


Quantum Information Processing

A hands-on primer

*Emanuel Knill, Raymond Laflamme, Howard N. Barnum, Diego A. Dalvit, Jacek J. Dziarmaga,
James E. Gubernatis, Leonid Gurvits, Gerardo Ortiz, Lorenza Viola, and Wojciech H. Zurek*



Quantum information processing, Science of—The theoretical, experimental and technological areas covering the use of quantum mechanics for communication and computation
—Kluwer Encyclopedia of Mathematics, Supplement II

Research conducted in the last few decades has established that quantum information, or information based on quantum mechanics, has capabilities that exceed those of traditional “classical” information. For example, in communication, quantum information enables quantum cryptography, which is a method for communicating in secret. Secrecy is guaranteed because eavesdropping attempts necessarily disturb the exchanged quantum information without revealing the content of the communication. In computation, quantum information enables efficient simulation of quantum physics, a task for which general-purpose, efficient, classical algorithms are not known to exist. Quantum information also leads to efficient algorithms for factoring large numbers, which is believed to be difficult for classical computers. An efficient factoring algorithm would break the security of commonly used public-key cryptographic codes used for authenticating and securing Internet communications. Yet another application of quantum information improves the efficiency with which unstructured search problems can be solved. Quantum unstructured search may make it possible to solve significantly larger instances of optimization problems, such as the scheduling and traveling salesman problems.

Because of the capabilities of quantum information, the science of quantum information processing is now a prospering, interdisciplinary field focused on better understanding the possibilities and limitations of the underlying theory, on developing new applications of quantum information, and on physically realizing controllable quantum devices. The purpose of this primer is to provide an elementary introduction to quantum information processing (see Part II), and then to briefly explain how we hope to exploit the advantages of quantum information (see Part III). These two sections can be read independently. For reference, we have included a glossary of the main terms of quantum information (see page 33).

When we use the word “information,” we generally think of the things we can talk about, broadcast, write down, or otherwise record. Such records can exist in many forms, such as sound waves, electrical signals in a telephone wire, characters on paper, pit patterns on an optical disk, or magnetization on a computer hard disk. A crucial property of information is that it is fungible: It can be represented in many different physical forms and easily converted from one form to another without changing its meaning. In this sense, information is independent of the devices used to represent it but requires at least one physical representation in order to be useful.

We call the familiar information stored in today’s computers classical or deterministic to distinguish it from quantum information. It is no accident that classical information is the basis of all human knowledge. Any information passing through our senses is best modeled by classical discrete or continuous information. Therefore, when considering any other kind of information, we need to provide a method for extracting classically meaningful information. We begin by recalling the basic ideas of classical information in a way that illustrates the general procedure for building an information-processing theory.

Part I: Classical Information

The basic unit of classical deterministic information is the bit, an abstract entity or system that can be in one of the two states symbolized by 0 and 1. At this point, the symbols for the two states have no numeric meaning. That is why we have used a font different from that for the numbers 0 and 1. By making a clear distinction between the bit and its states, we emphasize that a bit should be physically realized as a system or device whose states correspond to the ideal bit's states. For example, if you are reading this primer on paper, the system used to realize a bit is a reserved location on the surface, and the state depends on the pattern of ink (0 or 1) in that location. In a computer, the device realizing a bit can be a combination of transistors and other integrated-circuit elements with the state of the bit determined by the distribution of charge.

In order to make use of information, it must be possible to manipulate (or process) the states of information units. The elementary operations that can be used for this purpose are called gates. Two one-bit gates are the **not** and **reset** gates. Applying the **not** gate to a bit has the effect of flipping the state of the bit. For example, if the initial state of the bit is 0, then the state after applying **not** is **not** (0) = 1. We can present the effect of the gate in the following form:

Initial State		Final State
0	→	not (0) = 1 , and
1	→	not (1) = 0 .

(1)

The **reset** gate sets the state to 0 regardless of the input:

Initial State		Final State
0	→	reset (0) = 0 , and
1	→	reset (1) = 0 .

(2)

By applying a combination of **not** and **reset** gates, one can transform the state of a bit in every possible way.

Information units can be combined to represent more information. Bits are typically combined into sequences. The states of such a sequence are symbolized by strings of state symbols for the constituent bits. For example, a two-bit sequence can be in one of the following four states: 00, 01, 10, and 11. The different bits are distinguished by their position in the sequence.

The one-bit gates can be applied to any bit in a sequence. For example, the **not** gate applied to the second bit of a three-bit sequence in the state 011 changes the state to 001.

One-bit gates act independently on each bit. To compute with multiple bits, we need gates whose action can correlate the states of two or more bits. One such gate is the **nand** (“not and”) gate, which acts on two bits in a bit sequence. Its effect is to set the state of the first bit to 0 if both the first and the second bit are 1; otherwise, it sets it to 1. Here is what happens when **nand** is applied to two consecutive bits:

Initial State		Final State
00	→	nand (00) = 10 ,
01	→	nand (01) = 11 ,
10	→	nand (10) = 10 , and
11	→	nand (11) = 01 .

(3)

The **nand** gate can be applied to any two bits in a sequence. For example, it can be applied to the fourth and second bits (in this order) of four bits, in which case the initial state 1101 is transformed to 1100, setting the fourth bit to 0.

Other operations on bit sequences include adding a new bit to the beginning (prepend) or end (append) of a sequence. The new bit is always initialized to 0. It is also possible to discard the first or last bit regardless of its state. Versions of these operations that are conditional on the state of another bit may also be used. An example is the conditional append operation: “If the k^{th} bit is in the state 0, then append a bit.”

The gates just introduced suffice for implementing arbitrary state transformations of a given bit sequence. Instructions for applying gates in a particular order are called a circuit. An important part of investigations in information processing is to determine the minimum resources required to perform information-processing tasks. For a given circuit, the two primary resources are the number of gates and the total number of bits used. The circuit complexity of a desired transformation is the minimum number of gates needed to implement it.

The model of computation defined by the ability to apply gates in a fixed sequence is called the circuit model. Classical computation extends the circuit model by providing a means for repeating blocks of instructions indefinitely or until a desired condition is achieved. In principle, it is possible to conceive of a general-purpose computer as a device that repeatedly applies the same circuit to the beginnings of several bit sequences. In this article, we take for granted a traditional programmable computer based on classical information. Thus, a quantum algorithm is a program written for such a computer with additional instructions for applying gates to quantum information. The computational power of this model is equivalent to that of other general-purpose models of quantum computation, such as quantum Turing machines (Yao 1993).

For an introduction to algorithms and their analysis, refer to Thomas Cormen et al. (1990). Christos Papadimitriou wrote (1994) a useful textbook on computational complexity with an introduction to classical computation and computational machine models.

Part II: Quantum Information

The foundations of an information-processing theory can be constructed by the procedure we followed in the previous section:

1. Define the basic unit of information.
2. Give the means for processing one unit.
3. Describe how multiple units can be combined.
4. Give the means for processing multiple units.
5. Show how to convert the content of any of the extant units to classical information.

Note that the last step was not required for classical information processing.

In this section, we follow the general procedure for defining an information-processing theory to introduce quantum information processing. A simple example that exhibits the advantages of quantum information is given in the section “The Parity Problem” on page 21. A version of the quantum factoring algorithm is described immediately following this article in “From Factoring to Phase Estimation” on page 38.

The Quantum Bit

The fundamental resource and basic unit of quantum information is the quantum bit (qubit), which behaves like a classical bit enhanced by the superposition principle (see discussion in this section). From a physical point of view, a qubit is represented by an ideal two-state quantum system. Examples of such systems include photons (vertical and horizontal polarization), electrons and other spin-1/2 systems (spin-up and -down), and systems defined by two energy levels of atoms or ions. From the beginning, the two-state system played a central role in studies of quantum mechanics. It is the simplest quantum system, and in principle, all other quantum systems can be modeled in the state space of collections of qubits.

From the information-processing point of view, a qubit’s state space contains the two “logical,” or computational, states $|0\rangle$ and $|1\rangle$. The so-called “ket” notation for these states was introduced by Paul Dirac, and its variations are widely used in quantum physics. One can think of the pair of symbols $|$ and \rangle as representing the qubit system. Their content specifies a state for the system. In this context, 0 and 1 are system-independent state labels. When, say, 0 is placed within the ket, the resulting expression $|0\rangle$ represents the corresponding state of a specific qubit.

The initial state of a qubit is always one of the logical states. Using operations to be introduced later, we can obtain states that are superpositions of the logical states. Superpositions can be expressed as sums $\alpha|0\rangle + \beta|1\rangle$ over the logical states with complex coefficients. The complex numbers α and β are the amplitudes of the superposition. The existence of such superpositions of distinguishable states of quantum systems is one of the basic tenets of quantum theory and is called the superposition principle. Another way of writing a general superposition is as a vector:

$$\alpha|0\rangle + \beta|1\rangle \leftrightarrow \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \quad (4)$$

where the two-sided arrow is used to denote the correspondence between expressions that mean the same thing.

The qubit states that are superpositions of the logical states are called pure states: A superposition $\alpha|0\rangle + \beta|1\rangle$ is a pure state if the corresponding vector has length 1, that

is, $|\alpha|^2 + |\beta|^2 = 1$. Such a superposition or vector is said to be normalized. (For a complex number given by $\gamma = x + iy$, one can evaluate $|\gamma|^2 = x^2 + y^2$. Here, x and y are the real and imaginary part of γ , and the symbol i is a square root of -1 , that is, $i^2 = -1$. The conjugate of γ is $\bar{\gamma} = x - iy$. Thus, $|\bar{\gamma}|^2 = \gamma\bar{\gamma}$. Here are a few examples of states given in both the ket and vector notation:

$$|\psi_1\rangle = \frac{(|0\rangle + |1\rangle)}{\sqrt{2}} \leftrightarrow \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}, \quad (5)$$

$$|\psi_2\rangle = \frac{3}{5}|0\rangle - \frac{4}{5}|1\rangle \leftrightarrow \begin{pmatrix} 3/5 \\ -4/5 \end{pmatrix}, \quad \text{and} \quad (6)$$

$$|\psi_3\rangle = \frac{i3}{5}|0\rangle - \frac{i4}{5}|1\rangle \leftrightarrow \begin{pmatrix} i3/5 \\ -i4/5 \end{pmatrix}. \quad (7)$$

The state $|\psi_3\rangle$ is obtained from $|\psi_2\rangle$ by multiplication with i . It turns out that two states cannot be distinguished if one of them is obtained by multiplying the other by a phase $e^{i\theta}$. Note how we have generalized the ket notation by introducing expressions such as $|\psi\rangle$ for arbitrary states.

The superposition principle for quantum information means that we can have states that are sums of logical states with complex coefficients. There is another, more familiar type of information, whose states are combinations of logical states. The basic unit of this type of information is the probabilistic bit (pbit). Intuitively, a pbit can be thought of as representing the as-yet-undetermined outcome of a coin flip. Since we need the idea of probability to understand how quantum information converts to classical information, we briefly introduce pbits.

A pbit's state space is a probability distribution over the states of a bit. One very explicit way to symbolize such a state is by using the expression $\{p:0, (1-p):1\}$, which means that the pbit has probability p of being 0 and $1-p$ of being 1. Thus, a state of a pbit is a probabilistic combination of the two logical states, where the coefficients are nonnegative real numbers summing to 1. A typical example is the unbiased coin in the process of being flipped. If tail and head represent 0 and 1, respectively, the coin's state is $\{1/2:0, 1/2:1\}$. After the outcome of the flip is known, the state collapses to one of the logical states 0 and 1. In this way, a pbit is converted to a classical bit. If the pbit is probabilistically correlated with other pbits, the collapse associated with learning the pbit's logical state changes the overall probability distribution by a process called conditioning on the outcome.

A consequence of the conditioning process is that we never actually "see" a probability distribution. We only see classical deterministic bit states. According to the frequency interpretation of probabilities, the original probability distribution can only be inferred after one looks at many independent pbits in the same state $\{p:0, (1-p):1\}$: In the limit of infinitely many pbits, p is given by the fraction of pbits seen to be in the state 0. As we will explain, we can never see a general qubit state either. For qubits, there is a process analogous to conditioning. It is called measurement and converts qubit states to classical information.

Information processing with pbits has many advantages over deterministic information processing with bits. One advantage is that algorithms are often much easier to design and

analyze if they are probabilistic. Examples include many optimization and physics simulation algorithms. In some cases, the best available probabilistic algorithm is more efficient than any known deterministic algorithm. An example is an algorithm for determining whether a number is prime or not. It is not known whether every probabilistic algorithm can be derandomized efficiently. There are important communication problems that can be solved probabilistically but not deterministically. For a survey of these algorithms, see Rajiv Gupta (1994a).

What is the difference between bits, pbits, and qubits? One way to visualize the difference and see the enrichment provided by pbits and qubits is shown in Figure 1.

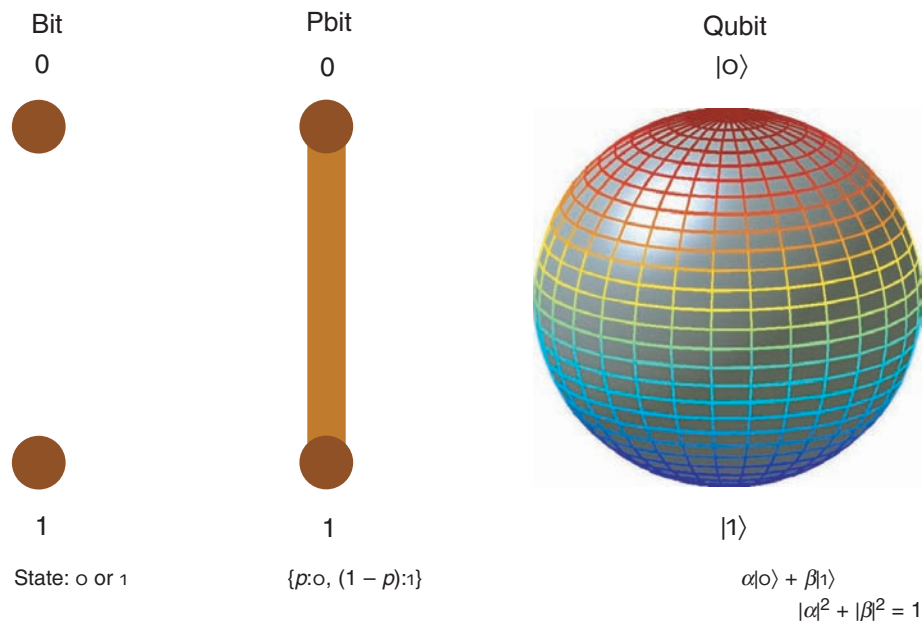


Figure 1. Comparing State Spaces of Different Information Units

The states of a bit correspond to two points. The states of a pbit can be thought of as convex combinations of a bit's states and therefore can be visualized as lying on the line connecting the two bit states. A qubit's pure states correspond to the surface of the unit sphere in three dimensions, where the logical states correspond to the poles. This representation of qubit states is called the Bloch sphere. The explicit correspondence is discussed at the end of the section "Mixtures and Density Operators." Also refer to the definition and use of the Bloch sphere in the article "NMR and Quantum Information Processing" on page 226. There, the correspondence between the pure states and the sphere is physically motivated and comes from a way of viewing a spin-1/2 system as a small quantum magnet. Intuitively, a state is determined by the direction of the north pole of the magnet.

Processing One Qubit

The quantum version of the **not** gate for bits exchanges the two logical states; that is, using ket notation,

$$\text{not}(\alpha|0\rangle + \beta|1\rangle) = \alpha|1\rangle + \beta|0\rangle = \beta|0\rangle + \alpha|1\rangle. \quad (8)$$

In vector notation, this equation becomes

$$\mathbf{not} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \beta \\ \alpha \end{pmatrix}. \quad (9)$$

Another way of expressing the effect of **not** is by multiplying the vector by a matrix representing **not**,

$$\mathbf{not} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \beta \\ \alpha \end{pmatrix}, \quad (10)$$

so that we can identify the action of **not** with the matrix

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

An even simpler gate is the one that does nothing. We call it the **noop** gate, and its matrix form is the identity matrix, as shown in the following equation:

$$\mathbf{noop} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}. \quad (11)$$

The **noop** and **not** gates are reversible. In other words, we can undo their actions by applying other gates. For example, the action of the **not** gate can be undone by another **not** gate. The action of every reversible quantum gate can be represented by matrix multiplication, where the matrix has the additional property of preserving vector lengths. Such matrices are called unitary and are characterized by the equation $A^\dagger A = \mathbb{1}$, where A^\dagger is the conjugate transpose of A and $\mathbb{1}$ is the identity matrix. (The conjugate transpose of a matrix is computed by flipping that matrix across the main diagonal and conjugating the complex numbers). For gates represented by a matrices, the unitarity condition is necessary and sufficient for ensuring that pure states get mapped to pure states.

Because qubit states can be represented as points on a sphere, reversible one-qubit gates can be thought of as rotations of the Bloch sphere. This is why such quantum gates are often called rotations. As explained in detail on page 232 in the article “NMR and Quantum Information Processing”, rotations around the x -, y -, and z -axis are in a sense generated by the three Pauli matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \text{and} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (12)$$

each of which represents a one-qubit gate. For example, a rotation around the x -axis by an angle ϕ is given by $e^{-i\sigma_x\phi/2} = \cos(\phi/2)\mathbb{1} - i\sin(\phi/2)\sigma_x$. To obtain this identity, one can use the power series for e^A , $e^A = \sum_{k=0}^{\infty} (1/k!)A^k$, and exploit the fact that $\sigma_x^2 = \mathbb{1}$ to simplify the expression. Here are some gates that can be defined with the help of rotations:

$$90^\circ \text{ } x\text{-rotation: } \mathbf{rotx}_{90^\circ} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix},$$

$$90^\circ \text{ } y\text{-rotation: } \mathbf{roty}_{90^\circ} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix},$$

$$\phi \text{ } z\text{-rotation: } \mathbf{rotx}_\phi = \begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix}, \text{ and}$$

$$\text{Hadamard gate: } \mathbf{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (13)$$

The rotation gates often show up in controlling spins or ions with radio-frequency pulses or lasers. The Hadamard gate is used primarily by quantum programmers. It can be expressed as a product of a 90° y -rotation and σ_z .

To check directly that the rotation gates are reversible, one can determine their inverses. In this case and as expected, the inverse of a rotation is the rotation around the same axis in the opposite direction. For example, the inverses of the \mathbf{roty}_{90° and \mathbf{rotx}_ϕ gates are given by

$$\mathbf{roty}_{-90^\circ} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}, \text{ and } \mathbf{rotx}_{-\phi} = \begin{pmatrix} e^{i\phi/2} & 0 \\ 0 & e^{-i\phi/2} \end{pmatrix}. \quad (14)$$

Another useful property of the rotation gates is that the angles add when rotations are applied around the same axis. For example, $\mathbf{rotx}_\phi \mathbf{rotx}_\theta = \mathbf{rotx}_{\phi+\theta}$.

The Bra-Ket Notation for Logic Gates. The ket notation can be extended so that we can write gates in a compact form that readily generalizes to multiple qubits. To do so, we have to introduce expressions such as $\langle\psi| = \alpha\langle 0| + \beta\langle 1|$. This is called the “bra” notation. The terminology comes from the term “bracket.” The bra is the left, and the ket is the right part of a matched pair of brackets. From the vector point of view, $\langle\psi|$ corresponds to the row vector (α, β) . Note that a column vector multiplied by a row vector yields a matrix. In the bra-ket notation, this corresponds to multiplying a ket $|\psi\rangle$ by a bra $\langle\phi|$, written as $|\psi\rangle\langle\phi|$. Since this represents an operator on states, we expect to be able to compute the effect of $|\psi\rangle\langle\phi|$ on a state $|\phi\rangle$ by forming the product. To be able to evaluate such products with one-qubit bras and kets, we need the following two rules: distributivity and inner-product evaluation.

Distributivity

You can rewrite sums and products using distributivity. For example,

$$\left(\frac{3}{5}\langle 0| + \frac{4}{5}\langle 1| \right) i|1\rangle = \frac{i3}{5}\langle 0||1\rangle + \frac{i4}{5}\langle 1||1\rangle. \quad (15)$$

Observe that we can combine the amplitudes of terms, but we cannot rearrange the order of the bras and kets in a product.

Inner-Product Evaluation

The product of a logical bra and a logical ket is evaluated according to the identities

$$\begin{aligned}\langle 0|0\rangle &= 1 \\ \langle 0|1\rangle &= 0, \\ \langle 1|0\rangle &= 0, \text{ and} \\ \langle 1|1\rangle &= 1.\end{aligned}\tag{16}$$

It follows that for logical states, if a bra multiplies a ket, the result cancels unless the states match, in which case the answer is 1. Applying inner-product evaluation to Equation (15) results in

$$\frac{i3}{5}\langle 0|1\rangle + \frac{i4}{5}\langle 1|1\rangle = \frac{i3}{5}0 + \frac{i4}{5}1 = \frac{i4}{5}.\tag{17}$$

To simplify the notation, we can omit one of the two vertical bars in products such as $\langle a|b\rangle$ and write $\langle a|b\rangle$.

To understand inner-product evaluation, think of the expressions as products of row and column vectors. For example,

$$\langle 0|1\rangle \leftrightarrow \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = 0.\tag{18}$$

That is, as vectors, the two states $|0\rangle$ and $|1\rangle$ are orthogonal. In general, if $|\phi\rangle$ and $|\psi\rangle$ are states, then $\langle\phi|\psi\rangle$ is the inner product, or “overlap,” of the two states. In the expression for the overlap, we compute $\langle\phi|$ from $|\phi\rangle = \bar{\alpha}|0\rangle + \bar{\beta}|1\rangle$ by conjugating the coefficients and converting the logical kets to bras: $\langle\phi| = \alpha\langle 0| + \beta\langle 1|$. In the vector representation, this is the conjugate transpose of the column vector for $|\phi\rangle$, so the inner product agrees with the usual one. Two states are orthogonal if their overlap is zero. We write $|\phi\rangle^\dagger = \langle\phi|$ and $\langle\phi|^\dagger = |\phi\rangle$.

Every linear operator on states can be expressed with the bra-ket notation. For example, the bra-ket expression for the **noop** gate is **noop** = $|0\rangle\langle 0| + |1\rangle\langle 1|$. To apply **noop** to a qubit, you multiply its state on the left by the bra-ket expression

$$\begin{aligned}\text{noop}(\alpha|0\rangle + \beta|1\rangle) &= (|0\rangle\langle 0| + |1\rangle\langle 1|)(\alpha|0\rangle + \beta|1\rangle) \\ &= |0\rangle\langle 0|(\alpha|0\rangle + \beta|1\rangle) + |1\rangle\langle 1|(\alpha|0\rangle + \beta|1\rangle) \\ &= \alpha|0\rangle\langle 0|0\rangle + \beta|0\rangle\langle 0|1\rangle + \alpha|1\rangle\langle 1|0\rangle + \beta|1\rangle\langle 1|1\rangle \\ &= \alpha|0\rangle 1 + \beta|0\rangle 0 + \alpha|1\rangle 0 + \beta|1\rangle 1 \\ &= \alpha|0\rangle + \beta|1\rangle.\end{aligned}\tag{19}$$

One way to think about an operator such as $|a\rangle\langle b|$ is to notice that, when it is used to operate on a ket expression, the $\langle b|$ picks out the matching kets in the state, which are

then changed to $|a\rangle$. For example, we can write the **not** operation as **not** = $|0\rangle\langle 1| + |1\rangle\langle 0|$.

The coefficients of the $|a\rangle\langle b|$ in a bra-ket representation of a gate correspond to matrix entries in the matrix representation. The relationship is defined by

$$\alpha_{00}|0\rangle\langle 0| + \alpha_{01}|0\rangle\langle 1| + \alpha_{10}|1\rangle\langle 0| + \alpha_{11}|1\rangle\langle 1| \leftrightarrow \begin{pmatrix} \alpha_{00} & \alpha_{01} \\ \alpha_{10} & \alpha_{11} \end{pmatrix}. \quad (20)$$

Two Quantum Bits

Some states of two quantum bits can be symbolized by the juxtaposition (or multiplication) of the states of each quantum bit. In particular, the four logical states $|0\rangle|0\rangle$, $|0\rangle|1\rangle$, $|1\rangle|0\rangle$, and $|1\rangle|1\rangle$ are acceptable pure states for two quantum bits. In these expressions, we have distinguished the qubits by position (first or second). It is easier to manipulate state expressions if we explicitly name the qubits, say, A and B. We can then distinguish the kets by writing, for example, $|\psi\rangle_A$ for a state of qubit A. Now, the state $|0\rangle|1\rangle$ can be written with explicit qubit names (or labels) as

$$|0\rangle_A |1\rangle_B = |1\rangle_B |0\rangle_A = |01\rangle_{AB} = |10\rangle_{BA}. \quad (21)$$

Having explicit labels allows us to unambiguously reorder the states in a product of states belonging to different qubits. We say that kets for different qubits “commute.”

So far, we have seen four states of two qubits, which are the logical states that correspond to the states of two bits. As in the case of one qubit, we can use the superposition principle to get all the other pure states. Each state of two qubits is therefore of the form

$$\alpha|00\rangle_{AB} + \beta|01\rangle_{AB} + \gamma|10\rangle_{AB} + \delta|11\rangle_{AB}, \quad (22)$$

where α , β , γ , and δ are complex numbers. Again, there is a column vector form for the state,

$$\begin{pmatrix} \alpha \\ \beta \\ \gamma \\ \delta \end{pmatrix}, \quad (23)$$

and this vector has to be of unit length, that is, $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$. When using the vector form for qubit states, one has to be careful about the convention used for ordering the coefficients.

Other examples of two-qubit states in ket notation are the following:

$$\begin{aligned}
|\psi_1\rangle_{AB} &= \frac{1}{\sqrt{2}}(|0\rangle_A + |1\rangle_A)|1\rangle_B \ . \\
|\psi_2\rangle_{AB} &= \frac{1}{\sqrt{2}}(|0\rangle_A - |1\rangle_A) \frac{1}{\sqrt{2}}(|0\rangle_B + i|1\rangle_B) \\
&= \frac{1}{2}(|00\rangle_{AB} + i|01\rangle_{AB} - |10\rangle_{AB} - i|11\rangle_{AB}) \ . \\
|\psi_3\rangle_{AB} &= \frac{1}{\sqrt{2}}(|00\rangle_{AB} + |11\rangle_{AB}) \ . \\
|\psi_4\rangle_{AB} &= \frac{1}{\sqrt{2}}(|01\rangle_{AB} - |10\rangle_{AB}) \ .
\end{aligned} \tag{24}$$

The first two of these states have the special property that they can be written as a product $|\phi_1\rangle_A|\phi_2\rangle_B$ of a state of qubit A and a state of qubit B. The second expression for $|\psi_2\rangle$ shows that the product decomposition is not always easy to see. Such states are called product states. The last two states, $|\psi_3\rangle_{AB}$ and $|\psi_4\rangle_{AB}$, are two of the famous Bell states. They have no such representation as a product of independent states of each qubit. They are said to be entangled because they contain a uniquely quantum correlation between the two qubit systems. Pbits can also have correlations that cannot be decomposed into product states, but the entangled states have additional properties that make them very useful. For example, if Alice and Bob each have one of the qubits that together are in the state $|\psi_3\rangle_{AB}$, they can use them to create a secret bit for encrypting their digital communications (see the article “Quantum State Entanglement” on page 52).

Processing Two Qubits

The simplest way of modifying the state of two qubits is to apply one of the one-qubit gates. If the gates are expressed in the bra-ket notation, all we need to do is add qubit labels so that we know which qubit each bra or ket belongs to. For example, the **not** gate for qubit B is written as

$$\mathbf{not}^{(B)} = |0\rangle_B^B \langle 1| + |1\rangle_B^B \langle 0| \ . \tag{25}$$

The labels for bra expressions occur as left superscripts. To apply expressions like this to states, we need one more rule, namely, commutation.

Commutation

Kets and bras with different labels can be interchanged in products (they commute). This property is demonstrated by the following example:

$$\begin{aligned}
(|0\rangle_B^B \langle 1|) |01\rangle_{AB} &= |0\rangle_B^B \langle 1| |0\rangle_A |1\rangle_B \\
&= |0\rangle_A |0\rangle_B^B \langle 1| |1\rangle_B \\
&= |0\rangle_A |0\rangle_B^B \langle 1| |1\rangle_B \\
&= |0\rangle_A |0\rangle_B = |00\rangle_{AB} \ .
\end{aligned} \tag{26}$$

Note that we cannot merge the two vertical bars in expressions such as ${}^B\langle 1||0\rangle_A$ because the two terms belong to different qubits. The bars can only be merged when the expression is an inner product, which requires that the two terms belong to the same qubit.

With the rules for bra-ket expressions in hand, we can apply the **not** gate to one of our Bell states to see how it acts:

$$\begin{aligned}
 \mathbf{not}^{(B)} \frac{1}{\sqrt{2}} (|00\rangle_{AB} + |11\rangle_{AB}) &= (|0\rangle_B^B \langle 1| + |1\rangle_B^B \langle 0|) \frac{1}{\sqrt{2}} (|00\rangle_{AB} + |11\rangle_{AB}) \\
 &= \frac{1}{\sqrt{2}} \left(|0\rangle_B^B \langle 1| (|00\rangle_{AB} + |11\rangle_{AB}) + |1\rangle_B^B \langle 0| (|00\rangle_{AB} + |11\rangle_{AB}) \right) \\
 &= \frac{1}{\sqrt{2}} (|0\rangle_B^B \langle 1||00\rangle_{AB} + |0\rangle_B^B \langle 1||11\rangle_{AB} + |1\rangle_B^B \langle 0||00\rangle_{AB} + |1\rangle_B^B \langle 0||11\rangle_{AB}) \\
 &= \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B^B \langle 1||0\rangle_B + |1\rangle_A |0\rangle_B^B \langle 1||1\rangle_B + |0\rangle_A |1\rangle_B^B \langle 0||0\rangle_B + |1\rangle_A |1\rangle_B^B \langle 0||1\rangle_B) \\
 &= \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B 0 + |1\rangle_A |0\rangle_B 1 + |0\rangle_A |1\rangle_B 1 + |1\rangle_A |1\rangle_B 0) \\
 &= \frac{1}{\sqrt{2}} (|1\rangle_A |0\rangle_B + |0\rangle_A |1\rangle_B) = \frac{1}{\sqrt{2}} (|01\rangle_{AB} + |10\rangle_{AB}) .
 \end{aligned} \tag{27}$$

The effect of the gate was to flip the state symbols for qubit B, which results in another Bell state.

The gate $\mathbf{not}^{(B)}$ can also be written as a 4×4 matrix acting on the vector representation of a two-qubit state. However, the relationship between this matrix and the one-qubit matrix is not as obvious as for the bra-ket expression. The matrix is

$$\mathbf{not}^{(B)} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{28}$$

which swaps the top two and bottom two entries of a state vector.

One way to see the relationship between the one- and two-qubit representations of the gate $\mathbf{not}^{(B)}$ is to notice that because the **noop** gate acts as the identity and because we can act on different qubits independently, $\mathbf{noop}^{(A)} \mathbf{not}^{(B)} \equiv \mathbf{not}^{(B)}$. The matrix for $\mathbf{not}^{(B)}$ can be expressed as a Kronecker product (\otimes) of the matrices for **noop** and **not**:

$$\begin{aligned}
\text{noop}^{(A)} \text{not}^{(B)} &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\
&= \begin{pmatrix} 1 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} & 0 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\ 0 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} & 1 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \end{pmatrix} \\
&= \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.
\end{aligned} \tag{29}$$

The Kronecker product of two matrices expands the first matrix by multiplying each entry by the second matrix. A disadvantage of the matrix representation of quantum gates is that it depends on the number and order of the qubits. However, it is often easier to visualize what the operation does by writing down the corresponding matrix.

One cannot do much with one-bit classical gates. Similarly, the utility of one-qubit gates is limited. In particular, it is not possible to obtain a Bell state starting from $|\text{OO}\rangle_{AB}$ or any other product state. We therefore need to introduce at least one two-qubit gate not expressible as the product of two one-qubit gates. The best-known such gate is the controlled-not (**cnot**) gate. Its action can be described by the statement, “if the first bit is 1, flip the second bit; otherwise, do nothing.” The bra-ket and matrix representations for this action are

$$\begin{aligned}
\text{cnot}^{(AB)} &= |\text{O}\rangle_A^A \langle \text{O}| + |1\rangle_A^A \langle 1| \left(|\text{O}\rangle_B^B \langle 1| + |1\rangle_B^B \langle \text{O}| \right) \\
&= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.
\end{aligned} \tag{30}$$

The **cnot** gate is reversible because its action is undone if a second **cnot** is applied. This outcome is easy to see by computing the square of the matrix for **cnot**, which yields the identity matrix. As an exercise in manipulating bras and kets, let us calculate the product of two **cnot** gates by using the bra-ket representation:

$$\begin{aligned}
\text{cnot}^{(AB)} \text{cnot}^{(AB)} &= \left(|\text{O}\rangle_A^A \langle \text{O}| + |1\rangle_A^A \langle 1| \left(|\text{O}\rangle_B^B \langle 1| + |1\rangle_B^B \langle \text{O}| \right) \right) \\
&\quad \times \left(|\text{O}\rangle_A^A \langle \text{O}| + |1\rangle_A^A \langle 1| \left(|\text{O}\rangle_B^B \langle 1| + |1\rangle_B^B \langle \text{O}| \right) \right). \tag{31}
\end{aligned}$$

The first step is to expand this expression by multiplying out. Expressions such as $|\mathbf{o}\rangle_A^A \langle \mathbf{o} | | \mathbf{1} \rangle_A^A \langle \mathbf{1} |$ cancel because of the inner-product evaluation rule ${}^A \langle \mathbf{o} | \mathbf{1} \rangle_A = 0$. One can also reorder bras and kets with different labels and rewrite $|\mathbf{o}\rangle_A^A \langle \mathbf{o} | | \mathbf{o} \rangle_A^A \langle \mathbf{o} | = |\mathbf{o}\rangle_A^A \langle \mathbf{o} |$ to get

$$\begin{aligned}
 \mathbf{cnot}^{(AB)} \mathbf{cnot}^{(AB)} &= |\mathbf{o}\rangle_A^A \langle \mathbf{o} | + |\mathbf{1}\rangle_A^A \langle \mathbf{1} | \left(|\mathbf{o}\rangle_B^B \langle \mathbf{1} | + |\mathbf{1}\rangle_B^B \langle \mathbf{o} | \right) \left(|\mathbf{o}\rangle_B^B \langle \mathbf{1} | + |\mathbf{1}\rangle_B^B \langle \mathbf{o} | \right) \\
 &= |\mathbf{o}\rangle_A^A \langle \mathbf{o} | + |\mathbf{1}\rangle_A^A \langle \mathbf{1} | \left(|\mathbf{o}\rangle_B^B \langle \mathbf{o} | + |\mathbf{1}\rangle_B^B \langle \mathbf{1} | \right) \\
 &= |\mathbf{o}\rangle_A^A \langle \mathbf{o} | + |\mathbf{1}\rangle_A^A \langle \mathbf{1} | \mathbf{noop}^{(B)} \\
 &\equiv |\mathbf{o}\rangle_A^A \langle \mathbf{o} | + |\mathbf{1}\rangle_A^A \langle \mathbf{1} | \\
 &= \mathbf{noop}^{(A)} \\
 &\equiv 1 .
 \end{aligned} \tag{32}$$

Here we used the fact that, when the bra-ket expression for **noop** is applied to the ket expression for a state, it acts the same as (here denoted by the symbol \equiv) multiplication by the number 1.

Using Many Quantum Bits

To use more than two, say, five qubits, we can just start with the state $|\mathbf{o}\rangle_A |\mathbf{o}\rangle_B |\mathbf{o}\rangle_C |\mathbf{o}\rangle_D |\mathbf{o}\rangle_E$ and apply gates to any one or two of these qubits. For example, $\mathbf{cnot}^{(DB)}$ applies the **cnot** operation from qubit D to qubit B. Note that the order of D and B in the label for the **cnot** operation matters. In the bra-ket notation, we simply multiply the state with the bra-ket form of $\mathbf{cnot}^{(DB)}$ from the left. One can express everything in terms of matrices and vectors, but now the vectors have length $2^5 = 32$, and the Kronecker product expression for $\mathbf{cnot}^{(DB)}$ requires some reordering to enable inserting the operation so as to act on the intended qubits. Nevertheless, to analyze the properties of all reversible (that is, unitary) operations on these qubits, it is helpful to think of the matrices because a lot of useful properties about unitary matrices are known. One important result from this analysis is that every matrix that represents a reversible operation on quantum states can be expressed as a product of the one- and two-qubit gates introduced so far. We say that this set of gates is universal.

For general-purpose computation, it is necessary to have access to arbitrarily many qubits. Instead of assuming that there are infinitely many from the start, it is convenient to have an operation to add a new qubit, namely, **add**. To add a new qubit labeled X in the state $|\mathbf{o}\rangle_X$, apply $\mathbf{add}^{(X)}$ to the current state. This operation can only be used if there is not already a qubit labeled X. To implement the $\mathbf{add}^{(X)}$ operation in the bra-ket notation, we multiply the ket expression for the current state by $|\mathbf{o}\rangle_X$.

Qubit Measurements

In order to classically access information about the state of qubits, we use the measurement operation **meas**. This is an intrinsically probabilistic process that can be applied to any extant qubit. For information processing, one can think of **meas** as a subroutine or function whose output is either 0 or 1. The output is called the measurement outcome. The probabilities of the measurement outcomes are determined by the current state. The state of the qubit being measured is collapsed to the logical state corresponding to the outcome. Suppose we have just one qubit, currently in the state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Measurement of this qubit has the effect

$$\text{meas}(\alpha|0\rangle + \beta|1\rangle) = \begin{cases} 0: |0\rangle & \text{with probability } |\alpha|^2 \\ 1: |1\rangle & \text{with probability } |\beta|^2 \end{cases} . \quad (33)$$

The classical output is given before the new state for each possible outcome. This measurement behavior explains why the amplitudes have to define unit length vectors: Up to a phase, they are associated with square roots of probabilities.

For two qubits, the process is more involved. Because of possible correlations between the two qubits, the measurement affects the state of the other one too, similar to conditioning for pbits after one “looks” at one of them. As an example, consider the state

$$|\psi\rangle_{AB} = \frac{2}{3}|01\rangle_{AB} + \frac{i2}{3}|10\rangle_{AB} + \frac{1}{3}|00\rangle_{AB} . \quad (34)$$

To figure out what happens when we measure qubit A, we first rewrite the current state in the form $\alpha|0\rangle_A|\phi_0\rangle_B + \beta|1\rangle_A|\phi_1\rangle_B$, where $|\phi_0\rangle_B$ and $|\phi_1\rangle_B$ are pure states for qubit B. It is always possible to do that. For the example given in Equation (34),

$$\begin{aligned} |\psi\rangle_{AB} &= \frac{2}{3}|0\rangle_A|1\rangle_B + \frac{1}{3}|0\rangle_A|0\rangle_B + \frac{i2}{3}|1\rangle_A|0\rangle_B \\ &= |0\rangle_A \left(\frac{2}{3}|1\rangle_B + \frac{1}{3}|0\rangle_B \right) + |1\rangle_A \frac{i2}{3}|0\rangle_B \\ &= \frac{\sqrt{5}}{3}|0\rangle_A \left(\frac{2}{\sqrt{5}}|1\rangle_B + \frac{1}{\sqrt{5}}|0\rangle_B \right) + \frac{i2}{3}|1\rangle_A|0\rangle_B , \end{aligned} \quad (35)$$

$$\text{so } \alpha = \frac{\sqrt{5}}{3}, \beta = \frac{i2}{3}, |\phi_0\rangle_B = \frac{2}{\sqrt{5}}|1\rangle_B + \frac{1}{\sqrt{5}}|0\rangle_B, \text{ and } |\phi_1\rangle_B = |0\rangle_B .$$

The last step required pulling out the factor of $\sqrt{5/3}$ to make sure that $|\phi_0\rangle_B$ is properly normalized for a pure state. Now, that we have rewritten the state, the effect of measuring qubit A can be given as follows:

$$\text{meas}^{(A)}\left(\alpha|\phi_0\rangle_A|\phi_0\rangle_B + \beta|\phi_1\rangle_A|\phi_1\rangle_B\right) = \begin{cases} 0: |\phi_0\rangle_A|\phi_0\rangle_B & \text{with probability } |\alpha|^2 \\ 1: |\phi_1\rangle_A|\phi_1\rangle_B & \text{with probability } |\beta|^2 \end{cases}. \quad (36)$$

For the example, the measurement outcome is 0 with probability $5/9$, in which case the state collapses to $|\phi_0\rangle_A(1/\sqrt{5}|\phi_0\rangle_B + 2/\sqrt{5}|\phi_1\rangle_B)$. The outcome is 1 with probability $4/9$, in which case the state collapses to $|\phi_1\rangle_A|\phi_1\rangle_B$. The probabilities add up to 1 as they should.

The same procedure works for figuring out the effect of measuring one of any number of qubits. Say we want to measure qubit B among qubits A, B, C, D, currently in state $|\psi\rangle_{ABCD}$. First, rewrite the state in the form $\alpha|\phi_0\rangle_B|\phi_0\rangle_{ACD} + \beta|\phi_1\rangle_B|\phi_1\rangle_{ACD}$, making sure that the ACD superpositions are pure states. Then, the outcome of the measurement is 0 with probability $|\alpha|^2$ and 1 with probability $|\beta|^2$. The collapsed states are $|\phi_0\rangle_B|\phi_0\rangle_{ACD}$ and $|\phi_1\rangle_B|\phi_1\rangle_{ACD}$, respectively.

Probabilities of the measurement outcomes and the new states can be calculated systematically. For example, to compute the probability and state for outcome 0 of $\text{meas}^{(A)}$ given the state $|\psi\rangle_{AB}$, one can first obtain the unnormalized ket expression $|\phi'_0\rangle_B = {}^A\langle\phi_0|\psi\rangle_{AB}$ by using the rules for multiplying kets by bras. The probability is given by $p_0 = {}^B\langle\phi'_0|\phi'_0\rangle_B$, and the collapsed, properly normalized pure state is

$$\frac{|\phi_0\rangle_A|\phi'_0\rangle_B}{\sqrt{p_0}} = \frac{|\phi_0\rangle_A {}^A\langle\phi_0|\psi\rangle_{AB}}{\sqrt{p_0}}. \quad (37)$$

The operator $P_0 = |\phi_0\rangle_A {}^A\langle\phi_0|$ is called a projection operator or projector for short. If we perform the same computation for the outcome 1, we find the projector $P_1 = |\phi_1\rangle_A {}^A\langle\phi_1|$. The two operators satisfy $P_a^2 = P_a$, $P_a^\dagger = P_a$, and $P_0 + P_1 = \mathbb{1}$. In terms of the projectors, the measurement's effect can be written as follows:

$$\text{meas}^{(A)}|\psi\rangle_{AB} = \begin{cases} 0: P_0|\psi\rangle_{AB}/\sqrt{p_0} & \text{with probability } p_0 \\ 1: P_1|\psi\rangle_{AB}/\sqrt{p_1} & \text{with probability } p_1 \end{cases}, \quad (38)$$

where $p_0 = {}^{AB}\langle\psi|P_0|\psi\rangle_{AB}$ and $p_1 = {}^{AB}\langle\psi|P_1|\psi\rangle_{AB}$. In quantum mechanics, any pair of projectors satisfying the properties given above is associated with a potential measurement whose effect can be written in the same form. This is called a binary von Neumann, or projective, measurement.

Mixtures and Density Operators

The measurement operation reads out information from qubits to pbits. What if we discard the pbit that contains the measurement outcome? The result is that the qubits are in a probabilistic mixture of two pure states. Such mixtures are a generalization of pure states. The obvious way to think about a mixture is that we have a probability distribution over pure quantum states. For example, after discarding the pbit and qubit A in Equation (36), we can write the state of B as $\rho = \{|\alpha|^2:|\phi_0\rangle_B, |\beta|^2:|\phi_1\rangle_B\}$, using the notation for probability distributions introduced earlier.

Mixtures frequently form when irreversible operations are used, such as measurement. Except for measurement, the quantum gates we have introduced so far are reversible and therefore transform pure states to pure states so that no mixtures can be formed. One of the fundamental results of reversible classical and quantum computation is that there is no loss in power in using only reversible gates. Specifically, it is possible to change a computation that includes irreversible operations to one that accomplishes the same goal, has only reversible operations, and is efficient in the sense that it uses at most polynomial additional resources. However, the cost of using only reversible operations is not negligible. In particular, for ease of programming and, more important, when performing repetitive error-correction tasks (see the article on this subject on page 188), the inability to discard or reset qubits can be very inconvenient. We therefore introduce additional operations that enable resetting and discarding.

Although resetting has a so-called thermodynamic cost (think of the heat generated by a computer), it is actually a simple operation. The **reset** operation applied to qubit A can be thought of as the result of first measuring A, then flipping A if the measurement outcome is $|1\rangle$, and finally discarding the measurement result. Using the notation of Equation (36), the effect on a pure state $|\psi\rangle_{AB}$ is given by

$$\text{reset}^{(A)}|\psi\rangle_{AB} = \left\{ |\alpha|^2 : |\phi_0\rangle_A |\phi_0\rangle_B, |\beta|^2 : |\phi_1\rangle_A |\phi_1\rangle_B \right\} . \quad (39)$$

To apply **reset** to an arbitrary probability distribution, you apply it to each of that distribution's pure states and combine the results to form an expanded probability distribution. The **discard**^(A) operation is **reset**^(A) followed by discarding qubit A. In the expression for the state after **reset**^(A), therefore, all the $|\phi\rangle_A$ are removed. It is an important fact that every physically realizable quantum operation, whether reversible or not, can be expressed as a combination of **add** operations, gates from the universal set, and **discard** operations.

The representation of mixtures using probability distributions over pure states is redundant. That is, many probability distributions are physically indistinguishable. A nonredundant description of a quantum state can be obtained if density operators are used. The density operator for the mixture ρ in Equation (39) is given by

$$\hat{\rho} = |\alpha|^2 |\phi_0\rangle_B \langle\phi_0| + |\beta|^2 |\phi_1\rangle_B \langle\phi_1| . \quad (40)$$

The general rule for calculating the density operator from a probability distribution is the following: For each pure state $|\phi\rangle$ in the distribution, calculate the operators $|\phi\rangle\langle\phi|$ and sum them weighted by their probabilities.

There is a way to apply gates to the density operators defined by states. If the gate acts by the unitary operator U , then the effect of applying it to $\hat{\rho}$ is given by $U\hat{\rho}U^\dagger$ where U^\dagger is the conjugate transpose of U . (In the bra-ket expression for U , U^\dagger is obtained by replacing all complex numbers by their conjugates, and terms such as $|\phi\rangle\langle\phi|$, by $|\phi\rangle\langle\phi|$.)

The relationship between a qubit's state space and a sphere can be explained in terms of the qubit's density operators. In matrix form, this operator is a 2×2 matrix, which can be written uniquely as a sum $(\mathbb{1} + x\sigma_x + y\sigma_y + z\sigma_z)/2$. One can check that, if the density operator $|\psi\rangle\langle\psi|$ for a qubit's pure state is written as such a sum,

$$|\psi\rangle\langle\psi| = (\mathbb{1} + x\sigma_x + y\sigma_y + z\sigma_z)/2, \quad (41)$$

then the vector (x, y, z) thus obtained is on the surface of the unit sphere in three dimensions. In fact, for every vector (x, y, z) on the unit sphere, there is a unique pure state satisfying Equation (41). Since the density operators for mixtures are arbitrary, convex (that is, probabilistic) sums of pure states, the set of (x, y, z) thus obtained for mixtures fills out the unit ball. The rotations introduced earlier modify the vector (x, y, z) in the expected way, by rotation of the vector around the appropriate axis. See the article "NMR and Quantum Information Processing" on page 232 for more details.

Quantum Computation

The model of computation defined by the one- and two-qubit gates and the operations **add**, **meas**, and **discard** qubits is called the quantum network model. A sequence of instructions for applying these operations is called a quantum network. Quantum computation extends the network model by providing a means for repeating blocks of instructions. Such means can be specified by a formal machine model of computation. There are several such models of classical and quantum computers. One of the best known is the Turing machine, which has a quantum analogue, the quantum Turing machine. This model has its uses for formal studies of computation and complexity but is difficult to program. Fortunately, as mentioned in Part I, there is no loss of computational power if the means for repeating instructions is provided by a classical computer that can apply gates and other operations to qubits. A general quantum algorithm is a program written for such a computer.

There are three practical methods that can be used to write quantum networks and algorithms. The first is to use the names for the different operations and algebraically multiply them. The second is to draw quantum networks, which are pictorial representations of the sequence of steps in time, somewhat like flowcharts without loops. The third is to use a generic programming language enhanced with statements for accessing and modifying quantum bits. The first two methods work well as long as the sequence is short and we do not use many operations that depend on measurement outcomes or require loops. They are often used to describe subroutines of longer algorithms presented either in words or by use of the third method.

To see how to use the different methods and also to illustrate the power of quantum computation, we work out a short quantum algorithm that solves the so-called parity problem.

The Parity Problem. Given is a “black-box” quantum operation $\mathbf{BB}^{(ABC)}$ that has the following effect when applied to a logical basis state:

$$\mathbf{BB}^{(ABC)}|a_A a_B\rangle_{AB}|a_C\rangle_C = |a_A a_B\rangle_{AB}|a_C \oplus (b_A a_A \oplus b_B a_B)\rangle_C, \quad (42)$$

where b_A and b_B are 0 or 1. The actual values of b_A and b_B are unknown. The problem is to determine what b_A and b_B are by using the black box only once.

The terminology and definition of the operation $\mathbf{BB}^{(ABC)}$ require explanation. In computation, we say that an operation is a black box, or an “oracle,” if we have no access whatsoever to how the operation is implemented. In a black-box problem, we are promised that the black box implements an operation from a specified set of operations. In the case of the parity problem, we know that the operation is to add one of four possible parities (see below). The problem is to determine that parity by using the black box in a network. Black-box problems serve many purposes. One is to study the differences between models of computation, just as we are about to do. In fact, black-box problems played a crucial role in the development of quantum algorithms by providing the first and most convincing examples of the power of quantum computers (Bernstein and Vazirani 1993, Simon 1994). Some of these examples involve generalizations of the parity problem. Another purpose of black-box problems is to enable us to focus on what can be learned from the input/output behavior of an operation without having to analyze its implementation. Focusing on the input/output behavior is useful because, in many cases of interest, it is difficult to exploit knowledge of the implementation in order to determine a desirable property of the operation. A classical example is the well-known satisfiability problem, in which we are given a classical circuit with one output bit and we need to determine whether there is an input for which the output is 1. Instead of trying to analyze the circuit, one can look for and use a general-purpose black-box search algorithm to find the satisfying input.

In the definition of the effect of $\mathbf{BB}^{(ABC)}$, the operation \oplus is addition modulo 2, so $1 \oplus 1 = 0$, and all the other sums are as expected. As the state symbols have a numeric meaning now, we will use the number font for states. To see what \mathbf{BB} does, suppose that b_A and b_B are both 1. Then \mathbf{BB} adds (modulo 2) the parity of the logical state in AB to the logical state of C. The parity of a logical state is 0 if the number of 1s is even and 1 if it is odd. The action of \mathbf{BB} for this example is given by

$$\begin{aligned} \mathbf{BB}^{(ABC)}|00\rangle_{AB}|0\rangle_C &= |00\rangle_{AB}|0\rangle_C \\ \mathbf{BB}^{(ABC)}|01\rangle_{AB}|0\rangle_C &= |01\rangle_{AB}|0 \oplus 1\rangle_C \\ &= |01\rangle_{AB}|1\rangle_C \\ \mathbf{BB}^{(ABC)}|10\rangle_{AB}|1\rangle_C &= |10\rangle_{AB}|1 \oplus 1\rangle_C \\ &= |10\rangle_{AB}|0\rangle_C \\ \mathbf{BB}^{(ABC)}|11\rangle_{AB}|0\rangle_C &= |11\rangle_{AB}|0\rangle_C \end{aligned} \quad (43)$$

The action of the black box is extended to superpositions by linear extension. This means that to apply **BB** to a superposition of the logical states, we simply apply it to each logical summand and add the results. Different values of b_A and b_B correspond to different parities. For example, if $b_A = 1$ and $b_B = 0$, then the parity of the state in A is added to the state in C. In this sense, what is added is the parity of a subset of the two qubits AB. Thus, one way of thinking about the problem is that we wish to find out which subset's parity the black box is using.

We can give an algorithm that solves the parity problem using each of the three methods for describing quantum networks. Here is an algebraic description of a solution, **qparity**^(ABC), given as a product of quantum gates that involves one use of the black box. We defer the explanation of why this solution works until after we show how to represent the algorithm pictorially, using quantum networks.

$$\mathbf{qparity}^{(ABC)} = \mathbf{meas}^{(B)} \mathbf{H}^{(B)} \mathbf{meas}^{(A)} \mathbf{H}^{(A)} \mathbf{BB}^{(ABC)} \mathbf{H}^{(C)} \mathbf{not}^{(C)} \mathbf{add}^{(C)} \mathbf{H}^{(B)} \mathbf{add}^{(B)} \mathbf{H}^{(A)} \mathbf{add}^{(A)} . \quad (44)$$

The output of the algorithm is given by the classical outputs of the measurements of qubit A, which yield b_A , and of qubit B, which yield b_B . As is conventional, in writing products of linear operators, the order of application in Equation (44) is right to left, as in a product of matrices applied to a column vector. This order of terms in a product is, however, counterintuitive, particularly for operations to be performed sequentially. It is therefore convenient to use left to right notation, as is done in describing laser or radio-frequency pulse sequences, and to put dots between gates to indicate left to right order:

$$\mathbf{qparity}^{(ABC)} = \mathbf{add}^{(A)} . \mathbf{H}^{(A)} . \mathbf{add}^{(B)} . \mathbf{H}^{(B)} . \mathbf{add}^{(C)} . \mathbf{not}^{(C)} . \mathbf{H}^{(C)} . \mathbf{BB}^{(ABC)} . \mathbf{H}^{(A)} . \mathbf{meas}^{(A)} . \mathbf{H}^{(B)} . \mathbf{meas}^{(B)} . \quad (45)$$

In this representation, the first operation is **add**^(A), the second is **H**^(A) (the Hadamard gate on qubit A), and so on.

The algebraic specification of the algorithm as products of gates does not make it easy to see why the algorithm works. It is also difficult to see which operations depend on each other. Such dependencies are used to determine whether the operations can be parallelized. Quantum networks make these tasks simpler. The quantum network for the above sequence is shown in Figure 2.

To understand how the quantum network illustrated in Figure 2 solves the parity problem, we can follow the states as the network is executed from left to right, using the indicated checkpoints. Using vector notation for the states, at checkpoint 1, the state is

$$|\psi\rangle_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad (46)$$

where we used Kronecker product notation to denote the states of A, B, and C in this order. In the next time step, the network involves applying Hadamard gates—see Equation (13)—to A and B and a **not** gate—see Equation (9)—to C. At checkpoint 2, this operation results in the state

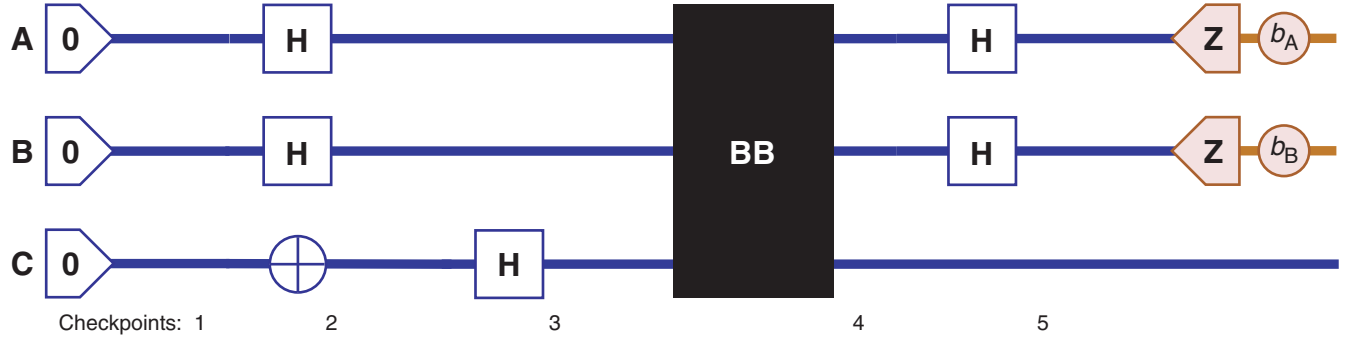


Figure 2. Quantum Network for Solving the Parity Problem

A quantum network has a line (horizontal in this case) for each qubit. The line can be thought of as the timeline for the qubit and is shown in blue. Each gate is drawn as a box, circle, or other element intercepting the lines of the qubits it acts on. In this case, time runs from left to right. Each qubit's timeline starts at the point where it is added. In this example, the qubits' timelines end when they are measured, at which point a classical bit (brown timeline) containing the measurement outcome is introduced. The operation BB is illustrated as a black box. The numbers underneath the network refer to checkpoints used to explain how the network solves the parity problem.

$$|\psi\rangle_2 = \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \otimes \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (47)$$

Next, a Hadamard gate is applied to C, so that at checkpoint 3, we have

$$|\psi\rangle_3 = \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \otimes \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix} \otimes \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix}. \quad (48)$$

The next event involves applying the black box. To understand what happens, note that the effect of the black box can be described as, “apply **not** to C if the parity according to b_A and b_B of the logical state of AB is 1.” The current state of C is such that, if **not** is applied, only the sign changes:

$$\begin{aligned} \text{not} \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix} &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix} \\ &= -\begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix}. \end{aligned} \quad (49)$$

Now, AB is in a superposition of each of the logical states, and conditional on the logical state and the (hidden) parity, the sign changes. As a result, although the state of C does not change, a phase is “kicked back” to AB. A generalization of this effect is at the heart of Alexei Kitaev’s version of Peter Shor’s quantum factoring algorithm (see the article “From Factoring to Phase Estimation” on page 38). At the next checkpoint, and after some arithmetic to check which logical states change sign, we can write the state as

$$|\psi\rangle_4 = \begin{pmatrix} 1/\sqrt{2} \\ (-1)^{b_A}/\sqrt{2} \end{pmatrix} \otimes \begin{pmatrix} 1/\sqrt{2} \\ (-1)^{b_B}/\sqrt{2} \end{pmatrix} \otimes \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix}. \quad (50)$$

Notice that qubits A and B are in orthogonal states for different values of b_A and b_B . It suffices to apply the Hadamard transform again to A and B to get

$$|\psi\rangle_4 = \begin{pmatrix} 1-b_A \\ b_A \end{pmatrix} \otimes \begin{pmatrix} 1-b_B \\ b_B \end{pmatrix} \otimes \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix}. \quad (51)$$

Measurements of A and B now reveal the previously unknown b_A and b_B .

As can be seen, the visual representation of a quantum network eases the tasks of following what happens. This is why it is used extensively for presenting basic subroutines and algorithms in quantum computation. A guide to the commonly used network elements is given in Table I.

When designing or describing complicated algorithms for quantum computers, providing everything in terms of quantum networks can become difficult, particularly when an important part of the algorithm consists of computations that are best done on a classical computer. For example, a full description of Shor's algorithm for factoring integers (see the article "From Factoring to Phase Estimation" on page 38) includes a significant amount of classical preprocessing, which determines choices made in the quantum algorithm, and classical postprocessing, which computes a factor from the measured result by a continued fraction algorithm. For such algorithms, one can use a programming language similar to Pascal, BASIC, or C enhanced with statements to access quantum bits and to apply quantum operations. For algorithm design, computer scientists often use a semiformal language called pseudocode (Cormen et al. 1990). With a simple extension called quantum pseudocode, the algorithm for the parity problem can be written as follows:

BBPARITY(BB)

Input: Access to a quantum black box **BB** that acts on three qubits by adding a parity function of the first two qubits to the third

Output: The two bits b_A and b_B of the parity function

foreach $i \in \{A, B, C\}$

$\lceil a_i \rceil \leftarrow |0\rangle$

C: Initialize three one-qubit registers $\lceil a_i \rceil$, $i = A, B, C$.

The corner bracket annotation declares a_i as a quantum register.

end

$\lceil a_C \rceil \leftarrow \sigma_x \lceil a_C \rceil$

foreach $i \in \{A, B, C\}$

$\lceil a_i \rceil \leftarrow \mathbf{H} \lceil a_i \rceil$

end

$\lceil a \rceil \leftarrow \mathbf{BB} \lceil a \rceil$

C: $\lceil a \rceil$ refers to the three-qubit register consisting of the $\lceil a_i \rceil$.

foreach $i \in \{A, B\}$

$\lceil a_i \rceil \leftarrow \mathbf{H} \lceil a_i \rceil$

```

 $b_i \leftarrow \text{meas } \lceil a_i \rceil$ 

end

return  $b_A, b_B$ 

end

```

Any classical programming language can be extended with statements to access and manipulate quantum registers.

Now, that we have looked at the quantum solution to the parity problem, let us consider the question of the least number of black-box applications required by a classical algorithm: Each classical use of the black box can only give us one bit of information. In particular, one use of the black box with input $a_A a_B$ reveals only the parity of $a_A a_B$ according to the hidden parameters b_A and b_B . Each use of the black box can therefore only help us distinguish between two subsets of the four possible parities. At least two uses of the black box are therefore necessary. Two uses are also sufficient. To determine which of the four parities is involved, use the black box first with input $a_A a_B = 10$ and then with input $a_A a_B = 01$. As a result of this argument, one can consider the parity problem as a simple example of a case in which there is a more efficient quantum algorithm than is possible

Table I. Quantum Network Elements










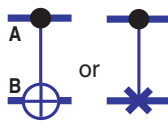

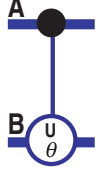

Gate Names and Their Abbreviations	Gate Symbols	Algebraic Form	Matrix Form
Add/Prepare, add		If applied to an existing qubit $\{ 0\rangle\langle 0 , 0\rangle\langle 1 \}$ (operator mixture)	$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ or $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$
Measure, meas		$\{0: 0\rangle\langle 0 , 1: 1\rangle\langle 1 \}$	$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ or $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$
Not, not , σ_x	 or 	$ 0\rangle\langle 1 + 1\rangle\langle 0 $	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$
Hadamard, H		$e^{-i\sigma_y\pi/4}\sigma_z$	$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$
Phase Change, S ($e^{i\phi}$)		$e^{i\phi/2}e^{-i\sigma_z\phi/2}$	$\begin{pmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{pmatrix}$

Table I. (Continued)

Gate Names and Their Abbreviations	Gate Symbols	Algebraic Form	Matrix Form
z-Rotation, \mathbf{Z}_ϕ		$e^{-i\sigma_z\phi/2}$	$\begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix}$
y-Rotation, \mathbf{Y}_θ		$e^{-i\sigma_y\theta/2}$	$\begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$
x-Rotation, \mathbf{X}_θ		$e^{-i\sigma_x\theta/2}$	$\begin{pmatrix} \cos(\theta/2) & -i\sin(\theta/2) \\ -i\sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$
Controlled not, cnot		$ 0\rangle_A \langle 0 + 1\rangle_A \langle 1 \sigma_x^{(B)}$ $e^{-i\sigma_z^{(A)}\pi/4} e^{-i/2(1 - \sigma_z^{(A)}) \sigma_x^{(B)}\pi/2}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$
zz-Rotation, $(\mathbf{ZZ})_\theta$		$e^{-i\sigma_z^{(A)} \sigma_z^{(B)} \theta/2}$	$\begin{pmatrix} e^{-i\theta/2} & 0 & 0 & 0 \\ 0 & e^{i\theta/2} & 0 & 0 \\ 0 & 0 & e^{i\theta/2} & 0 \\ 0 & 0 & 0 & e^{-i\theta/2} \end{pmatrix}$
Controlled Rotation, cU $_\theta$		$ 0\rangle_A \langle 0 + 1\rangle_A \langle 1 e^{-i\sigma_U^{(B)} \theta/2}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & & e^{-i\sigma_U\theta/2} \\ 0 & 0 & & \end{pmatrix}$
Toffoli Gate, c²not		$\mathbb{1} - 11\rangle_{AB} \langle 11 + 11\rangle_{AB} \langle 11 \sigma_x^{(C)}$	

classically. However, it is worth noting that the comparison is not entirely fair: A truly classical oracle answering parity questions or implementing the black box on the states of classical bits is useless to a quantum algorithm. To take advantage of such an algorithm, it must be possible to use superpositions that are not implicitly collapsed. Collapse can happen if the oracle makes a measurement or otherwise “remembers” the question that it was asked.

Resource Accounting

When trying to solve a problem using quantum information processing, an important issue is to determine what physical resources are available and how much of each resource is needed for the solution. As mentioned before, in classical information, the primary resources are bits and operations. The number of bits used by an algorithm is its space requirement; the number of operations used, its time requirement. If parallel computation is available, one can distinguish between the total number of operations (work) and the number of parallel steps (time).

When quantum information processing is used, the classical resources are still relevant for running the computer that controls the quantum system and performs any preprocessing and postprocessing tasks. The main quantum resources are analogous to the classical ones: Quantum space is the number of qubits needed, and quantum time, the number of quantum gates. Because it turns out that reset operations have a thermodynamic cost, one can count irreversible quantum operations separately. This accounting of the resource requirements of algorithms and of the minimum resources needed to solve problems forms the foundation of quantum complexity theory.

As a simple example of resource accounting, consider the algorithm for the parity problem. No classical computation is required to decide which quantum gates to apply or to determine the answer from the measurement. The quantum network consists of a total of 11 quantum gates (including **add** and **meas** operations) and one oracle call (the application of the black box). In the case of oracle problems, one usually counts the number of oracle calls first, as we have done in discussing the algorithm. The network is readily parallelized to reduce the time resource to 6 steps.

Part III: Advantages of Quantum Information

The notion of quantum information as explained in this primer was established in the 1990s. It emerged from research focused on understanding how physics affects our capabilities to communicate and process information. The recognition that usable types of information need to be physically realizable was repeatedly emphasized by Rolf Landauer, who proclaimed that “information is physical” (1991). Beginning in the 1960s, Landauer studied the thermodynamic cost of irreversible operations in computation (1961). Charles Bennett showed that, by using reversible computation, this cost can be avoided (1973). Limitations of measurement in quantum mechanics were investigated early by researchers such as John von Neumann (1932a and 1932b) and later by Alexander Holevo (1973b) and Carl Helstrom (1976). Holevo introduced the idea of quantum communication channels and found bounds on their capacity for transmitting classical information (1973a). Initially, most work focused on determining the physical limitations placed on classical information processing. The fact that pairs of two-level systems can have correlations not possible for classical systems was proved by John Bell (1964). Subsequently, indications that quantum mechanics offers advantages to information processing came from Stephen Wiesner’s studies of cryptographic applications in the late 1960s. Wiesner’s work was not recognized, however, until the 1980s,

when Bennett, Gilles Brassard, Seth Breidbart, and Wiesner introduced (1982) the idea of quantum cryptography, which can be used to communicate in secret.

Initially, the term quantum computation was mostly used to refer to classical computers realized with quantum mechanical systems. In the 1980s, Paul Benioff (1980), Richard Feynman (1982), and Yuri Manin (1980) introduced the idea of a quantum computer based on quantum information. They noted that the apparent exponential complexity of simulating quantum mechanics on a classical computer might be overcome if one could use a computer based on quantum mechanics. A formal model of quantum Turing machines was soon defined by David Deutsch (1985), who later also introduced quantum networks (1989). Deutsch and Richard Jozsa (1992) were the first to introduce a black-box problem that could be solved deterministically on a quantum computer in fewer steps than on a classical computer.

In spite of suggestions that it could lead to large efficiency improvements in simulating physics, quantum information processing was still largely an academic subject. Based on work by Ethan Bernstein and Umesh Vazirani (1993) that formalized quantum complexity theory, Dan Simon (1994) showed that, for black-box problems, quantum computers can be exponentially more efficient than classical deterministic or probabilistic computers, giving the first indication of a strong advantage for quantum information processing. It was Shor's algorithm for factoring large integers (1994 and 1997) that finally convinced a larger community that quantum information was more than just a tool for realizing classical computers. This change in attitude was in no small part due to the fact that the security of commonly used cryptographic protocols is based on the difficulty of factoring.

At that point, it was still generally believed that the fragility of quantum states made it unlikely for reasonably large quantum computers to be realized in practice. But the discovery by Shor (1995) and Andrew Steane (1996) that quantum error correction was possible soon changed that view (for an introductory overview, see the article on quantum error correction on page 188).

Because the usefulness and realizability of quantum information has been recognized, the science of quantum information processing is a rapidly growing field. As quantum information becomes increasingly accessible by technology, its usefulness will be more apparent. The next few sections discuss what we currently know about applications of quantum information processing. Refer to Michael Nielsen and Isaac Chuang (2001) as a useful reference text on quantum computation and information with historical notes.

Quantum Algorithms

Shor's factoring algorithm, which precipitated much of the current work in quantum information processing, is based on a quantum realization of the fast Fourier transform. The most powerful version of this technique is now represented by the phase estimation algorithm of Kitaev (1995) as formalized by Richard Cleve et al. (1998). (For an explanation, see the article "From Factoring to Phase Estimation" on page 38.) The best-known application of quantum factoring is cryptanalysis, where it allows efficiently breaking the currently used public-key cryptographic codes. Whether there are any constructive applications of quantum factoring and its generalizations remains to be determined. For users of public-key cryptography, a crucial question is, "How long can public-key codes based on factoring continue to be used safely?" To attempt an answer to this question, one can note that to break a code with a typical key size of 1000 bits requires more than 3000 qubits and 10^8 quantum gates, which is well out of reach of current technology. However, it is conceivable that a recording of encrypted information transmitted in 2000 can be broken in the next "few" decades.

Shor's quantum factoring algorithm was not the first with a significant advantage over classical algorithms. The first proposed quantum algorithms with this property were for simulating quantum mechanical systems. These algorithms simulate the evolution of a reasonably large number of interacting quantum particles—for example, the electrons and nuclei in a molecule. The algorithms' outputs are what would be measurable physical quantities of the system being simulated. The known methods for obtaining these quantities on classical computers scale exponentially with the number of particles, except in special cases.

The idea of using quantum computers for simulating quantum physics spurred the work that eventually led to the quantum factoring algorithm. However, that idea did not have the broad scientific impact that the quantum factoring algorithm had. One reason is that, because of its cryptographic applications, factoring is a heavily studied problem in theoretical computer science and cryptography. Because so many people have tried to design efficient algorithms for factoring and failed, the general belief that factoring is hard for classical computers has a lot of credibility. In contrast, a quantum physics simulation has no simple formulation as an algorithmic problem suitable for study in theoretical computer science. Furthermore, many researchers still believe that the physically relevant questions can be answered with efficient classical algorithms, requiring only more cleverness on the part of algorithm designers. Another reason for the lack of impact is that many of the fundamental physical quantities of interest are not known to be efficiently accessible even on quantum computers. For example, one of the first questions about a physical system with a given Hamiltonian (energy observable) is, "What is the ground-state energy?" It is known that the ability to efficiently answer this question for physically reasonable Hamiltonians leads to efficient algorithms for hard problems, such as the traveling salesman or the scheduling problems. In spite of occasional claims to the contrary, an efficient quantum solution to these problems is widely considered unlikely.

Most quantum algorithms for physics simulations are based on a direct emulation of a quantum mechanical system's evolution. The focus of the original proposals by Feynman and others was on how to implement the emulation using a suitable formulation of general-purpose quantum computers. After such computers were formalized by Deutsch, the implementation of the emulation was generalized and refined by Seth Lloyd (1996), Wiesner (1996), and Christof Zalka (1998). The ability to emulate the evolution of quantum systems is actually widely used by classical Monte Carlo algorithms for simulating physics. In those algorithms, state amplitudes are, in effect, represented by expectations of random variables that are computed during the simulation. As in the case of quantum algorithms for physics emulation, Monte Carlo algorithms efficiently evolve the representation of the quantum system. The inefficiency of the classical algorithm arises only in determining a physical quantity of interest. In the case of Monte Carlo algorithms, the measurement of a physical quantity suffers from the so-called sign problem, often resulting in exponentially large, random errors that can be reduced only by repeating the computation exponentially many times. In contrast, the quantum algorithms for emulation can determine many (but not all) of the interesting physical quantities with polynomially bounded statistical errors. How to efficiently implement measurements of these quantities has been the topic of more recent work in this area, much of which is based on variants of the phase-estimation algorithm (Terhal and DiVincenzo 2000, Knill and Laflamme 1998, Abrams and Lloyd 1999, Ortiz et al. 2001, Miquel et al. 2002).

Although several researchers have suggested that there are interesting quantum physics simulations that can be implemented with well below 100 qubits, one of the interesting problems in this area of research is to come up with a specific simulation algorithm using small numbers of qubits and quantum gates, an algorithm that computes an interesting physical quantity not easily obtainable using available classical computers.

Another notable algorithm for quantum computers, unstructured quantum search, was described by Lov Grover (1996). Given is a black box that computes a binary function f on inputs x with $0 \leq x < N$. The function f has the property that there is a unique input a for which $f(a) = 1$. The standard quantum version of this black box implements the transformation $\hat{f}|x\rangle|b\rangle = |x\rangle|b \oplus f(x)\rangle$, where b is a bit and $b \oplus f(x)$ is computed modulo 2. Unstructured quantum search finds a quadratically faster, that is, in time of order $N^{1/2}$, than the best classical black-box search, which requires time of order N . The context for this algorithm is the famous $P \neq NP$ conjecture, which is captured by the following algorithmic problem: Given is a classical circuit C that computes an output. Is there an input to the circuit for which the circuit's output is 1? Such an input is called a satisfying input or assignment. For any given input, it is easy to check the output, but an efficient algorithm that finds a satisfying input is conjectured to be impossible. This is the $P \neq NP$ conjecture. Generalizations of Grover's search algorithm—amplitude amplification (Brassard et al. 1998)—allow finding satisfying inputs faster than naive, classical search does, which tries every possible input in some, possibly random, order. It is worth noting, however, that if sufficient classical parallelism is available, quantum search loses many of its advantages.

The three algorithms just described capture essentially all the known algorithmic advantages of quantum computers. Almost all algorithms that have been described are applications of phase estimation or of amplitude amplification. These algorithms well justify developing special-purpose quantum information-processing technology. Will general-purpose quantum computers be useful? More specifically, what other algorithmic advantages do quantum computers have?

Quantum Communication

Quantum communication is an area in which quantum information has proven (rather than conjectured) advantages. The best-known application is quantum cryptography, whereby two parties, Alice and Bob, can generate a secret key using a quantum communication channel (for example, photons transmitted in optical fiber) and an authenticated classical channel (for example, a telephone line). Any attempt at learning the key by eavesdropping is detected. A quantum protocol for generating a secret key is called a quantum-key-exchange protocol. There are no equally secure means for generating a secret key by using only classical deterministic channels. Few quantum operations are needed to implement quantum key exchange, and as a result, there are working prototype systems (Hughes et al. 2000, Townsend 1998, Ribordy et al. 2001). To overcome the distance limitations (tens of kilometers) of current technology requires the use of quantum error correction and hence more demanding quantum technology.

Quantum key exchange is one of an increasing number of multiparty problems that can be solved more efficiently with quantum information. The area of research concerned with how several parties at different locations can solve problems while minimizing communication resources is called communication complexity. For quantum communication complexity (Cleve and Burhman 1997), the communication resources include either shared entangled qubits or a means for transmitting quantum bits. A seminal paper by Howard Burhman, Cleve, and Wim van Dam (2000) shows how the nonclassical correlations present in maximally entangled states lead to protocols based on such states that are more efficient than any classical deterministic or probabilistic protocol achieving the same goal. Ran Raz (1999) showed that there is an exponential improvement in communication resources for a problem in which Alice and Bob have to answer a question about the relationship between a vector known to Alice and a matrix known to Bob. Although this problem is artificial, it suggests that there are potentially useful advantages to be gained from quantum information in this setting.

Quantum Control

According to Moore's law of semiconductor technology, the size of transistors is decreasing exponentially, by a factor of about .8 every year. If this trend continues, then over the next few decades, devices will inevitably be built whose behavior will be primarily quantum mechanical. For the purpose of classical computation, the goal is to remove the quantum behavior and stabilize classical information. But quantum information offers an alternative: It is possible to directly use quantum effects to advantage. Whether or not this alternative is useful (and we believe it is), the ideas of quantum information can be used to systematically understand and control quantum mechanical systems.

The decreasing size of semiconductor components is a strong motivation to strive for better understanding the behavior of condensed-matter quantum mechanical systems. But there is no reason to wait for Moore's law: There are a rapidly increasing number of experimental systems in which quantum mechanical effects are being used and investigated. Examples include many optical devices (lasers, microwave cavities, entangled photon pairs), nuclear magnetic resonance with molecules or in solid state, trapped ion or atom systems, Rydberg atoms, superconducting devices (Josephson junctions and SQUIDs), and spintronics (electron spins in semiconductor devices). Many of these systems are being considered as candidates for realizing quantum information processing. Yet, regardless of the future of quantum information processing, there is ample motivation for studying these systems.

Outlook

The science of quantum information processing is promising a significant impact on how we process information, solve algorithmic problems, engineer nanoscale devices, and model fundamental physics. It is already changing the way we understand and control matter at the atomic scale, making the quantum world more familiar, accessible, and understandable. Whether or not we do most of our everyday computations by using the classical model, it is likely that the physical devices that support these computations will exploit quantum mechanics and integrate the ideas and tools that have been developed for quantum information processing. ■

Acknowledgment

We thank Nikki Cooper and Ileana Buican for their extensive encouragement and editorial help.

Further Reading

- Abrams, D. S., and S. Lloyd. 1999. Quantum Algorithm Providing an Exponential Speed Increase for Finding Eigenvalues and Eigenvectors. *Phys. Rev. Lett.* **83**: 5162.
- Barenco, A., C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, et al. 1995. Elementary Gates for Quantum Computation. *Phys. Rev. A* **52**: 3457.
- Bell, J. S. 1964. On the Einstein-Podolsky-Rosen Paradox. *Phys.* **1**: 195.
- Benioff, P. 1980. The Computer as a Physical System: A Microscopic Quantum Mechanical Hamiltonian Model of Computers as Represented by Turing Machines. *J. Stat. Phys.* **22**: 563.
- Bennett, C. H. 1973. Logical Reversibility of Computation. *IBM J. Res. Dev.* **17**: 525.
- Bennett, C. H., G. Brassard, S. Breidbart, and S. Wiesner. 1982. Quantum Cryptography, or Unforgeable Subway Tokens. In *Advances in Cryptology: Proceedings of Crypto '82*. Edited by D. Chaun, R. L. Rivest, and A. T. Sherman, p. 267. New York: Plenum Press.



Emanuel (Manny) Knill received his Ph. D. in pure mathematics from the University of Colorado at Boulder in 1991. Since 1992, he has been with Los Alamos National Laboratory. Manny has worked on various aspects of quantum information processing since 1995. (Photo of Augustus)

Contact Information

E. Knill: knill@lanl.gov

R. Laflamme: laflamme@iqc.ca

H. Barnum: barnum@lanl.gov

D. Dalvit: dalvit@lanl.gov

J. Dziarmaga: ufjacekd@th.if.uj.edu.pl

J. Gubernatis: jg@lanl.gov

L. Gurvits: gurvits@lanl.gov

G. Ortiz: g_ortiz@lanl.gov

L. Viola: lviola@lanl.gov

W. Zurek: whz@lanl.gov

- Bernstein, E., and U. Vazirani. 1993. Quantum Complexity Theory. In *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing 1993*, p. 11. New York: ACM Press.
- Bolker, E. D. 1970. *Elementary Number Theory: An Algebraic Approach*. New York: W. A. Benjamin, Inc.
- Brassard, G., P. Hoyer, and A. Tapp. 1998. Quantum Counting. In *Automata, Languages and Programming, Proceedings of ICALP '98*, Vol. 1443 of *Lecture Notes in Computer Science*. Edited by K. G. Larsen, S. Skyum, and G. Winskel, p. 820. Berlin: Springer Verlag.
- Buhrman, H., R. Cleve, and W. Van Dam. 2000. Quantum Entanglement and Communication Complexity. *SIAM J. Comput.* **30**: 1829.
- Cleve, R., and H. Buhrman. 1997. Substituting Quantum Entanglement for Communication. *Phys. Rev. A* **56**: 1201.
- Cleve, R., A. Ekert, C. Macchiavello, and M. Mosca. 1998. Quantum Algorithms Revisited. *Proc. R. Soc. London, Ser. A* **454**: 339.
- Cormen, T. H., C. B. Leiserson, and R. L. Rivest. 1990. *Introduction to Algorithms*. Cambridge, MA: MIT Press.
- Deutsch, D. 1985. Quantum Theory, The Church-Turing Principle and the Universal Quantum Computer. *Proc. R. Soc. London, Ser. A* **400**: 97.
- . 1989. Quantum Computational Networks. *Proc. R. Soc. London, Ser. A* **425**: 73.
- Deutsch, D., and R. Jozsa. 1992. Rapid Solution of Problems by Quantum Computation. *Proc. R. Soc. London, Ser. A* **439**: 553.
- Ekert, A. 1998. From Quantum Code-Making to Quantum Code-Breaking. In *The Geometric Universe*, p. 195. Oxford: Oxford University Press.
- Feynman, R. P. 1982. Simulating Physics with Computers. *Int. J. Theor. Phys.* **21**: 467.
- Griffiths, R. B., and C.-S. Niu. 1996. Semiclassical Fourier Transform for Quantum Computation. *Phys. Rev. Lett.* **76**: 3228.
- Grover, L. K. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the 28th Annual ACM Symposium on the Theory of Computation*, p. 212, New York: ACM Press.
- Gupta, R., S. A. Smolka, and S. Bhaskar. 1994. On Randomization in Sequential and Distributed Algorithms. *ACM Comput. Surveys* **26**: 7.
- Hardy, G. H., and E. M. Wright. 1979. *An Introduction to the Theory of Numbers*. Fifth edition. London: Oxford University Press.
- Helstrom, C. W. 1976. *Quantum Detection and Estimation Theory. Mathematics in Science and Engineering* Vol. 123. New York: Academic Press.
- Holevo, A. S. 1973a. Bounds for the Quantity of Information Transmitted by a Quantum Communication Channel. *Probl. Inf. Transm.* **9**: 177.
- . 1973b. Statistical Problems in Quantum Physics. In *Proceedings of the Second Japan-USSR Symposium on Probability Theory, Lecture Notes in Mathematics* Vol. 330. Edited by G. Maruyama and Y. V. Prokhorov, p. 104. Berlin: Springer Verlag.
- Hughes, R. J., G. L. Morgan, and C. G. Peterson. 2000. Quantum Key Distribution Over a 48-km Optical Fibre Network. *J. Modern Opt.* **47**: 533.
- Kitaev, A. Yu. 1995. Quantum Measurements and the Abelian Stabilizer Problem. [Online]: [http://eprints.lanl.gov. \(quant-ph/9511026\)](http://eprints.lanl.gov. (quant-ph/9511026)).
- Knill, E., and R. Laflamme. 1998. On the Power of One Bit of Quantum Information. *Phys. Rev. Lett.* **81**: 5672.
- Landauer, R. 1961. Irreversibility and Heat Generation in the Computing Process. *IBM J. Res. Dev.* **5**: 183.
- . 1991. Information is Physical. *Phys. Today* **44**: 22.
- Lloyd, S. 1996. Universal Quantum Simulators. *Science* **273**: 1073.
- Manin, Y. I. 1980. *The Computable and the Not Computable*. Moscow: Sovetskoye Radio. (In Russian).
- Miquel, C., J. P. Paz, M. Saraceno, E. Knill, R. Laflamme, and C. Negrevergne. 2002. Interpretation of Tomography and Spectroscopy as Dual Forms of Quantum Computations. *Nature* **418**: 59.
- Nielsen, M. A., and I. L. Chuang. 2001. *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.

- Ortiz, O., J. E. Gubernatis, E. Knill, and R. Laflamme. 2001. Quantum Algorithms for Fermionic Simulations. *Phys. Rev. A* **64**: 022319.
- Papadimitriou, C. H. 1994. *Computational Complexity*. Reading, MA: Addison-Wesley.
- Raz, R. 1999. Exponential Separation of Quantum and Classical Communication Complexity. In *Proceedings of the 31st Annual ACM Symposium on the Theory of Computation (STOC)*, p. 358. El Paso, TX: ACM Press.
- Ribordy, O., J. Brendel, J.-D. Gautier, N. Gisin, and H. Zbinden. 2001. Long-Distance Entanglement-Based Quantum Key Distribution. *Phys. Rev. A* **63**: 012309.
- Shor, P. W. 1994. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, p. 124. Los Alamitos, CA: IEEE Press.
- . 1995. Scheme for Reducing Decoherence in Quantum Computer Memory. *Phys. Rev. A* **52**: 2493.
- . 1997. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* **26**: 1484.
- Simon, D. R. 1994. On the Power of Quantum Computation. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, p. 116. Los Alamitos, CA: IEEE Press.
- Steane, A. 1996. Multiple Particle Interference and Quantum Error Correction. *Proc. R. Soc. London, Ser. A* **452**: 2551.
- Terhal, B. M., and D. P. DiVincenzo. 2000. Problem of Equilibration and the Computation of Correlation Functions on a Quantum Computer. *Phys. Rev. A* **61**: 022301.
- Townsend, P. D. 1998. Quantum Cryptography on Optical Fiber Networks. *Opt. Fiber Tech.: Mat., Dev., Sys.* **4**: 345.
- von Neumann, J. 1932a. Der Messprozess. Ch. VI. In *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer Verlag.
- . 1932b. "Messung und Reversibilität." Allgemeine Betrachtungen. Ch. V. In *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer Verlag.
- Wiesner, S. 1983. Conjugate Coding. *Sigact News* **15**: 78.
- . 1996. Simulations of Many-Body Quantum Systems by a Quantum Computer. [Online]: [http://eprints.lanl.gov. \(quant-ph/9603028\)](http://eprints.lanl.gov. (quant-ph/9603028)).
- Yao, A. 1993. Quantum Circuit Complexity. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, p. 352. Los Alamitos, CA: IEEE Press.
- Zalka, C. 1998. Simulating Quantum Systems on a Quantum Computer. *Proc. R. Soc. London, Ser. A* **454**: 313.

Glossary

- Algorithm.** A set of instructions to be executed by a computing device. What instructions are available depends on the computing device. Typically, instructions include commands for manipulating the contents of memory and means for repeating blocks of instructions indefinitely or until a desired condition is met.
- Amplitude.** A quantum system with a chosen orthonormal basis of "logical" states $|i\rangle$ can be in any superposition $\sum_i \alpha_i |i\rangle$ of these states, where $\sum_i |\alpha_i|^2 = 1$. In such a superposition, the complex numbers α_i are called the amplitudes. Note that the amplitudes depend on the chosen basis.
- Ancillas.** Helper systems used to assist in a computation involving other information systems.
- Bell basis.** For two qubits A and B, the Bell basis consists of the four states $1/\sqrt{2}(|00\rangle_{AB} \pm |11\rangle_{AB})$ and $1/\sqrt{2}(|01\rangle_{AB} \pm |10\rangle_{AB})$.
- Bell states.** The members of the Bell basis.
- Bit.** The basic unit of deterministic information. It is a system that can be in one of two possible states, 0 and 1.

Bit sequence. A way of combining bits into a larger system whose constituent bits are in a specific order.

Bit string. A sequence of 0s and 1s that represents a state of a bit sequence. Bit strings are the words of a binary alphabet.

Black box. A computational operation whose implementation is unknown. Typically, a black box implements one of a restricted set of operations, and the goal is to determine which of these operations it implements by using it with different inputs. Each use of the black box is called a “query.” The smallest number of queries required to determine the operation is called the “query complexity” of the restricted set. Determining the query complexity of sets of operations is an important area of computational complexity.

Bloch sphere. The set of pure states of a qubit represented as points on the surface of the unit sphere in three dimensions.

Bra. A state expression of the form $\langle\psi|$ considered to be the conjugate transpose of the ket expression $|\psi\rangle$.

Bra-ket notation. A way of denoting states and operators of quantum systems with kets (for example, $|\psi\rangle$) and bras (for example, $\langle\phi|$).

Circuit. A combination of gates applied to information units in a prescribed order. To draw circuits, one often uses a convention for connecting and depicting gates. See also “network.”

Circuit complexity. The circuit complexity of an operation on a fixed number of information units is the smallest number of gates required to implement the operation.

Classical information. The type of information based on bits and bit strings and more generally on words formed from finite alphabets. This is the information used for communication between people. Classical information can refer to deterministic or probabilistic information, depending on the context.

Computation. The execution of the instructions provided by an algorithm.

Computational states. See “logical states.”

Computer. A device that processes information.

Density matrix or operator. A representation of pure and mixed states without redundancy. For a pure state $|\psi\rangle$, the corresponding density operator is $|\psi\rangle\langle\psi|$.

A general density operator is a probabilistic combination $\sum_i \lambda_i |\psi_i\rangle\langle\psi_i|$, with $\sum_i \lambda_i = 1$.

Deterministic information. The type of information that is based on bits and bit strings. Deterministic information is classical, but it explicitly excludes probabilistic information.

Distinguishable states. In quantum mechanics, two states are considered distinguishable if they are orthogonal. In this case, a measurement exists that is guaranteed to determine which of the two states a system is in.

Efficient computation. A computation is efficient if it requires, at most, polynomially many resources as a function of input size. For example, if the computation returns the value $f(x)$ on input x , where x is a bit string, then it is efficient if there exists a power k such that the number of computational steps used to obtain $f(x)$ is bounded by $|x|^k$, where $|x|$ is the length (number of bits) of x .

Entanglement. A nonclassical correlation between two quantum systems most strongly exhibited by the maximally entangled states, such as the Bell states for two qubits, and considered to be absent in mixtures of product states (which are called separable states). Often, states that are not separable are considered to be entangled. However, nearly separable states do not exhibit all the features of maximally entangled states. As a result, studies of different types of entanglement are an important component of quantum information theory.

Gate. An operation applied to information for the purpose of information processing.

Global phase. Two quantum states are indistinguishable if they differ only by a global phase. That is, $|\psi\rangle$ and $e^{i\phi}|\psi\rangle$ are in essence the same state. The global phase difference is the factor $e^{i\phi}$. The equivalence of the two states is apparent from the fact that their density matrices are the same.

Hilbert space. An n -dimensional Hilbert space consists of all complex n -dimensional vectors. A defining operation in a Hilbert space is the inner product. If the vectors are thought of as column vectors, then the inner product $\langle x, y \rangle$ of x and y is obtained by forming the conjugate transpose x^\dagger of x and calculating $\langle x, y \rangle = x^\dagger y$. The inner product induces the usual squared norm $|x|^2 = \langle x, x \rangle$.

Information. Something that can be recorded, communicated, and computed with. Information is fungible; that is, its meaning can be identified regardless of the particulars of the physical realization. Thus, information in one realization (such as ink on a sheet of paper) can be easily transferred to another (for example, spoken words). Types of information include deterministic, probabilistic, and quantum information. Each type is characterized by information units, which are abstract systems whose states represent the simplest information of each type. The information units define the “natural” representation of the information. For deterministic information, the information unit is the bit, whose states are symbolized by 0 and 1. Information units can be put together to form larger systems and can be processed with basic operations acting on few of them at a time.

Inner product. The defining operation of a Hilbert space. In a finite dimensional Hilbert space with a chosen orthonormal basis $\{e_i : 1 \leq i \leq n\}$, the inner product of two vectors $x = \sum_i x_i e_i$ and $y = \sum_i y_i e_i$ is given by $\sum_i \bar{x}_i y_i$. In the standard column representation of the two vectors, this is the number obtained by computing the product of the conjugate transpose of x with y . For real vectors, that product agrees with the usual “dot” product. The inner product of x and y is often written in the form $\langle x, y \rangle$. Pure quantum states are unit vectors in a Hilbert space. If $|\phi\rangle$ and $|\psi\rangle$ are two quantum states expressed in the ket-bra notation, their inner product is given by $(|\phi\rangle)^\dagger |\psi\rangle = \langle \phi | \psi \rangle$.

Ket. A state expression of the form $|\psi\rangle$ representing a quantum state. Usually, $|\psi\rangle$ is thought of as a superposition of members of a logical state basis $|i\rangle$. One way to think about the notation is to consider the two symbols $|$ and \rangle as delimiters denoting a quantum system and ψ as a symbol representing a state in a standard Hilbert space. The combination $|\psi\rangle$ is the state of the quantum system associated with ψ in the standard Hilbert space via a fixed isomorphism. In other words, one can think of $\psi \leftrightarrow |\psi\rangle$ as an identification of the quantum system’s state space with the standard Hilbert space.

Linear extension of an operator. The unique linear operator that implements a map defined on a basis. Typically, we define an operator U on a quantum system only on the logical states $U : |i\rangle \rightarrow |\psi_i\rangle$. The linear extension is defined by $U(\sum_i \alpha_i |i\rangle) = \sum_i \alpha_i |\psi_i\rangle$.

Logical states. For quantum systems used in information processing, the logical states are a fixed orthonormal basis of pure states. By convention, the logical basis for qubits consists of $|0\rangle$ and $|1\rangle$. For larger dimensional quantum systems, the logical basis is often indexed by integers, $|0\rangle$, $|1\rangle$, $|2\rangle$, and so on. The logical basis is often called the computational basis, or sometimes, the classical basis.

Measurement. The process used to extract classical information from a quantum system. A general projective measurement is defined by a set of projectors P_i , satisfying $\sum_i P_i = \mathbb{1}$ and $P_i P_j = \delta_{ij} P_i$. Given the quantum state $|\psi\rangle$, the outcome of a measurement with the set $\{P_i\}_i$ is one of the classical indices i associated with a projector P_i . The index i is the measurement outcome. The probability of outcome i is $p_i = |P_i |\psi\rangle|^2$, and given outcome i , the quantum state “collapses” to $P_i |\psi\rangle / \sqrt{p_i}$.

Mixture. A probabilistic combination of the pure states of a quantum system. Mixtures can be represented without redundancy with density operators. Thus, a mixture is of the form $\sum_i \lambda_i |\psi_i\rangle\langle\psi_i|$, with $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$ being the probabilities of the states $|\psi_i\rangle$. This expression for mixtures defines the set of density operators, which can also be characterized as the set of operators ρ satisfying $\text{tr}(\rho) = 1$ and for all $|\psi\rangle$, $\langle\psi|\rho|\psi\rangle \geq 0$ (“positive semidefinite operator”).

Network. In the context of information processing, a network is a sequence of gates applied to specified information units. Networks can be visualized as displaying horizontal lines that denote the timeline of an information unit. The gates are represented by graphical elements that intercept the lines at specific points. A realization of the network requires applying the gates to the information units in the specified order (left to right).

Operator. A function that transforms the states of a system. Operators may be restricted depending on the system’s properties. For example, in talking about operators acting on quantum systems, one always assumes that they are linear.

Oracle. An information processing operation that can be applied. A use of the oracle is called a query. In the oracle model of computation, a standard model is extended to include the ability to query an oracle. Each oracle query is assumed to take one time unit. Queries can reduce the resources required for solving problems. Usually, the oracle implements a function or solves a problem not efficiently implementable by the model without the oracle. Oracle models are used to compare the power of two models of computation when the oracle can be defined for both models. In 1994, for example, Dan Simon showed that quantum computers with a specific oracle O could efficiently solve a problem that had no efficient solution on classical computers with access to the classical version of O . At the time, this result was considered the strongest evidence for an exponential gap in power between classical and quantum computers.

Overlap. The inner product between two quantum states.

Pauli operators. The Hermitian matrices σ_x , σ_y , and σ_z acting on qubits, which are two-level quantum systems. They are defined in Equation (12). It is often convenient to consider the identity operator to be included in the set of Pauli operators.

Polynomial resources. To say that an algorithm computing the function $f(x)$, where x is a bit string, uses polynomial resources (in other words, is efficient) means that the number of steps required to compute $f(x)$ is bounded by $|x|^k$ for some fixed k . Here, $|x|$ denotes the length of the bit string x .

Probabilistic bit. The basic unit of probabilistic information whose state space consists of all probability distributions over the two states of a bit. The states can be thought of as describing the outcome of a biased coin flip before the coin is flipped.

Probabilistic information. The type of information obtained by extending the state spaces of deterministic information to allow arbitrary probability distributions over the deterministic states. This is the main type of classical information with which quantum information is compared.

Probability amplitude. The squared norm of an amplitude with respect to a chosen orthonormal basis $\{|i\rangle\}$. Thus, the probability amplitude is the probability with which the state $|i\rangle$ is measured in a complete measurement that uses this basis.

Product state. For two quantum systems A and B, product states are of the form $|\psi\rangle_A |\phi\rangle_B$. Most states are not of this form.

Program. An algorithm expressed in a language that can be understood by a particular type of computer.

Projection operator. A linear operator P on a Hilbert space that satisfies $P^2 = P^\dagger P = P$. The projection onto a subspace V with orthogonal complement W is defined as follows: If $x \in V$ and $y \in W$, then $P(x + y) = x$.

Pseudocode. A semiformal computer language intended to be executed by a standard random-access machine, which is a machine model with a central processing unit and access to a numerically indexed unbounded memory. This machine model is representative of the typical one-processor computer. Pseudocode is similar to programming languages such as BASIC, Pascal, or C but does not have specialized instructions for human interfaces, file management, or other “external” devices. Its main use is to describe algorithms and enable machine-independent analysis of the algorithms’ resource usage.

Pure state. A state of a quantum system that corresponds to a unit vector in the Hilbert space used to represent the system’s state space. In the ket notation, pure states are written in the form $|\psi\rangle = \sum_i \alpha_i |i\rangle$, where the $|i\rangle$ form a logical basis and $\sum_i |\alpha_i|^2 = 1$.

Quantum information. The type of information obtained when the state space of deterministic information is extended by normalized superpositions of deterministic states. Formally, each deterministic state is identified with one of an orthonormal basis vector in a Hilbert space, and normalized superpositions are unit-length vectors expressible as complex linear sums of the chosen basis vectors. It is convenient to extend this state space further by permitting probability distributions over the quantum states (see the entry for “mixtures”). This extension is still called quantum information.

Qubit. The basic unit of quantum information. It is the quantum extension of the deterministic bit, which implies that its state space consists of the unit-length vectors in a two-dimensional Hilbert space.

Readout. A method for obtaining human-readable information from the state of a computer. For quantum computers, readout refers to a measurement process used to obtain classical information about a quantum system.

Reversible gate. A gate whose action can be undone by a sequence of gates.

Separable state. A mixture of product states.

States. The set of states for a system characterizes the system’s behavior and possible configurations.

Subspace. For a Hilbert space, a subspace is a linearly closed subset of the vector space.

The term can be used more generally for a system Q of any information type:

A subspace of Q or, more specifically, of the state space of Q is a subset of the state space that preserves the properties of the information type represented by Q .

Superposition principle. One of the defining postulates of quantum mechanics according to which if states $|1\rangle, |2\rangle, \dots$ are distinguishable, then $\sum_i \alpha_i |i\rangle$ with $\sum_i |\alpha_i|^2 = 1$ is a valid quantum state. Such a linear combination is called a normalized superposition of the states $|i\rangle$.

System. An entity that can be in any of a specified number of states. An example is a desktop computer whose states are determined by the contents of its various memories and disks. Another example is a qubit, which can be thought of as a particle whose state space is identified with complex, two-dimensional, length-one vectors. Here, a system is always associated with a type of information that determines the properties of the state space. For example, for quantum information, the state space is a Hilbert space. For deterministic information, it is a finite set called an alphabet.

Unitary operator. A linear operator U on a Hilbert space that preserves the inner product. That is, $\langle Ux, Uy \rangle = \langle x, y \rangle$. If U is given in matrix form, then this expression is equivalent to $U^\dagger U = \mathbb{1}$.

Universal set of gates. A set of gates that satisfies the requirement that every allowed operation on information units can be implemented by a network of these gates. For quantum information, it means a set of gates that can be used to implement every unitary operator. More generally, a set of gates is considered universal if, for every operator U , there are implementable operators V arbitrarily close to U .