From Cbits to Qbits: Teaching computer scientists quantum mechanics

N. David Mermin
Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853-2501

(Received 22 July 2002; accepted 26 September 2002)

A strategy is suggested for teaching mathematically literate students, with no background in physics, just enough quantum mechanics for them to understand and develop algorithms in quantum computation and quantum information theory. Although the article as a whole addresses teachers of physics well versed in quantum mechanics, the central pedagogical development is addressed directly to computer scientists and mathematicians, with only occasional asides to their teacher. Physicists uninterested in quantum pedagogy may be amused (or irritated) by some of the views of standard quantum mechanics that arise naturally from this unorthodox perspective. © *2003 American Association of Physics Teachers.*

[DOI: 10.1119/1.1522741]

I. COMPUTER SCIENCE AND QUANTUM MECHANICS

> These ‘‘bras’’ and ‘‘kets’’—they’re just vectors!
> —Newly enlightened computer scientist[1]

There is a new audience for the teaching of quantum mechanics whose background and needs require a new style of quantum pedagogy. The audience consists of computer scientists. Compared with the usual students in an introductory quantum mechanics course, they are mathematically sophisticated, but are often ignorant of and uninterested in physics. They want to understand the applications of quantum mechanics during the past dozen years to information processing, and their focus is exclusively on algorithms (software), not engineering (hardware).

Although the obstacles to quantum computers becoming a viable technology are formidable, the profound consequences of quantum mechanics for the theory of computation discovered during the past decade ought to be part of the intellectual equipment of every computer scientist, if only because it provides dramatic proof that the abstract analysis of computation cannot be divorced from the physical means available for its execution. Future computer scientists ought to learn quantum mechanics.

But how much quantum mechanics? In December 2001 I was at a conference on quantum computation and information at the Institute for Theoretical Physics in Santa Barbara. At lunch one day I remarked to the Director of the ITP that I spent the first four or five lectures of my course[2] on quantum computation teaching the necessary quantum mechanics to the computer scientists in the class. His response was that any application of quantum mechanics that could be taught after only a four hour introduction to the subject could not have serious intellectual content. After all, he remarked, it takes any physicist years to develop a feeling for quantum mechanics.

It’s a good point. Nevertheless, it is a fact that computer scientists and mathematicians with no background in physics have been able quickly to learn enough quantum mechanics to understand and contribute importantly to the theory of quantum computation, even though quantum computation repeatedly exploits the most notoriously paradoxical features of the subject. There are three main reasons for this:

First, a quantum computer—or, more accurately, the abstract quantum computer that one hopes some day to be able to realize—is an extremely simple example of a physical system. It is discrete, not continuous. It is made out of a finite number of units, each of which is the simplest possible kind of quantum mechanical system, a 2-state system, whose possible behavior is highly constrained and easily analyzed. Much of the analytical complexity of learning quantum mechanics is connected to mastering the description of continuous (infinite-state) systems in $(3+1)$-dimensional spacetime. By restricting attention to discrete transformations acting on collections of 2-state systems, one can avoid much suffering (and lose much wisdom, none of it—at least at this stage of the art—relevant to the theory of quantum computation.)

Second, the most difficult part of learning quantum mechanics is to get a good feeling for how the abstract formalism can be applied to actual phenomena in the laboratory. Such applications almost invariably involve formulating oversimplified abstract models of the real phenomena, to which the quantum formalism can effectively be applied. The best physicists have an extraordinary intuition for what features of the actual phenomena are essential and must be represented in the abstract model, and what features are inessential and can be ignored. It takes years to develop such intuition. Some never do. The theory of quantum computation, however, is only concerned with the abstract model—the easy part of the problem.

Third, to understand how to *build* a quantum computer, or to study what physical systems are promising candidates for realizing such a device, you must indeed have many years of experience in quantum mechanics and its applications under your belt. But if you only want to know what such a device is capable of doing in principle, then there is no reason to get involved in the really difficult physics of the subject. The same holds for ordinary (‘‘classical’’) computers: one can be a masterful practitioner of computer science without having the foggiest notion of what a transistor is, not to mention how it works.

So although the approach to quantum mechanics for computer scientists sketched below is focused and limited in scope, it is neither oversimplified nor incomplete, for the special task for which it is designed. (There is, however, an isolated subset of quantum-computational theory called adiabatic quantum computation that uses the quantum system more like an analogue than a digital computer, and does require a somewhat broader view of quantum theory.)

II. CLASSICAL BITS

The first step in teaching quantum mechanics to computer scientists is to reformulate the language of conventional (classical) computation in an unorthodox manner that introduces much of the quantum formalism in an entirely familiar setting

To begin, we need a term for a physical system that can exist in two unambiguously distinguishable states, which are used to represent 0 and 1. Often such a system is called a *bit*, but this can obscure the important distinction between the abstract bit (0 or 1) and the physical system used to represent it. If one could establish a nomenclature for the field at this late date, I would argue for the term Cbit for a classical physical system used to represent a bit, in parallel with the term Qbit for its quantum generalization. Unfortunately, the orthographically preposterous term *qubit* currently holds sway for the quantum system,[3] while *bit* is used indiscriminately for both the classical system and the abstract bit. Because clear distinctions between bits, Cbits, and Qbits are crucial in the exposition that follows, I shall use this unfashionable terminology. It is inspired by Paul Dirac's early use of *c-number* and *q-number* to describe classical variables and their generalizations to quantum-mechanical operators. (''Cbit'' and ''Qbit'' are preferable to ''c-bit'' and ''q-bit,'' because the terms themselves often appear in hyphenated constructions.)

It can be fruitful, even on the strictly classical level, to represent the two states of a Cbit by a pair of orthonormal 2-vectors, denoted by the symbols

$$|0\rangle \qquad |1\rangle. \tag{1}$$

This notation for vectors also goes back to Dirac. (For reasons too silly to go into, he called such vectors *kets*, a terminology that has survived to the present day.)

To do nontrivial computation requires more than one Cbit. It is convenient (and, as we shall see in a moment, even natural) to represent the four states of two Cbits as four orthogonal vectors in four dimensions, formed by the tensor products of two such pairs:

$$|0\rangle\otimes|0\rangle \qquad |0\rangle\otimes|1\rangle \qquad |1\rangle\otimes|0\rangle \qquad |1\rangle\otimes|1\rangle. \tag{2}$$

One often omits the $\otimes$, writing (2) in the more compact, but equivalent form,

$$|0\rangle|0\rangle \qquad |0\rangle|1\rangle \qquad |1\rangle|0\rangle \qquad |1\rangle|1\rangle, \tag{3}$$

or, more readably,

$$|00\rangle \qquad |01\rangle \qquad |10\rangle \qquad |11\rangle, \tag{4}$$

or, most compactly of all, using the decimal representation of the 2-bit number represented by the pair of Cbits,

$$|0\rangle_2 \qquad |1\rangle_2 \qquad |2\rangle_2 \qquad |3\rangle_2. \tag{5}$$

The subscript 2 is necessary in this last form, because in going from binary to decimal, we lose the information of how many Cbits the vector describes, making it necessary to indicate in some other way whether $|3\rangle$ means $|11\rangle=|3\rangle_2$ or $|011\rangle=|3\rangle_3$ or $|0011\rangle=|3\rangle_4$, etc.

As this last remark illustrates, one represents the states of n Cbits as the 2^n orthonormal vectors in 2^n dimensions,

$$|x\rangle_n, \quad 0\leqslant x<2^n, \tag{6}$$

given by the n-fold tensor products of n mutual orthogonal pairs of orthogonal 2-vectors. Thus, for example,

$$\begin{aligned}|19\rangle_6=|010011\rangle&=|0\rangle|1\rangle|0\rangle|0\rangle|1\rangle|1\rangle\\ &=|0\rangle\otimes|1\rangle\otimes|0\rangle\otimes|0\rangle\otimes|1\rangle\otimes|1\rangle.\end{aligned} \tag{7}$$

That the tensor product is a convenient and highly appropriate way to represent multi-Cbit states becomes clear if one expands the vectors representing each Cbit as column vectors,

$$|0\rangle\leftrightarrow\begin{pmatrix}1\\0\end{pmatrix}, \quad |1\rangle\leftrightarrow\begin{pmatrix}0\\1\end{pmatrix}. \tag{8}$$

The corresponding column vectors for tensor products are

$$\begin{pmatrix}y_0\\y_1\end{pmatrix}\begin{pmatrix}z_0\\z_1\end{pmatrix}\leftrightarrow\begin{pmatrix}y_0z_0\\y_0z_1\\y_1z_0\\y_1z_1\end{pmatrix}, \tag{9}$$

$$\begin{pmatrix}x_0\\x_1\end{pmatrix}\begin{pmatrix}y_0\\y_1\end{pmatrix}\begin{pmatrix}z_0\\z_1\end{pmatrix}\leftrightarrow\begin{pmatrix}x_0y_0z_0\\x_0y_0z_1\\x_0y_1z_0\\x_0y_1z_1\\x_1y_0z_0\\x_1y_0z_1\\x_1y_1z_0\\x_1y_1z_1\end{pmatrix}, \tag{10}$$

etc.

Thus, for example, the 8-dimensional column vector representing $|5\rangle_3$ is given by

$$|5\rangle_3=|101\rangle=|1\rangle|0\rangle|1\rangle=\begin{pmatrix}0\\1\end{pmatrix}\begin{pmatrix}1\\0\end{pmatrix}\begin{pmatrix}0\\1\end{pmatrix}=\begin{array}{cc}\left(\begin{array}{c}0\\0\\0\\0\\0\\1\\0\\0\end{array}\right) & \begin{array}{l}\scriptstyle 0\\ \scriptstyle 1\\ \scriptstyle 2\\ \scriptstyle 3\\ \scriptstyle 4\\ \scriptstyle 5\\ \scriptstyle 6\\ \scriptstyle 7\end{array}\end{array} \tag{11}$$

which has a 0 in every entry except for a 1 in the entry labeled by the integer 5 that the three Cbits represent. (Label the entries by counting down 0,1,2... from the top. The small numerals on the extreme right in (11) make this labeling explicit.) This general rule for the column vector representing $|x\rangle_n$, 1 in position x and 0 everywhere else, is the obvious generalization to n Cbits of the form for a 1-Cbit column vector. It is an automatic consequence of standard tensor-product notation.

III. OPERATIONS ON CBITS

In quantum computation almost all operations on Qbits are reversible. (An example of an irreversible operation is Erase: $|0\rangle\rightarrow|0\rangle,|1\rangle\rightarrow|0\rangle$. It is irreversible because one cannot reconstruct the input from the output: it has no inverse.) The single exception is the operation or process called ''measurement'' described in Sec. VI. Measurement plays no role in classical computation (or, perhaps more accurately, a role so trivial that it is not recognized explicitly as a part of the computational process). Because Cbit states turn out to be a (tiny) subset of Qbit states, our reformulation of classical bits and what can be done with them need only consider reversible operations on the Cbits.

There are just two reversible operations on a single Cbit: (1) Do nothing (identity operator $\mathbf{1}$):

$$\mathbf{1}|0\rangle=|0\rangle, \quad \mathbf{1}|1\rangle=|1\rangle. \tag{12}$$

(2) Flip it (flip operator $\mathbf{X}$):

$$\mathbf{X}|0\rangle=|1\rangle, \quad \mathbf{X}|1\rangle=|0\rangle. \quad (\sigma_x) \tag{13}$$

(I have indicated in parentheses the standard physicists' notation; quantum computer scientists prefer $\mathbf{X}$ to σ_x.)

Less trivial reversible operations are available on two Cbits. One can, for example, exchange the values of the bits they represent (swap operator $\mathbf{S}$):

$$\mathbf{S}|xy\rangle=|yx\rangle. \tag{14}$$

In manipulating such muti-Cbit operations, it is useful to have a compact notion for the action on a many-Cbit state of operations that act on only a single one of the Cbits. One labels the Cbits by integers 0,1,2,... (starting with zero on the right) associated with the power of 2 that each Cbit represents. Thus if x has the binary expansion $x=8x_3+4x_2+2x_1+x_0$, then

$$\begin{aligned}|x\rangle_4=|x_3x_2x_1x_0\rangle&=|x_3\rangle|x_2\rangle|x_1\rangle|x_0\rangle\\&=|x_3\rangle\otimes|x_2\rangle\otimes|x_1\rangle\otimes|x_0\rangle.\end{aligned} \tag{15}$$

An operation that acts only on Cbit #2 is

$$\mathbf{X}_2=\mathbf{1}\otimes\mathbf{X}\otimes\mathbf{1}\otimes\mathbf{1}. \tag{16}$$

Clearly the form with a subscript indicating which of the four Cbits is subject to the flip operation $\mathbf{X}$ is more transparent than the explicit form of the operator tensor product on the right. The subscript notation is unavoidable when large numbers of Cbits are involved. From the definition of the operator tensor product it follows that, as desired,

$$\mathbf{X}_2[|x_3\rangle\otimes|x_2\rangle\otimes|x_1\rangle\otimes|x_0\rangle]=|x_3\rangle\otimes[\mathbf{X}|x_2\rangle]\otimes|x_1\rangle\otimes|x_0\rangle. \tag{17}$$

It is possible to build up meaningful multi-Cbit operations out of single-Cbit operations that, although formally well defined, act an on individual Cbit in a way that has no meaningful classical interpretation. Here, for example, is a meaningless operation on one Cbit which can be used to build up meaningful multi-Cbit operations:

$$\mathbf{Z}|0\rangle=|0\rangle, \quad \mathbf{Z}|1\rangle=-|1\rangle. \quad (\sigma_z) \tag{18}$$

The action of $\mathbf{Z}$ on the state $|1\rangle$, multiplying it by -1, although mathematically well defined on the 2-dimensional 1-Cbit vector space, produces a vector that has no meaning within the context of Cbits. Only the two vectors $|0\rangle$ and $|1\rangle$ have meaning as the two distinguishable states of the Cbit used to represent 0 and 1. Indeed, the introduction of a 2^n-dimensional vector space when we are only interested in a single set of 2^n orthonormal basis vectors could be viewed as extravagant conceptual overkill, except, perhaps, for the pleasing structure introduced by the column-vector representation of the tensor product. The only classically meaningful reversible operations on n Cbits are the $(2^n)!$ different permutations of the 2^n basis vectors.

Nevertheless, a meaningless 1-Cbit operation like $\mathbf{Z}$ can acquire classical meaning when used in conjunction with other such meaningless operations in a multi-Cbit context. As an important example, notice that the 2-Cbit operation $\frac{1}{2}(\mathbf{1}+\mathbf{Z}_1\mathbf{Z}_0)$ acts as the identity on the 2-Cbit states $|0\rangle|0\rangle$ and $|1\rangle|1\rangle$, while giving 0 (another classically meaningless output) when acting on $|0\rangle|1\rangle$ or $|1\rangle|0\rangle$. The operation $\frac{1}{2}(\mathbf{1}-\mathbf{Z}_1\mathbf{Z}_0)$, on the other hand, acts as the identity on $|0\rangle|1\rangle$ and $|1\rangle|0\rangle$, while giving 0 on $|0\rangle|0\rangle$ and $|1\rangle|1\rangle$. Evidently both are projection operators in the full vector space spanned by all the 2-Cbit basis states. (More precisely the projection operators are their linear extensions to the full space from the basis on which they are defined. Quite generally any operation whose action is defined only on the classical basis states can be identified with its linear extension to the whole vector space.)

Because the operation $\mathbf{S}_{10}$, which exchanges the values of Cbits 1 and 0, acts as the identity if their state is $|00\rangle$ or $|11\rangle$ and as the double-flip operator $\mathbf{X}_1\mathbf{X}_0$ if their state is $|01\rangle$ or $|10\rangle$, we are led to the following operator representation of $\mathbf{S}_{10}$:

$$\mathbf{S}_{10}=\tfrac{1}{2}(\mathbf{1}+\mathbf{Z}_1\mathbf{Z}_0)+\mathbf{X}_1\mathbf{X}_0\tfrac{1}{2}(\mathbf{1}-\mathbf{Z}_1\mathbf{Z}_0), \tag{19}$$

or

$$\mathbf{S}_{10}=\tfrac{1}{2}(\mathbf{1}+\mathbf{Z}_1\mathbf{Z}_0+\mathbf{X}_1\mathbf{X}_0-\mathbf{Y}_1\mathbf{Y}_0), \tag{20}$$

where

$$\mathbf{Y}=\mathbf{XZ}. \quad (-i\sigma_y) \tag{21}$$

(Note that 1-Qbit operators acting on different Qbits (like $\mathbf{X}_1$ and $\mathbf{Z}_0$) commute even though the 1-Qbit operators ($\mathbf{X}$ and $\mathbf{Z}$) do not commute when acting on the same Qbit.) I digress to remark that this "classical" derivation of the exchange operator is simpler and more transparent than the standard quantum mechanical derivation, which invokes the full-blown theory of angular momentum.

Another important example of a 2-Cbit operation is the controlled-NOT or reversible XOR:

$$\mathbf{C}_{10}|x\rangle|y\rangle=(\mathbf{X}_0)^x|x\rangle|y\rangle=|x\rangle|y\oplus x\rangle, \tag{22}$$

(where $\oplus$ denotes addition modulo 2), which flips Cbit 0 (the *target* Cbit) if and only if Cbit 1 (the *control* Cbit) has the value 1. We can build this operation out of 1-Qbit projections,

$$\mathbf{C}_{10}=\tfrac{1}{2}(\mathbf{1}+\mathbf{Z}_1)+\mathbf{X}_0\tfrac{1}{2}(\mathbf{1}-\mathbf{Z}_1)=\tfrac{1}{2}(\mathbf{1}+\mathbf{Z}_1+\mathbf{X}_0-\mathbf{X}_0\mathbf{Z}_1). \tag{23}$$

In this form one sees a curious symmetry: interchanging the operations $\mathbf{X}$ and $\mathbf{Z}$ has the effect of exchanging the roles of target and control Cbit, converting $\mathbf{C}_{10}$ to $\mathbf{C}_{01}$.

A classically meaningless operation that can be used to perform just this interchange is the *Hadamard transform*

$$\mathbf{H}=\frac{1}{\sqrt{2}}(\mathbf{X}+\mathbf{Z})=\frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}. \tag{24}$$

This transform takes the Cbit states $|0\rangle$ and $|1\rangle$ into the two classically meaningless linear combinations $1/\sqrt{2}\ (|0\rangle\pm|1\rangle)$. Because

$$\mathbf{X}^2=\mathbf{Z}^2=\mathbf{1}, \quad \mathbf{XZ}=-\mathbf{ZX}, \tag{25}$$

it follows that

$$\begin{aligned}&\mathbf{H}^2=\tfrac{1}{2}(\mathbf{X}+\mathbf{Z})^2=\mathbf{1},\\&\mathbf{HX}=(\mathbf{X}+\mathbf{Z})\mathbf{X}/\sqrt{2}=\mathbf{Z}(\mathbf{X}+\mathbf{Z})/\sqrt{2}=\mathbf{ZH},\end{aligned} \tag{26}$$

and therefore

$$\mathbf{HXH}=\mathbf{Z}, \quad \mathbf{HZH}=\mathbf{X}. \tag{27}$$

Consequently. we can use four classically meaningless operations $\mathbf{H}$ to achieve a classically meaningful task: interchanging the role of target and control Cbits:

$$\mathbf{C}_{01}=(\mathbf{H}_1\mathbf{H}_0)\mathbf{C}_{10}(\mathbf{H}_1\mathbf{H}_0).$$

IV. QUANTUM BITS

We have represented the 2^n states of n Cbits as a basis of 2^n orthonormal vectors in a 2^n-dimensional vector space constructed as the n-fold tensor product of n 2-dimensional vector spaces. Although the only classically meaningful operations on these vector spaces consist of permutations of these classical basis vectors, we have been able to construct such operations, or reveal relations among them, by introducing classically meaningless operations that multiply basis vectors by scalars (in particular 0 or -1) or (like the Hadamard transform (24)) take them into non-trivial linear combinations. One such construction, the form (20) of the exchange operator, would achieve an even more pleasing form were we to introduce $\sqrt{-1}$, replacing $\mathbf{Y}$ with $i\mathbf{Y}$. This would also restore another symmetry, because $\mathbf{X}=\sigma_x$, $i\mathbf{Y}=\sigma_y$, and $\mathbf{Z}=\sigma_z$ are all hermitian.

One is reminded of arithmetic before the introduction of $\sqrt{-1}$. By introducing the ''meaningless'' quantity i, we are able to achieve great simplifications among certain relations connecting purely ''meaningful'' real numbers. The bold next step is to declare the meaningless to be meaningful too, taking full advantage of the expanded number system.

A major part of quantum mechanics consists of an analogous expansion of the notion of the state of a Cbit, called in this extended setting a quantum bit or *Qbit*. We democratically expand the set of meaningful states from the 2^n special orthonormal states, known in this broader setting as the *classical basis* (or, in the prevailing but less informative terminology, the *computational basis*) to arbitrary unit vectors from the entire vector space consisting of all linear combinations (called *superpositions*) of classical basis states with complex coefficients (called *amplitudes*).

Thus the general state of a single Qbit is a superposition of the two classical-basis states

$$|\psi\rangle=\alpha|0\rangle+\beta|1\rangle, \tag{28}$$

where the amplitudes α and β are complex numbers constrained only by the normalization condition

$$|\alpha|^2+|\beta|^2=1. \tag{29}$$

The general state of n Qbits has the form

$$|\Psi\rangle=\sum_{0\leqslant x<2^n}\alpha_x|x\rangle_n, \tag{30}$$

with complex amplitudes constrained only by the normalization condition

$$\sum_{0\leqslant x<2^n}|\alpha_x|^2=1. \tag{31}$$

Physics offers many examples of physical systems—Qbits—whose natural description is in terms of states that are precisely these peculiar generalizations of the states of classical bits that expand the constrained set of classical basis vectors to the entire complex vector space that they span. The most elementary physical examples are the polarization states of a photon or the spin states of a spin-$\frac{1}{2}$ particle. For an understanding of quantum-computational algorithms it is no more important to know about the detailed physics of such systems than it is to know about the detailed physics of transistors for an understanding of classical algorithms.

We shall return momentarily to the consequences of a set of Qbits having such nonclassical states, but the first thing to note is that by expanding the set of states from the classical basis vectors to arbitrary unit vectors in the entire complex vector space spanned by the classical basis, we have already introduced one of the most profound differences between Cbits and Qbits:

The most general possible state of two Cbits has the form

$$|\Psi\rangle=|x_1\rangle|x_0\rangle. \tag{32}$$

This can be described as a state in which Cbit #1 has the state $|x_1\rangle$ and Cbit #0, the state $|x_0\rangle$: each individual Cbit has a state of its own. On the other hand, the most general possible state of two Qbits has the form

$$\begin{aligned}|\Psi\rangle&=\alpha_3|3\rangle_2+\alpha_2|2\rangle_2+\alpha_1|1\rangle_2+\alpha_0|0\rangle_2\\&=\alpha_3|1\rangle|1\rangle+\alpha_2|1\rangle|0\rangle+\alpha_1|0\rangle|1\rangle+\alpha_0|0\rangle|0\rangle.\end{aligned} \tag{33}$$

If each Qbit had a state of its own, this 2-Qbit state would be, under the obvious generalization of the rule for multi-Cbit states, the tensor product of those two 1-Qbit states. The 2-Qbit state would thus have the general form

$$\begin{aligned}|\psi\rangle|\phi\rangle&=(\alpha|1\rangle+\beta|0\rangle)(\gamma|1\rangle+\delta|0\rangle)\\&=\alpha\gamma|1\rangle|1\rangle+\alpha\delta|1\rangle|0\rangle+\beta\gamma|0\rangle|1\rangle+\beta\delta|0\rangle|0\rangle.\end{aligned} \tag{34}$$

But the state $|\Psi\rangle$ in Eq. (33) cannot have this form unless $\alpha_3\alpha_0=\alpha_2\alpha_1$.

So in a general multi-Qbit state each individual Qbit has no state of its own. This is the first major way in which Qbits differ from Cbits. States of n Qbits in which no subset of fewer than n have states of their own are called *entangled*. Generic n-Qbit states are entangled. The amplitudes in the expansion (30) have to satisfy special constraints for the state to be a tensor product of states associated with fewer than n Qbits.

V. OPERATIONS ON QBITS

Quantum algorithms are constructed of operations that act linearly on the state of n Qbits, while preserving the normalization condition (31). The linear norm-preserving operators on a complex vector space are the *unitary* operators. So the basic ingredients of a quantum algorithm are unitary operators on the 2^n-dimensional complex space:

$$|\Psi\rangle\rightarrow\mathbf{U}|\Psi\rangle,\quad \mathbf{U}\ \text{unitary}. \tag{35}$$

The classical operations—permutations of the 2^n classical basis vectors (more precisely, linear extensions of the permutations from the basis on which they are defined to the whole space)—are special cases of such operators.

The problem of how to implement physically such unitary transformations is a question of quantum-computational engineering, just as the question of how to produce permutations of the values of a collection of Cbits is a question of classical-computational engineering. All that need concern the designer of quantum-computational software, however, is that unitary transformations constitute the full field of available operations (except for measurement, as described below in Sec. VI). For practical reasons—software designers should be willing to take into account constraints suggested by engineering practicalities—the available set of unitary transformations is usually restricted to those that can be built up out of products of unitary transformations, each of which act only on single Qbits or only on pairs of Qbits, and an important part of the ingenuity of quantum programming is devoted to how best to build up more interesting transformations as products of these basic units. (In the case of Cbits this is quite straightforward: an arbitrary permutation of n Cbits can be expressed as a product of 2-Cbit transpositions—swap operators.)

So if we view the 2^n states of n classical bits as the 2^n orthonormal basis vectors $|x\rangle_n$ in a 2^n-dimensional vector space, and the reversible operations we can perform on the Cbits as simply the permutations of these basis vectors, then the generalization to n quantum bits is extremely simple: the states of Qbits consist of all the normalized complex linear combinations of the classical basis vectors, and the reversible operations we can perform on the Qbits consist of all unitary transformations. The classical states and classical operations are a very small subset of the quantum states and quantum operations.

It looks as if the extension from Cbits to Qbits opens up an enormously richer landscape of computational possibilities. Although the state of one Cbit is specified by a single bit of information, specifying the state of one Qbit requires infinitely many bits of information: two complex numbers constrained only by the normalization condition (29). And instead of being limited to shuffling a finite collection of Cbit states through permutation, one can act on Qbits with a continuous collection of unitary transformations. Because it is no more complex a matter to prepare a given state for Qbits than it is for Cbits and because it is no more complex a matter to implement a broad range of unitary transformations on Qbits than it is to implement permutations on Cbits, the extension from Cbits to Qbits would appear to bring us to a new level of computational power.

But there is a catch! Qbits suffer from a major limitation that does not afflict Cbits. Although their state contains vast amounts of information, given n Qbits in some state $|\Psi\rangle$, there is nothing you can do to the Qbits that enables you to learn what $|\Psi\rangle$ is. There is thus no way to extract anything like the huge amount of information contained in the amplitudes α_x.

What, then, are Qbits good for? How can we exploit their greater flexibility to do anything useful at all?

VI. MEASUREMENT: HOW TO SQUEEZE INFORMATION OUT OF QBITS

A. The Born rule

The very limited possibilities for extracting information is the second major way in which Qbits differ from Cbits. If we have n Cbits in the general classical state $|x\rangle_n$, finding out what the state is—learning the number x—is unproblematic. Indeed, it is so straightforward that the act of learning the state is generally not even regarded as a formal part of the computation. One simply looks (on a display or a printout). Importantly, the state of the Cbits is unaltered by this acquisition of information. Once the computer has ceased to operate on the Cbits, their state remains $|x\rangle_n$ whether or not anybody takes the trouble to ascertain the particular value of x.

Things could not be more different for Qbits. If one has n Qbits in the state

$$|\Psi\rangle_n = \sum_x \alpha_x |x\rangle_n, \tag{36}$$

there is nothing one can do to them to learn the values of the amplitudes α_x. There is only one way to extract any information from the Qbits: to *measure* them. Measuring n Qbits consists of subjecting them to a device that produces (at a display or a printout) an integer x in the range $0 \leqslant x < 2^n$. The only link between the state $|\Psi\rangle$ one may have labored to impose on the Qbits and the value of x revealed by the measurement is this: the probability of getting the output x is just $p_x = |\alpha_x|^2$, where α_x is the amplitude of $|x\rangle_n$ in the expansion (36) of $|\Psi\rangle_n$. This connection between amplitudes and the probabilities of measurement outcomes is known as the *Born rule*, after the physicist Max Born. The condition that the states be unit vectors is thus the condition that the sum of the probabilities of all the possible measurement outcomes should be 1.

You might think that by measuring repeatedly, one could at least get some good statistics on the distribution of the magnitudes $|\alpha_x|$, but this possibility of additional partial information about $|\Psi\rangle$ is ruled out by a second fundamental proviso of the Born rule: once the value x has been indicated by the measurement, the state of the n Qbits is no longer $|\Psi\rangle_n$, but $|x\rangle_n$. The postmeasurement state $|x\rangle_n$ contains no trace of the information present in the premeasurement state $|\Psi\rangle$ (beyond revealing that $\alpha_x \neq 0$) and is nothing more than the classical state associated with the value of x indicated by the measuring device.

Physicists, in a nomenclature that invites misinterpretation, like to say that the state $|\Psi\rangle_n$ *collapses* or *is reduced* to the state $|x\rangle_n$ by the measurement. The conservative way to put it is simply to specify the relation between the states immediately before and immediately after the measurement, in a way that suggests no mechanism for the change of state, confers no objective status on it, and makes no commitment to what (if anything) a change in state implies about what (if anything) has happened to the Qbits themselves.

You might wonder how we can learn anything at all of computational interest under these wretched conditions. The general trick is to produce, through a cunningly constructed unitary transformation, a superposition (36) in which most of the amplitudes α_x are zero or very close to zero, with useful information being carried by any of the values of x that have a significant probability of being indicated by the measurement. It is also important to be seeking information that, once possessed, can easily be confirmed (for example, the factors of a large number) so that one is not misled by the occasional irrelevant low probability outcome.

Clearly the action of a measurement on the state of n Qbits is irreversible: any state Ψ_n with non-zero amplitude α_x is capable of becoming the state $|x\rangle_n$ after a measurement. There is no way to reconstruct the input from the output. Measurement is, however, the only irreversible operation on Qbits. All other operations are unitary.

The Born rule contains, as a special case, the unproblematic character of extracting information from Cbits. If the state $|\Psi\rangle$ of n Qbits happens to be one of the 2^n classical-basis states $|x_0\rangle_n$ then $\alpha_x = 0$, $x \neq x_0$, and $\alpha_{x_0} = 1$. So the result of measuring the Qbits is x_0 with probability 1. The second proviso of the Born rule then requires that the state of the Qbits is $|x_0\rangle_n$ after the measurement—that is, the post-measurement state continues to be what it was before the measurement. The statistical, state-altering character of the outcome of a measurement of n Qbits in a general state becomes the deterministic, state-preserving, unproblematic classical extraction of information when the state is one of the 2^n classical states.

A technical remark for physicists: In this approach to quantum mechanics, it is useful to restrict the term ''measurement'' to what a broader and more conventional use of the term would characterize as ''measurement in the classical basis.'' Because measurement in any other basis could be accomplished by applying an appropriate unitary transformation—one that takes the basis of interest into the classical basis—followed by measurement in the classical basis, this restriction of the scope of the term ''measurement'' does not preclude more general possibilities.

B. Generalization of the Born rule to partial measurements

There is a generalization of the Born rule, not often explicitly noted in quantum-mechanics texts, that is needed whenever some but not all of the Qbits are measured, as often happens in a quantum computation. Suppose we have $m+n$ Qbits, and we decide to measure only m of them. By representing the $m+n$ bit number z as x,y, the concatenation of the m and n bit binary strings representing x and y, we can write the state of the $m+n$ Qbits as

$$|\Psi\rangle_{m+n}=\sum_{x,y}\alpha_{x,y}|x,y\rangle_{m+n}. \tag{37}$$

Suppose that we decide to measure only the m Qbits on the left. (The rule for the more general choice of which Qbits to measure is the obvious generalization of the one enunciated below.) The generalized Born rule states that the measurement will indicate x, $0\leqslant x<2^m$, with probability

$$p_x=\sum_{0\leqslant y<2^n}|\alpha_{x,y}|^2, \tag{38}$$

and that after the value of x is indicated, the state of the $m+n$ Qbits is changed from $|\Psi\rangle_{m+n}$ to $|x\rangle_m|\Phi_x\rangle_n$, where

$$|\Phi_x\rangle_n=p_x^{-1/2}\sum_y \alpha_{x,y}|y\rangle_n. \tag{39}$$

If one immediately follows a measurement of the m Qbits on the left, with a measurement of the remaining n Qbits on the right, then this measurement ought to be tantamount to directly measuring all $m+n$ Qbits. And indeed, if one applies the generalized Born rule twice—first to the measurement of the m Qbits on the left and then to the measurement of the remaining n on the right—one recovers the ordinary Born rule.

Although the generalized Born rule does not follow from the ordinary Born rule, it is equivalent to the ordinary Born rule supplemented by two very reasonable further conditions:

(1) Suppose that between time t and t', no unitary transformations act on the m Qbits on the left, but arbitrary unitary transformations may act on the n Qbits on the right—that is, the only unitary transformations acting on the $m+n$ Qbits between t and t' are of the form $\mathbf{U}=\mathbf{1}_m\otimes\mathbf{V}_n$. Then the statistical distribution of outcomes if all $m+n$ Qbits are measured at time t' is unaltered if the m Qbits on the left are measured at any earlier time between t and t'. Informally, once the computer ceases from further action on any group of Qbits, you do not have to wait to the end of the full computation before measuring those Qbits.

(2) For a group of n Qbits to be in the state $|\Phi\rangle$ means nothing more (or less) than this: If the Qbits are measured after the application of an arbitrary unitary transformation $\mathbf{V}$, then the distribution of measurement outcomes will be that specified by the Born rule for n Qbits in the state $\mathbf{V}|\Phi\rangle$.

The most important principles formulated in Secs. II–VI are in Table I, which summarizes the relevant features of Qbits by contrasting them to the analogous features of Cbits. In the table I have introduced the term ''Bit,'' with an upper-case B, to mean ''Qbit or Cbit'' (in contrast to ''bit,'' with a lower-case b, which means ''0 or 1'').

VII. CAUTIONARY REMARKS AND QUASI-PHILOSOPHICAL REFLECTIONS

A. An important warning

It is extremely important to avoid a tempting misinterpretation—a gross oversimplification—of quantum superpositions of classical states, as illustrated by the following simple example:

A Qbit in the state $|\psi\rangle=1/\sqrt{2}|0\rangle+1/\sqrt{2}|1\rangle$ is *not* the same as a Qbit that is either in state $|0\rangle$ or state $|1\rangle$ with equal probability, even though in either case a measurement will indicate 0 or 1 with equal probability. To see that the two cases are inherently different, suppose a Hadamard transform $\mathbf{H}=1/\sqrt{2}(\mathbf{X}+\mathbf{Z})$ is applied to the Qbit just before the measurement is made. Because

$$\mathbf{H}|0\rangle=\frac{1}{\sqrt{2}}(|0\rangle+|1\rangle),\quad \mathbf{H}|1\rangle=\frac{1}{\sqrt{2}}(|0\rangle-|1\rangle), \tag{40}$$

Table I.

CLASSICAL versus QUANTUM BITS	**Cbits**	**Qbits**
States of n Bits	$|x\rangle_n$, $0\leqslant x<2^n$	$\Sigma\alpha_x|x\rangle_n$, $\Sigma|\alpha_x|^2=1$
Subsets of n Bits	Always have states	Generally have no states
Reversible operations on states	Permutations	Unitary transformations
Can state be learned from Bits?	Yes	No
To get information from Bits	Just look	Measure
Information acquired	x	x with probability $|\alpha_x|^2$
State after information acquired	Same: still $|x\rangle$	Different: now $|x\rangle$

in the second case, whether the initial state is $|0\rangle$ or $|1\rangle$, the measurement after $\mathbf{H}$ is applied will continue to indicate 0 or 1 with equal probability. But in the first case, in which the initial state is $|\psi\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$, we have $\mathbf{H}|\psi\rangle = |0\rangle$ so the measurement after $\mathbf{H}$ is applied must necessarily indicate 0.

A Qbit in a superposition of classical-basis states is distinctly different from a Qbit that is in one or the other of those classical states with probabilities given by the squared moduli of the corresponding amplitudes. Superpositions have no classical interpretation. They are *sui generis*, an intrinsically quantum-mechanical construct, whose meaning derives only from the rules that characterize the reversible operations (unitary) that can be performed on them and the available means (measurement) for extracting information from them.

B. Meaning of the quantum state

People have been arguing about the meaning of the quantum state ever since the concept first appeared, with no indication that we are getting any closer to a consensus. These conceptual issues are unimportant for an understanding of quantum computation which only requires one to know how states are built up from other states (by appropriate unitary transformations) and how information can be extracted from Qbits in a given state (by measurement, according to the Born rules).

The initial state on which the unitary transformations operate is usually a classical-basis state $|x\rangle_n$. Such a state can be unambiguously identified as the post-measurement state of n Qbits after a measurement that indicated the value x. From this point of view the computational process begins and ends with a measurement, and the entire role of the state of the Qbits at any stage of a succession of unitary transformations is to encapsulate the probability of the outcomes, should the final measurement be made at that stage of the process, or to enable one to calculate new outcome probabilities, should further unitary transformations be applied before the measurement.

The notion that the state of n Qbits is simply a convenient compact mathematical device for calculating the correlations between the outcomes of two measurements on those Qbits, between which an arbitrary unitary transformation may have been applied, is often associated with the constellation of ideas about quantum mechanics called the *Copenhagen interpretation*. It is to be contrasted with the notion that the state of n Qbits is an objective physical property of those Qbits, in the same strong sense that we can view the state of n Cbits—the unique value x that they represent—as an objective property of those Cbits. People who regard the quantum state as objective in this sense tend to make a fuss about the fact that there are two quite different ways in which Qbits can change: deterministically and continuously (if one builds each unitary transformation out of many infinitessimal ones) via unitary transformations, and statistically and discontinuously via measurements. This dichotomy looses its content if one replaces ‘‘Qbits’’ by ‘‘the state of Qbits,’’ and recognizes that the state is nothing but a catalog of how different unitary transformations will result in different distributions of measurement outcomes—classical basis states, which alone can be viewed as objective.

Another pitfall of taking their state to be an objective property of the Qbits is that one can then succumb to the temptation to believe that the application of a series of unitary transformations to the Qbits implements a physical computation of all the resulting amplitudes α_x. The clue that the amplitudes have not all been calculated lies in the fact, noted above, that given the Qbits there is nothing whatever you can do with them to reveal the values of those amplitudes.

There are nevertheless some who believe that all the amplitudes α_x have the status of objective physical quantities, inaccessible though those quantities may be. Such people then wonder how that vast number of high-precision calculations (10^{30} different amplitudes if you have 100 Qbits) could all have been physically implemented. Those who ask such questions like to provide sensational but fundamentally silly answers involving vast numbers of parallel universes, invoking a point of view known as the *many worlds* interpretation of quantum mechanics. My own opinion is that, imaginative as this vision may appear, it is symptomatic of a lack of a much more subtle kind of imagination, which can grasp the exquisite distinction that quantum physics has forced upon us, between quantum states and objective properties.

C. Where’s Planck’s constant?

Where’s h-bar? Where is h-bar?!
—Disgruntled quantum optician.[4]

Like my disapproving colleague, some physicists may be appalled to have finished what purports to be an exposition of quantum mechanics—indeed, of applied (well, *gedanken* applied) quantum mechanics—without ever having run into Planck’s constant. How can this be?

The answer goes back to my first reason why enough quantum mechanics to understand quantum computation can be taught in a mere four hours. We are interested in discrete (2-state) systems and discrete (unitary) transformations. But Planck’s constant only appears in the context of continuously infinite systems (for example, position eigenstates) and continuous families of transformations (for example, time development) that act on them. Its role is to relate the conventional units in which we measure space and time, to the units in which it is natural quantum-mechanically to take the generators of the unitary transformations that produce translations in space or time.

If we are not interested in location in continuous space and are only interested in global rather than infinitesimal unitary transformations, then $\hbar$ need never enter the story. The engineer, who must figure out how to implement unitary transformations acting over time on Qbits located in different regions of physical space, must indeed deal with $\hbar$ and with Hamiltonians that generate the unitary transformations out of which the computation is built. But the designer of algorithms for the finished machine need only deal with the resulting unitary transformations, from which $\hbar$ has disappeared as a result, for example, of judicious choices by the engineers of the times over which the interactions that produce the unitary transformations act.

Deploring the absence of $\hbar$ from expositions of quantum computer science is rather like complaining that the I-V curve for a p-n junction never appears in expositions of classical computer science. It is to confuse computer *science* with computer *engineering*.

VIII. THAT IS ALL YOU NEED TO KNOW

Armed with the contents of Secs. II–VI, one is ready to embark on the exposition of quantum computer science. To be sure, there will be times when it is convenient to expand upon the minimal formalism developed above. But such expansions, for example the introduction of *bras* (as linear functionals on the kets), the introduction of density matrices, or the useful connection between **X**, **Y**, and **Z** and the group of 3-dimensional rotations, are all technical mathematical refinements within the basic structure of the complex vector space of Qbit states. They require no new physical principles for their development.

Sections II–VI provide all the quantum mechanics one needs to develop fully the factorization algorithm of Peter Shor, the search algorithm of Lov Grover, and their later generalizations (see for example, Ref. 2 and references cited therein). Only in developing the very important subject of quantum error correction is it necessary to introduce a new physical assumption, that the formalism developed to describe Qbits—quantum states, unitary transformations, the Born rules—describes not only Qbits, but anything else in the world that the Qbits might happen to interact with.

If this far from modest extension of the scope of the formalism proves too big a pill for computer scientists to swallow, one can compromise with a more limited model of error correction, in which the computer contains large numbers of extraneous Qbits. Ideally, these irrelevant Qbits are not coupled to the Qbits of interest, in the sense that all unitary transformations act only on the Qbits of interest or only (unimportantly and uninterestingly) on the extraneous Qbits. But unfortunately, there is a small amount of unintended coupling between the two sets of Qbits—unitary transformations whose action is not restricted to either the relevant or irrelevant Qbits—whose disruptive action on the relevant Qbits it is the task of error correction to undo. One can then remark, as an aside, that parts of the world outside the computer (or computationally irrelevant internal degrees of freedom of the computer) that cannot be be perfectly isolated from the parts that do the computation can always be well modeled as just such collections of extraneous Qbits.

A detailed view of how to erect the edifice of quantum computation on this foundation can be found in Chapters 2–5 of my lecture notes.[2] Chapter 6 describes a few further topics in the broader area of quantum information that can be built on this same foundation. I do not delve into these matters here because the subject of this essay has been how to teach computer scientists quantum mechanics—not quantum computation. I have therefore tried to restrict references to the computational applications of quantum mechanics to those that motivate the quantum-mechanical formalism, and those that address in broad general terms broad general questions that the formalism gives rise to (such as ‘‘How can this possibly lead to anything useful?’’)

ACKNOWLEDGMENTS

Supported by the National Science Foundation, Grant No. PHY0098429.

[1] A speaker at the second AQIP (Algorithms in Quantum Information Theory) workshop, Chicago, January 1999.

[2] Lecture notes and homework assignments can be found at `www.ccmr.cornell.edu/~mermin/qcomp/CS483.html`. The pedagogical approach sketched below is fleshed out in Chapter 1. Alternative introductions to quantum mechanics in the context of quantum computation have been given by Eleanor G. Rieffel and Wolfgang Polak, ‘‘An introduction to quantum computing for non-physicists,’’ quant-ph/9809016, and by Michael A. Nielsen and Isaac L. Chuang in their excellent textbook, *Quantum Computation and Quantum Information* Cambridge University Press, 2000.

[3] *Qubit* seems to have been used first in print by Benjamin Schumacher, ‘‘Quantum coding,’’ Phys. Rev. A **51**, 2738–2747 (1995). A brief history of the term can be found in the Acknowledgments of this paper Although *qubit* honors the English rule that *q* should be followed by *u*, it ignores the equally powerful requirement that *qu* should be followed by a vowel. My guess is that it has gained acceptance because it visually resembles an ancient English unit of distance, the homonymic *cubit*. To see its ungainliness with fresh eyes, imagine that Dirac had written *qunumber* instead of *q-number*, or that one erased transparencies and cleaned one’s ears with *Qutips*.

[4] Private communication to the author at the International Conference on Quantum Information, Rochester, June 2001.