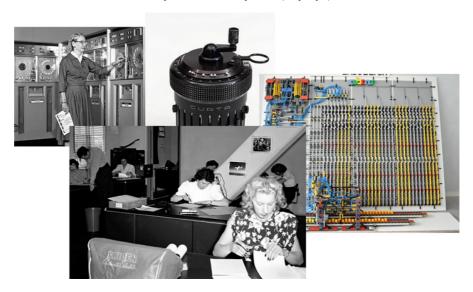
Let's talk a bit about what actually makes up a quantum computer. To facilitate this, let's briefly talk about what the term "computer" means. In a very broad sense, a computer is a device that takes data as input and does some sort of operations on that data.

COMPUTER

A computer is a device that takes data as input and does some sort of operations on that data.

There are many examples of what we have called a computer, see for some examples.

Figure 1.1. Several examples of different kinds of computers, including a mainframe operated by Rear Admiral Hopper, a room of humans working to solve flight calculations, a mechanical calculator, and a LEGO-based Turing machine. Each computer can be described by the same mathematical model as computers like cell phones, laptops, and servers.



When we say the word "computer" in conversation, though, we tend to mean something more specific. Often, we think of a computer as an electronic device like the one we are currently writing this book on (or that you might be using to read this book!). For any resource up to this point that we have made a computer out of, we can model it with classical physics — that is, in terms of Newtons's laws of motion, Newtonian gravity, and electromagnetism.

Following this perspective, we will refer to computers that are described using classical physics as classical computers. This will help us tell apart the kinds of computers we're used to (e.g.: laptops, phones, bread machines, houses, cars, and pacemakers) from the computers that we're learning about in this book.

Specifically, in this book, we'll be learning about quantum computers. By the way we have formulated the definition for classical computers, if we just replace the term *classical physics* with *quantum physics* we have a suitable definition for what a quantum computer is!

QUANTUM COMPUTER

A quantum computer is a device that takes data as input and does some sort of operations on that data, which requires the use of quantum physics to describe this process..

The distinction between classical and quantum computers is precisely that between classical and quantum physics. We will get in to this more later in the book, but the primary difference is one of scale: our everyday experience is largely with objects that are large enough and hot enough that even though quantum effects still exist, they

don't do much on average. Quantum physics still remains true, even at the scale of everyday objects like coffee mugs, bags of flour, and baseball bats, but we can do a very good job of describing how these objects interact using physical laws like Newton's laws of motion.

DEEP DIVE: What happened to relativity?

Quantum physics applies to objects that are very small and very cold or well-isolated. Similarly, another branch of physics called *relativity* describes objects that are large enough for gravity to play an important role, or that are moving very fast — near the speed of light. We have already discussed classical physics, which could also be said to describe things that are **neither** quantum mechanical nor relativistic. So far we have primarily been comparing classical and quantum physics, raising the question: why aren't we concerned about relativity? Many computers use relativistic effects to perform computation; indeed, global positioning satellites depend critically on relativity.

That said, all of the computation that is implemented using relativistic effects can also be described using purely classical models of computing such as Turing machines. By contrast, quantum computation cannot be described as faster classical computation, but requires a different mathematical model. There has not yet been a proposal for a "gravitic computer" that uses relativity to realize a different model of computation, so we're safe to set relativity aside in this book.

Quantum computing is the art of using small and well-isolated devices to usefully transform our data in ways that cannot be described in terms of classical physics alone. We will see in the next chapter, for instance, that we can generate random numbers on a quantum device by using the idea of *rotating* between different states. One way to build quantum devices is to use small classical computers such as digital signal processors (DSPs) to control properties of exotic materials.

PHYSICS AND QUANTUM COMPUTING The exotic materials used to build quantum computers have names that can sound intimidating, like "superconductors" and "topological insulators." We can take solace, though, from how we learn to understand and use classical computers. Modern processors are built using materials like semiconductors, but we can program classical computers without knowing what a semiconductor is. Similarly, the physics behind how we can use superconductors and topological insulators to build quantum computers is both a fascinating subject, but it's not required for us to learn how to program and use quantum devices.

Quantum operations are applied by sending in small amounts of microwave power and amplifying very small signals coming out of the quantum device.

Other qubit devices may differ in the details of how they are controlled, but what remains consistent is that all quantum devices are controlled from and read out by classical computers and control electronics of some kind. After all, we are ultimately interested in classical data, and so there must eventually be an interface with the classical world.

TIP

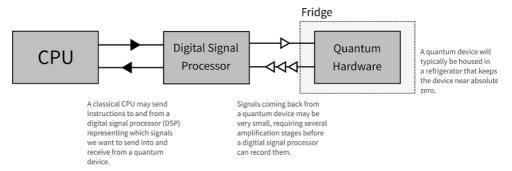
For most quantum devices, we need to keep them very cold and very well isolated, since quantum devices can be very susceptible to noise.

By applying quantum operations using embedded classical hardware, we can

©Manning Publications Co. To comment go to liveBook

manipulate and transform quantum data. The power of quantum computing then comes from carefully choosing which operations to apply in order to implement a useful transformation that solves a problem of interest.

Figure 1.2. An example of how a quantum device might interact with a classical computer through the use of a digital signal processor (DSP). The DSP sends low-power signals into the quantum device, and amplifies very low-power signals coming back to the device.



1.7 HOW WILL WE USE QUANTUM COMPUTERS?

It is important to understand both the potential and the limitations of quantum computers, especially given the hype surrounding quantum computation. Many of the misunderstandings underlying this hype stem from extrapolating analogies beyond where they make any sense — all analogies have their limits, and quantum computing is no different in that regard.

TIP

If you've ever seen descriptions of new results in quantum computing that read like "we can teleport cats that are in two places at once using the power of infinitely many parallel universes all working together to cure cancer," then you've seen the danger of extrapolating too far from where analogies are useful.

Indeed, any analogy or metaphor used to explain quantum concepts will be wrong if you dig deep enough. Simulating how a quantum program acts in practice can be a great way to help test and refine the understanding provided by analogies. Nonetheless, we will still leverage analogies in this book, as they can be quite helpful in providing intuition for how quantum computation works.

One especially common point of confusion regarding quantum computing is the way in which users will leverage quantum computers. We as a society now have a particular understanding of what a device called a *computer* does. A computer is something that you can use to run web applications, write documents, and run simulations to name a few common uses. In fact, classical computers are present in every aspect of our lives, making it it easy to take computers for granted. We don't always even notice what is and isn't a computer. Cory Doctorow made this observation by noting that "your car is a computer you sit inside of" (https://www.youtube.com/watch?v=iaf3Sl2r3jE).

Quantum computers, however, are likely to be much more special-purpose. Just as not all computation runs on graphical processing units (GPUs) or field-programmable gate arrays (FPGAs), we expect quantum computers to be somewhat pointless for some tasks.

IMPORTANT

Programming a quantum computer comes along with some restrictions, so classical computers will be preferable in cases where there's no particular quantum advantage to be found.

Classical computing will still be around and will be the main way we communicate and interact with each other as well as our quantum hardware. Even to get the classical computing resource to interface with the quantum devices we will also need in most cases a digital to analogue signal processor as shown in .

Moreover, quantum physics describes things at very small scales (both size and energy) that are well-isolated from their surroundings. This puts some hard limitations to what environments we can run a quantum computer in. One possible solution is to keep our quantum devices in cryogenic fridges, often near absolute 0 K (-459.67 °F, or -273.15 °C). While this is not a problem at all to achieve in a data center, maintaining a dilution refrigerator isn't really something that makes sense on a desktop, much less in a laptop or a cell phone. For this reason, it's very likely that quantum computers will, at least for quite a while after they first become commercially available, be used through the cloud.

Using quantum computers as a cloud service resembles other advances in specialized computing hardware. By centralizing exotic computing resources in data centers, it's possible to explore computing models that are difficult for all but the largest users to deploy on-premises. Just as high-speed and high-availability Internet connections have made cloud computing accessible for large numbers of users, you will be able to use quantum computers from the comfort of your favorite WiFi-blanketed beach, coffee shop, or even from a train as you watch majestic mountain ranges off in the distance.

1.7.1 Exotic cloud computing resources

- Specialized gaming hardware (PlayStation Now, Xbox One).
- Extremely low-latency high-performance computing (e.g. Infiniband) clusters for scientific problems.
- Massive GPU clusters.
- Reprogrammable hardware (e.g. Catapult/Brainwave).
- Tensor processing unit (TPU) clusters.
- High-permanence high-latency archival storage (e.g. Amazon Glacier).

1.8 WHAT CAN'T QUANTUM COMPUTERS DO?

Like other forms of specialized computing hardware, quantum computers won't be good at everything. For some problems, classical computers will simply be better suited to the task. In developing applications for quantum devices, it's helpful to note what tasks or problems are out of scope for quantum computing.

The short version is that we don't have any hard-and-fast rules to quickly decide between which tasks are best run on classical computers, and which tasks can take advantage of quantum computers. For example, the storage and bandwidth requirements for Big Data—style applications are very difficult to map onto quantum devices, where you may only have a relatively small quantum system. Current quantum computers can only record inputs of no more than a few dozen bits, a limitation that will become more relevant as quantum devices are used for more demanding tasks. Although we expect to eventually be able to build much larger quantum systems than we can now, classical computers will likely always be preferable for problems which require large amounts of input/output to solve.

Similarly, machine learning applications that depend heavily on random access to large sets of classical inputs are conceptually difficult to solve with quantum computing. That said, there may be other machine learning applications exist that map much more naturally onto quantum computation. Research efforts to find the best ways to apply quantum resources to solve machine learning tasks are still ongoing. In general, problems that have small input and output data sizes, but large amounts of computation to get from input to output are good candidates for quantum computers.

In light of these challenges, it might be tempting to conclude that quantum computers *always* excel at tasks which have small inputs and outputs, but have very intense computation between the two. Notions like "quantum parallelism" are popular in media, and quantum computers are sometimes even described as using parallel universes to compute.

NOTE

The concept of "parallel universes" is a great example of an analogy that can help make quantum concepts understandable, but that can lead to nonsense when taken to its extreme. It can be sometimes helpful to think of the different parts of a quantum computation as being in different universes that can't affect each other, but this description makes it harder to think about some of the effects we will learn in this book, such as interference. When taken too far, the "parallel universes analogy" also lends itself to thinking of quantum computing in ways that are closer to a particularly pulpy and fun episode of a sci-fi show like Star Trek than to reality.

What this fails to communicate, however, is that it isn't always obvious how to use quantum effects to extract useful answers from a quantum device, even if it appears to contain the desired output. For instance, one way to factor an integer classically is to list each *potential* factor, and to check if it's actually a factor or not.

1.8.1 Factoring N classically.

- Let i = 2.
- Check if the remainder of N / i is zero.
 - o If so, return that i factors N.
 - o If not, increment i and loop.

We can speed this classical algorithm up by using a large number of different classical computers, one for each potential factor that we want to try. That is, this problem can

be easily parallelized. A quantum computer can try each potential factor within the same device, but as it turns out, this isn't *yet* enough to factor integers faster than the classical approach above. If you run this on a quantum computer, the output will be one of the potential factors chosen at random. The actual correct factors will occur with probability about $1/\sqrt{N}$, which is no better than the classical algorithm above.

As we'll see in Chapter 11, though, we can improve this by using other quantum effects, however, to factor integers with a quantum computer faster than the best-known classical factoring algorithms. Much of the heavy lifting done by Shor's algorithm is to make sure that the probability of measuring a correct factor at the end is much larger than measuring an incorrect factor. Cancelling out incorrect answers in this way is where much of the art of quantum programming comes in; it's not easy or even possible to do for all problems we might want to solve.

To understand more concretely what quantum computers can and can't do, and how to do cool things with quantum computers despite these challenges, it's helpful to take a more concrete approach. Thus, let's consider what a quantum program even **is**, so that we can start writing our own.

1.9 WHAT IS A PROGRAM?

Throughout this book, we will often find it useful to explain a quantum concept by first re-examining the analogous classical concept. In particular, let's take a step back and examine what a classical program is.

PROGRAM

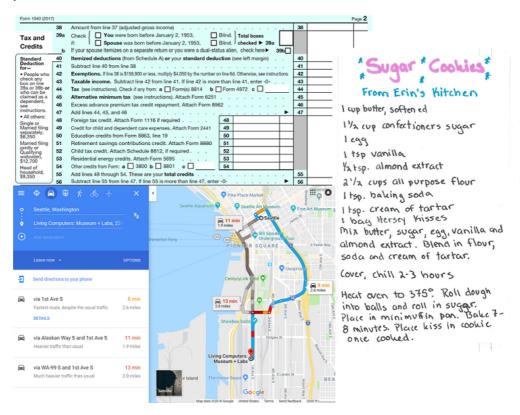
A program is a sequence of instructions that can be interpreted by a classical computer to perform a desired task.

Examples of classical programs

- Tax forms
- Map directions
- Recipes
- Python scripts

We can write classical programs to break down a wide variety of different tasks for interpretation by all sorts of different computers.

Figure 1.3. Examples of classical programs. Tax forms, map directions, and recipes are all examples in which a sequence of instructions is interpreted by a classical computer such as a person.



Let's take a look at what a simple "hello, world" program might look like in Python:

```
>>> def hello():
... print("Hello, world!")
...
>>> hello()
Hello, world!
```

At its most basic, this program can be thought of as a sequence of instructions given to the Python interpreter, which then executes each instruction in turn to accomplish some effect — in this case, printing a message to the screen.

We can make this way of thinking more formal by using the dis module provided with Python to *disassemble* hello() into a sequence of instructions:

©Manning Publications Co. To comment go to liveBook

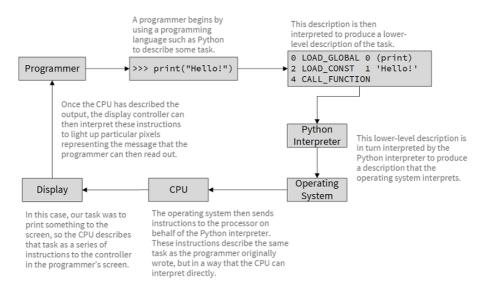
```
6 POP_TOP
8 LOAD_CONST Ø (None)
10 RETURN_VALUE
```

NOTE

You may get different output on your system, depending on what version of Python you're using.

Each line consists of a single instruction that is passed to the Python virtual machine; for instance, the LOAD GLOBAL instruction is used to look up the definition of the print function. The print function then called bv the CALL FUNCTION instruction. The Python code that we wrote above was compiled by the interpreter to produce this sequence of instructions. In turn, the Python virtual machine executes our program by calling instructions provided by the operating system and the CPU.

Figure 1.4. An example of how a classical computing task is repeatedly described and interpreted.



At each level, we have a *description* of some task that is then *interpreted* by some other program or piece of hardware to accomplish some goal. This constant interplay between description and interpretation motivates calling Python, C, and other such programming tools *languages*, emphasizing that programming is ultimately an act of communication.

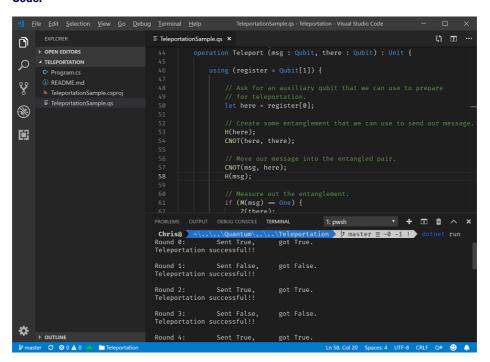
In the example of using Python to print "Hello, world!," we are effectively communicating with Guido von Rossum, the founding designer of the Python language. Guido then effectively communicates on our behalf with the designers of the operating system that we are using. These designers in turn communicate on our behalf with Intel, AMD, ARM, or whomever has designed the CPU that we are using, and so

forth. As with any other use of language to communicate, our choice of programming language affects how we think and reason about programming. When we choose a programming language, the different features of that language and the syntax used to express those features mean that some ideas are more easily expressed than others.

1.10 WHAT IS A QUANTUM PROGRAM?

Like classical programs, quantum programs consist of sequences of instructions that are interpreted by classical computers to perform a particular task. The difference, however, is that in a quantum program, the task we wish to accomplish involves controlling a quantum system to perform a computation.

Figure 1.5. Writing a quantum program with the Quantum Development Kit and Visual Studio Code.

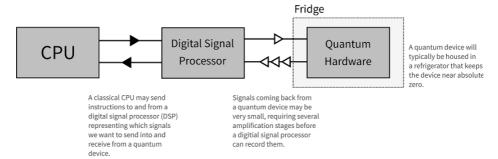


The instructions available to classical and quantum programs differ according to this difference in tasks. For instance, a classical program may describe a task such as loading some cat pictures from the Internet in terms of instructions to a networking stack, and eventually in terms of assembly instructions such as mov (move). By contrast, quantum languages like Q# allow programmers to express quantum tasks in terms of instructions like M (measure).

Figure 1.6. An example of how a quantum device might interact with a classical computer through the use of a digital signal processor (DSP). The DSP sends low-power signals into the

©Manning Publications Co. To comment go to <u>liveBook</u>

quantum device, and amplifies very low-power signals coming back to the device.



When run using quantum hardware, these programs may instruct a digital signal processor such as that shown in to send microwaves, radio waves, or lasers into a quantum device, and to amplify signals coming out of the device.

If we are to achieve a different ends, however, it makes sense for us to use a language that reflects what we wish to communicate! We have many different classical programming languages for just this reason, as it doesn't make sense to use only one of C, Python, JavaScript, Haskell, Bash, T-SQL, or any of a whole multitude of other languages. Each language focuses on a subset of tasks that arise within classical programming, allowing us to choose a language that lets us express how we would like to communicate that task to the next level of interpreters.

Quantum programming is thus distinct from classical programming almost entirely in terms of what tasks are given special precedence and attention. On the other hand, quantum programs are still interpreted by classical hardware such as digital signal processors, so a quantum programmer writes quantum programs using classical computers and development environments.

Throughout the rest of this book, we will see many examples of the kinds of tasks that a quantum program is faced with solving or at least addressing, and what kinds of classical tools we can use to make quantum programming easier. We will build up the concepts you need to write quantum programs chapter by chapter, you can see a road map of how these concepts will build up in.

Domain areas like chemistry simulation, cryptography, and **Quantum Applications** machine learning are potential applications for quantum computers. There are a variety of quantum Libraries / Techniques development techniques and Q# libraries available that you can learn and leverage to build your own quantum programs. Just like learning any other programming language, when Q# Language learning Q# it is important to learn things like syntax, patterns, and types for Q#. No matter the hardware or **Intrinsic Operations** simulator you use, there are some intrinsic operations that vour quantum device can perform, like operations that Quantum To help build understanding as can generate superposition Simulators well as test programs, and entanglement that will be simulators for quantum used by your quantum Hardware hardware are a core piece of programs. the quantum development

Figure 1.7. This book will try to build up the concepts you need to write quantum programs.

1.11 SUMMARY

stack.

In this chapter you learned:

- Recognize the significance of quantum computing in modern society,
- Predict what kinds of problems a quantum computer may be good at solving,
- and recognize the similarities and differences between programming for a quantum computer vs. a classical computer.

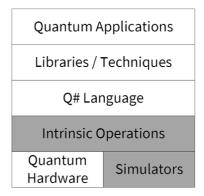
Qubits: The building blocks

This chapter covers:

- · Why random numbers are an important resource.
- What is a qubit?
- · What are the basic operations we can perform on a qubit?
- · How to program a quantum random number generator in Python.

In this chapter, we are going to start to get our feet wet with some quantum programming concepts. The main concept we will explore is the qubit, the quantum analogue of a classical bit. We use qubits as an abstraction or model to describe the new kinds of computing that are possible with quantum physics. To help learn about what qubits are and how we interact with them, we will use an example of how they are being used today: random number generation. While we can build up much more interesting devices from these qubits, the simple example of a quantum random number generator (QRNG) will be a good way to get familiar with the qubit!

Figure 2.1. In this Chapter, we will be covering topics in simulating quantum hardware and intrinsic operations for the quantum devices.



2.1 WHY DO WE NEED RANDOM NUMBERS?

Humans like certainty. We like it when we press a key on our keyboard and it does the same thing every time. However, there are some contexts in which we do want randomness.

2.1.1 Things some humans like to use randomness for

- Playing games
- Simulating complex systems (e.g.: stock market)
- Picking secure secrets (e.g.: passwords and cryptographic keys)

In all of these situations where we want randomness, we can describe the "chances" for each outcome. Being random events means that describing the chances is all we can say about the situation until the die is cast (or the coin is flipped or the password is reused). When we describe the "chances" of each example, we say things like:

2.1.2 Statements about probability

- if I roll this die, then I will see a six with probability 1 out of 6,
- if I flip this coin, then I will see a heads with probability 1 out of 2.

We can also describe cases where the probabilities aren't the same for every outcome we could measure. In Wheel of FortuneTM (), the probability that **if** we spin the wheel **then** we will get a \$1,000,000 prize is much smaller than the probability that **if** we spin the wheel, **then** we will go bankrupt.

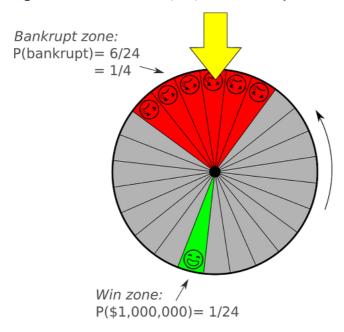


Figure 2.2. Probabilities of \$1,000,000 and Bankrupt in Wheel of Fortune™

Like in game shows, there are also may contexts in computer science where randomness is critical, especially when security is required. If we want to keep some information private, then cryptography lets us do so by combining our data with random numbers in different ways. If our random number generator isn't very good — that is to say if an attacker can predict what numbers we use to protect our private data — then cryptography doesn't help us very much. We can also imagine using a poor random number generator to run a raffle or a lottery; an attacker who figures out how our random numbers are generated can take us straight to the bank.

WHAT ARE THE ODDS

You can lose a lot of money by using random numbers that your adversaries can predict. Just ask the producers of Press Your Luck!, a popular 1980s game show. A contestant found that he could predict where their new electronic "wheel" would land, letting him win more than \$250,000 in today's money. Read more at: https://priceonomics.com/the-man-who-got-no-whammies/.

As it turns out, quantum mechanics lets us build some really unique sources of randomness. If we build them right, the randomness of our results is guaranteed by *physics*, not an assumption about how long it would take for a computer to solve a difficult problem. This means that for a hacker or adversary they would have to break the laws of physics to break the security! Now, this does not mean that you should use quantum random numbers for everything, humans are still the weakest link in security infrastructure

21

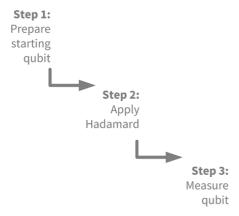
DEEP DIVE: Computational Security and Information Theoretic Security

Some ways of protecting private information rely on assumptions about what problems are easy or hard for an attacker to solve. For instance, the RSA algorithm is a commonly used encryption algorithm, and is based on the difficulty of finding prime factors for large numbers. RSA is used on the web and in other contexts to protect user data, under the assumption that adversaries can't easily factor very large numbers. So far, this has proven to be a rather good assumption, but it is entirely possible that a new factoring algorithm is discovered, undermining the security of RSA. New models of computation like quantum computing also change how reasonable or unreasonable it is to make computational assumptions like "factoring is hard." As we'll see in Chapter 11, a quantum algorithm known as *Shor's algorithm* allows for solving some kinds of cryptographic problems much faster than classical computers, challenging the assumptions that are commonly used to promise computational security.

By contrast, if an adversary can only ever randomly guess at secrets, even with very large amounts of computing power, then a security system provides much better guarantees about its ability to protect private information. Such systems are said to be *informationally secure*. In the next chapter, we will see that generating random numbers in a hard-to-predict fashion allows us to implement an informationally secure procedure called a one-time pad.

This gives us some confidence that we can use them for vital tasks, such as to protect private data, run lotteries, and to play Dungeons and Dragons[™]. Simulating how quantum random number generators work lets us learn many of the basic concepts underlying quantum mechanics, so let's jump right in and get started!

Figure 2.3. Quantum random number generator algorithm.



One great way to get started is to look at an example of a quantum program that generates random numbers. Let's call it a quantum random number generator or QRNG for short. Don't worry if the algorithm doesn't make a lot of sense right now, we'll explain the different pieces as we go through the rest of the chapter.

2.1.3 Quantum random number generator algorithm

- Ask the quantum device to allocate a qubit.
- Apply a Hadamard instruction to our qubit.
- Measure our qubit and return the result.

©Manning Publications Co. To comment go to <u>liveBook</u>

In the rest of the chapter, we'll develop a Python class QuantumDevice to let us write programs that implement algorithms like the one above. Once we have a QuantumDevice class, we'll be able to write out QRNG as a Python program similar to classical programs that you're used to.

NOTE

Please see Appendix A for instructions on how to setup Python on your device to run quantum programs.

Listing 2.1. qrng.py: A quantum program that generates random numbers. Note that

this sample will not run until you have written the simulator in this chapter 😥

def qrng(device : QuantumDevice) -> bool: **2 3 4** with device.using qubit() as q:

return q.measure()

- 1)Our quantum programs are written just like the classical programs that you're used to. In this case, we're using Python, so our quantum program is a Python function qrng that implements a quantum random number generator.
- Quantum programs work by asking quantum computing hardware for qubits, quantum analogues of bits that we can use to perform computations.
- 3 Once we have a qubit, we can issue instructions to that qubit. Similarly to assembly languages, these instructions are often denoted by short abbreviations; we'll see what h() stands for later on in this chapter.
- 4 To get data back from our qubits, we can measure them. In this case, half of the time, our measurement will return True, and the other half of the time, we'll get back False.

That's it!

Four steps and we've just created our first quantum program. This QRNG returns true or false. In Python terms, this means that you get a 1 or a 0 each time you run grng. It's not a very sophisticated random number generator but the number it returns is truly random.

To run the grng program above, we'll need to give our function a QuantumDevice that can give us access to qubits and that implements the different instructions we can send to qubits. Though we need only one qubit, to start, we're going to build our own quantum computer simulator. Existing hardware could be used for this modest task, but what we will look at later will be beyond the scope of available hardware. It will run locally on a laptop or desktop and behave in the same way as actual quantum hardware. The simplest kind of device that we can consider is a simulator that runs locally on a laptop or desktop, and that behaves in the same way as actual quantum hardware. Throughout the rest of the chapter, we'll build up the different pieces we need to write our own simulator and to run grng.

2.2 WHAT ARE CLASSICAL BITS?

When learning about the concepts of quantum mechanics, it can often be helpful to step back and re-examine *classical* concepts in a way that makes it easier to make the connection to how that concept is expressed in quantum computing. With that in mind, let's step back and take another look at what *bits* are.

Suppose that you'd like to send your dear friend Eve an important message, such as . How can we represent our message in a way that it can be easily sent?

We might start by making a list of every letter and symbol that that we could hope to use in writing down messages. Thankfully, the Unicode Consortium (https://unicode.org/) has already done this for us, and has assigned a *code* to a very wide variety of the characters we use across the world to communicate with each other. For instance, I is assigned the code 0049, while is given the code A66E, is denoted by 2E0E, and by 1F496. These codes may not seem too helpful at first glance, but they're useful recipes for how to send each symbol as a message. If we know how to send two messages, let's call them 0 and 1, these recipes let us build up more complicated messages like a, and as sequences of 0 and 1 messages:

0	0000	8	1000
1	0001	9	1001
2	0010	A	1010
3	0011	В	1011
4	0100	C	1100
5	0101	D	1101
6	0110	E	1110
7	0111	F	1111

Now we can send whatever we want if we know how to send just two messages to Eve, a 0 and a 1 message. Using these recipes, our message of "" becomes 0001 1111 0100 1001 0110 or unicode 1F496.

TIP

Don't send 0001 1111 0100 1001 0100 by mistake, or Eve will get a ♥ from you!

We call each of the two messages "0" and "1" a bit.

NOTE

To distinguish bits from the quantum bits that we'll see throughout the rest of the book, we'll often emphasize that we're talking about classical bits.

When we use the word bit, we generally mean one of two things:

• Any physical system that can be completely described by answering one true/false question.

• The information stored in such a physical system.

For example, padlocks, light switches, transistors, left or right spin on a curveball, and wine in wine glasses can all be thought of as bits, as we can use all of them to send or record messages:

Table 2.1. Table Examples of bits

Label	Padlock	Light Switch	Transistor	Wine Glass	Baseball
0	Unlocked	Off	Low voltage	Has white wine	Rotating to the left
1	Locked	On	High voltage	Has red wine	Rotating to the right

These examples are all bits, because we can fully describe them to someone else by answering a single true/false question. Put differently, each example lets us send either a 0 or a 1 message. Like all conceptual models, a "bit" has its limitations — how would we describe a rosé wine, for instance?

That said, a "bit" is a useful tool because we can describe ways of interacting with bits that are independent of how we actually build the bit.

2.2.1 What Can We Do With Classical Bits?

Now that we have a way of describing and sending classical information, what can we do to process and modify our information? We describe the ways that we can process information in terms of *operations*, which we define as the ways of describing how a model can be changed or acted upon.

To visualize the NOT operation, let's imagine labeling two points as "0" and "1".

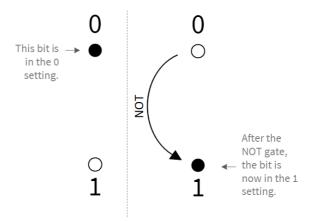
Figure 2.4. We depict a classical bit as a black dot in either the 0 or 1 position.



The NOT operation is then any transformation which turns "0" bits into "1" bits and vice versa. In classical storage devices like hard drives (and even floppy disks!), a NOT gate flips the magnetic field that stores our bit value. We can think of NOT as

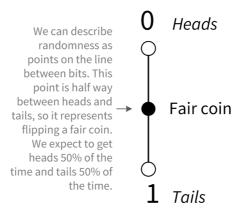
implementing a 180° rotation between the "0" and "1" points we drew above.

Figure 2.5. The classical NOT operation flips a 0 bit to a 1 bit and vice versa.



Visualizing classical bits this way also lets us extend our notion of bits slightly to include a way to describe *random* bits (which will be helpful later). If we have a *fair coin* (that is, a coin that lands heads "half" of the time and "tails" the other half), then it wouldn't be correct to call that coin a "0" or a "1". We only know what bit value our coin bit has if we set it with a particular side face up on a surface, or we can flip it for a random bit value. Every time we flip a coin, we know that eventually it will land and we will get a heads or tails. Whether it lands heads or tails is governed by a probability called the *bias* of the coin. We have to pick a side of the coin to describe the bias, which is easy to phrase as a question like this: What is the probability the coin will land heads? Thus a fair coin would have a bias of 50% because it lands with the value heads half of the time, which is mapped to the bit value 0 in .

Figure 2.6. Extending the concept of a bit to describe a coin, which has a probability of being found in either "0" or "1" each time it is flipped.



Using this visualization, we can take our previous two dots indicating the bit values "0" and "1" and connect them with a line on which we can plot the bias of our coin. It becomes easier to see that a NOT operation (which still works on our new probabilistic bit) doesn't do anything to a fair coin. If "0" and "1" occur with the same probability, then it doesn't matter if we rotate a "0" to a "1" and a "1" to a "0", we'll still wind up with "0" and "1" having the same probability.

What if our bias is not in the middle? If we know that someone is trying to cheat by using a weighted or modified coin that almost always lands on "heads", we could say the bias of the coin is 90% and could plot it on our line bit by drawing a point much closer to "0" than to "1."

STATE

We say that the point on a line where one would draw each classical bit is the state of that bit.

Let's consider a scenario:

Say I want to send you a bunch of bits stored using padlocks, what is the cheapest way I could do so?

One approach is to mail a box containing many padlocks that are either open or closed and hope that they arrive in the same state that I sent them in. On the other hand, we can both agree that all padlocks start off initially in the "0" (unlocked) state, and I can send you instructions on which padlocks to close. This way, you can go buy your own padlocks, and I only need to send a **description** of how to prepare those padlocks using classical NOT gates. Sending a piece of paper or even just an email is way cheaper than mailing a box of padlocks!

This illustrates a principle we will rely on throughout the book: The state of a physical system can also be described in terms of instructions for how prepare

that state. Thus, the operations allowed on a physical system also define what states are possible.

Though it may sound completely trivial, there is one more thing that we can do with classical bits that will turn out to be critical to how we understand quantum computing: we can look at them. If I look at a padlock and conclude that "aha! that padlock is unlocked~!," then I can now think of my brain as a particularly squishy kind of bit. The "0" message is stored in my brain by my thinking "aha! that padlock is unlocked~!," while a "1" message might have been stored by my thinking "ah, well, that padlock is locked ..." In effect, by looking at a classical bit, I have *copied* it into my brain. We say that the act of *measuring* the classical bit copies that bit.

More generally, modern life is built all around the ease with which we copy classical bits by looking at them. We copy classical bits with truly reckless abandon, measuring many billions of classical bits every second that we copy data from our video game consoles to our TVs.

On the other hand, if a bit is stored as a coin, then the process of measuring involves flipping it. Now measuring doesn't quite copy the coin, as I might get a different measurement result the next time I flip. If I only have one measurement of a coin, I can't conclude what the probability of getting a heads or tails was. We didn't have this ambiguity with our padlock bits, because we knew the state of our padlocks was either "0" or "1". If I measured a padlock and found it to be in the "0" state, I would know that unless I did something to the padlock, it would always be in the "0" state.

The situation isn't precisely the same in quantum computing, as we'll see later on in the chapter. While measuring classical information is cheap enough that we complain about precisely how many billions of bits a \$5 cable can let us measure, we will have to be much more careful with how we approach quantum measurements.

2.2.2 Abstractions are our friend

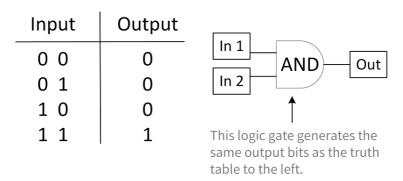
Regardless of how we physically build a bit, we can thankfully represent them in the same way in both math and in code. For instance, Python provides the bool type (short for Boolean, in honor of the logician George Boole), which has two valid values: True and False. We can then represent transformations on bits such as NOT and OR as operations acting on bool variables. Importantly, we can specify a classical operation by describing how that operation transforms each possible input, which is often called a *truth table*.

TRUTH TABLE

A table describing the output of a classical operation for every possible combination of inputs is called that operation's truth table.

Figure 2.7. Truth table for the logical operation AND.

Truth tables are one way to show what happens to classical bits in functions or logical circuits.



For example, we can find the truth table for the NAND (short for NOT-AND) operation in Python by iterating over combinations of True and False:

Listing 2.2. Using python to print out a truth table for NAND >>> from itertools import product >>> for inputs in product([False, True], repeat=2): output = not (inputs[0] and inputs[1]) print(f"{inputs[0]}\t{inputs[1]}\t->\t{output}") False False -> True False True -> True True False True -> False

TRUTH TABLES ALL THE WAY DOWN

Describing an operation as a truth table holds for more complicated operations as well; in principle even an operation like addition between two 64-bit integers can be written as a truth table. This isn't very practical, though, as a truth table for two 64-bit inputs would have 2^{128} \approx \times 10^{38} entries, and would take 10^{40} bits to write down. By comparison, recent estimates put the size of the entire Internet at closer to 10^{27} bits.

Much of the art of classical logic and hardware design is making *circuits* which can provide very compact representations of classical operations, rather than relying on potentially massive truth tables. In quantum computing we use the name *unitary operators* for similar truth tables for our quantum bits, which we will expand on as we go along.

In summary,

- Classical bits are physical systems which can be in one of two different *states*,
- Classical bits can be manipulated through *operations* to process information,
- The act of measuring a classical bit makes a copy of the information contained in the state.

NOTE

In the next section, you'll use linear algebra to learn about qubits, the basic unit of information in a quantum computer. If you need a refresher on linear algebra, this would be a great time to take a detour to visit Appendix B. We'll be right here when you get back!

2.3 QUBITS: STATES AND OPERATIONS

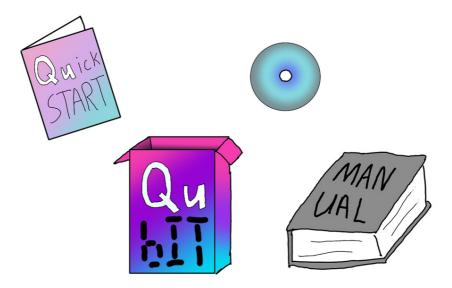
Just like classical bits are the most basic unit of information in a classical computer, **qubits** are the basic unit of information in a quantum computer. They can be physically implemented by systems that have 2 states, just like classical bits, but that behave according to the laws of quantum mechanics, which allows for some behaviors that classical bits are not capable of. Let's treat qubits like you would any other fun and new computer part: plug it in and see what happens!

EXAMPLE 2.3. SCENARIO

Say you just got a brand new computer part for your desktop. You have a thick manual, a quick start guide, and a driver CD. What steps do you take to figure out how to set it up correctly and verify that it works ok?

The easiest approach would to be to just plug it in and hope it works! But some very hard working technical writers and devs worked to put together that quick start guide and system requirements... We will treat the rest of this chapter as a sort of quick start guide to qubits, so let's see if they work!

Figure 2.8. Our brand new quantum development (the CD they provide is just a novelty coaster).



SIMULATED QUBITS

For almost all of this book, we won't be using actual qubits, but will be using classical simulations of qubits. This lets us learn how quantum computers will work, and to get started programming small instances of the kinds of problems that quantum computers can solve, even if we don't yet have access to the kinds of quantum hardware we'll need to solve practical problems.

The trouble with this is that simulating qubits on classical computers takes an exponential amount of classical resources in the number of qubits, such that the most powerful classical computing services can simulate up to about 40 qubits before having to simplify or reduce the types of quantum programs being run. For comparison, current commercial hardware maxes out at about 70 qubits at the time of this writing. Devices with that many qubits are extremely difficult to simulate with classical computers, but currently available devices are still too noisy to complete most useful computational tasks.

Imagine having to write a classical program with only 40 classical bits to work with! While 40 bits is quite small compared to the gigabytes that we are used to working with in classical programing, there are still some really interesting things we can do with only 40 qubits, and that help us prototype what an actual quantum advantage might look like.

2.3.1 State of the qubit

Looking through the listed system requirements for our quantum hardware, the first thing listed is *qubits*, followed by a bunch of stuff like "initial state: $|0\rangle$ ". This is more helpful than it might first seem, as our manual is telling us about how to describe and initialize our new qubits. We have used locks, baseballs, and other classical systems to represent our classical bit values of 0 or 1. There are many physical systems we can use to act as our qubit, and *states* are the "values" our qubit can have.

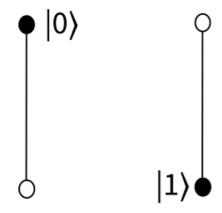
Similar to the 0 and 1 states of classical bits, we can write labels for quantum states.

The qubit states that are most similar to the classical "0" and "1" are $|0\rangle$ and $|1\rangle$, as we draw in . These are referred to as "ket 0" and "ket 1" respectively. With this in mind, the " $|0\rangle$ " in our manual tells us what state our qubits start off in when we pull them out of the box.

KET?

The term "ket" comes from a kind of whimsical naming that is seen in quantum computing owing its history to a particularly silly pun. As we'll see more of when we look at measurements, there's another kind of object called a bra that gets written like $\langle 0 |$. When you put a bra and a ket together, you get a pair of brackets $\langle \rangle$. The use of bras and kets to write out math for quantum mechanics is often called Dirac notation after Paul Dirac, who both invented the notation and the truly groan-worthy pun that we're now stuck with. We will see more of this style of whimsey throughout the book.

Figure 2.9. Bracket notation for qubits.



One thing to be mindful of, though, is that a state is a convenient model use to predict how a qubit behaves, not some inherent property of the qubit itself. This distinction becomes especially important when we consider measurement later in the chapter — as we will see, we cannot directly measure the state of a qubit.

IMPORTANT

In real systems, we will never be able to extract or perfectly learn the state of a qubit given a finite number of copies.

Don't worry if this doesn't all make sense yet, we'll see plenty of examples as we go through the book. What's important to keep in mind for now is that qubits aren't states.

If we want to simulate how a baseball moves once it's thrown, we might start by writing down it's current location, how fast it's going and in what direction, which way it's spinning, and so forth. That list of numbers helps us represent a baseball on a piece of paper or in a computer so that we can predict what that baseball will do, but we wouldn't say that the baseball is that list of numbers. To get our simulation started,

we'd have to take a baseball we're interested in and *measure* where it is, how fast it's going, and so forth.

We say that the full set of data we need to accurately simulate the behaviour of a baseball is the *state* of that baseball. Similarly, the state of a qubit is the full set of data we need to simulate it and to predict what results we'll get when we measure it. Just as we need to update the state of a baseball in our simulator as it goes along, we'll update the state of a qubit when we apply operations to it.

THE MAP IS NOT THE TERRITORY

One way to remember this subtle distinction is that a qubit is described by a state and but it is not true that a qubit is a state.

Where things get a little more subtle is that while we can measure a baseball without doing anything to it other than copying some classical information around, as we'll see throughout the rest of the book, we can't perfectly copy the quantum information stored in a qubit — when we measure a qubit, we have an effect on its state. This can be sometimes confusing, as we record the full state of a qubit when we simulate it, such that we could look at the memory in our simulator whenever we want. There's nothing we can do with actual qubits that lets us look at their state, so if we "cheat" by looking at the memory of a simulator, we won't be able to run our program on real hardware.

Put differently, while it can be useful for debugging our classical simulators as we are building them to look at states directly, we have to make sure we are only writing algorithms based on information we could plausibly learn from real hardware.

CHEATING WITH OUR EYES SHUT

As mentioned above, when we are using a simulator, the simulator must store the state of our qubits internally — this is why simulating quantum systems requires is so difficult. Every qubit could in principle be correlated with every other qubit, so we need exponential resources in general to write down the state in our simulator (we'll see more about this in Chapter 4). If we "cheat" by looking directly at the state stored by a simulator, then we can only ever run our program on a simulator, not on actual hardware. We'll see in later chapters how to cheat more safely, by using assertions and by making cheating unobservable.

2.3.2 The game of Operations

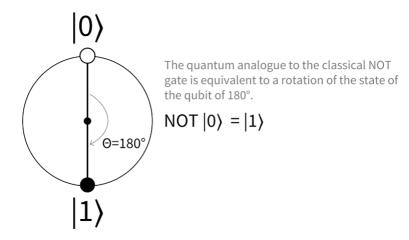
Now that we have names for these states, let's show how to represent the information they contain. With classical bits we can record the information contained in the bit at any time as simply a value on a line: 0 or 1. This worked because the only operations we could do consisted of flips (or 180° rotations) on this line. Quantum mechanics allows us to apply more kinds of operations to qubits, including rotations by less than 180°. That is, qubits differ from classical bits is in what operations we can do with them.

IMPORTANT

While operations on classical bits are logical operations that can be made by combining NOT, AND, and OR in different ways, quantum operations consist of rotations.

For instance, if we want to turn the state of a qubit from [0] to [1] and vice-versa, the quantum analogue of a NOT operation, we rotate the qubit clockwise by an angle of 180° .

Figure 2.10. A visualization of the quantum equivalent of a NOT operation operating on a qubit in the $|0\rangle$ state, leaving the qubit in the $|1\rangle$ state.



We have seen how rotation by 180° is the analogue to a NOT gate, but what other rotations can we do?

REVERSIBILITY

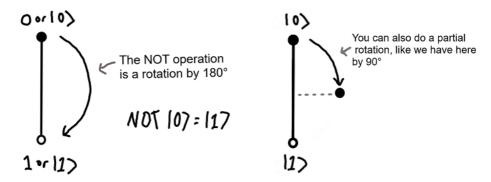
When we rotate a quantum state, we can always get back to the same state we started with by rotating backwards. This property, known as reversibility, turns out to be fundamental to quantum computing. With the exception of measurement, which we'll learn more about later in this chapter, all quantum operations must be reversible.

Not all of the classical operations that we're used to are reversible, though. Operations like AND and OR aren't reversible as they are typically written, so they cannot be implemented as quantum operations without a little bit more work. We'll see how to do this in Chapter 6 when we introduce the "uncompute" trick for expressing other classical operations as rotations.

On the other hand, classical operations like XOR can easily be made reversible, so we can write them out as rotations using a quantum operation called the "controlled NOT" operation, as we will see in Chapter 4.

If rotate a qubit in the $[|0\rangle]$ state clockwise by 90° instead of 180°, we get a quantum operation that we can think of as a "square root" of a NOT operation.

Figure 2.11. Rotating a state by less than 180 degrees.



In the same way as we earlier defined the square root \sqrt{x} of a number x as being a number y such that $y^2 = x$, we can define the square root of a quantum operation. If we apply a 90° rotation twice, we get back the NOT operation, so we can think of the 90° rotation as the square root of NOT.

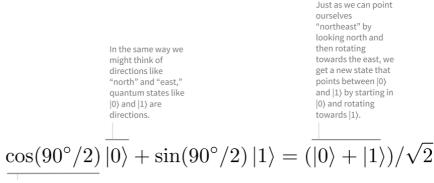
HALVES AND HALVE-NOTS

Every field has it stumbling blocks. Ask a graphics programmer whether positive y means "up" or "down," for instance. In quantum computing, the rich history and interdisciplinary nature of the field sometimes comes across as a double-edged sword in that each different way of thinking about quantum computing comes with its own conventions and notations.

One way this manifests is that it's really easy to make mistakes with where to put factors of two. In this book, we've chosen to follow the convention used by Microsoft's Q# language.

We now have a new state that is neither [0] nor [1], but an equal combination of them both. In precisely the same sense that we can describe "northeast" by adding the directions "north" and "east," we can write this new state as shown in .

Figure 2.12. Rotating a state by 90°.



We can rotate between states using the same math we use to describe rotations of map directions; we just have to watch out for factors of 2.

For example, if we want to rotate the |0⟩ state by 90°, we use the cosine and sine functions to find the new state.

|+⟩ AND |-⟩ We call this state the $|\rangle=(|0\rangle+|1\rangle)/\sqrt{2}|$ state (due to the sign between the terms). We say that the $|\rangle$ state is a _superposition_ of $|\mathbf{0}\rangle$ and $|\mathbf{1}\rangle$. If the rotation was by a -90° (anti-clockwise), then we call the resulting state stem $[|-\rangle=(|0\rangle-|1\rangle)/\sqrt{2}$ instead. Try writing out the rotations above using -90° to see that you get $|-\rangle$!

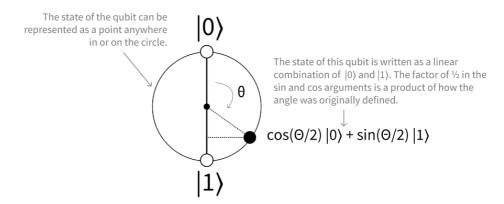
2.3.3 A mouthful of math

At first glance, something like $|+\rangle = (|0\rangle + |1\rangle) / \sqrt{2}$ would be terrible to have to say out loud, making it rather useless in conversation. In practice, however, quantum programmers often take some shortcuts when speaking outloud or sketching things out at the whiteboard.

For instance, the " $\sqrt{2}$ " part always has to be there, since vectors representing quantum states always have to be length one; that means we can sometimes be a little bit casual and write things like " $|+\rangle = |0\rangle + |1\rangle$," relying on our audience to remember to divide by $\sqrt{2}$. If we're giving a talk, or discussing quantum computing over some nice tea, we may say this as "ket plus is ket 0 plus ket 1," but the reuse of the word "plus" gets a little confusing without bras and kets to help. To emphasize verbally that addition allows us to represent superposition, we might say "the plus state is an equal superposition of zero and one" instead.

The state of a qubit can be represented as a point on a circle that has two labeled states on the poles: $|0\rangle$ and $|1\rangle$ More generally, we will picture rotations by arbitrary angles θ between qubit states as follows in.

Figure 2.13. A visualization of the state of a qubit.



Mathematically, we can write the state of any point on the circle that represents our qubit as $\cos(\theta/2)$ $|0\rangle + \sin(\theta/2)$ $|1\rangle$, where $|0\rangle$ and $|1\rangle$ are different ways of writing the vectors [[1], [0]] and [[0], [1]] respectively.

One way to think of ket notation is as giving names to vectors that we commonly use. When we write $|0\rangle = [[1], [0]]$, we're saying that [[1], [0]] is important enough that we name it after "0." Similarly, when we wrote that $|+\rangle = [[1], [1]] / \sqrt{2}$, we gave a name to the vector representation of a state that we will use all the time throughout the book.

Another way to say this would be a qubit is generally the *linear combination* of the vectors of $|0\rangle$ and $|1\rangle$ with coefficients that describe the angle that $|0\rangle$ would have to be rotated to get to the state. To write this in a way that is useful for programming, we can write out how rotating a state affects each of the $|0\rangle$ and $|1\rangle$ states:

Let's look at rotating $|0\rangle$ by an angle θ again, and see how we can write it out even when we don't know what θ is.

As before, we start by writing out the rotation using sines and cosines.

Sometimes, using matrix notation instead of Dirac notation (kets) is more helpful. We can use that $|0\rangle = [[1], [0]]$ here, since both are ways of writing down the same qubit state. $\frac{1}{\cos(\theta/2)} |0\rangle + \sin(\theta/2) |1\rangle = \cos(\theta/2) \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \sin(\theta/2) \begin{bmatrix} 0 \\ 1 \end{bmatrix}$

$$\frac{1}{\cos(\theta/2)} |0\rangle + \sin(\theta/2) |1\rangle = \cos(\theta/2) \begin{bmatrix} 1\\0 \end{bmatrix} + \sin(\theta/2) \begin{bmatrix} 0\\1 \end{bmatrix}$$

$$= \begin{bmatrix} \cos(\theta/2)\\\sin(\theta/2) \end{bmatrix}$$
Similarly, |1\rangle = [[0], [1]].

Once we've written each state using matrix notation, we can just add the corresponding elements together.

For instance, we get $\cos(\theta/2)$ for the first row, since $\cos(\theta/2)$ + 0 = $\cos(\theta/2)$.

TIP This is precisely the same as when we used a basis of vectors earlier to represent a linear function as a matrix.

There are other quantum operations that we will learn about in this book, but these are the easiest to visualize as rotations. Here is a table summarizing the states we have learned to create from these rotations:

Table 2.2. Table showing state labels, expansions in Dirac notation, and representations as vectors.

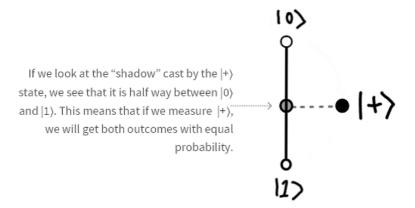
State label	Dirac notation	Vector representation
0)	0	[[1], [0]]
1>	1>	[[0], [1]]
+>	$(0\rangle + 1\rangle) / \sqrt{2}$	[[1 / \sqrt{2}], [1 / \sqrt{2}]]
->	$(0\rangle - 1\rangle) / \sqrt{2}$	$[[1/\sqrt{2}], [-1/\sqrt{2}]]$

2.3.4 Measuring Qubits

When we actually want to retrieve the information stored in a qubit, we need to measure it. Ideally, we would like a measurement device that would let us directly read out all the information about the state at once. As it turns out, this is not possible by the laws of quantum mechanics, as you'll see more about in Chapters 3 and 4. That said, measurement *can* allow us to learn information about the state relative to particular directions in the system. For instance, if we have a qubit in the $|0\rangle$ state and we look to see if it is in the $|0\rangle$ state, we'll always get that it is. On the other hand, if we have a qubit in the $|0\rangle$ state and we look to see if it is in the stem: $|0\rangle$ state, we'll get a "0"

outcome with 50% probability. This is because the stem:[$|\rangle$] overlaps equally with the $|0\rangle$ and $|1\rangle$ states, such that we'll get both outcomes with the same probability.

Figure 2.14. The $|+\rangle$ state overlaps equally with both $|0\rangle$ and $|1\rangle$, because the "shadow" it casts is exactly in the middle.

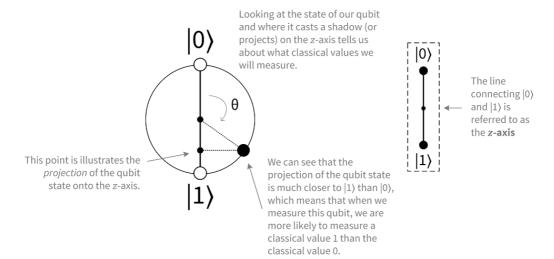


IMPORTANT

Measurement outcomes of qubits are classical bit values! Put differently, whether we measure a classical bit or a qubit, our result is always a classical bit.

Most of the time, we will choose to measure whether we have a $|0\rangle$ or a $|1\rangle$; that is, we'll want to measure along the line between the $|0\rangle$ and $|1\rangle$. For convenience, we give this axis a name, calling it the Z-axis. We can visually represent this by *projecting* our state vector onto the Z-axis (see), using the inner product we saw earlier. Think of shining a flashlight from where we draw the state of a qubit back onto the Z-axis; the probability for getting a 0 or 1 result is determined by shadow the state leaves on the Z-axis.

Figure 2.15. A visualization of a quantum measurement, which can be thought of as projecting the state along a particular direction.



DEEP DIVE: Why isn't measurement linear?

It may seem odd, after having made such a big deal of the linearity of quantum mechanics, that we immediately introduce measurement as being non-linear. If we're allowed non-linear operations like measurement, can we also implement other non-linear operations like cloning qubits?

The short version is that while everyone agrees on the math behind measurement, there's still a lot of philosophical discussion about the best way to understand why quantum measurement acts the way that it does. These discussions fall under the name of *quantum foundations*, and attempt to do more than simply understand what quantum mechanics is and what it predicts, by also understanding *why*. For the most part, foundations explores different ways to *interpret* quantum mechanics. In the same way that we can understand classical probability by considering counterintuitive thought experiments such as game show strategies or how casinos can win even from games that seem to lose money, quantum foundations develops new interpretations through small thought experiments that probe at different aspects of quantum mechanics. Thankfully, some of the results from quantum foundations can help us make sense of measurement.

In particular, one critical observation is that you can always make quantum measurements linear again by including the state of the measurement apparatus into your description; we'll see some of the mathematical tools needed to do so in Chapters 4 and 6. When taken to its extreme, this observation leads to interpretations such as the *many-worlds interpretation*. The many-worlds interpretation solves the interpretation of measurement by insisting that we only consider states that include measurement devices, such that the apparent nonlinearity of measurement doesn't really exist.

At the other extreme, we can interpret measurement by noting that the nonlinearity in quantum measurement is precisely the same as in a branch of statistics known as Bayesian inference. Thus, quantum mechanics only appears nonlinear when we forget to include that there is an agent performing the measurement who then learns from each result. This observation leads to thinking of quantum mechanics not as a description of the world, but as a description of what we know about the world.

Though these two kinds of interpretations disagree at a philosophical level, both offer different ways of resolving how a linear theory such as quantum mechanics can sometimes appear to be non-linear. Regardless of which interpretation helps you to understand the interaction between measurement and

the rest of quantum mechanics, you can take solace in that the results of measurement are always described by the same math and by the same simulations. Indeed, relying on simulations (sometimes sarcastically called the "shut up and calculate" interpretation) is the oldest and most celebrated of all interpretations.

What this squared length of each projection *represents* is the probability that the state you are measuring would be found along that direction. If you have a qubit in the $|0\rangle$ state and try to measure it along the direction of the $|1\rangle$ state, you will get a probability of zero, because the states are opposite each other when we draw them on a circle. Thinking in terms of pictures, the $|0\rangle$ state has no projection onto the $|1\rangle$ state — in the sense of, $|0\rangle$ doesn't leave a shadow on $|1\rangle$.

TIP

If something happens with probability 1, then that event always occurs. If something happens with probability 0, then that event is impossible. For example, the probability that a typical 6-sided die rolls a "7" is zero, since that roll is impossible. Similarly, if a qubit is in the $|0\rangle$ state, getting a "1" result from a Z-axis measurement is impossible, since $|0\rangle$ has no projection onto $|1\rangle$.

If you have a $|0\rangle$ and try to measure it along the $|0\rangle$ direction, however, you will get a probability of 1 because the states are parallel (and of length 1 by definition). Let's walk through what measuring a state that is neither parallel nor perpendicular would look like.

EXAMPLE

Say you had a qubit in state $(|0\rangle + |1\rangle)/\sqrt{2}$ (same as stem: $|1\rangle$ from our table), and you wanted to measure it or project it along the Z-axis. Then, we can find the probability that the classical result will be a 1 by projecting $|1\rangle$ onto $|1\rangle$.

We can find the projection of one state onto another by using the *inner product* between their vector representations. In this case, we write the inner product of $|+\rangle$ and $|1\rangle$ as $\langle 1|+\rangle$, where $\langle 1|$ is the transpose of $|1\rangle$, and where butting the two bars against each other indicates taking the inner product. We can write this out as follows.

NOTE

In the next chapter, we'll see that $\langle \mathbf{1} |$ is the conjugate transpose of $|\mathbf{1} \rangle$, but we'll set that aside for now.

To compute the projection, we start by writing down the "bra" that we want to project onto.

 $\langle 1 | (|0\rangle + |1\rangle) / \sqrt{2}$

Next, we distribute the bra.

 $=(1/\sqrt{2})(\langle 1|0\rangle + \langle 1|1\rangle)$

We can write each bra-ket pair as an inner product between two vectors.

 $=(1/\sqrt{2})([0],[1])\cdot[[1],[0]]+[[0],[1]]\cdot[[0],[1]]$

Calculating each inner product makes things a lot simpler!

 $=(1/\sqrt{2})(0+1)$

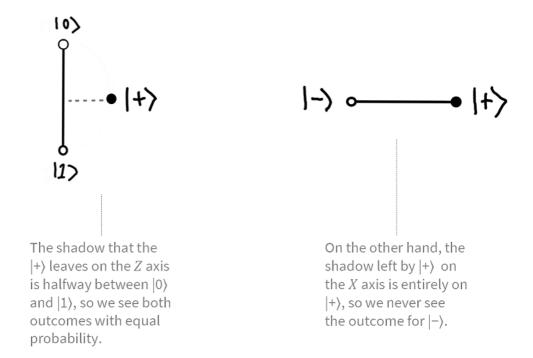
We now have the overlap between $|0\rangle$ and $|1\rangle$.

 $=1/\sqrt{2}$.

To turn this projection into a probability, we square it, getting that the probability of observing a "1" outcome when we prepare a $|+\rangle$ state is $\frac{1}{2}$.

We will often project onto the Z axis because it is convenient in many real experiments, but we could have also measured along the X-axis, to see if we have a $|+\rangle$ or a stem:[\left|-\right\rangle] state. Measuring along the X axis, we would have gotten stem:[$|+\rangle$ with certainty and would never have gotten \left|-\right\rangle.

Figure 2.16. Measuring $|+\rangle$ along the X axis always results in $|+\rangle$.



NOTE

We can get a fully certain measurement out come only because we know the "right" direction to measure ahead of time in this case — if we are simply handed a state with no information about what the "right" measurement direction is, we cannot predict any measurement outcome perfectly.

2.3.5 Generalizing measurement: basis independence

Sometimes you may not know how your qubit was prepared so you will not know how to measure the bits properly. More generally, any pair of states that don't overlap (that are opposite poles) defines a measurement in the same way. The actual outcome of a measurement is a classical bit value that indicates which pole the state is aligned with when we perform the measurement.

More general measurements still

Quantum mechanics allows for much more general kinds of measurements — we'll see a few of these as we go along, but mostly we focus in this book on the case of checking between two opposite

poles. This choice is a pretty convenient way of controlling most quantum devices, and can be used in almost any of the commercial platforms for quantum computing that are currently available.

Mathematically, we use notation like "(measurement | state)" to represent measuring a qubit. The left component (measurement | is called a *bra*, and we have seen the *ket* part on the right already. So together they are called a braket!

Bras are very similar to *kets*, except that to switch from one to the other you have to take the transpose (turn rows to columns and vice versa) of the bra or ket you have,

$$|0\rangle^T = [[1], [0]]^T = [[1, 0]].$$

Another way to think of this is that taking the transpose turns column vectors (kets) into row vectors (bras).

NOTE

Since we're only working with real numbers for now, we won't need to do anything else to go between kets and bras, but when we work with complex numbers in the next chapter, we'll need the complex conjugate as well.

Bras let us write down measurements, but to see what measurements actually do, we need one more thing at our disposal: a rule for how to use a bra and a ket together to get the *probability* for seeing that measurement result. In quantum mechanics, measurement probabilities are found by looking at the length of the projection or shadow that the ket for a state leaves on a bra for a measurement. We know from our experience from the party that we can find projections and lengths using inner products. In Dirac notation, the inner product of a bra and a ket is written as (measurement | state), giving us just the rule we need.

For example, if we have prepared a state $|+\rangle$] and we want to know the probability that we observe a "1" when we measure in the Z basis, then projecting in the way we saw with figure 2.15; we can find the the length we need. The projection of stem:[$|+\rangle$ onto $\langle 1|$ tells us that we see a "1" outcome with probability $\Pr(1|+) = |\langle 1|+\rangle|^2 = |\langle 1|0\rangle + \langle 1|1\rangle|^2/2 = |0+1|^2/2 = 1/2$. Thus, 50% of the time, we'll get a "1" outcome. The other 50% of the time, we'll get a "0" outcome.

BORN'S RULE

If we have a quantum state |state| and we perform a measurement along the (measurement| direction, we can write the probability that we will observe "measurement" as our result as

$$Pr(measurement|state) = |\langle measurement|state\rangle|^2$$

In words, the probability is the square of the magnitude of the inner product of the measurement bra and the state ket.

This expression is called Born's rule.

In table 2.3, we've listed several other examples of using Born's rule to predict what classical bits we will get when we measure qubits.

Table 2.3. Table Examples of using Born's rule to find measurement probabilities

If we prepare	and we measure	then we see that outcome with probability.
0>	(0)	$ (0 0) ^2 = 1$
0>	(1	$ \langle 1 \mid 0 \rangle ^2 = 0$
0>	(+	$ (+ \mid 0) ^2 = ((0 \mid + (1 \mid) \mid 0) / \sqrt{2} ^2 = (1 / \sqrt{2} + 0)^2 = 1 / 2$
+>	(+	$ (4 + +) ^2 = ((0 + (1) (0) + 1)) / 2 ^2 = (0 0) + (1 0) + (0 1) + (1 1) ^2 / 4 = 1 + 0 + 0 + 1 ^2 / 4 = 2^2 / 4 = 1$
+>	(-	$ (- +) ^2 = 0$
- 0>	(0)	$ -\langle 0 \mid 0 \rangle ^2 = -1 ^2 = 1^2$
- +>	(-	

TIP

In , we used that $\langle 0 \mid 0 \rangle = \langle 1 \mid 1 \rangle = 1$ and $\langle 0 \mid 1 \rangle = \langle 1 \mid 0 \rangle = 0$. (Try checking this for yourself!) When two states have an inner product of zero, we say that they are orthogonal (or perpendicular). That $|0\rangle$ and $|1\rangle$ are orthogonal makes a lot of calculations easier to do quickly.

We now have made it to the bottom of the Qubit quick start guide! Let's review the requirements we needed to satisfy to make sure we had working qubits.

QUBIT

A qubit is any physical system satisfying three properties:

- The system can be perfectly simulated given knowledge of vector of numbers (the "state").
- The system can be transformed using quantum operations (e.g.: rotations).
- Any measurement of the system produces a single classical bit of information, following Born's rule.

Anytime we have a qubit (a system with the above three properties) we can describe it using the same math or simulation code, without further reference to what kind of system we are working with. This is similar to how we need not know whether a bit is defined by the direction of a pinball's motion or the voltage in a transistor in order to write down NOT and AND gates, or to write software which uses those gates to do interesting computation.

Phase

In the last two rows of , we saw that multiplying a state by a phase of -1 didn't affect measurement probabilities. This isn't a coincidence at all, but points to one of the more interesting things about qubits. Because Born's rule only cares about the squared absolute value of the inner product of a state and a measurement; multiplying a number by (-1) doesn't affect its absolute value. We call numbers such as +1 or -1, whose absolute value is equal to 1, phases. In the next Chapter when we work more with complex numbers, we'll see a lot more about phases.

For now, though, we say that multiplying an entire vector by -1 is an example of applying a **global phase** while changing from $| \rangle$ to $| - \rangle$ is an example of applying a relative phase between $| 0 \rangle$ and $| 1 \rangle$. While global phases don't ever affect measurement results, there's a big difference between the states $| \rangle = (| 0 \rangle + | 1 \rangle) / \sqrt{2}$ and $| - \rangle = (| 0 \rangle - | 1 \rangle) / \sqrt{2}$: the coefficients in front of $| 0 \rangle$ and $| 1 \rangle$ are the same in $| + \rangle$ and are different by a phase of (-1) in $| - \rangle$. We will see much more of the difference between these two concepts in Chapters 3, 4, 6, and 7.

NOTE

Similarly to how we use the word "bit" to mean both a physical system that stores information and the information stored in a bit, we will also use the word "qubit" to mean both a quantum device and the quantum information stored in that device.

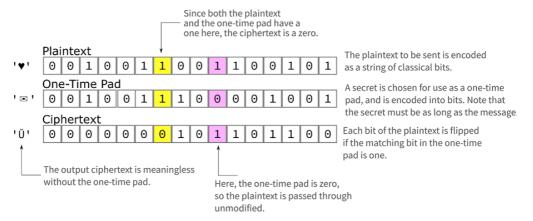
2.3.6 Simulating qubits in code

The quantum quick start guide lists cryptography as an initial application to check individual qubits in your new device.

Suppose you would like to keep your **v** for Eve a secret lest anyone else finds out. How can you scramble up your message to Eve so that only she can read it?

We'll explore this application more in the next Chapter, but the most basic step we need for any good *encryption* algorithm is a source of random numbers that's difficult to predict. Let's write down exactly how we would combine our secret and random bits to make a secure message to sent to Eve. In , we show an example of how if both Eve and I know the same secret sequence of random classical bits, we can use that sequence to communicate securely. At the start of the chapter, we saw how we could write the message, or *plaintext*, that we want to send to Eve (in this case, "") as a string of classical bits. The one-time pad is a sequence of random classical bits that will act as a way to scramble or encrypt our message. This scrambling is done by taking the bitwise XOR of the message and one-time pad bits for each position in the sequence. This then produces a sequence of classical bits called the *ciphertext*. To anyone else trying to read our message, the ciphertext will just look like random bits. For example, it's impossible to tell if a bit in the ciphertext is "1" because of the plaintext or the one-time pad.

Figure 2.17. An example of how to use random bits to encrypt secrets, even over the Internet or another untrusted network.



Now you might ask how to get the random bit strings for our one-time pad? We can make our own quantum random number generator with qubits! It may seem odd but we will now simulate qubits with classical bits to make our quantum random number generator. The random numbers it will generate won't be any more secure than the computer we use to do our simulation, but it lets us get a good start in understanding qubits and how they work.

Let's send Eve our message! In the same way as a classical bit can be represented in code by the values True and False, we've seen that we can represent the two qubit states $|0\rangle$ and $|1\rangle$ as *vectors*. That is, qubit states are represented in code as lists of lists of numbers.

- We use the NumPy library for Python to represent vectors, as NumPy is highly optimized and will make our lives much easier.
- 2We name our variable ket0 after the notation $|0\rangle$, in which we label qubit states by the "ket" half of $\langle \rangle$ "brakets."

3 NumPy will print out 2 \times 1 vectors as columns.

As we saw above, we can construct other states such as $|+\rangle$] by using linear combinations of $|0\rangle$ and $|1\rangle$. In exactly the same sense, we can use NumPy to add the vector representations of $|0\rangle$ and $|1\rangle$ to construct the vector representation of stem: $|+\rangle$:

- We can see NumPy using vectors to store the \left|+\right\rangle state, which is a linear combination of |0⟩ and |1⟩.
- ②We will see the number 0.70710678 a lot in this book, as it is a rather good approximation to $\sqrt{2}$, the length of the vector [[1], [1]].

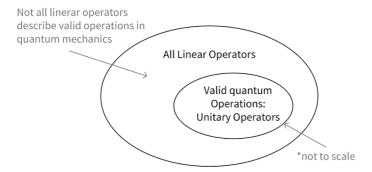
In classical logic, if we wanted to simulate how an operation would transform a list of bits, we could use a *truth table*. Similarly, since quantum operations other than measurement are always linear, to simulate how an operation transforms the state of a qubit, we can use a matrix that tells us how each state is transformed.

LINEAR OPERATORS AND QUANTUM OPERATIONS

Describing quantum operations as linear operators is a good start, but not all linear operators are valid quantum operations! If we could implement an operation described by a linear operator such as 2×1 (that is, twice the identity operator), then we would be able to violate that probabilities are always numbers between zero and one. We also require that all quantum operations other than measurement are reversible, as this is a fundamental property of quantum mechanics.

It turns out that the operations realizable in quantum mechanics are described by matrices U whose inverses U^{-1} can be computed by taking the conjugate transpose, $U^{-1} = U^{+}$. Such matrices are called unitary matrices.

Figure 2.18. Visualizing types of valid quantum operations.



One particularly important quantum operation is called the *Hadamard operation*, which transforms $|0\rangle$ to stem: $|+\rangle$ and $|1\rangle$ to $|-\rangle$. As we saw above, measuring $|\rangle$ along

the Z-axis gives us either a "0" or a "1" result with equal probability. Since we wanted random bits in order to send secret messages, this makes the Hadamard operation really useful for us in making our QRNG.

Using vectors and matrices, we can define the Hadamard operation by making a table of how it acts on the $|0\rangle$ and $|1\rangle$ states, as shown in .

Table 2.4. Table Representing the Hadamard operation as a table.

Input state	Output state
0>	$ +\rangle = (0\rangle + 1\rangle) / \sqrt{2}$
1>	$ -\rangle = (0\rangle - 1\rangle) / \sqrt{2}$

Because quantum mechanics is linear, this is a fully complete description of the Hadamard operation!

In matrix form, we write down as H = np.array([[1, 1], [1, -1]]) / np.sqrt(2).

Listing 2.6. Defining the Hadamard operation

• We define a variable H to hold the matrix representation H of the Hadamard operation that we saw in .
We'll need H throughout the rest of this chapter, so it's helpful to define it here.

HADAMARD OPERATION

The Hadamard operation is a quantum operation which can be simulated by the linear transformation

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

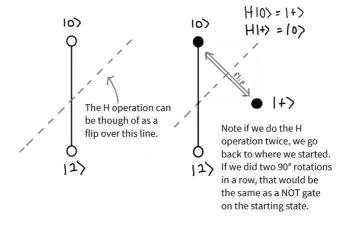
Any operation on quantum data can be written as a matrix in this way. If we wish to transform $|0\rangle$ to $|1\rangle$ and vice versa (the quantum generalization of the classical NOT operation that we saw earlier, corresponding to a 180° rotation), we do the same thing as we did to define the Hadamard operation.

Listing 2.7. Representing the quantum NOT gate

- **1** The quantum operation corresponding to the classical NOT operation is typically called the *X* operation; we represent the matrix for *X* with a Python variable X.
- 2 We can confirm that X transforms $|0\rangle$ to $|1\rangle$. The NumPy method all() returns True if every element of X @ ket0 == ket1 is True; that is, if every element of the array X @ ket0 is equal to the corresponding element of ket1.
- 3 The X operation doesn't do anything to H $|0\rangle$, since X will swap $|0\rangle$ and $|1\rangle$ and H\ $|0\rangle$ is already a sum of the two kets: $(|0\rangle+|1\rangle)/\sqrt{2}=(|1\rangle+|0\rangle)/\sqrt{2}$. We can confirm this by using the @ operator again to multiply X by a Python value representing the state stem: $|+\rangle=H$ $|0\rangle$. We can express that value as H @ ket0.

Returning to the map analogy, we can think of H as a *reflection* about the \angle direction.

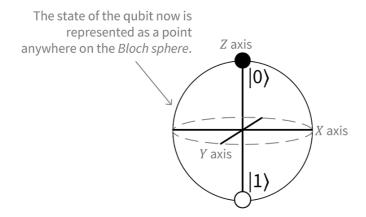
Figure 2.19. The H operation as a reflection or flip about \mathbb{Z} .



The third dimension awaits!

For qubits, the map analogy helps us understand how to write down and manipulate the states of single qubits. So far, however, we've only looked at those states that can be written down using real numbers. In general, quantum states can use complex numbers. If rearrange our map a bit and make it three-dimensional, then we can include complex numbers without any problem. This way of thinking about qubits is called the *Bloch sphere*, and can be a very useful way of thinking about quantum operations as rotations and reflections, as we'll see more of in Chapter 5.

Figure 2.20. Visualizing qubit states as points on a sphere.



DEEP DIVE: Infinitely many states?

It may seem from that there are infinitely many different states of a qubit. For any two different points on a sphere, we can always find a point that's "between" them. While this is true, it can also be a little bit misleading. Thinking of the classical situation for a moment, a coin which lands heads 90% of the time is distinct from a coin that lands heads 90.000000001% of the time. In fact, we can always make a coin whose bias is "between" the bias of two other coins in this way. Flipping a coin can only ever give us one classical bit of information, though. On average, it would take is about 10^{23} flips to tell a coin that lands heads 90% of the time apart from one that lands heads 90.000000001% of the time. For all intents and purposes, we can treat these two coins as identical because we cannot do an experiment which reliably tells them apart. Similarly for quantum computing, there are limits to our ability to tell apart the infinitely many different quantum states that we recognize from the Bloch sphere picture.

The fact that a qubit has infinitely many states is not what makes it unique. Sometimes people say that a quantum system can be "in infinitely many states at once", which is why they say quantum computers can offer speedups. **THIS IS FALSE!** As pointed out above, we can't distinguish states that are very close together so the "infinitely many" part of the statement can't be what gives our quantum computer an advantage. We will talk more in the upcoming chapters about the "at once" part, but suffice it to say it is not the number of states that our qubit can be in that makes quantum computers cool!

2.4 PROGRAMMING A WORKING QRNG

Now that we have a few quantum concepts to play with, let's apply what we've learned to program a quantum random number generator (QRNG) so that we can send ♥s without a worry. We are going to build a quantum random number generator that returns either a 0 or a 1.

Random bits or random numbers?

It may seem limiting that our random number generator can only output one of two numbers, either 0 or 1. Quite to the contrary, though, this is enough to generate random numbers in the range 0 to N for any positive integer N. It's easiest to see this starting with the special case that N is $2^n - 1$ for some positive integer n, in which case we simply write down our random numbers as n-bit strings. For example, we can make random numbers between 0 and 7 by generating three random bits n0, n1, and n2, then returning n4 n2 + n3 returning 4 n3.

The case is slightly more tricky if N isn't given by a power of two, in that we have "left over" possibilities that we need to deal with. For instance, if we need to roll a six-sided die, but only have an eight-sided die on hand (maybe we played a Druid last time at RPG night), then we need to decide what to do when that die rolls either a 7 or an 8. The best thing we can do if we want a fair six-sided die is to simply reroll when that happens. Using this approach, we can build arbitrary fair dice from coin flips — handy for whatever game we want to play. Long story short, we aren't limited by having just two outcomes from our RNG!

As with any quantum program, our quantum random number generator program will be a sequence of instructions to a device that performs operations on a qubit. In pseudocode, a quantum program for implementing a QRNG consists of three instructions:

2.4.1 QRNG

- 1. Prepare a qubit in the state $|0\rangle$.
- 2. Apply the Hadamard operation to our qubit, so that it is in the state $|+\rangle = H|0\rangle$.
- 3. Measure the qubit to get either a 0 or 1 result with 50/50 probability.

Figure 2.21. Steps to writing the QRNG program we want to test out from the new hardware kit.

```
Step 1: Prepare starting qubit H|0\rangle \Rightarrow |+\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle)
Step 2: Apply Hadamard Step 3: Measure qubit OR 1 if state was |0\rangle
OR 1 if state was |1\>

Each outcome happens 50 % of the time, just like if you just flipped a coin!
```

That is, we want a program that looks something like the following:

Listing 2.8. Example pseudocode for a QRNG program

```
def qrng():
    q = Qubit()
    H(q)
    return measure(q)
```

Using matrix multiplication, we can use a classical computer like a laptop to simulate how grng () would act on an ideal quantum device.

Our qrng program calls into a software stack that abstracts away whether we're using a

classical simulator or an actual quantum device.

Classical Computer Quantum Program A programmer writes Q# or Python program on their One way of implementing classical computer that sends the interface for a quantum instructions to an interface for a device is by using a classical Quantum Device Interface quantum device. computer to simulate quantum mechanics. Simulator Backend This interface defines what Such simulations make instructions are available for a heavy use of linear quantum device. For instance, algebra, such as matrix Linear Algebra Package a machine might expose an multiplication. Most of the interface with "Hadamard" or time this is provided by a "measure" instructions. backend like LAPACK or Intel's Math Kernel Library. Quantum Computer Dilution Fridge Quantum computers are very susceptible to noise so we Digital Signal Processor The quantum program have to isolate them from the sends instructions to a environment. One easy way to digital signal processor do so is to keep them very or other classical cold, often near absolute zero! Quantum Device hardware used to control a quantum system.

Figure 2.22. An example of what a software stack for a quantum program might look like.

There are a lot of parts we see here to the stack, but don't worry we will talk about them as we go. For right now we are focusing on the top section (labeled "Classical Computer"), and will start by writing code for a quantum program as well as a simulator backend in Python.

NOTE

In Chapter 6, we'll pivot to making use of the simulator backend provided with Microsoft's Quantum Development Kit instead.

With this view of a software stack in mind, then, we can write our simulation of a QRNG by first writing a QuantumDevice class with abstract methods for allocating qubits, performing operations, and measuring qubits. We can then implement this class with a simulator and then call into that simulator from grng().

To design the interface for our simulator in a way that looks like, let's list out what we need our quantum device to be able to do:

2.4.2 Quantum device interface requirements.

• Users must be able to allocate and return qubits.

Listing 2.9. An example of specifying an interface into a quantum device as a set of

```
abstract methods.

class QuantumDevice(metaclass=ABCMeta):
    @abstractmethod
    def allocate_qubit(self) -> Qubit:
        pass

    @abstractmethod
    def deallocate_qubit(self, qubit: Qubit):
        pass

    @contextmanager

def using_qubit(self):
        qubit = self.allocate_qubit()
        try:
            yield qubit
        finally:
            qubit.reset()
            self.deallocate_qubit(qubit)

4
```

- 1Any implementation of a quantum device must implement this method, allowing users to obtain qubits.
- When users are done with a qubit, implementations of the deallocate_qubit will allow users to return the qubit back to the device.
- 3We can provide a Python context manager to make it easy to allocate and deallocate qubits safely.
- The context manager makes sure that no matter what exceptions are raised, each qubit is reset and deallocated before being returned to the classical computer.

The qubits themselves then can expose the actual transformations that we need:

2.4.3 Qubit interface requirements.

- Users must be able to perform Hadamard operations on qubits.
- Users must be able to measure qubits to get out classical data.

Listing 2.10. An example of specifying an interface into the qubits on a quantum device

as a set of abstract methods. class Qubit(metaclass=ABCMeta): @abstractmethod def h(self): pass @abstractmethod def measure(self) -> bool: pass @abstractmethod def reset(self): pass 3

The h method can be implemented to transform a qubit in place (not making a copy) using the Hadamard operation np.array([[1, 1], [1, -1]]) / np.sqrt(2).

- 2 The measure method can be implemented to allow users to measure qubits and extract classical data.
- 3 The reset method makes it easy for users to prepare the qubit from scratch again.

With this in place, we can return to our definition of grng using these new classes.

Listing 2.11. qrng.py def qrng(device: QuantumDevice) -> bool: with device.using_qubit() as q: q.h() return q.measure()

If we implement the QuantumDevice interface with a class called SingleQubitSimulator, then we can pass this to qrng to run our QRNG implementation on a simulator.

```
Listing 2.12. qrng.py

if __name__ == "__main__":
    qsim = SingleQubitSimulator()
    for idx_sample in range(10):
        random_sample = qrng(qsim)
        print(f"Our QRNG returned {random_sample}.")
```

We now have everything we write our SingleQubitSimulator. We start by defining a couple of constants for the vector $|0\rangle$ and the matrix representation of the Hadamard operation H.

- 1 Since we'll be using |0⟩ a lot in our simulator it helps to define a constant for it.
- 2 Similarly, we'll use the Hadamard matrix *H* to define how the Hadamard operation transforms states, so we define a constant for that as well.

Next, we define what a simulated qubit looks like. From the perspective of a simulator, a qubit wraps a vector that stores the current state of the qubit. We use a NumPy array to represent our qubit's state.

```
self.reset()

def h(self):
    self.state = H @ self.state

def measure(self) -> bool:
    pr0 = np.abs(self.state[0, 0]) ** 2
    sample = np.random.random() <= pr0
    return bool(0 if sample else 1)

def reset(self):
    self.state = KET_0.copy()</pre>
```

- **1** As a part of the Qubit interface, we ensure that the reset method prepares our qubit in the $|0\rangle$ state. We can use that when we create the qubit to make sure that qubits always start in the correct state.
- 2) The Hadamard operation can be simulated by applying the matrix representation H to the state that we're storing at the moment, then updating to our new state.
- (3) We stored the state of our qubit as a vector, so we know that the inner product with |0⟩ is simply the first element of that vector. For instance, if the state is np.array([[a], [b]]) for some numbers a and b, then the probability of observing a 0 outcome is |a|². We can find this using np.abs(a) ** 2. This gives us the probability that a measurement of our qubit returns 0.
- **4** To turn the probability of getting a 0 into a measurement result, we generate a random number between 0 and 1 using np.random.random and check if it's less than pr0.
- **5** Finally, we return out to the caller a 0 if we got a 0 and a 1 if we got a 1.

WHAT RANDOM NUMBER CAME FIRST: 0 OR

In making this QRNG, we'll have to call a classical random number generator. This may feel a bit circular, but it comes about because our classical simulation is just that: a simulation. A simulation of a quantum random number generator won't be any more random than the hardware and software we use to implement that simulator.

That said, the quantum program qrng.py itself does not need to call a classical RNG, but calls into the simulator. If we were to run qrng.py on an actual quantum device, the simulator and hence the classical RNG would be substituted out for operations on the actual qubit. At that point, we would have a stream of random numbers that would be impossible to predict thanks to the laws of quantum mechanics.

Running our program, we now get the random numbers we expected!

```
$ python qrng.py
Our QRNG returned False.
Our QRNG returned True.
Our QRNG returned True.
Our QRNG returned False.
Our QRNG returned False.
Our QRNG returned False.
Our QRNG returned False.
Our QRNG returned True.
Our QRNG returned False.
Our QRNG returned True.
```

Congratulations! You've not only written your first quantum program, but you've also written a simulation backend and used it to run your quantum program in the same way as you'd run on an actual quantum computer.

DEEP DIVE: Schrödinger's Cat

You may have already seen or heard of the quantum program above, but under a very different name. Often, the QRNG program is described in terms of the "Schrödinger's cat" thought experiment: a cat is in a closed box with a vial of poison that will be released if a particular random particle decays. Before you open the box to check, how do you know if it is alive or dead?

The [state] of the entire system would express this by having in it the living and dead cat (pardon the expression) mixed or smeared out in equal parts.

- Erwin Schrödinger

Historically, Schrödinger proposed this description in 1935 to express his view that some implications of quantum mechanics are "ridiculous" by means of a thought experiment that highlights how counterintuitive these implications are. Such thought experiments, known as *gedankenexperiment*, are a celebrated tradition in physics, and can help us understand or critique different theories by pushing them to extreme or absurd limits.

In reading about Schrödinger's cat nearly a century later, however, it's helpful to remember everything that's happened in the intervening years. Since his original letter, the world has seen:

- War on a scale never before imagined,
- The first steps that humanity has taken to explore beyond our own planet,
- The rise of commercial jet travel,
- The understanding and first effects of anthropogenic climate change,
- A fundamental shift in how we communicate (television all the way through the Internet),
- A wide availability of affordable computing devices, and
- The discovery of a wondrous variety of subatomic particles.

Put simply, the world we live in isn't the same world in which Schrödinger tried to make sense of quantum mechanics. We have a lot of advantages in trying to understand, none the least of which being that we can quickly get our hands on quantum mechanics by programming simulations using classical computers. For example, the h instruction we saw earlier puts our qubit in a similar situation as the cat in the gedankenexperiment above, but with the advantage that it's much easier to experiment with our program than with a thought experiment. Throughout the rest of the book, we'll make use of our quantum programs to learn the parts of quantum mechanics we need to write quantum algorithms.

2.5 SUMMARY

In this chapter you learned:

- Recognize classical and quantum bits (qubits),
- Predict how different quantum operations transform qubits with linearity,
- program a qubit simulator that can simulate quantum random number generators.

Sharing secrets with Quantum key distribution

This chapter covers:

- · Recognize the implications that quantum resources have for security
- · Programming a simulator in Python for a quantum key distribution protocol
- Implementing the quantum NOT operation

Figure 3.1. In this Chapter, we will be covering topics in simulating quantum hardware and intrinsic operations for the quantum devices.

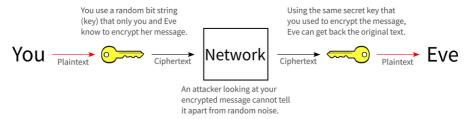
Quantum Applications		
Libraries / Techniques		
Q# Language		
Intrinsic Operations		
Quantum Hardware	Simulators	

3.1 ALL'S FAIR IN LOVE AND ENCRYPTION

In the previous Chapter, we saw that we could use random numbers as a resource to send secret messages like "" to our friends. We also saw that by using quantum bits or qubits, we could generate truly random numbers that are impossible for any adversary to predict.

That's only half the story, though, because we need to share those random numbers with our friends if we want to use the quantum random numbers to communicate securely with them. Those random numbers (often called a *key*) can be use with *encryption* algorithms which combine the randomness of the key with information people want to keep secret in such a way that only someone else with the *key* can see the information. We can see in how two people could use a key (here random binary string) to encrypt and decrypt messages between themselves.

Figure 3.2. Mental model for how you and Eve might use encryption to communicate secretly, even over the Internet or another untrusted network.



In this Chapter, we will see that quantum technologies can help us with encryption or other cryptographic tasks, by letting us securely *distribute* our secret keys. There are classical methods for sharing random keys (e.g. RSA), but they have different guarantees about the security of the sharing. In short, using *quantum key distribution* (QKD) is *provably* secure, whereas classical key distribution methods are often *computationally* secure. This difference doesn't matter for most use cases, but if you are a government, activist group, bank, journalist, spy, or any other group where information security is a life-and-death matter, this is a huge deal.

EXAMPLE 3.1. COMPUTATIONAL VS. PROVABLE SECURITY

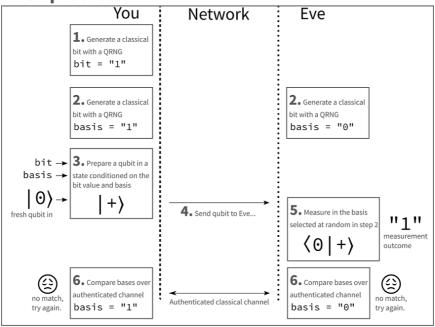
Provable security for our cryptographic protocols is the dream. A method or protocol for a cryptographic task is *provably secure* if we can write an proof showing it is secure using no assumptions about an adversary; i.e. that they can have all the time and computing power in the universe and still our protocol is secure! Most of our current cryptographic infrastructure is *computationally secure* which guarantees the security of a method or protocol with reasonable assumptions about the capabilities of an adversary. The designer or user of the protocol can choose thresholds for what finite computer resources looks like (e.g. largest current super computer or all the computers on the planet) and what a reasonable time is (e.g. 100 years, 10000 years, the age of the universe).

When we share a key with QKD, it does not guarantee that the key will get to the other person. This is because someone can always do a denial of service attack (e.g. cut the optical fiber between sender and receiver), which is the same for any other classical protocol. A good analogy for what QKD can promise is similar to the tamper proof seal on food products. When a peanut butter company wants to ensure that when you open the jar it is exactly as it was when it left the factory floor, they put one of these tamper proof seals on the container. The company makes a promise that if the seal makes it to you (the consumer) intact, the peanut butter will be good, and that no third party has done anything to it. Transmitting a cryptographic key with a QKD protocol is like putting a tamper proof seal on the bits in transit. If someone tries to compromise the key in transit, the receiver will know and not use that key. Sealing the bits in transit does *not*, however, guarantee that the bits make it to the receiver.

There are many protocols that we can use to implement the general QKD scheme. In this chapter we will be working with one of the most common QKD protocols BB84, but there are many others as well that we won't have time to get in to. We will build up to this throughout this chapter, but you can see in the steps to the BB84 protocol.

 $\textbf{Figure 3.3. Timing diagram for the BB84 protocol}, \, \textbf{a particular variant of a QKD protocol}. \\$

BB84 protocol



QKD is an example of a quantum program that uses a single qubit, as well as a spin-off technology from quantum computing. What makes it attractive to develop is that we already have the hardware to implement it today! There are a number of companies

Repeat steps 1-6 until you have as much key as you need.

that have for the last ~15 years have been commercially selling QKD hardware, however the important next steps for the technology involve hardware and software security vetting of these systems.

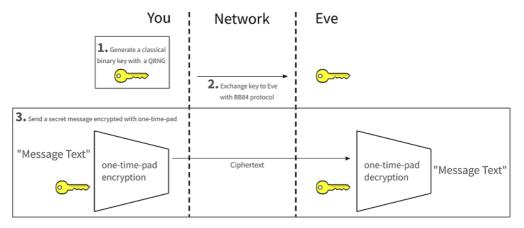
TRY THIS AT HOME, BUT NOT WITH YOU SECRETS! The examples we are implementing/using here in this book will simulate provably secure protocols. Given we are not running the examples on quantum devices, they are not provably secure. Even when implementing these protocols with real quantum hardware, these security proofs do nothing to stop side channel attacks or social engineering from separating you from your key we will talk more about these proof later in this chapter when we talk about the No-cloning theorem.

Let's dive into how QKD works! For our purposes here, let's say you and Eve are the two people from the previous chapter that want to exchange a key so you can send secret messages The scenario is as follows.

You wish to send a secret message to your friend. Using your quantum random number generator from Chapter 2 and the quantum key distribution protocol BB84, and one-time-pad encryption, design a program to send messages that can be provably secure.

You can visualize the scenario as a kind of timing diagram like.

Figure 3.4. Your scenario for this chapter: Sending a secret message to Eve with BB84 and one-time-pad encryption.



To start off, note that the key you need to send is a string of classical bits. How can we use qubits to send those classical bits? We will start by learning how to encode classical information in qubits, and then learn the specific steps of the BB84 protocol. In the next section, we will look at a new quantum operation that will help you encode classical bits with qubits.

3.2 QUANTUM NOT OPERATIONS

If we have some classical information, say a single binary bit, how can we encode this classical information with a quantum resource like a qubit? Take a look at the

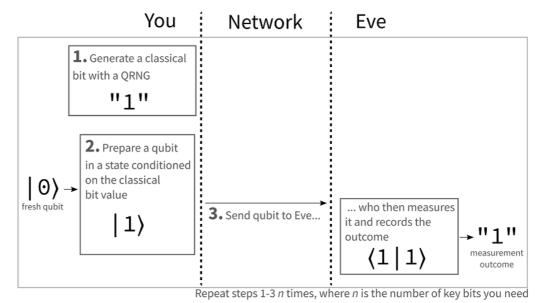
following algorithm.

3.2.1 Algorithm for sending a random classical bit string encoded in qubits.

- Using a quantum random number generator to generate a random key bit to send.
- Starting with a qubit in the $|0\rangle$ state and then prepare it in a state that represent that bit value from step 1; here you use $|0\rangle$ if the classical bit was a 0, and $|1\rangle$ if the classical bit was a 1.
- That prepared qubit is sent to Eve who then measures it and records the classical bit value they get.
- Repeat steps 1-3 until you and Eve have as much key as you want (usually dictated by the cryptographic protocol you want to use after).

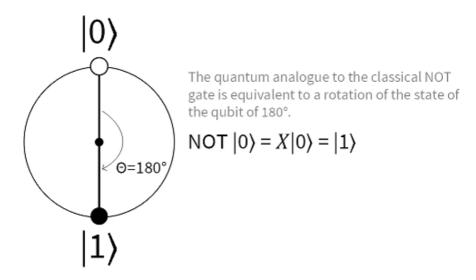
You can see a timing diagram for this algorithm in.

Figure 3.5. A visualization of the algorithm for sending a classical bitstring with qubits.



Now we need a way to switch the qubit from $|0\rangle$ to $|1\rangle$, so we need another quantum operation in our toolbox. In step 2, you can use a *quantum* NOT operation, which is similar to our classical NOT operation, that rotates the qubit from $|0\rangle$ to $|1\rangle$.

Figure 3.6. A visualization of the quantum equivalent of a NOT operation operating on a qubit in the $|0\rangle$ state, leaving the qubit in the $|1\rangle$ state.



We refer to this quantum NOT operation as the x operation. Step 2 from above then could be re-written as follows:

2\. If your classical bit from step 1 was 0, do nothing; if it was a 1, apply a quantum NOT operation (a.k.a. x operation) to your qubit.

This algorithm works 100% of the time because when Eve goes to measure the qubit they receive, the $|0\rangle$ and $|1\rangle$ states can be perfectly distinguished with a measurement in the Z-axis. It may seem like you and Even have done a lot of work just to share some random classical bits, but we will see how adding some quantum behaviors to this basic protocol will make it more useful! Let's look at how we could implement this in code:

Listing 3.2. You and Eve exchange classical bits via qubits, encoding the message with

the |0| and |1| states.

def prepare_classical_message(bit: bool, q: Qubit) -> None:
 if bit:
 q.x()

def eve_measure(q: Qubit) -> bool:
 return q.measure()

def send_classical_bit(device: QuantumDevice, bit: bool) -> None:
 with device.using_qubit() as q:
 prepare_classical_message(bit, q)
 result = eve_measure(q)



- 1 To prepare our qubit with the classical bit we want to send, we need as input the bit value and a qubit to use. This function does not return anything because the consequences of the operations we apply to our qubit are tracked in the single qubit simulator itself.
- 2 If you're sending a 1, you can use the NOT operation x to prepare q in the |1) state because the x operation will rotate |0) to |1) and vice versa.
- 3 This seems silly to separate measuring as another function given it's one line, but we will change up how Eve will measure the qubit in the future so this is a helpful setup.
- 4 We can check that measuring q gives the same classical bit as you sent.

The simulator you wrote in the previous chapter *almost* has what you need to implement this. You just need to add an instruction corresponding to the \times operation. The \times instruction can be represented with a matrix X, just as we represented the h instruction using the matrix H. In the same way as we wrote down H in Chapter 2, we can write down the matrix X as

$$X = \left(\sum_{m=1}^{\infty} 0 \& 1 \setminus 1 \& 0 \right).$$

Exercise 3.1: truth tables and matrices

In Chapter 2, we saw that unitary matrices play the same role in quantum computing that truth tables play in classical computing. We can use that to figure out what the matrix X has to look like in order to represent the quantum NOT operation, x. Let's start by making a table of what the matrix X has to do to each input state in order to represent what the x instruction does:

Input	Output
0)	1>
1>	0>

This table tells us that if we multiply the matrix X by the vector $|0\rangle$, we need to get $|1\rangle$, and similarly that $X|1\rangle = |0\rangle$.

Either by using NumPy or by hand, check that the matrix

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

matches what we have in our truth table above.

Let's go on and add the functionality we need to our simulator in order to run the above snippet. We will be working with the simulator you wrote in the previous chapter, but if you need a refresher, you can find the code in the GitHub repo for this book: https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp First, we need to update the interface for our quantum device, by adding a new method our qubit must have.

Listing 3.3. Adding the x operation to the interface specification for a qubit.

class Qubit(metaclass=ABCMeta):
 @abstractmethod

```
def h(self): pass

@abstractmethod
def x(self): pass

@abstractmethod
def measure(self) -> bool: pass

@abstractmethod
def reset(self): pass
```

• We can model implementing the quantum NOT operation after the h operation from Chapter 1.

Listing 3.4. Adding the x operation to the implementation of a qubit simulator.

Now that our interface for a qubit knows that we want an implementation of the x operation, let's add that implementation!

```
KET_0 = np.array([
    [1],
    [0]
], dtype=complex)
H = np.array([
    [1, 1],
    [1, -1]
], dtype=complex) / np.sqrt(2)
X = np.array([
    [0, 1],
    [1, 0]
], dtype=complex) / np.sqrt(2)
class SimulatedQubit(Qubit):
    def __init__(self):
        self.reset()
    def h(self):
        self.state = H @ self.state
    def x(self):
                                                                            2
        self.state = X @ self.state
    def measure(self) -> bool:
        pr0 = np.abs(self.state[0, 0]) ** 2
        sample = np.random.random() <= pr0</pre>
```

1 Let's add the a variable X to store the matrix *X* that we need to represent the x operation.

return bool(0 if sample else 1)

self.state = KET 0.copy()

def reset(self):

2 Just like the h function defined above, we want to implement the quantum operation x by just applying the matrix stored in X to the state vector.

3.3 SHARING CLASSICAL BITS WITH QUBITS

Awesome! Now let's try out using our upgraded Python qubit simulator to share a secret classical bit with a qubit. This is not quite the same as a quantum key distribution protocol yet, but it serves as a good foundation for the types of function and steps that our end goal QKD protocol has.

Open an IPython session where you are keeping the code for your simulator by running ipython in your terminal. After importing the Python files, create an instance of the single qubit simulator and generate a random bit to use as the classical bit we want to send. (Good thing we have a quantum random number generator!) Then using a fresh qubit, prepare it based on the classical bit value you want to send Eve. Eve then measures the qubit, and we can see if you both have the same classical bit value!

Listing 3.5. Using our single qubit simulator to send classical bits with qubits >>> qrng_simulator = SingleQubitSimulator() >>> key_bit = int(qrng(qrng_simulator)) >>> qkd_simulator = SingleQubitSimulator() >>> with qkd_simulator.using_qubit() as q: ... prepare_message(key_bit, q) ... print(f"You prepared the classical key bit: {key_bit}") ... eve_measurement = int(eve_measure(q)) ... print(f"Eve measured the classical key bit: {eve_measurement}") ... You prepared the classical key bit: 1 Eve measured the classical key bit: 1

- 1 We need a simulated qubit to use for our quantum random number generator.
- 2 Re-using the qrng function that we wrote in the previous chapter, we can generate a random classical bit to use for our key.
- **3**We are using a new qubit simulator instance here for the key exchange, strictly speaking we don't need to. You will see in the next chapter how to expand the simulator to work with multiple qubits.
- You encode your classical bit in the qubit provided by qkd_simulator. If the classical bit was 0, you do nothing to qkd_simulator, and if the classical bit was 1 then you use the x method to change the qubit to the |1) state.
- See measures the qubit from gkd simulator and then stores the bit value as eve measurement.

Our example of secret sharing with qubits above should be deterministic, which is to say that every time you prepare and send a bit, Eve will correctly measure the same value. Is this secure? If you suspect it is not secure, you are definitely on to something:) In the next section we will discuss the security of our prototype secret sharing scheme, and look at ways you can improve it!

3.4 A TALE OF TWO BASES

You and Eve now have a way of sending classical bits using qubits, but what happens if an adversary gets a hold of that qubit? They could just use the measure instruction to get the same classical data that Eve does. That's a huge problem and would

reasonably make you wonder why one would use qubits to share keys in the first place.

Thankfully, quantum mechanics offers us a way to make this exchange more secure! What are some modifications we could make to our protocol above? We could have, for instance, decided to represent a classical "0" message with a qubit in the $|+\rangle$ state and a "1" message with a qubit in the $|-\rangle$ state.

Listing 3.6. You and Eve exchange classical bits via qubits, but encode the message

with the $|+\rangle / |-\rangle$ states.

```
def prepare_classical_message_plusminus(bit: bool, q: Qubit) -> None:
    if bit:
        q.x()
    q.h()

def eve_measure_plusminus(q: Qubit) -> bool:
    q.h()
    return q.measure()

def send_classical_bit_plusminus(device: QuantumDevice, bit: bool) -> None:
    with device.using_qubit() as q:
        prepare_classical_message_plusminus(bit, q)
        result = eve_measure_plusminus(q)
        assert result == bit
```

- ① Everything up to this line of prepare_classical_message_plusminus was the same as we were using before with prepare_classical_message. Applying the Hadamard gate at this point rotates the |0⟩ / |1⟩ states to |+⟩ / |-⟩ states.
- ②Here we use the h operation to rotate our $|+\rangle$ / $|-\rangle$ states back to the $|0\rangle$ / $|1\rangle$ states because our measure operation is defined to only measure the $|0\rangle$ / $|1\rangle$ states correctly. Another way of thinking about this is that you are rotating the measurement to match the rotation we are currently working in $(|+\rangle$ / $|-\rangle$). It's all a matter of perspective!

Now you have two different ways of sending qubits that you and Eve could use when sending qubits (see for a summary). These two different ways of sending messages we call *bases*, which each contain two completely distinguishable (orthogonal) states. This is similar to the previous chapter where we looked at map directions (ex. North and West) which defined a convenient *basis* for describing directions. Here, we have used the $|0\rangle$ and $|1\rangle$ states as one basis (called the Z basis), and $|+\rangle$ and $|-\rangle$ as another (called the X basis). The names for these bases refer to the axis along which you can perfectly distinguish the states (see).

Figure 3.7. Now in addition to using the ${\it Z}$ basis to encode a classical bit on a qubit, we can use the ${\it X}$ basis.

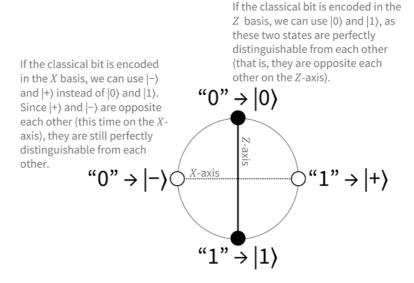


Table 3.1. Different classical messages we want to send, and how to encode them in the ${\cal Z}$ and ${\cal X}$ bases.

	"0" message	"1" message
"0" (or Z) basis	0)	$ 1\rangle = X 0\rangle$
"1" (or X) basis	$ +\rangle = H 0\rangle$	$ -\rangle = H 1\rangle = HX 0\rangle$

NOTE

In quantum computing there is never really a correct basis, so much as there are convenient bases that we choose to use by convention.

If you both don't know which way of sending you are using for a particular bit, you both have a problem. What happens if we mix sending our messages in Z basis and X basis? Good news, you can use your simulator to try it out and see what happens.

Listing 3.7. You and Eve exchange classical bits via qubits, but are not using the same

```
basis.

def prepare_classical_message(bit: bool, q: Qubit) -> None:
    if bit:
        q.x()

def eve_measure_plusminus(q: Qubit) -> bool:
```

```
q.h()
  return q.measure()

def prepare_classical_message(bit: bool, q: Qubit) -> None:
  if bit:
        q.x()

def eve_measure_plusminus(q: Qubit) -> bool:
        q.h()
        return q.measure()

def send_classical_bit_wrong_basis(device: QuantumDevice, bit: bool) -> None:
    with device.using_qubit() as q:
        prepare_classical_message(bit, q)
        result = eve_measure_plusminus(q)
        assert result == bit, "Two parties do not have the same bit value"
```

- 3 Here you will be using the method we saw before to prepare your qubit in the Z basis basis by using the h method.
- 4 Eve measures in the X basis because she does a Hadamard gate on her qubit before measuring.
- The function does not return anything, so if you and Eve end up with key bits that don't match, it will raise an error.

Listing 3.8. Sending qubits in the Z basis, and measuring in the X basis

```
>>> qsim = SingleQubitSimulator()
>>> send_classical_bit_wrong_basis(qsim, 0)
AssertionError: Two parties do not have the same bit value
```

• We will just pick our bit value is 0, and you may have to run this line a few times before you get the above error.

You can try this out experimentally, and you will find that you get the above AssertionError (the key exchange failed) about half of the time. Why is that? To start with, Eve is measuring in the X basis, so she can only tell $|+\rangle$ and $|-\rangle$ apart perfectly. What will she measure if she is not given a perfectly distinguishable state for her basis (like here where she is given a $|0\rangle$)? We can write the $|0\rangle$ state in the X basis as:

$$|0\rangle = (|+\rangle + |-\rangle)/\sqrt{2}$$

Recall that in Chapter 2 we defined $|+\rangle$ in a similar way by adding $|0\rangle$ and $|1\rangle$ together. The $|+\rangle$ state was also called a *superposition* of the $|0\rangle$ and $|1\rangle$ states.

NOTE

Any time a state can be written as a linear combination of states like this, it is considered to be a superposition of the states that are added together.

Exercise 3.2: verify that $|0\rangle$ is a superposition of $|+\rangle$ and $|-\rangle$.

Try using what we learned about vectors in the previous chapter to verify that $|0\rangle = (|+\rangle + |-\rangle) / \sqrt{2}$, either by hand or using Python. Hint: recall that $|+\rangle = (|0\rangle + |1\rangle) / \sqrt{2}$ and that $|-\rangle = (|0\rangle - |1\rangle) / \sqrt{2}$.

Now to calculate the actual measurement with Born's rule from Chapter 2. Recall we can calculate the probability of a measurement outcome by measuring a particular state with the expression:

$$\Pr(\text{measurement}|\text{state}) = |\langle \text{measurement} \mid \text{state} \rangle|^2$$

Writing out the measurement of the $|0\rangle$ state in the X basis you can see that we will get 0 (or $|+\rangle$) half of the time and 1 (or $|-\rangle$) the other half.

$$\Pr\left(\langle +|0\rangle\right) = \left|\langle +|0\rangle\right|^2$$

$$= \left(\left(\langle +|+\rangle + \langle +|-\rangle\right)/\sqrt{2}\right)^2$$

$$= (1+0)^2/2$$

$$= 1/2$$

Exercise 3.3: measuring qubits in different bases

Using the example above as a guide,

- calculate what the probability of getting the measurement outcome |-> when measuring the |0> state in the |-> direction.
- Also calculate what the probability of getting the |-> measurement outcome with the input state of |1>.

That tells us that if Eve does not know the right basis to measure in, then the measurements she makes are as good as randomly guessing. This is because in the wrong basis, the qubit is in a superposition of the two states that define the basis. One 'key' to how QKD works is that without the right additional information (the basis the qubit is encoded in) any measurement of the qubit is basically useless. Now to ensure our security, we have to make it difficult for an adversary to learn that extra information to know the right basis to measure in. The QKD protocol we will look at next has a solution for this, and a proof (out of scope here) that describes the chance that the attacker has any information about the key!

3.5 QUANTUM KEY DISTRIBUTION: BB84

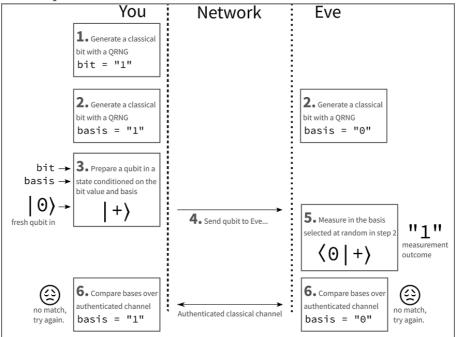
We have now seen how to share keys in two different bases, and what happens if you and Eve don't use the same basis. You might again ask why we use this to make sharing the key to our secret keys more secure? There are a wide variety of different QKD protocols, each with specific advantages and use cases (not unlike RPG character classes). The most *common* protocol for QKD is called BB84, named for an appropriately cryptic encoding of the two authors initials and the year it was published (Bennet and Brassard 1984).

BB84 is very similar to what we have worked out so far to share keys, but has one important difference in how you and Eve choose your bases. In BB84, both parties choose their basis randomly (and independently), which means they will end up using the same basis 50% of the time. You can see a figure with the steps of the BB84

protocol in.

Figure 3.8. The steps in the BB84 protocol, a particular version of a QKD protocol.

BB84 protocol



Repeat steps 1–6 until you have as much key as you need.

As a consequences of randomly choosing bases, you and Eve also have to do some communication over authenticated* classical channels (like the internet) to be able to take the keys you each have and transform them into a key that you believe is identical to the key their partner has. This is because this is real life, and when the qubits are exchanged, it will be possible for both the environment and third party individuals to manipulate or modify the state of the qubit.

Key expansion

There is one detail we glossed over in our description of the classical communication channel that you and Eve have, namely in that it must be *authenticated*. That is, when you send classical messages to Eve as a part of running BB84, it's OK if someone else can read them, but you need to make sure it's really Eve that you're talking to. To prove that someone wrote and sent a particular message you actually **already** have to have some form of shared secret that you can use to validate someone else's identity. So we have to already have some shared secret with the other person in BB84. This secret can be smaller than the message you're trying to send, so BB84 is techically more of a *key expansion* protocol.

Table 3.2. What state you should send for each random message and basis choice

	"0" message	"1" message	
"0" (or <i>Z</i>) basis	0)	$ 1\rangle = X 0\rangle$	
"1" (or <i>X</i>) basis	$ +\rangle = H 0\rangle$	$ -\rangle = H 1\rangle = HX 0\rangle$	

3.5.1 Steps of the BB84 protocol:

- You choose a random one-bit message to send by sampling your QRNG
- You and Eve each choose a random basis with your respective QRNGs (no communication between the them)
- You prepare a qubit in the randomly selected basis, representing your randomly selected message (see).
- You send your prepared qubit in the quantum channel to Eve.
- Eve measures the qubit when it arrives, performing the measurement in her randomly selected basis, and recording the classical bit outcome.
- Communicate on an authenticated classical channel with Eve and share which bases you used for preparing and measuring the qubit. If they match, keep the bit and add it to the key.
- Repeat steps 1–6 until you have as much key as you need.

An error free world

Since we are simulating the BB84 protocol, we know that the qubit Eve will receive is exactly the same as what you sent. BB84 more realistically will be done in batches where n qubits are exchanged first, and then a round of sharing the basis values (error correction happen). You also have to at the end shrink the key even further with privacy amplification algorithms to account for the fact that an eavesdropper could have gotten partial information from the errors you detected. We omitted these steps in our implementation of BB84 to keep things simple, however the steps are critical for real-world security \Box

Let's jump in and implement the BB84 QKD protocol in Python! We will start by writing a function that will will run the BB84 protocol (assuming lossless transmission) for a single bit transmission. That does not guarantee that we get one key bit from this run however, if you and Eve choose different bases then that exchange will have to be thrown out.

First, it is helpful to setup some functions that will help us simplify how we write out the full BB84 protocol. You and Eve need to do things like sample random bits and prepare and measure the message qubit, which are separated here for clarity.

Listing 3.9. Some helper functions before we get to the full BB84 key exchange.

```
def sample_random_bit(device: QuantumDevice) -> bool:
    with device.using_qubit() as q:
        q.h()
    result = q.measure()
```

```
1
        q.reset()
    return result
def prepare message qubit(message: bool, basis: bool, q: Qubit) -> None: 2
    if message:
        q.x()
    if basis:
        q.h()
def measure message qubit(basis: bool, q: Qubit) -> bool:
   if basis:
        q.h()
    result = q.measure()
    q.reset()
    return result
def convert to hex(bits: List[bool]) -> str:
    return hex(int(
        "".join(["1" if bit else "0" for bit in bits]),
   ))
```

- ①sample_random_bit is almost the same as our qrng function before, except here we will reset the qubit after measuring as we know we want to be able to use it more than once.
- 2 Here the qubit is encoded with the key bit value in the randomly selected basis.
- Similar to sample_random_bit after Eve measures the message qubit, she should reset it because in the simulator we will be reusing it for the next exchange.
- To help condense the display of long binary keys, a helper function is used to convert the representation to a shorter hex string.

Listing 3.10. BB84 protocol for sending a single classical bit between two parties.

```
def send_single_bit_with_bb84(
   your_device: QuantumDevice,
   eve_device: QuantumDevice
) -> tuple:

[your_message, your_basis] = [
        sample_random_bit(your_device) for _ in range(2)
]

eve_basis = sample_random_bit(eve_device)

with your_device.using_qubit() as q:
        prepare_message_qubit(your_message, your_basis, q)

# QUBIT SENDING...

eve_result = measure_message_qubit(eve_basis, q)

return ((your_message, your_basis), (eve_result, eve_basis))

1
```

• Here you can randomly choose a bit value and basis using your modified qrng from before, here the sample random bit function.

- 2 Eve needs to randomly choose a basis with their own qubit, which is why she is using a separate QuantumDevice.
- 3 With all the preparation done, you then need to prepare your qubit to send to Eve.
- Since all of our computation happens inside a simulator on your computer, nothing needs to be done to "send" the qubit from you to Eve. In real life however, this is where all the bad stuff happens: errors, loss, even attempted eavesdropping.
- **5** Now Eve has your qubit, and measures it in the randomly selected basis she chose earlier.
- 6 This function returns the key bit values and bases you and Eve would have at the end of this one round.

Qubits and No-cloning

From what we've seen so far, it seems like our adversary could cheat by eavesdropping on the qubits in the quantum channel and making copies. Here's the gig: The eavesdropper (called Bob here) first would need to (without detection):

Steps for Bob to eavesdrop on you and Eve

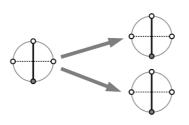
- 1. Copy qubits as they are being sent between you and Eve, and then store them.
- 2. Next, while you and Eve finish the classical part of the protocol, Bob listens to the bases they announce and keep track of ones you both chose the same.
- 3. For the qubits corresponding to bits where you and Eve used the same basis, Bob would measure their copies of the qubits in the same basis as well.

Ta-da! You, and Eve and Bob would all have the same key! If this seems like a problem, you are right. Don't worry though, quantum mechanics has the solution. It turns out the problem with Bob's plan is in step 1, where they need to make *identical* copies of the qubits that you and Eve are exchanging. The good news is that making an exact copy of a qubit without knowing what it was beforehand is forbidden by quantum mechanics. The rule that qubits cannot be identically copied without prior knowledge of the state is called the *No-cloning theorem* and is stated as follows.

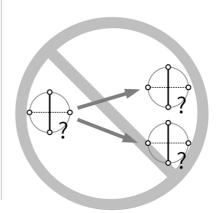
NO-CLONING THEOREM No quantum operation can perfectly copy the state of an arbitrary qubit onto another qubit.

Figure 3.9. The No-cloning theorem visualized.

If we have some classical description of the state of the qubit, then we can make copies of it. Here we have a qubit in the $|1\rangle$ state and we can make "copies" of it by preparing more by taking a $|0\rangle$ and applying the x operation.



If we have no prior information about a qubit, then we **cannot** make perfect copies of it.



You will be able to do the simple proof of this in the next chapter, once we learn how to describe the state of more than one qubit \bigcirc .

Alternately, if Bob could measure a qubit without disturbing it, he could get around needing a copy of the qubits he intercepts. This is not possible because once you measure a qubit, it "collapses" or changes in a way that can be detectable to Eve. So measuring in transit is not something Bob can do without being detected, so his eavesdropping would fail.

Now, exchanging one classical key bit is not going to be sufficient to send a whole key so now we need to use the above technique to send multiple bits.

Listing 3.11. BB84 protocol for exchanging key with Eve until a specified amount of key

1 At this point, you and Eve can publicly announce the bases you each used to measure this bit. If everything worked right, your results should agree whenever your bases agree. We'll check that here with an assert.

The key is now in the bag, so we can move on to using the key and the one-time-pad encryption algorithm to send a secret message!

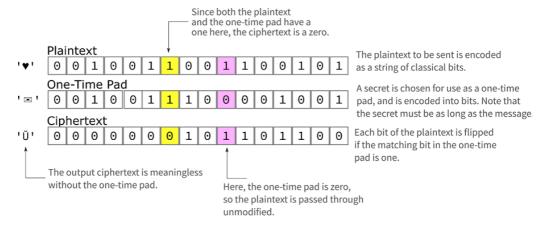
3.6 USING OUR SECRET KEY TO SEND SECRET MESSAGES

You and Eve have now sorted out how to use the BB84 protocol to share a random, binary key generated by a QRNG. The last step now is to use this key to share a secret message with Eve! You and Eve had previously decided that the best encryption protocol to use is a one-pad to send your secret messages. This turns out to be one of the most secure encryption protocols, and given you are sharing keys in one of the most secure ways possible it makes sense to keep up that standard!

Let's say that you were trying to send Eve that you like Python, so the message you want to send is " Since we are using a binary key, we need to convert the representation of our unicode message to binary, which is the following lengthy list of bits.

This binary representation of our message is our *message text* and now we want to combine that with a key to get a ciphertext that is safe to send over the network. Once you have the key from the BB84 protocol (at least as long as your message) next you need to use a one-time-pad encryption scheme to then encode your message. We saw this encryption technique already in Chapter 2, see for a quick refresher.

Figure 3.10. An example of one-time-pad encryption which uses random bits to encrypt secret messages.



To implement this, you will need to use a classical bitwise XOR (the ^ operator in Python) to combine the message and your key to create the ciphertext that you can safely then send to Eve. To decrypt your message Eve will do the same bitwise XOR operation with the ciphertext and her key (which should be the same as yours). This will give her back the message because any time you XOR a bitstring with another one twice, you will be left with the original bitstring. Here is what this would look like in Python.

Listing 3.12. BB84 protocol for exchanging key with Eve until a specified amount of key

is reached.

```
def apply_one_time_pad(message: List[bool], key: List[bool]) -> List[bool]:
    return [
        message_bit ^ key_bit
        for (message_bit, key_bit) in zip(message, key)
]
```

1 The ^ operator is a bitwise XOR in Python. This applies a single bit of our key as a one-time-pad to our message text.

Exercise 3.4: one-time-pad encryption

If you had the ciphertext 10100101 and the key 00100110, what was the message that was originally sent?

Let's put it all together and share our message (" De un uning the bb84.py file we have been building up!

Listing 3.13. Sharing a secret message with Eve with BB84 and one-time-pad

encryption.

```
if name == " main ":
   print("Generating a 96-bit key by simulating BB84...")
   key = simulate bb84(96)
   print(f"Got key
                                              {convert_to_hex(key)}.")
   message = [
       1, 1, 0, 1, 1, 0, 0, 0,
       0, 0, 1, 1, 1, 1, 0, 1,
       1, 1, 0, 1, 1, 1, 0, 0,
       1, 0, 0, 1, 0, 1, 1, 0,
       1, 1, 0, 1, 1, 0, 0, 0,
       0, 0, 1, 1, 1, 1, 0, 1,
       1, 1, 0, 1, 1, 1, 0, 0,
       0, 0, 0, 0, 1, 1, 0, 1,
       1, 1, 0, 1, 1, 0, 0, 0,
       0, 0, 1, 1, 1, 1, 0, 1,
       1, 1, 0, 1, 1, 1, 0, 0,
       1, 0, 1, 1, 1, 0, 1, 1
   print(f"Using key to send secret message: {convert to hex(message)}.")
```

```
encrypted_message = apply_one_time_pad(message, key)
print(f"Encrypted message:
{convert_to_hex(encrypted_message)}.")

decrypted_message = apply_one_time_pad(encrypted_message, key)
print(f"Eve decrypted to get:
{convert_to_hex(decrypted_message)}.")
```

Listing 3.14. Running the full solution to the scenario for the chapter.

```
$ python bb84.py
Generating a 96-bit key by simulating BB84...
Took 170 rounds to generate a 96-bit key.
Got key: 0xb35e061b873f799c61ad8fad.
Using key to send secret message: 0xd83ddc9dd83ddc0bb.
Encrypted message: 0x6b63da8d5f02a591b9905316.
Eve decrypted to get: 0xd83ddc9dd83ddc0bb.

5
```

- 1 Since your and Eve's bases will agree roughly half of the time, if should take about two rounds of BB84 for each bit of key you want to generate.
- 2The exact key you generate will be different every time you run the BB84 simulation that's a huge part of the point of the protocol, after all!
- When we combine our secret message with the key we got above, using the key as a one-time-pad, our message is completely scrambled.
- 5 When Eve uses the same key, she gets our original secret message back!

3.7 SUMMARY

In this chapter you learned:

- Recognize the implications that quantum resources have for security
- Implement the quantum NOT operation
- Program a simulator in Python for a quantum key distribution protocol

Quantum key distribution is one of the most important spin-off technologies from quantum computing and has huge potential impact for our security infrastructure. While it is currently quite easy to setup QKD for parties that are relatively close to each other (< 200 km), there are significant challenges to deploying a global system for QKD. Usually the physical system used in QKD is a photon and it is hard to send single particles of light long distances without losing them.

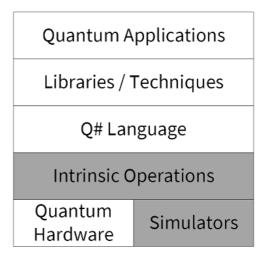
Now that you have built up a single qubit simulator and programmed some single qubit applications, you are now ready to start playing around with multiple qubits! In the next chapter we will take the simulator you have built and add features to allow it to simulate multiple qubit states and use it to play nonlocal games with Eve. \square

Nonlocal games: working with multiple Qubits

This chapter covers:

- Simulate state preparation, operations, and measurement results for multiple qubits,
- · Program a simulator for multiple qubits leveraging the QuTiP Python package and tensor products,
- Recognize the proof quantum mechanics is consistent with our observations of the universe by simulating experimental results.

Figure 4.1. In this Chapter, we will be covering topics in simulating quantum hardware and intrinsic operations for the quantum devices.



4.1 NONLOCAL GAMES

At this point, we have seen how single-qubit devices can be programmed to accomplish useful tasks such as random number generation and quantum key distribution. The most exciting computational tasks, however, require using multiple qubits together. In this chapter, you will learn about nonlocal games: a way to validate our quantum mechanical descriptions of the universe with friends with multi-qubit systems.

We will dive into a new Python package called QuTiP that will allow us to program quantum systems faster and has some cool built-in features for simulating quantum mechanics. Then you will learn how to leverage QuTiP and program a simulator for *multiple* qubits and see how that changes (or doesn't!) the 3 main tasks for our qubits: state preparations, operations, and measurement.

4.2 WHAT ARE NONLOCAL GAMES?

We have all played games of one type or another, whether sports, board games, video games or roleplaying games. Games are one of the best ways we have to explore new worlds, to test our limits of strength, endurance and understanding. Turns out Eve loves to play games, and the latest encrypted message from her was the following text.

"Hi player! I am keen to play a game called CHSH. It is a nonlocal game were we play with a referee. I'll send the instructions in the next message. STOP"

What makes the game proposed by Eve *nonlocal* is the fact that the players are (sadly) not in the same place while playing the game. The players participate in the game by sending and receiving messages with a central referee but don't have a chance to talk to each other while playing the game. What is really cool about it is that by playing it we can show that classical physics just doesn't cut it to describe the results we get in these games with particular strategies. The particular winning strategy we will look at here involves the players sharing a pair of qubits before the game starts. We will dive into just what entangling two qubits means as we go through this chapter, but let's just start with describing the full rules of our nonlocal game.

NOTE

A referee adjudicating a nonlocal game can ensure that they players don't communicate by separating the players by a large enough distance that no light from one player could reach the other before the game ends.

4.3 TESTING QUANTUM PHYSICS ITSELF: THE CHSH GAME

The nonlocal game the Eve has suggested playing is called *the CHSH game*. [1]

You get the next message from Eve and it contains the following information, as well the figure shown in .

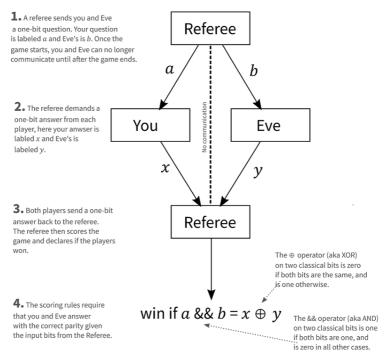


Figure 4.2. The CHSH game, a nonlocal game with two players and a referee.

The CHSH game is comprised of two players, and a referee. You can play as many rounds of the game as you like, and each round has 3 steps. As Eve mentioned in her first message, once a round starts, the players cannot communicate and make their own (possibly pre-planned) decisions.

4.3.1 Steps for one round of the CHSH game

The referee starts the round by giving you and Eve each a one classical bit. The referee choses these bits *independently* and at *uniformly random*, so you could get a 0 or 1 each with 50% probability, and same for Eve. This means that there are four possible ways that the referee can start the game (your bit, Eve's bit): (0,0), (0,1), (1,0), or (1,1).

You and Eve must each **independently** decide on a single classical bit to give back to the referee as a response.

The referee then calculates the parity (XOR) of your and Eve's classical bit responses. As listed in , in three out of the four cases, you and Eve must respond with *even* parity (your answers must be equal) in order to win, while in the fourth case, your answers must be different. These are definitely *unusual* rules, but not too bad compared to some multi-day board games.

Table 4.1. Win conditions for the CHSH game.

Your Input	Eve's Input	Response parity to win	
0	0	Even	
0	1	Even	
1	0	Even	
1	1	Odd	

• We can expand on to get all of the possible outcomes of the game, see .

Table 4.2. All possible states of the CHSH game with win conditions. Input bits come from the referee, and both of you respond to the referee as well.

Your Input	Eve's Input	Your Response	Eve's Response	Parity	Win?
0	0	0	0	Even	Yes
0	0	0	1	Odd	No
0	0	1	0	Odd	No
0	0	1	1	Even	Yes
0	1	0	0	Even	Yes
0	1	0	1	Odd	No
0	1	1	0	Odd	No
0	1	1	1	Even	Yes
1	0	0	0	Even	Yes
1	0	0	1	Odd	No
1	0	1	0	Odd	No
1	0	1	1	Even	Yes
1	1	0	0	Even	No
1	1	0	1	Odd	Yes
1	1	1	0	Odd	Yes
1	1	1	1	Even	No

Let's look at some Python code to simulate this game. Since your and Eve's responses to the referee are allowed to depend on the message that the referee gives you, you can represent each player's actions as a "function" that the referee calls.

Exercise 4.1: Umpire state of mind

Since the referee is purely classical, we'll model them as using classical random number generators. This leaves open the possibility, though, that you and Eve could cheat by guessing the referee's questions. A

possible improvement might be to use the QRNGs from Chapter 2. Modify the code sample in so that the referee can ask questions of you and Eve by measuring a qubit that starts off in the |+> state.

Listing 4.1. Python implementation for the CHSH game

```
import random
from functools import partial
from typing import Tuple, Callable
import numpy as np
from interface import QuantumDevice, Qubit
from simulator import Simulator
Strategy = Tuple[Callable[[int], int], Callable[[int], int]]
def random_bit() -> int:
    return random.randint(0, 1)
def referee(strategy: Callable[[], Strategy]) -> bool:
    you, eve = strategy()
    your input, eve input = random bit(), random bit()
    parity = 0 if you(your_input) == eve(eve_input) else 1
    return parity == (your_input and eve_input)
def est win probability(strategy: Strategy,
                        n games: int = 1000) -> float:
    return sum(
       referee(strategy)
        for idx_game in range(n_games)
    ) / n_games
```

- 1 Here we are declaring a new type Strategy to define the tuple of functions that represent your and Eve's one bit functions that represent your individual strategies. Using Python's typing module lets us document that a value of type Strategy is a tuple of two functions, each of which takes an int and returns an int. We can think of these functions as representing what you and Eve each do with the bits given to you by the referee.
- 2 The classical random number generator the referee will use.
- 3 You and Eve confer before the game begins and decide on a strategy. The strategy function will assign a function to you and eve one-bit functions that represent what you will do based on your input.
- 4 The referee picks two random bits, one for each player.
- **5** Give each player their random bit, then compute the parity of their responses.
- 6 Check to see if the players won.
- To finish your implementation, you will need a function that runs the CHSH game many times in a row and checks how often you and Eve win. Here, you can use Python's built-in sum function to count the number of times referee returns True for a particular strategy. Divide by the number of games that you played then gives you an estimate for the probability that your and Eve's strategy wins the CHSH game.

Note that in we don't have a definition yet for the input strategy to the referee. Now that we have the rules of the game implemented in Python, let's talk *strategy* and get to implementing a classical strategy for playing the CHSH game.

ГΙР

It's helpful to choose variable naming conventions that make it obvious what role each variable plays in your code. Above, we chose to use the prefix n_ in the variable n_games to indicate that the variable refers to a number or size, and have used the prefix idx_ to refer to the index for each individual game. Much like driving, it's good if your code is predictable.

4.4 CLASSICAL STRATEGY

The simplest strategy for both you and Eve to pursue is to ignore your inputs entirely. Looking at, if both of you agreed before the game that you would never change your outputs (i.e. always return 0), you will win 75% of the time (this does assume the referee chooses the random bits for each player uniformly).

Table 4.3. Best classical strategy for the CHSH game, where you both always respond with 0 with win conditions.

Your Input	Eve's Input	Your Response	Eve's Response	Parity	Win?
0	0	0	0	Even	Yes
0	1	0	0	Even	Yes
1	0	0	0	Even	Yes
1	1	0	0	Even	No

If we were to write this strategy as a Python function, we would have the code in .

Listing 4.2. An easy, constant strategy for players of a CHSH game. Also the best

strategy you can use with only classical resources.

```
def constant_strategy() -> Strategy:
    return (
        lambda your_input: 0,
        lambda eve_input: 0
)
```

Listing 4.3. Testing how often we could expect to win a round of CHSH if we use the constant_strategy.

• Note that you may get slightly more or less than 75% when you try this. This is because the win probability is estimated using a finite number of rounds (in stats, this is called a binomial distribution). For this example, we'd expect error bars of about 1.5%.

Ok, so thats an easy strategy, but is there anything more clever you can do? Given that you and Eve only have classical resources, sadly this is **provably** the best you can do. Short of cheating **(e.g.** communicating or guessing the referee's inputs) you cannot

win this game more than 75% of the time on average.

This all leads up to the obvious question: what if you and Eve could use qubits? What would be your best strategy then, and how often would you win? What does it say about our understanding of the universe if we have proof that you cannot win CHSH more than about 75% of the time, and then we find a way that we *can* beat that win rate? As you might guess, we can do better than 75% win rate playing CHSH if the players share quantum resources, i.e. have qubits. Later in the chapter we will get into quantum based strategies for CHSH, but spoiler we are going to need to simulate more than one qubit.

4.5 WORKING WITH MULTIPLE QUBIT STATES

So far in this book, we've only worked with one qubit at a time. To play a nonlocal game, for example, each player will need their own qubit. This raises the question, how do things change when the system we are considering has more than one qubit? The main difference is that we cannot describe each qubit individually and have to think in terms of a state that describes the whole system.

IMPORTANT

When describing a group or register of qubits, you generally cannot just describe each qubit individually. The most useful quantum behaviors can only be seen when you describe the state of a group or register of qubits.

The next section will help relate this system level view with a similar, classical programming concept of a register.

4.6 REGISTERS

Suppose that we have a *register* of classical bits; that is, a collection of many classical bits. We can index through each bit in that register and look at it's value independently, even though they are still a part of that register. The contents of the register can represent a more complex value, like bits that together represent a Unicode character (as we saw in Chapter 3), but this higher level interpretation is not necessary.

When we store information in a classical register, the number of different states of that register grows very rapidly as we add more bits. If we have three bits, for example, there are eight different states that our register can be in; see for an example. We say for the classical register state 101 that the zeroth bit is a 1, the first is 0, and the second is 1. When these values are concatenated together, they give us the string 101.

Listing 4.4. Using Python to write out all the possible states of a classical three bit

register.

```
>>> from itertools import product
>>> ", ".join(map("".join, product("01", repeat=3)))
'000, 001, 010, 011, 100, 101, 110, 111'
```

If we have four bits, we can store one of 16 different states; if we have n bits, we can

store one of 2ⁿ. We say that the number of different possible states of a classical register *grows exponentially* with the number of bits. The bit strings output by show the actual data in the register for each state of the register. They also serve as convenient labels for one of 8 possible messages that we can encode with 3 classical bits.

What does all this have to do with qubits? We saw in Chapters 2 and 3 that any state of classical bit also describes a qubit state. The same holds for registers of qubits as well. For instance, the three-bit state "010" describes the three-qubit state \left|010\right\rangle.

TIP

In Chapter 2, we saw that qubit states described by classical bits in this way are called computational basis states; here, we'll use the same term for states of multiple qubits that are described by strings of classical bits.

Just like with single qubits, however, the state of a register of multiple qubits can also be made up by adding different qubit states together. In the exact same way that we can write down $|+\rangle$ as $(|0\rangle + |1\rangle)/\sqrt{2}$ to get another valid qubit state, our three-qubit register can be in a wide variety of different states.

4.6.1 Example three-qubit states

- $(|010\rangle + |111\rangle)/\sqrt{2}$
- $(|001\rangle + |010\rangle + |100\rangle)/\sqrt{3}$

WHY THE SQUARE ROOTS?

We'll see more as we go along, but just as we needed the square root of 2 to make the measurement probabilities work out for $|+\rangle=(|0\rangle+|1\rangle)/\sqrt{2}$, we need to divide by $\sqrt{2}$ and

 $\sqrt{3}$ in our examples above to make sure all the probabilities for each measurement are realistic; i.e. add up to one.

This example of the linearity of quantum registers is called the *superposition principle*.

SUPERPOSITION PRINCIPLE

The superposition principle tells us that we can add together two different states of a quantum register together to get another valid state.

To write down the state of a quantum register in a computer, we'll again use vectors, just as we did in Chapter 2. The main difference is how many numbers we list in each vector. Let's look at what it looks like to write down the state of a two-qubit register on a computer. For example, the vector for the two-qubit state $(|00\rangle + |11\rangle)/\sqrt{2}$ can also be written as the state $(1 \times |00\rangle + 0 \times |01\rangle + 0 \times |10\rangle + 1 \times |11\rangle)/\sqrt{2}$. If we make a list of what we had to multiply each computational basis state by in order to get the state we wanted, we have precisely the information we need to write down in our vector. In listing 4.5, we've written down $(|00\rangle + |11\rangle)/\sqrt{2}$ as a vector.

Listing 4.5. Using Python to write out an example of a two-qubit state.

- 1 We start the same way, using the np.array function to make a new vector.
- 2 Each entry in this vector describes a different computational basis state. This entry tells us that we have to multiply |00⟩ by 1.
- **3** Similarly, this entry tells us how much of the $|01\rangle$ state we need to add to get the state that we want.
- **4** Finally, we divide by $\sqrt{2}$ to make sure all the measurement probabilities work out, just as we did with the $|+\rangle$ state in Chapters 2 and 3.

The numbers in the vector from are coefficients that we multiply by each of the computational basis states which we then add together to make a new state. These coefficients are also called the *amplitudes* for each of the basis states in the sum.

THINKING WITH DIRECTIONS

Another way to think about this example is by thinking back to the maps that we saw in Chapter 2. Each different computational basis state tells us about a direction that a qubit state can be pointed in. We can think of the state of a two-qubit state as a direction in four dimensions instead of the two-dimensional maps we saw in Chapter 2. Since this book is two-dimensional rather than four-dimensional, we unfortunately can't draw a picture here, but sometimes thinking of vectors as directions can be more helpful than thinking of vectors as lists of numbers.

4.7 WHY IS IT HARD TO SIMULATE QUANTUM COMPUTERS?

We have seen above that as the number of bits grow, the number of different states a register can be in grows exponentially. While this won't be a problem for us in playing a nonlocal game with only two qubits, you'll want to use more than just two qubits as you proceed through this book.

When you do, you'll also have exponentially many different computational basis states for your quantum register, meaning that you will need exponentially many different amplitudes in your vectors as you grow the size of your quantum register. To write down the state of a 10-qubit register, you'll need to use a vector that is a list of $2^{10} = 1024$ different amplitudes. For a 20-qubit register, you'll need about a million amplitudes in our vectors. By the time we get to 60 qubits, you'll need about 1.15×10^{18} numbers in our vectors. That's about one amplitude for each grain of sand on the planet.

Table 4.4. How much memory is required to store a quantum state?

# of qubits	# of amplitudes	Memory	Size comparison	
1	2	128 bits		
2	4	256 bits		
3	8	512 bits		
4	16	1 kilobit		
8	256	4 kilobytes	Tap-enabled credit card	
10	1024	16 kilobytes		
20	1,048,576	16 megabytes		
26	67,108,864	1 gigabyte	Raspberry Pi RAM	
28	268,435,456	4 gigabytes	iPhone XS Max RAM	
30	1,073,741,824	16 gigabytes	Laptop or desktop RAM	
40	1,099,511,627,776	16 terabytes		
50	1,125,899,906,842,624	16 petabytes		
60	1,152,921,504,606,846,976	16 exabytes		
80	1,208,925,819,614,629,174,706,176	16 yottabytes	Approximate size of the internet	
410	2.6 × 10 ¹²³	4.2 × 10^{124} bytes	Computer the size of the universe	

In , we summarize what this exponential growth means for us as we try to simulate quantum computers using classical computers like phones, laptops, clusters, or cloud environments. This table shows that even though it's very challenging to reason about quantum computers using classical computers, you can pretty easily reason about small examples. With a laptop or desktop, you can simulate up to about 30 qubits without too much hassle. As you'll see throughout the rest of the book, this is more than enough to understand how quantum programs work, and to understand how quantum computers can be used to solve interesting problems.

DEEP DIVE: Are quantum computers exponentially more powerful?

You may have heard that the number of different numbers we have to keep track of to simulate a quantum computer using a classical computer is why quantum computers are more powerful, or that a quantum computer can store that much information. This isn't exactly true, though. A mathematical theorem known as *Holevo's theorem* tells us that a quantum computer made up of 410 qubits can store

at most 410 classical bits of information, even if it would take us a classical computer about the size of the entire universe to write down the state of that quantum computer.

Put differently, **just because it's hard to simulate a quantum computer doesn't mean that it does something** *useful.* We'll see throughout the rest of the book that it takes a bit of art to figure out how to use a quantum computer to do solve useful problems.

Moreover, we run into exactly the same problem in some classical applications! If we want to keep track of a classical probability distribution, for instance, we would need to write down a number for each possible classical state.

4.8 TENSOR PRODUCTS FOR STATE PREPARATION

Describing quantum registers as vectors describing computational basis states is all well and good, but even if we know the state we want to get to, we need to know how to prepare it. For example, if one player in a nonlocal game has a qubit in the $|0\rangle$ state and the other has a qubit in the $|1\rangle$ state, we can combine those two single-qubit states in a straight-forward way to describe the state of the game as $|01\rangle$. What does it mean to "combine" the states of two (or more) qubits? We can do this by adding one more concept to our mathematics toolbox, called the *tensor product*.

In the same way that we used the product function in above to combine labels for a three classical bit register, we can use the concept of a tensor product, written as $^{\otimes}$, to combine the quantum states for each qubit together to make a state that describes multiple qubits. The output of product was a list of all possible states of those three classical bits. The output of a tensor product is similarly, a state that lists all computational basis states for a quantum register. You can use NumPy to compute tensor products for you; NumPy provides an implementation of the tensor product as the function np.kron, as shown in listing 4.6.

Why kron?

The name np.kron may seem odd for a function that implements tensor products, but the name is short for a related mathematical concept called the "Kronecker product." NumPy's use of kron as short for "Kronecker" follows the convention used in MATLAB, R, Julia and other scientific computing platforms.

Listing 4.6. Using NumPy to create a 2 qubit register state by combining $|0\rangle$ and $|1\rangle$

>>> import numpy as np >>> ket0 = np.array([[1], [0]])

with a tensor product to make |01>.

- (1) We start by defining vectors for the single-qubit states |0\) and |1\), just as we did in Chapters 2 and 3.
- **2**We can build the vector for $|0\rangle \otimes |1\rangle$ by calling NumPy's implementation of the tensor product. For historical reasons, the tensor product is represented in NumPy by the kron function.

- 3 The vector returned by np.kron has a 1 for the entry corresponding to the $|01\rangle$ computational basis state and zeroes everywhere else, so we recognize this vector as being the state $|01\rangle$.
 - This example shows us that |0⟩ ⊗ |1⟩ = |01⟩. That is, if we have the state of each qubit individually, we can combine them by using the tensor product to describe the state of the whole register.

We can combine as many qubits as we want this way; say you had four qubits, all in the $|0\rangle$ state. The register of all four qubits could be described as $|0\rangle \otimes |0\rangle \otimes |0\rangle$.

```
>>> import numpy as np
>>> ket0 = np.array([[1], [0]])
>>> from functools import reduce
>>> reduce(np.kron, [ket0] * 4)
array([[1],
       [0],
       [0],
       [0],
        [0],
        [0],
        [0],
       [0],
        [0],
        [0],
        [0],
       [0],
        [0],
        [0],
       [0],
       [0]])
```

- 1 The reduce function provided with the Python standard library lets us apply a two-argument function like kron between each element in a list. Here, we use reduce instead of np.kron(ket0, np.kron(ket0, np.kron(ket0, ket0))).
- **2** We get back a four-qubit state vector representing $|0\rangle \otimes |0\rangle \otimes |0\rangle \otimes |0\rangle = |0000\rangle$.

WE CAN'T ALWAYS GO THE OTHER WAY! The two-qubit state $(|00\rangle+|11\rangle)/\sqrt{2}$ that we saw earlier cannot be written as the tensor product of two single-qubit states. Multiple-qubit states that can't be written out as tensor products are called entangled. We'll see a lot more about entanglement throughout the rest of this Chapter, and in the rest of the book.

4.9 TENSOR PRODUCTS FOR QUBIT OPERATIONS ON REGISTERS

Now that we know how to use the tensor product to combine quantum states together, what is np.kron actually doing? In essence, the tensor product is a table of every different way of combining its two arguments, as shown in .

Figure 4.3. The tensor product of two matrices shown step by step.

Let's consider a simple example, and The tensor product is take the tensor denoted using the ⊗ product of two operator, just like products are indicated with x. matrices A and B: $C = A \otimes B$ $= \begin{bmatrix} 1 \times B & 3 \times B \\ 5 \times B & 7 \times B \end{bmatrix}$ To start, you first need to make copies of B for each element in A. It's a bit like making a big matrix by tiling $7 \times B$ element in A. It's a bit like making a big matrix by tiling B. $B = \begin{bmatrix} 2 & 4 \\ 6 & 8 \end{bmatrix}$ $[1 \times 2 \quad 1 \times 4 \quad 3 \times 2 \quad 3 \times 4]$ After that, you can expand each copy of B to get the pairs of numbers you need to 7 × 4 multiply together to find the $5 \times 6 \quad 5 \times 8 \quad 7 \times 6$ big matrix. At this point, you can 18 24 find the answer by 28 using ordinary (a.k.a. scalar) multiplication. Each element of the resulting big matrix tells us the product of one element from A and one element from B. That is, the tensor product $A \otimes B$ is made up of all possible products of elements from A and B.

The same tensor product shown in of two matrices in Python is also shown in .

- 1) The matrices A and B are just arbitrary 2 × 2 matrices that we are using as examples here.
- 2 As we saw above, np.kron is NumPy's implementation of the tensor product.
- 3 Note that the order of the arguments to the tensor product matters; although both np.kron(A, B) and np.kron(B, A) contain the same information, the entries in each are ordered quite differently!

Using the tensor product between two matrices, we can find how different quantum operations transform the state of a quantum register. We can also understand how a

quantum operation transforms the state of multiple qubits by taking the tensor product of two matrices instead, letting us understand how your and Eve's moves in a nonlocal game affect your shared state.

For example, we know that we can write $^{|1\rangle}$ as $^{X|0\rangle}$ (that is, the result of applying the x instruction to an initialized qubit). This also gives us another way to write out multiple-qubit states like the $^{|01\rangle}$ state we saw earlier. In this case, we can get $^{|01\rangle}$ by applying an x instruction *only* to the second qubit of a two-qubit register. Using the tensor product, we can find a unitary matrix to represent this:

Listing 4.8. Calculating the tensor product of two matrices.

```
>>> import numpy as np
>>> I = np.array([[1, 0], [0, 1]])
>>> X = np.array([[0, 1], [1, 0]])
>>> IX = np.kron(I, X)
>>> IX
array([[0, 1, 0, 0],
       [1, 0, 0, 0],
       [0, 0, 0, 1],
       [0, 0, 1, 0]])
>>> ket0 = np.array([[1], [0]])
>>> ket00 = np.kron(ket0, ket0)
>>> ket00
array([[1],
       [0],
       [0],
       [0]])
>>> IX @ ket00
array([[0],
       [1],
       [0],
       [0]])
```

- We start by defining a matrix that represents doing nothing to the first qubit, known as the identity matrix 1. Since 1 is hard to write in Python, we use I instead.
- 2 Next, we define the unitary matrix *X* that lets us simulate the x instruction.
- **3** We can combine the two using the tensor product $1 \otimes X$.
- **4** The matrix $1 \otimes X$ consists of two copies of X, representing what happens to the second qubit for each possible state of the first qubit.
- **S**Let's see what happens when we use $\mathbb{1} \otimes X$ to simulate how the x instruction transforms the second qubit in a two-qubit register. We'll start with that register in the $|00\rangle = |0\rangle \otimes |0\rangle$ state.
- (6) We recognize the state we get back as the |01⟩ state from earlier in this section. As expected, that is the state we get by flipping the second qubit from |0⟩ to |1⟩.

Exercise 4.2: Hadamard operation on a two gubit register.

How would you prepare a $|0\rangle$] state? First, what vector would you use to represent the two-qubit state $|+0\rangle = |+\rangle \otimes |0\rangle$? You have an initial two qubit register in the $|00\rangle$ state. What operation should you apply to get the state you want?

HINT: try ($H \otimes \mathbb{1}$) if you are stuck!

DEEP DIVE: Finally proving the no-cloning theorem

Learning that operations on multiple qubits are also represented by unitary matrices lets us finally prove the No-cloning Theorem that we've seen a few times so far. The key insight is that that cloning a state isn't linear, and thus cannot be written as a matrix.

As with many proofs in mathematics, the proof of the No-cloning Theorem works by *contradiction*. That is, we assume the opposite of the theorem, and then show that we get something false as a result of that assumption.

Without further ado, then, we start by assuming that we have some wondrous instruction clone that can perfectly copy the state of its qubit. For instance, if we have a qubit q1 whose state starts in $|1\rangle$, and a qubit q2 whose state starts in $|0\rangle$, then after calling q1.clone(q2), we would have the register $|11\rangle$.

Similarly, if q1 starts in $|\rangle$, then q1.clone(q2) should give us a register in the state $|++\rangle = |+\rangle \otimes |+\rangle$. The problem comes in reconciling what q1.clone(q2) should do in these two cases. We know that any quantum operation other than measurement must be linear, so let's give the matrix that lets us simulate clone a name; C seems pretty reasonable.

Using C, we can break down the case that we want to clone $|+\rangle$ into the case in which we want to clone $|0\rangle$ plus the case in which we want to clone $|1\rangle$. We know that $C|+0\rangle = |++\rangle$, but we also know that $C|+0\rangle = C(|00\rangle + |10\rangle)/\sqrt{2} = (C|00\rangle + C|10\rangle)/\sqrt{2}$. Since `clone` needs to clone $|0\rangle$ and $|1\rangle$ as well as $|+\rangle$, we know that $C|00\rangle = |00\rangle$ and $C|10\rangle = |11\rangle$. That gives us that $C|00\rangle + C|10\rangle/\sqrt{2} = (|00\rangle + |11\rangle)/\sqrt{2}$, but we concluded earlier that $C|+0\rangle = |++\rangle = (|00\rangle + |01\rangle + |10\rangle + |11\rangle)/2$

We thus have a contradiction, and can conclude that we went wrong at the very first step, where we assumed that clone could exist! Thus, we have shown the No-cloning Theorem.

One important thing to note from this argument is that you can always copy information from one qubit to another *if you know the right basis*. The problem came in when we didn't know if we should copy information about $|0\rangle$ vs $|1\rangle$ or $|++\rangle$ vs $|-\rangle$, as we could copy either $|0\rangle$ or $|+\rangle$, but not both. This isn't a problem classically, as we only ever have the computational basis to work with.

4.10 QUTIP OF THE ICEBERG

Up to this point, we've used NumPy to write our qubit simulator, SingleQubitSimulator(). This is very helpful, as without NumPy, we'd need to write our own matrix analysis functions and methods. It is often convenient, however, to rely on Python packages with special support for quantum concepts, building on the excellent numerical support provided by NumPy and SciPy (an extension to the NumPy numerical capabilities).

4.11 QUANTUM OBJECTS IN QUTIP

One particularly useful package is the QuTiP (Quantum Toolbox in Python http://www.qutip.org) package, which provides built-in support for representing states and measurements as bras and kets, respectively, and for building matrices to represent quantum operations.

Just as np.array is at the core of NumPy, all of our use of QuTiP will center around the Qobj class (short for "quantum object"). This class encapsulates vectors and matrices, providing additional metadata and useful methods that will make it easier for

us to improve our simulator. You can see in an example of creating a Qobj from a vector, where it keeps track of some metadata:

4.11.1 Selected metadata that the Qobj class tracks:

- data holds the array representing the Qobj,
- dims the size of our quantum register. You can think of it as a way of keeping track of how we record the qubits we are dealing with,
- shape keeps the dimension of the original object we used to make the Qobj. Similar to the np.shape attribute,
- type what the Qobj represents (a state = ket, a measurement = bra, or an operator = oper).

Figure 4.4. Properties of the Qobj class.

the Hadamard operation.

Let's try importing QuTiP and asking it for the Hadamard operation, see .

NOTE

Make sure as you run things that you are in the right conda env, for more information see Appendix A.

Listing 4.9. Using the QuTiP Python package to load the package's representation of

```
>>> import qutip as qt
>>> H = qt.hadamard_transform()
>>> H
Quantum object: dims = [[2], [2]], shape = (2, 2), type = oper, isherm = True 
Qobj data =
[[ 0.70710678     0.70710678]
[ 0.70710678     -0.70710678]]
```

1 Note that QuTiP prints out some diagnostic information about each Qobj instance along with the data itself. Here, for instance, type = oper tells you that H represents an operator (a more formal term for the matrices we've seen so far), along with some information about the dimensions of the operator

represented by H. Finally, the isherm = True output tells you that H is an example of a special kind of matrix called a Hermitian operator. We'll see more about why Hermitian operators are helpful in Chapter 9.

We can make new instances of <code>Qobj</code> in much the same way as we made NumPy arrays, by passing in Python lists to the <code>Qobj</code> initializer, see .

Listing 4.10. Using a vector to initialize a Qobj that represents the state of a qubit.

```
>>> import qutip as qt
>>> ket0 = qt.Qobj([[1], [0]])
>>> ket0
Quantum object: dims = [[2], [1]], shape = (2, 1), type = ket  2
Qobj data =
[[1.]
    [0.]]
```

- ①One key difference between creating Qobj instances and arrays is that when we create Qobj instances, we always need two levels of lists. The outer list is a list of rows in the new Qobj instance.
- **2** QuTiP prints some metadata about the size and shape of the new quantum object, along with the data contained in the new object. In this case, the data for the new Qobj that you constructed has two rows, each with one column. We identify that as the vector or ket that we use to write down the $|0\rangle$ state.

Exercise 4.3: Creating the Qobj for other states

How would you create a Qobj to represent the $|1\rangle$ state? How about the $|+\rangle$ or $|-\rangle$ state? If you need to check back to Simulating qubits in code section of Chapter 2 for what vectors represent those states.

Where QuTiP really helps out, though, is that it provides us with a lot of nice shorthand for the kinds of objects we need to work with in quantum computing. For instance, we could have also made ket0 in the above sample by using the QuTiP basis function, see .

Listing 4.11. Using built-in features of QuTiP to easily create the $|0\rangle$ and $|1\rangle$ states.

1 The basis function takes two arguments. The first tells QuTiP that we want a qubit state; since a qubit can be in either of two different states, you can pass a 2 here to indicate a single qubit. It's 2 for a

single qubit because the length of a vector that is needed to represent a single qubit is 2. The second argument tells QuTiP which basis state you want. Since we want $|0\rangle$, we pass a 0 here.

- 2 Note that we get exactly the same output here as in the previous example.
- **3** We can also construct a quantum object for $|1\rangle$ by passing a 1 instead of a 0.

...basis?

As we have seen before, the states $|0\rangle$ and $|1\rangle$ make up the computational basis for a single qubit. The QuTiP function basis gets its name from this definition, as it makes quantum objects to represent computational basis states.

There are more things in heaven and earth, than your qubits.

It may seem a little odd that we had to tell QuTiP that we wanted a qubit. After all, what else *could* we want? As it turns out, quite a bit (yes, pun very much so intended)!

There are many other ways to represent classical information than bits, such as *trits*, which have three possible values. We tend not to see classical information represented using anything other than bits when we write programs, though, as it's very useful to pick a convention and stick with it. Things other than bits still have their uses, though, in specialized domains such as telecommunications systems.

In the exact same fashion, quantum systems can have any number of different states, so that we can have qutrits, qu4its, qu5its, qu17its, and so forth, collectively known as qudits. While representing quantum information using qudits other than qubits can be useful in some cases, and can have some very interesting mathematical properties, qubits give us all we need to dive into quantum programming, so we'll focus on them until we get to Part III and meet our first fermions.

Exercise 4.4: Using qt.basis for multiple qubits

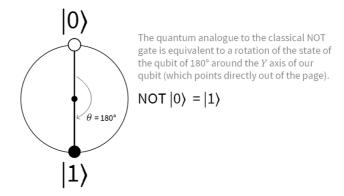
How could you use the qt.basis function to create a two qubit register in the $|10\rangle$ state? How could you create the $|001\rangle$ state? Remember that the second argument to qt.basis is an index to the computational basis states we saw earlier.

QuTiP also provides a number of different functions for making quantum objects to represent unitary matrices. For instance, we can make a quantum object for the X matrix by using the sigmax function:

```
>>> import qutip as qt
>>> qt.sigmax()
Quantum object: dims = [[2], [2]], shape = (2, 2), type = oper, isherm = True
Qobj data =
[[0. 1.]
    [1. 0.]]
```

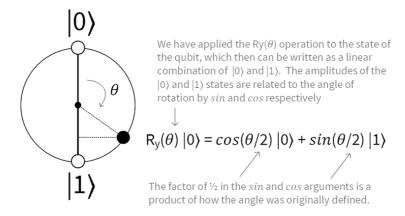
As we saw in Chapter 2, the matrix for sigmax represents a rotation of 180° ().

Figure 4.5. A visualization of the quantum equivalent of a NOT operation operating on a qubit in the $|0\rangle$ state, leaving the qubit in the $|1\rangle$ state.



QuTiP also provides a function ry to represent rotating by whatever angle we like instead of 180° like the x operation. We have actually already seen the operation that ry represents in Chapter 2, when we considered rotating $|0\rangle$ by an arbitrary angle θ . See for a refresher on the operation we now know as ry.

Figure 4.6. A visualization of QuTiP function ry which corresponds to a variable rotation of θ around the Y axis of our qubit (which points directly out of the page).



Now that we have a few more single-qubit operations down, how can we easily simulate multi-qubit operations in QuTiP? We can use QuTiP's tensor function to quickly get up and running with tensor products to make our multi-qubit registers and operations, as we show in listing 4.12.

Listing 4.12. Tensor products in QuTiP

```
>>> import qutip as qt
>>> psi = qt.basis(2, 0)
                               2
>>> phi = qt.basis(2, 1)
>>> qt.tensor(psi, phi)
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
                               3
[[ 0.]
[ 1.]
[ 0.]
[ 0.]]
>>> H = qt.hadamard_transform(1)
>>> I = qt.qeye(2)
>>> qt.tensor(H, I)
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[ 0.70710678 0.
                        0.70710678 0.
[ 0.
                                   0.707106781
             0.70710678 0.
-0.7071067811
[ 0.
             0.70710678 0.
>>> (
       qt.tensor(H, I) * qt.tensor(psi, phi) -
. . .
       qt.tensor(H * psi, I * phi)
. . .
...)
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[ 0.]
[ 0.]
[ 0.]
[ 0.]]
```

- **1** Set psi to represent $|\psi\rangle = |0\rangle$.
- **2** Set phi to represent $|\phi\rangle = |1\rangle$.
- 3 After calling tensor, QuTiP tells us the amplitudes for each classical label in

```
|\psi\rangle\otimes|\phi\rangle=|0\rangle\otimes|1\rangle=|01\rangle , using the same order as .
```

- 4 Set H to represent the Hadamard operation discussed above.
- **5** You can use the qeye function provided by QuTiP to get a copy of a Qobj instance representing the identity matrix that we first saw in . Since identity matrices are often written using the letter "I," many scientific computing packages use the name eye as a bit of a pun to refer to the identity matrix.
- **6** The unitary matrices representing quantum operations combine using tensor products in the same way as states and measurements.
- The same apply a unitary to a state then take the tensor product, we get the same answer as if we applied the tensor product and then the unitary. Here we can tell these two approaches are give the same answer because their difference is 0. In math, we would say that for any unitary operators U and V and for any states $|\psi\rangle$ and $|\phi\rangle$, $(U|\psi\rangle)\otimes(V|\phi\rangle) = (U\otimes V)(|\psi\rangle\otimes|\phi\rangle)$.

NOTE

For a list of all of the built-in states and operations QuTiP has, see http://qutip.org/docs/latest/guide/guide-basics.html#states-and-operators

4.12 UPGRADING THE SIMULATOR

The goal now is to use QuTiP to upgrade our single-qubit simulator to a multi-qubit simulator with some of the features of QuTiP. We will do this by adding a few features to our single-qubit simulator from Chapters 2 and 3.

The most significant change we'll need to make to our simulator from previous chapters is that we can no longer assign a state to each qubit. Rather, we must assign a state to the entire *register* of qubits in our device, since some of the qubits may be entangled with each other.

Let's jump into making the modifications necessary to separate the concept of the state to the device level.

NOTE

To see the code we wrote earlier, as well as the samples for this chapter see the GitHub repo for the book: https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp.

To review, we have two files for our simulator: the interface (interface.py), and the simulator itself (simulator.py). The device interface (QuantumDevice) defines a way of interacting with an actual or simulated quantum device, which is represented in Python as an object that lets us allocate and deallocate qubits.

We won't need anything new for the QuantumDevice class in the interface in order to model our CHSH game, since we'll still need to allocate and deallocate qubits. Where we can actually add features is in the Qubit class provided along with our SingleQubitSimulator in simulator.py.

Now we need to consider what, if anything, needs to change in our interface for a Qubit we allocate from the QuantumDevice. In Chapter 2, we saw that the Hadamard operation was useful for rotating qubits between different bases to make a QRNG. Let's build on this by adding a new method to Qubit to allow quantum programs to send a new kind of rotation instruction that we will need to use the quantum strategy for CHSH.

Listing 4.13. The interface we have built in Chapter 3 for our Qubit, but adding a

new ry operation which rotates our state about the Y-axis.

class Qubit(metaclass=ABCMeta): @abstractmethod def h(self): pass @abstractmethod def x(self): pass @abstractmethod def ry(self, angle: float): pass

```
@abstractmethod
def measure(self) -> bool: pass
@abstractmethod
def reset(self): pass
```

1 This is the abstract method ry which takes an argument angle to specify how far to rotate the qubit around the *Y*-axis.

That should then cover all changes we need to make to our Qubit and QuantumDevice interface for playing CHSH with Eve. Now we need to address what changes we need to make to the simulator.py to actually allow it to allocate, operate, and measure multi-qubit states.

The main changes to our Simulator class which implements a QuantumDevice are that we now need some attributes to track how many qubits it has and the overall state of the register. shows these changes as well as an update to allocation and deallocation methods.

Listing 4.14. The implementation of Simulator, a class that represents a device that

represents a multi-qubit version of a QuantumDevice from our interface. class Simulator(QuantumDevice): capacity: int available_qubits: List[SimulatedQubit] register state: qt.Qobj def __init__(self, capacity=3): self.capacity = capacity self.available_qubits = [SimulatedQubit(self, idx) for idx in range(capacity) self.register state = qt.tensor(qt.basis(2, 0) for _ in range(capacity) def allocate qubit(self) -> SimulatedQubit: if self.available_qubits: return self.available qubits.pop() def deallocate qubit(self, qubit: SimulatedQubit): self.available_qubits.append(qubit)

- We have changed the name from SingleQubitSimulator to Simulator to indicate it is more generalized. Here, that means we can actually simulate multiple qubits with it.
- 2 The more general Simulator class now needs a few attributes, the first being capacity which represents the number of qubits it is capable of simulating.
- 3available_qubits is a list containing the qubits the Simulator is using.
- Aregister state uses the new QuTiP Qobj to represent the state of the entire simulator.

- **⑤**A list comprehension allows us to make a list of available qubits by calling SimulatedQubit with the indices from the range of capacity.
- **6** register_state is initialized by taking the tensor product of a number of copies of the |0⟩ state equal to the capacity of the simulator. The *[...] notation will take the generated list and turn it into a sequence of arguments for gt.tensor.
- The allocate_qubit and deallocate_qubit methods are the same as our previous simulator from Chapter 3.

PEER NOT INTO THE BOX, MORTAL!

In just the same way that we used NumPy to represent the state of a simulator, the register_state property of our newly upgraded simulator uses QuTiP to predict how each instruction has transformed the state of our register. When we write quantum programs, though, we do so against the interface in , which doesn't have any way to let us access register_state.

We can think of the simulator as being a kind of black box that encapsulates the notion of a state. If our quantum programs were able to look inside that box, then they would be able to cheat by copying that information in ways forbidden by the no-cloning theorem. This means that for a quantum program to be correct, we cannot look inside the simulator to see its state.

In this Chapter, we'll cheat a little bit, but in the next Chapter, we'll see how to fix that up to make sure our programs can be run on actual quantum hardware.

We also will add a new *private* method to our Simulator that allows us to apply operations to specific qubits in our device. This will let us write methods on our qubits that send operations back to the simulator to be applied to the state of an entire register of qubits.

TIP

Python is not strict about keeping methods or attributes private, but we will prefix this method name with an underscore to indicate it is meant for use in the class only.

Listing 4.15. One additional method for the Simulator class that allows us to apply

operations to specific qubits in our simulator.

The private method _apply takes an input unitary of type Qobj that represents a unitary operation to be applied, and a list of int to indicate the indices of the available_qubits list where you want to apply the operation. For now, that list will only ever contain one element, since we're only implementing single-qubit gates in our simulator. We'll relax this in the next Chapter, though.

- 2 If the operation is only to be applied to one qubit, we can use QuTiP to generate the matrix that we would need to apply the operation to the right qubit given the representation of the entire register by applying 1 everywhere else. This is done for us automatically by the gate_expand_1toN function.
- 3 Now that we have the right matrix to multiply our entire register_state by, we can update the value of that register accordingly.

Now let's get to the implementation of SimulatedQubit, the class that represents how we simulate a single qubit, given that we know it is part of a device that has multiple qubits. The main difference between the single and multi-qubit versions of SimulatedQubit is that we need each qubit to remember their "parent" device and their location or id in that device so that we can associate the state with the register and not each qubit. This is important as we will see in the next section when we want to measure qubits in a multi-qubit device.

Listing 4.16. The start of the class SimulatedQubit which implements a qubit from a

multi-qubit device.

```
class SimulatedQubit(Qubit):
    qubit_id: int
    parent: "Simulator"

def __init__(self, parent_simulator: "Simulator", id: int):
        self.qubit_id = id
        self.parent = parent_simulator

def h(self) -> None:
        self.parent._apply(H, [self.qubit_id])

def ry(self, angle: float) -> None:
        self.parent._apply(qt.ry(angle), [self.qubit_id])

def x(self) -> None:
        self.parent._apply(qt.sigmax(), [self.qubit_id])

3
```

- ①Now to initialize a qubit we will need the the name of the parent simulator (so we can easily associate it) and the index of the qubit in the simulator's register. __init__ then sets those attributes and resets the qubit to the |0⟩ state.
- 2 To implement the h operation now, we ask the parent of our SimulatedQubit (which is an instance Simulator) to use the _apply method to generate the right matrix that would represent the operation on the complete register, and then update the register state.
- **3** We can also pass the parameterized qt.ry operation from QuTiP to _apply to rotate our qubit about the *Y*-axis by an angle angle.

Great! Now we are almost done upgrading our simulator to use QuTiP and support multiple qubits. We will tackle simulating measurement on multi-qubit states, in the next section.

4.13 MEASURING UP: HOW CAN WE MEASURE MULTIPLE QUBITS?

TIP

This section is one of the hardest sections in the book, please don't worry if things don't make a lot of sense the first time around.

In some sense, measuring multiple qubits works the same way we're used to from measuring single-qubit systems. We can still use Born's rule to predict the probability of any particular measurement outcome. For example, let's return to the $(|00\rangle+|11\rangle)/\sqrt{2}$ state that we've seen a few times now. If we were to measure a pair of qubits in that state, we would get either "00" or "11" as our classical outcomes with equal probability, since both have the same amplitude, namely $1/\sqrt{2}$.

Similarly, we'll still demand that if we measure the same register twice in a row, we get the same answer. If we get the "00" outcome, for instance, we know that qubits are left in the $|00\rangle = |0\rangle \otimes |0\rangle$ state.

Where this gets a little bit trickier, however, is if we measure *part* of a quantum register without measuring the whole thing. Let's look at a couple examples to see how that could work. Again taking $(|00\rangle + |11\rangle)/\sqrt{2}$ as an example, if we measure *only* the first qubit and we get a "0", we know that we need to get the same answer again the next time we measure. The only way this can happen is if the state transforms to $|00\rangle$ as a result of having observed "0" on the first qubit.

On the other hand, what happens if we measure the first qubit out of a pair of qubits in the $|++\rangle$ state? It's helpful to first refresh our memory as to what $|++\rangle$ looks like when written as a vector.

```
>>> import qutip as qt
>>> ket_plus = qt.hadamard_transform() * qt.basis(2, 0)
>>> ket_plus
Quantum object: dims = [[2], [1]], shape = (2, 1), type = ket
Qobj data =
[[0.70710678]
[0.70710678]]
>>> register_state = qt.tensor(ket_plus, ket_plus)
>>> register_state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[0.5]
[0.5]
[0.5]
[0.5]
```

- ①Let's start again by writing down |+> as H|0>. In QuTiP, we'll use the hadamard_transform function to get a Qobj instance to represent H, and we'll use basis(2, 0) to get a Qobj representing |0>.
- 2We can print out ket_plus to get a list of the elements in that vector; as before, we call each of these elements an amplitude.
- 3 To represent the state $|++\rangle$, we use that $|++\rangle = |+\rangle \otimes |+\rangle$.

4 This output tells us that |++⟩ has the same amplitude each of the four computational basis states |00⟩, |01⟩, |10⟩, and |11⟩, just as ket_plus has the same amplitude for each of the computational basis states |0⟩ and |1⟩.

Suppose that we measure the first qubit and get a "1" outcome. To make sure we get the same result the next time we measure, the state after measurement can't have any amplitude on $|00\rangle$ or $|01\rangle$. If we only keep the amplitudes on $|10\rangle$ and $|11\rangle$ (the third and four rows of the vector we calculated above), then we get that the state of our two qubits becomes $(|10\rangle + |11\rangle)/\sqrt{2}$.

WHERE DID THE √2 COME FROM?

We included a $\sqrt{2}$ above to make sure that all our measurement probabilities still sum to one when we measure the second qubit. In order for Born's rule to make any sense, we always need the sum of the squares of each amplitude to sum to one.

There's another way to write this state, though, that we can check using QuTiP:

1 Recall that we can write $|+\rangle$ as $H|0\rangle$.

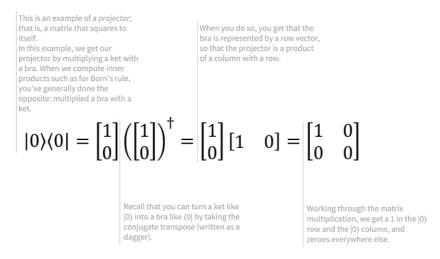
This tells us that if we only keep the parts of the state $|++\rangle$ that are consistent with getting a "1" outcome from measuring the first qubit, then we get $|1+\rangle = |1\rangle \otimes |+\rangle$. That is, nothing happens at all to the second qubit in this case!

Exercise 4.5: Measuring the other qubit

In the example where our two qubits start off in the $\frac{|++\rangle}{|++\rangle}$ state, suppose we measured the second qubit instead. Check that no matter what result we get, nothing happens to the state of the first qubit.

To work out measuring part of a register more generally, we can use another concept from linear algebra called *projectors*. A projector is the product of a state vector (the "ket" or $|\rangle$ part of a bra-ket) and a measurement (the "bra" or $\langle |$ part of a bra-ket), and represents our requirement that *if* a certain measurement outcome occurs, *then* we must transform to a state consistent with that measurement. See for a quick example of a single-qubit projector; defining projectors on multiple quibts works in exactly the same way.

Figure 4.7. An example of a projector acting on a single-qubit state.



$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \alpha \\ 0 \end{bmatrix}$$
 When we multiply an arbitrary state vector by this projector, it projects out or filters out all but the part corresponding to the $|0\rangle$ computational basis state. What you're left with might no longer be a valid state, as it might not have the right length. The length tells us about how much of the original state vector was picked out by our projector. In your simulator, you'll use this to find the probability for each measurement outcome, giving us another way to compute Born's rule.

In QuTiP, we write the bra corresponding to a ket by using the <code>.dag()</code> method (short for dagger, a call back to mathematical notation we saw in). Thankfully, even if the math isn't that straightforward, it winds up not being that bad to write in Python, as we can see in .

Listing 4.17. A method for the Simulator class that allows us to measure individual

```
def measure(self) -> bool:
    projectors = [
        qt.circuit.gate_expand_1toN(
            qt.basis(2, outcome) * qt.basis(2, outcome).dag(),
            self.parent.capacity,
            self.qubit_id
        )
        for outcome in (0, 1)
    ]
    post_measurement_states = [
        projector * self.parent.register_state
        for projector in projectors
```

```
probabilities = [
          post_measurement_state.norm() ** 2
          for post_measurement_state in post_measurement_states
]
sample = np.random.choice([0, 1], p=probabilities)
self.parent.register_state = post_measurement_states[sample].unit()
return int(sample)

def reset(self) -> None:
    if self.measure(): self.x()
```

- We start by using QuTiP to make a list of projectors, one for each possible measurement outcome. Don't worry too much about this for now, we'll get more practice with this concept as we go along.
- ②Just as you did in , you can use the gate_expand_1toN function provided by QuTiP again here to expand each single-qubit projector into a projector that acts on the whole register.
- 3 Next, we use each projector to pick out the parts of a state that are consistent with each measurement outcome.
- The length of what each projector picks out (written as the .norm() method in QuTiP) tells us about the probability of each measurement outcome.
- **5**Once we have the probabilities for each outcome, we can pick an outcome by using NumPy.
- **(i)** Next, you can use the .unit() method built-in to QuTiP to make sure that whatever state we get, its measurement probabilities still sum up to one so that you're ready for the next measurement.
- **7** Finally, you can use the new measurement method to implement the reset method from Chapters 2 and 3. Here, if the result of a measurement is |1⟩, then flipping with an x instruction resets back to |0⟩.

4.14 CHSH: QUANTUM STRATEGY

Now that we have expanded our simulator to handle multiple qubits, let's see how we can simulate a *quantum* based strategy for our players that will result in a win probability higher than could be possible with any classical strategy! See for a reminder of how the CHSH game is played.

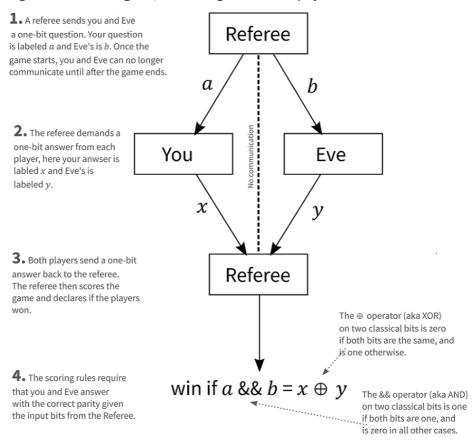


Figure 4.8. The CHSH game, a nonlocal game with two players and a referee.

You and Eve now have quantum resources, so let's start with the simplest option: you each have one qubit allocated from the same device. We'll use our simulator to implement this strategy, so this isn't really a test of quantum mechanics itself so much as that our simulator agrees with quantum mechanics.

NOTE

We can't simulate the players being truly nonlocal, as the parts of the simulator need to communicate in order to emulate quantum mechanics. Faithfully simulating quantum games and quantum networking protocols in this manner is exposes a lot of interesting classical networking topology questions that are well beyond the scope of this book. If you're interested in simulators intended more for use in quantum networking than quantum computing, we recommend looking at the SimulaQron project (http://www.simulaqron.org/) for more information.

Let's see how often you and Eve can win if you each start off with a single qubit, and if those qubits start off in the $(|00\rangle + |11\rangle)/\sqrt{2}$ state that we've seen a few times so far in this Chapter. Don't worry about how to prepare this state, we'll see how to do that in the next Chapter. For now, let's just see what you can do with qubits in that state once

you have them.

Using these qubits, we can form a new quantum strategy for the CHSH game we saw at the start of the chapter. The trick is that you and Eve can each apply operations to each of your qubits once you get your respective messages from the referee. As it turns out, ry is a very useful operation for this strategy. It lets you and Eve trade off slightly between how often you win when the referee asks for you both to output the same answers (the 00, 01, and 10 cases), in order to do slightly better when you need to output different answers (the 11 case), as shown in and in .

Figure 4.9. Rotating qubits to win at CHSH

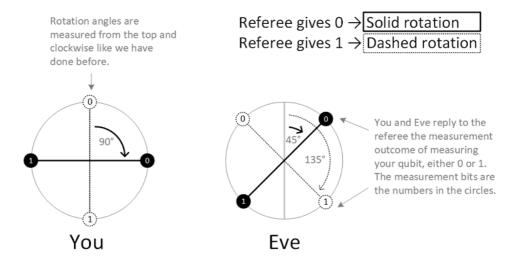


Table 4.5. Rotations you and Eve will do to your qubits as a function of the input bit you receive from the referee. Note they are all ry rotations, just by different angles (converted to radian for ry).

Input from referee Your Rotation		Eve's Rotation	
0	ry(90 * np.pi / 180)	ry(45 * np.pi / 180)	
1	ry(0)	ry(135 * np.pi / 180)	

• Don't worry if these angles look random, we can check to see that they work using our new simulator! In, we've used the new features we added to the simulator to write out a quantum strategy.

Listing 4.18. A new quantum strategy for playing the CHSH game where you and Eve

```
start with a pair of qubits.
import qutip as qt
def quantum_strategy(initial_state: qt.Qobj) -> Strategy:
```

```
shared_system = Simulator(capacity=2)
shared_system.register_state = initial_state
your_qubit = shared_system.allocate_qubit()
eve_qubit = shared_system.allocate_qubit()

shared_system.register_state = qt.bell_state()
your_angles = [90 * np.pi / 180, 0]
eve_angles = [45 * np.pi / 180, 135 * np.pi / 180]

def you(your_input: int) -> int:
    your_qubit.ry(your_angles[your_input])
    return your_qubit.measure()

def eve(eve_input: int) -> int:
    eve_qubit.ry(eve_angles[eve_input])
    return eve_qubit.measure()

return you, eve
```

- 1 To start the quantum strategy, we need to create a QuantumDevice instance where we will simulate our qubits
- 2 Labels can be assigned to each qubit as we allocate them to the shared system
- **3** We will cheat a little here to just set the state of our qubits to the entangled state $(|00\rangle + |11\rangle) / \sqrt{2}$, we will see in the next chapter how to prepare this state from scratch, and why the function to prepare this state is called bell state.
- ♠The angles for the rotations you and Eve need to do based on your input from the referee.
- **⑤**Your strategy for playing the CHSH game starts with you rotating your qubit based on the input classical bit from the referee.
- **6** The classical bit value your strategy returns is the bit value you get when you measure your qubit.
- 7 Eve's strategy is similar to yours, just uses different angles for her initial rotation.
- ③Just like our classical strategy, quantum_strategy returns a tuple of functions that represent your and Eve's individual actions.

Now that we have implemented a Python version of the quantum_strategy, let's see how often we can win with using our CHSH game <code>est_win_probability</code> function in.

Listing 4.19. We can run our CHSH game win estimator with the new quantum_strategy.

Note that this is is a higher win probability than with the classical strategy!

```
>>> est_win_probability(quantum_strategy)
0.832
```

①Note that you may get slightly more or less than 85% when you try this. This is because the win probability is estimated under the hood using a binomial distribution. For this example, we'd expect error bars of about 1.5%.

Awesome, you and Eve can start winning the CHSH game more frequently than any other classical players! What this strategy shows us, though, is an example of how states like $(|00\rangle + |11\rangle)/\sqrt{2}$ are an important resource provided by quantum mechanics.

NOTE

States like $(|00\rangle+|11\rangle)/\sqrt{2}$ are called entangled, as they can't be written as the tensor product of single-qubit states. We'll see many more examples of entanglement as we go along, but entanglement is of the most amazing and fun things that we get to use in writing quantum programs!

As we saw in this example, entanglement allows us to create correlations in our data that can be used to our advantage when we want to get useful information out of our quantum systems.

THE SPEED OF LIGHT IS STILL A THING

If you've read about relativity (if you haven't, no worries!), you may have heard that it's impossible to send information faster than the speed of light. It may seem what you've seen about entanglement so far that quantum mechanics violates this, but as it turns out, entanglement can never be used to communicate a message of your choosing all on its own. You always need to send something else along with using your entanglement. This means that the speed of light still constrains how fast information can travel through the universe — phew!

Far from being strange or weird, though, entanglement is a direct result of what we've already learned about quantum computing so far: it is a direct consequence of quantum mechanics being **linear**. If we can prepare a two qubit register in the $|00\rangle$ state and in the $|11\rangle$ state, then we can also prepare a state in a linear combination of the two, such as $(|00\rangle + |11\rangle) / \sqrt{2}$.

Since entanglement is an direct result of the linearity of quantum mechanics, the CHSH game also gives us a great way to check that quantum mechanics is really correct (or to the best our data can show). Let's go back to that win probability in again. If you do an *experiment*, and you see something like an 83.2% win probability, that tells you that your experiment couldn't have been purely classical, since we know a classical strategy can *at most* win 75% of the time. This experiment has been done many times throughout history, and is part of how we know that our universe isn't just classical — that we need quantum mechanics to describe it.

NOTE

In 2015, one experiment even had the two players in the CHSH game separated by over a kilometer!

Self-testing: An application for nonlocal games

This hints at another application for nonlocal games: if we can play *and win* a nonlocal game with Eve, then that means along the way, we must have built something that we can use to send quantum data. This sort of insight leads to ideas known as *quantum self-testing*, where we make parts of a device play nonlocal games with other parts of the device in order to make sure that it works correctly.

The simulator that you wrote in this chapter gives you everything you need to see how those kinds of experiments work, so that we can plow ahead to use quantum mechanics and our qubits to do awesome stuff, armed with the knowledge that quantum mechanics really is how our universe works.

4.15 SUMMARY

• In this chapter you learned:

- Simulate state preparation, operations, and measurement results for multiple qubits,
- Program a simulator for multiple qubits leveraging the QuTiP Python package,
- Recognize the proof quantum mechanics is consistent with our observations of the universe by simulating experimental results.

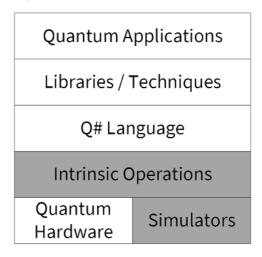
The name "CHSH" comes from the initials of the researchers who originally created the game, Clauser, Horne, Shimony, and Holt. You can find the original if you are interested: https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.23.880).

Teleportation and entanglement: Moving quantum data around

This chapter covers:

- Implement a quantum program to move data around a quantum computer using classical and quantum control
- · Recognize a new way of visualizing single qubit operations called the Bloch sphere
- · Predict the output of two-qubit operations, and Pauli operations

Figure 5.1. In this Chapter, we will be covering topics in simulating quantum hardware and intrinsic operations for the quantum devices.



5.1 MOVING QUANTUM DATA

Just as with classical computing, sometimes in a quantum computer you have some data *here* that you would very much appreciate being somewhere over *there*. Classically, this is an easy problem to solve by copying data around, but as we saw in Chapters 3 and 4, the No-Cloning Theorem means we in general **can't** copy data stored in qubits.

Moving data classically

In some parts of classical computing, we run into the same problem of not being able to copy information for very different reasons. Copying data in a multi-threaded application can introduce subtle race conditions, while performance considerations can prompt us to reduce the amount that we copy data. The solution embraced by many classical languages (e.g.: C++11 and Rust) is to instead focus on *moving* data. Thinking in terms of moving data is helpful in quantum computing, though we'll implement moves in a very different way.

So what can you do if you want to move data around in a quantum device? Thankfully, there's a number of different ways to move quantum data instead of copying it. In this chapter, we'll see a few of these approaches and will add the last couple features that we need to our simulator to implement them. Let's get to sharing quantum information then!

Suppose Eve has some qubits that encode some data she'd like to share with you.

Hey player! I have some quantum information I want to share with you, can I send it over to you?

```
-- Eve
```

Here, Eve is referring to the swap instruction — it's a bit different than the instructions we've seen so far in that it operates on *two* qubits at once. By contrast, every operation that you've seen so far only operates on one qubit at a time.

Looking at what swap does, the name is pretty descriptive, because it literally swaps the state of two qubits in the same register. For example, say you had two qubits in the $|01\rangle$ state. If you use the swap instruction on both qubits, the result would be your register is now in the $|10\rangle$ state. Let's look at for an example of using the swap matrix built-in to QuTiP.

Listing 5.1. Using the built in swap function in QuTiP, and an example of using it on the

$|+0\rangle$ to get the $|0\rangle+\rangle$ state.

```
>>> import qutip as qt
>>> ket_0 = qt.basis(2, 0)
>>> ket_plus = qt.hadamard_transform() * ket_0
>>> initial_state = qt.tensor(ket_plus, ket_0)
>>> initial_state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
```

```
2
[[0.70710678]
[0.
 [0.70710678]
[0.
       ]]
>>> swap_matrix = qt.swap()
>>> swap_matrix * initial_state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[0.70710678]
[0.70710678]
[0.
[0.
            ]]
>>> qt.tensor(ket_0, ket_plus)
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[0.70710678]
 [0.70710678]
 [0.
[0.
            11
```

- ①We'll start by using qt.basis, qt.hadamard_transform, and qt.tensor together to define a variable for old friend from last chapter, the state vector $|+0\rangle$. This state makes a nice example for seeing what the swap instruction does.
- 2 As we saw in Chapter 4, this state has equal amplitudes on the |00⟩ and |10⟩ computational basis states.
- 3 We can get a copy of the unitary matrix for the swap instruction by calling qt.swap.
- The same way that we've simulated single-qubit operations, we can multiply our state by the unitary matrix for swap to find the state of our two-qubit register after applying a swap instruction.
- (5) When we do so, we see that we end up in a superposition between |00⟩ and |01⟩ instead of between |00⟩ and |10⟩.
- **6** You can quickly check that the result of using the swap instruction on a register of two qubits that start off in the $|+0\rangle$ state is $|0+\rangle$.

Looking at , you can see that the swap instruction did pretty much what its name suggests. In particular, swap took two qubits that started off in the state $|+0\rangle$ to the $|0+\rangle$ state. More generally, you can read off what the swap instruction does by looking at the unitary matrix we used to simulate it above (see listing 5.2).

Listing 5.2. Looking at the unitary matrix for the swap instruction

```
>>> import qutip as qt
>>> qt.swap()
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True

Qobj data =
[[1. 0. 0. 0.]

[0. 0. 1. 0.]

[0. 1. 0. 0.]

[0. 0. 0. 1.]
```

- 4
- ①Note that the unitary matrix we use to simulate the swap instruction is a 4 × 4 matrix. This is because it acts on two-qubit states, meaning that it has to define what happens to each of the four possible computational basis states that a two-qubit register can be in.
- 2 Each column of this unitary matrix tells us what happens to one of the computational basis states; for instance, the first column tells us that the |00⟩ state is mapped to the vector [[1], [0], [0], [0]], which we recognize as |00⟩. That is, the swap instruction does nothing to qubits that start off in the |00⟩ state.
- 3)On the other hand, we can see that that the |01) and |10) states are swapped by the swap instruction.
- 4 Finally, the swap instruction also leaves (11) alone.

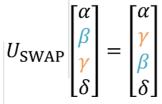
IMPORTANT

The unitary matrix for the swap instruction cannot be written as the tensor product of any two single-qubit unitary matrices. That is, you can't understand what swap does by considering one qubit at a time — you need to work out what it does to state of the pair of qubits that the swap instruction acts on.

As shown in , you can see what swap does in general, no matter what state our two qubits start off in.

Figure 5.2. The two qubit operation swap exchanges the states of two qubits in a register.

Since the swap instruction acts on two-qubit register, the state of each such register is represented by a four-element vector. Just as you might use x or y as variables in algebra, we tend to use Greek letters like α and β to stand for arbitrary amplitudes. Here, α , β , γ , and δ are the amplitudes for each of the $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$ computational basis states, respectively.



We can simulate how the swap instruction changes the state of a register by multiplying the state of that register by a *unitary matrix*. As always, this matrix isn't the swap instruction, but a tool we can use to simulate swap.

After running the swap instruction, the amplitudes for $|01\rangle$ and $|10\rangle$ are swapped, while the amplitudes for $|00\rangle$ and $|11\rangle$ stay the same, since swapping $0 \rightleftharpoons 0$ or $1 \rightleftharpoons 1$ doesn't do anything.

Remember that in Chapter 2, you saw that a unitary matrix is a lot like a truth table. That is, unitary matrices like what we get back from qt.swap are useful in that they

help us simulate what the swap instruction does. Just as a classical adder isn't its truth table, though, it's helpful to remember that these unitary matrices aren't quantum programs, but tools that we use to simulate how quantum programs work.

Exercise 5.1: swap the second and third qubits in a register.

Say you have a register with 3 qubits, in the state $|01+\rangle$. Using QuTiP, write out this state, and then swap the second and third qubits so your register is in the $|0+1\rangle$ state.

HINT: since nothing will happen to the first qubit, make sure to tensor an identity matrix to qt.swap to build up the correct operation for your register.

At this point, though, Eve is positively *dying* waiting to send us her qubits. Let's go on and add what we need to our simulator, so as to not keep her waiting any longer!

5.2 SWAPPING OUT OUR SIMULATOR

The simulator you worked on in Chapter 4 needs just a couple of tweaks to be able to use two qubit operations like swap. The changes we will need to make are:

- modify apply to work with two qubit operations,
- add the swap and other two-qubit instructions, and
- add the rest of the single qubit rotation instructions.

As you saw in Chapter 4, if you have a matrix that acts on a single-qubit register, you can use QuTiP to apply that to a register with an arbitrary number of qubits by using the <code>gate_expand_ltoN</code> function. This makes sure to tensor in the identity operators everywhere else *but* the qubits you're working with.

In the same way, you can call QuTiP's <code>gate_expand_2toN</code> function to turn two-qubit unitary matrices into matrices that you can use to simulate how two-qubit operations like <code>swap</code> transform the state of a whole register. Let's go on and add that into our simulator now:

TIP

We've made a couple small changes to the code in this Chapter to help make printed outputs look a little nicer. All of these changes, along with all the samples for this and other chapters, are on the GitHub repo for this book at https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp.

Listing 5.3. Using gate_expand_2toN to apply two-qubit unitary matrices to register

```
self.register_state = matrix * self.register_state
```

- 1 To simulate two-qubit operations, we need two indices for qubits in the register; one for each qubit that our instruction acts on.
- 2 The call to gate_expand_2toN looks very similar to our call to gate_expand_1toN, except that we pass a 4 × 4 matrix instead of a 2 × matrix.

We saw that QuTiP provides the function swap to give us a copy of the unitary matrix that simulates the swap instruction. This can be used to pretty quickly add the swap instruction to your simulator using the changes you made to Simulator. apply; see.

Listing 5.4. Adding a swap instruction to your simulator.

- To get the 4 × 4 unitary matrix we need to pass to _apply, we just use the qt.swap function we've seen a few times so far in this Chapter.
- 2 Next, we just need to make sure we pass the indices for both qubits that we want to swap between. This will make it so that gate_expand_2toN correctly applies the unitary matrix for our new swap instruction to the state of an entire register.

While we're working with the simulator anyway, let's add one more instruction to let us print out the simulator's state more easily without having to access its internals, see.

Listing 5.5. Adding a dump instruction to the simulator allows you to more conveniently

print the state of the register.

TIP

```
def dump(self) -> None:
    print(self.register_state)
```

This way, you can ask the simulator to help you out in debugging out quantum programs, but in a way that can be safely stripped out for devices that don't support that (e.g.: actual quantum hardware).

Remember that a qubit is not a state, a state is just a convenient way of representing how the quantum system will behave.

With both these changes in place, you're now all set to use the swap instruction. Let's use it to repeat our experiment before, where we swapped two qubits starting off in the $|0+\rangle$ state to transform them into the $|+0\rangle$ state, see .

Listing 5.6. Test out your simulator's new swap instruction on the |0+) state.

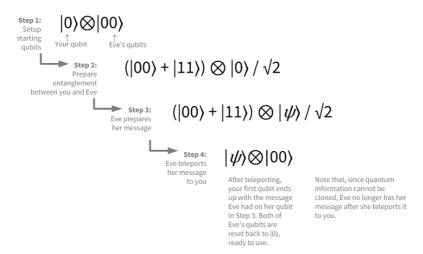
```
>>> from simulator import Simulator
>>> sim = Simulator(capacity=2)
>>> with sim.using_register(n_qubits=2) as (you, eve):
                                                                       0
        eve.h()
        sim.dump()
. . .
        you.swap(eve)
. . .
        sim.dump()
. . .
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket 2
Oobi data =
[[0.70710678]
[0.70710678]
 Γ0.
 [0.
            ]]
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket 3
Qobj data =
[[0.70710678]
 [0.
 [0.70710678]
```

- Since we'll be working a lot with multiple-qubit registers in this Chapter, we've added a new convenience method that lets us allocate several qubits at once. See the code samples for this book for details.
- **2** The first dump comes from our first call to sim.dump, and confirms that eve.h() prepared your qubits in the $|0+\rangle$ state.
- 3 After calling you.swap(eve), the next dump confirms for us that we ended up with the information Eve had prepared on her qubit, so that your qubit ends up in the |+> state, with Eve's qubit ending the way yours started: in the |0> state.

Great, you now have a way to share quantum data with Eve! Well, at least as long as you're sharing a single quantum device, so that you can apply the swap instruction to both your qubits at the same time.

What happens if we want to share quantum information between devices? Thankfully, quantum computing gives us a way for you to send each other your qubits by only communicating classical data, so long as you both start with some entanglement between your qubits. Like much in quantum computing, this technique is given a whimsical name, in this case, "quantum teleportation." Don't let the name fool you, however. When you get right down to it, teleportation uses what we learned in Chapter 4 to let you share quantum data in a useful way. shows a list of the steps in a teleportation program.

Figure 5.3. The steps to the teleportation program.



What's really neat about teleportation is that, while you and Eve still need to do some two-qubit operations between your respective qubits, Eve can decide what data she wants to send you *after* you've done those operations. This means you could prepare the entangled qubits before you needed to exchange quantum data, and just use them as needed.

Using the simulator that you've developed throughout the past several chapters, you might write up teleportation with a quantum program like that shown in .

```
Listing 5.7. A teleportation program in Python.

def teleport(msg : Qubit, here : Qubit, there : Qubit) -> None:
    here.h()
    here.cnot(there)

    msg.cnot(here)
    msg.h()

if msg.measure(): there.z()
    if here.measure(): there.x()

    msg.reset()
    here.reset()
```

You may notice a few new instructions in this program, though. In the rest of this Chapter, you'll see the other pieces you need to get up and running with quantum teleportation using your simulator.

5.3 WHAT OTHER TWO-QUBIT GATES ARE THERE?

As you may guess, swap is not the only two-qubit operation. Indeed, as you can see

from , in order to get teleportation working, we need to add another two-qubit instruction to your simulator, called <code>cnot</code>. The <code>cnot</code> instruction does something similar to <code>swap</code>, except that it switches the $|10\rangle$ and the $|11\rangle$ computational basis states instead of the $|01\rangle$ and $|10\rangle$ states. Another way to think of this is that <code>cnot</code> flips the second qubit *controlled* on the state of the first qubit being $|1\rangle$. Indeed, this is where the name <code>cnot</code> comes from: it's shorthand for "controlled-NOT."

TIP

We often call the first qubit passed to a cnot instruction the control qubit, and the second qubit the target qubit. As we will see in Chapter 7, though, there is a bit of subtlety to these names.

Let's jump right in and see how the cnot instruction works by applying it to that lovely example, the $|+0\rangle$ state.

```
>>> import qutip as qt
>>> ket_0 = qt.basis(2, 0)
>>> ket_plus = qt.hadamard_transform() * ket_0
>>> initial state = qt.tensor(ket plus, ket 0)
>>> qt.cnot() * initial_state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[0.70710678]
 [0.
 [0.
 [0.70710678]]
>>> qt.cnot()
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[1. 0. 0. 0.]
 [0. 1. 0. 0.]
 [0. 0. 0. 1.]
[0. 0. 1. 0.]]
```

- 1 QuTiP provides the unitary matrix for the cnot instruction as the qt.cnot function.
- ②Applying the cnot instruction to two qubits that start off in the $|+0\rangle = (|00\rangle + |10\rangle)/\sqrt{2}$, we see that cnot transformed the $|10\rangle$ part of our input state to $|11\rangle$ and did nothing to the first qubit in the $|1\rangle$ state, leaving our qubits in the state $(|00\rangle + |11\rangle)/\sqrt{2}$.
- 3 Looking at the matrix for cnot, we can see that it maps the |10\) computational basis state to |11\) and vice versa, just as we expected from the description.

IMPORTANT

The cnot instruction is not the same as if statements in classical programming languages, in that a cnot instruction preserves superposition. If we wanted to use an if statement, we'd have to measure the control qubit, such that it would collapse any superposition on the control qubit. We'll actually use both cnot instructions and if statements conditioned on measurement results when we write out our teleportation program at the end of this Chapter — both can be useful!

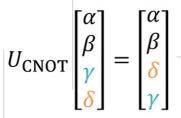
In Chapters 8 and 9, we'll see more about how controlled operations differ from if statements.

In , we can see how the cnot instruction acts on two-qubit states in general. For now, though, we recognize the output state that we got in by acting cnot on two qubits in

the $|+0\rangle$ state as the _entangled_ state we needed to play the CHSH game in Chapter 4, $(|00\rangle|11\rangle)/\sqrt{2}$.

Figure 5.4. The two qubit operation cnot exchanges the states of two qubits in a register.

The state of each register acted on by a cnot instruction is represented by a four-element vector. These elements tell us the amplitudes for each of the |00⟩, |01⟩, |10⟩, and |11⟩ computational basis states, respectively.



Just as we simulated swap by using a unitary matrix, we can also simulate the controlled-NOT (cnot) instruction by using a unitary matrix.

After running the cnot instruction, the amplitudes for $|10\rangle$ and $|11\rangle$ are swapped, while the amplitudes for $|00\rangle$ and $|01\rangle$ stay the same.

We can think of this instruction as doing nothing when the first qubit is in the $|0\rangle$ state, and as flipping the second qubit when the first one is in $|1\rangle$. Unlike a classical if statement, which requires measurement, the cnot instruction is *linear*, and thus preserves superposition.

This means we have everything we need to write a quantum program that entangles two qubits that start off in the $|00\rangle$ state. All we need to do is add the cnot instruction to our simulator, the same way as we added swap above.

```
def cnot(self, target: Qubit) -> None:
    self.parent._apply(
        qt.cnot(),
        [self.qubit_id, target.qubit_id]
)
```

Now we can write a program to prepare two qubits in an entangled pair.

```
>>> from simulator import Simulator
>>> sim = Simulator(capacity=2)
>>> with sim.using_register(2) as (you, eve):
...     eve.h()
...     eve.cnot(you)
...     sim.dump()
...
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
```

```
[[0.70710678]
[0. ]
[0. ]
[0.70710678]]
```

At this point, it's helpful to pause a moment (sorry, Eve!) and reflect on what you just did. In Chapter 4, when you simulated playing the CHSH game with Eve, you had to "cheat" by assuming that you and Eve could have access to two qubits that magically start off in the entangled state $(|00\rangle+|11\rangle)/\sqrt{2}$. Now, though, we see exactly how you and Eve can prepare that entanglement by running another quantum program before playing CHSH. The h instruction prepares the superposition that you need, while the new cnot instruction allows you to prepare entanglement with Eve. This entanglement "shares" that superposition across your two qubits. (Sharing is caring, after all.)

Just like preparing entanglement between you and Eve was how you needed to get ready for the CHSH game in Chapter 4, it will be the first step you need for Eve to teleport her quantum data to you. This makes cnot a very important instruction for us going forward.

Getting back to Eve, though, the next step that you'll need in order for her to teleport her data to you is that you'll need to do one of four different single-qubit operations to *decode* the quantum data that she sends you (recall), so let's look at those next.

5.4 ALL THE SINGLE (QUBIT) ROTATIONS

The last thing you will need to program quantum teleportation is to apply a correction based on some classical data that Eve sends you. To do so, you'll need couple of new single-qubit instructions. For that, it's helpful to revisit the pictures we've been using to depict quantum instructions as rotations, because we may have been cheating a little. We have depicted our qubits so far as any position on a circle, but in reality we are missing a dimension for our model of a qubit. The state of a *single* qubit is represented any point on the surface of a sphere.

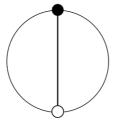
SINGLE QUBITS ONLY!

This (and the previous) way of visualizing the state of a qubit only works if that qubit is not entangled with any other qubits. Another way of saying this is we cannot easily visualize a multi-qubit state. Even trying to visualize the state of a two-qubit register with entanglement would take drawing pictures in 7 dimensions. While "7D" might be nice for advertising your ride at Niagara Falls, it is much harder draw useful pictures that way.

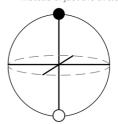
The circle you are familiar with was really just a slice through the sphere, and all the rotations we were doing resulted in states still on that circle. See to see how our previous model for the qubit compares to the Bloch sphere.

Figure 5.5. Comparison of our previous model for a qubit and the Bloch sphere, a better general model for a single qubit.

In Chapters 2, 3, and 4, you saw that the state of a single-qubit can be thought of as a point on a circle.



In this Chapter, we'll see that single-qubit states can also be rotated "out of the page," giving you an entire sphere instead of just the circle.



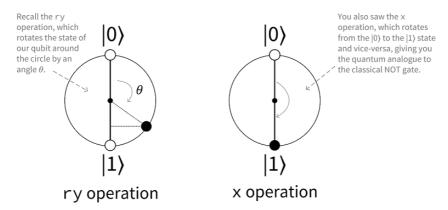
You might have inferred from the fact that we had shown the *Z*-axis and the *X*-axis that the *Y*-axis was probably hiding somewhere!

You may recall that when we first introduced the vector representation of a qubit state in Chapter 2, that the amplitudes in each vector were *complex numbers*. In the rest of this Chapter, we'll see that when you use rotation instructions to transform the state of a single qubit, we'll get complex numbers in general. Complex numbers are an incredibly useful tool for keeping track of rotations, and thus play a large role in quantum computing. Primarily, they help us understand the angles and phases between different quantum states. Don't worry if you're a little rusty with complex numbers, as you'll get plenty of chance practice with them throughout the rest of the book.

5.5 RELATING ROTATIONS TO COORDINATES: THE PAULI OPERATIONS

Let's do a quick review of a couple of the single qubit operations we have seen so far, x and ry.

Figure 5.6. Illustrations of what x and ry do to a qubit, both operations you have seen in previous chapters.



Now that we know the state of our qubit actually can be rotated on the surface of a sphere, what other rotations can help us rotate the state out of the plane? We can add a rotation about the line between the $|+\rangle$ and $|-\rangle$ states. This line is conventionally called the X-axis to distinguish it from the Z-axis that connects the $|0+\rangle$ and $|1\rangle$ states. The rx function in QuTiP gives us a Qobj encapsulating the rotation matrix for a X-axis rotation:

Listing 5.8. Using the QuTiP built-in function qt.sigmaz.

- 1 Due to how floating-point numbers work in classical computers, sometimes simulations of rotations can result in very small numbers like 10⁻¹⁷ where we would expect 0. QuTiP Qobj instances have a method tidyup to help make matrices more readable when this happens.
- 2 Up to a coefficient of -i (written in Python as -1j), rotating by 180° about the *X*-axis results in the x (NOT) instruction that we first saw in Chapter 2.

This snippet illustrates something very important: the x operation is precisely what we get by rotating around the X axis by an angle of 180° (\pi).

GLOBAL PHASE

As noted in the callouts for the above snippet, you can check that qt.rx(np.pi) is actually off by a factor of -i from qt.sigmax(). That factor is an example of something called a global phase. As we will see shortly, global phases cannot affect the results of measurement. Thus, qt.rx(np.pi) and qt.sigmax() are different unitary matrices that represent the same operation. We'll get more practice with global and local phases in Chapters 7 and 8.

By analogy, we call rotating by 180° about the Z-axis a z operation. In chapter 3,

QuTiP provided you with the qt.sigmax function which allow you to simulate the x instruction. Similarly, qt.sigmaz provides us with the unitary matrices we need to simulate the z instructions. See for an example of using qt.sigmaz.

Listing 5.9. Using the QuTiP built-in functions qt.rz and qt.sigmaz.

- ①Note that we've just included the coefficient (that is, the global phase) of -i by multiplying by i right away; this works since -i x i = -(-1) = 1. Cancelling out the global phase in this way will make it easier to read the output below.
- **2** As promised, up to the coefficient of -i, the z instruction applies a 180° rotation about the Z-axis.

In the same way that the x operation flips between $|0+\rangle$ and $|1\rangle$ while leaving $|1\rangle$ and $|-\rangle$ alone, the 'Z' operation flips between $|+\rangle$ and $|-\rangle$ while leaving $|0+\rangle$ and $|1\rangle$ alone.

shows a truth table like the ones we have made before for the Hadamard operation in Chapter 2. By looking at the truth table, you can confirm that, for any input state if you do a Z operation twice on it you will be back to where you started. Another way to say that is, Z squares to the identity operation $\mathbb{1}$, in the same way that $X^2 = \mathbb{1}$.

Table 5.1. Table Representing the z instruction as a table.

Input state	Output state
0)	0>
1>	- 1 >
+>	->
->	+>

NOTE

We listed four rows in , but we only need two rows to completely specify how Z acts for any input. The other two rows serve to emphasize that we can choose between defining Z by its action on $|0+\rangle$ and $|1\rangle$ or by its action on $|+\rangle$ and $|-\rangle$.

Exercise 5.2: Practice using rz and z

Suppose you prepare a qubit in the $|-\rangle$ state and apply a z rotation. If you measure along the X axis, what would you get? What would you measure if you apply two z rotations? If you had to implement those same two rotations with rz, what angles would you use?

We can define one more rotation in the same way, namely the rotation about an axis coming "out of the page." This axis connects the states $(|0\rangle+i|1\rangle)/\sqrt{2}=R_x(\pi)\,|0\rangle$ and $(|0\rangle-i|1\rangle)/\sqrt{2}=R_x(\pi)\,|1\rangle$, and is conventionally called the Y-axis. A 180° rotation about the Y-axis both flips bit labels $(|+\rangle\leftrightarrow|-1\rangle)$ and phases $(|+\rangle\leftrightarrow|-1\rangle)$, but leaves alone the two states along the Y-axis.

Exercise 5.3: Truth table for sigmay.

Use the qt.sigmay() function to make a table similar to, but for the y instruction.

PAULI MATRICES

Together, the three matrices X, Y, and Z representing the x, y, and z operations are called the Pauli matrices in honor of physicist Wolfgang Pauli. The identity matrix 1 is sometimes included as well, representing the "do nothing" or identity operation.

Playing rock-paper-scissors with the Pauli matrices

The Pauli matrices have a number of useful properties that we'll make use of throughout the rest of the book. Many of these properties make it easy to work out different equations involving the Pauli operators.

For example, if you multiply X and Y together, you get iZ, but if you multiply YX instead, you get iZ back:

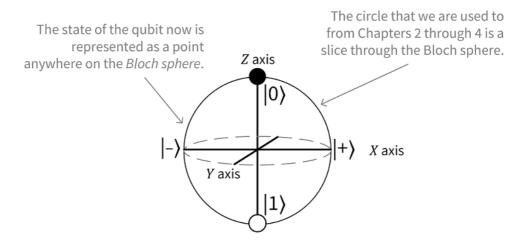
```
>>> qt.sigmax() * qt.sigmay()
Quantum object: dims = [[2], [2]], shape = (2, 2), type = oper, isherm = False
Qobj data =
[[0.+1.j 0.+0.j]
      [0.+0.j 0.-1.j]]
>>> qt.sigmay() * qt.sigmax()
Quantum object: dims = [[2], [2]], shape = (2, 2), type = oper, isherm = False
Qobj data =
[[0.-1.j 0.+0.j]
      [0.+0.j 0.+1.j]]
```

Similarly, YZ = iX and ZX = iY, but ZY = -iX and XZ = -iY. One way to remember this is as thinking of X, Y, and Z as playing a little game of rock-paper-scissors: X "beats" Y, Y "beats" Z, and Z "beats" X in turn.

We can think of these matrices as establishing a kind of coordinate system for qubit states, called the *Bloch sphere*. As shown in , the *X*- and *Z*-axes form the circle that you've seen in the book thus far, while the *Y*-axis comes out of the page.

DESCRIBING STATES WITH PAULI MEASUREMENTS Any single-qubit state can be entirely specified up to a global phase by the measurement probabilities for X, Y, and Z measurements. That is, if I tell you the probability of getting a "1" outcome for each of the three Pauli measurements that I could perform, then you can use that information to write down a state vector that is identical to mine, up to global phase. This makes the analogy to points in three dimensions a very good analogy for thinking about single-qubit states.

Figure 5.7. The Bloch sphere, in all its spherical glory.

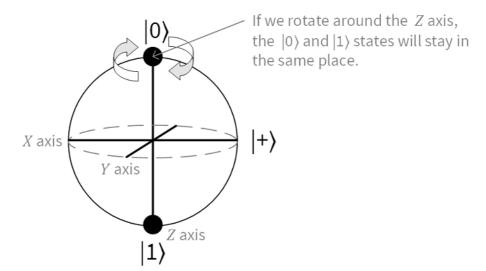


THE IS HAVE IT

The states at the ends of the Y-axis are usually labeled $|i\rangle$ and $|-i\rangle$, but are not often used on their own. We will just stick to the labeled states we were using before: $|0\rangle$, $|1\rangle$, $|+\rangle$, and $|-\rangle$.

With this picture in mind, it's easier to see why some rotations don't affect the results of measurements. For instance, as we illustrate in , the Bloch sphere picture helps us understand what happens if we rotate $|0\rangle$ about the Z axis.

Figure 5.8. The Bloch sphere illustrating how an rz rotation leaves the |0| state unchanged.



In the same way that the North Pole on a globe stays in the same spot no matter how much we spin the globe, if we rotate a state about an axis parallel to that state, there is no observable effect on our qubit. We can also see this effect come out of the math as well, see .

- 1 We start by defining a variable to represent the state |0).
- **2** Next, we introduce a new state $|\psi\rangle$ that's a 60° (π / 3 in radians) rotation of $|0\rangle$ about the Z axis.
- 3 The resulting state is $|\psi\rangle = [\cos(60^{\circ}/2) i \sin(60^{\circ}/2)] |0\rangle = [\sqrt{3}/2 i/2] |0\rangle$.
- **1** To check that this rotation doesn't do anything to our measurement results, we can check what we get from Born's rule. Let's start by writing down a measurement $\langle 0| = |0\rangle^+$. Recall that in QuTiP, we write out the "dagger" operator $^+$ by calling the .dag method of Qobj instances.

- Since |0⟩ is represented by a column vector, taking the conjugate transpose to get ⟨0| gives us a row vector.
- **6** Taking the inner product $(0 \mid \psi)$, we can compute Born's rule $Pr(0 \mid \psi) = |(0 \mid \psi)|^2$. Note that we need to index by [0, 0], since QuTiP represents the inner product of $|0\rangle$ with $|\psi\rangle$ as a 1 × 1 matrix.
- As before, we note that the probability of observing a "0" when measuring along the Z axis hasn't changed.

TIP

When writing down states, $|\psi\rangle$ is often used as just an arbitrary name, similar to how x is often used to represent an arbitrary variable in algebra.

Exercise 5.4: verify that applying rz doesn't change | 0).

We've only checked that one measurement probability is still the same, but maybe the probabilities have changed for X or Y measurements. To fully check that the global phase doesn't change anything, prepare the same state and rotation as in and check that the probabilities of measuring the state along the X or Y axis aren't changed by applying an xz instruction.

In general, you can always multiply a state by a complex number whose absolute value is 1 without changing the probabilities of any measurement. Any complex number z = a + bi can be written as $z = re^{i\theta}$ for real numbers r and θ , where r is the absolute value of z and where θ is an angle. When r = 1, we have a number of the form $e^{i\theta}$, which we call a **phase**. We then say that multiplying the state by a phase applies a *global phase* to that state.

IMPORTANT

No global phase can ever be detected by any measurement.

The states $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ are in every conceivable way two different ways of describing the exact same state. There is no measurement that one can do *even in principle* to learn about global phases. On the other hand, we've seen that we can tell apart states like $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|-\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$ that differ only in the *local* phase of the $|1\rangle$ computational basis state.

Taking a step back, let's summarize what you've learned about the x, y, and z instructions so far, and about the Pauli matrices we use to simulate those instructions. We've seen that the x instruction flips us between $|0+\rangle$ and $|1\rangle$, while the z instruction flips us between the $|+\rangle$ and $|-\rangle$ states. Put differently, the x instruction flips bits, while the z instruction flips phases.

Looking at the Bloch sphere, you can see that rotating about the Y-axis should do both of these. You can also see this by using that Y = -iXZ, as is straightforward to check with QuTiP. We summarize what each Pauli instruction does in .

Table 5.2. Table Pauli matrices as bit and phase flips.

Instruction	Pauli matrix	Flips bits ($ 0\rangle \leftrightarrow 1\rangle$)?	Flips phases ($ +\rangle \leftrightarrow -\rangle$)?
(no instruction)	1	No	No
x	X	Yes	No
У	Y	Yes	Yes
Z	Z	No	Yes

Exercise 5.5: Hmm, that rings a Bell...

The $(|00\rangle+|11\rangle)/\sqrt{2}$ state that we've seen a few times now isn't the only example of an entangled state. In fact, if you pick a two-qubit state at random, it is almost certainly going to be entangled. Just as the computational basis $\{|00\rangle,|01\rangle,|10\rangle,|11\rangle\}$ is a particularly useful set of unentangled states, there's a set of four particular entangled states known as the *Bell basis* after physicist John Stewart Bell.

Table 5.3. The four Bell states

Name	Expansion in computational basis
$ \beta_{00}\rangle$	$(00\rangle + 11\rangle) / \sqrt{2}$
$ \beta_{01}\rangle$	$(00\rangle - 11\rangle) / \sqrt{2}$
β ₁₀ ⟩	$(01\rangle + 10\rangle) / \sqrt{2}$
$ \beta_{11}\rangle$	$(01\rangle - 10\rangle) / \sqrt{2}$

Using what you've learned about the cnot instruction and the Pauli instructions (x, y, and z), write programs to prepare each of the four Bell states in the table above.

HINT: Table 5.2 should be very helpful in this exercise.

We finish our discussion of single-qubit operations by adding instructions to our Qubit interface and simulator for the X, Y and Z operations.

```
Listing 5.11. simulator.py

def rx(self, theta: float) -> None:
    self.parent._apply(qt.rx(theta), [self.qubit_id])

def ry(self, theta: float) -> None:
    self.parent._apply(qt.ry(theta), [self.qubit_id])

def rz(self, theta: float) -> None:
    self.parent._apply(qt.rz(theta), [self.qubit_id])

def x(self) -> None:
    self.parent._apply(qt.sigmax(), [self.qubit_id])

def y(self) -> None:
    self.parent._apply(qt.sigmay(), [self.qubit_id])
```

```
def z(self) -> None:
    self.parent._apply(qt.sigmaz(), [self.qubit_id])
```

- We can implement the rotation instructions rx, ry, and rz using the corresponding QuTiP functions qt.rx, qt.ry, and qt.rz to get copies of the unitary matrices we need to simulate each instruction.
- **2** QuTiP uses the notation \sigma_x instead of *X* for the Pauli matrices. Using this notation, the function sigmax() returns a new Qobj representing the Pauli matrix *X*. This way, we can implement the x, y, and z instructions corresponding to each Pauli matrix.

No one can be told what the matrix is. You have to see it for yourself.

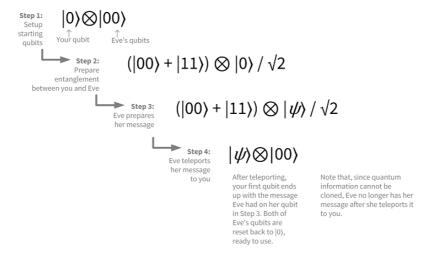
We've talked a lot about matrices thus far in Part I. A **lot**. It's tempting to say that quantum programming is all about matrices, and that qubits are really just vectors. In reality, though, matrices are how we **simulate** what a quantum device does. We'll see more in Part II, but quantum programs don't manipulate matrices and vectors at all — they manipulate classical data such as what instructions to send to a quantum device, and what to do with the data we get back from devices. For instance, if we have an instruction that we run on a device, there's no simple way to see what matrix we should use to simulate that instruction — rather, you have to reconstruct that matrix from many repeated measurements using a technique called *process tomography*.

When you write down a matrix, whether in code or on a piece of paper, you're implicitly simulating a quantum system. If that really bakes your noodle, don't worry, this will make a lot more sense as you go through the rest of the book.

5.6 TELEPORTATION

OK, now we have everything we need to write out what teleportation looks like as a quantum program. As a quick review, is what we want this program to do.

Figure 5.9. Recall the steps to the teleportation program.



We will assume you can prepare some entangled qubits while they are in the same

device, and that you and Eve have a means of classical communication that you can use to signal the correct correction to use.

We can now use the features we added to our simulator in this chapter to implement the teleportation program, see .

Listing 5.12. Quantum teleportation, all in just a few lines of Python. from interface import QuantumDevice, Qubit from simulator import Simulator def teleport(msg: Qubit, here: Qubit, there: Qubit) -> None: here.h() here.cnot(there) # ... msg.cnot(here) msg.h() if msg.measure(): there.z() if here.measure(): there.x() msg.reset() here.reset()

- ①Your teleport function will take two qubits as input: the qubit that you want to move (msg), and where you want it to be moved to (there). You'll also need one temporary qubit, which we call here. We'll presume by convention that both here and there start in the $|0+\rangle$ state.
- 2 We need to start off with some entanglement between here and there. We can use our old friend, the h instruction, together with our new friend, the cnot instruction.
- 3 This is the only instruction that we'll use in this program that needs to act on both here and there; after running this, you can send Eve your qubit, and both of you can run the rest of the program with only classical communication.
- **4** At this point in the program, here and there are in the $(|00\rangle + |11\rangle)/\sqrt{2}$ state that we first saw in Chapter 4.
- **5**The next step is to run the program we used to prepare the $(|00\rangle + |11\rangle)/\sqrt{2}$ state backwards, but on the msg and here qubits that live entirely on your device instead. We can think of running a preparation backwards as a kind of measurement, such that these steps set us up to measure the quantum message you're trying to send Eve in an entangled basis.
- **6** When we actually do that measurement, we get classical data back out that we can use to send Eve. Once she has that data, she can use the x and z instructions to decode the quantum message.
- Now that you're done with your qubits, it's good to put each of them back into $|0+\rangle$ so that they're ready to be used again. This doesn't affect the state of there, though, as you've only reset your qubits, not the one you gave to Eve!

C WHAT WE DID THERE?

If we didn't need to send Eve our classical measurement results as a part of teleportation, then we could use teleportation to send both classical and quantum data faster than the speed of light. Just like you couldn't communicate with Eve when you played the CHSH game in the previous chapter, the speed of light means that you need to communicate with Eve classically in order to use entanglement to send quantum data. In both cases, entanglement can help us communicate, but it doesn't let you communicate all on its own: you always need some other kind of communication as well.

To see that this actually works, You can prepare something on your qubit, send it to Eve, and then she can undo your preparation on her qubit. Thus far, the messages you and Eve have been sending have been classical, but here the message is *quantum*. You can and will measure the quantum message to get a classical bit out, but you can also use the quantum message you get from Eve like any other quantum data, for example, you can apply whatever rotations and other instructions you like.

What is this good for?

Sending quantum data may not seem that much more useful than sending classical data; after all, sending classical data has gotten us a lot of neat things thus far. By contrast, applications for sending quantum data tend to be a bit more niche at the moment.

That said, that we can move quantum data around is a really useful example to help us understand how quantum computers work. The ideas that you developed in this Chapter aren't often that *directly* useful, but will help you build up really great stuff going forward.

Let's say you prepare a *quantum* message to you by using the operation msg.ry(0.123), see for how you can teleport this message to Eve.

Listing 5.13. A sample program that uses teleportation to move quantum data around,

sending the message 0.123 to Eve.

```
if __name__ == "__main__":
    sim = Simulator(capacity=3)
    with sim.using_register(3) as (msg, here, there):
        msg.ry(0.123)
        teleport(msg, here, there)

        there.ry(-0.123)
        sim.dump()

4
```

- 1 As before, you'll allocate a register of qubits, and give each qubit a name.
- 2 Next, you prepare a message to send to Even. Here, we've shown using a particular angle as the message, but it really could be anything.
- 3 You can then call the teleportation program that you wrote earlier to move the message you prepared onto Eve's qubit.
- If Eve then undoes your rotation by rotating by the opposite angle, you can check that the output of the dump instruction that we added above tells us that the register we allocated is back in the |000⟩ state. This shows that your teleportation worked!

When you run this program, you'll get output similar to that in.

To verify that the teleportation worked, if Eve undoes the instruction that you did on her qubit (there.ry(0.123)), she should get back the $|0+\rangle$ state that you started with. With teleportation, you and Eve were able to send the quantum information, using entanglement and classical communication.

```
Quantum object: dims = [[2, 2, 2], [1, 1, 1]], shape = (8, 1), type = ket
Qobj data =
[[1.]
   [0.]
   [0.]
   [0.]
   [0.]
   [0.]
   [0.]
   [0.]
   [0.]
```

Exercise 5.6: What if it just didn't do anything?

Try changing your operation or Eve's operation to convince yourself that you only get a |000) state at the end if you undo the same operation that Eve applied to her qubit.

Now you can brag to all your friends (and make whatever sci-fi references you want) that you can do teleportation. Hopefully you see why this is *not* the same as getting beamed down to a planet from orbit, and that when you teleport a message it does not communicate faster than the speed of light.

In this chapter you learned:

- Implement a quantum program to move data around a quantum computer using classical and quantum control
- Recognize a new way of visualizing single qubit operations called the Bloch sphere
- Predict the output of two-qubit operations, and Pauli operations

5.7 PART I: CONCLUSION

You have made it to the end of Part I, but sadly our qubits are in another castle (a). Getting through this Part is no mean feat, as we were refreshing a number of topics such as linear algebra and complex numbers, all while introducing a whole host of new quantum concepts. Certainly there will probably still be some questions or parts you are a bit shakey on, and that's ok. We will be using and practicing these skills to start developing more complicated quantum programs for cool applications like chemistry and cryptography.

Before we move onto that, though, give yourself a pat on the back, you've done a lot so far! Let's summarize some of what you've already accomplished:

- Refreshed your linear algebra and complex number skills
- Learned about what a qubit is as well as what you can do with one
- Built up a multiple-qubit simulator in Python

- Wrote a number of quantum programs for tasks like quantum key distribution (QKD), playing nonlocal games, and even quantum teleportation
- Learned braket notation for the states of quantum systems

Your Python simulator will continue to be a useful tool when trying to understand what is happening when we get to some of the larger applications. For Part II, we will switch to primarily using Q# as our tool of choice to write our quantum programs. There are a number of reasons we will talk about in the next chapter for why we will choose to write these more advanced quantum programs in Q# over Python, but the main reasons are speed and extensibility. Plus, you can even use Q# from Python or with a Q# kernel for Jupyter, so you can use whatever development environment you like best!

See you in Part II!

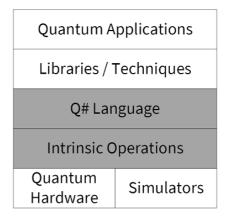
Changing the odds: An introduction to O#

This chapter covers:

- Using the Quantum Development Kit to write quantum programs in Q#,
- · How to use Jupyter Notebook to work with Q#,
- How to run Q# programs using a classical simulator.

Up to this point, we've used Python to implement our own software stack to simulate quantum programs. Moving forward, though, we'll be writing more intricate quantum programs that will benefit from specialized language features that are hard to implement by embedding our software stack inside Python. Especially as we explore quantum algorithms, it's helpful to have a language tailor-made for quantum programming at our disposal. In this chapter, we'll get started with Q#, Microsoft's domain-specific language for quantum programming, included with the Quantum Development Kit.

Figure 6.1. In this Chapter, we will be covering topics in both the Q# language and intrinsic operations for quantum devices.



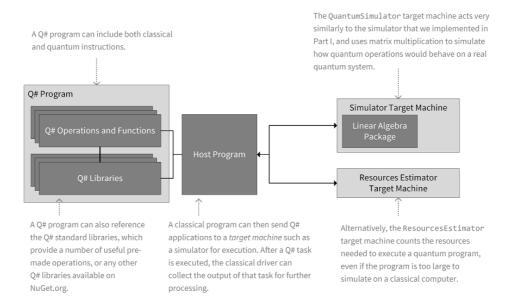
6.1 INTRODUCING THE QUANTUM DEVELOPMENT KIT

The Quantum Development Kit provides a new language, Q#, for writing quantum programs and simulating them using classical resources. Quantum programs written in Q# are run by thinking of quantum devices as a kind of accelerator, similar to how you might run code on a graphics card.

TIP

If you've ever used a graphics card programming framework like CUDA or OpenCL, this is a very similar model.

Figure 6.2. Q# software stack on a classical computer.



Let's take a look at this software stack for Q#.

Our Q# program itself consists of operations and functions that instruct quantum and classical hardware to do certain things. There are also a number of libraries that are provided with Q# that have helpful, pre-made operations and functions to use in our programs.

Once the Q# program is written, we need a way for it to pass instructions to the hardware. A classical program, sometimes called a "driver" or a "host program," is responsible for allocating a target machine and running a Q# operation on that machine.

The Quantum Development Kit provides a plugin for Jupyter Notebook called IQ# that makes it easy to get started with Q# by providing host programs automatically for us. In Chapter 8, we'll see how to write host programs using Python and C#, but for now we'll focus on Q# itself. See Appendix A for instructions on setting up your Q# environment to work with Jupyter Notebook.

Using the IQ# plugin for Jupyter Notebook, we can use one of two different target machines to run Q# code. The first is the QuantumSimulator target machine, which is very similar to the Python simulator that we have been developing. It will be a lot faster than our Python code at simulating our qubits.

The second is the ResourcesEstimator target machine which will allow us to estimate how many qubits and quantum instructions we would need to run it, without having to fully simulate it. This is especially useful for getting an idea of the resources you would need to run a Q# program for your application, as we'll see when we look at larger Q# programs later on in the book.

Figure 6.3. Getting started with IQ# and Jupyter Notebook



To get a sense for how everything works, let's start by writing out a purely classical Q# "hello, world" application. First, start Jupyter Notebook by running the following in a

terminal:

jupyter notebook

This will automatically open a new tab in your browser with the home page for your Jupyter Notebook session. From the **New** \downarrow menu, select "Q#" to make a new Q# notebook. Type the following into the first empty cell in the notebook and press Control + Enter or \Re + Enter to run it.

```
function HelloWorld() : Unit {
    Message("Hello, classical world!");
}
```

- 1 This line defines a new function which takes no arguments, and returns the empty tuple, whose type is written as Unit.
- 2 The Message function tells the target machine to collect a diagnostic message. The QuantumSimulator target machine prints all diagnostics to the screen, so we can use Message in the same way as print in Python.

WATCH OUT FOR SEMICOLONS!

Unlike Python, Q# uses semicolons rather than newlines to end statements. If you get a lot of compiler errors, make sure you remembered your semicolons.

You should get a response back listing that the Helloworld function was successfully compiled. To run our new function, we can use the %simulate command in a new cell.

%simulate HelloWorld

A BIT OF CLASSICAL MAGIC

The %simulate command we used above is an example of a magic command, in that it's not actually a part of Q# itself, but is an instruction to the Jupyter Notebook environment. If you're familiar with the IPython plugin for Jupyter, you may have used similar magic commands to tell Jupyter how to handle Python plotting functionality. The magic commands we use in this book all start with % to make them easy to tell apart from Q# code.

In this example, %simulate allocates a target machine for us and sends a Q# function or operation to that new target machine. In Chapter 8, we'll see how to accomplish something similar using Python and C# host programs, instead of using Jupyter Notebook.

The Q# program is sent to the simulator, but in this case, the simulator just runs the classical logic, since there's no quantum instructions to worry about yet.

Exercise 6.1

Change the definition of HelloWorld to say your name instead of "classical world," and then run %simulate again using your new definition.

6.2 FUNCTIONS AND OPERATIONS IN Q#

Now that we have the Quantum Development Kit up and running with Jupyter Notebook, let's use Q# to write some quantum programs. Back in Chapter 2, we saw

that one useful thing to do with a qubit is to generate random numbers one classical bit at a time. Revisiting that application makes a great place to start with Q#, especially since random numbers are useful if you want to play games.

138

Long ago in Camelot, Morgana le Fay shared our love for playing games. Being a clever mathematician with skills well beyond her own day, Morgana was even known to use qubits from time to time as a part of her games. One day, as Sir Lancelot lay sleeping under a tree, Morgana trapped him and challenged him to a little game: each of them must try to guess the outcome of measuring one of Morgana's qubits.

Two sides of the same... qubit?

In Chapter 2, we saw how we can generate random numbers one bit at a time by preparing and measuring qubits. That is, qubits can be used to implement *coins*. We'll use the same kind of idea in this Chapter as well, thinking of a coin as a kind of interface that allows its user to "flip" it and get out a random bit. That is, we can implement the coin interface by preparing and measuring qubits.

If the result of measuring along the Z axis is a 0, then Lancelot wins their game and gets to return to Genevieve. If the result is a 1, though, Morgana wins and Lancelot has to stay and play again. Notice the similarity to our QRNG program from before. Just as in chapter 2, we'll measure a qubit to generate random numbers, this time for the purpose of playing a game. Of course, Morgana and Lancelot could have also flipped a more traditional coin, but where is the fun in that?

6.2.1 MORGANA'S SIDE GAME

Prepare a qubit in the |0) state

Apply the Hadamard operation (recall that the unitary operator H takes $|0\rangle$ to $|+\rangle$ and vice versa)

Measure the qubit in the Z axis. If the measurement result is a 0, then Lancelot can go home. Otherwise, he has to stay and play again!

Sitting at a coffee shop watching the world go by, we can use our laptops to predict will happen in Morgana's game with Lancelot by writing a quantum program in Q#. Unlike the ClassicalHello function that we wrote above, our new program will need to work with qubits, so let's take a moment to see how to do so with the Quantum Development Kit.

The primary way that we interact with qubits in Q# is by calling *operations* that represent quantum instructions. For instance, the H operation in Q# represents the Hadamard instruction we saw in Chapter 2. To understand how these operations work, it's helpful to understand the difference between Q# operations and the functions that we saw in the ClassicalHello example above.

Functions in Q# represent *predictable* classical logic, things like mathematical functions (Sin, Log). Functions always return the same output when given the same input.

Operations in Q# represent code that can have side effects, such as sampling random

numbers, or issuing quantum instructions which modify the state of one or more qubits.

This separation helps the compiler figure out how to automatically transform your code as a part of larger quantum programs; we'll see more about this later.

Another perspective on functions versus operations

Another way of thinking of the difference between functions and operations is that functions compute things, but cannot cause anything to happen. No matter how many times we call the square root function Sqrt, nothing about our Q# program has changed. By contrast, if we run the X operation, then an X instruction is sent to our quantum device, which causes a change in the state of the device. Depending on the initial state of the qubit that the X instruction was applied to, we can then tell that the X instruction has been applied by measuring the qubit. Because functions don't do anything in this sense, we can always predict their output exactly given the same input.

One important consequence is that functions cannot call operations, but operations can call functions. This is because you can have an operation which is not necessarily predictable call a predictable function and you still have something that may or may not be predictable. However, a predictable function cannot call a potentially unpredictable operation and still be predictable.

We'll see more about the difference between Q# functions and operations as we use them throughout the rest of the book.

Since we want quantum instructions to have an effect on our quantum devices (and on Lancelot's fate), all quantum operations in Q# are defined as operations (hence the name). For instance, suppose that Morgana and Lancelot prepare their qubit in the |+> state using the Hadamard instruction. Then we can predict the outcome of their game by writing out the quantum random number generator (QRNG) example from Chapter 2 as a Q# operation.

THERE MAY
BE SIDE
EFFECTS TO
THIS
OPERATION...

When we want to send instructions to our target machine to do something with our qubits, we need to do so from an operation, since sending an instruction is a kind of side effect. That is, when we run an operation, we aren't just computing something, we're doing something. Running an operation twice isn't the same as running it once, even if we get the same output both times. Side effects aren't deterministic or predictable, and so we can't use functions to send instructions on how to manipulate our qubits.

In , we'll do just that, starting by writing an operation called NextRandomBit to simulate each round of Morgana's game. Note that since NextRandomBit needs to work with qubits, it has to be an operation and not a function. We can ask the target machine for one or more fresh qubits with the using block.

ALLOCATING QUBITS IN O#

The using statement is one of the only two ways we can ask the target machine for qubits. There's no limit to the number of using statements that we can have in our Q# programs, other than the number of qubits that each target machine can allocate. At the end of each using block, the qubits then go back to the target machine, so that one way to think of using blocks is to make sure that each qubit that is allocated is "owned" by a particular operation. This makes it impossible to "leak" qubits within a Q# program, which is very helpful given that qubits are likely to be very expensive resources on actual quantum hardware.

Q# offers one other way to allocate qubits, known as borrowing. Unlike when we allocate qubits with using statements, the borrowing statement lets us borrow qubits that are owned by different operations without knowing what state they start in. We won't see much of borrowing in this book, but the borrowing statement works very similarly to the using statement in that it makes it impossible for us to forget that we've borrowed a qubit.

By convention, all qubits start off in the $|0\rangle$ state right after we get them, and we promise the target machine that we'll put them back into the $|0\rangle$ state at the end of the block so that they're ready for the target machine to give to the next operation that needs them.

Listing 6.1. Simulating one round of Morgana's game using Q#

```
operation NextRandomBit() : Result {
    using (qubit = Qubit()) {
        H(qubit);
        let result = M(qubit);
        Reset(qubit);
        return result;
    }
}
```

- 1 This time, because we want to use a qubit, we declare an operation instead of a function. Since our operation needs to return a result to its caller, we denote by changing the return type to the Q# type Result.
- 2 The using keyword in Q# asks the target machine for one or more qubits. Here, we ask for a single value of type Qubit, which we store in the new variable qubit.
- 3 Quantum operations such as the Hadamard operation can be found in the Microsoft.Quantum.Intrinsic namespace. For instance, we can call Hadamard using the Microsoft.Quantum.Intrinsic.H operation. After calling H, qubit is in the $H|0\rangle=|+\rangle$ state.
- **4** Next, we use the M operation to measure our qubit in the Z basis, saving the result to the result variable we declared earlier. Since we are in an equal superposition of $|0\rangle$ and $|1\rangle$, result will be either Zero or One with equal probability.
- § Before returning our qubit to the target machine, we use the Microsoft.Quantum.Intrinsic.Reset operation to return it to the $|0\rangle$ state. Since we've already stored the classical data we got from our measurement into the result variable, we can safely reset the qubit without losing any information that we care about.
- **6** We finish our operation by returning the measurement result back to to the caller.

Exercise 6.2

Use the <code>%simulate</code> magic command to run the <code>NextRandomBit</code> operation a few times; do you get the results you'd expect?

Next, we need to see how many rounds it takes for Lancelot to get the zero he needs to go home. Let's write an operation to play rounds until we get a zero. Since this operation simulates playing Morgana's game, we'll call it PlayMorganasGame.

Listing 6.2. Simulating many rounds of Morgana's game using Q#

```
operation PlayMorganasGame() : Unit {
   mutable nRounds = 0;
   mutable done = false;
   repeat {
      set nRounds = nRounds + 1;
      set done = (NextRandomBit() == Zero);
   }
   until (done)
   fixup {}

Message($"It took Lancelot {nRounds} turns to get home."); 5
}
```

- 1 All Q# variables are immutable by default we can use the mutable keyword to declare a variable that we can change later with the set keyword. Here, we start by initializing a mutable variable indicating how many rounds have already passed, and a mutable variable we'll use to exit the loop.
- Q# allows operations to use a kind of loop called a "repeat-until-success" (RUS) loop. Unlike a while-loop, RUS loops also allow us to specify a "fixup" that runs if the condition to exit the loop isn't met. Though we've included the fixup {} block here to illustrate the concept, you can also omit fixup blocks when they're empty.
- 3 Inside our loop, we call the QRNG that we wrote above as the NextRandomBit operation. We check to see if the result is a Zero (that is, if Lancelot wins and can leave), and if so, set done to be true.
- 4 If we got a Zero, then we can stop the loop.
- **⑤** Finally, we use Message again to print the number of rounds to the screen. To do so, we use \$"" strings which, similar to \$"" strings in C# and f"" strings in Python, let us include variables in the diagnostic message by using {} placeholders inside the string.

Why Do We Need to Reset Qubits?

In Q#, when we allocate a new qubit with using, we promise the target machine that we will put it back in the |0) state before we deallocate it. At first glance, this seems rather unnecessary, as the target machine could just reset the state of qubits when they are deallocated — after all, we will often simply call the Reset operation at the end of a using block.

It is important to note, though, that the Reset operation works by making a measurement in the Z basis and flipping the qubit with an X operation if the measurement returns One. In many quantum devices, measurement is much more expensive than other operations, such that if we can avoid calling Reset we can reduce the cost of our quantum programs. Especially given the limitations of medium-term devices, this kind of optimization can be critical in making a quantum program practically useful.

Later in the chapter, we will see examples of where we know the state of a qubit when it needs to be deallocated, such that we can "unprepare" the qubit instead of measuring it.

We can run this new operation with the <code>%simulate</code> command in a very similar fashion as the <code>ClassicalHello</code> example. When we do so, we can see how long Lancelot has to stay:

Listing 6.3. Output from running the Qrng application

```
In []: %simulate PlayMorganasGame
It took Lancelot 1 turns to get home.
Out[]: ()
```

Looks like Lancelot got lucky that time! Or perhaps unlucky, if he was bored of hanging round the table in Camelot.

6.3 PASSING OPERATIONS AS ARGUMENTS

Suppose in Morgana's game above, we were interested in sampling random bits with non-uniform probability. After all, Morgana didn't promise Lancelot *how* she prepared the qubit that they measure; she can keep him playing longer if she makes a biased coin with their qubit instead of a fair coin.

The easiest way to modify Morgana's game is to, instead of calling H directly, take as an input an operation representing what Morgana does to prepare for their game. To take an operation as input, you need to write down the *type* of the input, just as you can write down qubit: Qubit to declare an input qubit of type Qubit. Operation types are indicated by thick arrows (=>) from their input type to their output type. For instance, H has type Qubit => Unit since H takes a single qubit as input and returns an empty tuple as its output.

TIP

In Q#, functions are denoted by thin arrows (->), while operations are denoted by thick arrows (=>).

Listing 6.4. Using operations as inputs in order to predict Morgana's game.

- 1 This time, we've added a new input called statePreparation to NextRandomBit that represents the operation we want to use to prepare the state we use as a coin. In this case, Qubit => Unit is the type of any operation which takes a single qubit and returns the empty tuple type Unit.
- 2 Within NextRandomBit, the operation passed as statePreparation can be called in the same way as any other operation.
- 3 The Q# standard libraries provide the Microsoft.Quantum.Measurement.MResetZ as a convenience for measuring and resetting a qubit in one step. This is equivalent to the set result = M(qubit); Reset(qubit); return result; statements we saw in the previous example, but requires one less measurement to perform. You'll see more about this operation later on the chapter, as well as how to use a shorter name when calling this operation.

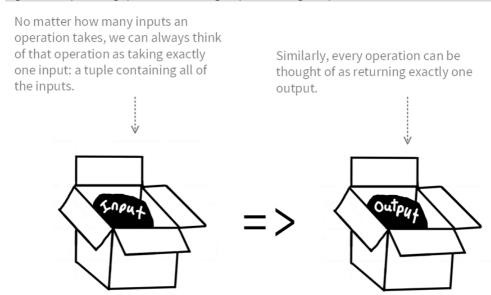
Exercise 6.3

What's the type of your new definition of NextRandomBit?

Tuple-In Tuple-Out

All functions and operations in Q# take a single tuple as an input and return a single tuple as an output. For instance, a function declared as function Pow(x:Double, y:Double): Double {...} takes as input a tuple (Double, Double), and returns a tuple (Double) as its output. This works because of a property known as singleton-tuple equivalence. For any type 'T, the tuple ('T) containing a single 'T is equivalent to 'T itself. In the example of Pow, this means that we can think of the output as a tuple (Double) that is equivalent to Double.

Figure 6.4. Representing operations with a single input and a single output



With this in mind, a function or operation that returns no outputs can be thought of as returning a tuple with no elements, (). The type of such tuples is called Unit, similar to other tuple-based languages such as F#. If we think of a tuple as a kind of box, then this is distinct from void as used in C, C++, or C# because there still is something there, namely a box with nothing in it.

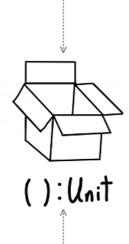
NOTE

In Q#, we always return a box, even if that box is empty.

There's no meaning in Q# to a function or operation that returns "nothing." For more details, see Section 7.2 of Get Programming with F#.

Figure 6.5. Unit versus void

If an operation that doesn't take any inputs or that doesn't return any outputs, then we represent that with an empty tuple, written in Q# as ().



Just like any other value in Q#, the empty tuple has a type. The type of () is called Unit.

Most quantum operations that don't involve a measurement will return Unit.



In other languages, like C, C++, C#, or Java, you'll sometimes see the keyword void used to indicate that a function or method doesn't return anything. Unlike Unit, void represents that there isn't a value at all, not even an empty box.

In this example, we see that NextRandomBit treats its input statePreparation as a "black box." The only way to learn anything about Morgana's preparation strategy is to *run* it.

Put differently, we don't want to do anything with with statePreparation that implies we know what it does or what it is. The only way that NextRandomBit can interact with statePreparation is by calling it, passing it a Qubit to act on.

This allows us to reuse the logic in NextRandomBit for many different kinds of state preparation procedures that Morgana might use to cause Lancelot a bit of trouble. For example, suppose she wants a biased coin that returns a One ¾ of the time and a Zero ¼ of the time. Then, we might run something like the following to predict this new strategy:

Listing 6.5. Passing different state preparation strategies to

©Manning Publications Co. To comment go to liveBook

the PlayMorganasGame example.

```
open Microsoft.Quantum.Math;

operation PrepareQuarterCoin(qubit : Qubit) : Unit {
   Ry(2.0 * PI() / 3.0, qubit);
}
```

- Oclassical math functions such as Sin, Cos, Sqrt, and ArcCos, as well as constants like PI() are provided by the Microsoft.Quantum.Math namespace, so we open it as well as the intrinsics.
- 2 The Ry operation implements the Y-axis rotation that we saw in Chapter 2. Q# uses radians rather than degrees to express rotations, so this is a rotation of 120° about the Y-axis. Thus, if qubit starts in |0⟩, this prepares qubit in the state $R_y(-120^\circ) |0\rangle = \sqrt{3/4} |0\rangle + \sqrt{1/4} |1\rangle$, such that the probability of observing 1 when we measure is $\sqrt{3/4}^2 = 3/4$.

We can make this example even more general, allowing Morgana to specify an arbitrary bias for her coin (which is implemented by their shared qubit):

Listing 6.6. Passing operations to implement PlayMorganasGame with arbitrary coin

biases.

```
operation PrepareBiasedCoin(morganaWinProbability : Double, qubit : Qubit) : Unit {
   let rotationAngle = -2.0 * ArcCos(Sqrt(morganaWinProbability));
   Ry(rotationAngle, qubit);
}

operation PrepareMorganasCoin(qubit : Qubit) : Unit {
    PrepareBiasedCoin(0.62, qubit);
}
```

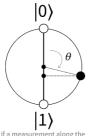
- We need to find out what angle we rotate the input qubit by in order to get the right probability of seeing a Zero as our result. This takes a little bit of trigonometry, see the sidebar below for the details.
- 2 This operation has the right type signature (Qubit => Unit) and we can see that the probability Morgana will win each round is 62%.

Working out the trigonometry

As we've seen a number of times, quantum computing deals extensively with *rotations*. To figure out what angles we need for our rotations, we need to rely on a little bit on a branch of mathematics for describing rotation angles, known as trigonometry (literally, the study of triangles). For instance, as we saw in Chapter 2, rotating $|0\rangle$ by an angle θ about the Y axis results in a state $\cos(-\theta/2)|0\rangle + \sin(-\theta/2)|1\rangle$. We know we want to choose θ such that $\cos(-\theta/2) = \sqrt{62\%}$, so that we get a 62% probability of getting a Zero result. That means we need to "undo" the cosine function to figure out what θ needs to be. In trigonometry, the inverse of the cosine function is called the arccosine function, and is written \arccos. Taking the arccosine of both sides of $\cos(-\theta/2) = \sqrt{62\%}$ gives us. We can cancel out the arccos and cos to find a rotation angle that gives us what we need, $\frac{-\theta}{2} = \arccos(\sqrt{62\%})$. Finally, we multiply both sides by -2 to get the equation we used in line 1 of .

Figure 6.6. How Morgana can choose θ to control how her game plays out

If a measurement along the Z axis results in a Zero, then Lancelot wins and can go home.

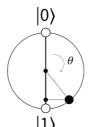


Z axis results in a One, then Morgana wins, and Lancelot has to play another round.

TIP

Morgana can control how probable a One outcome is by choosing θ , the angle that she rotates about the Y axis.





On the other hand, the closer Morgana picks to π (180°), the closer the projection of the state onto the Z axis is to $|1\rangle$.

If she picks θ to be exactly π , then the result will always be a One: she can keep Lancelot playing indefinitely.

This is somewhat unsatisfying, though, in that the operation PrepareMorganasCoin introduces a lot of boilerplate just to lock down the value of 0.62 for the input argument headsProbability to PrepareBiasedCoin. If Morgana changes her strategy to have a different bias, then using this approach, we'll need another new boilerplate operation to represent it. Taking a step back, let's look at what PrepareMorganasCoin actually does. It starts with an operation PrepareBiasedCoin: (Double, Qubit) => Unit, and wraps it into an operation of type Qubit => Unit by locking down the Double argument to 0.62. That is, it removes one of the arguments to PrepareBiasedCoin by fixing the value of that input to 0.62.

Thankfully, Q# provides a convenient shorthand for making new functions and operations by locking down some (but not all!) of the inputs. Using this shorthand, known as *partial application*, we can rewrite the above in a more readable form:

Listing 6.7. Using partial application to make it easier to vary Morgana's strategy.

let flip = NextRandomBit(PrepareBiasedCoin(0.62, _));

The _here indicates that a part of the input to PrepareBiasedCoin is missing. We say that PrepareBiasedCoin has been partially applied. Whereas PrepareBiasedCoin had type (Double, Qubit) => Unit, because we filled in the Double part of the input, PrepareBiasedCoin(0.62, _) has type Qubit => Unit, making it compatible with our modifications to NextRandomBit.

Partial application in Q# is similar to functools.partial in Python and the _ keyword in Scala.

Another way to think of partial application is as a way to make new functions and operations by specializing existing functions and operations:

```
function BiasedPreparation(headsProbability : Double) : (Qubit => Unit) {
    return PrepareBiasedCoin(headsProbability, _);
}
```

- Here, the output type of BiasedPreparation is an operation that takes a Qubit and returns the empty tuple. That is, BiasedPreparation is a function that makes new operations!
- 2We make the new operation by passing along headsProbability, but leaving a blank (_) for the target qubit. This gives us an operation that takes a single Qubit and substitutes in the blank.

It may seem a bit confusing that BiasedPreparation returns an operation from a function, but this is completely consistent with the split between functions and operations described above, since BiasedPreparation is still predictable. In particular, BiasedPreparation (p) always returns the same operation for a given p, no matter how many times you call the function. We can assure ourselves that this is the case by noticing that BiasedPreparation only partially applies operations, but never calls them.

Exercise 6.4

Partial application works for functions as well as operations! Try it out by writing a function Plus that adds two integers, n and m, and another function PartialPlus that takes an input n and returns a function that adds n to its input.

```
HINT: You can get started using the following code snippet as a template.
function Plus(n : Int, m : Int) : Int {
    // fill in this part
}
function PartialPlus(n : Int) : (Int -> Int) {
    // fill in this part
}
```

6.4 PLAYING MORGANA'S GAME IN Q#

With first-class operations and partial application at the ready, we can now make a more complete version of Morgana's game.

EXAMPLE 6.8. THE Q# STANDARD LIBRARIES

The Quantum Development Kit comes with a variety of different standard libraries that we'll see throughout the rest of the book. In , for example, we make use of an operation MResetZ that both measures a qubit (similar to M) and resets it (similar operation to Reset). This offered the Microsoft Quantum Measurement namespace, one of the main standard libraries that comes with the Quantum Development Kit. A full list of the operations and functions available that namespace can he found at https://docs.microsoft.com/qsharp/api/qsharp/microsoft.quantum.measurement. For now, though, don't worry too much about it; we'll see more of the Q# standard libraries as we go.

Listing 6.9. Complete listing of Q# operations for the

biased PlayMorganasGame example open Microsoft.Quantum.Math; open Microsoft.Quantum.Measurement; operation PrepareBiasedCoin(winProbability : Double, qubit : Qubit) : Unit { let rotationAngle = 2.0 * ArcCos(Sqrt(1.0 - winProbability)); Ry(rotationAngle, qubit); operation NextRandomBit(statePreparation : (Qubit => Unit)) : Result { using (qubit = Qubit()) { statePreparation(qubit); return MResetZ(qubit); } } operation PlayMorganasGame(winProbability : Double) : Unit { mutable nRounds = 0; mutable done = false; let prep = PrepareBiasedCoin(winProbability, _); repeat { set nRounds = nRounds + 1; set done = (NextRandomBit(prep) == Zero); until (done) fixup {} Message(\$"It took Lancelot {nRounds} turns to get home.");

- We start by opening namespaces from the Q# standard library to help with classical math, and to help with measuring qubits.
- 2 The rotation angle chooses the bias the coin has.
- 3 Here we use the operation we passed in as statePreparation and apply it to the qubit.
- **1** The MResetZ operation is defined in the Microsoft.Quantum.Measurement namespace that we open at the beginning of the sample. It measures the qubit in the Z basis and then applies what operations are needed to return the qubit to the $|0\rangle$ state.
- **5** We use partial application to specify the bias for our state preparation procedure, but not the target qubit. While PrepareBiasedCoin has type (Double, Qubit) => Unit, PrepareBiasedCoin(0.2, _) "fills in" one of the two inputs, leaving an operation with type Qubit => Unit, as expected by EstimateBias.

Providing documentation for Q# functions and operations

Documentation can be provided for Q# functions and operations by writing small specially formatted text documents in triple-slash (///) comments before a function or operation declaration. These documents are written in Markdown, a simple text formatting language used on sites like GitHub, Azure DevOps, Reddit, and Stack Exchange, and by site generators like Jekyll. The information in /// comments is shown when hovering over calls to that function or operation, and can be used to make API references similar to those at https://docs.microsoft.com/quantum/.

```
Different parts of /// comments are indicated with section headers, for example /// #
Summary. For example, we may document the PrepareBiasedCoin operation from with the
following:
/// # Summary
/// Prepares a state representing a coin with a given bias.
111
/// # Description
/// Given a qubit initially in the |0\rangle state, applies operations
/// to that qubit such that it has the state \sqrt{p} |0\rangle + \sqrt{(1 - p)} |1\rangle,
/// where p is provided as an input.
/// Measurement of this state returns a One Result with probability p.
111
/// # Input
/// ## winProbability
/// The probability with which a measurement of the qubit should return One.
/// ## qubit
/// The qubit on which to prepare the state \sqrt{p} | 0 + \sqrt{1 - p} | 1.
operation PrepareBiasedCoin(
        winProbability : Double, qubit : Qubit
) : Unit {
    let rotationAngle = 2.0 * ArcCos(Sqrt(1.0 - winProbability));
    Ry(rotationAngle, qubit);
   When using IQ#, you can look up documentation comments by using the? command. For instance,
you can look up the documentation for the X operation by running X? in an input cell.
   For a full reference,
see https://docs.microsoft.com/quantum/language/statements#documentation-comments.
```

To estimate the bias of a particular state preparation operation, we can run the PlayMorganasGame operation repeatedly and count how many times we get a Zero.

Let's pick a value for winProbability and run the PlayMorganasGame operation with that value to see how long it Lancelot will be stuck.

Listing 6.10. Running PlayMorganasGame to see how long Lancelot is stuck.

Try playing around with different values of winProbability! Note that if Morgana really tips the scales, we can confirm that it will take Lancelot quite a long time to make it back to Genevieve.

Listing 6.11. Output of varying Morgana's strategy.

6.5 SUMMARY

In this chapter you learned:

- how to use the Quantum Development Kit to write quantum programs in Q#, and
- how to use Jupyter Notebook to run your quantum programs with a quantum simulator.

In the next chapter, we'll build on these skills by going back to Camelot to find our first example of a quantum algorithm, the Deutsch–Jozsa algorithm.

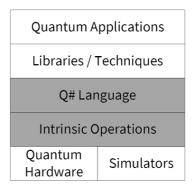
What is a Quantum Algorithm?

This chapter covers:

- · What is a quantum algorithm?
- · How to design oracles to represent classical functions in quantum programs
- · A first example of a quantum algorithm
- · Several useful quantum programming techniques

One important application for quantum algorithms is in obtaining speedups for solving problems where we need to search over inputs to a function that we're trying to learn about. Such functions could be obfuscated (such as hash functions), or could be computationally difficult to evaluate (common in studying mathematical problems). In either case, applying quantum computers to such problems requires us to understand how we program and provide input to quantum algorithms. To learn how to do so, we'll program up and run an implementation of an algorithm known as the *Deutsch–Jozsa algorithm*, which will let us learn properties of unknown functions quickly using quantum devices.

Figure 7.1. In this Chapter, we will be covering topics in both the Q# language and intrinsic operations for quantum devices.



7.1 CLASSICAL AND QUANTUM ALGORITHMS

ALGORITHM

noun: a step-by-step procedure for solving a problem or accomplishing some end. —Merriam-Webster Dictionary

When we talk about classical programming, we sometimes say that a program implements an *algorithm*; that is, a sequence of steps that can be used to solve some problem. If we want to sort a list, for example, we can talk about the quicksort algorithm independently of what language or operating system we are using. We often specify these steps at a high level. In the quicksort example, we might list the steps as something like the following.

7.1.1 Quicksort algorithm

- If the list to be sorted is empty or only has one element, return it as-is.
- Pick an element of the list to be sorted, called the *pivot*.
- Separate all other elements of the list into those that are smaller than the pivot, and those that are larger.
- Quicksort each new list recursively.
- Return the first list, then the pivot, and finally the second list.

These steps then serve as a guide for writing an implementation in a particular language of interest. Say we want to write the quicksort algorithm in Python:

```
Listing 7.1. An example implementation of the quicksort algorithm.

def quicksort(xs):
    if len(xs) > 1:
        pivot = xs[0]
        left = [x in xs[1:] if x <= pivot]
        right = [x in xs[1:] if x > pivot]
        return quicksort(left) + [pivot] + quicksort(right)
    else:
```

©Manning Publications Co. To comment go to liveBook

return xs

- 1 We check for the base case by checking if there's at least two elements in the list.
- 2 We pick the first element to be our pivot for step 2.
- 3 Next, we write out Python code that builds two new lists as described in step 3.
- **4** Finally, we concatenate everything back together as described in steps 4 and 5.

A well-written algorithm can help guide how to write implementations by making the steps that must be executed clear. Quantum algorithms are the same in this respect: they list the steps that we need to perform in any implementation.

QUANTUM PROGRAM

A quantum program is an implementation of a quantum algorithm, consisting of a classical program that sends instructions to a quantum device in order to prepare a particular state or measurement result.

As we saw in Chapter 6, when we write a Q# program, we are writing a classical program which sends instructions to one of several different target machines on our behalf, as illustrated as , returning measurements back to our classical program.

The art of quantum programming

We cannot copy quantum states, but if they resulted from running a program, we can tell someone else what steps they need to take to prepare the same states. As we saw above, quantum programs are a special kind of classical program, so we can copy them with reckless abandon. As we will see throughout the rest of the book, any quantum state can either be approximated or written out exactly by the output of a quantum program that starts with only copies of the $|0\rangle$ state. For example, in Chapter 2, we prepared the initial state $|+\rangle$ of a QRNG by a program consisting of a single $\mathbb H$ instruction.

Put differently, we can think of a program as being a recipe for how to prepare a qubit. Given a qubit, we cannot determine what recipe was used to prepare it, but we can copy the recipe itself as much as we like.

Whereas the steps in executing quicksort instruct the Python interpreter to compare values and to move values around in memory, the steps in a Q# program instruct our target machines to apply rotations and measurements to qubits in a device. As shown in , we can use a host program to send Q# applications to each different target machine to run. For now, we'll keep using the IQ# plugin for Jupyter Notebook as our host program; in the next chapter, we'll see how to use C# to write our own host programs as well.

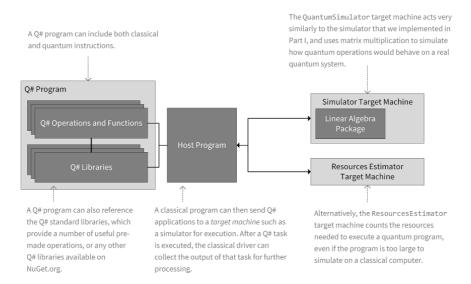


Figure 7.2. The Quantum Development Kit software stack.

Most of the time in this book, we'll be interested in simulating our quantum programs, so we'll use the QuantumSimulator target machine. This simulator works very similarly to the ones we developed in Chapters 2 and 4, as it executes instructions such as the Hadamard instruction H by multiplying quantum states with unitary operators like H.

TIP

As in previous chapters, we can use fonts to distinguish instructions like H from the unitary matrices like H that we use to simulate those instructions.

The ResourcesEstimator target machine allows you to not run a quantum program, but to get estimates on how many qubits it *would* take to run it. This is useful for larger programs that cannot be classically simulated or run on available hardware, to help you learn how many qubits it will take; we'll see more about this target machine later.

Since Q# applications send instructions to the target machines that we use to run them, it's easy to reuse Q# code later across different target machines that share the same instruction set. The QuantumSimulator target machine, for instance, uses the same instructions that we expect actual quantum hardware to take once it becomes available, so that we can test Q# programs on simulators now using small instances of problems, and can then run the same programs on quantum hardware later.

What remains in common across these different target machines and applications is that we need to write the program that sends instructions to the target machine in order to accomplish some goal. Our task as quantum programmers is thus to make sure that these instructions have the effect of solving some useful problem.

TIP

The way that we use simulators to test Q# programs is a bit similar to the way we use simulators to test programs for other specialized hardware like field-programmable gate arrays (FPGAs), or the way we use emulators to test applications for mobile devices from our desktops and laptops. The main difference is that we can only a classical computer to simulate a quantum computer for a very small numbers of qubits, or for restricted kinds of programs.

This is much easier to do when we have an algorithm guiding us to organize the steps that need to happen in both the classical and quantum devices. In developing new quantum algorithms, we can make use of the quantum effects, such as entanglement, which we saw in Chapter 4.

TIP

In fact, to get any advantage from our quantum hardware we must use the unique quantum properties of the hardware, or else we just have a more expensive and slower classical computer.

7.2 DEUTSCH-JOZSA ALGORITHM: MODERATE IMPROVEMENTS FOR SEARCHING

So what might make a good example of a *quantum* algorithm that takes advantage of our shiny new quantum hardware? We learned in Chapters 4 and 6 that thinking about games can often help, and this is no exception. To find our game for this chapter, let's take a trip back to Camelot, where Merlin finds himself facing a test...

7.3 LADY OF THE (QUANTUM) LAKE

Merlin, the famous and wise wizard, has just encountered Nimue, the lady of the lake. Nimue, seeking a capable mentor for the next King of England, has decided to test Merlin to see if he is up to the task. Two bitter rivals, Arthur and Mordred, are vying for the throne, such that if Merlin is to accept Nimue's task, he must choose whom to mentor as king.

For her part, Nimue does not care who becomes king, so long as Merlin can give them sage council. What Nimue is concerned about is whether Merlin, the appointed instructor for the new king, will be reliable and consistent in his leadership.

Since Nimue shares our love of games, she has decided to play a game with Merlin to test if he will be a good mentor or not. Nimue's game, Kingmaker, tests to see if Merlin is *consistent* in his role as advisor to the king. To play Kingmaker, Nimue gives Merlin the name of one of the two bitter rivals for the throne, to which Merlin must respond with whether Nimue's candidate should be the true heir to the throne or not.

7.3.1 Kingmaker game rules

- In each round, Nimue asks Merlin a single question of the form "Should *potential heir* be the king?"
- Merlin must answer either "yes" or "no," giving no additional information.

Each round gives Nimue more information about the realm of mortals, so her objective

is to ask as few questions as is needed to catch Merlin out if he is not trustworthy.

7.3.2 Nimue's objectives

- Verify that Merlin will be a good mentor to the new King of England.
- Ask as few questions as possible to verify.
- Avoid learning who Merlin will say yes to mentoring.

At this point, Merlin has 4 possible strategies:

7.3.3 Merlin's strategies:

- Say "yes" when asked if Arthur should be king, and "no" otherwise (good mentor)
- Say "yes" when asked if Mordred should be king, and "no" otherwise (good mentor)
- Say "yes" regardless of who Nimue asks about (bad mentor)
- Say "no" regardless of who Nimue asks about (bad mentor)

One way to think of Merlin's strategies is by using the concept of a truth table once again. Suppose, for instance, Merlin has decided to be singularly unhelpful and deny any candidates to the throne. Then we might write this down using the truth table in .

Table 7.1. Truth table for one possible strategy of Kingmaker: Merlin always says no.

Input (Nimue)	Output (Merlin)
"Should Morded be king?"	"No."
"Should Arthur be king?"	"No."

At this point, Nimue would be right to complain about Merlin's wisdom as a mentor! Merlin has not been consistent with his charge to choose between Arthur and Mordred. While Nimue may not care whom Merlin picks, he surely must pick *someone* to mentor and prepare for the throne.

Nimue needs a strategy to determine if Merlin has either strategy 1 or 2 (good mentor) or if Merlin is playing according to 3 or 4 (bad mentor) in as few rounds of the game as possible. She could just ask both questions: "Should Mordred be king?" and "Should Arthur be king?" and then compare his answers, but this would result in Nimue knowing for sure who he chose to be king. After all, with each question Nimue learns more about the mortal affairs of the kingdom — how distasteful!

While it would seem Nimue's game is doomed to force her to learn his choice of heir, she is in luck. This being a quantum lake, we'll see throughout the rest of this chapter that Nimue can ask a *single* question that will tell her *only* if Merlin is committed to his role as mentor, and not whom he has chosen.

Since we don't have a quantum lake at our disposal, let's try to model what Nimue is doing with quantum instructions in Q# on our classical computer and then simulate it. Let's represent Merlin's strategy by a classical function f, which takes Nimue's

question as an input x. That is, we'll write f(Arthur) to mean "what Merlin answers when asked if Arthur should be king." Note that since Nimue will only ask one of two questions, which question she asks is an example of a bit. Sometimes it's convenient to write out that bit using the labels "0" and "1", while sometimes it's helpful to label Nimue's input bit using the Boolean values "False" and "True." After all, "1" would be a pretty strange answer to a question like "should Mordred be king?"

Table 7.2. Encoding Nimue's question as a bit

Nimue's question	Representation as a bit	Representation as a Boolean
"Should Mordred be king?"	0	False
"Should Arthur be king?"	1	True

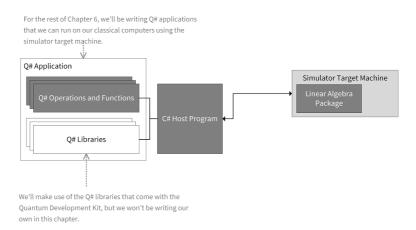
Using bits, we write f(0) = 0 to mean that Nimue if asks Merlin "should Mordred be king," his answer will be no.

If she didn't have any quantum resources, to be sure of what Merlin's strategy is, Nimue would have to try both inputs to f; that is, she'd have to ask both questions to Merlin. Trying all the inputs would give us Merlin's full strategy, which as noted, Nimue is not really interested in.

To solve this, instead of having to ask Merlin about both Mordred and Arthur, we can implement a quantum algorithm in Q# that uses quantum effects to learn whether Merlin is a good mentor or not by asking him only **one** question. Using the simulators provided with the Quantum Development Kit, we can even run our new Q# program on our laptops or desktops!

In the rest of this chapter, we'll learn an example of how to write this quantum algorithm, called the Deutsch-Jozsa algorithm (see).

Figure 7.3. Where we will be working in the Q# software stack for this chapter.



Let's try and sketch out what our quantum program will look like.

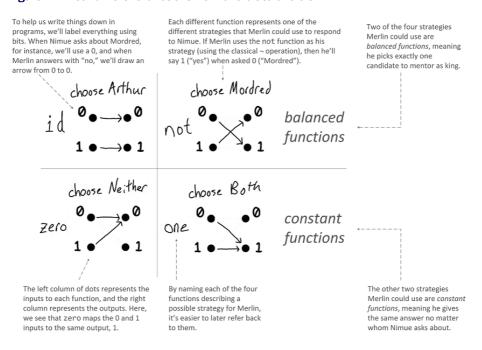
The possible inputs and outputs for f (Merlin's strategy) are True and False. We can write down a truth table for f using the inputs and outputs we get when we call f. For instance, if f is the classical NOT operation (often denoted \neg), then we will observe that f(True) is False and vice versa. As shown in , using a classical NOT operation as a strategy in our game corresponds to picking Mordred to be king.

Table 7.3. Truth table for the classical NOT operation

Input	Output
True ("Should Arthur be king?")	False("No.")
False ("Should Mordred be king?")	True ("Yes.")

There are four possible options for the definition of our function f, each of which represents one of the four strategies available to Merlin, as summarized in .

Figure 7.4. Four different functions from one bit to one bit.



Two of these functions, labeled id and not for convenience, send each of the 0 and 1 inputs to different outputs; we call these functions *balanced*. In our little game, they represent the cases in which Merlin picked exactly one person to be king.

Table 7.4. Classifying Merlin's strategies as either constant or balanced.

Merlin's strategy	Function	Туре	Passes Nimue's challenge?
Choose Arthur	id	Balanced $(f(0) \neq f(1))$	Yes
Choose Mordred	not	Balanced $(f(0) \neq f(1))$	Yes
Choose neither	zero	Constant $(f(0) = f(1))$	No
Choose both	one	Constant $(f(0) = f(1))$	No

On the other hand, the functions that we label as zero and one are each *constant* functions, since they send both inputs to the same output. Constant functions then represent strategies in which Merlin is being decidedly useless, as he's either picked both to be king (a good way to start a bad war), or because he's picked neither.

Classically, if we want to know if a function is constant or balanced (whether Merlin is a bad or good mentor, respectively), we have to learn the entire function by building up its truth table. Remember, Nimue wants to ensure Merlin is a reliable mentor. If Merlin is following a strategy represented by a constant function, then he will not be a good mentor. Looking at the truth tables for the id and one functions, and respectively, we can see how these describe when Merlin is following a strategy that will let him be either a good or bad mentor.

Table 7.5. Truth table for the id function, an example of a balanced function.

Input	Output
True ("Should Arthur be king?")	True ("Yes")
False ("Should Mordred be king?")	False("No")

Table 7.6. Truth table for the one function, an example of a constant function.

Input	Output
True ("Should Arthur be king?")	True ("Yes")
False ("Should Mordred be king?")	True ("Yes")

The difficulty that Nimue faces in trying to learn whether Merlin is a good or bad mentor (that is, whether f is balanced or constant) is that the quality of Merlin's mentorship is a kind of *global* property of his strategy. There's no way to look at a single output of f and conclude anything about what f would output for different inputs. If we only have access to f, then Nimue is stuck: she must reconstruct the entire truth table to decide whether Merlin's strategy is constant or balanced.

On the other hand, if we can represent Merlin's strategy as a part of a quantum program, then we can use the quantum effects we've learned about so far in the book. Using quantum computing, Nimue can learn *only* if his strategy is constant or balanced, without having to learn exactly which strategy he's using.

Since we are not interested in all the additional info the truth table provides beyond whether Merlin is a good or bad mentor, using quantum effects can help us learn what we care about more directly. With our quantum algorithm we can do this with one call of the function, and without needing to learn any additional information we are not interested in. By not asking for all of the details of the truth table, but only looking for more general properties of our function, we can best utilize our quantum resources.

The power of quantum computing

If we want to use a *classical* computer to learn whether a function is constant or balanced, we have to solve a harder problem first, namely of identifying exactly which function we have. By contrast, quantum mechanics lets us solve only the problem we care about (constant vs balanced) without solving the harder problem a classical computer has to solve.

This is an example of a pattern we'll see throughout the book, in which quantum mechanics lets us specify less powerful algorithms than we can express classically.

To do so, we will use the Deutsch–Jozsa algorithm, which uses a *single query* to our quantum representation of Merlin's strategy to learn he is a good or bad mentor. The advantage here isn't terribly practical (a savings of only one question) but that's OK, we'll see more practical algorithms later on in the book. For now, the Deutsch–Jozsa algorithm is a great place to start learning how to implement quantum algorithms, and even more importantly, to start learning what tools we can use to understand what quantum algorithms do.

7.4 ORACLES: REPRESENTING CLASSICAL FUNCTIONS IN OUANTUM ALGORITHMS

Let's see what things look like from Nimue's quantum lake. As we plunge in for a swim, we face a somewhat immediate question: how can we implement the function f that represents Merlin's strategy with qubits? From the previous section, we saw that the classical function f is our description of a strategy that Merlin uses to play each round of Kingmaker. Since f is classical, it's easy to translate this back into a set of actions that Merlin will take: Nimue gives Merlin a single classical bit (her question), and Merlin gives Nimue a classical bit back (his answer).

To avoid meddling in the affairs of mortals, Nimue now wants to use the Deutsch–Jozsa algorithm instead. Since she lives in a quantum lake, Nimue can easily allocate a qubits to give Merlin instead. Lucky for us, Merlin knows how to communicate with qubits, but we still need to figure out what Merlin will do with Nimue's qubits in order to act on his strategy.

The trouble is that we can't pass qubits to the function f that we use to represent Merlin's strategy up above, but f takes and returns classical bits, not qubits. For Merlin to use his strategy to guide what he does with Nimue's qubits, we want to turn Merlin's

strategy *f* into a kind of quantum program known as an *oracle*.

Conveniently for us, Merlin plays the role of an oracle pretty well.

MERLIN AND REVERSIBILITY

From the T. H. White treatment of Merlin, we learn he lives life backwards in time. We'll represent that by making sure that everything Merlin does is unitary. As we saw in Chapter 2, one consequence of this is that the transformations Merlin applies are reversible. In particular, Merlin won't be able to measure Nimue's qubits, since measurement is not reversible. That privilege is Nimue's alone.

To understand what we need to do to model Merlin's actions as an oracle we have to figure out two things:

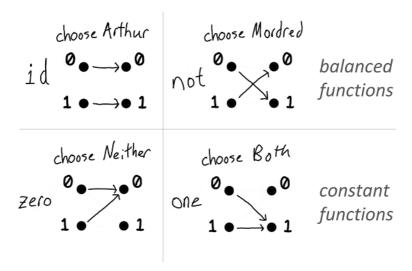
- What transformation should Merlin apply to Nimue's qubits based on his strategy?
- What quantum operations will Merlin need to apply to implement that transformation?

Unitary matrices and truth tables

Another way of saying what we need to do in step 1 is that we need to find a *unitary matrix* that represents what Merlin does, similarly to how we used classical functions like f to represent what Merlin did when Nimue gave him classical bits. As we saw in Chapter 2, unitary matrices are to quantum computing as truth tables are to classical computing: they tell us what the effect of a quantum operation is for every possible input. Once we found the right unitary, step 2 is where we'll figure out what sequence of quantum operations we can do that will be described by that unitary.

To complete step ①, we need to turn functions like f into unitary matrices, so let's start by recapping what f can be. The possible strategies Merlin could use are represented by each of the functions id, not, zero, and one (see).

Figure 7.5. Four different functions from one bit to one bit.



For the two balanced functions, id and not in figure 7.5, we can answer the first question easily. Quantum programs for id and not can be implemented as rotation operations, so it is easy to turn them into quantum operations. The quantum NOT operation, for instance, is a rotation of 180° around the *X* axis, exchanging the $|0\rangle$ and $|1\rangle$ states with each other.

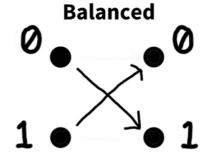
Recall from Chapter 4 that the quantum operation X, represented by the unitary matrix $X = (^{\circ}_{1})^{\circ}_{0}$, is used to apply a rotation of 180° around the X axis. This operation implements a quantum NOT: since $X \mid 0 \rangle = \mid 1 \rangle$ and $X \mid 1 \rangle = \mid 0 \rangle$, we can write using the ¬ (NOT) operator from Chapter 2 as $X \mid x \rangle = \mid \neg x \rangle$.

While any rotation can be undone by rotating the same amount in the opposite direction, we run into more problems with the constant functions zero and one. Neither zero nor one can be implemented directly as rotations, so we have a bit more work to do. For instance, if f is zero, then the outputs f(0) and f(1) are both 0. If we only have the output 0, we cannot tell whether we got that output from giving f a 0 or a 1 as input (see).

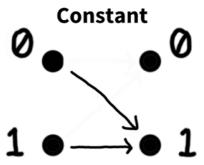
NOTE Once we apply zero or one, we have lost any information about the input.

Figure 7.6. Why can't we reverse the constant zero or one functions?

Example 1: We can see with the way the inputs and outputs are mapped that if we know what the output of the function, then we can trace it back to what the input was!



Example 2: We can see that with constant functions that though we know the output, it is ambigous what was the input. It could have been either 0 or 1 but we can't know...



Since one and zero are both irreversible, and valid operations on qubits are reversible, Merlin will need another way of representing functions like f in quantum algorithms such as the one for Nimue's challenge. On the other hand, if we can represent Merlin's strategy by a reversible classical function instead of f, it will be much easier to write down a quantum representation of his strategy.

7.4.1 Our strategy for representing classical functions as quantum oracles

- Find a way to represent our irreversible classical function by a reversible classical function
- Write down a transformation on quantum states using our reversible classical function
- Figure out what quantum operations we can do that result in that transformation

Let's use the tried and true approach of guessing and checking to see if we can design a valid reversible classical function! The easiest way to figure out whether we were given a 0 or a 1 as an input is to just record it somewhere, so let's make a new function that returns two bits instead of one:

First attempt: Record and keep the input

g(x) = (x, f(x))

As before, we'll represent functions by drawing arrows

For example, if Merlin uses the strategy one (that is, says "yes" to Nimue no matter what she asks), then f(x) = 1, and g(x) = (x, f(x)) = (x, 1).

This gets a lot closer, since we can now tell whether we started with a 0 or 1 input, but we're not quite there, since g has two outputs and one input (see).

Figure 7.7. First attempt: keeping input and output with g(x).

from possible inputs to the corresponding outputs. The one function, for example, sends both 0 and 1 to 1. Original function: First attempt: f(x) = 1g(x) = (x, f(x)) = (x, 1)(0,0)These two possible outputs have no inputs that (0,1)generate them, so we can't reverse this function! (1,0)(1,1)

To use g as a strategy, Merlin would have to give back Nimue more qubits than she gave him, but she is the keeper of both swords and qubits. More technically, reversing

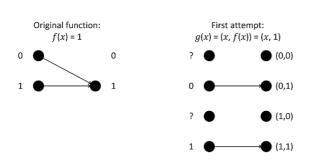
©Manning Publications Co. To comment go to liveBook

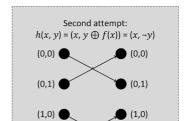
g would thus destroy information, as it would take two inputs and return one output!

Trying one more time, let's define a new classical function h that takes two inputs and returns two outputs, h(x, y). Let's again consider the example where we describe Merlin's strategy by the function f(x) = 1. Since g got us nearly there, we'll choose h such that h(x, 0) = g(x). We saw from the our first attempt that when Merlin uses the strategy f(x) = 1, then g(x) = (x, 1), so we have that h(x, 0) = (x, 1). Now we just need to define what happens when we pass y = 1 to h. If we want h to be reversible, we need that h(x, 1) is assigned to something other than (x, 1). One way to do this is to let h(x, y) = (x, y), so that $h(x, 1) = (x, 0) \neq (x, 1)$.

This choice is especially convenient since applying h twice gets us back our original input, $h(h(x, y)) = h(x, \neg y) = (x, \neg \neg y) = (x, y)$.

Figure 7.8. Second attempt: h(x, y) which is reversible and has the same number of inputs and outputs!





(1,1)

(1,1)

Hurray! By having 2 input bits and 2 output bits our function is reversable. Here, that means we can map every output to a unique input, and vice versa.

Now that we know how to make a *reversible* classical function from each strategy, let's finish by making a quantum program from our reversible function. In the case of one, we saw that h flips its second input, $h(x, y) = (x, \neg y)$. Thus, we can write a quantum program that does the same thing as our reversible classical function simply by flipping the second of two input qubits. As we saw in Chapter 4, we can do this using the x instruction, since $X|x\rangle = |\neg x\rangle$.

Example 7.2. Generalizing our results

More generally, we can make a reversible quantum operation, in precisely the same way we made reversible classical functions h by flipping an output bit based on the output of the irreversible function f. We can define the unitary matrix (that is, the quantum analogue of a truth table) U_f for f for each input state in exactly the same way.

Figure 7.9. Constructing reversible classical functions and unitary matrices from irreversible classical functions

$$h(x,y) = (x, y \oplus f(x)) \qquad U_f|x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle$$

We can make a new reversible classical function h from an irreversible function f by flipping a bit based on the output of f.

To define h, we specify what it does for arbitrary classical bits x and y.

In exactly the same way, we can define a unitary matrix $\mathcal{U}_{\mathrm{f}}.$

Much like we defined h by saying what it did for each classical bit x and y, we can say what U_f does for input qubits in the labeled by classical bits; that is, in the $|0\rangle$ and $|1\rangle$ states.

Defining U_f in this way makes it easy to undo the call to f, since applying U_f twice gives us the identity 1 (that is, the unitary matrix for the "do nothing" instruction). When we define a unitary in this way by applying a function f conditionally to the labels for qubit states, we call this new operation an *oracle*.

ORACLE

An oracle is a quantum operation represented by a unitary matrix \mathbf{U}_f that transforms its input state as

$$U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$$

The symbol \oplus represents the exclusive or operator from regular boolean logic.

Now all that's left is to figure out what sequence of instructions we need to send to implement each unitary U_f . We've seen above how what instructions we need to implement an oracle for one, namely an x instruction on the second qubit. Now let's look at how to write oracles for other possible functions f. That way, Merlin will know what he should do no matter what his strategy is.

DEEP DIVE: Why is it called an oracle?

So far, we've seen a few different examples of the kind of whimsical naming that quantum computing owes to its physics history, like "bra," "ket," and "teleportation." It's not just physicists who like to have a bit of fun, though! One branch of theoretical computer science called *complexity theory* explores what is possible to do efficiently even in principle given different kinds of computing machines. You may have heard, for instance, of the P versus NP problem, a classic conundrum in complexity theory that asks whether problems in P are as difficult to solve as those NP or not. The complexity class P is the group of questions for which there exists a way to answer them with an algorithm that takes polynomial time. By contrast, NP is the group of questions for which we can check a potential answer in polynomial time, but we don't know if we could come up with an answer from scratch in polynomial time.

Many other problems in complexity theory are posed by introducing small games or stories to help researchers remember what definitions to use where. Our own little story about Merlin and Nimue is a nod to this tradition. In fact, one of the most celebrated stories in quantum computing is called MA for "Merlin-Arthur." Problems in the class MA are thought of using a story in which Arthur gets to ask an Merlin, an all-powerful but untrustworthy wizard, some set of questions. A yes/no decision problem is in MA if, whenever the answer is "yes," there exists a proof that Merlin can give Arthur and that Arthur can easily check using a P machine and a random number generator.

The name "oracle" fits into this kind of storytelling, in that any complexity class A can be turned into a new complexity class AB by allowing A machines to solve a B problem in a single step, as though they were consulting an oracle. Much of the history of problems like the Deutsch–Jozsa problem stems from trying to understand how quantum computing affects computational complexity, so many of the naming conventions and much of the terminology has been adopted into quantum computing itself.

For more on complexity theory and how it relates to quantum computing, black holes, free will, and Greek philosophy, check out *Quantum Computing Since Democritus* (ISBN: 978-0521199568).

In general, finding a sequence of instructions by starting from a unitary matrix is a mathematically difficult problem known as *unitary synthesis*. That said, in this case, we can figure it out by substituting each of Merlin's strategies f into our definition for U_f and identifying what instructions will have that effect — we can guess and check in the same way that we did to turn the one function into an oracle.

Let's try this out for the zero function:

Exercise 7.1: Try out writing an oracle!

What would the oracle operation U_f be for f if f was zero?

Solution: Let's work it out one step at a time.

From the definition for U_f , we know that $U_f |xy\rangle = |x\rangle |y \oplus f(x)\rangle$.

Substituting in zero for f, f(x) = 0, we get that $U_f |xy\rangle = |x\rangle |y \oplus 0\rangle$.

We can use that $y \oplus 0 = y$ to simplify this a little bit further, getting that $U_{I}(xy) = |x\rangle |y\rangle$.

At this point, we notice U_f does nothing to its input state, so we can implement it by... doing nothing.

The function f = id is slightly more subtle than the zero and one cases because $y \oplus f(x)$ cannot be simplified to not depend on x. As summarized in , we need $U_f|x\rangle|y\rangle = |x\rangle|y\oplus f(x)\rangle = |x\rangle|y\oplus x\rangle$. That is to say, we need the action of the oracle on the input state $(|x\rangle|y\rangle)$ to leave x alone and replace y with the exclusive or of y and y.

Another way to think of this output is to is recall that $y \oplus 1 = \neg y$, so that that when x = 1, we need to flip y. This is precisely how we defined the controlled-NOT (CNOT) instruction in Chapter 5, so we recognize that when f is id, U_f can be implemented by applying a CNOT.

This leaves us with how to define the oracle for f = not. Just as the oracle for id flips the output (target) qubit when the input (control) qubit is in the $|1\rangle$ state, the same argument gives us that we need our oracle for not to flip the second qubit when the input qubit is in $|0\rangle$. The easiest way to do this is to first flip the input qubit with an x instruction, apply a CNOT instruction, and then undo the first flip with another x.

To review all of the oracles we have learned who to define, we collected all our work in the above section in .

Table 7.7. Table Oracle outputs for each one-bit function f

Function Name	Function	Output of Oracle
zero	f(x)=0	$ x\rangle y\oplus 0\rangle = x\rangle y\rangle$

©Manning Publications Co. To comment go to liveBook

Function Name	Function	Output of Oracle	
one	f(x)=1	$ x\rangle y \oplus 1\rangle = x\rangle \neg y\rangle$	
id	f(x) = x	$ x\rangle y\oplus x\rangle$	
not	$f(x) = \neg x$	$ x\rangle y\oplus \neg x\rangle$	

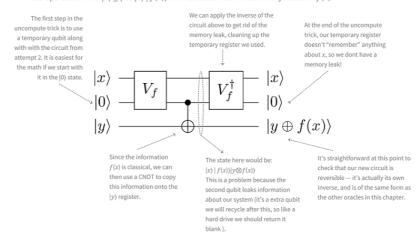
The uncompute trick: Turning functions into quantum oracles

As it stands, it may seem as though it takes a lot of work to design U_f for each function f. Thankfully, there's a nice trick that lets us build an oracle, starting with a somewhat simpler requirement.

Recall that above, we attempted to define a reversible version of f by returning (x, f(x)) as output when given (x, 0) as input. Similarly, suppose that we are given a quantum operation V_f that correctly transforms $|x\rangle |0\rangle$ to $|x\rangle |f(x)\rangle$. Then, we can always make an oracle U_f by using one additional qubit, and by calling V_f twice using a technique called the "uncompute trick," as shown in .

Figure 7.10. Using the "uncompute trick" to turn an operation that only works when you add an extra |0) input qubit to an operation that can be used as an oracle.

The uncompute trick: If we add another input $|0\rangle$, then an operation which gives us a way to make the operation from $|x\rangle|y\rangle$ to $|x\rangle|f(x)\rangle$ is reversible! This works for any function f(x).



While this doesn't especially help in the case of Deutsch–Jozsa, it shows us that the concept of an oracle is a very general one, as it's often much easier to start with an operation of the form V_f .

NOTE

The oracle construction also works for multiple qubit functions. As a thought exercise, if we have a function $f(x_0, x_1) = x_0$ AND x_1 , how would the oracle U_f transform an input state $|x_0x_1y\rangle$? We'll see in later chapters how to code up this oracle.

We've thus used the oracle representation to solve the problem that functions like zero and one cannot be represented as rotations. With that dealt with, we can continue on to actually write the rest of the algorithm that Nimue uses to challenge Merlin.

DEEP DIVE: Other ways to represent functions as oracles

This isn't the only way we could have defined U_f . Merlin could have also flipped the sign of Nimue's input x when f(x) is one,

$$U_f |x\rangle = (-1)^{f(x)} |x\rangle$$

This turns out to be a more useful representation in some cases, such as in gradient descent algorithms. These algorithms are common in machine learning, and minimize functions by searching along directions in which a function changes the fastest. For more information, see Chapter 4.10 of *Grokking Deep Learning*.

Picking the right way for a particular application to represent classical information such as subroutine calls within a quantum algorithm is a part of the art of quantum programming. For now, we will use the definition of "oracle" introduced above.

With the oracle representation of f in hand, the first few steps of the Deutsch-Jozsa algorithm can be written out in the same sort of pseudocode that we used to write quicksort earlier:

7.4.2 Deutsch-Jozsa Algorithm

- Prepare two qubits labeled control and target in the $|0\rangle \otimes |0\rangle$ state.
- Apply operations to the control and target qubits to prepare the following state: |+⟩⊗|-⟩.
- Apply the oracle U_f to the input state $|+\rangle \otimes |-\rangle$. Recall that $[U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$.
- Measure the control qubit in the X basis; if we observe a 0, then the function is constant, otherwise the function is balanced.

Measuring a qubit in the X basis always returns a 0 or 1 just like if we measured in the Z basis. Recall from Chapter 3 that if the state of the qubit is $|+\rangle$, we always get a 0 when we measure in the X basis, while we always get a 1 if the qubit is in $|-\rangle$.

Let's look at a figure illustrating these steps:

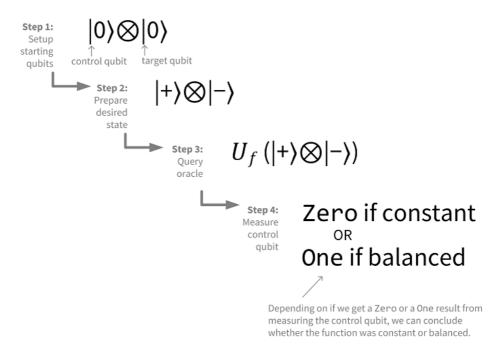


Figure 7.11. Steps to the Deutsch-Jozsa Algorithm.

We'll see at the end of the chapter **why** this algorithm works, but let's jump right in start programming it up. To do so, we'll use the Q# language provided by the Quantum Development Kit, since this will make it much easier for us to see the structure of a quantum algorithm from its source code.

7.5 SIMULATING THE DEUTSCH-JOZSA ALGORITHM IN Q#

We tried out in chapter 6 passing operations as arguments in Q# programs. We can use the same approach of passing operations as inputs to pass oracles to help us predict how Nimue's challenge will turn out. To do so, recall that there are four possible functions that we can consider for this problem, each representing a possible strategy that Merlin could use, as shown in .

 Table 7.8. Representing one-bit functions as two-qubit oracles.

Function Name	Function	Output of Oracle	Q# Operation
zero	f(x)=0	$ x\rangle y\oplus 0\rangle = x\rangle y\rangle$	NoOp(control, target);
one	f(x) = 1	$ x\rangle y \oplus 1\rangle = x\rangle \neg y\rangle$	X(target);

©Manning Publications Co. To comment go to liveBook

Function Name	Function	Output of Oracle	Q# Operation	
id	f(x) = x	$ x\rangle y \oplus x\rangle$	CNOT(control, target)	
not	$f(x) = \neg x$	$ x\rangle y\oplus \neg x\rangle$	<pre>X(control); CNOT(control, target); X(control);</pre>	

If we represent each function f(x) by an oracle (quantum operation) that maps $|x\rangle |y\rangle$ to $|x\rangle |y\oplus f(x)\rangle$, then we can identify each of the functions zero, one, id, and not from .

Each of the four oracles above translates immediately into Q#:

Listing 7.3. Oracles.qs: Q# operations for each Deutsch-Jozsa oracle

```
namespace DeutschJozsa {
    open Microsoft.Quantum.Intrinsic;

    operation ApplyZeroOracle(control : Qubit, target : Qubit) : Unit {
     }

    operation ApplyOneOracle(control : Qubit, target : Qubit) : Unit {
         X(target);
    }

    operation ApplyIdOracle(control : Qubit, target : Qubit) : Unit {
         CNOT(control, target);
    }

    operation ApplyNotOracle(control : Qubit, target : Qubit) : Unit {
         X(control);
         CNOT(control, target);
         X(control);
    }
}
```

Can't We Just Look at the Source Code?

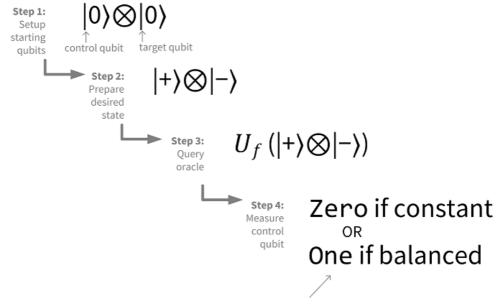
In Oracles.qs, we wrote the source code out for each of the four single-qubit oracles.qs, we wrote the source code out for each of the four single-qubit oracles.ApplyZeroOracle, ApplyOneOracle, ApplyIdOracle, and ApplyNotOracle. Looking at that source code, we can tell whether each is constant or balanced without having to call it at all, so why should we worry with the Deutsch-Jozsa algorithm? Thinking from Nimue's perspective, she doesn't necessarily have the source code that Merlin uses to apply operations to her qubits. Even if she does, Merlin's ways are inscrutable, such that she may not be able to easily predict what Merlin does even given the source code that he uses.

Practically speaking, while it's hard to obfuscate a two-qubit oracle all that much, the Deutsch–Jozsa algorithm demonstrates a technique that is useful more generally. For example, we might have access to the source code for some operation, but it is a mathematically or computationally difficult problem to extract the answer to a question about that operation. All cryptographic hash functions have this property by design, whether they're used to ensure that a file has been downloaded correctly, an application has been signed by a developer, or as a part of growing a blockchain through mining for collisions.

We'll see an example in Chapter 10 where we can use techniques such as those developed in the Deutsch-Jozsa algorithm to more quickly ask questions about such functions.

With these oracles now implemented in Q#, we can write the entirety of the Deutsch–Jozsa algorithm (as well Nimue's strategy for Kingmaker)!

Figure 7.12. Steps to the Deutsch-Jozsa Algorithm.



Depending on if we get a Zero or a One result from measuring the control qubit, we can conclude whether the function was constant or balanced.

Listing 7.4. Algorithm.qs: Q# operation implementing the Deutsch-Jozsa algorithm

©Manning Publications Co. To comment go to liveBook

- This using statement asks the target machine to give us two qubits, control and target, each starting in the |0⟩ state.
- 2 These lines prepare the input state $|+-\rangle = (|00\rangle |01\rangle + |10\rangle |11\rangle) / 2$ on control and target, as shown in Step $\bigcirc{1}$ of .
- 3 Here, we call the oracle that we were given as the input argument oracle. Note that the oracle is only called once!
- 4 We know that the target qubit is still in |->, so we can undo the X(target); H(target); sequence of operations above to reset the target qubit.
- **5** Finally, we can measure whether the control qubit is in $|+\rangle$ or $|-\rangle$, corresponding to Zero or One results in the X basis. Similarly to the MResetZ operation provided by the Q# standard libraries, the MResetX operation performs the desired X-basis measurement and resets the measured qubit to $|0\rangle$.

Now as we want to make sure our implementation is good, so let's test it!

Measurement results in Q#

We've now seen both the MResetX and MResetZ operations in Q#, which measure and reset a qubit in the X and Z bases, respectively. Both of these operations return a Result value, which seems a little bit confusing at first. After all, an X basis measurement tells us whether we were in the $|+\rangle$ or the $|-\rangle$ state, so why does Q# use the labels Zero and One?

Table 7.9. Conventions used for X- and Z-basis measurement results in Q#.

Result value	X basis	Z basis
Zero	(+	(0)
One	(-	(1

We'll see in more detail on this later, but the short version is that a value of type Result tells us how many about how many phases of (-1) are applied to a state by different instructions. For example, $Z \mid 1 \rangle = - \mid 1 \rangle = (-1)^1 \mid 1 \rangle$, while $X \mid - \rangle = (-1)^1 \mid - \rangle$. Since in both cases, we raise (-1) to the power of 1, $\mid 1 \rangle$ and $\mid - \rangle$ are assigned to the One result when we measure in the Z and X bases, respectively. Similarly, since $Z \mid 0 \rangle = (-1)^0 \mid 0 \rangle$, we assign $\mid 0 \rangle$ to the Zero result when we measure Z.

We said earlier that Nimue would like to learn as little as she can about the affairs of humankind. Thankfully, she only asks Merlin to do something with her qubits once, where we call <code>oracle(control, target)</code>. Nimue only gets one classical bit of information out, from the call to <code>MResetx</code>, which isn't enough for her to tell the difference between the <code>id</code> strategy (Merlin selects Arthur to mentor as king) and the <code>not</code> strategy (Merlin selects Mordred to mentor).

To make sure that she can still learn what she actually cares about, whether Merlin's strategy is constant or balanced, we can use the Fact function provided with the Q# standard libraries to test that our implementation works. Fact takes two boolean variables as the first two arguments, checks to see if they are equal, and if not issues a message.

TIP

Later, we'll see how to use these assertions to write unit tests for quantum libraries.

Listing 7.5. Algorithm.qs: Q# operation testing that the Deutsch-Jozsa algorithm works

as intended

- This line runs the Deutsch–Jozsa algorithm for the case in which Merlin uses the zero strategy. We pass the oracle for zero as the ApplyZeroOracle operation that we wrote above. Since zero isn't a balanced function, we expect ChecklfOraclelsBalanced(ApplyZeroOracle) to return false as its output; this expectation can be checked using the Fact function.
- We can do exactly the same thing for the one strategy, this time calling CheckIfOracleIsBalanced(ApplyOneOracle) instead.
- 3 If all four assertions passed, then we can be sure that our program for the Deutsch–Jozsa algorithm works regardless of which strategy Merlin uses.

If we run this using the %simulate magic command, we can confirm that by using the Deutsch-Jozsa algorithm, Nimue is able to learn exactly what she wanted to learn about Merlin's strategy.

```
In [ ]: %simulate RunDeutschJozsaAlgorithm
All tests passed!
```

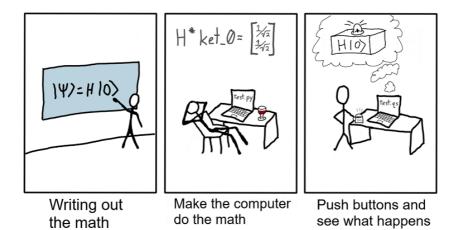
7.6 EXPLORING THE DEUTSCH-JOZSA ALGORITHM BY EXAMPLE

Now that we've implemented the Deutsch-Jozsa algorithm for ourselves, it's helpful to take a step back and see how it works, as we'll learn some valuable quantum programming techniques along the way.

Trying things out!

In Chapters 2 and 4, we learned to use NumPy and QuTiP to simulate how the states of qubits are transformed when we send instructions to a quantum computer. We effectively used these packages to do math for us to find out what would happen to our quantum states. This is like the "make the computer do the math" approach in .

Figure 7.13. Three different approaches of learning how a quantum program or algorithm works.



Now that we're programming larger algorithms in Q#, we can use both the what we learned before and a bit of the right box (push all the buttons!) to help us predict what a particular operation will do.

These three approaches when used together are powerful problem solving tools when learning quantum programing! If you get stuck using one approach you can always try another to see if that helps.

Recall the four steps to Deutsch-Jozsa algorithm in as we just programmed in the previous section.

Listing 7.6. The four steps in the Deutsch-Jozsa algorithm.

```
// 1. Prepare the input state |+-).
H(control);
X(target);
H(target);

// 2. Apply the oracle.
oracle(control, target);

// 3. Undo the preparation on the target qubit.
H(target);
X(target);

// 4. Finally, we measure in the X basis.
set result = MResetX(control);
```

The key to understanding how the Deutsch-Jozsa algorithm works is understanding the step where we call the oracle, <code>oracle(control, target)</code>. Before we can get to that, though, we need to understand step 1, where we prepare our input to <code>oracle</code>.

7.7 STEP 1. PREPARING THE INPUT STATE FOR DEUTSCH-JOZSA

Let's use Python to try and understand what is happening when we prepare our |+--> state. The operations we used to prep the input state in Q# were:

Listing 7.7. State prep for the Deutsch-Jozsa algorithm.

```
H(control);
X(target);
H(target);
```

Each of the operations applied above are single qubit gates, so we can consider what happens to each qubit independently. Let's look at what happens to the control qubit after the Hadamard operation.

Listing 7.8. Using QuTiP to model the control qubit state preparation.

- **1** While H in Q# is an instruction, hadamard_transform in QuTiP gives us a unitary matrix that we can use to simulate how the H instruction transforms states.
- **2** 1 / $\sqrt{2} \approx 0.707$, so this output tells us that $H = \begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 \end{pmatrix} / \sqrt{2}$.
- **3** In QuTiP, we can get the vector for the $|0\rangle$ state by calling basis(2, 0). The 2 tells QuTiP we want a qubit (necessary dimension of $|0\rangle$), while the 0 tells us we want the state to have the value $|0\rangle$. Since $|+\rangle = H|0\rangle$, this sets control state to be $|+\rangle$.
- 4 Using that $1/\sqrt{2} \approx 0.707$, we read this as telling us that $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$.

So that's a pretty easy one, the control qubit is now in the $|+\rangle$ state. Now let's look at preparing the target qubit in .

Listing 7.9. Using QuTiP to model the target qubit state preparation.

```
>>> target_state = H * (qt.sigmax() * qt.basis(2, 0))
>>> target_state
Quantum object: dims = [[2], [1]], shape = (2, 1), type = ket
Qobj data =
[[ 0.70710678]
[-0.70710678]]
```

- 1 We repeat the same H operation as before, but this time on $X|0\rangle = |1\rangle$.
- **2** QuTiP tells us that $|-\rangle = (|0\rangle |1\rangle) / \sqrt{2}$. That is, the same as $|+\rangle$, but with the sign of $|1\rangle$ flipped.

Now that we have see how we prepare each qubit, let's have QuTiP help us write the

©Manning Publications Co. To comment go to liveBook

state of our input register.

Listing 7.10. Using QuTiP to model the combined input state.

```
>>> register_state = qt.tensor(control_state, target_state)
>>> register_state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket 2
Qobj data =
[[ 0.5]
[-0.5]
[ 0.5]
[ 0.5]
[ -0.5]
```

- ① As we saw in Chapter 5, we combine the states of different qubits to get the state of an entire register of qubits using the tensor function.
- **2** QuTiP tells us that $|+\rangle \otimes |-\rangle = |+-\rangle = (|00\rangle |01\rangle + |10\rangle |11\rangle) / 2$. That is, we have an equal superposition over all four possible computational basis states, with a minus sign in front of computational basis states where the target qubit is in the $|1\rangle$ state.

NOTE

As we saw in Chapter 5, when writing down the state of a multi-qubit system, tensor products can get a little bit verbose. Thus, we'll often write down multi-qubit states like $|0\rangle \otimes |1\rangle$ by concatenating their labels inside a single ket, as in $|01\rangle$. Similarly, $|+-\rangle$ is the same as $|+\rangle \otimes |-\rangle$

7.8 STEP 2. APPLYING THE ORACLE

Having prepared our input, let's get back to the core of the Deutsch-Jozsa algorithm, where we call into our oracle .

Listing 7.11. Applying the oracle in the Deutsch-Jozsa algorithm.

```
oracle(control, target);
```

In the same way that we can understand operations like H(control) by writing down the state of the control qubit and applying the unitary operator H to that state, we can understand what the oracle U_f does by analyzing its action on the state we pass to it.

Our quantum oracle operates on two qubits, which raises a question of how we should interpret each of those qubits. In the original classical case, the interpretation of the input and output classical bits from f was clear: Nimue asked a one-bit question and got a one-bit answer. To understand what each qubit does for us, recall that we saw that when we went to a reversible classical function that we also needed two inputs: the first input acted like the question we asked in the irreversible case, while the second input gave us somewhere to put the answer (see for a reminder).

Figure 7.14. Constructing reversible classical functions and unitary matrices from irreversible classical functions

$$h(x,y) = (x, y \oplus f(x)) \qquad U_f|x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle$$

We can make a new reversible classical function h from an irreversible function f by flipping a bit based on the output of f.

To define h, we specify what it does for arbitrary classical bits x and y.

In exactly the same way, we can define a unitary matrix \boldsymbol{U}_f .

Much like we defined h by saying what it did for each classical bit x and y, we can say what U_f does for input qubits in the labeled by classical bits; that is, in the $|0\rangle$ and $|1\rangle$ states.

We can roughly think of the oracle in the same way: the first qubit (control in the snippet above) represents our question, while the second qubit (target) gives us somewhere for Merlin to apply his answer. This interpretation makes sense when control starts off in either the $|0\rangle$ or $|1\rangle$ state, but how can we interpret the case above where we pass qubits in the $|+-\rangle$ state to the oracle? In this case, our control qubit starts off in the $|++\rangle$ state, but f(+) doesn't make any sense. Since f is a classical function, its input has to be either 0 or 1... we can't pass + to the classical function f. It may seem like we would be at a dead end, but thankfully there's a way to figure it out.

Quantum mechanics is linear, which means we can always understand what a quantum operation does by breaking it down into its action on a representative set of states.

TIP As we saw in Chapter 2, a set of states that can be used this way is called a basis.

To understand what our oracle does when the control qubit is in the $|+\rangle$ state, we can use the fact that $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ to break the oracle's action down into what it does to $|0\rangle$ plus its action on $|1\rangle$, then sum both parts back together (making sure to divide by $\sqrt{2}$ at the end). This helps because instead being confused trying to understand what "f(+)" means, we can reduce the action of U_f to cases we do know how to compute like f(0) and f(1)!

COMPUTATIONAL BASIS STATES

Expanding the action of a quantum operation in terms of how it acts on $|0\rangle$ and $|1\rangle$ is very common in quantum programming. Given how useful this is, we often give a special name to these two input states, and call $|0\rangle$ and $|1\rangle$ the computational basis to set it apart from other bases we might use, like $|+\rangle$ and $|-\rangle$.

Using linearity to understand quantum operations isn't limited to a single qubit, as we'll see in the rest of the chapter. For two qubits, for instance, the computational basis for two qubits then consists of the states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$.

If we have even more (say, five) qubits, we can write out states like $|1\rangle \otimes |0\rangle \otimes |0\rangle \otimes |1\rangle \otimes |0\rangle$ as strings in the same way, getting $|10010\rangle$. We can write down the computational basis for five qubits in the same way as $\{|00000\rangle, |00001\rangle, |00010\rangle, ..., |11110\rangle, |11111\rangle$.

More generally, if we have n qubits, then the computational basis consists of all strings of n classical bits, each as the label of a ket. Put differently, the computational basis for a multiple qubit system is made up of all tensor products of $|0\rangle$ and $|1\rangle$; that is, all states labelled by a string of classical bits.

With this approach of breaking down how the oracle works, let's look at some examples of the oracles we have implemented earlier.

EXAMPLE 1: TD ORACLE

Suppose we were given an oracle that implemented the strategy where Merlin chooses Arthur as king. Recall the classical one bit function that represents this strategy is id.

Table 7.10. Representing one-bit functions as two-qubit oracles

Function Name	Function	Output of Oracle	Q# Operation	
id	f(x) = x	$ x\rangle y\oplus x\rangle$	CNOT (control,	target)

From , we know that this means that U_f is implemented by the CNOT instruction, so let's see what that does to register state.

TIP | Recall that the controlled NOT instruction flips its second qubit if the first qubit is in |1).

Listing 7.12. Using QuTiP to compute how the id oracle transforms its input state. >>> cnot = qt.cnot(2, 0, 1) >>> cnot Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True Qobj data = [[1. 0. 0. 0.] [0. 1. 0. 0.] [0. 0. 1.] [0. 0. 1. 0.] >>> register_state = cnot * register_state >>> register_state Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket Qobj data =

```
[[ 0.5]
[-0.5]
[-0.5]
[ 0.5]]
```

- We can ask QuTiP for a matrix that will let us simulate the CNOT instruction by using the cnot function. Here, the 2 indicates that we want to simulate CNOT on a two-qubit register, the 0 indicates that the 0th qubit is our control, and the 1 indicates that the 1st qubit is our target.
- 2 Remember that unitary operators are for quantum computing what truth tables are for classical logic. Each row in this table tells us what happens to a computational basis state.
- 3 For example, this row 2 (zero-indexed), which we write as 10 in binary. Thus, this row is the vector that we'll get out if our input is |10), and tells us that the CNOT instruction leave our qubits in |11) (3 in decimal, hence there's a 1 in the third column).
- 4 QuTiP tells us that the register with our control and target qubits is now in the state ($|00\rangle |01\rangle |10\rangle + |11\rangle / 2$.

Now that we have worked out the action of the id oracle, let's look at what the not oracle does to our input state.

EXAMPLE 2: THE NOT ORACLE

Let's repeat the analysis using the not oracle, the other balanced function. The oracle representing Merlin choosing Mordred is implemented with a series of x and CNOT operations as in .

Table 7.11. Representing the one-bit function not as a two-qubit oracle.

Function Name	Function	Output of Oracle	Q# Operation
not	$f(x) = \neg x$	$ x\rangle y\oplus \neg x\rangle$	<pre>X(control); CNOT(control, target); X(control);</pre>

Let's now jump to Python to see how to break down the operation of the not oracle.

```
Listing 7.13. Using QuTiP again, now with the not oracle.

>>> control_state = H * qt.basis(2, 0)
>>> target_state = H * qt.basis(2, 1)
>>> register_state = qt.tensor(control_state, target_state)
>>> I = qt.qeye(2)
>>> X = qt.sigmax()
>>> oracle = qt.tensor(X, I) * qt.cnot(2, 0, 1) * qt.tensor(X, I)
>>> oracle
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[0. 1. 0. 0.]
[1. 0. 0. 0.]
[1. 0. 0. 0.]
[0. 0. 1. 0.]
[0. 0. 1. 0.]
[0. 0. 0. 1.]
```

```
>>> register_state = oracle * register_state
>>> register_state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[-0.5]
    [ 0.5]
    [ 0.5]
    [ -0.5]]
```

- 1) We prepare the control and target qubits in the |+-) state in exactly the same way as we did before.
- ②As in Chapter 4, it's helpful to define some variables I and X as shorthand for the identity matrix (qt.qeye) and the matrix representing the X operation, respectively.
- 3 This time our oracle is the not oracle, which we implement with the sequence of instructions X(control); CNOT(control, target); X(control); as per .
- 4 The unitary operator for the oracle operation looks a bit different this time: it flips the target qubit when the control qubit is a |0⟩.
- **5** For instance, row 0 (00 in binary) tells us that |00) is transformed to |01).
- (6) Similarly, row 2 (10 in binary) tells us that |10) is transformed to |10); the oracle leaves that input alone.
- **7** The state after applying the oracle is $(-|00\rangle + |01\rangle + |10\rangle |11\rangle) / 2 = (-1)|+-\rangle$, precisely the same as before, aside from a global phase of -1.

Looking at these two examples, we got the exact same output state, except the signs are all flipped. This means that if we multiplied on of the state vectors by a -1, they would both be the same. Multiplying an entire vector by a constant can be referred to as adding a global phase. Since global phases cannot be observed through any measurements, this means that we got *exactly* the same information out from applying the id and not oracles. We learned nothing about whether we applied id or not, and if we were able to compare the vectors right here we would only know that we applied a balanced oracle.

Now let's see for comparison what the register looks like after we apply an oracle representing a *constant* function.

EXAMPLE 3: THE ZERO ORACLE

Once more with feeling, let's use Python to break down how an oracle representing the constant function zero works. We want to to this with the zero oracle here so that we can show what goes differently when we apply an oracle representing a constant function. This oracle is especially easy to apply, since it consists of applying no instructions at all.

Table 7.12. Representing the one-bit function zero as two-qubit oracles

Function Name	Function	Output of Oracle	Q# Operation
zero	f(x)=0	$ x\rangle y\oplus 0\rangle = x\rangle y\rangle$	(empty)

Listing 7.14. Computing how the zero oracle transforms its input state.

```
>>> control_state = H * qt.basis(2, 0)
>>> target_state = H * qt.basis(2, 1)
>>> register state = qt.tensor(control state, target state)
>>> X = qt.sigmax()
>>> oracle = qt.tensor(I, I)
>>> oracle
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[1. 0. 0. 0.]
                                                                           a
[0. 1. 0. 0.]
[0. 0. 1. 0.]
[0. 0. 0. 1.]]
>>> register_state = oracle * register_state
>>> register state
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[ 0.5]
 [-0.5]
 [ 0.5]
[-0.5]]
```

- ① Doing nothing on the control qubit and nothing on the target qubit can be simulated by doing nothing on the entire register; thus, we get out the two-qubit identity matrix $1 \otimes 1$ for the zero oracle.
- 2) Here, we see our first difference from the previous listing. Before, we used the id oracle and got that the output state was $(|00\rangle |01\rangle |10\rangle + |11\rangle) / 2$. Using the zero oracle, we see that the output state is $(|00\rangle |01\rangle + |10\rangle |11\rangle) / 2$. These two states differ by more than a global phase, as there's no scalar that we can multiply by to turn the id output into the zero output.

Here now where the minus signs are on our state vector are different. There is no number we can multiply the entire vector by to change it into either [[0.5], [-0.5], [-0.5], [-0.5], [0.5], [0.5], [0.5], [-0.5]]. To see how this difference leads to us being able to tell for sure whether we had a balanced or constant oracle, let's continue to the next step of the Deutsch–Jozsa algorithm.

Exercise: Try out the one oracle

See if you can use the python tricks we have used above to work out how the state of the target and control qubits change when the one oracle is applied.

7.9 STEPS 3 AND 4. UNDO THE PREPARATION ON THE TARGET OUBIT AND MEASURE.

At this point, it is much easier to make sense of the output if we undo the steps we used to prepare $|+-\rangle$, so that everything is back in the computational basis ($|00\rangle \dots |11\rangle$). To review, has the state vectors for all four oracles (three of which we worked out above).

Table 7.13. The vectors representing the register after applying the various oracles

Function Name	State of the register after applying the oracle
zero	[[0.5], [-0.5], [0.5], [-0.5]]
one	[[-0.5], [0.5], [-0.5], [0.5]]
id	[[0.5], [-0.5], [-0.5], [0.5]]
not	[[-0.5], [0.5], [0.5], [-0.5]]

Now we want to undo our preparation steps on the target qubit.

Why do we do "unprepare" the target qubit?

Recall from Chapter 6 that we need to reset our qubits to the $|0\rangle$ state before returning them to the target machine. At this point, our target qubit is always in the $|-\rangle$ state, **no matter which oracle we use**. This means that after applying the oracle, we know exactly how to put it back to $|0\rangle$. As we saw in Chapter 6, this helps us avoid an additional measurement, which can be expensive on some quantum devices.

Note that we can safely return the target qubit to $|-\rangle$ without affecting the results when we measure the control qubit, as the oracle call is the only two-qubit operation in the Deutsch–Jozsa algorithm. As we saw in Chapter 4, doing single-qubit operations on one qubit can't on their own affect results on another qubit; otherwise, we would be able to send information faster than light!

Let's try this out by undoing the preparation on the register from the oracle representing the id function.

Listing 7.15. Putting the oracle output back in the computational basis for the oracle

```
representing id.
>>> I = qt.qeye(2)
>>> register_state_id = qt.cnot(2,0,1) *
... (qt.tensor(H * qt.basis(2, 0), H * (qt.sigmax() * qt.basis(2, 0))))
>>> register_state_id = qt.tensor(I, H) * register_state_id
>>> register_state_id
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Oobi data =
[[ 0.
 [ 0.70710678]
 [ 0.
 [-0.70710678]]
>>> register state id = qt.tensor(I, qt.sigmax()) * register state id
>>> register_state_id
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[ 0.70710678]
 [ 0.
 [-0.70710678]
             11
>>> qt.tensor(H * qt.basis(2, 1), qt.basis(2, 0)) == register_state_id
True
```

```
>>> register_state_id = qt.tensor(H, I) * register_state_id
>>> register_state_id
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[0.]
[0.]
[1.]
[0.]]
```

- 1 It's helpful to define a shorthand for the identity matrix 1, which we use to represent what happens to a qubit when we don't apply any instruction to it. In QuTiP, we can get the identity matrix with the qeye function (the name stems from that the identity matrix is often written *I*, which is then pronounced as "eye").
- 2 Here we reproduce the register for the oracle representing the id function, just after the oracle was applied.
- 3 Since we're transforming a two-qubit state now, we need to say what happens to each qubit to get our matrix out. We can do this using the tensor function again.
- **4** The output's now much easier to read: the register is in the state ($|01\rangle |11\rangle$) / $\sqrt{2}$.
- **5** In our Q# program, we used the X instruction to return the target qubit to |0) before releasing it. We can simulate that by using the *X* operator, implemented by the QuTiP function sigmax().
- **6** Since the X instruction flips its argument, applying the X matrix gives us the state ($|00\rangle |10\rangle$) / $\sqrt{2}$.
- **7** We can use QuTiP to confirm that another way of writing ($|00\rangle |10\rangle$) / √2 is $(H|1\rangle) ⊗ |0\rangle = |-0\rangle$.
- ③In our Q# program, we used the MResetX operation from the Q# standard libraries to measure in the X basis. An X-basis measurement returns Zero when its argument is in $|+\rangle$ and returns a One result when its argument is in $|-\rangle$. Thus, we can simulate the MResetX operation by applying H and then measuring in the Z-basis.
- Doing so, we read that the state of our register immediately before measuring has a 1 in row 2 (10 in binary), so we conclude that the state of our register is |10). Measuring the first qubit thus will give us a One with certainty.

Looking at the final vector in , we can see that that represents the state |10). If we were to measure the control qubit from that state, we would get the classical bit one 100% of the time. In Algorithm.qs, we returned to the user that an oracle is balanced if we measured one on the control qubit, so we correctly conclude that id is a balanced oracle! The fact that we will measure one on the control bit every time is really cool.

QUANTUM ALGORITHMS NEED NOT BE RANDOM

Though some quantum algorithms can be random, like the QRNG example from Chapter 2 or Morgana's and Lancelot's game from Chapter 6, they don't need to be. In fact, the Deutsch–Jozsa algorithm is deterministic: we get the same answer every time we run it.

Table 7.14. The vectors representing the state of the register after applying the various oracles.

Function	Register state just before	Result of measuring control qubit along
Name	measurement	Z

Function Name	Register state just before measurement	Result of measuring control qubit along ${\cal Z}$
zero	[[1], [0], [0], [0]]	Zero
one	[[-1], [0], [0], [0]]	Zero
id	[[0], [0], [1], [0]]	One
not	[[0], [0], [-1], [0]]	One

From these examples, we have made some important observations:

 Applying an oracle to a control and target qubit can affect the state of the control qubit.

7.10 REFLECTING BACK

Phew, we've taken a couple pretty big steps here!

- We've used classical reversible functions to model Merlin's strategy in a way that we can write it down as a quantum oracle.
- We've used Q# and the Quantum Development Kit to implement the Deutsch– Jozsa algorithm and test that we can learn Merlin's strategy with a single oracle call, and
- We've used Python and QuTiP to model the Deutsch–Jozsa algorithm and see what happens to Nimue's qubits when she gives them to Merlin.

At this point, it's helpful to reflect back on what we learned from taking a splash in Nimue's quantum lake, as the techniques we used in this Chapter will be helpful throughout the rest of the book.

7.11 SHOES AND SOCKS: APPLYING AND UNDOING QUANTUM OPERATIONS

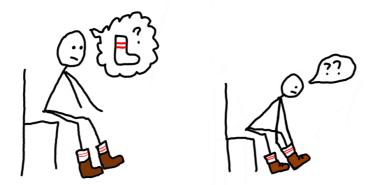
The first pattern that's helpful to reflect on is one that you might have noticed in Algorithm.qs. Let's take another look at the order in which operations were applied to the target qubit:

Listing 7.16. A sequence of instructions from the Deutsch-Jozsa algorithm for

the target qubit. // ... X(target); H(target); oracle(control, target); H(target); X(target); // ...

One way to think of this sequence is that the X(target); H(target); instructions prepare target in the |-> state, while the H(target); X(target); instructions "unprepare" |->, returning target back to the |0> state. We have to reverse the order due to what's often called the "shoes and socks" principle. If you want to put on shoes and socks, you'll have better results if you put on your socks first, but if you want to take them off, you need to take your shoes off first.

Figure 7.15. You can't take your socks off before your shoes.



The Q# language makes it easier to do shoes-and-socks kind of transformations of your code, using a feature called *functors*. Functors allow us to easily describe new variants of an operation that we have already defined. Let's jump right into an example and introduce a new operation, PrepareTargetQubit, that encapsulates the X(target); H(target); sequence:

```
Listing 7.17. Separating out the state preparation from .

operation PrepareTargetQubit(target : Qubit) : Unit is Adj {
    X(target);
    H(target);
}
```

• By writing is Adj as a part of the signature, we tell the Q# compiler to automatically compute the inverse operation — that is, the adjoint — of this operation.

We can then call the inverse operation generated by the compiler using Adjoint, one of the two functors provided by Q# (we'll see the other one in Chapter 8):

Listing 7.18. Using the Adjoint keyword to apply instructions following shoes-and-socks.

```
PrepareTargetQubit(target);
oracle(control, target);
Adjoint PrepareTargetQubit(target);
1
```

Adjoint PrepareTargetQubit applies the Adjoint functor to PrepareTargetQubit, giving back an operation that "undoes" PrepareTargetQubit. Following shoes-and-socks sorts of thinking, this new operation works by first calling Adjoint H(target); and then Adjoint X(target);.

SELF-ADJOINT OPERATIONS

In this case, X and Adjoint X are the same operation, since flipping a bit and then flipping it again always gets you back to where you started. Put differently, X undoes X. Similarly, Adjoint H is the same as H, so the above snippet gives us the sequence H(target); X(target);. We say that the instructions X and H are self-adjoint.

Not all operations are their own adjoints, though! For instance, Adjoint Rz(theta, $_$) is the same operation as Rz(-theta, $_$).

In more practical terms, the Adjoint functor on the operation U is the same as the operation that reverses or undoes the effects of U. The name "adjoint" refers to that the conjugate transpose U^+ of a unitary matrix U is called the adjoint of U. The Adjoint keyword in Q# guarantees that if an operation U is described by the unitary U, then if Adjoint U exists, it is described by U^+ .

The pattern of performing an instruction is so commonly used that the Q# standard libraries provides the ApplyWith operation to express this pattern of doing and then undoing an operation.

THE Q# STANDARD LIBRARY

The ApplyWith operation is provided by the Microsoft.Quantum.Canon namespace in the Q# standard library. Much like the standard library in other languages, the Q# standard library provides many of the basic tools you'll need for writing programs in Q#. As you go through the rest of the book, you'll see lots of ways that the Q# standard library can help make your life as a quantum developer easier.

Using ApplyWith and partial application, we can rewrite the using block of CheckIfOracleIsBalanced in a compact way:

Listing 7.19. Using ApplyWith and partial application to help with shoes-and-socks

ordering.

```
H(control);
ApplyWith(PrepareTargetQubit, oracle(control, _), target);
set result = MResetX(control);
```

The ApplyWith operation in the sample above will automatically apply the adjoint of PrepareTargetQubit after oracle(control, _) is done. Note that the _ is used for partially applying the oracle to the control qubit.

Let's expand step by step to see how it all works. The call to ApplyWith applies its first argument, then applies its second argument, then the adjoint of its first argument all to the qubit supplied in the last argument:

Listing 7.20. Expanding ApplyWith in .

```
H(control);
PrepareTargetQubit(target);
(oracle(control, _))(target);
Adjoint PrepareTargetQubit(target);
set result = MResetX(control);
```

The partial application on line 3 can then be replaced by substituting target in for the :

Listing 7.21. Resolving partial application in .

```
H(control);
PrepareTargetQubit(target);
oracle(control, target);
Adjoint PrepareTargetQubit(target);
set result = MResetX(control);
```

Using operations like ApplyWith is helpful for reusing common patterns in quantum programming, and in particular for making sure we don't forget to take an Adjoint in a large quantum program.

Q# also provides another way of representing the shoes-and-socks pattern using blocks of statements instead of passing around operations. For example, you could also write using the within and apply keywords instead, as we've done in .

Listing 7.22. Using within and apply to help with shoes-and-socks ordering.

```
H(control);
within {
    PrepareTargetQubit(target);
} apply {
    oracle(control, target);
}
set result = MResetX(control);
```

Both forms can be useful in different contexts, so feel free to use whichever works best for you!

7.12 USING HADAMARD INSTRUCTIONS TO FLIP CONTROL AND TARGET

One way that we can use the "shoes-and-socks" kinds of thinking from the previous section is to change which qubits play the role of a control and a target in instructions like CNOT.

In understanding how this works, it's important to keep in mind with quantum computing is that quantum instructions transform the entire state of the registers that they act upon. In cases like the Deutsch-Jozsa algorithm, this means that the control qubit can be affected by applying gates to the control and target qubits together — not

just the target qubit. This is an example of a more general pattern, that the control and target of a CNOT operation swap roles when we apply a CNOT instruction in the X basis instead of the Z (computational) basis.

To see this, let's look at the unitary operator (the quantum analogue to classical truth tables) for what happens if we use H instructions to transform to the X basis, apply a CNOT instruction, and then use more H instructions to go back to the Z basis:

Listing 7.23. Using QuTiP to check that H flips between the control and target of

```
a CNOT.
>>> import qutip as qt
>>> H = qt.hadamard transform()
>>> HH = qt.tensor(H, H)
>>> HH
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[ 0.5 0.5 0.5 0.5]
 [ 0.5 -0.5 0.5 -0.5]
 [ 0.5 0.5 -0.5 -0.5]
 [ 0.5 -0.5 -0.5 0.5]]
>>> HH * qt.cnot(2, 0, 1) * HH
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[1. 0. 0. 0.]
 [0. 0. 0. 1.]
 [0. 0. 1. 0.]
 [0. 1. 0. 0.]]
>>> qt.cnot(2, 1, 0)
Quantum object: dims = [[2, 2], [2, 2]], shape = (4, 4), type = oper, isherm = True
Qobj data =
[[1. 0. 0. 0.]
                                                                       6
 [0. 0. 0. 1.]
 [0. 0. 1. 0.]
[0. 1. 0. 0.]]
```

- ① It's helpful to define a shorthand for the unitary operator $H \otimes H$ that simulates the sequence of instructions H(control); H(target);.
- **2** Looking at the unitary operator $H \otimes H$, we see that $|00\rangle$ is transformed to $(|00\rangle + |01\rangle + |10\rangle + |11\rangle) / 2$, a uniform superposition over all four computational basis states.
- 3 This line gives us the unitary operator representing H(control); H(target); CNOT(control, target); H(control); H(target);. We can think of this sequence of instructions as applying a CNOT in the *X* basis instead of the *Z* basis.
- The unitary operator for this sequence looks a bit like CNOT, but with some of the rows flipped around... what happened?
- **⑤**To try and figure out what applying H instructions to each qubit did to the CNOT instruction, let's look at the unitary operator for CNOT(target, control).
- **6** Reversing the role of the control and target qubits in a call to the CNOT instruction gives us exactly the same unitary operator as using H instructions to apply a CNOT instruction in the *X*

basis. Figure 7.16. Changing which qubits are the "control" and "target" of a CNOT instruction using Hadamard instructions.

Let's look at two different programs (written as circuits) that each use a CNOT instruction.

As we saw in Chapter 4, the black dot indicates the control qubit for the CNOT instruction.

H

H

Similarly, the ⊕ symbol indicates the target for a CNOT instruction.

In the second program, we don't apply any H instructions — just a CNOT instruction with the control and target swapped.

These two programs are completely equivalent. There is no way to tell which sequence of instructions has been applied, as both programs have exactly the same effect.

From the above calculation, we can conclude that CNOT(target, control) does precisely the same thing as H(control); H(target); CNOT(control, target); H(control); H(target);. In the same way that the H(target) finished bases, H(target) instructions can flip between using a qubit as a control and as a target.

7.13 PHASE KICKBACK

With these techniques in mind, we're now equipped to explore what makes the Deutsch-Jozsa algorithm do its thing, namely a quantum programming technique called *phase kickback*. This technique is what let us write the CheckIfOracleIsBalanced operation so that it works for several different oracles, while revealing only the one bit we wanted to know (whether Merlin was acting as a good mentor or not).

To see how the Deutsch–Jozsa algorithm uses phase kickback to work in *general*, let's go back to our three ways of thinking, and use math to predict what happens when we call any oracle. Recall that we defined the oracle U_f that we constructed from each classical function f such that, for all classical bits x and y, $U_f|xy\rangle = |x\rangle |f(x) \oplus y\rangle$.

Above, we used x and y to represent classical bits that label two-qubit states; this is another example of using the computational basis to reason about how quantum programs behave.

Let's start the same way as we did in our QuTiP programs above, by expanding the input state $|+-\rangle = |+\rangle \otimes |-\rangle$ in the computational basis. Starting by expanding the state of the control qubit, we have that $|+-\rangle = |+\rangle \otimes |-\rangle = (|0\rangle + |1\rangle) / \sqrt{2} \otimes |-\rangle = (|0-\rangle + |1-\rangle) / \sqrt{2}$. As before, we can check our math using QuTiP (see).

Listing 7.24. Using QuTiP to check that $(|0-\rangle + |1-\rangle) / \sqrt{2} = |+-\rangle$.

```
>>> import qutip as qt
>>> from numpy import sqrt
>>> H = qt.hadamard transform()
>>> ket 0 = qt.basis(2, 0)
>>> ket_1 = qt.basis(2, 1)
>>> ket plus = H * ket 0
>>> ket minus = H * ket 1
>>> qt.tensor(ket_plus, ket_minus)
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[ 0.5]
 [-0.5]
 [ 0.5]
 [-0.5]]
>>> (
        qt.tensor(ket_0, ket_minus) +
        qt.tensor(ket 1, ket minus)
...) / sqrt(2)
Quantum object: dims = [[2, 2], [1, 1]], shape = (4, 1), type = ket
Qobj data =
[[0.5]]
 [-0.5]
 [ 0.5]
[-0.5]]
```

- 1 We'll start by writing out some useful shorthand notation.
- **2**Both vectors are the same, telling us that $(|0-\rangle + |1-\rangle) / \sqrt{2}$ is another way of writing out $|+-\rangle$.

Next, as we saw in Chapter 2, we can use linearity to predict how U_f transforms this input state.

THE MATRIX REVISITED

We've also implicitly used linearity earlier in this section as well, when we used matrices to model how the Deutsch-Jozsa algorithm works. As described in Chapter 2, matrices are one way of writing down linear functions.

Since U_f is a unitary matrix, we know that for *any* states $|\psi\rangle$ and $|\varphi\rangle$, and for *any* numbers α and β , $U_f(\alpha|\psi\rangle + \beta|\phi\rangle) = \alpha U_f |\psi\rangle + \beta U_f |\phi\rangle$. Using this property with the computational basis, we have that in the same way that $|-\rangle$ and $(|0-\rangle + |1-\rangle) / \sqrt{2}$ are the same state.

TIP Using our shorthand for multi-qubit states, $|+-\rangle = |+\rangle \otimes |-\rangle$, $|0-\rangle = |0\rangle \otimes |-\rangle$, and $|1-\rangle = |1\rangle \otimes |-\rangle$.

Figure 7.17. Applying linearity to understand how our oracle transforms the input state.

$$\frac{U_f \ket{+-} = (\ket{0-} + \ket{1-})/\sqrt{2}, \text{we can substitute that in here.}}{U_f \ket{+-} = U_f (\ket{0-} + \ket{1-})/\sqrt{2}}$$
 We want to know what happens when we apply our oracle to the input state that we've prepared, so we start by writing that down.
$$U_f \ket{0-} + U_f \ket{1-}/\sqrt{2}$$
 Using linearity, we can apply U_f to each term above (that is, to $\ket{0-}$) and $\ket{1-2}$.

Written like this, it's not immediately clear what advantage we've obtained by applying U_f to $|0-\rangle$ and $|1-\rangle$. Let's look at how the oracle operation applies to the first term by factoring out the control (first) qubit, so that we can consider the effect on the target qubit.

Doing so, we'll once again make use of linearity to understand how U_f works by passing one state at a time to our oracle. As we learned in Chapter 2, linearity is a very powerful tool that lets us break down even quite complicated quantum algorithms into pieces that we can understand and analyze more easily. In this case, we can understand how U_f acts on $|0-\rangle$ by using linearity (that is, by breaking $|0-\rangle$ down into a superposition between $|00\rangle$ and $|01\rangle$):

We start by using that
$$|-\rangle=(|0\rangle-|1\rangle)/\sqrt{2}$$
.
$$U_f\,|0-\rangle=U_f(|00\rangle-|01\rangle)/\sqrt{2}$$
 We can apply U_f to each term by linearity.
$$=(U_f\,|00\rangle-U_f\,|01\rangle)/\sqrt{2}$$
 The definition of an oracle tells us how U_f transforms states in the computational basis.
$$=(|0(0\oplus f(0))\rangle-|0(1\oplus f(0))\rangle)/\sqrt{2}$$
 XORing something with 0 does nothing.
$$=(|0f(0)\rangle-|0(1\oplus f(0))\rangle)/\sqrt{2}$$
 XORing something with 1 is the same as NOT.
$$=(|0f(0)\rangle-|0(\neg f(0))\rangle)/\sqrt{2}$$
 Since the control qubit is in $|0\rangle$ in both terms, we can factor it out.

For instance, if we're considering the zero function, then f(0) = 0.

©Manning Publications Co. To comment go to liveBook

 $= |0\rangle \otimes (|f(0)\rangle - |\neg f(0)\rangle) /\sqrt{2}$

Thus,
$$|f(0)\rangle = |0\rangle$$
 and $|\neg f(0)\rangle = |1\rangle$, so $U_f |0-\rangle = |0-\rangle$.

On the other hand, if f(0) = 1, then $U_f |0-\rangle = |0\rangle \otimes (|1\rangle - |0\rangle) / \sqrt{2} = -|0-\rangle$. That is, U_f flips the sign of $|0-\rangle$.

TIP
$$(|1\rangle - |0\rangle)/\sqrt{2}$$
 can also be written as $-|-\rangle$, or as $X|-\rangle$.

We can note then that U_f either rotates the target qubit by X or not depending on the value of f(0):

$$U_f |0-\rangle = (-1)^{f(0)} |0-\rangle / \sqrt{2}$$

We can use the exact same argument that we used above to understand what U_f does to $|0-\rangle$ to understand what U_f does in the case the that the control qubit is in the $|1\rangle$ state instead. Doing so, we get a phase of $(-1)^{f(1)}$ instead of $(-1)^{f(0)}$, so that $U_f |1-\rangle = (-1)^{f(1)} |1-\rangle$.

Using linearity again to combine the terms for the two states of the control qubit, we now know how U_f transforms the state of both qubits when the control qubit is in $|+\rangle$:

$$U_f |+-\rangle = \frac{1}{\sqrt{2}} ((-1)^{f(0)} |0-\rangle + (-1)^{f(1)} |1-\rangle)$$

The last step is to note that, as we saw in Chapter 4, we cannot observe *global* phases. Thus, we can factor out $(-1)^{f(0)}$ to express the output state in terms of $f(0) \oplus f(1)$, the question we were interested in to start, as shown in .

Figure 7.18. Working out the last couple steps to the Deutsch-Jozsa algorithm

By using linearity, and that
$$|+\rangle = (|0^-\rangle + |1^-\rangle)/\sqrt{2}$$
, we can apply U_f to both of the cases we've already computed, $U_f|0^-\rangle$ in both cases, we get a phase that depends on the value of our irreversible classical function, evaluated at the label of the control qubit; e.g., $U_f|0^-\rangle = (-1)^{f(1)} |1-\rangle$)

To finish our exploration of the Deutsch-Jozsa algorithm, we need to know what the easier to see that it's a global phase.

Next, we can factor out $f(0)$, so that it's easier to see that it's a global phase.

Doing so leaves us with a phase on the $|1^-\rangle$ part of the state that depends on $f(0) \oplus f(1)$, which is 0 for constant functions and is 1 for balanced functions.

$$= \frac{1}{\sqrt{2}} (-1)^{f(0)} (|0-\rangle + (-1)^{f(0)} \oplus f(1) |1-\rangle)$$
Finally, we notice that the target qubit is in the same state regardless of the state of the control qubit. Thus, we can factor it out, and safely measure or reset the target qubit without affecting the control.

If $f(0) \oplus f(1) = 0$ (constant), then the output state is $|+-\rangle$, but

if $f^{(0)} \oplus f^{(1)} = 1$ (balanced), then the output state is $|--\rangle$. With one call to U_f , we learn whether f was constant or balanced even though we do not learn what f(x) is for any particular input x.

One way to think of what happened when we applied U_f with the input qubit in the $|+\rangle$ state is that the state of the input qubit represents the question we are asking about f. If we ask the question $|0\rangle$, we get the answer f(0), while if we ask $|1\rangle$, we get the answer f(1). The question $|+\rangle$ then tells us about $f(0) \oplus f(1)$ without telling us about either f(0) or f(1) alone.

When we ask questions in superposition like this, however, the role of "input" and "output" isn't as immediately clear as it is classically. In particular, the $|0\rangle$ and $|1\rangle$ inputs both cause the output qubit to flip, while the $|+\rangle$ input causes the *input* qubit to flip, provided we start the output qubit in $|-\rangle$ state. In general, two-qubit operations like U_f transform the entire space of the qubits that it acts upon — our division into inputs and outputs is one of how we interpret the action of U_f .

PHASE KICKBACK

The fact that the state of the input qubit changed based on transformations defined on the output qubit is an example of a quantum effect known as phase kickback. In the next two chapters, we'll leverage phase kickback in order to explore new algorithms, such as those used in quantum sensing and quantum chemistry simulations.

Deep Dive: Extending Deutsch-Jozsa

While we only considered functions with one bit inputs here, the Deutsch-Jozsa algorithm will only ever need one query for any sized input/output of our function!

To encode a two-qubit function $f(x_0, x_1)$, we can introduce a three-qubit oracle $U_f \mid x_0 x_1 y \rangle = \mid x_0 x_1 \rangle \otimes f(x_0, x_1) \oplus y$. For example, consider $f(x_0, x_1) = x_0 \oplus x_1$. This function is balanced since f(0, 0) = f(1, 1) = 0 but f(0, 1) = f(1, 0) = 1. When we apply U_f to the three-qubit state $|+-\rangle = (|00\rangle + |01\rangle + |10\rangle + |11\rangle) \otimes |-\rangle$, we get $(|00\rangle - |01\rangle - |10\rangle + |11\rangle) \otimes |-\rangle = |---\rangle$. Using an X-basis measurement, we can tell this apart from a constant function like $f(x_0, x_1) = 0$, which will give us an output of $|+-\rangle$.

As long as we're promised that f is either constant or balanced, the same pattern holds no matter how many bits f takes as input: we can learn a single bit of data about how f behaves with a single call to U_f . Talk about O(1)! If you are not familiar with Big O notation, see $Grokking\ Algorithms$ by Aditya Bhargava.

7.14 SUMMARY

In this chapter you learned:

- to recognize quantum algorithms,
- to design *oracles* to represent classical functions within quantum algorithms,
- to predict the output of quantum programs by using QuTiP and linearity, and
- to recognize quantum programming techniques like phase kickback.

In the next chapter, we'll build on these skills by looking at how the **phase estimation algorithm** enables spin-off technologies like quantum sensors.

Quantum Sensing: It's Not Just A Phase

This chapter covers:

- Predict how quantum operations can be used to learn useful information about unknown operations
 with the quantum algorithm for phase kickback.
- · Create new types in Q#
- Run Q# code from a Python host program
- · Recognize important properties and behaviors of eigenstates and phase
- · Program controlled quantum operations in Q#

Figure 8.1. In this Chapter, we will be covering topics in both the Q# language and techniques for developing quantum algorithms.

Quantum Applications		
Libraries / Techniques		
Q# Language		
Intrinsic Operations		
Quantum Hardware	Simulators	

8.1 PHASE ESTIMATION: LEVERAGING USEFUL PROPERTIES OF OUBITS FOR MEASUREMENT

Throughout the book so far, you've seen that games can be a really helpful way to learn quantum computing concepts. In the previous Chapter, for instance, Lady Nimue's game with Merlin let you explore your first quantum algorithm, the Deutsch–Jozsa algorithm. In this Chapter you'll use another game to explore how to learn the phases of quantum states using *phase kickback*, the quantum development technique used by Deutsch–Jozsa and many other quantum algorithms.

For this Chapter's game, let's go back and see what Lancelot has been up to. While Nimue and Merlin were deciding the fate of kings, you find Lancelot and the court jester Dagonet playing a guessing game where he has "borrowed" some of Nimue's quantum tools. Since they've had a while to play, however, Dagonet has gotten a bit bored and wants to start incorporating quantum computing to make their game a bit more interesting.

8.2 PART AND PARTIAL APPLICATION

For Dagonet's new game, rather than having Lancelot guess a number, Dagonet has Lancelot guess what a quantum operation does to a single qubit by letting Lancelot call it with different inputs. Given all single qubit operations are rotations, this work really well for the game. Dagonet can pick a rotation angle about a particular axis, then Lancelot gets to input a number to Dagonet's operation which will change how to scale the rotation that Dagonet applies. What axis Dagonet picks doesn't really matter, as the game is to guess the rotation angle. For convenience here, Dagonet's rotations will always be around the Z axis. Finally, Lancelot can measure the qubit and use his measurement to try and guess what Dagonet's original rotation angle was. See for a flowchart of these steps.

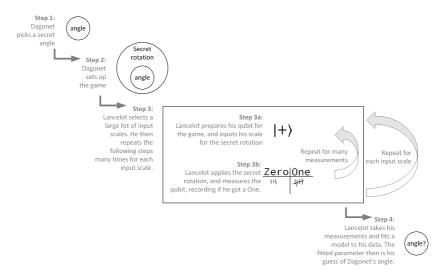


Figure 8.2. Playing Dagonet's guessing game

8.2.1 Steps for Dagonet and Lancelot's game

- 1. Dagonet picks a secret angle for a single qubit rotation operation
- 2. Dagonet prepares an operation for Lancelot to use that hides their secret angle, and allows Lancelot one additional input of a number (we'll call it a scale) that will get multiplied with the secret angle to give the total rotation angle of the operation.
- 3. Lancelot's best strategy for the game will be to select many scale values and estimate the probability of measuring One for each value. To do this, he will need to do steps a. and b. many times for each of the many scale values.
 - Prepare the |+> state and input the scale value in Dagonet's rotation. He uses the |+> state because he knows Dagonet is rotating around the Z axis, and for this state these rotations will result in a local phase change he can measure.
 - After preparing each |+> state, Lancelot can rotate it with the secret operation, measure the qubit, and record the measurement.
- 4. Lancelot will then have data relating his scale factor and the probability he measured a 1 for that scale factor. He can then fit this data in his head, and get Dagonet's angle from the fitted parameters (he *is* the greatest knight in the land). We can use Python to help us do the same!

Note that this is a game, as there is no way for Lancelot to directly measure this rotation with just a single measurement. If he could, it would violate the No-Cloning theorem, and he would transcend the laws of physics himself. As a Knight of the Round Table, Lancelot is bound not only by duty and honor, but also by the laws of physics, so he will have to play Dagonet's game by the rules.

DEEP DIVE: Learning the axis with Hamiltonian learning

In this book, we've focused on the case where Dagonet's rotation axis is known but we need to learn his angle. This case corresponds to a common problem in physics, where one is tasked with learning the Larmor precession of a qubit in a magnetic field. Learning Larmor precessions isn't just useful in building qubits, but allows for detecting very small magnetic fields and for building very precise sensors.

More generally, though, you can learn much more than a single rotation angle. The case in which the axis is unknown as well is an example of a general kind of problem called Hamiltonian learning, a rich area of research in quantum computing. In Hamiltonian learning, one reconstructs a physical model for a qubit or register of qubits using a game very similar to the one you'll explore throughout the rest of this Chapter.

Let's jump into prototyping this game in Q#! It will be helpful to have access to different parts of the Q# standard libraries, so you can start by adding the open statements below to the top of your Q# file, operations.qs. We show the open statements used in this example in .

Listing 8.1. Opening Q# namespaces needed for Lancelot's and Dagonet's game.

```
namespace PhaseEstimation {
    open Microsoft.Quantum.Intrinsic;
    open Microsoft.Quantum.Convert as Convert;
    open Microsoft.Quantum.Measurement as Meas;
    open Microsoft.Quantum.Arrays as Arrays;

// ...
```

- 1 All open statements in a Q# file come right after your namespace declaration.
- ②As before, opening Microsoft.Quantum.Intrinsic gives you access to all the fundamental instructions (e.g.: R1, Rz, X, and so forth) that you can send to a quantum device.
- 3 You can also give an alias to namespaces when you open them, similar to how you can alias Python packages and modules when you import them. Here, for instance, we've abbreviated Microsoft.Quantum.Convert so that we can later use type conversion functions in that namespace by prefixing them with Convert..
- 4 Similarly, you can make the MResetZ operation that you saw in previous Chapters available as Meas. MResetZ to document where the operation came from.
- **5**The last namespace you'll need to open in this Chapter is the Microsoft.Quantum.Arrays namespace, which provides a number of useful functions and operations for working with arrays.

In , you can see an example of the quantum operation Dagonet would need to implement the rotation he and Lancelot will play with.

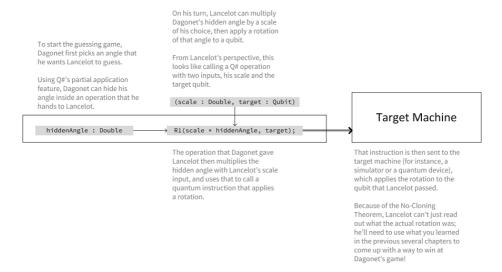
Listing 8.2. A quantum operation Dagonet could use to setup his part of the game.

- 1 To play the guessing game, we'll need a quantum operation that takes two classical arguments; one that Dagonet gets to pass, and one that Lancelot gets to pass.
- ②Like other rotations, our new rotation operation returns Unit (the type of the empty tuple ()), indicating that there's no meaningful output from the operation. The is Adj + Ctl part of the signature indicates that this operation supports the Adjoint functor that you first saw in Chapter 6, as well as the Controlled functor that you'll see later in this chapter. Don't worry too much about this part for now, though.
- 3 For the actual body of the operation, you can find the angle to rotate by by multiplying Dagonet's hidden angle angle with Lancelot's scale factor scale. The rotation operation R1 here is almost identical to the Rz operation that you've seen a few times so far now. The difference between R1 and Rz will become important later, when you add the Controlled functor.

IMPORTANT

When writing Q# in its own file (that is, not from a Jupyter Notebook), all operations and functions must be defined inside of a namespace. This helps keep your code organized, and makes it harder to conflict with what code may be written in the various libraries that you use in your quantum application. For brevity, we often won't show the namespace declarations, but the full code can always be found in the samples repository for this book: https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp

Figure 8.3. Playing Dagonet's guessing game with partial application



DEEP DIVE: Why not just measure the angle?

It may seem like Lancelot has to go through a lot of hoops to guess Dagonet's hidden angle. After all, what besides duty and honor is there stopping Lancelot from passing a scale of 1.0 and then just reading out the angle from the phase applied to his qubit? It turns out that the No-Cloning Theorem strikes again in this case, telling us that Lancelot can never that he phase from a single measurement.

The easiest way to see this is to pretend for a moment that Lancelot could do that, then see what goes wrong. Suppose that Dagonet hides an angle of $\pi/2$, and that Lancelot prepares his qubit in

199

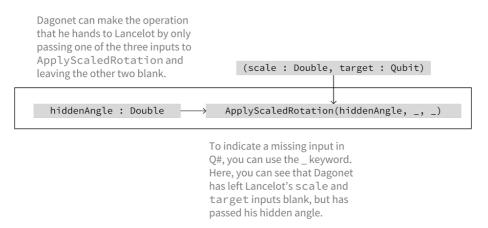
the $|+\rangle=\frac{1}{\sqrt{2}}(|0\rangle+|1\rangle)$ state. If Lancelot then passes 1 . 0 as his scale, his qubit ends up in the $\frac{1}{\sqrt{2}}(|0\rangle+i|1\rangle$ state. If Lancelot could directly measure the phase $e^{i\pi/2}=+i$ in order to guess Dagonet's angle from a single measurement, he could use that to prepare another copy of $\frac{1}{\sqrt{2}}(|0\rangle+i|1\rangle$ state, even though Lancelot didn't know the right basis to measure in. Indeed, Lancelot's magic measurement device should also work if Dagonet hides the angle π , in which case Lancelot winds up with a qubit in the $\frac{1}{\sqrt{2}}(|0\rangle-|1\rangle)$ state.

Put differently, if Lancelot could figure out Dagonet's angle by measuring phases directly, he could make copies of arbitrary states of the form $\frac{1}{\sqrt{2}}(|0\rangle+e^{i\phi}|1\rangle)$ for any angle Φ , without having to know Φ ahead of time. That rather badly violates the No-Cloning Theorem, so you can safely conclude that Lancelot will need to do a bit more work to win Dagonet's game.

Once you've defined an operation in this way, Dagonet can use the partial application feature of Q# that you first saw in the previous chapter to hide his input. Then Lancelot gets an operation that he can apply to his qubits, but not in a way he can directly see the angle he is trying to guess.

Using ApplyScaledRotation, Dagonet can easily make an operation for Lancelot to call. For instance, if Dagonet picks the angle 0.123, he can 'hide' it by giving Lancelot the operation ApplyScaledRotation(0.123, _, _). As with the examples of partial application that you've seen in Chapter 7, the _ indicates a slot for future inputs.

Figure 8.4. Partially applying ApplyScaledRotation to make an operation for Lancelot.



As shown in , since ApplyScaledRotation has type (Double, Double, Qubit) => Unit is Adj + Ctl, providing only the first input results in an operation of type (Double, Qubit) => Unit is Adj + Ctl. This means that Lancelot can provide an input of type Double, a qubit he wants to apply his operation to, and can use the Adjoint and functor that you saw in Chapter 6.

Now, the fact that we can see the value of the angle in the syntax above does *not* mean Lancelot can. Indeed, the only things that Lancelot can do with a partially-applied

operation or function are to call it, partially apply it further, or pass it to another function or operation. From Lancelot's perspective, ApplyScaledRotation(0.123, _, _) is a complete black box. Thanks to this partial application trick, he will just have an operation that takes his scale value and can be used to rotate a qubit.

We can make our lives as Q# developers a bit easier, though, by giving a name to the type of Lancelot's operation that's a bit easier to read than (Double, Qubit) => Unit is Adj + Ctl. In the next section, you can see how Q# lets us annotate the type signatures that you use to play Dagonet and Lancelot's guessing game.

8.3 USER-DEFINED TYPES

You have seen a bit already on how types play a role in Q#, particularly in the signatures for our functions and operations. You have also seen that both functions and operations are tuple-in, tuple-out. In this section you will see how you can build up your own types in Q# and why that might be handy.

In Q# (as well as many other languages) there are a number of types that are defined as a part of the language itself; types like Int, Qubit, and Result that we have seen already.

TIP

For a complete list of these basic types see the Q# language docs here: https://docs.microsoft.com/quantum/language/type-model#primitive-types

Building up from these basic types, you can make array types by adding a [] after the type. For example, in the game for this chapter you will likely need to input an array of doubles to represent Lancelot's multiple inputs to Dagonet's operation. You can use <code>Double[]</code> to indicate an array of <code>Double</code> values, see .

Listing 8.3. An example of defining an array type, here an array of Double`s of length

10.

let scales = new Double[10];

You can also define your own types in Q# with the newtype statement. This statement allows you to declare new *user-defined types* (often abbreviated "UDTs"). There are two main reasons why you may want to use UDTs.

8.3.1 Reasons you might define UDTs in Q#:

- convenience
- communicate intent

The first reason is a matter of *convenience*. Sometimes the type signature for your function or operation could get pretty long, so you can define your own type as a kind of shorthand. Another reason you may want to name your own type is so you can communicate *intent*. Say your operation takes a tuple of two <code>Double</code> values that represents a complex number. Defining a new type <code>Complex</code> could be useful to help

remind you and your teammates what that tuple represents.

Listing 8.4. How complex numbers are defined in the Q# runtime. namespace Microsoft.Quantum.Math { newtype Complex = (Real: Double, Imag: Double); 2

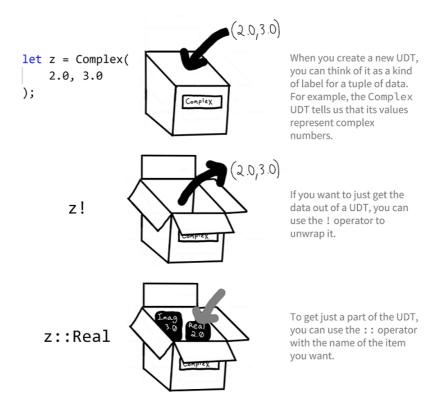
- ① The Quantum Development Kit provides a number of different functions, operations, and UDT to the various namespaces that make up the Q# libraries. For instance, complex numbers are implemented as a user-defined type in the Microsoft.Quantum.Math namespace; this lets you use this type by including the statement open Microsoft.Quantum.Math; in your quantum application.
- 2 The Complex type is defined as a tuple of two Double values, where the first item is named Real and the second is named Imag.

BACK TO THE SOURCE

The Quantum Development Kit is open source, so if you're curious you can always look up how various parts of the Q# language, runtime, compiler, and standard libraries all work. For example, the definition of the Complex user-defined type can be found in the src/Simulation/Intrinsic/Math/Types.qs file in Q# runtime repository at https://github.com/microsoft/qsharp-runtime/

As shown in , there's two ways to get the different items back out of a user-defined type.

Figure 8.5. Using the :: and the ! operator with UDTs.



You can either use named items together with the :: operator, or you can "unwrap" the user-defined type with the ! operator to get back to the original type wrapped by the UDT.

```
function TakesComplex(complex : Complex) : Unit {
   let realFromNamedItem = complex::Real;
   let (real, imag) = complex!;
}
```

- ① Since the Complex UDT is defined with a named item called Real, you can access that item as ::Real to get back the real part of our input.
- 2 Alternatively, since Complex was defined as wrapping a tuple of type (Double, Double), the unwrap operator! takes you back to the real and imaginary parts of complex without the UDT wrapper.

TIP Both ways of working with UDTs are useful in different cases, but most of the time we'll stick with using named items and the :: operator in this book.

Once a new UDT has been defined, it can also act as a way to instantiate a new instance of that type. For example, the Complex type will act as a function that will

create a new complex number with the input of a tuple of two <code>Double</code> values (see an example of this in). This is similar to Python where types are also functions that create instances of that type.

Listing 8.5. Creating a complex number with the user defined type Complex.

```
let imaginaryUnit = Complex(0.0, 1.0);
```

① Defining a user-defined type with newtype also defines a new function with the same name as that type that returns values of your new user-defined type. For instance, we can call Complex as a function with two Double inputs representing the real and imaginary parts of the new Complex value. Here, we've defined a Complex value representing 0 + 1.0 i (0+1j in Python notation), also known as the imaginary unit.

Exercise 8.1

In Chapter 4, you used Python type annotations to represent the concept of a *strategy* in the CHSH game. User-defined types in Q# can be used in a similar fashion. Give it a go by defining a new UDT for CHSH strategies and then use your new UDT to wrap the constant strategy from Chapter 4.

HINT: Your and Eve's parts of the strategy can each be represented as operations that take a Result and output a Result. That is, as operations of type Result \Rightarrow Result.

For the game in this chapter, we have defined a new UDT both so that we can label how we intend to use it, and as a convenient shorthand for the operation type that Lancelot gets as his part of the guessing game. In, you can see the definition of this new type, called ScalableOperation, as a tuple with one named input called Apply.

Listing 8.6. Setting up the quantum guessing game

```
newtype ScalableOperation = (
        Apply: ((Double, Qubit) => Unit is Adj + Ctl)
);

function HiddenRotation(hiddenAngle : Double) : ScalableOperation {
    return ScalableOperation(ApplyScaledRotation(hiddenAngle, _, _));
}
```

- 1 You can declare a new user-defined type with the newtype statement by giving a name for your new type, and defining the underlying type that your new type is based on.
- 2 You can give names to the various items in a user-defined type using the same syntax as to declare the signature for an operation or a function. Here, for example, your new UDT has a single item named Apply that allows for calling the operation wrapped by ScalableOperation.
- 3 Once defined, you can use your new UDT just like any other type. Here, you've defined a function that outputs values of type ScalableOperation.
- You can easily make output values by calling ScalableOperation with the operation to be wrapped in your new UDT. In this example, you can create new instances of ScalableOperation by using the same partial application of ApplyScaledRotation that you saw earlier in the Chapter.

WHAT A CAPITAL IDEA!

When you define inputs to functions and operations in Q#, those inputs have names that start with lowercase letters. In , however, the named item <code>Apply</code> in <code>ScalableOperation</code> starts with an initial upper-case letter. This is because the inputs to a function or operation only have meaning within that callable, while named items mean something more broadly; you can use the capitalization of inputs and named items to make it obvious where to look for the definitions of each.

The function <code>HiddenRotation</code> defined in helps us in implementing Lancelot and Dagonet's game, by giving us a way for Dagonet to hide his angle. Calling <code>HiddenRotation</code> with Dagonet's angle will return a new <code>ScalableOperation</code> that Lancelot can then call to gather the data he needs to make a guess at the hidden angle.

Figure 8.6. Playing Dagonet's guessing game with partial application and user-defined types



With some new types and a way for Dagonet to hide his angle, let's continue implementing the rest of the game! You have everything you need to go on to the next step: estimating the probability for each measurement that you make during Lancelot's and Dagonet's game. This is very similar to how you would estimate the probability of flipping a coin, see .

Figure 8.7. Lancelot's estimation is similar to esimating the outcome of a coin flip.

Suppose that someone hands you a coin, and you'd like to know the *bias* of that coin. That is, the probability that your shiny new coin lands heads.

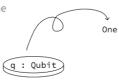


One way to do that is to flip your coin a bunch, and make a table of how many times it lands heads and how many times it lands tails.



If you get 3 heads and 5 tails, you could very reasonably estimate that the probability of getting heads is 3 / 8 = 37.5%. With more flips, you could get a more accurate estimate.

The same idea holds if you're trying to estimate the probability of getting a One from measuring a qubit.



Zero	0ne
11/	44

For example, suppose you prepare a qubit in the state $|\psi\rangle$ and measure it, repeating the whole process eight times. If you get five One results, you would estimate that $|\langle 1|\psi\rangle|^2 = 5/8 = 62.5\%$.

- ①To play the game, Lancelot needs to estimate the probability of measuring a |1⟩ at each given scale that he picks. Thus, his operation needs to take in a Double value representing what scale he picks to run the operation given to him by Dagonet.
- Next, Lancelot needs to pick a number of times to measure his qubit to get his estimate of the probability.
- 3 The last input Lancelot needs to take is a value of the ScalableOperation user-defined type that you declared earlier in the chapter. This input represents the operation that Dagonet gives Lancelot.
- **4** As output, Lancelot will want to return an estimated probability, so you can declare that as a <code>Double</code> output.
- (3) To keep track of the number of |1) outcomes that have been observed so far, you can define a mutable variable. As before, initializing the variable with the Integer value 0 sets the type of the mutable variable as well as its initial value
- **6** For each measurement, you'll need to allocate a qubit that's the target for Dagonet's operation.
- Next, you can use the within and apply keywords to use the shoes-and-socks kind of pattern that you first learned about in Chapter 7.
- (3) Since Lancelot and Dagonet agreed that the Z-axis should be the rotation axis for their game, Lancelot can prepare his target qubit in the |+) state so that Dagonet's rotation does something. Here, you can implement Lancelot's strategy by using the `H` operation to prepare a qubit in the |+) state.
- **9**Once you have prepared the input to Dagonet's operation, you can call it by using ::Apply to unwrap the ScalableOperation UDT.
- ①At this point, the within/apply block has made sure that Lancelot's qubit is back in the right axis. You can count how many times the final measurement returns a One result by adding either 1 or 0 to nones. Here, the ?| ternary operator (much like the if ... else operator in Python or the ?: operator in C, C++, and C#) provides a convenient way to increment nones.
- ①To get your final estimate as to the probability of measuring |1), you can take the fraction of how many times you got a one over how many times you measured. Here, Convert.IntAsDouble helps you to get the types right returning a floating-point number out.

AN OPERATION BY ANY OTHER NAME

You may have noticed that operations tend to be named using verbs, while functions tend to be named as nouns. This helps you to remember the distinction that you saw in Chapter 6, namely that a function **is** something while an operation **does** something. Being consistent with names can help you make sense of how a quantum program works, so Q# uses conventions like this throughout the language and libraries.

With this in place, you can now write out an operation that runs the whole game, and

returns everything Lancelot needs to guess Dagonet's hidden angle.

- 1 Dagonet can start off by making a new ScalableOperation value that hides his angle using the HiddenRotation function that you wrote earlier.
- 2 The ForEach operation in Microsoft.Quantum.Arrays (which you gave the shorthand name Arrays above) takes an operation and applies it to every element of the scales array.
- 3To get the operation we pass to ForEach, we use partial application lock down how many measurements Lancelot does at each different scale, and what the hidden operation he was given by Dagonet. This leaves an operation that only takes the scale as an input.
- When you pass scales as the second input to ForEach, each element of scales will be substituted into the partial application slot above.

FUNCTIONS AND OPERATIONS REDUX

It may seem funny that ForEach acts like map in Python and other languages, when Q# also has Microsoft.Quantum.Arrays.Mapped.The critical difference is that ForEach takes an operation, while Mapped takes a function.

For Lancelot to actually make sense of all the data he gets out of his Q# program, though, it might help to use some good old classical data science techniques. Since Python is great at that, running your new Rungame operation from a Python host program can be a great way for you to help Lancelot out.

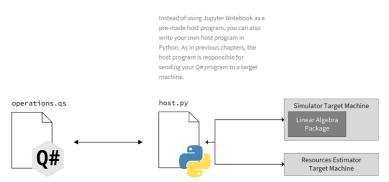
8.4 RUN, SNAKE, RUN: RUNNING Q# FROM PYTHON

In previous Chapters, you ran your Q# code in a Jupyter Notebook with a Q# kernel. In this Chapter we want to look at a different way to run Q# code: from Python. Calling Q# from Python can be helpful in different scenarios, especially if you want to preprocess data before using it in Q#, or if you want to visualize your resulting output from your quantum program.

Let's start actually writing the files that will implement Dagonet and Lancelot's game. To try out the Q# and Python interop, we will use a Python host program to run the Q#

program. This means we will have two files for our game, our operations.qs file and a host.py file that we will directly use to run the game. Let's dive into the host.py file to see how we can interact with Q# from Python.

Figure 8.8. Using host programs written in Python



All of the interoperable functionality we will need between Python and Q# is provided by the qsharp Python package.

TIP Reminder that there are full installation instructions for the qsharp Python package in Appendix A.

Once you have the qsharp package, you can import it just like any other Python package. Let's look at a small sample Python file where we can see this in action.

- 1 The qsharp Python package needs to be imported just like any other Python package.
- 2 You can use the qsharp.compile Python function to take a string containing Q# code and compile it for use in the Python file, the callable result of the compilation here being called prepare qubit.
- 3 Just like a regular Q# file, you need to include open statements to use the different parts of the Q# standard library.

- 4 The operation this string of Q# code describes simply prepares a qubit in the |0⟩ state, and uses DumpMachine to show what the target machine knows about that qubit. Given you are using the simulator, DumpMachine shows the values the simulator is using to record the state of the qubit.
- (5) When this Python script is run, we want to also use the callable defined as prepare_qubit. To do so, you can use the simulate method from the qsharp package which will run previously compiled Q# code snippets.

Let's try running the qsharp-interop.py script we just looked at in.

Listing 8.10. Running the previous example of a Python script containing Q# code.

TIP If you are running code from a Q# Jupyter Notebook the output of the Q# snippet will look different, see for an example.

From the output in , you can see that it did indeed prepare a $|0\rangle$ state, as the only term in the output that has a coefficient of 1.0 is the $|0\rangle$ state.

The qsharp Python package will also look for Q# operations or functions defined in *.qs files in the same directory as your Python program. In your case, as you proceed through the rest of this chapter, you'll be adding things to a Q# file called operations.qs. This is a pretty convenient way for you to start your host.py file for the game.

Listing 8.11. The beginning of our host.py file for simulating the phase estimation

```
import qsharp
from PhaseEstimation import RunGame, RunGameUsingControlledRotations

from typing import Any
import scipy.optimize as optimization
import numpy as np

BIGGEST_ANGLE = 2 * np.pi
```

- 1) First, you need to import the Q# Python package.
- The loaded <code>qsharp</code> package then allows you to import operations and functions from namespaces in Q# files in the same directory as <code>host.py</code>. Here, you import <code>RunGame</code> and <code>RunGameUsingControlledRotations</code> operations from operations.qs; this automatically creates Python objects representing each Q# operation that you import. You saw <code>RunGame</code> above, and will see <code>RunGameUsingControlledRotations</code> shortly.

3 The rest of the imports will help us with type hinting in Python, visualizing the results of your Q# simulation, and fitting measurement data to get Lancelot's final guess.

Now that we have imported and setup our Python file, let's write run game at scales: the function that will actually call the Q# operations.

Listing 8.12. The Python function that will call our Q# operations.

- Here the return type hint is set to Any which tells Python to not worry about type checking the return value of this function.
- 2 Here, Dagonet chooses the hidden angle that he wants Lancelot to guess.
- 3 The return for run_game_at_scales is conditioned on control, which allows you to choose between two simulations you will develop for this game. Don't worry about the control = True case for now, you'll see more about that later in the Chapter.
- 4 The Python objects created when qsharp importing these operations from Q# have a method called simulate, which takes the required arguments to the Q# operations and pass them along to the Q# target machine (here the simulator). When results are finished, they are returned by the Python function run game at scales.

This Python file should be runnable as a script, so we also need to define $__{main}$ as well. This is where we can do what Lancelot does in his head by using our host program in Python to take the measurements and scales and fit them to a model for Dagonet's rotation. The best model for how the rotation angle changes the measurement results is given by where Θ is Dagonet's hidden angle and where \mathrm{scale} is Dagonet's scale factor:

$$\Pr(1) = \sin(\frac{\theta * \text{scale}}{2})^2$$

Exercise 8.2

This model can be found if you use Born's rule! We have put the definition from Chapter 2 below, see if you can plot the resulting value as a function of Lancelot's scale using Python. Does your plot look like a trigonometric function?

Equation 8.1. Born's rule $\Pr(\text{measurement}|\text{state}) = |\langle \text{measurement}| \text{ state} \rangle|^2$ HINT: For Lancelot's measurements, the $\frac{\langle \text{measurement}|}{\langle \text{measurement}|} \text{ part of Born's rule is given by }^{\langle 1|}. \text{ Immediately before measuring, his qubit will be in the state }^{HR_1(\theta * \text{scale})H|0\rangle}. \text{ You can simulate the R1 operation in QuTiP by using the matrix form in the Q# reference at https://docs.microsoft.com/qsharp/api/qsharp/microsoft.quantum.intrinsic.r1.}$

Once you have that model and data, you can use the scipy.optimize function from the SciPy Python package to fit your data to the model. The value it finds for the Θ parameter is Dagonet's hidden angle! Check out for an example of how to pull this all together.

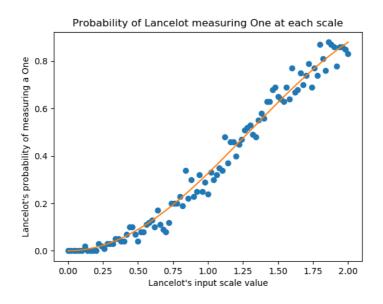
```
Listing 8.13. The code that will run when we run host.py as a script.
if name == " main ":
   import matplotlib.pyplot as plt
   scales = np.linspace(0, 2, 101)
   for control in (False, True):
       data = run game at scales(scales, control=control)
       def rotation_model(scale, angle):
           return np.sin(angle * scale / 2) ** 2
       angle_guess, est_error = optimization.curve_fit(
           rotation_model, scales, data, BIGGEST_ANGLE / 2,
           bounds=[0, BIGGEST_ANGLE]
       print(f"The hidden angle you think was {angle guess}!")
       plt.figure()
       plt.plot(scales, data, 'o')
       plt.title("Probability of Lancelot measuring One at each scale")
       plt.xlabel("Lancelot's input scale value")
       plt.ylabel("Lancelot's probability of measuring a One")
       plt.plot(scales, rotation_model(scales, angle_guess))
   plt.show()
```

- 1 This script will plot the data and fitted results so you will need to import the friendly matplotlib.
- 2Lancelot's list of inputs to the game (that is, his scales) are generated as just a regularly spaced, sequential list of numbers from np.linspace.
- 3This script will run both versions of the game simulation so they can be compared; don't worry about the control = True case for now, we'll come back to that shortly.
- ① data stores the result from the Q# simulation that runs from Python in run game at scales.
- This model represents the operation that is done on the qubit, i.e. its a rotation of an unknown angle where the angle can be multiplied by scale which is Lancelot's input to the game. Lancelot will take that data he gets back from his measurements and fit it to a model, with the goal of fitting the angle parameter from the data.

- **6** You can use a standard scipy function called optimization.curve_fit that takes a model of a function, inputs, measured data, and an initial guess to try and fit all the parameters of the model. This is what Lancelot needs to do to convert the data he collected into an actual angle to guess.
- Validating the fit found by optimization.curve_fit is important so you can plot both the data and the fitted model to see if it looks right.
- 8 This displays the plots with the data and fit in new window.

Now that you have a host program you can use to run the whole game, you can see that Lancelot does a pretty reasonable job at figuring out what angle Dagonet hid in his Q# operation. By taking different measurements and using classical data science techniques, Lancelot can estimate the phase that Dagonet's operation applies to his qubits. Running host.py should generate 2 pop up windows that show plots of the measurement probabilities as a function of Lancelot's scale for two strategies he can use. The first is the approach we have already outlined above, the latter you will implement in the last section of the chapter. See for an example for what you should see.

Figure 8.9. An example of one of the two plots that should pop up when you run host.py



Since the SciPy fitting packages are not perfect, sometimes the fitted parameter it finds won't be right. Just try running it a few times and hopefully the fitting algorithm will do better the next time. If you have any questions about the plotting package matplotlib check out these other titles from Manning: Data Science Bootcamp Chapter 2, and Data Science with Python and Dask Chapters 7 and 8.

You can see from these plots, that you were able to fit Lancelot's data fairly well. That

means that the fitted value that you find in angle_guess is a pretty good approximation to Dagonet's hidden angle!

There's still one more nagging problem with Lancelot's strategy, though: every time he performs a measurement, he needs to prepare the right input to pass to Dagonet's operation. In this particular game, that may not be much of a problem, but as you explore bigger applications of this game in the next chapter, it can be expensive to prepare the input register in the right state every time. Thankfully, you can use *controlled operations* to reuse the same inputs over and over again, as you'll see in the rest of the chapter.

You have seen examples of controlled operations already (like CNOT), but it turns out many other quantum operations can also be applied conditionally which can be very useful. Controlled operations along with the last new quantum computing concept you will need (eigenstates) will help you to implement a technique you have already seen at the end of the last chapter, called phase kickback.

TIP

There will be a lot of discussion about local and global phases in next few sections. Recall that global phase is a complex coefficient that can be factored out of all of the terms of your state, and cannot be observed. If you need a refresher on phase, check out Chapters 4 and 5.

8.5 EIGENSTATES AND LOCAL PHASES

By now, you've seen that the x quantum operation allows us to flip bits $|0\rangle \leftrightarrow |1\rangle$) and that the z operation allows us to flip phases $|+\rangle \leftrightarrow |-\rangle$). Both of these operations, though, only apply global phases to some input states. As you saw in previous Chapters, we can't actually learn anything about global phases, so understanding what states each operation leaves alone is important to understanding what we can learn by applying that operation.

For example, let's revisit the z operation. In , you can see what happens when we try to use z not to flip a qubit between the $|+\rangle$ and $|-\rangle$ states, but on an input qubit in the $|0\rangle$ state.

Listing 8.14. Applying the \mathbb{Z} instruction to a qubit in the $|0\rangle$ state.

```
using (qubit = Qubit()) {
    Z(qubit);
    DumpRegister((), [qubit]);
    Reset(qubit);
}
```

- (1) As usual in Q#, we start by allocating a Qubit with a using block. This will provide us with a fresh qubit in the |0\) state.
- 2 We then apply a Z operation, such that the state of qubit is transformed to $Z|0\rangle = |0\rangle$.
- 3 To confirm that the Z operation didn't do anything, we can use the <code>DumpRegister</code> function to instruct the simulator to print out all of its diagnostic information--- in this case, the full state vector. If we run

this on a target machine other than a simulator, we won't get a state vector, but whatever other diagnostics that machine offers (e.g.: hardware IDs).

④ Finally, we reset the qubit before releasing it. This isn't strictly needed, since we know beforehand that that qubit is still in the |0⟩ state.

Figure 8.10. Output of running

The second line shows you that the coefficient of the $|1\rangle$ state is 0. Taken with what you learned from the first line, you have that the state being dumped is $1|0\rangle + 0|1\rangle = |0\rangle$.

Note that in the listing above, z didn't do anything to qubit since $Z|0\rangle = |0\rangle$. If we modify the above listing to prepare $|1\rangle$ instead by using x before z, we'll see something very similar.

Listing 8.15. Applying the z instruction to a qubit in the $|1\rangle$ state.

```
using (qubit = Qubit()) {
   X(qubit);
   Z(qubit);
   DumpRegister("1.txt", [qubit]);
   Reset(qubit);
}
```

- **1** As before, if you want to prepare a $|1\rangle$ state, you can use that $|1\rangle = X |0\rangle$.
- 2 We next repeat our experiment from above, but with a different input.
- 3 As before, we can write out the state of qubit to a text file, using that we're running on a simulator that keeps the state internally.

Listing 8.16. Output

1 This file represents the vector [[0], [-1]], or -|1) in Dirac notation.

The effect of applying the z operation to a $|1\rangle$ state was to flip the sign of the state of qubit. This is another example of a *global phase*, as you saw in Chapters 5 and 7.

GLOBAL PHASE

Whenever two states $|\Psi\rangle$ and $|\Phi\rangle$ are only different by a complex number $e^{i\Phi}$, $|\Phi\rangle=e^{i\Phi}|\Psi\rangle$, we say that $|\Psi\rangle$ and $|\Phi\rangle$ vary by a *global phase*. For example, $|0\rangle$ and $|0\rangle$ differ by a global phase of -1 = $e^{i\pi}$.

The global phase of a state doesn't affect any measurement probabilities, so we cannot ever detect whether we applied a z operation when its input is in either the $|0\rangle$ or $|1\rangle$ state. We can confirm this by using the AssertQubit operation, which

checks the probability of a particular measurement result.

- 1 This call to AssertQubit checks that measuring qubit will return the result Zero, and terminates the program if that's not the case.
- ②After this, qubit is in the -|0) state as opposed to the |0) state. That is, the Z operation applies a global phase to the state of qubit.
- 3 Calling AssertQubit again, we can check that the probability of obtaining a Zero result is still 1.
- 4 Printing a message lets us check that the quantum program proceeds past both assertions.

Running this snippet will simply print out Passed test!, since the calls to AssertQubit don't do anything in the case that the assertion succeeds. Using assertions like this lets us write unit tests that use simulators to confirm our understanding of how particular quantum programs will behave. On actual quantum hardware, since we can't actually do this kind of check due to the No-Cloning Theorem, assertions can be safely stripped out.

DON'T DEPEND ON ASSERTIONS!

Assertions can be really useful tools for writing unit tests and for checking the correctness of your quantum programs. That said, it's important to remember that they will be stripped out when running your program on actual quantum hardware, so don't use assertions to *make* your program run correctly.

Of course, this is also just good programming practice; assertions in clasical languages like Python can often be disabled for performance reasons, such that you can't rely on assertions always being there for you.

Identifying which quantum states are assigned global phases by an operation u gives us a way of understanding the behavior of that quantum operation. We call such states *eigenstates* of the operation u. If two operations have the same eigenstates and apply the same global phases to each of those eigenstates, there is no way to tell those two operations apart. Just like if two classical functions have the same truth table, you can't tell which one is which, no matter what state your qubits are in when you apply each operation. This means that we can understand operations not only by a matrix representation of them, but also by understanding what their eigenstates are and what global phase the operation applies to each. As you have seen, we cannot directly learn about the global phase of a qubit, so in the next section you will learn how you can use controlled versions of an operation to turn that global phase into a local one you can actually measure. For now though, let's summarize this with a more formal definition of what an eigenstate is.

EIGENSTATE

If after applying an operation U, the state of a register of qubits qs is only modified by a global phase, then we say that the state of that register is an *eigenstate* of U. For example, $|0\rangle$ and $|1\rangle$ are both eigenstates of the Z operation. Similarly, $|+\rangle$ and $|-\rangle$ are both eigenstates of X.

Try out an exercise to practice working with eigenstates.

Exercise 8.3

Try writing Q# programs that use AssertQubit and DumpMachine to verify that:

- |+ and |- are both eigenstates of the X operation.
- |0⟩ and |1⟩ are both eigenstates of the Rz operation, regardless of what angle you choose to rotate by.

For even more practice, try figuring out what the eigenstates of the Y and CNOT operations and writing a Q# program to verify your guesses!

HINT: You can find the vector form of the eigenstates of a unitary operation using QuTiP. For instance, the eigenstates of the Y operation are given by qt.sigmay().eigenstates(). From there, you can use what you learned about rotations in Chapters 4 and 5 to figure out which Q# operations prepare those states.

Don't forget you can always test if a particular state is an eigenstate of an operation by just writing a quick test in Q#!

Eigenstates are a very useful concept and can be leveraged in a variety of quantum computing algorithms. You will use them in the next section along with controlled operations to implement a quantum development technique called *phase kickback* which you saw a bit of at the end of Chapter 7.

DEEP DIVE: It's only proper

Eigenstates get their name from a concept used throughout linear algebra, known as eigenvectors. Just as an eigenstate is a state that is left alone by a quantum operation, an eigenvector is a vector that is preserved up to a scaling factor when multiplied by a matrix. That is, if for a matrix A, $A\vec{x} = \lambda x$ for some number Λ , then \vec{x} is an eigenvector of A. We say that Λ is the corresponding eigenvalue.

The prefix "eigen-", German for "proper" or "characteristic," points to that eigenvectors and eigenvalues help us to understand the properties or characteristics of matrices. In particular, if a matrix A commutes with its conjugate transpose (that is, if $AA^{\dagger} = A^{\dagger}A$), then it can be *decomposed* into projectors onto eigenvectors, each scaled by their eigenvalues,

$$A = \sum_{i} \lambda_{i} \vec{x}_{i} \vec{x}_{i}^{\dagger}$$

Since this condition always holds for unitary matrices, this means that we can always understand quantum operations by decomposing them into their eigenstates and the phases applied to each eigenstate. For example, $Z=|0\rangle \left<0|-|1\rangle \left<1|\right|$ and $X=|+\rangle \left<+|-|-\rangle \left<-|$.

This way of thinking about states and operations can often help us make sense of different quantum computing concepts. In fact, another way of thinking about the phase estimation game you're working on playing in this chapter is as an algorithm for learning the phases associated with each eigenstate! You'll

see in the next chapter that this connects especially well to some applications, such as to learning properties of chemical systems.

8.6 CONTROLLED APPLICATION: TURNING GLOBAL PHASES INTO LOCAL PHASES

From what you have seen and can test, *global* phases of states are unobservable, while *local* phases of states can be measured. For example, consider the

state $\frac{1}{\sqrt{2}}(-i|0\rangle-i|1\rangle)=\frac{-i}{\sqrt{2}}(|0\rangle+|1\rangle)$. There is no measurement you could do to differentiate that state from $(\frac{(|0\rangle+|1\rangle)}{\sqrt{2}}$). However, you could distinguish either of those two states from $\frac{(|0\rangle-|1\rangle)}{\sqrt{2}}$ as it differs by a _local_ phase; that is, one of the states has a stem:[in front of the |1⟩ and the other has a -.

TIP

If you want a refresher on phases and how to think of them as rotations, see Chapters 4 and 5. When you are using the simulator as your target machine, the output of <code>DumpMachine</code> and <code>DumpReqister</code> can also be used to help learn about the phases of states.

From the last section, you played around with eigenstates and saw the global phases of eigenstates can carry information about an operation, let's call it U. If you want to learn this global phase information about the eigenstates, then it seems like you're stuck. If Lancelot were to only prepare eigenstates of Dagonet's operation, he'd never be able to learn what angle Dagonet had hidden.

Quantum algorithms to the rescue! There's a very useful trick that you can apply to turn global phases applied by an operation U into *local* phases applied by a closely related operation. To see how this works, let's return to the CNOT operation. Recall from Chapter 5 that you can simulate CNOT using a unitary matrix,

EQUATION 8.2. SIMULATING THE CNOT OPERATION WITH A UNITARY

$$U_{\text{CNOT}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

When you first encountered this matrix in Chapter 5, you used the analogy between unitary matrices and classical truth tables to work out that the CNOT operation swaps the $|10\rangle$ and $|11\rangle$ states, but leaves qubits in the $|00\rangle$ and $|01\rangle$ states alone. That is, CNOT flips the state of its second qubit, *controlled* on the state of the first qubit. As show in , you can read the unitary matrix for the CNOT operation as describing a kind of "quantum if" statement: "if the control qubit is in the $|1\rangle$ state, then apply the X operation to the target qubit."

To use the CNOT operation in Q# you can try out the following snippet:

```
using (control = Qubit(), target = Qubit()) {
   H(control);
   X(target);

  CNOT(control, target);
  DumpMachine();

  Reset(control);
   Reset(target);
}
```

- 1 Prepare the control qubit in |+>.
- 2 Prepare the target qubit in |1).
- 3 Apply CNOT and print out the state of the simulator.

By thinking about CNOT as a quantum analogue to a conditional statement, you can write out its unitary matrix a bit more directly. In particular, you can look at the unitary matrix for the CNOT operation as a kind of "block matrix" that you can build up using the tensor product that you saw in Chapter 4.

$$U_{\text{CNOT}} = \begin{pmatrix} \mathbb{1} & 0 \\ 0 & X \end{pmatrix} = |0\rangle \langle 0| \otimes \mathbb{1} + |1\rangle \langle 1| \otimes X$$

Exercise 8.4

Verify that $|0\rangle\langle 0|\otimes \mathbb{1} + |1\rangle\langle 1|\otimes X$ is the same as the equation above.

HINT: You can verify this by hand, by using NumPy's np.kron function, or QuTiP's qt.tensor function. If you need a refresher, check out how you simulated teleportation in Chapter 5, or check out the derivation of the Deutsch-Jozsa algorithm in Chapter 7.

You can construct other operations following this pattern, such as the CZ (controlled-Z) operation:

$$U_{\rm CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} = |00\rangle \langle 00| + |01\rangle \langle 10| + |10\rangle \langle 10| - |11\rangle \langle 11|$$

Much like the CNOT operation did the same thing as an x operation (namely, applied a bit flip), but controlled on the state of another qubit, when its control qubit is in the |1| state, the CZ operation flips a phase in the same way as the Z operation.

Figure 8.11. Writing out unitary matrices for controlled operations

Start with the CNOT operation that you saw in Chapter 5. You can represent it by a unitary matrix. Since CNOT acts on two qubits (the control and the target qubit), its unitary matrix is a 4×4 matrix:

You can write down unitary matrices for other controlled operations in the same way. For instance, the Controlled Z operation (CZ for short) is represented by a matrix with the 11 and 2 matrices on its diagonal.

The upper-left part of this matrix tells you what the CNOT operation does when the control qubit is in the |0\) state.

$$U_{\text{CNOT}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

The lower-right part tells you what happens when the control qubit is in the |1) state. Note that this part of the matrix for CNOT is the same as the matrix for the X operation.

$$U_{\rm CZ} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

You can see an example of how this works in , let's see how controlling z works in practice by writing out some Q# to give CZ a try:

```
Listing 8.17. Snippet showing how to test out the Q# operation CZ.

using ((control, target) = (Qubit(), Qubit())) {
    H(control);
    X(target);
    CZ(control, target);
    DumpRegister("cz-output.txt", [control, target]);
    Reset(control);
    Reset(target);
}
```

- 1 Prepare the control qubit in |+).
- 2 Prepare the target qubit in |1).
- 3 Apply CZ and save out the resulting state.

```
Listing 8.18. Output
# wave function for qubits with ids (least to most significant): 0;1
|0⟩:
        0.000000 + 0.000000 i ==
                                                           [ 0.000000 ]
        0.000000 + 0.000000 i ==
|1>:
                                                           [ 0.000000 ]
12>:
       0.707107 + 0.000000 i ==
                                                           [ 0.500000 ]
0.00000 rad ]
      -0.707107 + 0.000000 i ==
                                      ******
                                                           [ 0.500000 ] ---
3.14159 rad ]
```

If you run this, the contents of cz-output.txt will show you that the final state of

the [control, target] register is $U_{\rm CZ} \left| +1 \right\rangle = \left| -1 \right\rangle$.

Exercise 8.5

Either by hand or using QuTiP, verify that the above output is the same as $|-1\rangle=|-\rangle\otimes|1\rangle$

NOTE: If you seem to get the right answer other than that the order of the qubits are swapped, note that DumpMachine uses a *little-endian* representation to order states. In little-endian, $|2\rangle$ is short-hand for $|01\rangle$, not $|10\rangle$. If this seems confusing, blame the x86 processor architecture...

That is, based on the state of the **target**, the state of the **control** changed as a result, just as you saw in Chapter 6 with the Deutsch-Jozsa algorithm! This is because the phase applied by z in the case where control was in the $|0\rangle$ state is not the same as when control was in the $|1\rangle$ state, an effect known as **phase kickback**. In Chapter 6, you used phase kickback with a pair of qubits in the $|+-\rangle$ state to tell if the CNOT operation had been applied or not. Here, you've seen that you can use the CZ operation (also known as the Controlled z operation) to learn about the global phase applied by the z operation.

IMPORTANT

Even though $|1\rangle$ is an eigenstate of the Z operation, $|+1\rangle$ is **not** an eigenstate of the CZ operation. This means that calling CZ on a register in the $|+1\rangle$ state has an observable effect!

Phase kickback is a common quantum programming technique, as it allows us to turn what would otherwise be global phases into a phase between the $|0\rangle$ and $|1\rangle$ branches of the control qubit. In the CZ example above, both the input state $|+\rangle|1\rangle$ and the output state $|-\rangle|1\rangle$ are product states, allowing us to measure the control qubit without affecting the target qubit.

THINK GLOBALLY, LEARN PHASES LOCALLY.

Note that a global phase difference between $|1\rangle$ and $Z|1\rangle = -|1\rangle$ became a *local* phase difference between $|1\rangle$ and $U_{CZ}|+1\rangle = |-1\rangle$. That is, by controlling the Z instruction on a qubit in the $|+\rangle$ state, we were able to learn what would have been a global phase without control.

```
using ((control, target) = (Qubit(), Qubit())) {
   H(control);
   X(target);

CZ(control, target);
   if (M(control) == One) { X(control); }

   // Now let's dump only the state of the target.
   DumpRegister("cz-target-only.txt", [target]);

   Reset(target);
}
```

- Prepare the control qubit in |+), and the target qubit in |1).
- ②Apply CZ and save out the resulting state. Before you dump the state of the target qubit, though, let's measure and reset the control qubit.
- 3 Fun fact: this is actually how the Reset operation is implemented in Q#.

4 You already reset control above, so we don't need to reset it again here.

① As expected, we get that the target qubit stays in the |1) state, ready to feed into another CZ operation.

8.7 CONTROLLING ANY OPERATION

Thinking back to Lancelot's and Dagonet's game, it'd be really helpful if you could help Lancelot reuse the qubit that he passes into Dagonet's operation, so that he doesn't have to re-prepare it every time. Thankfully, using controlled operations to implement phase kickback gives a hint as to how you could do it. In particular, when you used phase kickback in Chapters 6 to 7 to implement the Deutsch–Jozsa algorithm, the target qubit was in the |-) state both at the start and end of the algorithm. That means that Lancelot could re-use the same qubit for each round of his game, and not need to re-prepare each time. That didn't matter for Deutsch–Jozsa, since you only ran one round of Nimue's and Merlin's game, but it's exactly the right trick for Lancelot to win his game with Dagonet, so let's look at how you can help him use phase kickback.

The trouble is that, while phase kickback is a very helpful tool to have in your toolbox as a quantum developer, so far you have only seen how to use it with the x and z operations. We know for our game, Dagonet told Lancelot that he will use the R1 operation; is there a way you can use phase kickback to help here? The pattern you used to implement phase kickback in the previous section only required you to be able to control an operation, so what you will need is a way to control the op::Apply operation that Dagonet gives Lancelot. In Q#, this is as simple as writing Controlled op::Apply instead of op::Apply, thanks to the Controlled functor. Much like the Adjoint functor that you saw in Chapter 6, Controlled is a Q# keyword that modifies how an operation behaves, in this case to turn it into its controlled version.

TIP

Just like is Adj indicates that an operation can be used with Adjoint, is Ctl in the type of an operation indicates that it can be used with the Controlled functor. If you want to denote that an operation supports both, you can write is Adj + Ctl. For example, the type of the X operation is (Qubit \Rightarrow Unit is Adj + Ctl), letting you know that X is both adjointable and controllable.

Thus, to help Lancelot out, you can take the op::Apply(scale, target) line and make it Controlled op::Apply([control], (scale, target)) and you have the controlled version of R1.

While that does solve Lancelot's problem, it can be helpful to unpack what's

happening under the hood just a little bit more. Any unitary operation (that is, a quantum operation that doesn't allocate, deallocate, or measure qubits) can be controlled, the same way you controlled the z operation to get CZ, and as you controlled x to get CNOT. For instance, we can define a controlled-controlled-NOT (CCNOT, also known as Toffoli) operation as an operation that takes two control qubits and flips its target if *both* controls are in the $|1\rangle$ state. Mathematically, we write that the CCNOT operation transforms an input state $|x\rangle|y\rangle|z\rangle$ to the output $|x\rangle|y\rangle|zXOR(yANDz)\rangle$. We can also write down a matrix that lets us simulate the CCNOT operation:

Similarly, the controlled-SWAP operation (also known as the Fredkin operation) transforms its input states from $|1\rangle|y\rangle|z\rangle$ to $|1\rangle|z\rangle|y\rangle$, and leaves its input the same when the first qubit is in the state $|0\rangle$.

TIP

We can make a controlled-SWAP out of three CCNOT operations: CCNOT (a, b, c); CCNOT (a, c, b); CCNOT (a, b, c); is equivalent to Controlled SWAP ([a], (b, c));. To see this, note that we can also make the uncontrolled SWAP operation from three CNOT operations for the same reason that we can swap two classical registers in-place using a sequence of three classical XOR operations.

We can generalize this pattern for any unitary operation U (that is, any operation which neither allocates, deallocates, nor measures its qubits). In Q#, the transformation performed by using the Controlled functor adds a new input to an operation representing which qubits should be used as controls.

TIP

This is where the fact that Q# is a tuple-in tuple-out language comes in very handy. Since every operation takes exactly one input, for any operation $\tt U$, $\tt Controlled\ U$ takes the original input to $\tt U$ as its second input.

In fact, the CNOT and CZ operations are simply shorthand for appropriate calls to Controlled. In , you can see more examples of this pattern.

TIF

Just like Adjoint works on any operation that has is Adj in its type (as you saw in Chapter 6), the Controlled functor works on any operation that has is Ctl in its type.

Table 8.1. Some examples of controlled operations in Q#

Description	Shorthand	Definition
Controlled-NOT	CNOT(control, target)	Controlled X([control], target)
Controlled- Controlled-NOT (Toffoli)	CCNOT(control0, control1, target)	Controlled X([control0, control1], target)
Controlled-SWAP (Fredkin)	n/a	Controlled SWAP([control], (target1, target2))
Controlled Y	CY(control, target)	Controlled Y([control], target)
Controlled-PHASE	CZ(control, target)	Controlled Z([control], target)

As you saw with the CZ example, controlling operations in this way lets you turn global phases such as those applied to eigenstates into relative phases that we can learn through measurements.

More than that, by using controlled rotations to kickback phase onto the control register, you can also reuse the same target qubit over and over. When you applied CZ to a target register that was in an eigenstate of Z, that target register stayed in the same state, even though the control register changed. In the rest of this Chapter, you'll see how you can use that to finish Lancelot's strategy for his little game with Dagonet.

8.8 IMPLEMENTING LANCELOT'S BEST STRATEGY FOR THE PHASE ESTIMATION GAME

You now have everything you need to write out a slightly different strategy for Lancelot that will allow him to use controlled operations to re-use the same qubits. As noted before, this may not make a huge impact for Dagonet's game, but does for other applications of quantum computing.

For example, in the next Chapter, you'll see how problems in quantum chemistry can be solved using a very similar game to the one that Dagonet and Lancelot are playing. There, however, preparing the right input state can require calling a lot of different quantum operations, such that if you can preserve the target qubit for later use, you can gain quite a lot of performance.

Let's review briefly what Lancelot's previous steps were before.

224

8.8.1 Steps for Dagonet and Lancelot's game

- 1. Dagonet picks a secret angle for a single qubit rotation operation
- 2. Dagonet prepares an operation for Lancelot to use that hides their secret angle, and allows Lancelot one additional input of a number (we'll call it a scale) that will get multiplied with the secret angle to give the total rotation angle of the operation.
- 3. Lancelot's best strategy for the game will be to select many scale values and estimate the probability of measuring One for each value. To do this, he will need to do steps a. and b. many times for each of the many scale values.
 - Prepare the |+> state and input the scale value in Dagonet's rotation. He uses the |+> state because he knows Dagonet is rotating around the Z axis, and for this state these rotations will result in a local phase change he can measure.
 - o After preparing each |+> state, Lancelot can rotate it with the secret operation, measure the qubit, and record the measurement.
- 4. Lancelot will then have data relating his scale factor and the probability he measured a 1 for that scale factor. He can then fit this data in his head, and get Dagonet's angle from the fitted parameters (he *is* the greatest knight in the land). We can use Python to help us do the same!

The step that will need to change to leverage your newfound skills with *controlled* rotations is step 3. For step 3a, what will change is the *allocation* of the qubits. Rather than allocating, preparing and measuring one qubit per measurement, he can allocate one target qubit to rotate with Dagonet's black box, and instead allocate and measure control qubits. He still can repeat the measurements, but won't have to measure or re-prepare the target each time.

You could summarize these changes by re-writing step 3 like this:

- 1. Lancelot's best strategy for the game will be to select many scale values and estimate the probability of measuring One for each value. To do this, he will need to do steps a. and b. many times for each of the many scale values. He prepares one qubit the |1⟩ state to use as the target for all of his measurements as it's an eigenstate of the hidden rotation.
 - o Prepare a second control qubit in the |+> state,
 - Apply the new controlled version of the secret rotation with Lancelot's scale value, unprepare the control qubit and measure it, then record the measurement.

In your code, these changes can be accomplished by modifying the previous EstimateProbabilityAtScale operation. Since the rotation axis could be anything Dagonet chooses (here it's just Z axis for convience), Lancelot needs to know how to control an arbitrary rotation. You can do this with the Controlled functor before calling the ScalableOperation that is passed from Dagonet. The Controlled functor is very similar to the Adjoint functor, in that it takes an operation and returns a new operation. Controlled U(control, target) is an

example of the syntax that would allow you to apply U to your target qubit, controlled on one or more control qubits. Take a look at for how you can modify EstimateProbabilityAtScale to use the Controlled functor.

Listing 8.20. Operation EstimateProbabilityAtScaleUsingControlledRotation

```
operation EstimateProbabilityAtScaleUsingControlledRotations(
                        target : Qubit,
                        scale : Double,
                        nMeasurements: Int,
                        op : ScalableOperation)
: Double {
    mutable nOnes = 0;
    for (idx in 0..nMeasurements - 1) {
        using (control = Qubit()) {
            within {
                H(control);
            } apply {
                Controlled op::Apply(
                    [control],
                    (scale, target)
                );
            set nOnes += Meas.MResetZ(control) == One
                            ? 1
                             0;
        }
    return Convert.IntAsDouble(nOnes) /
           Convert.IntAsDouble(nMeasurements);
```

- Now instead of allocating and preparing the target register, our guessing operation takes the target register as an input and reuses it. Thus, we only need to allocate and prepare the control register; we can do so with an H operation, regardless of what rotation Dagonet is hiding.
- 2The only other change we need to make is to call Controlled op::Apply instead of op::Apply, passing our new control qubit along with the original inputs.

The other modification you will have to make (step five) is the operation that actually runs the game. Since using the controlled operation actually allows Lancelot to reuse the target qubit, it only needs to be allocated once at the beginning of the game. See for how you can implement this.

Listing 8.21. Implementing the operation RunGameUsingControlledRotations.

- Using EstimateProbabilityAtScaleUsingControlledRotations, you can now allocate the target qubit once, since you'll be using it over and over again through each guess.
- ②Using the X operation, you can prepare the target in the |1\rangle state, an eigenstate of the (uncontrolled)
 R1 operation that Dagonet hid his angle in. Since each measurement uses phase kickback to affect
 only the control register, this preparation can be done once before playing the game.

8.9 SUMMARY

In this Part of the book, you've had a lot of fun using Q# and quantum computing to help the various denizens of Camelot. Using a quantum random number generator written in Q#, you were able to help Morgana pull one over on poor Lancelot. At the same time, you were able to help Merlin and Nimue each play their respective roles in deciding the fate of kings, learning about the Deutsch–Jozsa algorithm and phase kickback all the while. With the land at peace and the fires in Castle Camelot burning down for the night, you saw how to use everything you learned to help Lancelot play another game, winning this time by guessing a quantum operation hidden by Dagonet.

Throughout your Camelot escapades, you picked up quite a few new tricks to help you on your way as a quantum developer:

You have learned as a quantum developer:

- What is a quantum algorithm, and how to implement it with the Quantum Development Kit and Q#,
- How to use Q# from Python and Jupyter Notebook,
- How to design *oracles* to represent classical functions in quantum programs,
- User defined types,
- Controlled operations, and
- Phase kickback.

Going forward, it's time to bring what you've learned from Camelot back home, and apply these new techniques to something a bit more practical. In the next Chapter, you'll see how quantum computing can help in understanding chemistry problems. Don't worry if you don't remember the periodic table, you'll will be working with

some collegues who know the chemistry side of things, and are looking for your help in using everything that you learned in this Part to upgrade their workflow with quantum technology.

In this chapter you learned:

- Predict how quantum operations can be used to learn useful information about unknown operations with the quantum algorithm for phase kickback.
- Create new types in Q#
- Run Q# code from a Python host program
- Recognize important properties and behaviors of eigenstates and phase
- Program controlled quantum operations in Q#

Solving Chemistry Problems With Quantum Computers

This chapter covers:

- · Recognize how quantum computers can be used to solve difficult chemistry simulation problems
- Implement qubit rotation operations about arbitrary Pauli axes using Exp. operation,
- · Implement time-dependent operations on qubits using the Trotter-Suzuki method, and
- · Create programs using Q# standard library features for phase estimation, decomposition, etc.

Figure 9.1. In this Chapter, we will be covering topics in both techniques for developing quantum algorithms and applications for quantum programs.

Quantum Applications		
Libraries / Techniques		
Q# Language		
Intrinsic Operations		
Quantum Hardware	Simulators	

9.1 REAL CHEMISTRY APPLICATIONS FOR QUANTUM COMPUTING

So far in this book, you've learned how to use quantum devices to do everything from chat with your friend Eve, to help decide the fate of kings. In this Chapter, though, you'll get the chance to do something a bit more *practical*.

As it turns out, your quantum chemist friend Marie has hit the limit of what her classical computer can do to help her model different chemical systems. The problems that Marie solves with computational chemistry techniques can help to combat climate change, understand new materials, and to improve energy usage across industries; if you can help her out by using Q#, that could have quite a lot of practical applications indeed. Thankfully, by using what you learned about estimating phases in Chapter 8, you can do just that, so let's jump in!

Better tasting through chemistry

Any candymaker can tell you the difference that temperature makes: cook sugar to the "soft crack" stage and you get taffy, but if you add in a bit more energy, you can make any number of other delightful confections ranging from toffee to caramels. Everything about sugar, its taste, how it looks, and how it pulls, all changes depending on the energy that we pour into it with our saucepan. To no small degree, if we understand how the shape of sugar molecules changes as we add energy into a sweet melting pot, we understand sugar itself.

We see this effect not just with candy, but throughout our lives. Water, steam and ice are differentiated by understanding what shapes H_2O can take, what shapes it can arrange itself in, as a function of energy. In many cases, we want to understand how the molecule arranges itself as a function of energy based on simulations rather than experiments. In this Chapter, we build on the techniques in the previous several chapters, showing how we can simulate the energy of chemical systems so that we can understand them as keenly as a candymaker understands their craft, and use these chemical systems to make our lives better— maybe even a little sweeter.

To get a feel for how this all works, you and Marie agree to start off looking at *molecular hydrogen*, or H₂, as this is a simple enough chemical system that you can compare what you learn from your quantum program back to what classical modeling tools are able to simulate. That way, as you use the same techniques to study larger molecules than what can be simulated classically, you have a great test case you can fall back on to make sure everything is correct.

The simulation within the simulation

In this Chapter, your work with Marie involves two different kinds of simulation: using a classical computer to simulate a quantum computer, and using a quantum computer to simulate a different kind of quantum system. You'll often even want to do both, in that it's helpful in building quantum chemistry applications to use a classical computer to simulate how a quantum computer simulates a quantum chemical system. That way, when you run your quantum simulation on actual quantum hardware, you can be assured that it works correctly.

As shown in figure 9.2, Marie will start things off by using her expertise in quantum chemistry to describe a problem that she's interested in solving with a quantum computer; in this case, understanding the structure of H₂. For the most part, these problems consist of learning properties of a special kind of matrix called

a *Hamiltonian*. Once you get a Hamiltonian from Marie, you can then write a quantum operation much like the one shown in to simulate it and learn things about it that Marie can then use to understand how different chemicals behave. Throughout the rest of this Chapter, you'll develop the concepts and understanding that you need to implement the steps in .

Figure 9.2. Overview of the steps you will develop in this chapter to help Marie learn the ground state energy of her molecule.

Step 1: Hamiltonian

Marie gives her problem to you as a *Hamiltonian*, a mathematical object that describes the energy levels of a quantum system.

Step 2: Decompose Terms

You then decompose Marie's Hamiltonian into the *Pauli basis*, just as you can break down a state into the computational basis.

Step 3: Simulate Terms

You next turn each term in the decomposed Hamiltonian into a *multi-qubit rotation* operation that's easy to run on a quantum computer.

Step 4: Trotter-Suzuki

To simulate the entire Hamiltonian that Marie gives you, you then recombine each multi-qubit rotation using the *Trotter–Suzuki decomposition*, rapidly alternating which rotation you apply to your qubits.

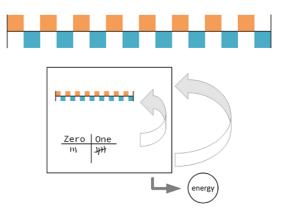
Step 5: Estimate Energy

To estimate the energy that Marie wants, you finish by playing the same phase estimation game that you helped Lancelot play in Chapter 8.

H

$$H = \sum_{i} H_{i}$$

let term = HamiltonianTerm[idx]; Exp(term::Axis, term::Angle, targets);



The following steps are what you will implement in this chapter for your Hamiltonian simulation algorithm.

9.1.1 Steps to simulate another quantum system, Marie's H $_2$ molecule, with your quantum computer:

- 1. Collaborate with Marie to figure out what Hamiltonian describes the energy levels in the system she is interested in, as well as an approximation of the ground (or lowest energy) state.
- 2. Prepare that approximation of the ground state, and use the Exp operation in Q# to implement evolution of the quantum system for each term of the Hamiltonian.
- 3. Using the Trotter-Suzuki decomposition implemented in the Q# function DecomposedIntoTimeStepsCA, simulate evolving your system under the action of all of the terms of the Hamiltonian at once by breaking the evolution into small steps.
- 4. Now that you have simulated the evolution of the system under the Hamiltonian, you can use phase estimation to learn about the change in phase of your quantum device.
- 5. You then make a final correction to the phase you estimate for the system and you then have the energy of the ground state for H₂.

Listing 9.1. Final Q# operation that can estimate the ground state energy of the H₂

molecule.

Without further ado, then, let's dive in to the first quantum concept you'll need to help Marie out: energy.

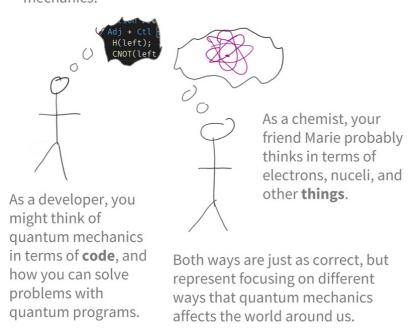
9.2 MANY PATHS LEAD TO QUANTUM MECHANICS

Thus far, you've learned about quantum mechanics using the language of computing: bits, qubits, instructions, devices, functions, and operations. Marie's way of thinking about quantum mechanics is very different, though. For her, quantum mechanics is a physical theory that tells her how subatomic particles like electrons behave. Thinking in terms of physics and chemistry, quantum mechanics is a theory about the *stuff* that everything around us is made out of.

232

Figure 9.3. Thinking about quantum mechanics in two very different ways.

There's a lot of different ways of thinking about quantum mechanics!



The two ways of thinking meet when it comes time to simulate how physical systems like molecules behave. You can use quantum computers to simulate how other quantum systems evolve and change in time. That is, quantum computing isn't about only physics or chemistry, but it can help you understand scientific problems like the ones Marie runs into.

Information is core to how you might think about quantum computation, but to the physics and chemistry way of thinking, quantum mechanics relies heavily on the concept of **energy**. Energy tells us how physical systems as varied as balls and compasses are affected by the world around them, giving us a consistent way to understand each of these different systems. In , you can see how the state of a ball on a hill and the state of a compass can both be described in the same way using the concept of energy.

As it turns out, energy doesn't just apply to classical systems like balls and compasses. Indeed, you can understand how *quantum* systems like electrons and nuclei will do by understanding the energy of different configurations. In quantum mechanics, energy is described by a special kind of matrix called a **Hamiltonian**. Any matrix which is its own adjoint can be used as a Hamiltonian, and Hamiltonians are *not* operations

themselves.

ADJOINT

Recall from Chapters 7 and 8 that the *adjoint* of a matrix A is its conjugate transpose, A^{\dagger} . This concept is closely related to the Adjoint keyword in Q#: if an operation op can be simulated by the unitary matrix U, then the operation Adjoint op can be simulated by U^{\dagger} .

In this chapter, you will be learning all the tools and techniques to figure our the energies of quantum systems that you have a Hamiltonian for. Often the process to get the Hamiltonians for systems will be a collaboration, but once you have it and a few more pieces of information you can estimate the energy of that system. This process is called *Hamiltonian simulation* and is critical to many different applications of quantum computing, including chemistry.

TIP

In fact, you've already seen a few examples of Hamiltonians in previous Chapters, actually: all of the Pauli matrices $(X,\,Y,\,$ and Z) are examples of Hamiltonians as well as being unitary matrices. Not all unitary matrices can be used as Hamiltonians, though! Most of the examples you'll see in this Chapter take some more work before we can apply them as quantum operations.

Marie is interested in understanding the energy of the bonds in her chemicals. So it makes sense if you both could come up with a Hamiltonian that could describe her molecule, then you can help estimate the energy she is interested in. In chemistry, this energy is often called the *ground state energy*, and the corresponding state is known as the *ground state* (or minimum energy state).

Once you have a Hamiltonian, the next step is to figure out how to construct the operations that will simulate how the quantum system changes in time as described by the Hamiltonian. In the next section, , you will learn how to describe the evolution of a quantum system under a Hamiltonian.

Then, with operators in hand that represent that Hamiltonian, the next challenge is to figure out how to simulate them on your quantum device. There won't likely just be a single operation built into the physical device that will exactly what you need, so you will need a way to decompose your operations for your Hamiltonian in terms of what your device can provide. In section 9.4, we will cover how you can take any operations and express them in terms of Pauli operations, which are commonly available as hardware instructions.

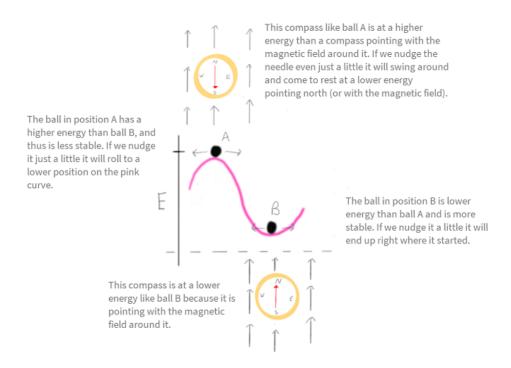
Now that you have your Hamiltonian expressed in as a sum of Pauli matrices, how do you simulate all of them on your system? Likely, you will have multiple terms that all sum together to represent the action of the Hamiltonian which don't necessarily commute. In , you will learn how to use the Trotter–Suzuki method to apply a little bit of each term in the operation to simulate evolving under the entire thing at once. Finally then you will have evolved your quantum system in a way that represents Marie's Hamiltonian!

Lastly, to work out the energy of the system described by the Hamiltonian you both found, you can use phase estimation to help Marie out! In section 9.6, you get to

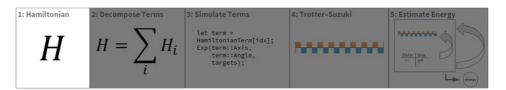
leverage the algorithm you learned in the last Chapter to learn about the phase applied to your qubits by simulating the Hamiltonian.

Let's get to it!

Figure 9.4. Using energy to understand how different physical systems are affected by their environments.



9.3 HAMILTONIANS DESCRIBE HOW QUANTUM SYSTEMS EVOLVE IN TIME.



To use a Hamiltonian to describe the energy of a physical or chemical system, you need to look at its eigenstates and their eigenvalues.

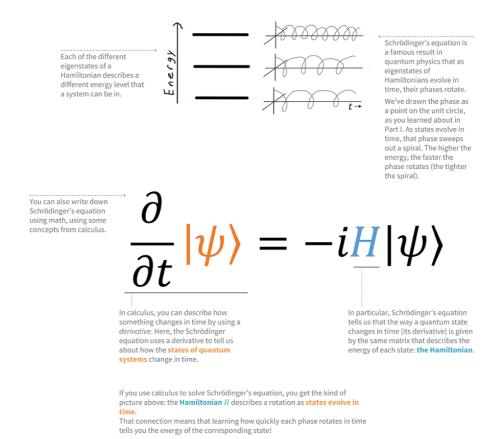
EIGENSTATES AND EIGENPHASES

Remember from Chapter 8 that, if a state $|\Psi\rangle$ is an eigenstate of an operation $\circ p$, then applying $\circ p$ to a register in the state $|\Psi\rangle$ at most applies a global phase to $|\Psi\rangle$. This phase is called the eigenvalue or the eigenphase corresponding to that eigenstate. Like all other global phases, this eigenphase cannot be directly observed, but you can use the Controlled functor that you learned about in Chapter 8 to turn that phase into a local phase.

Each eigenstate of a Hamiltonian is a state of constant energy; just like quantum operations don't do anything to eigenstates, a system that is in an eigenstate of its Hamiltonian will stay at that energy over time. The other property of eigenstates that you saw from the previous Chapter still holds here as well: the phase of each eigenstate still evolves in time.

The observation that the phases of eigenstates evolve in time is in fact the content of Schrödinger's equation, one of the most important equation in all of quantum physics. Schrödinger's equation tells us that as a quantum system evolves, each eigenstate of a Hamiltonian accumulates a phase proportional to its energy. Using math, you can write out Schrödinger's equation as shown in figure 9.5.

Figure 9.5. Schrödinger's equation, written out in mathematical notation.



The real experts were the friends we made along the way

TIP

How can it be that this is Chapter 9, and you're seeing the most important equation in quantum physics for the first time? Developing quantum applications may be closely connected to quantum physics, but it's not the same thing, and you don't need to be a physics expert to write quantum applications — you can be, if you're interested, but you don't have to be. In fact, Schrödinger's equation only even enters the picture here because you need it to understand how quantum computers may be used for practical impact.

Just like your friend Marie is an expert in quantum chemistry and not quantum computing, you don't need to know everything to do something awesome. That's what friends are for!

Schrödinger's equation relates how phases of different states evolve in time to the energy of those states. Since global phases are unobservable, and since eigenstates of Hamiltonians only acquire global phases as they evolve, Schrödinger's equation tells us that eigenstates of Hamiltonians don't evolve in time.

Schrödinger's equation is most critical for us in this chapter here because it relates the

energy of a system to phase, a very helpful connection given that in the last Chapter, you learned how to do phase estimation! There are other ways Schrödinger's equation is useful, one such being another way to look at implementing operations on our quantum systems.

One way to implement the rotations that you've seen so far in the book is to set up the right Hamiltonian and then... just wait. The time derivative $(\partial/\partial t)$ in Scrödinger's equation tells you that the way the state of your qubits rotate is entirely described by the energy associated with each state. For example, Scrödinger's equation tells us that if your Hamiltonian is $H = \omega Z$ for some number ω , and if you want to rotate by an angle Θ about the Z-axis, then you can let your qubit evolve for time $t = \Theta/\omega$.

Exercisee

Try writing other rotations that you've seen earlier in the book (e.g.: Rx and Ry) as Hamiltonians.

In listing 9.2, you can see a simple Q# operation that simulates evolving under the Hamiltonian $H = \omega Z$.

NOTE

In practice, a large part of the challenge of building a quantum computer is ensuring that qubits **don't** evolve except as instructed by a quantum program. It wouldn't be very useful if walking away from your quantum device for a moment meant that all your qubits were in totally different states when you got back. This is part of why as quantum developers, we tend to think at the level of instructions sent to a device — that is, quantum operations — and not directly in terms of Hamiltonians.

Switching gears for a moment and temporarily thinking in terms of Hamiltonians gives you a bit of what you need to start making progress on the problem that Marie asked for help with. After all, the problems that Marie works with are much easier to describe in that language. For example, in the previous Chapter you saw how you can learn about the phase applied by rotations like the one that Dagonet hid from Lancelot. You can also express that Dagonet's and Lancelot's game in terms of Hamiltonians, though. Suppose that Dagonet was hiding the rotation angle $2.1~\pi$; then, since his rotation was about the Z-axis, you could also describe that hidden rotation as a hidden Hamiltonian $H = -2.1\pi Z$.

NOTE

You need the - sign due to the - sign in Schrödinger's equation. Getting this wrong is about as common in quantum programming as off-by-one errors in other languages, so don't stress out if you forget once or twice, or even if you forget that pesky minus sign almost every time. You're still doing great.

Using this kind of a description, Lanceot's scale corresponds to how long he lets his qubits evolve under Dagonet's hidden Hamiltonian. While thinking in terms of a game makes it easier to write out quantum programs to learn Dagonet's hidden rotation, thinking in terms of a Hamiltonian makes it easier to map to the kinds of physical concepts that Marie is concerned with, such as field strengths and time.

Listing 9.2. Evolving under the Hamiltonian H = Z.

```
operation EvolveUnderZ(
    strength : Double,
    time : Double,
    target : Qubit
)
: Unit is Adj + Ctl {
    Rz(2.0 * strength * time, target); 4
}
```

- 1 You can write out an operation that simulates evolution under H = ωZ using Q#.
- 2 Here, ω says how large energies described the Hamiltonian are. For instance, if the Hamiltonian comes from a magnetic field, then ω tells us how strong that field is. This plays the role of Dagonet's hidden angle in the previous Chapter!
- 3 You can also choose how long a time you want to simulate the Hamiltonian for. This is analogous to Lancelot's scale from Chapter 8.
- The actual simulation here is just the one line, since rotations about the Z-axis are built into Q#. The only tricky part is that we need to multiply by 2.0, as Rz multiplies its angle by -1/2 by convention. Since Schrödinger's equation tells us that we need a minus sign, 2.0 gives us the angle we need to match figure 9.5.

Since Schrödinger's equation tells us that evolving Hamiltonians rotates quantum systems according to their energy, if you can **simulate** the Hamiltonian that Marie gives us, then you can play exactly the same phase estimation game as in Chapter 8 to learn the energy levels of that Hamiltonian.

Suppose that instead of $H = \omega Z$, Marie asks if you can simulate $H = \omega X$. Thankfully, Q# provides rotations about the X axis as well, so that you can modify the call to RZ in with a call to RX. Unfortunately, though not every Hamiltonian Marie is interested in is as simple as $H = \omega Z$ or $H = \omega X$, so let's look at what quantum development techniques you can use to simulate Hamiltonians that are a bit more difficult.

These aren't the Hamiltonians you are looking for

It is likely that when you start talking to Marie, that she will also be working on describing Hamiltonians for her system in her simulation and modeling software. However these are likely fermionic Hamiltonians which are different from the kind we are using here to describe how quantum devices change in time. As a part of the workflow for your collaboration you and Marie will likely need to use some tools like NWChem (http://www.nwchem-sw.org/) to actually convert between a Hamiltonian that describes how a chemical changes in time and how qubits change in time. It's out of scope for this book to look at these methods in detail, but there are great software tools that can help with this, check out the Quantum Development Kit documentation for detail if you're interested (https://docs.microsoft.com/quantum/libraries/chemistry/). This isn't a big deal for now, just a handy tip for when you are talking to your collaborators!

9.4 ROTATING AROUND ARBITRARY AXES WITH PAULI OPERATIONS



Moving up in complexity, perhaps Marie is interested in something that takes more than a single-qubit Hamiltonian to describe. If she gives you a Hamiltonian like $H = \omega X \otimes X$, what can you do to simulate it? Thankfully, what you learned about rotations in Part I is still useful for that kind of two-qubit Hamiltonian, as you can think of it as describing another kind of rotation.

UNSEEN

In the previous section, you saw that rotations like Rx, Ry, and Rz correspond to Hamiltonians like X, Y, and Z, respectively. You can think of two-qubit Hamiltonians like $X \otimes X$ as specifying an axis in much the same way; it turns out that there are 15 possible such orthogonal rotation axes for a two-qubit register, as opposed to the three dimensions that you get for a single-qubit register, so to draw things out as a picture, you'd need 13 more dimensions than paper normally comes with, making it somewhat difficult to draw out!

This rotation doesn't look like any of the built-in (that is, intrinsic) instructions that you've seen so far, though, so it might seem like you're stuck. As it turns out, though, you can still simulate this Hamiltonian using *single* qubit rotations like Rx, as long as you use some two-qubit operations on either side. In this section, you'll see how that works, and how Q# makes it easy to automate applying multiple-qubit rotations.

By way of getting started, let's look at some of the ways we can change what quantum operations do by surrounding them with other operations. You can always use math to reason about things, the way you saw in Chapter 7, but thankfully Q# also provides some nice testing functions and operations that can help you out. For example, in Chapter 7, you saw that surrounding a CNOT operation with H operations gives you a CNOT going the other direction; let's see how you can check that using Q#!

DON'T FORGET YOUR SHOES!

Recall that the within/apply block in applies the shoes-and-socks principle that you first learned about in Chapter 7. Most of the code examples in this section will make use of within/apply blocks to help keep track of our shoes-and-socks thinking.

Listing 9.3. Changing the control and target of a CNOT operation by using H operations

```
open Microsoft.Quantum.Diagnostics;

operation ApplyCNOT(register : Qubit[])
: Unit is Adj + Ctl {
    CNOT(register[0], register[1]);
}
```

```
poperation ApplyCNOTTheOtherWay(register : Qubit[])
: Unit is Adj + Ctl {
    within {
        ApplyToEachCA(H, register);
    } apply {
        CNOT(register[1], register[0]);
    }
}

operation CheckThatThisWorks() : Unit {
    AssertOperationsEqualReferenced(2, ApplyCNOT, ApplyCNOTTheOtherWay);
    Message("Woohoo!");
}
```

- 1) Since we'll use the AssertOperationsEqualReferenced operation in this snippet, you'll need to start by opening the Microsoft.Quantum.Diagnostics namespace. The operations and functons in this namespace help with testing and debugging quantum programs, and can be very useful in helping to make sure your programs work as intended.
- 2To compare the two ways of writing out a CNOT operation, we'll need each to be callable as an operation that takes an array of qubits representing a quantum register, so you can start off by writing an operation that just applies a CNOT operation whose control is the first qubit of a register and whose target is the second qubit.
- 3 To check the equivalence that you first saw in Chapter 7, you can write a second operation that reverses the control and target of a CNOT operation, but that uses a within/apply block to surround the call to CNOT with calls to the H operation.
- Finally, you can call the AssertOperationsEqualReferenced operation to check that both of the operations above do exactly the same thing. The first input to AssertOperationsEqualReferenced specifies how large a register each operation acts on, while the second and third inputs represent the operations being compared. If the operations do anything different whatsoever, the assertion will fail and the quantum program ends.
- If we get to this Message call, we know that the assertion in the previous line didn't fail, and thus that the operations ApplyCNOT and ApplyCNOTTheOtherWay do exactly the same thing to their input qubits. Put differently, if we see the message "Woohoo!", you can safely conclude that the two operations cannot be distinguished from each other by looking at what they do to the states of quantum registers.

IMPORTANT

Assertions like <code>AssertOperationsEqualReferenced</code> only make sense when run on a simulator, as they require violating the No-Cloning Theorem to run. On actual hardware, these kinds of assertions would be stripped out, just as running Python with the <code>-O</code> command line argument disables the <code>assert</code> keyword. That means that Q# assertions give us a way of cheating <code>safely</code>, since quantum programs using assertions do the same thing regardless of whether or not we cheat.

Exercise

Use QuTiP to verify that the two operations ApplyCNOT and ApplyCNOTTheOtherWay can be simulated by the same unitary matrix, and thus do the exact same thing.

Exercise

Just as you can use three classical XOR instructions to implement an in-place classical swap, you can use three CNOT operations to do the same thing as a single SWAP operation. In fact, the following Q# snippet does the same thing as SWAP (left, right):

```
CNOT(left, right);
CNOT(right, left);
CNOT(left, right);
```

Double-check that this is the same as SWAP (left, right), both by using AssertOperationsEqualReferenced, and by using QuTiP.

Extra credit: SWAP(left, right) is the same as SWAP(right, left), so the snippet above should work even if you start with CNOT(right, left) instead. Double-check that!

DEEP DIVE: The Choi-Jamiłkowski Isomorphism

The AssertOperationsEqualReferenced operaton in works using a neat piece of mathematics called the *Choi-Jamiłkowski isomorphism*, which says that any operation that can be simulated using a unitary matrix is perfectly equivalent to a particular state called its Choi state. This means that a simulator can effectively find the entire truth table for any adjointable operation (that is, any operation that has is Adj in its signature) by finding its Choi state.

The AssertOperationsEqualReferenced operation uses this concept to prepare a register of qubits in the Choi state for each of the operations passed as inputs. On a simulator, it's easy to cheat and check if two states are the same, even though the No-Cloning Theorem tells us we can't do that on an actual device.

When writing unit tests and other checks that quantum programs are correct, this can be a powerful technique for making use of classical simulators while still preventing cheating on actual hardware.

When you run CheckThatThisWorks, either in Jupyter Notebook (as you saw in Chapter 6), or at the command-line, you should see the message "Woohoo!" printed out, telling you that your Q# program ran past the call to AssertOperationsEqualReferenced. Since that assertion only passes if the two

operations that you give it do exactly the same thing for all possible inputs, that tells you that the equivalence you learned about in Chapter 7 works.

You can use the same logic to check how two-qubit operations like CNOT transform other operations. For example, transforming a call to X with calls to CNOT does the same thing as calling X multiple times, as demonstrated in .

Listing 9.4. Turning the x operation on a single qubit into an x operation on each qubit

```
of a register.
open Microsoft.Quantum.Diagnostics;
operation ApplyXUsingCNOTs(register : Qubit[])
: Unit is Adj + Ctl {
    within {
        ApplyToEachCA(
            CNOT(register[0], _),
            register[1...]
    } apply {
        X(register[0]);
operation CheckThatThisWorks() : Unit {
    AssertOperationsEqualReferenced(2,
        ApplyXUsingCNOTs,
        ApplyToEachCA(X, _)
                                                                 6
    Message("Woohoo!");
```

- ① Just as in the previous listing, you'll need two operations to pass as inputs to

 AssertOperationsEqualReferenced. Here, you can start by writing out a single call to the

 X operation, using a within/apply block to transform that with a sequence of calls to the CNOT operation.
- ②For the "socks" part of your within/apply block, you can write out the sequence of CNOT calls you need by using ApplyToEachCA together with the partial application technique that you learned about in Chapter 6.
- 3 This part of your call to ApplyToEachCA says to apply a CNOT operation controlled on the first qubit of a register to each element of an array of qubits.
- Next, you can specify the array whose elements you want to apply CNOT (register[0], _) on; that is, the array of targets that you'd like to use. Here, you can write out register[1...] to pick out all but the first (that is, the 0'th) element of the register `array. In Q#, register[1...] is shorthand for the array slice register[1..Length (register) 1]. Using this shorthand gives you a sequence of operations like CNOT (register[0], register[1]), CNOT (register[0], register[2]), and so forth.
- The "shoes" part of your within/apply block is a bit simpler, as it consists of just an X operation on the same qubit that you used as a control for your sequence of CNOT calls.

6 This time, instead of writing your own operation to compare to, you can compare instead to an X operation on each qubit in a register by using partial application again.

Exercise

Using QuTiP, check that when run on two-qubit registers, the two programs in can be simulated by the same unitary matrix and thus do the same thing to their input registers.

Exercise

Try modifying to see if both programs are equivalent when applied to more than two qubits.

NOTE: It can be pretty expensive to use <code>AssertOperationsEqualReferenced</code> for more than a few qubits.

You can also build up other interesting kinds of operations by using the within/apply concept. In particular, transforming a rotation by CNOT operations in the same way as lets you implement the kinds of multiple qubit rotations that Marie asked for at the beginning of this section. Using the DumpMachine and DumpRegister features that you learned about in Chapter 8, you can see that, in the same way that Rx applies an X-axis rotation between $|0\rangle$ and $|1\rangle$, you can implement a $(X \otimes X)$ -axis rotation between $|00\rangle$ and $|11\rangle$.

Listing 9.5. Turning the Rx operation on a single qubit into a multi-qubit rotation.

```
open Microsoft.Quantum.Diagnostics;
open Microsoft.Quantum.Math;

operation ApplyRotationAboutXX(angle : Double, register : Qubit[])
: Unit is Adj + Ct1 {
    within {
        CNOT(register[0], register[1]);
    } apply {
        Rx(angle, register[0]);
    }
}

operation DumpXXRotation() : Unit {
    let angle = PI() / 2.0;
    using (register = Qubit[2]) {
        ApplyRotationAboutXX(angle, register);
        DumpMachine();
        ResetAll(register);
    }
}
```

- ① Just as with , you'll use a call to CNOT as the "socks" part of a within/apply block. For simplicity, we've specialized to the two-qubit case in this listing, but you can use an ApplyToEachCA call in the same way again to work with registers of more than two qubits.
- 2 This time, instead of applying an X operation to the control qubit, you'll want to apply an X rotation about an arbitrary angle to the control qubit.

- **3** To check what your new ApplyRotationAboutXX operation does, you can start by asking your target machine for a two-qubit register with a using statement.
- $oldsymbol{\Phi}$ You can then apply your new rotation about the (X \otimes X)-axis to your new register to see what it does.
- Sas in Chapter 8, when run on a simulator, DumpMachine will print out the full state of the simulator, letting you check how your new rotation operation transformed the state of your register.
- (6) As usual, before releasing your register back to the target machine, you'll need to reset all of your qubits back into the |0) state.

Exercise

Try preparing your register in states other than |00\) before calling ApplyRotationAboutXX. Does your operation do what you expected?

HINT: Recall from Part I that you can prepare a copy of the $|1\rangle$ state by applying an X operation, and that you can prepare $|+\rangle$ by applying an H operation.

TIP

Your output from may look slightly different, as the IQ# kernel for Jupyter Notebook supports multiple different ways to label qubit states. By default, IQ# uses the "little-endian" convention, useful for arithmetic problems like the ones that you'll see in Chapter 11. To label qubit states using bitstrings like the ones you've seen so far in the book, run %config

dump.basisStateLabelingConvention = "Bitstring" from a new Jupyter Notebook
cell.

Figure 9.6. Output from running in a Jupyter Notebook.

Exercise

Try using DumpMachine to explore how the Rx operation acts on a single qubit, and compare to the two-qubit rotation about the $X \otimes X$ axis that you implemented in . How are the two rotation operations similar, and how do they differ? Compare rotating about the $X \otimes X$ axis with applying an Rx operation to each qubit in a two-qubit register.

In general, *any* rotation about an axis given by a tensor product of Pauli matrices (such as $X \otimes X$, $Y \otimes Z$ or $Z \otimes Z \otimes Z$) can be implemented by applying a single-qubit rotation transformed by a sequence of operations like CNOT and H. Finding what the right transformation is, however, can be a bit annoying, so Q# provides a nice built-in operation called $E\times D$ to help you out.

```
open Microsoft.Quantum.Diagnostics;
open Microsoft.Quantum.Math;
```

WARNING

The convention used by $\mathbb{E} \times \mathbb{P}$ and $\mathbb{R} \times$ to denote angles differ by a factor of -1 / 2. When using the $\mathbb{E} \times \mathbb{P}$ operation and single-qubit rotation operations in the same program, make sure to double-check all of your angles!



Using Exp, it's easy to simulate the Hamiltonian $H = \omega X \otimes X$, or even any other Hamiltonian made up of tensor products of Pauli matrices.

```
Listing 9.6. Using Exp to simulate evolving under X & X.

operation EvolveUnderXX(
    strength : Double,
    time : Double,
    target : Qubit
)
: Unit is Adj + Ctl {
    Exp([PauliX, PauliX], strength * time, target);
}
```

- ①Using what you've learned so far, you can now write out an operation to simulate evolving under a Hamiltonian proportional to (X⊗X), just as the operation that you wrote in simulated evolving under a Hamiltonian proportional to Z.
- 2 As before, you'll need a parameter to represent the strength of the Hamiltonian; that is, how large the energies described by your Hamiltonian are.
- (3) Next, you'll need a parameter that describes how long to simulate evolution for (analogous to Lancelot's scale parameter from Chapter 8).

1 Finally, you can ask for a rotation about the $(X \otimes X)$ -axis using the Exp operation provided by the Microsoft.Quantum.Intrinsic namespace. Here, you can specify $(X \otimes X)$ by the Q# value [PauliX, PauliX].

WARNING

It may be tempting to think that you can implement a two-qubit rotation about $Z \otimes Z$ by rotating the first qubit about Z, then rotating the second qubit about Z. These turn out to be very different operations, however.

$$R_z(\theta) \otimes R_z(\theta) = \begin{pmatrix} e^{-i\theta} & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{i\theta} \end{pmatrix} \qquad R_{zz}(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 & 0 & 0 \\ 0 & e^{i\theta/2} & 0 & 0 \\ 0 & 0 & e^{i\theta/2} & 0 \\ 0 & 0 & 0 & e^{-i\theta/2} \end{pmatrix}$$

One way to think of it is that a rotation about $Z \otimes Z$ is only sensitive to the *parity* of each computational basis state, so that $|00\rangle$ and $|11\rangle$ are each rotated by the same phase.

In any case, now that you have the Exp operation at your disposal, it's pretty easy to use it to write an operation that simulates each term in the Hamiltonian that Marie gives you, as shown in .

Listing 9.7. A Q# operation to simulate evolution under a single Hamiltonian term

```
operation EvolveUnderHamiltonianTerm(
    idxBondLength : Int,
    idxTerm : Int,
    stepSize : Double,
    qubits : Qubit[]
)
: Unit is Adj + Ctl {
    let (pauliString, idxQubits) = H2Terms(idxTerm);
    let coeff = (H2Coeff(idxBondLength))[idxTerm];
    let op = Exp(pauliString, stepSize * coeff, _);
    (RestrictedToSubregisterCA(op, idxQubits))(qubits);
}
```

- 1 The first input you'll need is an index to use to look up which of the Hamiltonians Marie gave you; remember, she wants to understand how the energy of an H₂ molecule changes with its bond length, so she'll give you a different Hamiltonian from each bond length.
- 2Next, you'll need an input saying which term of Marie's Hamiltonian you want to simulate evolution under.
- 3 You'll also need an input saying how long to simulate evolution for; that is, how long of a simulation step to take.
- ①Once you have all of that, you can get out the term from the Hamiltonian by using idxTerm together with the H2Terms function provided in the code repository for this book.
- **5** Next, you can get the coefficient of that term by using the H2Coeff function, also provided in the samples repository for this book.
- **6** To actually simulate evolution under that term, you can use Exp to do a rotation scaled by the simulation step size, just as you did with the EvolveUnderXX operation from .

Since not all terms affect all qubits, you can use the RestrictedToSubregisterCA operation provided with the Q# standard libraries to apply your call to Exp only to a subset of the qubits input.

In the next Section, you'll see how you can use this to simulate evolution under Marie's entire Hamiltonian.

9.5 MAKING THE CHANGE YOU WANT TO SEE IN THE SYSTEM

Now that you have seen how you can describe how a quantum device can change in time, using the concept of a Hamiltonian, a very natural question is how do you implement the particular Hamiltonian you want to simulate. Most quantum devices will have some operations that are easy for them to do. For instance, you saw in the previous section that it's straightforward to simulate evolution under any Hamiltonian that is given by a tensor product of Pauli matrices. That said, the Hamiltonian that you and Marie are interested in is likely *not* a built-in operation, but something that isn't directly available on your quantum computer.

TIP Usually it will be easy for devices to implement some of the Pauli operators, and maybe a few other operations; the game then becomes figuring out how to transform the operation you *want* into ones operations that the device can easily do.

If there is no easy button that will simulate evolution under your Hamiltonian, how can you implement a simulation of a particular Hamiltonian that you can apply to the qubits in your device?

Let's break it down. Literally. You learned all the way back in Chapter 2, you can describe a vector as a linear combination of *basis* vectors or directions. It turns out that you can do the same thing with matrices, and a really convenient basis for such is the Pauli operators.

TIP If you need a refresher as to what the Pauli matrices are, no worries, we've got you covered:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Just like you could describe any direction on the map with North and West, you can describe any matrix as a linear combination of Pauli matrices. For example,

$$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \frac{1}{2} \mathbb{1} + \frac{1}{2} Z$$

Similarly,

$$\begin{pmatrix} 2 & 3 \\ 4 & 5 \end{pmatrix} = \frac{1}{2} \left(7 \mathbb{1} + 7X - iY - 3Z \right)$$

The same holds for matrices acting on multiple qubits as well,

$$U_{\mathrm{SWAP}} = \frac{1}{2} (\mathbb{1} \otimes \mathbb{1} + X \otimes X + Y \otimes Y + Z \otimes Z)$$

Exercisee

Use QuTiP to verify the above equations.

HINT: You can use qt.qeye(2) to get a copy of $^{1/2}$, qt.sigmax() to get a copy of X, and so forth. To compute tensor products like $X \otimes X$, you can use qt.tensor.

This is good news, because you can then write the Hamiltonian you and Marie want to simulate as a linear combination of Pauli matrices; in the previous section, you saw that you can use Exp to easily simulate Hamiltonians that are made up only of tensor products of Pauli matrices. This makes the Pauli basis a very convenient one, such that it's likely that the workflow from Marie's chemistry tools will output the Hamiltonian for your quantum device in the Pauli basis already.

Let's look at the representation of the Hamiltonian that Marie wants you to simulate, using the Pauli basis to expand it out. Using her chemistry modeling skills, Marie can helpfully tell you that Hamiltonian is need you to simulate with your qubits is given by below, where each of a, b_0, ..., and b_4 is a real number that depends on what bond length she wants to simulate H₂ at.

Equation 9.1. Equation 9.1: Marie's Hamiltonian, given as a function of six real numbers.

$$H = a \mathbb{1} \otimes \mathbb{1} + b_0 \mathbb{Z} \otimes \mathbb{1} + b_1 \mathbb{1} \otimes \mathbb{Z} + b_2 \mathbb{Z} \otimes \mathbb{Z} + b_3 \mathbb{Y} \otimes \mathbb{Y} + b_4 \mathbb{X} \otimes \mathbb{X}$$

All the terms and coefficients Marie is using for the from the paper "Scalable Quantum Simulation of Molecular Energies," by O'Malley et. al. https://arxiv.org/abs/1512.06860. The exact coefficients depend on length of the bond between the hydrogen atoms, but all of these constants are helpfully typed in for you in the code repo for the

book: https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp.

With this representation of Marie's Hamiltonian in hand, it's time to figure out how to actually *use* it. There are six terms to this Hamiltonian, so which term should you apply first? Does the order matter? It unfortunately does often matter when simulating evolution of a system under a Hamiltonian what order the terms are used. In the next section, you will learn about a method that allows you to break up the evolution of the system into little steps to simulate evolution under all of the terms at once.

9.6 GOING THROUGH (VERY SMALL) CHANGES



At this point, it's helpful to take a step back and assess where you're at with helping Marie out. You've seen how to break arbitrary Hamiltonians down into sums of Pauli matrices, and how to use the Exp operation to simulate evolution under each term in that sum. In order to simulate arbitrary Hamiltonians, all that's left is to combine those simulations together to simulate the entire Hamiltonian. For that, you can use one more quantum computing trick called the *Trotter–Suzuki decomposition*.

Before getting into the details of the Trotter–Suzuki decomposition, though, let's go back to the map analogy that you've used throughout the book to break down linear algebra concepts. Suppose that you're exploring downtown Phoenix and decide to see what it feels like to go Northeast throughout the city. If you start out by going North several blocks, then go several blocks East, the route you trace out on a map won't look much like a diagonal line at all. On the other hand, if you switch going North and East each block, you'll trace out something much closer to a path that looks like it came out of the map of Minneapolis that you saw in Chapter 2. That is, you can simulate the way you might walk through Minneapolis even if you're stuck in Phoenix by quickly switching which way you walk.

Figure 9.7. If you're in downtown Phoenix, you can still simulate how you might walk through downtown Minneapolis by rapidly alternating directions.



In the previous section, you saw that just like states, the different terms in a Hamiltonian can be thought of as directions on a high-dimensional map. Tensor

products of Pauli matrices, such as $Z \otimes \mathbb{I}$ and $X \otimes Z$, play a similar role to the cardinal directions or axes of a map. When you try to simulate Marie's Hamiltonian, though, that doesn't point along a single axis, but along a kind of diagonal in that

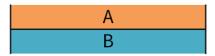
higher-dimensional space. That's where the Trotter-Suzuki decomposition comes in.

Just as when you quickly switch which direction you walk, your path looks more diagonal, you can rapidly switch between simulating different Hamiltonian terms. As illustrated in figure 9.8, the Trotter–Suzuki decomposition tells you that when you rapidly switch in this way, you approximately evolve under the sum of the different terms that you're simulating.

Figure 9.8. Using the Trotter–Suzuki decomposition to approximate evolving under two Hamiltonian terms at once.

Suppose that you have two quantum operations, A and B. Each of them simulates rotation under a different Hamiltonian term.

If you wanted to use them to simulate a combined Hamiltonian, it would be great if you could run both at the same time.



Sadly, you can only apply a single operation to a given set of qubits at a time. One solution might then be to use A, then to use B.



It turns out to give you a much better approximation, though, if you run A for a short time, then B for a short time, and keep alternating.



You could in principle write this out in Q# as a for-loop. In pseudocode, you might

©Manning Publications Co. To comment go to liveBook

251

have something like this:

Listing 9.8. Pseudocode for simulating a Hamiltonian term-by-term, using the Trotter-

Suzuki decomposition.

```
operation EvolveUnderHamiltonian(time, hamiltonian, register) {
    for (idx in 0..nTimeSteps - 1) {
        for (term in hamiltonian) {
            evolve under term for time / nTimeSteps
        }
    }
}
```

- ①Since this is **pseudocode**, let's not worry too much about types for a moment. This operation won't compile without types, but that's fine for now.
- ②For each step that you want to divide your simulation into (think of city blocks in Phoenix, or pixels on a screen), you'll need to do a bit of each Hamiltonian term.
- **3** Within each timestep, then, you can loop over every term that you need to simulate, and simulate each for one step.

Thankfully, Q# provides a standard library function that does precisely this for you: DecomposedIntoTimeStepsCA. In listing 9.9, we show how calling DecomposedIntoTimeStepsCA makes it easy to use the Trotter—Suzuki decomposition to simulate evolution under Marie's Hamiltonian.

Listing 9.9. Using DecomposedIntoTimeStepsCA to simulate Marie's entire

Hamiltonian.

```
operation EvolveUnderHamiltonian(
   idxBondLength : Int,
   trotterStepSize : Double,
   qubits : Qubit[]
)
: Unit is Adj + Ctl {
   let trotterOrder = 1;
   let op = EvolveUnderHamiltonianTerm(idxBondLength, _, _, _);
   (DecomposedIntoTimeStepsCA ((5, op), trotterOrder))(trotterStepSize, qubits);

6
}
```

- $oldsymbol{1}$ Just like <code>EvolveUnderHamiltonianTerm</code>, <code>EvolveUnderHamiltonian</code> will apply the appropriate Hamiltonian based on the coefficients for the desired bond length of the H_2 molecule that Marie asked for your help with.
- 2 Again, you also will have a step size that represents how long you want to simulate Hamiltonian evolution for.
- 3 The DecomposedIntoTimeStepsCA function supports higher-order Trotter—Suzuki decompositions than the first-order approximation that you've explored so far in this Chapter (represented by a trotterOrder of 1). In some cases, this feature can be useful in helping increase the accuracy of your simulation, but a trotterOrder of 1 works fine for our purposes.

©Manning Publications Co. To comment go to liveBook

- 4 This line uses partial application to fix the <code>idxBondLength</code> input to

 EvolveUnderHamiltonianTerm, but that leaves the <code>idxTerm</code>, <code>stepSize</code>, and <code>qubits</code> arguments blank. This will be useful as an input for <code>DecomposedIntoTimeStepsCA</code> in the next line.
- The first input to this function is a tuple of the number of terms to be simulated and an operation that can simulate evolution under each term in turn. The second input is the approximation order that you defined above. This function outputs an operation that can be used to automatically simulate evolution under the entire Hamiltonian, using the operation that simulates each term one-by-one, so you can go on and apply it to your qubits.

9.7 PUTTING IT ALL TOGETHER



Now that you have a better idea about what Hamiltonians are and how we can simulate evolution under them to understand how quantum systems change in time, you're ready to put together a program that will help Marie solve her question. As a reminder, Marie is a chemist who studies the ground state energies (a.k.a. lowest possible energies) of different chemicals. She has asked us for help trying to figure out what the ground state energies for the H₂ molecule with our quantum device. Because the hydrogen atoms that make up H₂ molecules are also quantum systems, it is much easier to simulate the behavior of H₂ with qubits than a classical computer.

TIP Quantum computers are so well suited to simulating the behavior of other quantum systems, this was arguably the first application ever proposed for quantum computing!

is a reminder of all of the steps or techniques you have learned in this Chapter to simulate the evolution of Marie's H₂ molecule in your quantum device.

Figure 9.9. Overview of the steps you will develop in this chapter to help Marie learn the ground state energy of her molecule.

Step 1: Hamiltonian

Marie gives her problem to you as a *Hamiltonian*, a mathematical object that describes the energy levels of a quantum system.

Step 2: Decompose Terms

You then decompose Marie's Hamiltonian into the *Pauli basis*, just as you can break down a state into the computational basis.

Step 3: Simulate Terms

You next turn each term in the decomposed Hamiltonian into a multi-qubit rotation operation that's easy to run on a quantum computer.

Step 4: Trotter-Suzuki

To simulate the entire Hamiltonian that Marie gives you, you then recombine each multi-qubit rotation using the *Trotter–Suzuki decomposition*, rapidly alternating which rotation you apply to your qubits.

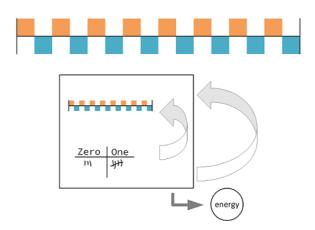
Step 5: Estimate Energy

To estimate the energy that Marie wants, you finish by playing the same phase estimation game that you helped Lancelot play in Chapter 8.

Н

$$H = \sum_{i} H_{i}$$

let term = HamiltonianTerm[idx]; Exp(term::Axis, term::Angle, targets);



So, you being a quantum developer, can collaborate with Marie to simulate the evolution of the H₂ molecule in time and calculate what the ground state energy is thanks to Schrödinger's equation. The key point to remember is that the energy levels that the H₂ molecule can be is correspond to the different *eigenstates* of the Hamiltonian.

Suppose that your qubits are in an eigenstate of the Hamiltonian. Then, simulating

evolution under that Hamiltonian won't change the state of your qubit register, *except* to apply a global phase that is proportional to the energy of that state. That energy tells you exactly what you need to solve Marie's problem, but global phases are unobservable. Thankfully, in the previous Chapter, you learned from Lancelot's and Dagonet's game all about how to turn global phases into something that you can learn with phase estimation — here is a great place to apply it! Summarizing, then, the steps that you need to do to collaborate with Marie and solve her problem are as follows:

9.7.1 Steps to find the ground state energy of H_2 .

- 1. Prepare the initial state that Marie gives you; in this case, she helpfully tells you to prepare |10⟩.
- 2. Take the Hamiltonian representing the system and break it up into little steps that can be simulated sequentially to represent the entire operation.
- 3. Apply each step representing the Hamiltonian to your initial state.
- 4. Use a phase estimation algorithm to learn about the accumulated global phase on your quantum state, which will be proportional to the energy.

You now have the skills and code from the previous sections in this chapter to pull this all together, so let's give it a go.

Starting with the Q# file (here called operations.qs to match what you've seen in previous Chapters), you can open some namespaces to utilize some pre-made functions and operations.

Listing 9.10. Namespaces needed from the QDK for helping Marie find the ground state

energy of H₂.

```
namespace HamiltonianSimulation {
    open Microsoft.Quantum.Intrinsic;
    open Microsoft.Quantum.Canon;
    open Microsoft.Quantum.Simulation;
    open Microsoft.Quantum.Characterization;
```

- **1** We have seen Microsoft.Quantum.Intrinsic and Microsoft.Quantum.Canon before, they have the basic utilities/helper functions and operations we will need.
- 2 Microsoft.Quantum.Simulation is a namespace for the QDK that has, as you might expect, utilities for simulating systems.
- Microsoft.Quantum.Characterization has easy to use implementations of the phase estimation algorithms that you developed in the previous chapter.

Next, you need to add the data Marie has about her molecule. You will find all of this typed out for you already in the sample file in the GitHub repo for this book: https://github.com/crazy4pi314/learn-qc-with-python-and-

<u>qsharp/blob/master/ch09/operations.qs</u>. The functions H2BondLengths, H2Coeff, and H2IdentityCoeff will be what you need (they are kinda long to reproduce here in

the text).

Once you have all of the coefficient data from Marie in your file, next you will need the actual terms/structure for the Hamiltonian that you will use with those coefficients you just added. In listing 9.11, you can see an outline of a function that returns terms of Marie's Hamiltonian expressed as Pauli operators, as well as an operation that will prepare your two-qubit register in the right state for this algorithm.

- 1 The H2Terms function makes it easy to construct the terms of Marie's Hamiltonian.
- 2 This function is really just a hard-coded list of tuples that describe the terms of the Hamiltonian. This first tuple says that the first term of the Hamiltonian is the PauliZ operation on the zeroth qubit.
- 3 This term of the Hamiltonian applies PauliZ to both the zeroth and the first qubit.
- (1) You will also need a way to prepare your qubits for the algorithm. Following Marie's advice, you put the first qubit in the |1) state, leaving the rest of the input qubits in the |0) state.

To take care of steps 2 and 3 of your quantum algorithm, you will need the operations you defined earlier EvolveUnderHamiltonianTerm and EvolveUnderHamiltonian.

Lastly, you can use the EstimateEnergy operation provided with the Quantum Development Kit to take the Trotter-Suzuki steps and automate their application as well as the phase estimation step. You can also use a built-in phase estimation operation that implements a better version of the phase estimation algorithm you learned about in the last Chapter.

GLOBAL PHASES AND THE CONTROLLED FUNCTOR Just as you saw in the previous Chapter, applying <code>EvolveUnderHamiltonian</code> doesn't do anything to qubits prepared in an eigenstate of Marie's Hamiltonian — indeed, that's the whole point! In Chapter 8, you were able to solve this by the using the <code>Controlled</code> functor to turn the global phase resulting from applying Dagonet's operation on a qubit prepared in an eigenstate into a local phase that can be observed, and by then using phase kickback to apply that phase to a control qubit. The <code>EstimateEnergy</code> operation provided by the Quantum Development Kit uses the exact same trick to learn what would otherwise be a global phase of your <code>EvolveUnderHamiltonian</code> operation. This means that it's critical that your operation supports the <code>Controlled</code> functor by adding is <code>Ctl</code> to the signature for each operation that you pass to <code>EstimateEnergy</code> as a part of helping out Marie.

Listing 9.12. Final Q# operation that can estimate the ground state energy of the H_2

- **1** This is it! The operation EstimateH2Energy takes the index of the molecule bond length and returns the energy of its ground state or lowest energy state.
- 2 You define that you will need two qubits to simulate this system.
- 3 Here you can set a scale parameter for the Trotter-Suzuki steps that will apply the terms of your Hamiltonian to your qubits.
- The trotterStep variable builds off of the operation you defined in the previous section,

 ApplyHamiltonian and gives a convenient name for the operation that will apply the terms of
 your Hamiltonian with your parameters for the Trotter—Suzuki steps.
- § You have built into the Microsoft.Quantum.Simulation library an operation that will take a specification for then number of qubits (nQubits), an operation to prepare your desired initial state (PrepareInitialState), how to apply your Hamiltonian (trotterStep), and what algorithm you want to use to estimate the phase resulting from applying your Hamiltonian (here you can leverage a built-in option provided by the Microsoft.Quantum.Characterization library RobustPhaseEstimation).
- **(6)** To make sure the units are right for the returned energy you have to divide by the trotter step size and add the energy from the identity term in the Hamiltonian.

Now, to actually run the algorithm! Since the ground state energy is a function of the molecule bond length, you can use a Python host to run the Q# algorithm and then plot the results as a function of the bond length.

Listing 9.13. Setting up the Python file for running your Hamiltonian simulation.

- **1** First, you need to import the Python package for Q# and then import the Q# namespace (HamiltonianSimulation) from your operations.qs file.
- 2 Just to make things easier, you can pull the list of bond lengths we can simulate for H₂ from the Q# function H2BondLengths.
- 3 The estimate_energy function is basically a Python wrapper for the EstimateH2Energy Q# operation, but that runs EstimateH2Energy a few times to make sure the energy estimated is actually the minimum.

Why do we need to run EstimateH2Energy multiple times?

The state $|01\rangle$ that Marie gave you isn't actually an eigenstate of any H_2 Hamiltonian, but something that she computed using an quantum chemistry approximation known as Hartree–Fock theory. Since quantum chemistry is her area of expertise, she's able to help out by providing you with approximations like that.

In practice, this means that when you run phase estimation using the tools provided by the Microsoft.Quantum.Characterization namespace, you aren't learning the energy of a particular eigenstate, but rather you're randomly projecting onto an eigenstate and learning its energy. Since your initial state is a pretty good approximation, most of the time you'll actually project onto the lowest energy state of Marie's Hamiltonian (that is, the ground state), but you can get unlucky and correctly learn the energy of the wrong eigenstate. Since you're looking for the smallest energy, running multiple times and taking the minimum makes it much more probable that you'll learn the energy that you want.

With everything setup in the Python host, all that's left is to write and run the main function.

Listing 9.14. The main program for your Hamiltonian simulation algorithm.

```
if __name__ == "__main__":
    import matplotlib.pyplot as plt

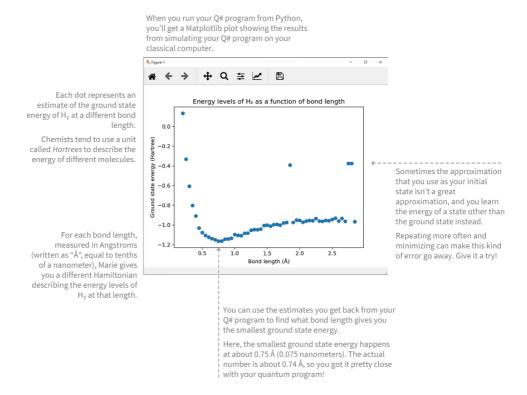
print(f"Number of bond lengths: {len(bond_lengths)}.\n")
    energies = [estimate_energy(i) for i in range(len(bond_lengths))]
    plt.figure()
    plt.plot(bond_lengths, energies, 'o')
    plt.title('Energy levels of H<sub>2</sub> as a function of bond length')
    plt.xlabel('Bond length (Å)')
    plt.ylabel('Ground state energy (Hartree)')
```

©Manning Publications Co. To comment go to <u>liveBook</u>

plt.show()

- Running host.py as a script will plot the estimated ground state energies from the quantum algorithm in Q#.
- 2 Here you are directly generating the list of estimated energies for each H₂ molecule bond length.
- 3 This and the next four lines setup the data and style for the plot.
- 4 Calling plt.show() should pop-up or return an image of the plot!

Figure 9.10. A sample of the plot that running your host.py should produce.



You can see in a sample of what running python host.py should output. This plot shows the results of your simulations for a variety of Hamiltonians for different bond lengths of the H₂ molecule. You can see that the energy of the lowest state is much higher when the bond length is short, and kind of levels off as the bonds get longer. The lowest energy possible should occur around a bond length of approximately 0.75 Å. If you look up, the stable (that is, equilibrium) bond length for hydrogen is 0.74 Å! So yeah, it turns out Marie's molecule is a pretty well known one, but you can see how you could follow this process not only for other chemicals, but also for simulating other quantum systems.

Congratulations, you have now implemented your first, practical application for a quantum computer! Of course, the actual chemical we were using here is fairly simple,

but this process holds for most other quantum systems you might want to simulate. In the next two chapters we will explore two other applications of quantum computers, unstructured searching with Grover's algorithm, and factoring numbers with Shor's algorithm.

9.8 SUMMARY

In this chapter you learned:

- Recognize how quantum computers can be used to solve difficult chemistry simulation problems
- Implement qubit rotation operations about arbitrary Pauli axes using Exp. operation,
- Implement time-dependent operations on qubits using the Trotter-Suzuki method, and
- Create programs using Q# standard library features for phase estimation, decomposition, etc.

Searching With Quantum Computers

This chapter covers:

- Implement a quantum algorithm for searching unstructured data.
- Recognize how the resource estimator for the QDK can help understand the costs for running particular algorithms.
- Implement reflections of quantum registers about states.

Figure 10.1. In this Chapter, we will be covering topics in both techniques for developing quantum algorithms and applications for quantum programs.

Quantum Applications	
Libraries / Techniques	
Q# Language	
Intrinsic Operations	
Quantum Hardware	Simulators

©Manning Publications Co. To comment go to liveBook

10.1 SEARCHING UNSTRUCTURED DATA

Suppose you want to search through some data, such as to find a contact's phone number. If that list of contacts is sorted by name, then it's pretty easy to find the phone number associated with a particular name by using a binary search:

Algorithm 10.1: Pseudocode for binary search

- 1. Pick a name/phone number pair in the middle of your list, call this pair your pivot.
- 2. If the pivot's name is the name you're looking for, return the pivot's phone number.
- 3. If the name you're looking for comes before the pivot's name, repeat the search on the first half of the list.
- 4. Else, if the name you're looking for comes after the pivot's name, repeat the search on the second half of the list.

Not just a character on Star Trek.

In this Chapter, we talk a lot about searching through data. That data can come in a lot of different forms.

Kinds of data

- Phone numbers
- Names of dogs
- Weather measurements
- Types of doorbells

What all of these have in common is that we can represent them on classical computers as strings of bits, using a variety of different conventions about how that representation should work.

Searching this way can all be done pretty quickly, and is key to how you can search databases full of information quickly. The problem is that in binary search, you critically depend on your list of names and phone numbers being sorted. If it's not sorted, binary searching simply doesn't work.

Put differently, to search your data quickly, you need some kind of structure to apply to your data, whether that structure is that your data is sorted, or some other kind of assumption that lets you avoid having to look at every single item. If you don't have any structure, the best thing you can do is to randomly look through your data until you find what you want. The steps listed in show pseudo code for how you would search a list without structure. You might get lucky, but on average, random searching is only ever twice as fast as looking at every single item.

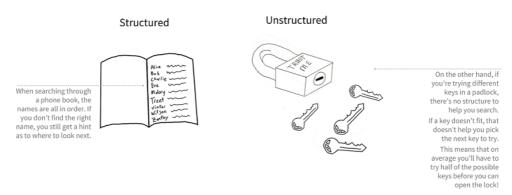
Algorithm 10.2: Pseudocode for searching unstructured lists

- 1. Pick a random element of your list.
- 2. If it's the right element, return it. Otherwise, pick a new element and repeat.

The fact that searching unstructured lists is hard is the basis of much of cryptography as well. In that case, rather than writing down the list explicitly, your task in trying to

break an encryption algorithm is to try different keys until one works. We can think of the decryption function as *implicitly* defining a list, where there's one special "marked item" that corresponds to the correct secret key.

Figure 10.2. Structured and unstructured searches



The pseudocode in could represent this decryption task. The random input you are picking is the key, which you use with the decryption "function" or algorithm to see if it decrypts the message.

Algorithm 10.3: Pseudocode for searching unstructured inputs to a function

- 1. Pick a random input.
- 2. Call your function with that input. If it worked, return your input.
- 3. Otherwise, pick a new random input and repeat.

If you could search unstructured lists more quickly, that would let you sort through databases, solve mathematical problems, or — yes — even break some kinds of classical encryption.

Perhaps surprisingly, if the function that defines your list can be written out as a quantum operation (using what you learned about oracles in Chapter 7), then you can use a quantum algorithm known as *Grover's algorithm* to find an input much faster than with .

NEED A REFRESHER?

You're getting towards the end of the book now, meaning you have the opportunity to put what you've learned throughout the book so far together. In particular, in this Chapter, you'll use what you learned about *oracles* from Nimue's and Merlin's game in Chapter 7 to represent the input to Grover's algorithm. If you need a refresher on what oracles are, no worries, Chapter 7 is there to help.

When you run Grover's algorithm, you're searching for one or more particular values of a function over all possible inputs to a function. If you want to search an unstructured list of data, you could consider defining a function that is responsible for looking up a particular entry in the list as the function that you want to search the

inputs to.

Consider the scenario where you need to decrypt a message in 1 minute. There are 2.5 million different keys you need to try, but only one works to decrypt the message, and trying each key at a time will take too long. You can use Grover's algorithm, and a function representing the problem like "does this cryptographic key decrypt a particular message?" to find the right key much quicker and without having to test each key individually! This is a lot like the padlock example from , where we think of the different possible keys as inputs.

We'll make the function you need for Grover's algorithm to represent this problem more precise when we revisit oracles later on the chapter, but it's really helpful to remember that when you search with Grover's algorithm, you're searching over *inputs to a function*, not a list of data. With that in mind, you can see the pseudocode for Grover's Algorithm in algorithm 10.4.

Algorithm 10.4: Pseudocode for performing unstructured search (Grover's Algorithm)

- 1. Allocate a register of qubits that is large enough to represent all of inputs to the function you are searching over.
- 2. Prepare the register in a uniform superposition state, i.e. all possible states have the same amplitude. This is because due to the type of problem, you don't have any additional information about which input might be the "correct" input, so this represents a uniform probability distribution or prior on the data.
- 3. Reflect the register about the marked state, or the state you are searching for. Here, a reflection just means picking out a particular state and flipping the sign on it, but you'll see more detail in the next Section, as well as how to implement reflections in Q#.
- 4. Reflect the register about the initial state (uniform superposition).
- 5. Repeat steps 3 and 4 enough times until the probability of measuring the item you are searching for is sufficiently high. Then measure the register. You can mathematically work out the optimal number of times you need to do this so that you maximize the correct answer.

You can see a figure representing these steps in figure 10.3.

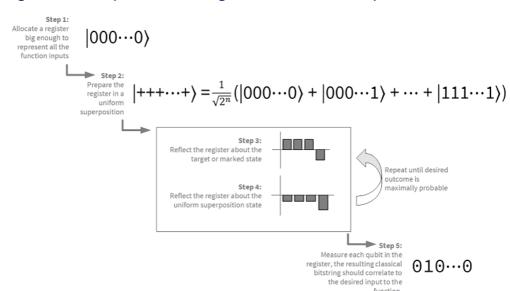


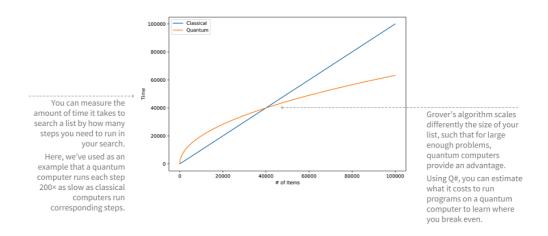
Figure 10.3. The steps to the Grover's Algorithm that searches over inputs to a function.

IMPORTANT

As you go through this Chapter, you'll see that one way to think of Grover's algorithm is as a kind of rotation between the states representing whether or not you've found the right marked item, a.k.a. the decryption key for our scenario. If you apply steps 3 and 4 above too many times, you'll rotate right past the state that you're looking for, so choosing the number of iterations is an important part of the algorithm!

In , we show an example of how the cost of classically searching an unstructured list might compare to using Grover's algorithm.

Figure 10.4. An example of how the time taken for searching unstructured lists might scale for classical and quantum computers



TIP One way to describe what's shown in figure 10.4 is to use a concept called **asymptotic complexity**. In particular, we say that a classical unstructured search requires O(N) function calls to search through N inputs, while Grover's algorithm requires $O(\sqrt{N})$ calls. Don't worry if this isn't familiar to you, but if you're curious to learn more about how to understand algorithms this way, check out Chapter 1 of *Grokking Algorithms* from Manning Publications.

As before, let's jump in and see what that code looks like. In listing 10.1, you can see an example of a Q# operation that uses Grover's algorithm to search an unstructured list for a marked item. Here, rather than the 2.5 million keys to try we will reduce the scope to 8 keys, and the marked item or correct key is indicated by an integer in the range 0 to 7. Yes, that would mean that maybe you could solve this just as fast on a classical computer. However, you will see at the end of the chapter how as the number of keys you need to search through grows, the number of steps or computations that are needed to find the key are much lower when using Grover's algorithm. Also, for our sample code here the function representing the decryption algorithm won't really do any decrypting, it will just act like it is playing a guessing game and return a boolean if it is given the right key. Implementing a particular decryption algorithm for this chapter is out of scope, and would likely need a bit of research to pull off. The goal here is to show how using Grover's to search through the inputs to a function allows for speedups to particular problems.

Listing 10.1. Operation that runs a specific example of Grover's Algorithm. operation RunGroverSearch(): Unit { let idxMarkedItem = 6; let markItem = ApplyOracle(idxMarkedItem, _, _); let foundItem = SearchForMarkedItem(3, markItem); Message(\$"marked {idxMarkedItem} and found {foundItem}."); 4

©Manning Publications Co. To comment go to <u>liveBook</u>

}

- 1) For this example, the item we are searching for is represented by the index 6.
- 2 You can use partial application to include the index of the marked item in the oracle that you will provide to the search algorithm.
- 3 Here you are actually running Grover's search algorithm on a register of three qubits and providing the oracle markItem that you defined above.
- 4 The last line here just issues a message to verify it found the right item.

If you run the sample in listing 10.1, you should get the output below ().

Listing 10.2. Output from running RunGroverSearch.

```
In [1]: %simulate RunGroverSearch
Out[1]: Marked 6 and found 6.
```

You can see from running the above sample of Grover's algorithm, that the decryption key you were looking for was the one indexed or marked by the number 6 and the algorithm found the key as number 6 as well.

Now, since the SearchForMarkedItem operation is really the meat of this example, let's look at it's implementation (see listing 10.3).

Listing 10.3. Writing out Grover's Algorithm as a Q# operation

```
operation SearchForMarkedItem(
    nQubits : Int,
    markItem : ((Qubit[], Qubit) => Unit is Adj)
3
)
: Int {
    using (qubits = Qubit[nQubits]) {
        PrepareInitialState(qubits);
        for (idxIteration in 0..NIterations(nQubits) - 1) {
            ReflectAboutMarkedState(markItem, qubits);
            ReflectAboutInitialState(PrepareInitialState, qubits);
        }
        return MeasureInteger(LittleEndian(qubits));
    }
}
```

- As usual, you'll want to start by defining a new operation, using the operation keyword.
- 2 The first input that your operation needs to take is the number of qubits that your list is defined on. Your quantum register needs to be big enough to index into that list. For example, if your list has 16 items, you need 4 qubits to get the 2^4 = 16 indices that you need to refer to each item in your list.
- 3 The next input is a representation of your search problem. Just as with the cryptography example above, you can implicitly define your search problem by an oracle that marks whether an item in your list is the right one or not.

- When you're done with your search, you'll have an index for where the marked item was. Defining your output as an Int lets you return that index back out.
- 5 To start the search, you'll need to allocate a big enough register to hold an index into your list.
- **6** Since you're searching an unstructured list, when you first start the search, all items are equally good places to look. You can represent this by preparing a uniform superposition over all indices into your list. You'll see more about how to do this below.
- The actual heart of Grover's algorithm comes down to repeatedly reflecting about your start state, and about the index for the item that you're looking for.
- ③Once you're done, measuring your qubit register tells you the index of the item that Grover's algorithm found. The Q# standard library provides a useful operation, MeasureInteger, that interprets measurement results as a classical integer. To use MeasureInteger, you can mark your register as encoding an integer in a little-endian register by using the

Microsoft.Quantum.Arithmetic.LittleEndian user-defined type.

USER-DEFINED TYPES

If you need a reminder about what user-defined types (UDTs) are and how to use them, check out Chapter 8, where you used them to help Lancelot and Dagonet play their angle-guessing game.

By this point in the book, you have almost all of the quantum concepts you need to make sense of listing 10.3. In the rest of this Chapter, you'll see how to use what you've learned to implement an example oracle that can be used to define a simple search problem, and how you can implement the two reflections that make up Grover's algorithm to solve that problem.

10.2 REFLECTING ABOUT STATES

You saw in and that there were two operations that you used repeatedly in the for loop, ReflectAboutInitialState and ReflectAboutMarkedState. Let's dig into how these operations helps us search the input of the function representing our decryption scenario.

Each of these operations is an example of a reflection about a particular state. This is an example of a new kind of quantum operation, but we can still simulate it with a unitary matrix as before. The term "reflecting about a state" means that when you have a register of qubits, you are picking out a particular state it could be in, and if it happens to be in that state, you flip the sign on it (a.k.a change the phase of that state). If you think this sounds like some of the controlled operations we looked at before, you are correct as you will be making use of controlled operations to implement these reflections.

10.3 REFLECTION ABOUT THE ALL-ONES STATE

Let's start by looking at a particularly useful example: a reflection about the all-ones state, $|11\cdots 1\rangle$. You can implement this reflection using the Controlled z operation that you first saw in Chapter 9.

CONTROLLED ISN'T JUST AN IF STATEMENT!

Remember that the <code>Controlled</code> functor isn't just a fancy <code>if</code> block, but rather can be used in superposition. For a review of how the <code>Controlled</code> functor works, check out Chapter 8, where you used it to help Lancelot and Dagonet play their game.

```
Listing 10.4. An operation that reflects about the |11\cdots 1\rangle state.
```

```
operation ReflectAboutAllOnes(register : Qubit[]) : Unit is Adj + Ctl {
   Controlled Z(Most(register), Tail(register));
}
```

There, the Controlled functor allows you to use the Z operation in a controlled fashion. Like other Controlled operations, Controlled Z takes two inputs: the first being the register that should be used as control qubits, and the second, the qubit that the Z operation will be applied to if all the qubits in the control register are in the |1) state. You can use the Most function from Microsoft.Quantum.Arrays to get all but the last qubit and Tail to get only the last qubit.

TIP

By using $Controlled\ Z$ together with Most and a Tail, your implementation works no matter how many qubits are in your register. This will be useful later, as you might need to different numbers of qubits to represent the data in your list.

Recall from Chapter 9 that the Controlled z operation applies a phase of -1 to the $|11\cdots 1\rangle$ state and does nothing to every other computational basis state. Thinking back to Chapter 2, where each computational basis state is a kind of direction, that means there's a single direction that gets flipped by Controlled z, while all other input states are left alone. That kind of picture is why we call operations that behave in this way *reflections*, though an actual graphical representation is tricky due to the number of dimensions that can be involved.

TIP

One way to see that the ${\tt Controlled}\ {\tt Z}$ operation flips the sign of a single input state, as described above, is to write down a unitary matrix that simulates ${\tt Controlled}\ {\tt Z}$. For example, with a single control qubit:

$$U_{\rm CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

For two control aubits:

Using what you've learned about unitary matrices throughout the rest of the book, these matrices make it clear that the input states |11\) and |111\), respectively, get flipped by - 1 while all other input states are left alone (get a phase of +1). The same pattern continues no matter how many control qubits you use with Controlled Z.

Exercise

Use DumpMachine to see how Controlled Z acts on the uniform superposition state |++...+>.

HINT: Recall that $|+\rangle = H|0\rangle$, so you can use the program `ApplyToEachCA(H, register)` to prepare $|+\cdots +\rangle$ on a register that starts off in the $|00\cdots 0\rangle$ state.

DEEP DIVE: Reflections are rotations?

Given our geometrical understanding of reflections, it would be natural to think of them as a kind of rotation by 180° degrees. That turns out to only be true because quantum states are vectors of complex numbers; if you only had access to real numbers, then you wouldn't be able to get a reflection by using rotations! You can see this by thinking back to how you described states and rotations in Chapters 2 and 3, namely as rotations of a 2D circle. If you pick up a two-dimensional object, flip it over, and put it back down, there's no way to transform it back to the way it was without picking it back up again. On the other hand, three-dimensional space gives you enough extra room that you can combine different rotations together to make a reflection. Since complex numbers give you your third axis when describing the states of qubits (namely, the Y-axis), that's what lets you do the reflections you need in Grover's algorithm as well!

10.4 REFLECTION ABOUT AN ARBITRARY STATE

Once you have a reflection about $|11\cdots 1\rangle$ under your belt, you can use that to reflect about other states as well. This is important as likely you will not be able to setup your oracle function representing the decryption algorithm so that the input or key you want is represented by the all ones input. Also recall for your sample code here, the oracle only implements a guessing game sort of decryption, and not any sort of real decryption algorithm.

The trick to reflecting about states other than all ones, is to turn whatever state you want to reflect about into the all-ones state, call ReflectAboutAllones, then undo the operation you used to map your reflection onto the all-ones state. You can describe any state by starting from the all-zeros state, so what you will need then is a way to go from the all-zeros to all-ones state where you can use the reflection you just learned. Check out for an example of how to prepare a register in the all-ones state from the all-zero state.

Listing 10.5. Preparing the all-ones state in a register.

```
operation PrepareAllOnes(register : Qubit[]) : Unit is Adj + Ctl {
    ApplyToEachCA(X, register);
}
```

① The ApplyToEachCA operation allows you to apply the first input (an operation) to each qubit in a register (the second input). Since registers are initially allocated in the |00 · · · 0⟩ state, applying X to every one should turn it into the |11 · · · 1⟩ state.

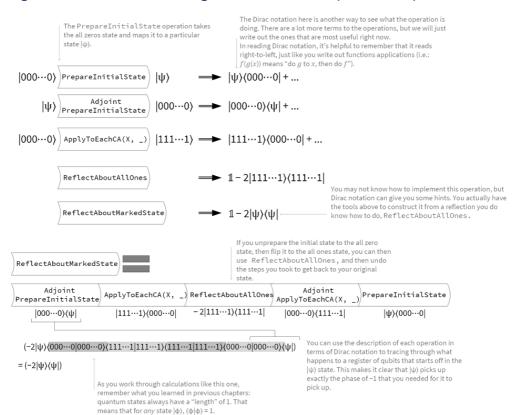


Figure 10.5. Reflections on reflecting about states with Dirac (a.k.a. bra-ket) notation.

For the next step, you need to consider the operation that prepares the state you want to reflect around. If you have an adjointable operation (is Adj) that prepares a particular state that you want to reflect about, then all you have to do is *unprepare* the state, prepare the all-ones state, reflect about the all-ones state ($|11 \cdots 1\rangle$), unprepare the all-ones state, and then *reprepare* the state you are trying to reflect about.

Why we love Dirac notation

In understanding the steps for , it is really helpful to think of what each operation in that sequence of steps does to its input state. Thankfully, Dirac notation (as you first encountered way back in Chapter 2) can be really helpful in writing out how unitary matrices transform different states, so that you can understand and predict what the corresponding Q# operations will do to your qubits.

For example, consider the Hadamard operation ${\tt H.}$ As you've seen throughout the book, ${\tt H}$ can be simulated by the unitary matrix

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1\\ 1 & -1 \end{pmatrix}.$$

```
This unitary matrix acts as a kind of truth table, telling you that H transforms the |0\rangle into the \frac{1}{\sqrt{2}} \frac{(|0\rangle + |1\rangle)}{(|0\rangle + |1\rangle)} state. Using Dirac notation, you can make this more clear by writing H = |+\rangle \langle 0| + |-\rangle \langle 1\rangle. By thinking of kets (|\cdot\rangle]) as denoting inputs and bras (\langle \cdot|), you can read this as saying "the `H` operation transforms |0\rangle into |+\rangle and |1\rangle." In figure 10.5, you can see how Dirac notation acts as a sort of visual language that tells you about the inputs and outputs for different unitary matrices, making it easier to understand how sequences of
```

Algorithm 10.5: How to reflect about an arbitrary state

O# operations work together.

- 1. Using the Adjoint functor, "un-prepare" your arbitrary state, mapping it to the all-zeros $|00\cdots 0\rangle$ state.
- 2. Prepare the all-ones state $|11\cdots 1\rangle$ state from the all-zeros state.
- 3. Use Controlled Z to reflect about $|11 \cdots 1\rangle$.
- 4. Un-prepare the all-ones state, mapping it back to the all-zeros state.
- 5. Prepare your state again, mapping the all-zeros state to your arbitrary state.

TIP

In , steps 1 and 5 cancel adjoints of each other, as are steps 2 and 4. Using what you learned about "shoes-and-socks" thinking, this makes the procedure of ideal for implementing using the within/apply feature of Q#! For a reminder of how this feature works, check out Chapter 7 where you used within/apply blocks to implement the Deutsch-Jozsa algorithm for Nimue and Merlin.

Since you don't have any prior notion as to what the right input to your oracle is, when you actually run Grover's algorithm, you'll want to start your search from the uniform superposition $|++\cdots+\rangle$ in order to represent that any input could be the right one. This gives you a chance to practice what you learned from to practice implementing reflections in Q#! Following the steps above to reflect about the initial state, you can implement the ReflectAboutInitialState operation that you used above in listing 10.3. In listing 10.6, we show how you can follow in a Q# operation.

operation PrepareInitialState(register : Qubit[]) : Unit is Adj + Ctl { ApplyToEachCA(H, register); } operation ReflectAboutInitialState(prepareInitialState : (Qubit[] => Unit is Adj), register : Qubit[]) : Unit { within { Adjoint prepareInitialState(register); PrepareAllOnes(register); } apply { ReflectAboutAllOnes(register); } }

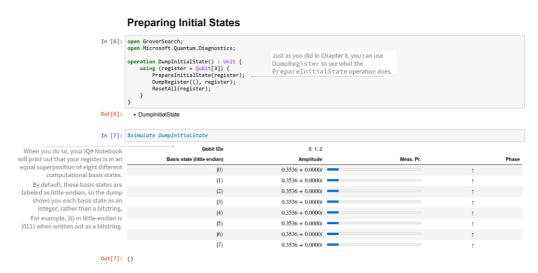
Listing 10.6. Reflecting about an arbitrary state, given a state preparation operation.

}

- ① Here, the uniform superposition state represents that you have no prior information for your search (after all, it is an unstructured search problem).
- 2 Following , to reflect about the initial state, you will need to provide an operation that prepares it.
- 3 Of course, you'll also need a register of qubits to apply your reflection to!
- 4 The within part of your within/apply block here does steps one and two from .
- The Adjoint functor here indicates you want to do the "reverse" or opposite of the operation that prepares the initial state. Another way to say that is if you started with the initial state, applying the Adjoint prepareInitialState operation would take you back to the $|00\cdots 0\rangle$ state.

You now have the code you need to reflect about that initial state, so how could you check to see if it did what you expect? When running the simulator target machine, you can use commands like <code>DumpRegister</code> that will then show you all of the information that it is using to simulate your register of qubits. You can see what the output of <code>DumpRegister</code> looks like after preparing a uniform superposition in .

Figure 10.6. Using DumpRegister to view the initial state prepared by your PrepareInitialState operation.



For reflecting about the marked state, the other reflection you'll need for Grover's algorithm, you'll need a slightly different approach. After all, you don't know how to prepare the marked state — that's the whole problem that you're using Grover's algorithm to solve! Thankfully, thinking back to Chapter 7, you can use what you learned from Nimue's and Merlin's game to implement a reflection about a state even when you don't know what that state is.

IMPORTANT

This is really the crux of Grover's algorithm: you can use your oracle to reflect about the marked state by calling it once with the right superposition of inputs. As you'll see in the next section, each of these reflections gives you some information about what the marked item is. By contrast, each classical function call can at most eliminate one possible input.

To see how that works in this case, let's first take a step back and look at what your marked state is. Since your list is defined by an oracle, you can write down a unitary operator O that lets you simulate that oracle. In Dirac notation:

$$O|x\rangle \otimes |y\rangle = |x\rangle \otimes \begin{cases} |\neg y\rangle & \text{xis the marked item} \\ |y\rangle & \text{otherwise} \end{cases}$$

In Chapter 7, you saw that applying an x operation to a qubit in the $|-\rangle$ applied a phase of -1, since $|-\rangle$ is an eigenstate of the x operation. Using the same trick here, we can write out what your oracle does when the flag qubit (the $|y\rangle$ register) is in the $|-\rangle$ state:

$$O|x\rangle \otimes |-\rangle = \begin{cases} -|x\rangle |-\rangle & \text{xis the marked item} \\ |x\rangle |-\rangle & \text{otherwise} \end{cases}$$

This is precisely the operation you need to implement the reflection! Thus, following what you learned in Chapter 7, you can implement it in the same way: simply apply your oracle to a qubit that starts off in the |-> state. This oracle then represents the decryption algorithm for the scenario, here simplified to basically a function that takes possible keys as input and just returns a bool indicating wether it is the correct key or not. We've shown an example Q# operation that uses this approach in listing 10.7.

```
Listing 10.7. Reflecting about a marked state.

operation ReflectAboutMarkedState(
    markedItemOracle : ((Qubit[], Qubit) => Unit is Adj),
    inputQubits : Qubit[])

: Unit is Adj {
    using (flag = Qubit()) {
        within {
            X(flag);
            H(flag);
        } apply{
            markedItemOracle(inputQubits, flag);
        }
    }
}
```

- ①Recall that in Q#, you can pass an operation as an input to another function or operation. Here, we've written out the operation for your item-marking oracle as having type ((Qubit[], Qubit) => Unit is Adj), indicating that it takes a register of qubits plus one additional qubit, and is adjointable.
- 2The second input you'll need to reflect about the marked state is a register that you want to apply your reflection to.

©Manning Publications Co. To comment go to <u>liveBook</u>

- 3 Once you have all your inputs, you'll need to allocate one additional qubit in order to apply your oracle. This qubit corresponds to the y register in the equations above. Often, a qubit that is used this way is called a "flag" qubit.
- In the same way that you used the H and X in Chapter 7 to prepare Nimue's qubit in the |-⟩ state, you can use that |-⟩ = HX|0⟩ here to prepare your flag qubit. Since this preparation is in a within/apply block, Q# will automatically put your qubit back into the |0⟩ for you by undoing the X and H operations. After all, as you saw in Chapter 7, applying your oracle leaves its target in the |-⟩ state
- **5** Finally, you can apply your oracle to use the Deutsch–Jozsa trick and apply a phase of -1 to the state marked by your oracle.

EXERCISE

In the above snippet, you can also write out H(flag); Z(flag);. Using either or both of QuTiP and AssertOperationsEqualReferenced, prove that these two ways of preparing your flag qubit give you the same reflection.

What's really amazing about this is that you were able to use the Deutsch-Jozsa trick to reflect about a state that was *implicitly* defined by your oracle! You didn't need to explicitly know what the marked state was in order to apply the reflection, perfect for use in an unstructured search.

In the next Section, you'll see how we combine the initial and marked state reflections to bring everything together and fully implement Grover's Algorithm, and find your key!

10.5 IMPLEMENTING GROVER'S SEARCH

Now that you have learned about rotating about states and revisited oracles it's time to put it all together to do some unstructured searching!

Let's start by reviewing all the steps we need to implement Grover's Algorithm in .

Steps for performing unstructured search (Grover's algorithm)

- 1. Allocate a register of qubits that is large enough to index the data set you are searching over.
- 2. Prepare the register in a uniform superposition state, i.e. all possible states have the same amplitude. This is because due to the type of problem, you don't have any additional information about the data set, so this represents a uniform probability distribution or prior on the data.
- 3. Reflect the register about the marked state, or the state you are searching for.
- 4. Reflect the register about the initial state (uniform superposition).
- 5. Repeat steps 3 and 4 enough times until the probability of measuring the item you are searching for is sufficiently high. Then measure the register. You can mathematically work out the optimal number of times you need to do this so that you maximize the probability of getting the marked item back.

You can see a figure representing these steps in

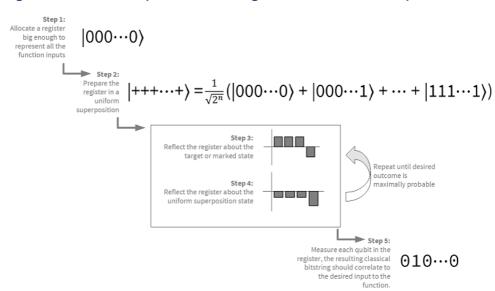


Figure 10.7. Recall the steps to the Grover's Algorithm that searches over inputs to a function.

IMPORTANT

If you go apply too many iterations of Grover's algorithm, then the amplitude of the state you want to measure will decrease. This is because each iteration is effectively a rotation; the trick is to stop that rotation at the right point. The trick to working out the trigonometry for the stopping criteria is to write the state of the register used in Grover's algorithm as a superposition of the unmarked and marked states. We won't cover the details of this derivation here, but please check

out https://docs.microsoft.com/quantum/libraries/standard/algorithms or Section 6.1.3 of Quantum Computation and Quantum Information by Nielsen and Chuang if you're interested in learning more about the math behind the scenes. It is already implemented for you in the samples repo for this book, but the formula is below if you want to try and program it in Q# yourself!

You have already implemented some of the operations you need for the full implementation in previous section. For example, you have implemented step 2 with the PrepareInitialState operation, and you've implemented the reflections in steps 3

the ReflectAboutMarkedState and ReflectAboutInitialState operations, respectively. What you will still need is a function to help you figure out how many times to loop over steps 3 and 4, as well as an implementation of the oracle that identifies if the item you are looking for.

Let's start with the function that helps define the stopping criteria for Grover's Algorithm, see listing 10.8.

Listing 10.8. Function that defines the stopping criteria for Grover's Algorithm.

function NIterations(nQubits : Int) : Int {

©Manning Publications Co. To comment go to <u>liveBook</u>

```
let nItems = 1 <<< nQubits;
let angle = ArcSin(1. / Sqrt(IntAsDouble(nItems)));
let nIterations = Round(0.25 * PI() / angle - 0.5);
return nIterations;
}</pre>
```

- 1 The <<< operator here is the left-bitshift operator, used to compute 2^{nQubits}, which represents the maximum number of items that can be indexed by a quantum register of length nQubits.
- 2) The next line figures out the effective rotation angle applied by each iteration of Grover's algorithm.
- 3 Using the effective rotation angle together with some trigonometry, you can figure out how many iterations you need to maximize the probability that you'll observe the marked item.

Now that you are able to calculate when to stop the loop in your implementation of Grover's Algorithm, the last thing you need is an oracle that can given the item you are looking for and a potential item from the data set, flip the phase on part of your register if the potential item is the item you are looking for.

For the purpose of an example, let's think of the oracle as representing a kind of guessing game. If someone is thinking of the number 4 and asks you to guess their number, that's an example of a kind of classical function:

$$f(x) = \begin{cases} 1 & x = 4 \\ 0 & \text{otherwise} \end{cases}$$

Classically, you'd have no better strategy than to try different inputs to funtil you tried x = 3. If you want to try using Grover's algorithm instead, using what you learned in Chapter 7, you know you'll need an operation that represents f:

$$U_{f}\left(\left|x\right\rangle \otimes\left|y\right\rangle \right)=\left|x\right\rangle \otimes \begin{cases} X\left|y\right\rangle & x=4\\ \left|y\right\rangle & \text{otherwise} \end{cases}$$

It's pretty easy to implement an operation that can be simulated by U f using the Q# function ControlledOnInt, provided as a part of the Q# standard libraries. Like the Controlled functor, the ControlledOnInt function allows you to control an operation on the state of another register. The difference while Controlled always all-ones controls the state |11 on the ControlledonInt function allows you to control on a different state, specified by an integer. For example, if Length (register) is 3, then (ControlledOnInt (4, flag) flips the state of flag whenever register is in the state |100), since 4 is written as 100 in little-endian notation.

EXERCISE

Try writing out what (ControlledOnInt(4, X)) (register, flag) does to the state of register + [flag], using either Dirac notation (check out Chapters 2 and 4 if you need a refresher) or by writing down a unitary matrix that can be used to simulate (ControlledOnInt(4, X)) acting on a three-qubit register and a flag qubit.

HINT: since (ControlledOnInt(4, X)) acts on four qubits in this example (three control qubits and a target qubit), your unitary matrix should be a 16x16 matrix.

Try doing the same, but for (ControlledOnInt(4, Z)).

Using the ControlledOnInt function, you can quickly write an oracle that flips the state of a flag qubit based on an input to that oracle, as shown in .

Listing 10.9. Operation that represents an oracle that marks the state you are looking

```
operation ApplyOracle(
   idxMarkedItem : Int,
   register : Qubit[],
   flag : Qubit)
: Unit is Adj + Ctl {
    (ControlledOnInt(idxMarkedItem, X))(register, flag);
}
```

- ①You can indicate the index of the item you are looking for as an integer with idxMarkedItem. Since your sample here will use three qubits, you can input any integer in the range 0 to 2³-1= 7.
- 2 Your oracle should flip its flag qubit whenever the input to the oracle is in the marked state, so you can use the ControlledOnInt function that you learned about above to apply an X controlled on the input register being in the right marked item.

With these two code snippets to add, you can return to the sample code you saw at the beginning, .

```
Listing 10.10. Writing out Grover's Algorithm as a Q# operation

operation SearchForMarkedItem(
    nQubits : Int,
    markItem : ((Qubit[], Qubit) => Unit is Adj)
)
: Int {
    using (qubits = Qubit[nQubits]) {
        PrepareInitialState(qubits);
        for (idxIteration in 0..NIterations(nQubits) - 1) {
            ReflectAboutMarkedState(markItem, qubits);
            ReflectAboutInitialState(PrepareInitialState, qubits);
        }

        return MeasureInteger(LittleEndian(qubits));
    }
}
```

©Manning Publications Co. To comment go to <u>liveBook</u>

- (1) As usual, you'll want to start by defining a new operation, using the operation keyword.
- The first input that your operation needs to take is the number of qubits that your list is defined on. Your quantum register needs to be big enough to index into that list. For example, if your list has 16 items, you need 4 qubits to get the $2^4 = 16$ indices that you need to refer to each item in your list.
- 3 The next input is the list you want to search. Just as with the cryptography example above, you can implicitly define your list by an *oracle* that marks whether an item in your list is the right one or not. You'll see more about this oracle in the rest of the Chapter.
- When you're done with your search, you'll have an index for where the marked item was. Defining your output as an Int lets you return that index back out.
- 5 To start the search, you'll need to allocate a big enough register to hold an index into your list.
- **6** Since you're searching an unstructured list, when you first start the search, all items are equally good places to look. You can represent this by preparing a uniform superposition over all indices into your list. You'll see more about how to do this below.
- The actual heart of Grover's algorithm comes down to repeatedly *reflecting* about your start state, and about the index for the item that you're looking for. In the next Section, you'll see more about what this means, and how you can implement each of the reflections you need.
- ③Once you're done, measuring your qubit register tells you the index of the item that Grover's algorithm found. The Q# standard library provides a useful operation, MeasureInteger, that interprets measurement results as a classical integer. To use MeasureInteger, you can mark your register as encoding an integer in a little-endian register by using the

Microsoft.Quantum.Arithmetic.LittleEndian user-defined type.

You have all the code you need now, so let's run an example, see listing 10.11.

Listing 10.11. Operation that runs a specific example of Grover's Algorithm.

```
operation RunGroverSearch() : Unit {
   let idxMarkedItem = 6;
   let markItem = ApplyOracle(idxMarkedItem, _, _);
   let foundItem = SearchForMarkedItem(3, markItem);
   Message($"marked {idxMarkedItem} and found {foundItem}.");
}
```

- 1) For this example, the item we are searching for is represented by the index 6.
- 2 You can use partial application to include the index of the marked item in the oracle that you will provide to the search algorithm.
- 3 Here you are actually running Grover's search algorithm on a register of three qubits and providing the oracle markItem that you defined above.
- 4 The last line here just issues a message to verify it found the right item.

If you run the sample then in you should get the output below ().

Listing 10.12. Output from running RunGroverSearch.

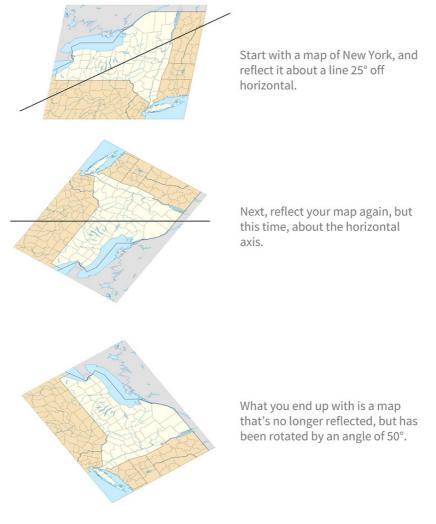
```
In [1]: %simulate RunGroverSearch
Out[1]: Marked 6 and found 6.
```

EXERCISE

Try changing the definition of the oracle to control on a different integer. Does this change the output when you run Grover's algorithm?

Congratulations, you can now use a quantum program to do unstructured searches! But what's actually going on here? The key insight from geometry that makes Grover's algorithm works is that when you reflect about two different axes, you get a rotation. In figure 10.8, you can see an example of how that works for maps.

Figure 10.8. How pairs of reflections can make a rotation



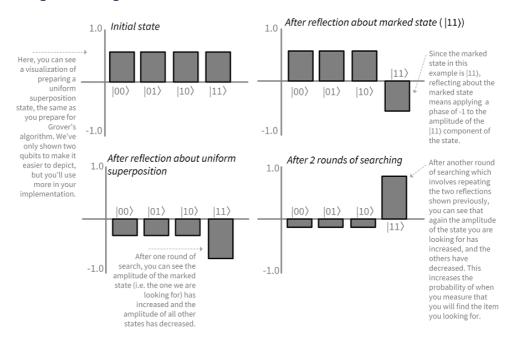
Map of New York @ NordNordWest, used under CC-BY-SA 3.0.

The same exact idea also works for quantum states. In Grover's algorithm, the initial

©Manning Publications Co. To comment go to liveBook

and marked state reflections combine into a single rotation from the unmarked states into the marked state. To understand how that works, you can use the techniques you've learned throughout the rest of the book to look at what's happening to the amplitudes of each state of the register as you go through the steps of the algorithm.

Figure 10.9. An schematic showing how the state you are searching over changes as you step through Grover's Algorithm.



What you can see from is that each round of reflections seems to *amplify* the the amplitude of the state that corresponds to the index you are looking for. By rotating between the unmarked and marked states, you can make the state of your qubits line up with the marked state that you want to find.

You can use the same kinds of ideas in other applications as well, as it turns out. Grover's algorithm is really an example of a more broad class of quantum algorithms that do what is called *amplitude amplification*.

EXERCISE

You've learned a lot about rotations so far in this book, so that can help a lot in understanding the rotation applied by each iteration of Grover's algorithm. Try modifying your implementation of Grover's algorithm to apply twice as many iterations, and use <code>DumpMachine</code> to look at the state that results. Does that seem like what you'd expect from applying a rotation twice?

More general amplitude amplification examples

Along with phase estimation, amplitude amplification is one of the most fundamental techniques used throughout quantum algorithms. In the 24 years since Grover's algorithm first introduced the concept of

amplitude amplification, a huge number of different variants have been developed to cover a wide range of different problems, such as when there are multiple marked items, when you want to optimize a function rather than find a marked item, or even for when you can only sometimes correctly prepare the initial state. A wide variety of these techniques are available under the Microsoft.Quantum.AmplitudeAmplification namespace in the Q# standard libraries, go take a look!

Before we close out this chapter, it will be helpful to have a short discussion on how doing searches like the one you just implemented on quantum hardware scales compared to using your classical hardware.

10.6 RESOURCE ESTIMATION

So how long does it take to run Grover's algorithm in practice? We mentioned before when we simplified our scenario from 2.5 million keys to 8 that there would be an advantage to using Grover's algorithm as the number of keys you needed to search through grew. That turns out to be a pretty complicated question — one could write several books on just that. In part, this is a complicated question as estimating resource requirements necessarily depends on a lot of different parts of your quantum computing stack.

For example, errors are common in quantum devices, so you need to use error correction to protect your computation as it runs. Which error correction method is used to protect your computation has a massive effect on what's required to run your program. There are entire conferences dedicated exclusively to finding better and better error correcting codes for just this reason.

Thankfully, Q# and the Quantum Development Kit provide us some of the tools we need to start getting a handle on what's required to actually run different quantum programs. Instead of running your program on a simulator that models how a real quantum computer would work, you can instead run your program on the *resources estimator*, which tells you how many of each kind of intrinsic operation you need to call, how many qubits your program needs, and how many of the quantum operations in your program can be called in parallel. Let's look at a small example, using what you learned about the Deutsch–Jozsa algorithm in Chapter 7.

©Manning Publications Co. To comment go to <u>liveBook</u>

```
using ((control, target) = (Qubit(), Qubit())) {
    H(control);

within {
    X(target);
    H(target);
} apply {
    oracle(control, target);
}

return MResetX(control) == One;
}

Out[2]: - CheckIfOracleIsBalanced
In [3]: operation RunDeutschJozsaAlgorithm() : Bool {
    return CheckIfOracleIsBalanced(ApplyNotOracle);
}
Out[3]: - RunDeutschJozsaAlgorithm
```

- This is the same ApplyNotOracle that you saw before, except that now it is using a within/apply flow instead.
- 2 Remember when using Q# Jupyter notebooks, you have to open namespaces in each cell you want to use them in.
- **3** The operation CheckIfOracleIsBalanced is the same as before, except that again a within/apply block has been used to replace the repeated H and X operations.
- In Q# notebooks, you will need an operation with no arguments to use with the %simulate or %estimate commands.

Figure 10.10. Output from running % estimate RunDeutschJozsaAlgorithm on the program in .

When you run the <code>%estimate</code> magic command in an IQ# Notebook, you'll get back a table much like the one shown above in . This table reports what the kinds of resources that the Quantum Development Kit estimated your program will take to run.

Table 10.1. Kinds of resources tracked by the %estimate magic command.

Resource kind	Description		
CNOT	How many times the CNOT operation is called.		
QubitCliffo rd	How many times the X, Y, Z, H, and S operations are called.		
R	How many times single-qubit rotation operations are called.		
Measure	How many measurement operations are called.		
T	How many times the T operation is called.		
Depth	How many T operations need to be called in a row on a single qubit.		
Width	How many qubits your program needs.		
Borrowed Width	How many qubits your program needs to be able to borrow; this is a more advanced technique than we'll cover in this book.		

TIP

You can also estimate resources from Python as well! Just use

the ${\tt estimate_resources}$ method instead of the ${\tt simulate}$ method that you learned about in previous Chapters.

As you can see from running <code>%estimate</code> in the notebook and , there are probably some categories that make sense like the width, how many measurements are made, and R for the number of single-qubit rotations are used. There are some others that are new, like counting ${\tt T}$ operations and depth. You have not seen ${\tt T}$ operations before, but they are just another kind of single qubit operation.

MEET MR.

Like most other operations you've seen so far in this book, the ${\scriptscriptstyle \mathbb{T}}$ operation can be simulated by a unitary matrix,

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & (1+i)/\sqrt{2} \end{pmatrix}$$

That is, T is a 45° ($\pi/4$) rotation about the Z-axis.

Another way of thinking of the T operation is as the fourth root of the Z operation that you've seen quite a bit of so far. Since 45°x4 = 180°, if you apply T four times in a row, that's an expensive way of applying the Z operation once.

In Q#, the T operation is made available as Microsoft.Quantum.Intrinsic.T, and has type $Qubit \Rightarrow Unit is Adj + Ctl$.

EXERCISE

Use <code>AssertOperationsEqualReferenced</code> to prove that applying the <code>T</code> operation four times does the same thing as applying <code>Z</code> once. There's another operation <code>S</code> that can be thought of as the square root of <code>Z</code> (a 90° rotation about the Z-axis); check that applying <code>T</code> twice is the same as applying <code>S</code> once.

What makes T operations somewhat special and thus worthy of focusing on so heavily when estimating resources is that they are expensive to use with error correcting methods that are needed when running on larger quantum devices. Most of the operations you've used so far are part of a group of operations called the *Clifford group*; operations in the Clifford group are easier to use with error correction. As noted before, getting into the details on error correction here is out of scope for this book, but in short, the more operations you have that are not Clifford operations, the harder it will be to implement your program on error-corrected hardware. Thus, it is important to count up the number of these "expensive" operations (like T) for running on our currently available hardware.

TIP

At a high level, the number of $\mathbb T$ operations that have to be applied in sequence (that is, that can't be run in parallel) is a pretty good approximation for how long a quantum program takes to run on error-corrected quantum computers. This is reported by the resources estimator as the "Depth" metric

So what are typical or exceptional values for the resources you can count with the resource estimator in Q#? For a simple program like RunDeutschJozsaAlgorithm, the resources required are very modest. Looking at , though, there's a lot of focus on the T operation, so let's delve in a little bit to see what that operation is and why it's so important to resource estimation. In , you can see the output from estimating the resources required to call the CCNOT operation that you learned about in Chapter 9.

EXERCISE

Why don't we need to reset the register of qubits allocated in ${\tt EstimateCcnotResources}$ as shown in ?

This output is a bit surprising, in that your tiny program required 10 CNOT operations, 5 single-qubit operations, and 7 T operations, even though you didn't call any of them directly. As it turns out, it's very difficult to apply operations like CCNOT directly in an error-corrected quantum program, so the Q# resources estimator first turns your program into some thing closer to what would actually be run on hardware, using calls to more basic operations like CNOT and T.

Figure 10.11. Output from estimating the resources required to call CCNOT.



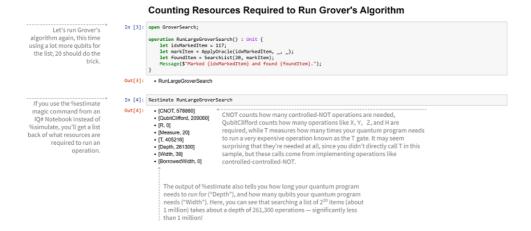
EXERCISE

How does the number of T operation calls change as you increase the number of control qubits? A rough trend is fine.

HINT: As you saw above, a controlled-NOT operation with an arbitrary number of qubits can be written as $Controlled\ X(Most(qs))$, Tail(qs));, using functions provided by the $Microsoft.Quantum.Arrays\ namespace.$

Where this comes in very handy is when you want to estimate what it would take for a program that is too large for you to simulate on a classical computer. In figure 10.12, you can see the output from running Grover's algorithm on a 20-qubit list (about 1 million items).

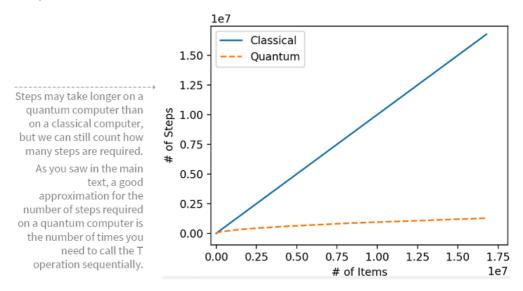
Figure 10.12. The result of running the resources estimator on Grover's algorithm.



If you run this for a variety of different list sizes, you get a curve like the one shown in figure 10.13. You can see that for our scenario with 2.5 million keys, the quantum number of steps is much much lower than the classical step cost. That's not the whole story, of course, since each step on a quantum computer will probably be much slower than a corresponding step on a classical computer, but it's a really good step towards

understanding what it would take to run different quantum programs in practice.

Figure 10.13. Output from estimating the resources required to run Grover's algorithm for a variety of different size lists.



10.7 SUMMARY

You have now learned how to combine the oracles you learned in chapter 7 and a new flavour of quantum operation (reflections) to be able to search over the inputs to functions. This was helpful in the scenario to help find decryption keys more quickly when there was a tight time constraint.

In the next Chapter, you'll use the skills you learned in this Chapter to answer one of the most important questions posed by quantum computing: how long would it take a quantum computer to break modern encryption?

In this chapter you learned:

- Implement a quantum algorithm for searching unstructured data.
- Recognize how the resource estimator for the QDK can help understand the costs for running particular algorithms.
- Implement reflections of quantum registers about states.

Arithmetic With Quantum Computers

This chapter covers:

- · Program with the Numerics library for Q#.
- Implement Shor's algorithm for factoring integers.
- Recognize the implications of quantum computing to security infrastructure.

Figure 11.1. In this Chapter, we will be covering an application for quantum computers, namely factoring numbers.

Quantum Applications			
Libraries / Techniques			
Q# Language			
Intrinsic Operations			
Quantum Hardware	Simulators		

11.1 FACTORING IN QUANTUM COMPUTING TO SECURITY

In Part I, you saw how quantum concepts can be applied to send data securely, using techniques such as quantum key distribution. Even without QKD, though, important data is shared secretly over the internet all the time. The internet is used to share payment data, personal health data, dating preferences, and even to organize political movements. In this Chapter, you'll take a look at how classical computers can protect your privacy, and how quantum computing changes how we make decisions about what tools we use to protect our data.

Starting off, let's take a look at the state of the art for securing data with classical computers. As it turns out, there's a lot of different problems in classical mathematics, some of which are really easy to solve (e.g.: "what's 2 + 2?"), while others are really hard to solve (e.g. "is P equal to NP?"). Between those two extremes, you get problems that are hard to solve unless someone gives you a hint, in which case they become easy. These problems tend to look more like *puzzles*, and can give us a great way to hide data: either you have to know the secret hint to solve them, or use huge amounts of computing time.

TIP

In Chapter 3, you saw quantum key distribution, a great way to share information securely that relies on quantum mechanics rather than on puzzles. You may not always be able to send qubits to your friends, though, so understanding how to use puzzles to communicate securely and privately still matters.

As you will see more later in the chapter, factoring numbers can be one of these puzzles that cryptographic algorithms can rely on for security. There are a number of very important algorithms/cryptographic protocols that are currently in use which rely on the seeming fact that it is hard for computers to solve puzzles that involve factoring large numbers. If you guessed quantum computers could help you factor large numbers, you're on the right track.

Enter Shor's algorithm. Using a classical computer, you can reduce the problem or puzzle of finding factors of integers to solving a kind of puzzle about how quickly functions repeat themselves when using a kind of arithmetic known as *modular arithmetic* (also known as "clock arithmetic," as we'll see more of later in the Chapter). If you use Shor's algorithm, estimating how quickly functions repeat themselves is precisely the kind of puzzle that you can solve easily on a quantum computer. Let's dig into the steps of Shor's and then look at an example of using it.

Example 11.1. Scenario: Factoring an integer N

Suppose that you're trying to factor the integer N, and that you know in advance that N has exactly two prime factors. Using Q#, implement Shor's algorithm to factor N.

COPRIME AND SEMIPRIME

As a helpful bit of terminology, we will say that two numbers which share no common factors other than 1 are coprime. For example, neither 15 nor 16 are prime, but 15 and 16 are coprime with respect to each other.

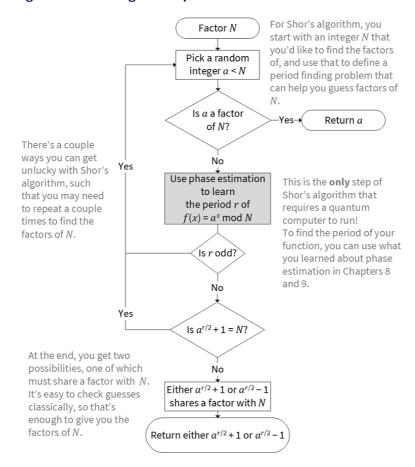
Similarly, we'll say that a number with exactly two prime factors is *semiprime*. For example, 15 is semiprime since it $15 = 3 \times 5$ and since both 3 and 5 are prime. On the other hand, 28 isn't semiprime since $28 = 4 \times 7 = 2 \times 2 \times 7$. Semiprime numbers come up often when considering cryptography, so it's often useful to make that assumption in your scenarios.

You can perform the steps in algorithm 11.1 (shown as a flowchart in figure 11.2) to use what you've learned about phase estimation in Chapters 8 and 9 together with some classical math to find the factors of N.

THE MOD SQUAD

In algorithm 11.1, you'll need one more bit of classical math, the mod operator. If you've not seen that operator before, don't worry, we'll go through that in a bit more detail later on in the Chapter.

Figure 11.2. Shor's algorithm represented as a flowchart.



Algorithm 11.1: Pseudocode for factoring an integer with Shor's algorithm.

- 1. Pick a random integer g, which we'll call the *generator*.
- 2. Check to see if the generator is accidentally a factor by checking if g and N are coprime. If they share a common factor, then you have a new factor of N, else continue with the rest of the algorithm.
- 3. Use iterative phase estimation to find the *frequency* of the classical function $f(x) = g^x \mod N$. The frequency tells you about how quickly f returns to the same value as x increases.
- 4. Use a classical algorithm known as the *continued fractions expansion* to convert the frequency from the previous step into a period (r). The period r should then have the property that f(x) = f(x + r) for all inputs x.
- 5. If the period r that you find is odd, go back to step one and make a new guess. If r is even, go to the next step.
- 6. Either $g^{\frac{z}{2}} 1$ or $g^{\frac{z}{2}} + 1$ shares a factor with N.

NOT ALL STEPS ARE QUANTUM

In algorithm 11.1, it is important to note that only step 3 involves any quantum computation. Most of the steps for Shor's algorithm are best suited for classical hardware, and are a good demonstration of how quantum hardware will likely be used. That is to say, quantum hardware and algorithms work well as *subroutines* for combined quantum-classical algorithms.

Now that you have seen the steps for Shor's algorithm, take a look at what the final implementation might look like in listing 11.2. The operation FactorSemiprimeInteger is the entry point to the algorithm and takes as input the integer that you want to factor, and returns it's two factors.

Listing 11.2. Operation that runs Shor's Algorithm on numbers that have only 2 prime

```
factors (other than 1 and itself).
operation FactorSemiprimeInteger(number : Int) : (Int, Int) {
    if (number % 2 == 0) {
        Message("An even number has been given; 2 is a factor.");
        return (number / 2, 2);
    mutable factors = (1, 1);
    mutable foundFactors = false;
    repeat {
        let generator = RandomInt(number - 2) + 1;
        if (IsCoprimeI(generator, number)) {
            Message($"Estimating period of {generator}...");
            let period = EstimatePeriod(generator, number);
            set (foundFactors, factors) = MaybeFactorsFromPeriod(
                generator, period, number
            );
        } else {
            let gcd = GreatestCommonDivisorI(number, generator);
            Message(
                $"We have guessed a divisor of {number} to be " +
```

```
$"{gcd} by accident. Nothing left to do."
);
set foundFactors = true;
set factors = (gcd, number / gcd);
}
until (foundFactors)
fixup {
    Message(
        "The estimated period did not yield a valid factor, " +
        "trying again."
);
}
return factors;

$
fixup {
    Message(
        "The estimated period did not yield a valid factor, " +
        "trying again."
);
}
```

- Before anything else, you'll want to check if you were asked to factor an even number, since then two must be a factor.
- 2 Following step 1 of algorithm 11.1, you start by picking a random number to define the periodic function that we'll use to factor number. If you got really lucky, then generator is already a factor, and we won't need a quantum computer to help us factor.
- 3In this Chapter, you'll learn how to write an EstimatePeriod operation to handle steps 3 and 4 of using what you learned about phase estimation.
- ①Once you have the period of your generator, you can use steps 5 and 6 of to turn that period into guesses as to the factors of number. You'll see how to write the MaybeFactorsFromPeriod function later in this Chapter.
- **6** If something went wrong (e.g.: your generator had an odd period), you can use a repeat/until loop to try again.
- **6** At the end, you can return out the two factors of number that you found using your quantum program!

Since this is the last Chapter of the book, you actually have all of the *quantum* concepts you need to know to understand what's happening in listing 11.2; all that's missing are the classical parts that connect what you've learned so far in the book to the problem of factoring semiprime numbers, as well as a few useful parts of the Q# libraries that can help you out. As mentioned before, there is only one step here that leverages the quantum technology and it does so by creating an oracle that implements the classical function that you want to learn about. Then by using a superposition state, applying the oracle and doing phase estimation then you can learn properties about the classical function, here the period. In the rest of the Chapter, you'll go through in detail, as well covering the last pieces that you need to run Shor's algorithm. The first piece that you'll need to work through is a bit of classical math known as *modular arithmetic*, so let's jump on in!

11.2 CONNECTING MODULAR MATH TO FACTORING

One place to find puzzles that can be used in security contexts is by looking how *modular arithmetic* works. Unlike normal arithmetic, in modular arithmetic, everything wraps back around like hours on a clock. For instance, if someone asks you what's two hours after 11 o'clock, you would get a very odd look if you responded

293

with "13 o'clock." Much more likely, this person was hoping to get an answer like "1 o'clock" instead! That is, if they don't do 24 hour time.

Using modular arithmetic, you can capture this idea by saying that $11 + 2 = 1 \mod 12$. In that equation, the mod 12 indicates that you want anything which goes past 12 to wrap around.

Figure 11.3. Using clocks to understand modular arithmetic.

Suppose that at 11 o'clock, someone asks you what time it will be in two hours.

If you said 13 o'clock, you'd get some strange looks, but if you said **1 o'clock**, that would be a lot more helpful!

You can think of (11 + 2) as having "wrapped around" the 12 on a clock; this is an example of **modular arithmetic**.



Of course, you don't have to do modular arithmetic only with 12-hour clocks. If clocks had 21 hours, for instance, then using modular arithmetic, $5^2 = 25 = 4$.



When arithmetic is allowed to wrap around like this, it can make it hard to figure out where different arithmetic calculations started out. If you're working with ordinary real numbers, for example, it's easy to calculate b if you're given a and a^b; you can take the logarithm of a^b to find b. If you try to solve the same problem in modular arithmetic, it can quickly get tricky. For instance, when computed \mod 21, the powers of 5 are 1, 5, 4, 20, 16, 17, 1, ... On the face of it, 5, 4 and 16 don't seem like powers of the same number, much less in increasing order, such that you have to check a lot more of a 21-hour clockface to work backwards from having taken an exponent.

Exercise 11.1

What are the powers of 11 when computed mod 21? How long does it take to loop back around to $11^0 = 1$?

Does it matter if you take the modulus by 21 at the end, or whether you compute the modulus at each step?

HINT: Either Python or O# works great for this, as both define the modulus operator %.

EXERCISE SOLUTIONS

All solutions for exercises in this book can be found in the companion code repo: https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp

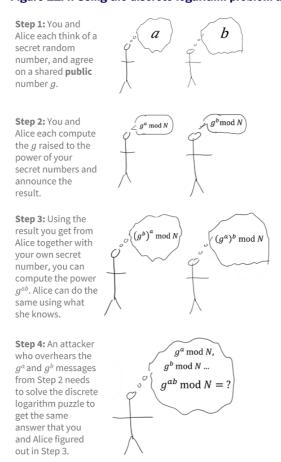
The observation that finding the power b given a^b mod N is hard already gives you a puzzle you can use to hide some data! This puzzle is commonly called the *discrete logarithm* problem. If Alice wants to share a secret with you, you can start by publicly agreeing on a small number like g = 13 and a big number like N = 71. You then each pick a secret number at random; suppose Alice picks a = 4 and you pick b = 5. Alice then sends you $a^b = 19$ and you send back $a^b = 19$ mod $a^b = 19$ mod $a^b = 19$ and Alice computes $a^b = 19$ mod $a^b =$

DON'T TRY THIS AT HOME

There's a lot of technical conditions to this protocol that go well beyond the scope of this book. Choosing g and N badly can undermine any security offered by this technique, making it trivial for an experienced attacker to break. It's also very easy to introduce bugs by doing this on your own, so please only consider this as a conceptual example!

If you're interested in learning more about the practical aspects of using these kinds of puzzles to keep your data secure, *Cryptography Engineering* by Niels Ferguson, Bruce Schneier, and Tadayoshi Kohno is a great book to continue learning from.

Figure 11.4. Using the discrete logarithm problem as a puzzle to hide secret messages.



Unlike QKD, this way of sharing secret data (known as the *Diffie–Hellman protocol*) relies on the assumption that the puzzle you and Alice used is hard to solve without a hint that your eavesdropper doesn't have access to. If someone can efficiently work out puzzles like solving g^a mod N for a given g and N, your data might as well be public.

One other puzzle that's commonly used to protect data today is called the RSA algorithm, and uses a bit more advanced classical math to make a puzzle out of

factoring really large integers. Just as you can break Diffie–Hellman by solving g^a mod N, you can break RSA by solving N = pq for p or q given only N. In the RSA puzzle, we call N a public key, and the factors p and q the private key; if you can factor N easily, you can get access to private keys just given public keys. With that in mind, you can make the scenario from earlier just a bit more precise.

Example 11.3. Scenario: Breaking RSA

Suppose that you know a public key N. Using Q#, implement Shor's algorithm to factor N to recover the private keys p and q.

Break it 'til you make it

The scenario above may seem to be a bit... nefarious when compared to those in previous Chapters. In practice, though, it's essential to understand attacks on the tools and protocols that you use to keep your data safe so that you can adjust your approach accordingly. If you use a cryptographic algorithm like RSA to protect data, implementing quantum attacks on that algorithm can help you understand how large a quantum device your attackers would need to compromise your data. After all, there's claims ranging from that there's no problem with RSA at all to that you should be immediately terrified now; separating between those extremes requires understanding the resources an attacker would need.

Put differently, by exploring a scenario that puts into the role of an attacker for a moment, we can understand how much quantum computing power would be needed to successfully attack RSA. We'll come back to this point at the end of the Chapter, but for now, it helps to think like an attacker. Understanding how classical cryptography could be attacked by quantum computers serves as a great scenario for you to practice applying your quantum computing skills!

It's worth being a bit cautious in using this example, though. The impact that quantum computing has on information security depends on what assumptions you make about classical algorithms, improvements in quantum algorithms, progress of quantum hardware development, how long you need forward secrecy to last for, and many other such concerns. Covering all of these well enough to make **responsible** decisions about how to best deploy cryptography would take more than is left in this book, unfortunately, so we recommend keeping in mind that the RSA scenario in this Chapter is an example, and not a complete analysis of the subject.

It turns out, though, that even though Diffie—Hellman and RSA look very different, you can use some classical math to turn the factoring puzzle of RSA into another example of figuring out how quickly you move around a clock face when you perform modular arithmetic (the Diffie—Hellman puzzle). That problem can then be solved easily on a quantum computer using what you learned in Chapter 9. To get a handle on how that all works, let's run through a quick example of using Shor's algorithm to factor a small integer so that you can see all of those parts at work.

11.3 EXAMPLE OF FACTORING WITH SHOR'S ALGORITHM

The steps listed for Shor's algorithm can seem very abstract, so before we get into how they all work, let's try working through an example using what you learned about modular arithmetic above. Say the number you want to factor is 21; real RSA public keys will be much larger, but let's use it for the sake of working through the math by hand. Trust us — it keeps the math much *much* easier.

11.3.1 Working through the steps of to factor 21.

- 1. Choose a random integer as a generator, let's say you got 11.
- 2. You can verify that that since 11 shares no common factors with 21, you can use it as the generator for the next step.
- 3. You can't do the quantum step in your head unfortunately, so you use the Q# iterative phase estimation operation to estimate the phase generated by applying an oracle that implements the classical function $f(x) = 11^x \mod 21$. It returns a phase Φ that you can convert to a frequency of 427 by the following

equation:
$$\frac{(\phi*2^9)/2)}{\pi}$$

- 4. You use the continued fractions algorithm on 427 to get a guess for what the period might be. Doing it by hand you get an estimate of 6 for the period.
- 5. The period you found is even so you can go to the next step.
- 6. Using the period 6 gives you that either 11^3 1 mod 21 = 7 or $11^3 + 1$ mod 21 = 9 shares a factor with 21. You can check and verify each possibility to confirm that 7 is indeed a factor of 21.

Exercise 11.2

Try step 6 from the process as above, but using 35 as the number to factor, 17 as your generator, and 12 as your period. Check that either or both of the answers that you get from step 6 share a common factor with 35.

Using either Python or O#, try the same with N = 143, g = 19, and the period r = 60.

NOTE: In the next section, you'll see how you can easily use a classical computer to factor a number when given another number that shares some of its factors.

While this is a lot of work to go through to factor a number as small as 21, 35, or 143, the exact same process also works for much larger integers, such as those that you might encounter trying to solve the puzzle that the RSA algorithm uses to protect data.

The rest of the Chapter will go through each of these steps in detail and show how they can work together to factor integers. To kick that process off, let's look at the classical math behind how period finding helps you factors integers, and how you can use Q# to implement that classical math.

11.4 CLASSICAL ALGEBRA AND FACTORING

With the concrete example of using Shor's algorithm above in mind, you can see how classical arithmetic and algebra help to take advantage of quantum computing. Before going over the core quantum part of the algorithm, it's helpful to explore the classical part a little bit more to understand why finding a the period of generator helps factor integers.

You may remember from algebra that for any number x, $x^2 - 1 = (x + 1)(x - 1)$. As it turns out, that works in modular (clock) arithmetic as well. If you find that the

period r of your generator g is even, - then that means there's an integer k - such that $g^r = g^{2k} \ mod \ N = 1$.

Subtracting one from each side, you get $(g^{2k} - 1) \mod N = 0$, - so that using $x^2 - 1 = (x + 1)(x - 1)$ gives you - $(g^k + 1)(g^k - 1) \mod N = 0$.

Why does this matter, though? If you have that $x \mod N = 0$, that tells you that x is a multiple of N. Thinking back to the clock analogy, 0, 12, 24, 36, and so forth are all equal to zero mod-12. Put differently, if $x \mod N = 0$, then there's some integer y such that x = yN. Using that with what you got from the period above, you know that there's some integer y such that $(g^k + 1)(g^k - 1) = yN$. If either $g^k - 1$ or $g^k + 1$ is a multiple of N, you haven't learned much, but in any other case, that tells you that either $g^k - 1$ or $g^k + 1$ has to share a factor with N.

To figure out whether g^k - 1 or g^k + 1 shares a factor with N, you can compute the *greatest common divisor* (also known as the GCD) of each guess with N. This is really straightforward to do with a classical computer using a technique called Euclid's algorithm.

NOTE

Since the GCD is so easy to compute classically, why do you need a quantum computer to help factor? By this point in Shor's algorithm, you've already narrowed down to two very good guesses as to potential factors, and are using the GCD just on those guesses. If you didn't have things narrowed down so well, you'd have to use the GCD on many, many more guesses to have a good chance at finding the factors of N. Even as easy as the GCD is to call, you still need a good way to narrow down to a good set of guesses first.

In Q#, you can compute the GCD using the GreatestCommonDivisorI function, as shown in where the code was run in a Q# Jupyter notebook.

```
Listing 11.4. Using GreatestCommonDivisorI to find the GCD between two integers.

In [1]: open Microsoft.Quantum.Math; open Microsoft.Quantum.Diagnostics;

function GcdExample(): Unit {
    let a = 2 * 3 * 113;
    let b = 2 * 3 * 5 * 13;
    let gcd = GreatestCommonDivisorI(a, b);
    Message($"The GCD of {a} and {b} is {gcd}.");

    EqualityFactI(gcd, 6, "Got the wrong GCD.");
}
Out[1]: - GcdExample
In [2]: %simulate GcdExample
The GCD of 678 and 390 is 6.
Out[2]: ()
```

①As usual, you'll need to open the right namespaces for the functions and operations that you want to use. In this case, the Q# function that computes the GCD of two integers is in the Microsoft.Quantum.Math namespace, so you can start by opening that namespace to make that functionality available.

- 2 Similarly, facts and assertions can be used by opening the Microsoft.Quantum.Diagnostics namespace.
- 3 Next, you can write a simple test case to see how the GCD works. For instance, 2, 3, 5, 13 and 113 are all prime numbers, so the only prime factors in common between 2 x 3 x 113 and 2 x 3 x 5 x 13 are 2 and 3. Thus, you would expect that the GCD of these two numbers is 2 x 3 = 6.
- **4** To compute the GCD, you can call GreatestCommonDivisorI from the Microsoft.Quantum.Math namespace that you opened earlier.
- **§** You can use the EqualityFactI function to confirm that the answer you got matches what you expected (namely, 6).
- Sa usual, you can use the IQ# magic command %simulate to run a function or operation on a simulator and print its output back into your notebook. Here, for instance, you get the output () back since GcdExample returns an output of type Unit.

Note the I at the end of the name GreatestCommonDivisorI. This tells you that GreatestCommonDivisorI works on inputs of type Int. When using Shor's algorithm in practice, N will be much, much larger than you can fit into an ordinary Int value, so Q# also provides another type called BigInt to help out.

To work with <code>BigInt</code> inputs, <code>Q#</code> also provides the <code>GreatestCommonDivisorL</code> function. Why <code>L</code> and not <code>B</code>? In this case, <code>L</code> stands for "long," helping disambiguate from other types starting with "B." such as <code>Bool</code>.

This convention is used throughout the rest of the Q# standard libraries as well. For instance, the equality fact that you used above compared two integers, and so was called EqualityFactI. The corresponding fact for comparing two big integers is called EqualityFactL, while the fact for comparing two Result values is called EqualityFactR.

Exercise 11.3

What's the GCD of 35 and 30? Does that help you find the factors of 35?

HINT: Think of this as step two of the previous exercise.

Putting everything together, if you have the period of your generator, let's take a look at for how you can use that to write out MaybeFactorsFromPeriod in Q#. The function name starts with maybe because there is a chance that the period that was found will not meet conditions necessary to learn something about the factors of the number.

Listing 11.5. Computing possible factors from a period.

```
let halfPower = ExpModI(generator, period / 2, number);

if (halfPower != number - 1) {
    let factor = MaxI(
        GreatestCommonDivisorI(halfPower - 1, number),
        GreatestCommonDivisorI(halfPower + 1, number)
    );
    return (true, (factor, number / factor));
} else {
    return (false, (1, 1));
}
else {
    return (false, (1, 1));
}
```

- 1 To compute possible factors from a period, you'll not only need to know the number N that you're trying to factor and the period r, but also the generator. Thus, your function will need to take all three as inputs.
- 2 If you got the unlucky case that either $g^{r/2} + 1$ or $g^{r/2} 1$ is a multiple of N, you won't be able to find any factors from period and will need to try another generator. To indicate that, your function can return a Bool indicating whether you successfully found a factor as well as the factors that you found if you were successful.
- 3 If the period is odd, then you can't use the $x^2 1 = (x + 1)(x 1)$ trick from above, and so you'll need another generator. Thus, you'll want to start by checking that period is even.
- **4** The Q# function Microsoft.Quantum.Math.ExpModI returns modular arithmetic exponentials of the form g^x mod N. You can use that to find $g^{r/2}$ mod N given g, r, and N.
- **⑤** You can check for the unlucky case that $g^{r/2} + 1$ is a multiple of N once you have computed $g^{r/2} \mod N$. If that check passes, then you know you have what you need to find an actual factor, so it's safe to continue.
- 6 As described above, the GCD tells you if one of your guesses has a common factor with N. If your guess has no common factors with N, though, we say that it is coprime with N, and the GCD will return 1. Using Microsoft.Quantum.Math.MaxI, you can try both guesses and take whichever guess gives you something other than 1.

Now that you know how to convert a period to potential factors, let's see the core of Shor's algorithm: using phase estimation to estimate the period of your generator. To do so, you'll use what you've learned throughout the rest of the book together with a couple new Q# operations to do arithmetic on a quantum computer, so let's jump in!

DEEP DIVE: Here's looking at Euclid.

Above, we used the <code>GreatestCommonDivisorI</code> function provided with the Q# standard libraries to compute the GCD of two integers. This function works by using Euclid's algorithm, which recursively attempts to divide an integer into another integer until no remainder is left over.

Suppose that you want to find the GCD of two integers a and b. Then you begin Euclid's algorithm by finding two additional integers q and r (short for "quotient" and "remainder") such that a = qb + r. It's straightforward to find q and r using integer division instructions, so this step isn't too hard on a classical computer. At that point, if r = 0, you're done: b is a divisor both of a and of itself, so there can't possibly be a larger common divisor. If not, you know that the GCD of a and b has to also be a divisor of r, so you

can recurse by finding the GCD of b and r instead. This process has to end eventually, since the integers you're looking for the GCD of get smaller and smaller as you go on, but never become negative.

To make things a bit more concrete, you can work through the example from listing 11.4, as shown in table 11.1.

Table 11.1. Using Euclid's algorithm to find the GCD of 678 and 390.

	b	r
78	390	288
90	288	102
88	102	84
02	84	18
4	18	12
8	12	6
	6	0
2	(answer)	(done)

11.5 QUANTUM ARITHMETIC

You have seen quite a few different parts of the Q# standard libraries by now, and with the focus of this chapter on arithmetic, it makes sense to introduce some functions and operations from the Microsoft.Quantum.Arithmetic namespace provided by the numerics library for Q#. This namespace, as you might guess, provides you with lots of helpful functions, operations, and types that make doing arithmetic in quantum systems easier. In particular, you can use implementations for things like adding and multiplying numbers represented in qubit registers with support for multiple qubit register encodings like BigEndian (where the least significant bit is on the left) and LittleEndian (where the least significant bit is on the right).

Let's look at some example code using the Q# numerics library to add two integers.

NOTE

There is a Q# notebook in the samples repo (https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp) that has all of these snippets already written out for you!

First, since the numerics package is not loaded by default, you will need to ask the Q# kernel to load it with the magic command %package. To make things easier, you can also turn off displaying small amplitudes from diagnostic outputs like DumpMachine, as shown in .

Listing 11.6. Loading the numerics package in a Q# notebook, and setting a display

preference for the diagnostic information. Your version numbers will likely vary.

- The %package magic command adds a new package into your IQ# session, making the functions, operations, and user-defined types implemented by that package available to you in your session. In this case, you can use the %package magic to load the Microsoft.Quantum.Numerics package, which provides additional operations and functions for working with numbers represented by registers of qubits.
- 2 After running %package, IQ# will report what packages are currently available in your IQ# session.
- The <code>%config</code> magic command sets various preferences for your current IQ# session. Here, for example, you can use <code>%config</code> to tell the <code>DumpRegister</code> and <code>DumpMachine</code> callables to leave out parts of each state vector that have very small amplitudes. That makes it much easier to visualize states on several qubits, as printing out each computational basis state can get unwieldy quickly.

TIP

When you're working with Q# from within Jupyter Notebooks, the IQ# kernel provides several other magic commands to help you write your quantum programs, in addition to the commands like simulate, package, and config that you've seen so far. For a complete list, check out the documentation

at https://docs.microsoft.com/gsharp/api/igsharp-magic/

11.6 ADDING WITH QUBITS

Now, let's get to coding an example that takes two integers and adds them while they are encoded in qubit registers. In listing 11.7, you can see an example of using the AddI operation to add the contents of two quantum registers to each other.

Listing 11.7. Example using the numerics library to add integers encoded in qubit

registers. In [3]: open Microsoft.Quantum.Arithmetic; open Microsoft.Quantum.Diagnostics; open Microsoft.Quantum.Math; operation AddCustom(num1 : Int, num2 : Int) : Int { let bitSize = BitSizeI(MaxI((num1, num2))) + 1; using ((reg1, reg2) = (Qubit[bitSize], Qubit[bitSize])) { let qubits1 = LittleEndian(reg1); let qubits2 = LittleEndian(reg2); ApplyXorInPlace(num1, qubits1); ApplyXorInPlace(num2, qubits2); 3

```
Message("Before addition:");
DumpRegister((), reg2);

AddI(qubits1, qubits2);

Message("After addition:");
DumpRegister((), reg2);

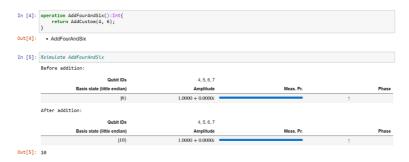
ResetAll(reg1);
return MeasureInteger(qubits2);
}
Out[3]: - AddCustom

6
```

- 1 You need the registers to be large enough to hold the largest possible sum of the two integers. If you double the number of qubits needed to represent the largest input and add 1, you know that be enough to hold the sum without overflow.
- ②Here, you can use the LittleEndian UDT provided by the Q# standard library to mark that each register of qubits is meant to interpreted as an integer using the little-endian encoding (also known as "least-significant order"). That is, when interpreting a LittleEndian register as an integer, that means to treat lowest qubit index as the least significant bit. For instance, to represent the integer 6 represented as a three-qubit quantum state in little-endian notation, you'd write |011), since 6 = 0 x 2⁰ + 1 x 2¹ + 1 x 2² = 2 + 4.
- 3The ApplyXorInPlace operation has the effect of encoding a LittleEndian representation of an integer into the qubit register. This works since x \oplus 0 = x whether x is 0 or 1.
- 4 Using the operation Addl loaded from the numerics package, you can add integers represented by the two input registers qubits1 and qubits2.
- **5**Lastly, you reset the first register so it can be deallocated, and then measure the register with the results. The MeasureInteger operation also resets the register after it is done as noted in the docs at https://docs.microsoft.com/qsharp/api/qsharp/microsoft.quantum.arithmetic.measureinteger, so you don't need to reset the qubits2 register as well.

You can see the output of running this snippet in figure 11.5.

Figure 11.5. Output of an example using the numerics library to add integers encoded in qubit registers.



11.7 MULTIPLYING WITH QUBITS IN SUPERPOSITION

Of course, if all you could do with arithmetic on a quantum computer was to add one pair of integers at a time, there wouldn't be any reason to do so using qubits instead of on your classical computer. Thankfully, the AddI and many other similar arithmetic operations work in superposition as well; a fact that you'll need to use these arithmetic operations with phase estimation in the next part of the Chapter. Before jumping straight to that, though, it's helpful to play around a little bit with what it means to add or multiply integers in superposition.

Here, you'll get to see how to apply what you've learned about superposition to arithmetic as well by using the MultiplyByModularInteger as an example. Shortly, you'll use the same operation to construct the oracle that you need for Shor's factoring algorithm later, so it's a pretty practical application as well!

First, let's look at an operation that you can use to prepare a register in a superposition of two integers. In , you can see how to do so using what you learned about the ApplyXorInPlace operation from the previous section and what you learned about the Controlled functor from Chapter 9.

```
Listing 11.8. Operation to prepare a register in a superposition of integers.
open Microsoft.Quantum.Arithmetic;
open Microsoft.Quantum.Diagnostics;
open Microsoft.Quantum.Math;
operation PrepareSuperpositionOfTwoInts(
    intPair : (Int, Int),
    register : LittleEndian,
: Unit is Adj + Ctl {
    using (ctrl = Qubit()) {
        H(ctrl);
        within {
            X(ctrl);
        } apply {
            Controlled ApplyXorInPlace(
                [ctrl],
                (Fst(intPair), register)
            );
        Controlled ApplyXorInPlace(
            [ctrl],
            (Snd(intPair), register)
        (ControlledOnInt(Snd(intPair), Y))(register!, ctrl);
    }
```

1 This operation takes a register and a pair of integers and prepares that register in a superposition of those integers in a LittleEndian encoding.

- ②The H operation puts your control qubit in a superposition, preparing your control qubit in the $|+\rangle=(|0\rangle+|1\rangle)/\sqrt{2}$ state so that when you control later operations on that qubit, they are also in superposition.
- 3As you saw with other uses of Controlled, controlled operations do something when their control registers are in the all-ones state (|11 ···1)). If you want to control on the zero state instead, then you can use the X operation to map the |0⟩ state to the |1⟩ state. By placing the call to X in a within/apply block, Q# will make sure to undo your call to X after you've applied your controlled operation. This is very similar to the ControlledOnInt function that you used in the previous chapter, and is how that function is actually implemented in the Q# standard libraries.
- Adding the Controlled functor here adds a new input representing the control register, as you learned about in Chapter 9. Following the control register is a tuple with the arguments for the original (uncontrolled) operation. In this case, that's the first integer that you want to prepare as a state, as well as the register that you want to prepare that state on.
- **5** Now you can do the same thing with the second integer in intPair and encode it in register, as controlled on the ctrl qubit.
- (a) Here, you can add some phase to one of the two branches of your superposition by using the Y operation to rotate, controlled on your control qubit being in the |1) state. While this is strictly unnecessary, it helps you see how the phase applied by the Controlled Y operation propagates through later steps.

Once you have a quantum register that represents the superposition of two integers, you can apply other arithmetic operations across that superposition. In listing 11.9, for instance, you can use the <code>DumpMachine</code> callable to see how the state of your register changes when using the <code>MultiplyByModularInteger</code> operation provided by the Q# standard libraries.

Listing 11.9. Example using the numerics library to multiply integers in superposition.

```
operation MultiplyInSuperpositionMod(
    superpositionInts: (Int, Int),
    multiplier : Int,
    modulus : Int
)
: Int {
    using (target = Oubit[BitSizeI(modulus)]) {
        let register = LittleEndian(target);
        PrepareSuperpositionOfTwoInts(superpositionInts, register);
        Message("Before multiplication:");
        DumpMachine();
        MultiplyByModularInteger(
            multiplier, modulus, register
        );
        Message("After multiplication:");
        DumpMachine();
```

```
return MeasureInteger(register);
}
```

- The first thing here is to figure out how big you need our registers to be, and allocate a register of qubits of that size. Here, since you are doing modular multiplication, the largest possible value our register has to hold is modulus.
- 2 You can use the operation that you defined in to prepare our target register in a superposition of integers.
- 3 The numerics package provides the MultiplyByModularInteger operation, which can take a LittleEndian register and multiply the value it represents by a classical multiplier modulo a given modulus
- The MeasureInteger operation measures the register and returns the classical Int that it represented (based on the encoding provided) and resets the register as you saw in listing 11.8.

If you run the code in listing 11.8 and in listing 11.9 the sample notebook you can see the following output in figure 11.6.

Figure 11.6. Output of multiplying 3 by 2 and 3 in superposition mod 8.



You can see in that the registers correctly show a superposition of 2 and 3 before multiplication, and then show a superposition of 1 and 6 after. What would you expect the correct output to be? If we multiplied as you would normally, it should be 6 and 9, but since we are doing arithmetic mod 8, the 9 goes to 1. When we then measure that register we will get 6 half of the time and 1 the other half since they are in an equal superposition as you can see from the state amplitude bars depicted in .

Exercise 11.4

Suppose that you have prepared a register in the state $\frac{1}{\sqrt{2}}|2\rangle + |7\rangle$, with each ket representing an integer in the little-endian encoding. What state would your register be in after multiplying by 5 modulo 9? Write a Q# program that uses DumpMachine to confirm your answer.

Bonus exercise 11.5

If you run the same program as in the previous exercise, but try to multiply by 3 modulo 9, you'll get an error; why? As a hint, consider what you learned in Chapter 7 has to be true of a classical function in order for it to be represented by a quantum oracle. If you're stuck, the answer is provided below, lightly hidden using https://rot13.com/.

ANSWER: Zhygvcylvat ol guerr zbq avar vf abg erirefvoyr. Sbe vafgnapr, obgu bar gvzrf guerr naq sbhe gvzrf guerr zbq avar tvir mreb, rira gubhtu bar naq sbhe nera'g rdhny zbq avar. Fvapr pynffvpny shapgvbaf unir gb or erirefvoyr va beqre gb or ercerfragrq ol dhnaghz bcrengvbaf, gur ZhygvcylolZbqhyneVagrtre envfrf na reebe va guvf pnfr.

IMPORTANT

Notice what you've done here: you've used a quantum program to multiply an integer represented by a quantum register by a classical integer. The computation happens entirely on the quantum device, and doesn't use any measurement; that means when your register started off in superposition, the multiplication happens in superposition as well.

11.8 MODULAR MULTIPLICATION IN SHOR'S ALGORITHM

Now that you have seen a bit of the numerics library, let's return to the factoring scenario. The main steps from where we will need to do some modular arithmetic and utilize the numerics library are implementing steps three and four (see algorithm 11.1). There are three operations that you can implement that will make short work of implementing those steps of Shor's algorithm and help you get to factoring integers.

The first operation you can take a look at from the sample code for this chapter is one that implements step three of, the EstimateFrequency operation ().

Algorithm 11.1 steps 3 and 4: Pseudocode for factoring an integer with Shor's algorithm.

- 1. Use iterative phase estimation to find the *frequency* of the classical function $f(x) = a^x \mod N$. The frequency tells you about how quickly f returns to the same value as x increases.
- 2. Use a classical algorithm known as the *continued fractions expansion* to convert the frequency from the previous step into a period (r). The period r should then have the property that f(x) = f(x + r) for all inputs x.

TIP

This operation makes use of the phase estimation operations provided with the Q# standard libraries and that you saw in Chapter 9. If you need a refresher, head back to Chapter 8 for an overview of phase estimation, or to Chapter 9 for how to run phase estimation with the standard libraries.

Listing 11.10. The operation EstimateFrequency uses phase estimation to help

```
estimate the frequency of a generator.
operation EstimateFrequency(
    inputOracle : ((Int, Qubit[]) => Unit is Adj+Ctl),
    nBitsPrecision : Int,
    bitSize : Int
)
: Int {
    using (register = Qubit[bitSize]) {
        let registerLE = LittleEndian(register);
        ApplyXorInPlace(1, registerLE);
        let phase = RobustPhaseEstimation(
            nBitsPrecision,
            DiscreteOracle(inputOracle),
            registerLE!
        );
        ResetAll(register);
        return Round(
            ((phase * IntAsDouble(2 ^ nBitsPrecision)) / 2.0) / PI()
        );
    }
```

- ① Since this is the primary step that leverages qubits, you need to allocate a register to use that is big enough to represent the modulus (since you are doing modular arithmetic and you won't have numbers larger than the modulus to represent).
- 2 The freshly allocated register has to specify how to encode the integers it will be representing so you can use the LittleEndian UDT from the numerics library to wrap your newly allocated register.
- The ApplyXorInPlace operation takes an Int and will XOR it with the integer stored in the register provided in the second argument. Since the integer stored in registerLE is 0, this operation has the effect of toggling the integer stored in this register to being the largest integer it can hold (|111...1)).
- You can reuse the RobustPhaseEstimation operation from the Q# standard libraries that you first used in Chapter 9. That operation takes an oracle, a quantum register, and an Int representing how many bits of precision you want the phase estimated to.
- Shere we need to wrap our oracle in a UDT DiscreteOracle that just makes it clear to RobustPhaseEstimation that we want the operation that we are passing in here to be interpreted as an oracle.
- 6 Once the phase estimation is done, you need to reset all the qubits in the register.

The phase you estimated is just that, a phase, and to convert that to a frequency you can use this equation to convert it to a frequency: $\frac{(\text{phase}*2^{\text{nBitsPrecision}-1})}{\pi}.$

Now that you have the scaffolding for <code>EstimateFrequency</code> under your belt, you can take a look at an operation that implements the oracle that you will need for this algorithm. The <code>ApplyPeriodFindingOracle</code> operation is just that, one that implements an operation that is structured like an oracle for the function $f(power) = generator^{power} \mod modulus$. See listing 11.11 for an implementation of <code>ApplyPeriodFindingOracle</code>.

Listing 11.11. Using the numerics library to implement an oracle that represents the

- 1 You can use the Fact function to do some input checking that the provided generator and modulus are coprime.
- ②The MultiplyByModularInteger operation is the same as what you used in listing 11.9 and is used here to help implement the action of this oracle, i.e. multiply the integer represented in the $f(\text{power}) = \text{generator}^{\text{power}} \mod \text{modulus}$ target register by
- $\label{eq:continuous}$ The Microsoft.Quantum.Math namespace has as you have seen earlier has the function $f(\text{power}) = \text{generator}^{\text{power}} \mod \text{modulus}$ ExpModI which allows you to easily calculate the expression
- **4** The qubit register that ApplyPeriodFindingOracle takes must specify how to interpret the register as an integer, so you can use the UDT LittleEndian to clarify this to the compiler.

The previous two operations you looked at form the basis for step three of algorithm 11.1. Now, you will need an operation that takes care of step four, where you need to convert the estimated frequency of the generator to a period. The operation EstimatePeriod in does just that, given a generator and a modulus it will repeat estimating the frequency using EstimateFrequency and use the continued fractions algorithm to ensure that the frequency estimated yields a valid period.

Listing 11.12. The EstimatePeriod operation takes the result

of EstimateFrequency and handles edge cases and converts it to a period with the

help of the operation PeriodFromFrequency.

```
operation EstimatePeriod(generator : Int, modulus : Int)
: Int {
    Fact(
        IsCoprimeI(generator, modulus),
        "`generator` and `modulus` must be co-prime"
    );
    let bitSize = BitSizeI(modulus);
    let nBitsPrecision = 2 * bitSize + 1;
    mutable result = 1;
    mutable frequencyEstimate = 0;
    repeat {
        set frequencyEstimate = EstimateFrequency(
            ApplyPeriodFindingOracle(generator, modulus, _, _),
            nBitsPrecision, bitSize
        );
        if (frequencyEstimate != 0) {
            set result = PeriodFromFrequency(
                frequencyEstimate, nBitsPrecision,
                modulus, result
            );
        } else {
            Message("The estimated frequency was 0, trying again.");
    } until(ExpModI(generator, result, modulus) == 1)
                                                                            ܣ
    fixup {
        Message(
                                                                            12
            "The estimated period from continued fractions failed, " +
            "trying again."
        );
    }
    return result;
```

- 1 You can use the Fact function to do some input checking that the provided generator and modulus are coprime.
- 2 The IsCoprimeI function from the Microsoft.Quantum.Math namespace makes checking if the generator and modulus are coprime easy.
- 3 The largest integer a register of qubits would need to hold is the modulus, so you can use BitSizeI to help calculate number of bits such that $modulus \leq 2^{of \ bits}$.
- **4** To use a floating-point number to represent a frequency of the form k / r where r is the period and where k is an arbitrary other integer, you need that floating-point number to have enough bits of precision to be a good approximation to k / r. Since neither k nor r can be larger than N, you just need one more bit than the number of bits you need to represent both k and r.
- The result mutable variable here will keep track of your current best guess for a period as you repeat the repeat block below.

- **6** You want to repeat the frequency estimation steps here as many times as necessary to ensure you have a viable period estimate to move forward with.
- You can call the EstimateFrequency operation that you looked at earlier and pass it the appropriate arguments.
- **3** The operation ApplyPeriodFindingOracle is partially applied here to ensure that EstimateFrequency can apply apply it to the right power and register values.
- ① If frequencyEstimate is 0, that is a trivial failure case, because a frequency of 0 has an undefined period ($\frac{1}{0}$). If you get 0 as your estimate, then the repeat block will just run again as the until condition will not be met in that case.
- The PeriodFromFrequency function call here captures step number four of which utilizes some built-in functions in the Microsoft.Quantum.Math namespace to use the continued fractions algorithm to calculate the period from the frequency.
- ①You know that you want to repeat the frequency estimation and period calculation until you have a period such that $\operatorname{generator^{result}} \operatorname{mod} \operatorname{modulus} = 1$.
- lf the until condition is not met, then the fixup block is run, which here just issues a message that it is going to try again.

With this last operation, you now have all the code you need to fully implement Shor's algorithm! In the next section, you will put it all together and explore the implication of this integer factoring algorithm.

11.9 PUTTING IT ALL TOGETHER

You now have learned and practiced all the skills that it takes to program and run Shor's algorithm. The quantum part to Shor's algorithm was pretty familiar, thanks to what you learned in Chapters 8 and 9 about phase estimation, and you worked through classical algebra that links the tasks of factoring numbers and finding the period of a generator. This is no small feat — you should be quite proud of yourself for making it this far in your quantum journey!

Continued fraction convergents

You may have noticed one other bit of classical math happening in Shor's algorithm that we haven't really touched on yet. In particular, the PeriodFromFrequency function is called on output you get from phase estimation before continuing.

Listing 11.13. The PeriodFromFrequency function used in Shor's algorithm.

```
function PeriodFromFrequency(
    frequencyEstimate : Int, nBitsPrecision : Int,
    modulus : Int, result : Int
)
: Int {
    let continuedFraction = ContinuedFractionConvergentI(
        Fraction(frequencyEstimate, 2^nBitsPrecision), modulus
    );
```

```
let denominator = AbsI(Snd(continuedFraction!));
return (denominator * result) / GreatestCommonDivisorI(
    result, denominator
);
}
```

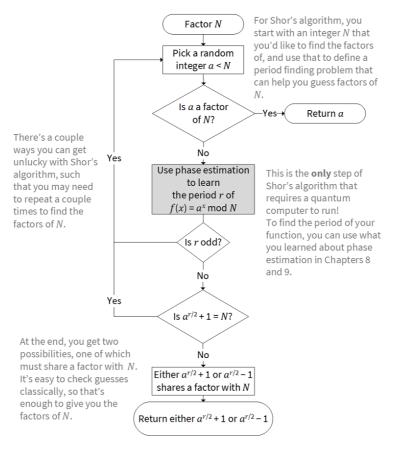
This is needed because phase estimation doesn't tell you *exactly* what you need. Instead of telling you how long it takes your function to spin around a clock (the period of your function), it tells you how fast your function spins around said clock; something more like a frequency. Unfortunately, you can't quite do something like taking the reciprocal of your frequency to get back to your period, as you're looking for the period as an integer. Thus, if you got a frequency estimate of f, then you need to search near $f / 2^n$ for the closest fraction of the form N / r in order to find your period r.

This is an entirely classical arithmetic problem, and has thankfully been well solved using a technique known as continued fraction convergents. This solution is made available in the Q# standard libraries by the ContinuedFractionConvergentI function, making it easy to go from the estimate you get from phase estimation to something about the period of your function.

Let's take a moment to review the FactorSemiprimeInteger operation that you saw at the beginning of the chapter in light of what you have now learned. If you need a refresher, check out figure 11.7.

313

Figure 11.7. Shor's algorithm represented as a flowchart.



Listing 11.14. Operation that factors semiprime integers with Shor's Algorithm. operation FactorSemiprimeInteger(number : Int) : (Int, Int) { if (number % 2 == 0) { Message("An even number has been given; 2 is a factor."); return (number / 2, 2); } mutable factors = (1, 1); mutable foundFactors = false; repeat { let generator = RandomInt(number - 2) + 1; if (IsCoprimeI(generator, number)) { Message(\$"Estimating period of {generator}..."); let period = EstimatePeriod(generator, number); set (foundFactors, factors) = MaybeFactorsFromPeriod(generator, period, number);

```
} else {
        let gcd = GreatestCommonDivisorI(number, generator);
            $"We have guessed a divisor of {number} to be " +
            $"{gcd} by accident. Nothing left to do."
        );
        set foundFactors = true;
        set factors = (gcd, number / gcd);
} until (foundFactors)
fixup {
   Message(
        "The estimated period did not yield a valid factor, " +
        "trying again."
    );
}
return factors;
                                                                        1
```

- 1 The first thing we have to check is if the integer to factor is even. If so, then 2 is as factor and we can stop early.
- 2 You can use the mutable variable factors to keep track of the factors that you find for number as you step through the algorithm.
- 3 You can use the mutable flag foundFactors to keep track of if you find any factors for number as you step through the algorithm.
- ◆ This line matches up with step 1 in algorithm 11.1, as RandomInt is used to select a random integer in the range from 1 to number − 1; you'll use this random integer as your generator.
- **⑤** This is step 2 in algorithm 11.1, where you verify that the generator is coprime to the integer that we want to factor, if it's not the else clause handles returning the common factor between the two.
- **6** This line covers steps 3 and 4 in algorithm 11.1, it will return a period that it computes with continued fractions from the frequency estimation inside EstimatePeriod.
- This function MaybeFactorsFromPeriod takes the estimated period and uses algebra to turn it into integers that might be potential factors of the integer we are trying to factor. Sometimes it will fail, so it returns a tuple of a boolean value indicating whether it succeeded as well as the tuple of the factors it finds when it succeeds.
- This else statement handles the case that the generator we guessed at the beginning has a common factor with the number we are trying to factor. We can use the Q# library function GreatestCommonDivisorI to find the shared factors and set foundFactors and factors appropriately.
- The until statement adds the condition for how long we should repeat the preceding block. Here we want to continue looking until we find the factors we are looking for.
- The fixup statement tells the program to do before repeating the main loop again, here we just have it message you to let it know that your Q# program will try again.
- finally, the program can return the tuple of factors for the input integer.

In listing 11.15, you can see what it looks like if you run this operation in an IQ# notebook to factor 21.

Listing 11.15. Output of running FactorSemiprimeInteger

```
In [1]: open IntegerFactorization;
        operation Factor21(): (Int, Int) {
            return FactorSemiprimeInteger(21);
Out[1]: - Factor21
In [2]: %simulate Factor21
       We have guessed a divisor of 21 to be 3 by accident. Nothing left to do.
Out[2]: (3, 7)
In [3]: %simulate Factor21
        Estimating period of 19...
       The estimated period from continued fractions failed, trying again.
Out[3]: (7, 3)
In [4]: %simulate Factor21
       Estimating period of 17...
       The estimated period did not yield a valid factor, trying again.
       We have guessed a divisor of 21 to be 3 by accident. Nothing left to do.
Out[4]: (3, 7)
```

1To call the FactorSemiprimeInteger operation from within a notebook, it's helpful to write a new operation that provides the input 21.

You can see in listing 11.15 that each time your code correctly returned the prime factors of 21, 3 and 7. The operation was run three times to show some potential different outcomes you might get. In In [2], when you tried to guess a generator, your call to RandomInt ended up returning an integer that was not coprime to 21. Thus, you were able to use GreatestCommonDivisorI to find a factorization. In In [3], your call to RandomInt selected a generator of 19, and then had to run the frequency estimation twice to ensure the continued fractions algorithm succeeded. Lastly, the final run on In [4], one full round of the period finding task completed but failed to yield a correct factor, and when it went to pick a new generator it guessed a factor by accident.

NOTE

Given the limitations of running this on simulators or small hardware devices, you will quite frequently guess the correct factors when selecting the generator. You'll also get unlucky more often with small integers, guessing trivial factors like 1. Both of these edge cases will happen less often as the number you are trying to factor grows in size.

Using a simulator on your laptop, your desktop, or in the cloud, you likely won't be able to factor anything terribly large with Shor's algorithm. For example, it would be quite challenging to factor a 30-bit number by simulating Shor's algorithm on a classical computer, but 40-bit numbers were already considered woefully insufficient at standing up against classical factoring algorithms in 1992. This might seem like it makes Shor's algorithm useless, but all that really tells you is that it's hard to use a classical computer to simulate large quantum programs; you saw why that was the case

in Chapters 4 and 5.

Indeed, since the same algorithm works for much larger numbers (using 4,096-bit keys isn't at all overkill for protecting personal data, for example), such as those commonly used to protect your data online, understanding how Shor's algorithm and other similar quantum algorithms work can help you understand the assumptions that go into modern uses of cryptography, and what else you'll need going forward.

What's next for privacy?

Given what you've learned about Shor's algorithm, it might seem like the cryptography that protects everything from your health records to your chat history is doomed. Thankfully, there's both quantum technologies such as quantum key distribution, as you learned about in Chapter 3, as well as new classical technologies meant to resist algorithms such as Shor's. The latter class of technologies, known as post-quantum cryptography, is the subject of much ongoing research and exploration.

As it turns out, Q# can play an important role in cryptography research by making it easier to understand how large a quantum computer you would need to attack a given cryptosystem. For instance, researchers at Google recently used Q# and Python to help improve the cost required to implement the modular multiplication step of Shor's algorithm (https://arxiv.org/abs/1904.07356), helping them estimate that 20 million qubits would be required to attack reasonable RSA instances using current quantum algorithms (https://arxiv.org/abs/1905.09749).

Similarly, https://github.com/Microsoft/grover-blocks is a great example of using Q# to understand how Grover's algorithm from the previous chapter affects symmetric-key algorithms like AES.

In both cases, Q# is a valuable tool for understanding how much quantum computing power an attacker would need to compromise current cryptosystems. Together with assumptions about the speed at which quantum algorithms and hardware will continue to improve, assumptions about how much quantum computing power will be feasible for attackers to purchase, and requirements for how long algorithms like RSA need to be secure in order to guarantee your privacy, the understanding developed by using Q# can help us to understand how quickly current cryptosystems need to be replaced. Like anything in information security, guaranteeing privacy against quantum attackers is a very complex topic, not to mention a topic without any easy answers. Thankfully, tools like Q# and the Quantum Development Kit help make the problem a bit more tractable.

11.10 WRAPPING UP

At this point, before we say goodbye, it's helpful to take a step back and appreciate how all the various skills you learned throughout the book came together in this Chapter to help you understand a real-world application for quantum computers. In Part I, you learned the basics about how we can describe, simulate, and the basic quantum effects that make quantum computing unique. You learned in Chapter 3 how to use single qubits and superposition to securely share cryptographic keys with quantum key distribution. In Chapters 4 and 5, you entangled multiple qubits to play games and move data around a quantum device. You even built your own quantum simulator in Python to help you implement these games and learn about the math that helps us describe quantum effects.

With all of those basics under your belt, in Part 2 you then started writing quantum algorithms to help the crew in Camelot play some games. In Chapter 6, you learned about Q#, a new programming language specifically designed to help you easily write

programs for a quantum computer. In Chapter 7, you implemented the Deutsch–Jozsa algorithm to help select a new king, but along the way, you also learned about oracles and how they can help you evaluate classical functions in a quantum program. You also developed your own phase estimation program in Chapter 8, where you learned how to manipulate phase and leverage it with phase kickback to learn things about operations in your quantum programs.

Having a new toolbox of quantum development techniques at your disposal, you tackled some of the most exciting applications of quantum computing. In Chapter 9, you learned about Hamiltonian simulation and how we can use the quantum systems in our quantum computers to simulate the energy levels in various chemicals. Chapter 10, you implemented Grover's algorithm to help search for information in unstructured data with amplitude amplification. In this Chapter, you used everything from Q# diagnostic functions and operations through to phase estimation, and everything from within/apply blocks through to oracle representations of classical functions to factor numbers on a quantum computer. Indeed, using what you learned in the rest of the book, most of the hard parts of writing out Shor's algorithm were the *classical* parts you needed to connect the problem of factoring numbers to one of period finding.

While this book did not exhaust all there is to learn about quantum computing — a lot has happened since 1985, after all! — what you've learned in this book gives you what you need to keep learning, keep exploring, and keep pushing forward with quantum computing. Using Python and Q# together, you have the tools you need to take part in one of the most exciting advances in computing, to help your peers and colleagues learn along with you, and to build a community that can put quantum computing to good use.

Go have fun!

What's next?

While there's still always more to learn about quantum computing, you now have what you need to start developing quantum applications using Python and Q# together. If you're interested in learning and doing more with quantum computing, here's some resources to help you take your next step!

- Q# Community (qsharp.community): An open-source community around quantum programming in Q#, including blogs, code repositories, and online meetups.
- Microsoft Quantum Docs (docs.microsoft.com/quantum): Full reference documentation for all things
 related to the Quantum Development Kit.
- arXiv (arxiv.org): An online repository for scientific papers and manuscripts, including a huge amount
 of research about quantum computing.
- Unitary Fund (unitary.fund): Grants and financial support for open-source quantum software, along
 with neat suggestions for open-source projects you can take on.
- Quantum Open Source Foundation (www.qosf.org): Foundation for the development of open-source quantum software, including a list of current projects, and resources for further learning.
- QCEthics (qcethics.org): Resources for ethics in quantum computing.
- O-Turn (q-turn.org): An inclusive quantum computing conference series.
- Quantum Algorithm Zoo (quantumalgorithmzoo.org): List of known quantum algorithms, with links to papers about each.

• Quantum Computing: A Gentle Introduction (ISBN 0262015064): More detail about the math behind the quantum algorithms that you learned about in this book.

Many universities and colleges also have courses or research programs that may be of interest as you continue exploring quantum computing. However you decide to continue, we hope you have fun and work to make the quantum computing community even more wonderful!

11.11 Summary:

- Program with the Numerics library for Q#.
- Implement Shor's algorithm for factoring integers.
- Recognize the implications of quantum computing to security infrastructure.

Installing Required Software

This chapter covers:

- · Use the samples for this book online with Binder, or with Visual Studio Online,
- How to install a Python environment that you can use to design your own quantum simulation stack,
 and
- How to install the Microsoft Quantum Development Kit, including a compiler for the O# language.

A.1 RUNNING SAMPLES ONLINE

If you'd like to try the samples in this book out without installing anything, there's two great options to do so:

- Binder (mybinder.org), a free service for exploring hosted Jupyter Notebooks, and
- Visual Studio Online, an Azure-hosted development environment.

A.1.1 Using Binder

To use Binder, simply go to https://mybinder.org/v2/gist/crazy4pi314/learn-qc-with-python-and-qsharp/master. It may take a few moments, by Binder will spin-up a new Jupyter Notebook installation, complete with the Python packages you need and Q# support.

The Binder service is intended for exploration only, and will **DELETE** your changes after approximately twenty minutes of inactivity. While Binder is a great way to get started, if you would like to continue developing quantum programs, it's helpful to either use Visual Studio Online or to install Python and the Quantum Development Kit locally on your machine.

A.1.2 Using Visual Studio Online

To use Visual Studio Online, go to http://online.visualstudio.com/ and press "Get Started." Once you login with a Microsoft account that is associated with an Azure subscription, you'll be prompted to create a new Visual Studio Online plan that can be used to bill Visual Studio Online services.

If you don't already have an Azure account, you can create one for free at https://azure.microsoft.com/free. For more details, check out Chapter 1 of Learn Azure in a Month of Lunches.

After selecting a plan, you can then create new Visual Studio Online *environments*. Each environment is a self-contained development environment complete with all the software you'll need to work on a project. The code samples for this book, provided at https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp, come with all the configuration needed to be used as a Visual Studio Online environment. Press "Create environment" from https://online.visualstudio.com/environments, then give your environment a new name. For the Git repository, paste in https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp, and then press "Create." After a few minutes, you should see a new Visual Studio Online tab open up with all of the samples for the book, along with all of the Python packages that you need, and the Quantum Development Kit.

In most ways, your new Visual Studio Online environment is just like running things locally on your system, but backed by a service running in Azure. There are a few differences to keep in mind, though.

For example, the Visual Studio Online environment configuration provided for use with this book differs a bit from a local installation in that it doesn't use the Anaconda distribution. As a result, you can skip the conda activate qsharp-book step when using Visual Studio Online.

As of this writing, Visual Studio Online is still in public preview, so that some bugs and service interruptions are to be expected. You can always find the latest information about Visual Studio Online at http://docs.microsoft.com/visualstudio/online.

A.2 INSTALLING A PYTHON ENVIRONMENT

In the first several chapters of the book, we make heavy use of Python as a tool to explore quantum programming. In doing so, we rely on several Python libraries that make it easier to write scientific programs. As a result, it's often easier to start from a Python distribution that packages Python along with other libraries and package management tools. We'll go through getting up and running with the Anaconda distribution, but you should be able to follow a similar procedure if you would prefer to use a distribution such as Enthought Python, or Python(x, y).

A.3 INSTALLING ANACONDA

To install Anaconda, follow the instructions

at https://docs.anaconda.com/anaconda/install/.

WARNING

At the time of this writing, the Anaconda distribution is provided with either Python 2.7 or 3.7. Python 2.7 has official reached end-of-life, and is provided for compatibility reasons only. We will assume Python 3.6 or later in this book, so make sure that you install a version of Anaconda that provides Python 3.

A.4 INSTALLING PYTHON PACKAGES WITH ANACONDA

Packages are a great way to collaborate and save time when you are trying to learn or develop new code. They are a way of collecting related code and wrapping it up in such a way it is easy to share to other machines. Packages can be installed on your machine with what is sensibly called a package manager, of which there are a few common options for Python. The reasons you might choose one manager over the other is each one has its own listing of packages and the package you want to install may only be known by a particular manager (depending on how the author deployed it). Let's start by looking the package managers we have already installed as a part of Anaconda.

As a default, Anaconda comes with two package managers pip and conda. Given we have installed Anaconda, the conda package manager has some additional features to pip that make it a good default choice for package management. conda has support for installing dependencies automatically when you install a package, it has the concept of environments (See sidebar) that are really helpful for creating dedicated Python sandboxes? for each project you are working on. A good general strategy is thus to install packages from conda when they are available, and to install packages from pip otherwise.

NOTE

The conda package manager can be used with most common command line environments, but if you would like to use conda with PowerShell, you will need version 4.6.0 or later. To check your version, run conda ——version. If you need to update, run conda update conda.

Conda Environments

So far, we've simply run commands like pip install and conda install to install new packages for our entire Anaconda installation. We may run into the situation, however, that the packages we need for two different projects contradict each other. To help isolate projects from each other, the Anaconda distribution provides conda env as a tool to help manage multiple environments. Each environment is a completely independent copy of Python with only the packages needed for a particular project or application. Environments may even use different versions of Python from each other, with one environment using 2.7 and another using 3.7. Environments are also great for collaborating with others, as you can send your teammates a single small text file, environment.yml, to tell their conda env how to create an environment that's identical to yours.

For more information:

https://docs.conda.io/projects/conda/en/latest/user-guide/tasks/manage-environments.html

Following that strategy, let's start by making a new environment with the packages we

©Manning Publications Co. To comment go to liveBook

need. The code samples for this book, found at https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp, come with a file called environment. yml that tells conda what packages you need in your new environment. Clone or download the code from the repository for this book, then run the following at your favorite command line:

Listing 12.1. Creating a new conda environment

conda env create -f environment.yml

This will create a new environment called qsharp-book using packages from the conda-forge channel, and will install Python (version 3.6 or later), the IPython interpreter, NumPy, the matplotlib plotting engine, Jupyter Notebook, and QuTiP into the new environment. The conda package manager will prompt you to confirm the list of packages that will be installed, press "y" and then "Enter," and grab a cup of coffee.

The environment.yml file provided at https://github.com/crazy4pi314/learn-qc-with-python-and-qsharp also installs the qsharp package, which provides integration between Python and the Quantum Development Kit (which you'll install in the next step). If you are using a distribution other than Anaconda, you'll need to run pip install qsharp manually.

Once conda has finished creating your new environment, let's give it a go. To test your new environment, first activate it:

Listing 12.2. Activating the new conda environment

conda activate gsharp-book

Once the qsharp-book environment has been activated, the python command should invoke the version installed into that environment. To check this, you can print the path to the python command for your environment. Run python, and then run the following at the Python prompt:

Listing 12.3. Testing the new conda environment

>>> import sys; print(sys.executable)"
C:\Users\Chris\Anaconda3\envs\qsharp-book\python.exe
1

1 Note that you might see a different path depending on your system.

If your environment was created successfully, you can use it either at the command line with IPython, or in your browser with Jupyter Notebook. To get started with IPython, run ipython from your command line: (Make sure that you've activated gsharp-book first!)

Listing 12.4. Using IPython from the quantum environment

```
$ ipython
In [1]: import qutip as qt
In [2]: qt.basis(2, 0)
Quantum object: dims = [[2], [1]], shape = (2, 1), type = ket
Qobj data =
[[1.]
  [0.]]
```

If you would like to learn more about using NumPy, continue reading in Chapter 2. If you would like to learn more about using QuTiP, continue reading in Chapter 4. If you would like to learn about the Quantum Development Kit, continue reading here.

For more information

- Anaconda docmentation: https://docs.anaconda.com/anaconda/
- NumPy documentation: http://numpy.scipy.org/
- Jupyter Notebook documentation: https://jupyter-notebook.readthedocs.io/en/stable/notebook.html

A.5 INSTALLING THE QUANTUM DEVELOPMENT KIT

TIP

The latest version of the installation instructions for the Quantum Development Kit can be found at https://docs.microsoft.com/quantum/install-guide/. To use the code samples provided with this book, you'll want to make sure to install Python and Jupyter Notebook support when following the Quantum Development Kit install guide.

The Quantum Development Kit from Microsoft is a set of tools for working with and programming in Q#, a new language for quantum programming.

NOTE

In this book we will focus on using Visual Studio Code, but the Quantum Development Kit can be used with any other text editor as well by following the command line instructions in the main text.

The Quantum Development Kit can also be used with Visual Studio 2019 or later by using the extension at https://marketplace.visualstudio.com/items?itemName=quantum.DevKit

Using the installation of Visual Studio Code from setting up your Python environment, we need to do a few things to use the Quantum Development Kit with C#, Python, and Jupyter Notebook.

- Install the .NET Core SDK,
- Install the project templates for Q#,
- Install the Quantum Development Kit extension for Visual Studio Code,
- Install Q# support for Jupyter Notebook, and
- Install the qsharp package for Python.

Meet the .NETs

At this point, answering the question of what .NET is has become a bit more complicated than it used to be. Historically, ".NET" was a reasonable shorthand for the .NET Framework, a virtual machine and

©Manning Publications Co. To comment go to <u>liveBook</u>

compiler infrastructure for any of the .NET languages (most notably, C#, F#, and Visual Basic .NET). The .NET Framework is Windows-only, but third-party reimplementations such as Mono exist for other platforms including macOS and Linux.

A couple years ago, though, Microsoft and the .NET Foundation open-sourced a new flavor of .NET called .NET Core. Unlike .NET Framework, .NET Core is cross-platform out of the box. \.NET Core is also much smaller, with much of the functionality separated out into optional packages. This makes it much easier to have multiple versions of .NET Core on the same machine, and for new .NET Core features to be introduced without compatibility issues.

This bifurcation of .NET into Framework and Core comes with its own foibles, however. To make .NET Core work better as a cross-platform programming environment, a few things in the .NET standard libraries were changed in ways that aren't entirely compatible with .NET Framework. To resolve this, the .NET Foundation introduced the concept of .NET Standard, a set of APIs offered by both .NET Framework and .NET Core. The .NET Core SDK can then be used to make libraries for either .NET Core or .NET Standard, and can build applications for .NET Core. Much of the libraries provided Quantum Development Kit target .NET Standard, so that Q# programs can be used from traditional .NET Framework applications or from new applications built using the .NET Core SDK.

Going forward, Microsoft has announced that the next version of .NET will based off of .NET Core, but will just be called ".NET 5," reducing much of the confusion.

Once you've done this, you have everything you need to write and run quantum programs written in Q#.

A.6 INSTALLING THE .NET CORE SDK

To install the .NET Core SDK, go to https://dotnet.microsoft.com/download and select your operating system from the selections near the top of the page.

A.7 INSTALLING THE PROJECT TEMPLATES

One thing that might be different than what you're used to is that both .NET Framework and .NET Core development center around the idea of a *project* that specifies how a compiler is invoked to make a new binary. For instance, a C# project (*.csproj) file tells the C# compiler what source files should get built, what libraries are needed, what warnings are turned on and off, etc. In this way, project files work similarly to makefiles or other build management systems. The big difference is in how project files on .NET Core reference libraries.

A project file can specify one or more references to *packages* on NuGet.org, a package repository for .NET Framework and .NET Core software. Each package can then provide a number of different libraries. When a project that depends on a NuGet package is built, the .NET Core SDK will automatically download the right package, and then will use the libraries in that package to build the project.

From the perspective of quantum programming, this allows for the Quantum Development Kit to be distributed as a small number of NuGet packages that can be installed not on a machine, but into each project. This makes it easy to use different versions of the Quantum Development Kit on different projects, or to include only the parts of the Quantum Development Kit that you need for a particular project. To help get started with a reasonable set of NuGet packages, the Quantum Development Kit is

provided with templates for creating new projects that reference everything you need.

To install the project templates, run the following command at your favorite terminal:

Listing 12.5. Installing the project templates for the Quantum Development Kit

dotnet new -i "Microsoft.Quantum.ProjectTemplates"

Once the project templates are installed, you can use them by running dotnet new again:

Listing 12.6. Creating a new project with the Quantum Development Kit project

templates

dotnet new console -lang Q# -o ProjectName 1

Make sure to replace ProjectName by the name of the project you would like to create.

A.8 INSTALLING THE VISUAL STUDIO CODE EXTENSION

NOTE

This section assumes that you installed Visual Studio Code with your Anaconda installation above. If you do not have Visual Studio Code, you can install it manually from https://code.visualstudio.com/.

To install the extension for Visual Studio Code, open a new Visual Studio Code window and press Ctrl+Shift+X (Windows and Linux) or \H+Shift+X to bring up the extensions sidebar. From the search bar, type in "Microsoft Quantum Development Kit," and press the "Install" button. Once Visual Studio Code has installed the extension, the "Install" button will change to a "Reload" button, which will close Visual Studio Code and reopen your window with the Quantum Development Kit extension installed.

Alternatively, press Ctrl+P or \Re +P to bring up the Go To pallette. In the pallette, type ext install quantum-quantum-devkit-vscode and press Enter.

In either case, once the extension has been installed, to use it, open a folder (Ctrl+Shift+O or \(\mathbb{H} + \mathbb{Shift} + \mathbb{O} \)) containing the Q# project you'd like to work on. At this point, you should have everything you need to get up and programming with the Quantum Development Kit!

For more information

- Quantum Development Kit documentation: https://docs.microsoft.com/quantum
- Using the dotnet command: https://docs.microsoft.com/en-us/dotnet/core/tools/dotnet
- Getting started with Visual Studio Code: https://code.visualstudio.com/docs/introvideos/basics

A.9 INSTALLING IQ# FOR JUPYTER NOTEBOOK

Run the following from your favorite command line:

©Manning Publications Co. To comment go to liveBook

dotnet tool install -g Microsoft.Quantum.IQSharp
dotnet iqsharp install

TIP On some Linux installations, you may need to run dotnet iqsharp install -- user instead.

This will make Q# available as a language for Jupyter Notebooks such as the ones used in Chapter 6.

Linear Algebra Refresher

This chapter covers:

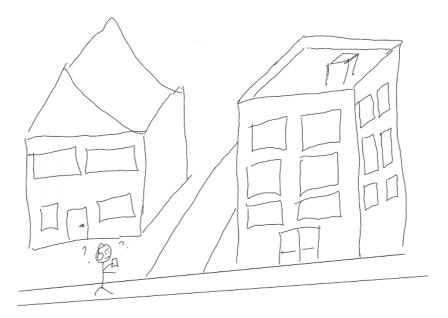
- · What vectors and matrices are,
- · How to work with vectors and matrices to represent linear functions,
- · How to use Python and NumPy to work with vectors and matrices.

B.1 APPROACHING VECTORS

One more thing we'll need before we can get to what qubits are is the concept of a *vector*.

Suppose a friend of yours is having people over to celebrate that they fixed their doorbell, and you'd very much like to find their house and celebrate the occasion with them. How can your friend help you find their home?

Figure B.1. Looking for your friend's house party...



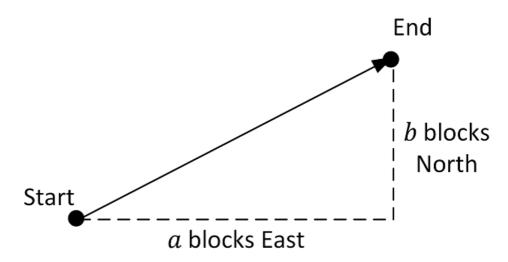
Vectors are a mathematical tool that can be used to represent a variety of different concepts — basically, anything that we can record by making an ordered list of numbers.

B.1.1 Examples of vectors

- Points on a map
- Colors of pixels in a display
- Damage elements in a computer game
- Velocity of an airplane
- Orientation of a gyroscope

For instance, if I'm lost in an unfamiliar city, someone can tell me where to go by giving me a vector that instructs me to take a blocks East and then b blocks North (we'll set aside the problem of routing around buildings). We write these instructions with the vector [[a], [b]].

Figure B.2. Vectors as coordinates.



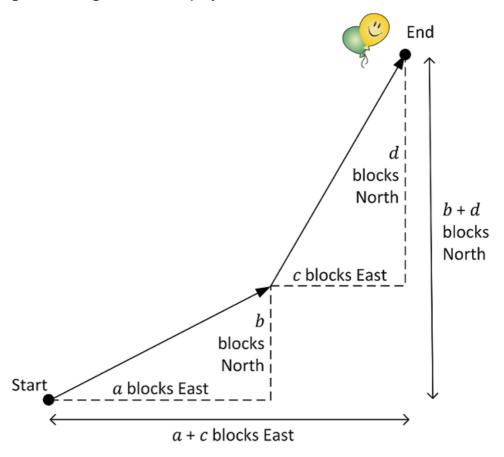
Like ordinary numbers, we can add different vectors together.

NOTE

An ordinary number is often called a scalar to distinguish it from a vector.

Using this way of thinking about vectors, we can think of this addition between vectors as being defined element-wise. That is, we interpret <code>[[a], [b]] + [[c], [d]]</code> as being instructions to go a blocks East, b blocks North, c blocks East, then finally d blocks North. Since it doesn't matter what order we step in, this is equivalent to taking <code>a + c</code> blocks East, then <code>b + d</code> blocks North, and so we write that <code>[[a], [b]] + [[c], [d]]</code> is <code>[[a + c], [b + d]]</code>.

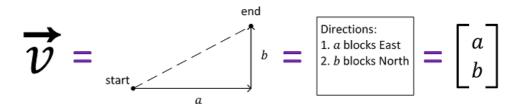
Figure B.3. Adding vectors to find a party.



VECTOR

A vector \vec{v} in d dimensions can be written as a list of d numbers. For instance, \vec{v} = [[2], [3]] is a vector in two dimensions.

Figure B.4. Drawn vectors have the same information as a list of directions, or a column of numbers.



©Manning Publications Co. To comment go to liveBook

Similarly, we can multiply vectors by ordinary numbers to transform vectors. I may be lost not just in any city, for example, but a city that uses meters instead of the feet that I'm used to. To transform a vector given in meters to a vector in feet, I'll need to multiply each element of my vector by about 3.28. Let's do this using a Python library called *NumPy* to help us manage how we represent vectors in a computer.

TIP

Full installation instructions are provided in Appendix A.

- ① Vectors are a special case of NumPy arrays. We create arrays using the array function, passing a list of the rows in our vector. Each row is then a list of the columns for vectors, we only ever have one column per row, but we'll have examples later where this isn't true.
- 2 Let's start with an example of going 30 meters East, then 50 meters North.
- 3 NumPy represents multiplication between scalars and vectors by the Python multiplication operator *.
- Printing out the result of multiplying, I see that I need to go 98.4 feet East, then 164 feet North.

Exercise B.1

What would 25 meters West and 110 meters North be in feet?

This structure makes it easier to communicate directions. If we didn't use vectors, then each scalar would need its own direction, and it would be critical to keep the directions and scalars together.

B.2 SEEING THE MATRIX FOR OURSELVES

As we'll see shortly, we can describe how qubits transform as we apply instructions to them in the same way that we describe transforming vectors, using a concept from linear algebra called a *matrix*. This is especially important as we consider transformations of vectors that are more complicated than adding or rescaling.

To see how to use matrices, let's return to the problem of finding the party — that doorbell isn't going to ring itself, after all! Up until now, we've simply assumed that the first component of each vector means East and the second means North, but someone could well have chosen another convention. Without a way of reconciling the transformation between these two conventions, I'll never find the party! Thankfully, not only will matrices help us model qubits later on in the chapter, they can help me find my way to my friends!

Happily, the transformation between listing North first and listing East first is simple to

©Manning Publications Co. To comment go to liveBook

implement: we swap the coordinates <code>[[a], [b]]</code> to obtain <code>[[b], [a]]</code>. Suppose that this swap is implemented by some function <code>swap</code>. Then, <code>swap</code> plays nicely with the vector addition that we saw above, in that <code>swap(v + w)</code> is always the same as <code>swap(v) + swap(w)</code>. Similarly, if we stretch a vector and then swap (that is, scalar multiplication), that's the same as if we had swapped and then stretched: <code>swap(a * v) = a * swap(v)</code>. Any function that has these two properties is a *linear* function.

LINEAR FUNCTION

A linear function is a function f such that f(ax + by) = af(x) + bf(y) for all scalars a and b and all vectors x and y.

Exercise B.2

Which of the following functions are linear?

- f(x) = 2x
- $f(x) = x^2$
- $f(x) = 2^x$

Linear functions are common in computer graphics and machine learning, as they include a variety of different ways of transforming vectors of numbers.

B.2.1 Examples of linear functions

- Rotations
- Scaling and stretching
- Reflections

What all of these linear functions have in common is that we can break them apart and understand them piece by piece. Thinking again of the map, if I'm trying to find my way to the party still (hopefully there's still some punch left) and the map I've been given has been stretched out by 10% in the North–South direction, and has been flipped in the East–West direction, that's not too hard to figure out. Since both the stretching out and flipping are linear functions, someone can set me on the right path by telling me what happened to the North–South direction and the East–West directions separately. In fact, we just did that at the beginning of this paragraph!

TIP

If you learn just one thing from this book, the most important take-away that we have to offer is that you can understand linear functions and thus quantum operations by breaking them up into components. We will see in the rest of the book that, since operations in quantum computing are described by linear functions, we can understand quantum algorithms by breaking them apart in the same way as we broke up our map example. Don't worry if that doesn't make a lot of sense at the moment, as it's a way of thinking that takes some getting used to.

This is because once I understand what happens to the North vector (let's call it [[1], [0]] as before), and to the West vector (let's call it [[0], [1]]), then I can figure out what happens to *all* vectors by using the linearity property. For example, if I am told there's a really pretty sight 3 blocks North and 4 blocks West of me, and I want to figure out where that is on my map, I can do so piece by piece:

- I need to stretch the North vector out by 10% and multiply it by 3, getting [[3.3], [0]].
- I need to flip the West vector and multiply it by 4, getting [[0], [-4]].
- I finish by adding what happens to each direction, getting [[3.3], [-4]].

Linear functions are pretty special!

In the example above, we were able to stretch our vectors using a linear function. This is because linear functions aren't sensitive to scale. Swapping North–South and East–West does the same thing to vectors, whether they're represented in steps, blocks, miles, furlongs, or parsecs. That's not true of most functions, though. Consider a function that squares its input, $f(x) = x^2$. The larger x is, the more it gets stretched out.

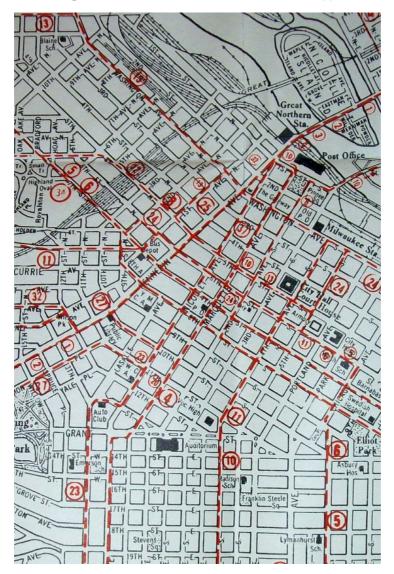
That linear functions work the same way no matter how large or small their inputs are is precisely what lets us break them down piece by piece: once we know how a linear function works at *any* scale, we know how they work at *all* scales.

Thus, I need to look 3.3 blocks North and 4 blocks East on my map.

TIP

Later, we'll see how the bits "0" and "1" can be thought of as directions or vectors, not too different from North or East. In the same way that North and East aren't the best vectors to help you understand Minneapolis, we'll find that "0" and "1" aren't always the best vectors to help you understand quantum computing.

Figure B.5. North and West aren't always the best directions to use if you want to understand where you're going. See this map of downtown Minneapolis, where a large section of the downtown grid is rotated to match the bend in the Mississippi river. Photo by davecito.



This way of understanding linear functions by breaking them apart piece by piece

©Manning Publications Co. To comment go to <u>liveBook</u>

works for rotations, too. If my map has the compass rotated by 45° clockwise (wow, I need a serious lesson in cartography), so that North becomes Northeast, and West becomes Northwest, then I can still figure out where things are piece by piece. Using the same example, the North vector now gets mapped to approximately [[0.707], [0.707]] on the map, and the West vector gets mapped to [[-0.707], [0.707]].

```
When we sum up what happens in the example above, we thus get 3 * [[0.707], [0.707]] + 4 * [[-0.707], [0.707]], which is equal to (3 - 4) * [[0.707], [0]] + (3 + 4) [[0], [0.707]], giving us [[-0.707], [4.95]].
```

It might seem that this has less to do with linearity and more to do with that North and West are somehow special. We could have, however, run through *exactly* the same argument but writing down Southwest as <code>[[1], [0]]</code> and Northwest as <code>[[0], [1]]</code>. This works because Southwest and Northwest are perpendicular to each other, allowing us to break down any other direction as a combination of Northwest and Southwest. Other than ease of reading a compass that we buy off a shelf, there's nothing that makes North or West special. If you've ever tried to drive around downtown Minneapolis (see figure B.5), it quickly becomes apparent that North and West aren't always the best way to understand directions!

Formally, we call any set of vectors that lets us understand directions by breaking them apart piece by piece in this way a *basis*.

NOTE

Technically, we'll be concerned here with what mathematicians call an *orthonormal basis*, as that's most often useful in quantum computing. All that means is that the vectors in a basis are perpendicular to all the other basis vectors and have a length of 1.

Let's try an example of writing a vector in terms of a basis. The vector \vec{v} = [[2],[3] can be written as $2\vec{b}_0 + 3\vec{b}_1$ using the basis $\vec{b}_0 = [[1],[0]]$ and $\vec{b}_1 = [[0],[1]]$.

BASIS

If any vector \vec{v} in d dimensions can be written as a sum of multiples of $\vec{b}_0, \vec{b}_1, \dots, \vec{b}_{d-1}$, we say that $\vec{b}_0, \vec{b}_1, \dots, \vec{b}_{d-1}$ are a **basis**. In two dimensions, one common basis is horizontal and vertical.

More generally, if we know the output of a function f for each vector in a basis, we can compute f for any input. This is similar to how we used truth tables to describe a classical operation by listing the outputs of an operation for each possible input.

Example 13.2. Problem solving with linearity

_Let's say f is a linear function that represents how our map is stretched and twisted, how could we find where we need to go? We want to compute the value f(np.array([[2], [3]])) (a somewhat arbitrary value) given our basis f(np.array([[1], [0]])) (horizontal) and f(np.array([[0], [1]])) (vertical)? We also know from looking parts of the map legend we see that the map warps the horizontal direction to np.array([[1], [1]]) and the vertical direction to np.array([[1], [-1]])

B2.2 Steps to compute f (np. array ([[2], [3]]))

- We use our basis, np.array([[1], [0]]) and np.array([[0], [1]]), to write that np.array([[2], [3]]) is equal to 2 * np.array([[1], [0]]) + 3 * np.array([[0], [1]]).
- Using this new way to write our input to the function, we want to compute f (2 * np.array([[1], [0]]) + 3 * np.array([[0], [1]])).
- Next, we use that f is linear to write f(2 * np.array([[1], [0]]) + 3 * np.array([[0], [1]])) as 2 * f(np.array([[1], [0]])) + 3* f(np.array([[0], [1]])):

- 1 First, we define variables horizontal and vertical to represent the basis we will use to represent [[2], [3]].
- **2** We can write [[2], [3]] by adding multiples of horizontal and vertical.
- We next define how f acts on horizontal and vertical by introducing new variables f_horizontal and f_vertical to represent f (horizontal) and f (vertical), respectively.
- 4 Because f is linear, we can define how it works for [[2], [3]] by replacing horizontal and vertical by the outputs f_horizontal and f_vertical.

```
Exercise B.3 Suppose that you have a linear function g such that g([[1],[0]])=[[2.3],[-3.1]] and g([[0],[1]])=[[-5.2],[0.7]].
```

Using this insight, we can make a table of how a linear function transforms each of its inputs. These tables are called *matrices*, and are complete descriptions of linear functions. If I tell you the matrix for a linear function, then you can compute that function for *any* vector. For example, the transformation from the North/East convention to the East/North convention for map directions transforms the instruction "go one unit North" from being written as [[1], [0]] to being written as [[0], [1]]]. Similarly, the instruction "go one unit East" goes from being written as [[0],

[1]] to being written as [[1], [0]]. If I stack up the outputs for both sets of instructions, I get the following matrix:

Listing 13.4. Stacking the vectors for swapping East/North conventions on a map

TIP This is a very important matrix in quantum computing as well! We will see much more of this matrix throughout the book.

To apply the linear function represented by a matrix to a particular vector, we multiply the matrix and the vector, as illustrated in figure B.6.

While the order in which we add vectors doesn't matter, the order in which we multiply matrices matters quite a lot. If we rotate our map by 90° and then look at it in the mirror, we'll get a very different picture than if we rotate what we see in the mirror by 90°. Both rotation and flipping are linear functions, and so we can write down a matrix for each; let's call them R and F, respectively. If we flip a vector \vec{x} , we get F \vec{x} . Rotating the output gives us RF \vec{x} , a very different vector than if we rotated first, FR \vec{x} .

Figure B.6. How to multiply a matrix by a vector: In this example, the matrix for f tells us that f([1], [0], [0]) is [1], [2], [9].

```
A matrix describing a linear
function f can be thought of as a
stack of a the outputs of f, one for
each column.
For instance, the first column tells
us that f([[1], [0], [0]])
is the column vector
[[1], [2], [9]], while the
second column tells us that
f([[0], [1], [0]]) is
[[3], [4], [8]].
                             1
   2
           4
                  6
                                                               = 2 \times 3 + 4 \times 1 + 6 \times 2
                                                               = 9 \times 3 + 8 \times 1 + 5 \times 2
Just as the first index of
                           The second matrix
                                                  Since the first factor had three rows
a matrix represents its
                           or vector being
                                                  and the second factor had one
rows and the second
                           multiplied is read
                                                  column, the product has three rows
index represents its
                           out in columns.
                                                  and one column.
columns, the first matrix
```

being multiplied is read out row-by-row.

Matrix multiplication formalizes the way that we computed f given its outputs for a

particular set of inputs by "stacking up" the outputs of $\mathfrak f$ for vectors like <code>[[1], [0], [0]]</code> and <code>[[0], [1], [0]]</code>, as illustrated in . While the actual sizes of the matrices and vectors may change, this idea that a matrix can describe a linear transformation stays the same. For the rest of this chapter, we will look linear transformations on vectors of length 2. We can think of each row (the outermost index in NumPy) of a matrix as how the function acts on a particular input.

DEEP DIVE: Why do we multiply functions?

When we multiply a matrix by a vector (or even by a matrix by a matrix), we're doing something that seems a bit odd at first. After all, matrices are another way of representing linear functions, so what does it mean to multiply a function by its input, let alone by another function?

To answer this, it's helpful to go back to ordinary algebra for a moment, in which we have that for any variables a, b, and c, a(b+c)=ab+ac. This property, known as the *distributive property* is fundamental to how multiplication and addition interact with each other. In fact, it's so fundamental, that the distributive property is one of the key ways we define what multiplication is — in number theory and other more abstract parts of math, researchers often work with objects known as *rings*, where all we really know about multiplication is that it distributes over addition. Though posed as an abstract concept, the study of rings and other similar algebraic objects has broad applications, especially in cryptography and error correction.

The distributive property looks very similar to the linearity property, though, that f(x + y) = f(x) + f(y). If we think of f as being a part of a ring, then the distributive property is identical to the linearity property.

Put differently, as much as programmers like to reuse code, mathematicians like to reuse *concepts*. Thinking of multiplying matrices together lets us treat linear functions in the many of the same ways as we're used to from algebra.

Thus, if we want to know the ith element of a vector x that has been rotated by a matrix M, we can find the output of M for each element in X, sum the resulting vectors, and take the i`th element. In NumPy, matrix multiplication is represented by the `@ operator.

NOTE

The code sample below only works in Python 3.5 or later.

Listing 13.5. Matrix multiplication with the @ operator

Exercise B.4

Let X be the matrix [[0, 1], [1, 0]], and let \vec{y} be the vector [[2], [3]]. Using NumPy, compute X \vec{y} and XX.

Why NumPy?

We could have written all of the matrix multiplication above out by hand, but there's a few reasons that it's very nice to work with NumPy instead. Most of the core of NumPy uses constant-time indexing, and is implemented in native code, such that it can take advantage of built-in processor instructions for fast linear algebra. Thus, NumPy will often be much, much faster than manipulating lists by hand. In , we show an example where NumPy can speed up multiplying even very small matrices by $10 \times .$ As we get to larger matrices in Chapters 4 and later, using NumPy over doing things by hand gives us even more of an advantage.

Listing 13.6. Timing NumPy evaluation for matrix multiplication

```
$ ipython
In [1]: def matmul(A, B):
            n_rows_A = len(A)
   . . . :
            n cols A = len(A[0])
            n_rows_B = len(B)
   . . . :
            n cols B = len(B[0])
            assert n_cols_A == n_rows_B
            return [
                [
                    sum(
                         A[idx_row][idx_inner] * B[idx_inner][idx_col]
                         for idx_inner in range(n_cols_A)
                    for idx_col in range(n_cols_B)
                for idx_row in range(n_rows_A)
            ]
   . . . :
   . . . :
In [2]: import numpy as np
In [3]: X = np.array([[0+0j, 1+0j], [1+0j, 0+0j]])
In [4]: Z = np.array([[1+0], 0+0], [0+0], -1+0])
In [5]: matmul(X, Z)
Out[5]: [[0j, (-1+0j)], [(1+0j), 0j]]
In [6]: X @ Z
                                                      8
Out[6]:
array([[ 0.+0.j, -1.+0.j],
       [ 1.+0.j, 0.+0.j]])
In [7]: %timeit matmul(X, Z)
10.3 µs ± 176 ns per loop (mean ± std. dev. of 7 runs, 100000 loops each)
In [8]: %timeit X @ Z
926 ns \pm 4.42 ns per loop (mean \pm std. dev. of 7 runs, 1000000 loops each)
```

- 1 This time, we'll use the IPython interpreter for Python, as it provides a few extra tools that are helpful in this example. Please see Appendix A for instructions on how to install IPython.
- 2 We start by finding the sizes of each matrix that we need to multiply. If we're representing matrices by lists of lists, then each element of the outer list is a row. That is, an n x m matrix has n rows and m columns when written out this way.
- 3 The inner dimensions of both matrices need to agree in order for matrix multiplication to make sense. Thinking of each matrix as representing a linear function, the first index (the number of rows) tells us how large each output is, while the second index (the number of columns) tells us how large each input is. Thus we need the outputs from the first function to be applied (the one on the right) to be of the same size as the inputs to the second function. This line checks that condition.

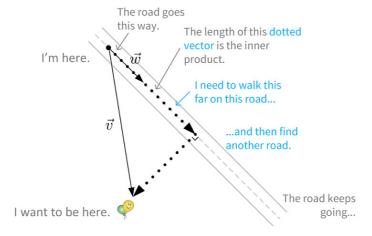
©Manning Publications Co. To comment go to <u>liveBook</u>

- ◆To actually compute the matrix product of A and B, we need to compute each element in the product and pack them into a list of lists.
- **5** We can find each element by summing over where the output from B is passed as input to A, similar to how we represented the product of a matrix with a vector in .
- **6** For comparison, we can import NumPy, which provides us with a matrix multiplication implementation that uses modern processor instructions to accelerate the computation.
- We'll initialize two matrices as NumPy arrays as test cases. We'll see much more about these two particular matrices throughout the book.
- 8 Matrix multiplication in NumPy is represented by the @ operator in Python 3.5 and later.
- **9** The %timeit "magic command" tells IPython to run a small piece of Python code many times and report the average amount of time that it takes.

B.3 PARTY WITH INNER PRODUCTS

There's one last thing we need to worry about in finding the party. Earlier, I said I was ignoring the problem of whether there was a road that would let me go in the direction I needed to, but this is a really bad idea when I'm wandering through an unfamiliar city. To make my way around, I need a way to evaluate how far I should walk along a given road to get where I'm going. Thankfully, linear algebra gives us a tool, the *inner product* to do just that. Inner products are a way of projecting one vector \vec{v} onto another vector \vec{v} , telling us how much of a "shadow" \vec{v} casts on \vec{v} .

Figure B.7. How to find a party with inner products.



We can compute the inner product of two vectors by multiplying their respective elements and summing the result. Note that this multiply-and-sum recipe is the same as what we do in matrix multiplication! Multiplying a matrix that has a single row with a matrix that has a single column does exactly what we need. Thus, to find the projection of \vec{v} onto \vec{w} , we need to turn \vec{v} into a row vector by taking its *transpose*,

written \vec{v}^{T} .

Example The transpose of $\vec{w} = \begin{pmatrix} 2 \\ 3 \end{pmatrix}$ is $\vec{w}^{\rm T} = \begin{pmatrix} 2 \\ 3 \end{pmatrix}$.

NOTE

Later, we'll see that we also need to take the complex conjugate of each element, but we'll set that aside for now.

In particular, the matrix product of \vec{v}^T (the transpose of \vec{v}) with \vec{w} gives us a 1×1 matrix containing the inner product we want. Suppose I need to go 2 blocks south and 3 blocks east, but I can only go on a road that points more south-southeast. Since I still need to travel south, this road helps me get where I need to go. But how far should I walk before this road stops helping?

SQUARE ROOTS AND LENGTHS

The square root of a number x is a number $y = \sqrt{x}$ such that we get x back when we square y, $y^2 = x$. We'll use the square root a lot throughout the book, as square roots are essential to finding the length of vectors. In computer graphics, for instance, quickly finding the lengths of vectors is essential to making games work

(see https://en.wikipedia.org/wiki/Fast_inverse_square_root for some fun history about how square roots are used in gaming).

Whether vectors describe how we get to parties, or as we'll see later, those vectors describe the information that a quantum bit represents, we'll use square roots to reason about their lengths.

- f 0 In this case, ec v is the vector describing where I need to go, namely two blocks north and three blocks east
- 2 If the road available points southeast, then it goes one block south for every block east that it goes.
- 3 We can find the length of a vector using Pythagoreas' theorem by taking the sum of the absolute values of each element, then taking the square root. In NumPy, this is done with the np.linalg.norm function, as the length of a vector is sometimes also called its norm.

- **4** The length of [[1], [-1]] is thus $\sqrt{(+1)^2 + (-1)^2} = \sqrt{2} \approx 1.4142$
- **5**Thus, when we define \vec{w} to be the direction southeast, we need to divide by $\sqrt{2}$.
- **6** Checking, we see that the length of \vec{w} is now approximately 1.
- **7** The transpose turns \vec{v}] = [[%u22122],[%u22123]]intothe"row"[[%u22122,%u22123]].
- 8 We can then multiply the transpose of \vec{v} with \vec{w} the same way as we multiplied matrices with vectors earlier.
- **9** Doing so, we see that I need to walk $1/\sqrt{2} \approx 0.707$ blocks along this road before it stops helping me to the party.

Exercise B.5

Given a vector [[2], [3]], find a vector that points in the same direction but with length 1.

HINT: You can either do this by using an inner product, or the np.linalg.norm function.

Finally we have made it to the party (only slightly late) and are ready to try out that new doorbell!