# Introduction

Why yet another book on quantum information theory? Like many lecturers we began writing this text because none of the alternatives seemed quite right. This book is aimed squarely at undergraduate physics students who want a brief but reasonably thorough introduction to the exciting ideas of quantum information, including its applications in computation and communication. It is based on a short course we have taught to fourth-year students at Oxford University since 2004; for the most part it only assumes knowledge of elementary quantum mechanics and linear algebra, and so could even be taught to third-year undergraduates. A brief revision guide to quantum mechanics is provided as an appendix, which should cover any minor points that have been missed.

As the title implies the book is structured in three parts, starting with the basics of quantum information and then applying this to quantum computation and quantum communication. Part I is self-contained, but contains only the barest hints of the exciting applications which attract many people to this field and so might prove unsatisfying on its own. Parts II and III draw heavily on Part I, but are largely independent of each other, and it would be perfectly possible to study only one of these two without the other.

As this text is aimed at physics undergraduates, we believe that it is vital to cover experimental techniques, rather than merely presenting quantum information as a series of abstract quantum operations. We have, however, concentrated on the basic ideas underlying each approach, rather than worrying about particular experimental details. Our aim is not to explain how quantum information processing can actually be achieved, but rather to provide the reader with enough of an introduction to start understanding more specialist sources. The choice of implementations described inevitably reflects our own research interests, but is broad enough to provide an introduction to many important fields. We have, however, almost completely neglected implementations with solid state devices.

In the first two parts of the book we have deliberately taken a cheerfully optimistic approach to quantum information, largely treating quantum systems in terms of a highly idealized picture. We do, of course, consider the problem of decoherence, and briefly describe methods that can be used for error correction, but in general we simply assume that these issues can be ignored. This allows us to address some interesting quantum algorithms without becoming bogged down in experimental details or notational complexities. This approach is, however, inappropriate for quantum communication: while quantum computations generally either work or don't work, communication protocols are dominated by considerations of reliability and efficiency. Furthermore, current experimental implementations are often quite inefficient, and cannot simply be treated as ideal. For this reason we begin Part III with a brief primer on information theory, and also take the opportunity to introduce some slightly more sophisticated notations and mathematical techniques.

The experimental sections will make more sense to a reader with some basic familiarity with atomic physics and optics, but this is not essential. The determined theorist could largely neglect the four experimental chapters describing atom and spin implementations, but would probably be happier with a more focused and rigorous text. The chapter on photon techniques is, however, essential reading before tackling Part III, as quantum communication protocols are closely linked to the photon technologies used to implement them.

Extensive references are included at the end of each chapter for readers who wish to take these ideas further; these largely concentrate on textbooks and review papers, rather than the primary literature, but we have also included a number of landmark papers and papers of particular pedagogical interest.

We have provided a range of exercises at the end of each chapter. Most of these should prove fairly straightforward, but worked answers are available. In some cases we do not prove certain key results, but leave these proofs as an exercise for the reader. Quantum information is a field best understood by doing these basic calculations which lead to familiarity with the underlying ideas.

Our text inevitably goes slightly beyond our course as originally taught, but we have endeavored to keep the range of topics covered very similar. We have, however, been tempted to add one major extension, in the form of a chapter on more advanced quantum algorithms. The material in this chapter is somewhat more challenging than the rest of the text, but even so only provides a very basic introduction to this fascinating field.

We are grateful to all our colleagues who have taught us many of the topics explained in this text. Any errors are, of course, entirely our own.

# PART I

# QUANTUM INFORMATION

# 1 Quantum bits and quantum gates

Classical information processing is performed using *bits*, which are just two-state systems, with the two states called 0 and 1. By grouping bits together we can represent arbitrary pieces of information, and by manipulating these bits we can perform arbitrary computations. The corresponding basic element used in quantum information is the quantum bit, or *qubit*. This is simply a quantum system with two orthonormal basis states, which we shall call $|0\rangle$ and $|1\rangle$.

There are many possible physical implementations of a qubit, such as spin states of electrons or atomic nuclei, charge states of quantum dots, atomic energy levels, vibrational states of groups of atoms, polarization states of photons, or paths in an interferometer. At this stage the physical implementation is not important: the idea of a qubit is to abstract the discussion away from physical details. Taking the standard approach of quantum information theory, we shall begin by not worrying too much about the properties of these states, or even what their energies are; we shall simply assume that they are eigenstates of the system's Hamiltonian with known eigenvalues (that is, known energies). This approach allows us to concentrate on the fundamental properties of the system, without considering all the tedious details.[1]

We can in principle perform classical information processing on our quantum system by using the two states $|0\rangle$ and $|1\rangle$ as our logical states 0 and 1 and proceeding in the usual fashion, giving rise to the field of *reversible computing*, which will be explored briefly in Part II. This, however, misses the point. A qubit is not confined to these two states, but can be found in arbitrary superposition states. Although it is not immediately obvious what a state like

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{1.1}$$

where $\alpha$ and $\beta$ are complex numbers, actually means in information processing terms, it is clear that quantum bits are in some sense more powerful than their classical equivalents. Quantum information processing is, of course, the art of exploiting these superposition states to perform information processing tasks which are impossible for classical systems. Just as the real power of classical information processing requires groups of bits, the real advantages of quantum information processing only become clear in systems with two or more qubits; for simplicity, however, we are confining ourselves to single isolated qubits at the moment.

Throughout this book we will assume that the reader is familiar with elementary quantum mechanics, and in particular with Dirac's notation for writing quantum states as kets, as

[1] A more careful approach is necessary when considering how efficiently quantum information protocols can actually be implemented, as will be seen in Part III.

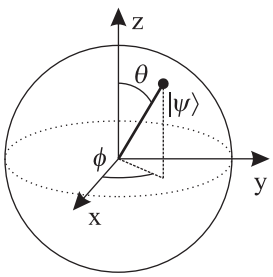


Fig. 1.1 The Bloch sphere: the state $|\psi\rangle$ of a single qubit can be represented as a point on the surface of a sphere with radius 1, or equivalently by the Bloch vector which points from the origin to this point. The state can then be completely described by its co-latitude $\theta$ (the angle between the Bloch vector and the $z$ axis) and azimuth angle $\phi$ (the angle between the projection of the Bloch vector into the $xy$ plane and the $x$ axis).

used above. No particularly sophisticated knowledge is necessary, and if the following two sections make sense then you probably know enough quantum mechanics to understand this book. A brief reminder of some key terms can be found in Appendix A.

## 1.1 The Bloch sphere

The enormous flexibility of a single qubit in comparison with a classical bit can be most clearly seen using the *Bloch sphere* description of a qubit (Figure 1.1). This also provides a simple but powerful way of visualizing the behavior of a qubit. We begin by looking again at the general state of a single qubit, equation (1.1), and noting that this ket must have unit norm, so that $|\alpha|^2 + |\beta|^2 = 1$. The fact that the state does not change under global phase shifts, so that $e^{i\gamma}|\psi\rangle$ is completely equivalent to $|\psi\rangle$ for *any* real number $\gamma$, means that we can always choose $\alpha$ to be *real*, and the normalization constraint is easily imposed by making $\alpha$ and $\beta$ depend on the cosine and sine of a single parameter. A particularly useful form is to write

$$|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\phi}\sin(\theta/2)|1\rangle, \tag{1.2}$$

where $0 \leq \theta \leq \pi$ and $0 \leq \phi < 2\pi$. Note that $\theta = 0$ corresponds to $|\psi\rangle = |0\rangle$, and $\theta = \pi$ corresponds to $|\psi\rangle = |1\rangle$; in these extreme cases the value of $\phi$ is irrelevant.

There is an obvious analogy between the variables $\theta$ and $\phi$ used above and those used in spherical polar coordinates. Clearly any ket $|\psi\rangle$ can be associated with a single point on the surface of a sphere of radius 1 with co-latitude and azimuth angles $\theta$ and $\phi$; this sphere is usually called the Bloch sphere. Alternatively (and entirely equivalently) a state can be represented as a unit vector (connecting the origin and a point on the Bloch sphere), called the Bloch vector.

**Example 1.1** The two basis states $|0\rangle$ and $|1\rangle$, which correspond to the states 0 and 1 of a classical bit, lie at the north and south poles of the Bloch sphere respectively, while a qubit can lie anywhere at all on the surface. Another important group of states is the set of equally weighted superpositions, with $|\alpha| = |\beta| = 1/\sqrt{2}$, which lie on the equator of the Bloch sphere, with the exact position determined by the relative phase of $\alpha$ and $\beta$. Unlike $|0\rangle$ and $|1\rangle$, these states have no classical interpretation.

Like any other quantum state, the state of a qubit will evolve under the influence of its Hamiltonian $\mathcal{H}$. The time-dependent Schrödinger equation

$$\mathrm{i}\hbar\frac{\partial}{\partial t}|\psi\rangle = \mathcal{H}|\psi\rangle \tag{1.3}$$

has the solution

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle \tag{1.4}$$

with

$$U(t) = \exp(-\mathrm{i}(\mathcal{H}/\hbar)t). \tag{1.5}$$

The evolution of quantum states can also be described using the compact notation

$$|\psi\rangle \xrightarrow{\mathcal{H}t} U|\psi\rangle. \tag{1.6}$$

Since $\mathcal{H}$ is Hermitian, the evolution operator $U$, usually called the *propagator*, must be unitary.

The discussion above assumes that the Hamiltonian is time-independent, that is it does not vary with time. This will not be true in a quantum computer, which is controlled by varying the Hamiltonian. In many cases, however, the Hamiltonian is *piecewise constant*, that is it has a constant value for some finite length of time, and is then replaced by a different constant value for another finite time period, and so on. In this case the evolution can be described using a series of propagators

$$|\psi\rangle \xrightarrow{\mathcal{H}_1 t_1} \xrightarrow{\mathcal{H}_2 t_2} \xrightarrow{\mathcal{H}_3 t_3} U_3 U_2 U_1 |\psi\rangle \tag{1.7}$$

with $U_1 = \exp[-\mathrm{i}(\mathcal{H}_1/\hbar)t_1]$ and so on. It is, of course, possible to combine the sequence of propagators into a single propagator, $U = U_3 U_2 U_1$. Note that the sequence of Hamiltonians is normally written with time running from left to right (that is the leftmost Hamiltonian is the first to be applied), while the sequence of propagators is always written from right to left, as the rightmost propagator is applied first. The situation is much more complicated when the Hamiltonian varies continuously with time; it is possible to write down a formal solution of the form of equation (1.7), but this is not generally a useful approach. For the moment this issue will simply be ignored.

The fact that any propagator describing the evolution of a quantum system is unitary has several significant consequences. Firstly it means that every propagator has an inverse, and so quantum evolution is *reversible*. (One exception to this general principle is *measurement*, which is discussed in more detail below.) Secondly unitary transformations are *length*

*preserving* and can in general be thought of as *rotations* of the vector describing the quantum state. Since qubits live on the Bloch sphere, the evolution of an isolated qubit under any Hamiltonian corresponds to a rotation of the vectors on the Bloch sphere.

## 1.2 Density matrices and Pauli matrices

It is frequently convenient to describe the state of a qubit using a vector, written using the basis states $|0\rangle$ and $|1\rangle$ (the computational basis). The basis states take the simple forms

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{and} \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{1.8}$$

(There is a potential ambiguity in any description of quantum bits, as to whether $|0\rangle$ and $|1\rangle$ are defined as shown here, or the other way round; fundamentally, of course, the choice does not matter, as long as one is consistent.) In this basis equation (1.1) can be written as

$$|\psi\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \tag{1.9}$$

while the corresponding bra can be written as

$$\langle\psi| = \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix}, \tag{1.10}$$

as a bra is the adjoint of the corresponding ket, and for a matrix the adjoint is the complex conjugate of the transpose.

Bras and kets are frequently combined by taking the *inner product*, such as

$$\langle\psi|\psi\rangle = \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \alpha^*\alpha + \beta^*\beta = 1 \tag{1.11}$$

but they can also be combined using the *outer product*

$$|\psi\rangle\langle\psi| = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix} = \begin{pmatrix} \alpha\alpha^* & \alpha\beta^* \\ \beta\alpha^* & \beta\beta^* \end{pmatrix}. \tag{1.12}$$

This outer product is called a *density matrix* description of the state. As we will see later, density matrices can provide a very useful approach to describe qubits whose states are at least partly unknown, called *mixed states*, but for the moment we will use them simply to explore an alternative to the ket notation for *pure states*.

It is obvious from the form of equation (1.12) that the density matrix describing a qubit is Hermitian, and has trace one; these are in fact general properties which apply to all density matrices. A two-by-two matrix can always be expanded as a weighted sum of four basic matrices (a matrix basis), and perhaps the most useful basis is provided by the Pauli matrices

$$\sigma_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -\mathrm{i} \\ \mathrm{i} & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{1.13}$$

where the usual set of three matrices has been extended to include the *identity matrix* $\sigma_0 = \mathbb{1}$. As the Pauli matrices are Hermitian, a density matrix can be written as

$$|\psi\rangle\langle\psi| = \tfrac{1}{2}\left(\sigma_0 + s_x\sigma_x + s_y\sigma_y + s_z\sigma_z\right), \tag{1.14}$$

where $s_x$, $s_y$ and $s_z$ are three *real* coefficients. This might seem excessive, as we know that any pure state can be described using only two real numbers ($\theta$ and $\phi$), but it is easily shown that $s_x$, $s_y$ and $s_z$ are not completely independent, being the three components of a vector of unit length; this vector is identical to the Bloch vector, discussed above.

Qubits can also be found in mixed states, which are just weighted sums of pure states of the form

$$\rho = \sum_n P_n|\psi_n\rangle\langle\psi_n|, \tag{1.15}$$

where $P_n \geq 0$ is the contribution of the pure state $|\psi_n\rangle\langle\psi_n|$ to the mixture (the probability of the pure state occurring in the mixture). Clearly such mixed states are Hermitian, and as the probabilities of the various contributions must sum to one ($\sum_n P_n = 1$), the density matrix must have trace one. It can be shown that any mixed state of a single qubit corresponds to a point *inside* the Bloch sphere.

It is useful to be able to calculate the evolution of states described using a density matrix rather than a ket vector. This problem can be addressed directly by solving the Liouville–von Neumann equation (the density matrix equivalent of the time-dependent Schrödinger equation), but it is simple to proceed by analogy. The evolution of a bra vector is clearly closely related to the evolution of the corresponding ket vector

$$(U|\psi\rangle)^\dagger = \langle\psi|U^\dagger, \tag{1.16}$$

and so the density matrix description of a pure state evolves according to

$$|\psi\rangle\langle\psi| \xrightarrow{\mathcal{H}_t} U|\psi\rangle\langle\psi|U^\dagger \tag{1.17}$$

and the linearity of the operations guarantees that a mixed state will evolve in the same way.

We have already noted that the Pauli matrices are Hermitian, and thus provide a natural basis for describing the density matrix corresponding to a qubit. In the same way, the fact that any Hamiltonian is Hermitian means that any Hamiltonian for a single qubit can be written as a weighted sum of the four Pauli matrices, equation (1.13), where the weights are *real*. This means that the Pauli matrices provide a natural language for describing single-qubit Hamiltonians as well as single-qubit states. Furthermore the Pauli matrices are unitary, and so correspond to possible propagators. As we shall see later, the Pauli matrices viewed as propagators correspond to important quantum logic gates. It might seem that using the Pauli matrices to describe quantum states, Hamiltonians, propagators, and logic gates will inevitably lead to confusion, but in practice such problems rarely occur.

The fact that the Pauli matrices are *both* unitary and Hermitian has the interesting consequence that

$$\sigma_\alpha^2 = \sigma_0, \tag{1.18}$$

where $\sigma_\alpha$ are the usual Pauli matrices, with $\alpha$ equal to $x$, $y$ or $z$. This observation can be combined with the series expansion of an exponential operator to show that

$$\exp(-\mathrm{i}\theta\,\sigma_\alpha) = \cos(\theta)\sigma_0 - \mathrm{i}\sin(\theta)\sigma_\alpha \tag{1.19}$$

without diagonalizing any matrices, making it easy to calculate many single-qubit propagators.

Finally we note that the propagator corresponding to a Hamiltonian that is some multiple of $\sigma_0$ is simply a global phase shift, which has no physical significance. In essence this occurs because adding multiples of $\sigma_0$ corresponds to moving the zero-point of the energy scale, which has no physical significance.

## 1.3 Quantum logic gates

The basic idea of quantum information processing is that information is stored in quantum bits and processed by quantum logic gates. Just as classical logic gates take classical bits from one state to another, so quantum logic gates take qubits from one state to another. This can be achieved by modifying the system's Hamiltonian, by applying additional *control fields* to the background Hamiltonian which underlies the system.

Applying Hamiltonians will cause qubits to evolve under unitary transformations, which are reversible. With classical bits there are only two reversible gates that act on a single bit: the NOT gate, which takes a bit in state 0 into state 1 and *vice versa*, and the IDENTITY gate, which just leaves the bit unchanged. (It may seem excessive to consider trivial gates such as IDENTITY, but the formalism works better if they are included.) There are also two irreversible gates, SET which sets a bit to 1 whatever its initial state, and CLEAR which sets a bit to 0. Clearly these two cannot be achieved with unitary transformations, and so we will neglect them for the moment.

Returning to the two reversible gates, we must first find unitary propagators that implement them. Clearly $\sigma_0$ will perform IDENTITY as

$$\begin{pmatrix}1&0\\0&1\end{pmatrix}\begin{pmatrix}1\\0\end{pmatrix} = \begin{pmatrix}1\\0\end{pmatrix} \quad \text{and} \quad \begin{pmatrix}1&0\\0&1\end{pmatrix}\begin{pmatrix}0\\1\end{pmatrix} = \begin{pmatrix}0\\1\end{pmatrix} \tag{1.20}$$

while $\sigma_x$ corresponds to NOT as

$$\begin{pmatrix}0&1\\1&0\end{pmatrix}\begin{pmatrix}1\\0\end{pmatrix} = \begin{pmatrix}0\\1\end{pmatrix} \quad \text{and} \quad \begin{pmatrix}0&1\\1&0\end{pmatrix}\begin{pmatrix}0\\1\end{pmatrix} = \begin{pmatrix}1\\0\end{pmatrix}. \tag{1.21}$$

We now have to find Hamiltonians which can give rise to these propagators. Obviously $\sigma_0$ can in principle be achieved simply by doing nothing at all, but in fact the IDENTITY gate is slightly more subtle than it might seem, as the state of the qubit will continue to evolve under the background Hamiltonian even when no additional control fields are applied, and unless the IDENTITY gate is instantaneous this background evolution must be considered. Achieving $\sigma_x$ is only slightly more difficult: using equation (1.19)

$$\exp(-\mathrm{i}\pi\sigma_x/2) = -\mathrm{i}\sigma_x \tag{1.22}$$