Copenhagen computation: How I learned to stop worrying and love Bohr

N. D. Mermin

To celebrate the 60th birthday of Charles H. Bennett, I 1) publicly announce my referee reports for the original dense coding and teleportation papers, 2) present a very economical solution to the Bernstein–Vazirani problem that does not even hint at interference between multiple universes, and 3) describe how I inadvertently reinvented the Copenhagen interpretation in the course of constructing a simple, straightforward, and transparent introduction to quantum mechanics for computer scientists.

1. Preface: Present at the birth

David DiVincenzo, Patrick Hayden, and Barbara Terhal [1] have designated me the "midwife of teleportation" in recognition of my having written a favorable referee's report on the discovery paper [2] and having advised the editors that the proposed terminology made sense. Though this honorific raises vexing biological questions—can something with six fathers and no mother be brought forth by a midwife?—I accept the title with pride. As midwife it seemed appropriate for me to read my referee's report at the Bennett sixtieth birthday symposium, attended, as it was, by all six fathers. I reproduce it here too, since it shows me to have had a taste for Copenhagen computation (about which more below) even before Chris Fuchs [3] got to work on me.

Referee's Report: Bennett et al., "Teleporting . . ." LZ4539

This is a charming, readable, thought-provoking paper. It presents a novel application of EPR correlations. The character of the quantum state (how much is inherent in the physical system, how much is a representation of our knowledge) is still an extremely elusive notion. This novel method for duplicating a quantum state somewhere else by a combination of quantum correlations and classical information will become an important one of the intellectual tools available to anybody trying to clear up this murkiness.

While hunting down the above report I discovered, to my amazement, that the year before I had also refereed the discovery paper on dense coding [4]. (I was under the impression that I had paid no attention whatever to dense coding until shortly before its deconstruction in 2002 [5].)

Bennett and Wiesner, "Communication via one-and two-particle . . ." LT4749

Your question was: Does this qualify as "strikingly different" enough to publish? I have never read anything like it, and I have read a lot on EPR, though far from everything ever written. So as far as I know it is different.

But strikingly? The argument is very simple, so shouldn't the point be obvious? After reading the paper I put it aside and spent the next week working hard on something totally unrelated. Every now and then I would introspect to see if some way of looking at the argument had germinated that reduced it to a triviality. None had. Last night I woke up at 3am, fascinated and obsessed with it. Couldn't get back to sleep. That's my definition of "striking".

So I say it's strikingly different and I say publish it.

0018-8646/04/$5.00 © 2004 IBM

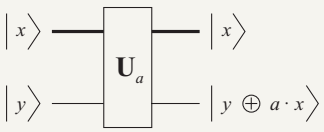

Figure 1

The black-boxed Bernstein–Vazirani subroutine $\mathbf{U}_a$. The heavy lines represent the n-Qbit input register; the light lines represent the 1-Qbit output register.

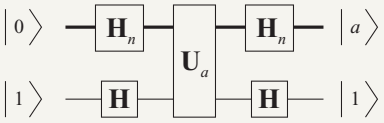

Figure 2

Quantum solution of the Bernstein–Vazirani problem. $\mathbf{H}_n$ is an n-fold tensor product of 1-Qbit Hadamards.

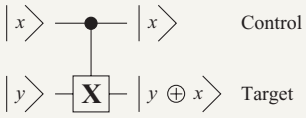

Figure 3

The 2-Qbit cNOT gate.

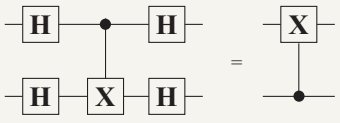

Figure 4

Hadamards can be used to interchange target and control Qbits.

Rereading these old reports reminded me that the myth that referees relish their power to reject papers is off the mark. Writing a favorable report for a good paper is sheer pleasure. Negative reports are no fun at all.

2. Prologue: Bernstein–Vazirani problem without parallel universes

The Bernstein–Vazirani problem presents one with a black-boxed subroutine, shown in **Figure 1**, whose action on $n + 1$ qubits is that of a unitary transformation $\mathbf{U}_a$ which takes the computational basis state $|x\rangle_n|y\rangle_1$ of an n-Qbit[1] input register and 1-Qbit output register into the state $|x\rangle_n|y \oplus x \cdot a\rangle_1$. Here $\oplus$ denotes addition modulo 2, $x \cdot a$ denotes the bitwise modulo 2 inner product of the two n-bit numbers x and a ($x \cdot a = x_{n-1}a_{n-1} \oplus \cdots \oplus x_1a_1 \oplus x_0a_0$), and a is some fixed but unknown n-bit integer with binary expansion $a = a_{n-1} \ldots a_1a_0$. The problem is to find the smallest number of invocations of the black box needed to learn a.

If the subroutine is applied to $x = 2^j$, the output register will be flipped if and only if $a_j = 1$, so a classical computer can determine a with n calls of the subroutine. Evidently there is no classical way to learn a with fewer than n calls, since one needs n independent linear relations among the bits of a. But with a quantum computer one can find a with just a single call of the subroutine, whatever the size of n.

This remarkable trick is accomplished by applying a Hadamard transformation,

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

to every one of the $n + 1$ Qbits both before and after the application of $\mathbf{U}_a$, as shown in **Figure 2**. If one initializes the input register to the state $|0\rangle_n$ and the output register to the state $|1\rangle$, then at the end of this process the input register is guaranteed to be in the n-Qbit state $|a\rangle_n$. So a can be learned by measuring each Qbit of the output register in the computational basis.

The conventional explanation for why this works goes like this:

1. Applying Hadamards to every Qbit of an input register initially in the n-Qbit state $|0\rangle_n$ results in a uniform superposition of all possible inputs:

$$\mathbf{H}_n|0\rangle_n = 2^{-n/2} \sum_{0 \le x < 2^n} |x\rangle_n . \tag{2}$$

2. Preparing the output register in the state $\mathbf{H}|1\rangle$ converts a bit-flip into a change of phase (specifically, a multiplication by -1).

[1] I use here the unorthodox spelling *Qbit* because it will be constantly juxtaposed to *Cbit*, a role the currently favored *qubit* cannot gracefully play.

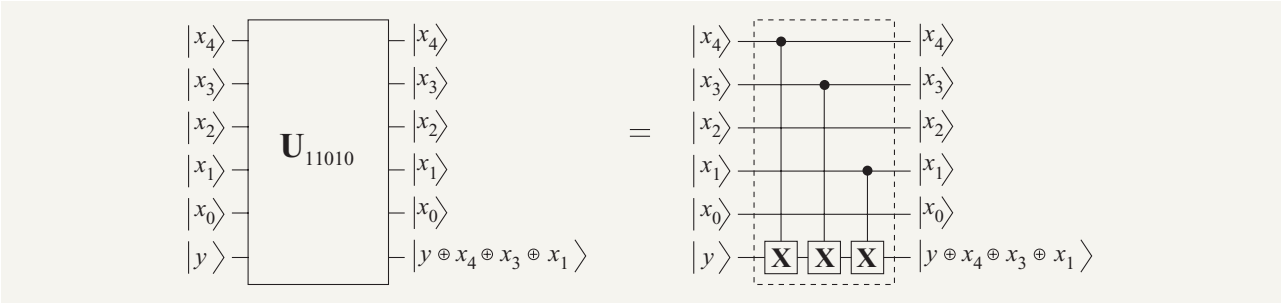

Figure 5

The black-boxed Bernstein–Vazirani oracle (shown for the case $n = 5$, $a = 11010$) behaves *as if* it contained a collection of cNOT gates.

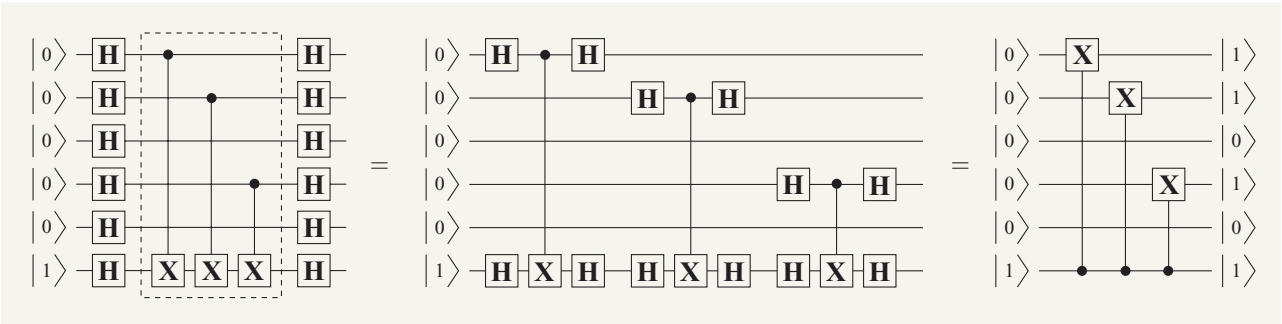

Figure 6

Because $\mathbf{U}_a$ behaves as a collection of cNOT gates, because Hadamards reverse the action of cNOT gates, and because the output register has been set to 1, we have a simple explanation for why a can be determined with only one invocation of $\mathbf{U}_a$.

3. Another application of Hadamards to the input register after the application of $\mathbf{U}_a$ introduces additional x-dependent phases according to the rule

$$\mathbf{H}_n|x\rangle_n = 2^{-n/2} \sum_{0 \le z < 2^n} (-1)^{x \cdot z} |z\rangle_n . \qquad (3)$$

4. A little arithmetic now reveals that the combined phases lead to complete destructive interference for every term characterizing the input register in the final superposition except for the single state $|a\rangle_n$.

This process is usually described as an application of massive quantum parallelism followed by destructive interference among all the unfavorable outcomes. People with overactive imaginations like to say that step 1 initializes a computer in each of 2^n parallel universes to each of the 2^n possible inputs. The remaining steps are cunningly designed to produce destructive interference among all those 2^n universes, in just such a way as to lead in every single universe to the presence of a in the input register at the end of the process.

There is, however, a much simpler way to understand why the circuit in Figure 2 behaves as advertised, which offers no hint of this metaphysical extravaganza. This approach merely notes that the effect of Hadamards on the basic 2-Qbit controlled-NOT (cNOT) gate, defined in **Figure 3**, is just to interchange the control and target Qbits, as shown in **Figure 4**. This follows from the fact that $\mathbf{H}^2 = \mathbf{1}$ and $\mathbf{HXH} = \mathbf{Z}$, where

$$\mathbf{X}|0\rangle = |1\rangle, \quad \mathbf{X}|1\rangle = |0\rangle, \quad \mathbf{Z}|0\rangle = |0\rangle, \quad \mathbf{Z}|1\rangle = -|1\rangle, \qquad (4)$$

and the fact that controlled-$\mathbf{Z}$ is symmetric under interchange of target and control Qbits.

The action of $\mathbf{U}_a$ shown in Figure 1 is identical to the action of a collection of cNOT gates—one for each nonzero bit of a. They all target the output register and are controlled by just those Qbits representing bits of x that correspond to nonzero bits of a. This is illustrated in **Figure 5** for $n = 5$ and $a = 11010$. Since sandwiching cNOT gates between Hadamards interchanges the control and target Qbits and since **H** is its own inverse, the magic of Bernstein–Vazirani follows at once, as shown in **Figure 6**, which makes perfect sense in just a single universe.

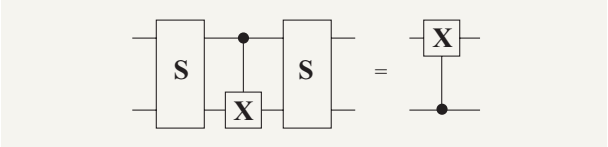

Figure 7

Classical SWAP gates also invert the action of a cNOT gate just as Hadamards do (Figure 4). But if swap gates are used to invert the cNOT gates in the final form of Figure 6, they get all tangled up with each other when one attempts to move them to the edges of the figure, as one can do with the Hadamards.

Notice that the only way quantum mechanics enters is through the ability of Hadamards to reverse the action of cNOT gates, as illustrated in Figure 4. Since this can also be done classically with 2-Qbit SWAP gates, as shown in **Figure 7**, the magic of quantum mechanics here lies entirely in the possibility it offers for reversing the roles of target and control Qbits using only 1-Qbit local operations. The power of Hadamards over classical SWAPs is that they can bring about the reversal without the need for any interaction between the two Qbits. If the six pairs of vertically separated Hadamards in the middle circuit of Figure 6 were vertically linked into irreducibly 2-Qbit gates, then they could no longer be moved to the extreme right and left of the circuit without leaving any traces in the central part, as in the circuit on the left in Figure 6.

3. How I invented the Copenhagen interpretation while teaching quantum mechanics to computer scientists

> *[I]n our description of nature the purpose is not to disclose the real essence of the phenomena but only to track down, so far as it is possible, relations between the manifold aspects of our experience.*
>
> —*Niels Bohr* [6]

For the past few years I have taught a course in quantum computation suitable for computer scientists having no background in physics [7]. My first challenge was to develop a minimalist introduction to quantum mechanics that straightforwardly conveyed in a few lectures everything a mathematically sophisticated student needed to know to understand discussions like, for example, that of the preceding section.

The advantage of teaching an approach to a subject as you develop it is that you get striking demonstrations of the ways in which it does and doesn't work. After several iterations and countless revisions, reorganizations, and reconstructions, the process seemed to be converging. It was only then that I realized that the unproblematic, no-nonsense, lucid, practical pedagogical approach that had so painfully evolved out of my clumsy initial attempts was nothing but the standard Copenhagen interpretation. What follows, therefore, is my vision of why quantum computation, far from demonstrating the existence of the multiverse, provides the simplest and most compelling example of a major application of quantum mechanics which the Copenhagen point of view fits like a glove.

We begin with a silly formulation of ordinary non-quantum *classical* computation, based on representing the integers less than N as orthonormal vectors in N dimensions:

$$0 \to \begin{pmatrix}1\\0\\0\\0\\\vdots\end{pmatrix},\ 1 \to \begin{pmatrix}0\\1\\0\\0\\\vdots\end{pmatrix},\ 2 \to \begin{pmatrix}0\\0\\1\\0\\\vdots\end{pmatrix},\ 3 \to \begin{pmatrix}0\\0\\0\\1\\\vdots\end{pmatrix},\ 4 \to \begin{pmatrix}0\\0\\0\\0\\1\\\vdots\end{pmatrix}, \ldots . \tag{5}$$

This clumsy form takes on a rather simpler structure if N is a power of 2, so we specialize to the case $N = 2^n$. When $n = 1$ we have only two such vectors, which we denote by a more compact pair of symbols due to Dirac:

$$\begin{pmatrix}1\\0\end{pmatrix} = |0\rangle, \qquad \begin{pmatrix}0\\1\end{pmatrix} = |1\rangle. \tag{6}$$

To manipulate these two numbers in a computer, it is necessary to represent them by a physical system having two distinguishable configurations. Continuing to follow Dirac, we call any such physical system a *Cbit*[2] ("C" for "classical"). The vectors $|0\rangle$ and $|1\rangle$ associated with these two configurations are called the *states* of the Cbit.

If we have two Cbits ($n = 2$), their four states conveniently decompose into the tensor product of two 1-Cbit states:

$$|0\rangle_2 = \begin{pmatrix}1\\0\\0\\0\end{pmatrix} = \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} = |0\rangle|0\rangle = |00\rangle,$$

$$|1\rangle_2 = \begin{pmatrix}0\\1\\0\\0\end{pmatrix} = \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = |0\rangle|1\rangle = |01\rangle,$$

$$|2\rangle_2 = \begin{pmatrix}0\\0\\1\\0\end{pmatrix} = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} = |1\rangle|0\rangle = |10\rangle,$$

[2] The term "c-bit" doesn't work because one often needs to talk, for example, about 2-Cbit states.

$$|3\rangle_2 = \begin{pmatrix}0\\0\\0\\1\end{pmatrix} = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = |1\rangle|1\rangle = |11\rangle. \tag{7}$$

The last two forms in each line provide some simpler notations for these 2-Cbit states. Pause to admire how the quantum-mechanical practice of representing the states of composite systems by the tensor product of the subsystem states emerges automatically from the trivial representation of integers introduced in (5).

The tensor product extends straightforwardly to many Cbits: States of n Cbits can be expressed as tensor products of n 1-Cbit states. For example,

$$|5\rangle_3 = \begin{pmatrix}0\\0\\0\\0\\0\\1\\0\\0\end{pmatrix} = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = |1\rangle|0\rangle|1\rangle = |101\rangle. \tag{8}$$

While the operation **X** defined in (4) makes perfect sense for Cbits (representing the logical NOT), the operation **Z** makes no sense at all, since we have assigned no meaning to the sign of the state-vector that describes a Cbit. Nevertheless, combinations of operators **Z** on pairs of Cbits can be classically meaningful. For example,

$$\tfrac{1}{2}(\mathbf{1} + \mathbf{Z} \otimes \mathbf{Z}) \text{ projects on } |0\rangle|0\rangle, \quad |1\rangle|1\rangle,$$

$$\tfrac{1}{2}(\mathbf{1} - \mathbf{Z} \otimes \mathbf{Z}) \text{ projects on } |0\rangle|1\rangle, \quad |1\rangle|0\rangle. \tag{9}$$

This leads directly to a representation of the SWAP operator **S** that takes the 2-Cbit state $|x\rangle|y\rangle$ to $|y\rangle|x\rangle$:

$$\mathbf{S} = \tfrac{1}{2}(\mathbf{1} + \mathbf{Z} \otimes \mathbf{Z}) + (\mathbf{X} \otimes \mathbf{X})\,\tfrac{1}{2}(\mathbf{1} - \mathbf{Z} \otimes \mathbf{Z})$$

$$= \tfrac{1}{2}(\mathbf{1} + \mathbf{Z} \otimes \mathbf{Z} + \mathbf{X} \otimes \mathbf{X} - \mathbf{Y} \otimes \mathbf{Y}), \quad \mathbf{Y} = \mathbf{XZ}. \tag{10}$$

Pause also to admire the simplicity of this classical derivation of the form of the quantum-mechanical exchange operator, compared with the standard derivation based on the full-blown quantum theory of angular momentum technology. Note also the further simplicity introduced into (10) by incorporating an additional factor of i into the definition of **Y** (which also makes it hermitian, like **X** and **Z**). With examples like this, one can motivate the utility of extending the notion of states to include multiplication by complex scalars, leading to the generalization from Cbits to *Qbits*.

Qbits are physical systems characterized by states which fully exploit the entire 2^n dimensional complex vector space spanned by the 2^n orthonormal Cbit states. Nature has been kind enough to present us with many examples of them. The general state $|\Psi\rangle$ of n Qbits is *any* unit vector:

$$|\Psi\rangle = \sum_{0\le x<2^n} a_x|x\rangle_n, \qquad \sum_{0\le x<2^n} |a_x|^2 = 1. \tag{11}$$

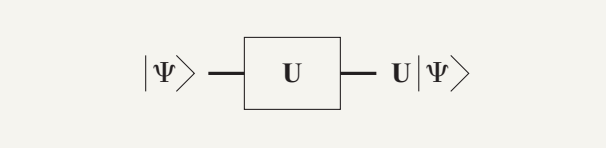

Figure 8

Circuit diagram representing a unitary gate **U** acting on n Qbits.

With one (extremely important) exception, all operations on Qbits are reversible. Since the exception ("measurement") has no nontrivial analog for Cbits, in comparing Qbits and Cbits it suffices to consider only reversible operations on Cbits. The only reversible operations on the 2^n Cbit states are their $(2^n)!$ possible permutations. But the general operation nature allows us to perform on n-Qbit states is *any* linear norm-preserving transformation,

$$|\Psi\rangle \to \mathbf{U}|\Psi\rangle, \qquad \mathbf{U} \text{ unitary}, \tag{12}$$

as shown schematically in **Figure 8**.

While Qbits are far more versatile than Cbits in their range of states and the operations one can perform on them, the usefulness of their versatility is highly constrained by one important difference between Qbits and Cbits. Learning the state $|x\rangle_n$ of n Cbits is unproblematic: One just looks to see which of the 2^n possible states $|x\rangle_n$ it is. In contrast, learning the state $|\Psi\rangle_n = \Sigma_x\, a_x|x\rangle_n$ associated with n Qbits is impossible. Given the Qbits, there is nothing one can do to them to reveal their state.

To extract any information from Qbits, one must "make a measurement." This consists of sending the Qbits through a "measurement gate." If the state of the Qbits is (11), then the measurement gate signals x with probability $p = |a_x|^2$. After x is signaled, the state associated with the Qbits must be taken to be $|x\rangle_n$. The manner in which an n-Qbit measurement gate operates and the fact that it represents the only way to extract information from the n-Qbits constitute the *Born rule*. The Born rule is illustrated in **Figure 9**.

As defined above, "measurement" is always in the computational basis. One loses nothing by this simplification, since measurement in any other basis can be described as measurement in the computational basis,

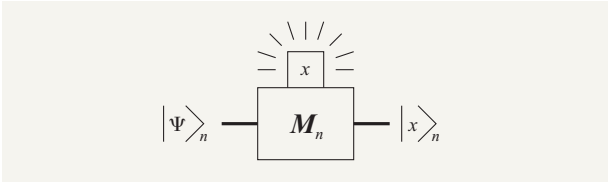

Figure 9

An n-Qbit measurement gate, acting as specified by the Born rule.

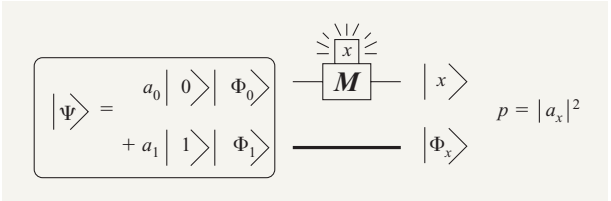

Figure 10

Action of a 1-Qbit measurement gate on a single one of n Qbits, according to the generalized Born rule.

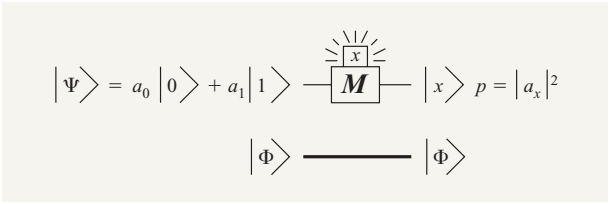

Figure 11

Simplification of Figure 10 when $|\Phi_0\rangle = |\Phi_1\rangle = |\Phi\rangle$.

preceded by an appropriate unitary transformation. What one gains, pedagogically, is the need to invoke only a single variety of measurement gate and, as noted below, only 1-Qbit measurement gates.

A somewhat stronger version of the Born rule plays a crucial role in quantum computation, though it is rarely explicitly mentioned in most standard quantum mechanics texts. The stronger form applies when one measures only a single one of n Qbits. The state $|\Psi\rangle$ of all n Qbits can always be represented in the form

$$|\Psi\rangle = a_0|0\rangle|\Phi_0\rangle + a_1|1\rangle|\Phi_1\rangle, \qquad |a_0|^2 + |a_1|^2 = 1, \tag{13}$$

where the Qbit to be measured appears on the left and where $|\Phi_0\rangle$ and $|\Phi_1\rangle$ are normalized but not necessarily orthogonal states of the $n - 1$ unmeasured Qbits. The generalized Born rule asserts that if only the single Qbit is measured, then the 1-Qbit measurement gate will indicate x (0 or 1) with probability $|a_x|^2$, after which the n-Qbit state will be the product state $|x\rangle|\Phi_x\rangle$, as illustrated in **Figure 10**.

To see that the gate acting on the measured Qbit in Figure 10 is indeed the $n = 1$ version of the n-Qbit measurement gate of Figure 9, note that Figure 10 simplifies to **Figure 11** when $|\Phi_0\rangle = |\Phi_1\rangle = |\Phi\rangle$. In this special case, the entangled input state in Figure 10 becomes an uncorrelated product in which both the measured Qbit and the remaining $n - 1$ Qbits have states of their own. The $(n - 1)$ unmeasured Qbits now take no part whatever in the process. Nothing acts on them and they do nothing but maintain their original state $|\Phi\rangle$. Their presence is irrelevant to the upper part of the figure, which is nothing more than the $n = 1$ version of Figure 10.

It is an elementary consequence of the generalized Born rule that the n-Qbit measurement gate of the ordinary Born rule can be constructed from n 1-Qbit measurement gates, as illustrated in **Figure 12**.

Although the generalized Born rule is stronger, it follows from the Born rule under two plausible assumptions: a) Once a Qbit ceases to interact with others and ceases to be acted on by unitary gates, it does not matter when it is measured. b) To assign a state to Qbits is to do nothing more than to specify the probabilities of subsequent measurement outcomes, possibly preceded by unitary gates. Since the generalized Born rule reduces the notion of measurement to a single black-boxed 1-Qbit measurement gate (and indeed, since Qbits can be measured one by one, one needs only a single specimen of such a measurement gate), by far the most economical introduction to quantum computation is to base it on a primitive concept of the 1-Qbit measurement gate and make explicit assumptions a) and b) above.

Since it is impossible to determine the state of a collection of Qbits from the Qbits themselves, how is one to associate with the registers of a quantum computer the initial states on which the unitary transformations subsequently act? The obvious, simplest, and conceptually most economical answer is to exploit the measurement gates themselves. One can initialize n Qbits to the state $|0\rangle_n$ by measuring each Qbit and applying **X** if and only if the measurement indicates 1, as illustrated in **Figure 13**.

Note the following features of this pedagogically motivated approach to quantum mechanics:

1. It relies on an irreducible primitive notion of a unique black-boxed 1-Qbit measurement gate. The measurement gate is the only irreversible circuit

element. It is defined by what it does, and what it does is to extract information from the Qbits in a form that is immediately accessible to *us*. There is no other way for us to obtain such information from the Qbits.[3]

2. Measurement plays a dual role. Output on a readable (classical) display not only ends the computation; it also provides, without any further complication, a straightforward way to *begin* the computation. Initialize every Qbit to the 1-Qbit state $|0\rangle$ by sending each through a measurement gate, and then do nothing or apply **X**, depending on whether the display shows 0 or 1.
3. Unlike the state of n Cbits, the state of n Qbits does not reside in the Qbits themselves: Presented with a bunch of Qbits, there is nothing one can do to them to reveal their state. Indeed, in general—for example if they share with other Qbits an entangled state—Qbits will not have a state of their own at all. To determine the state of Qbits (or whether they have one) one must ask Alice, who knows their history: what initial measurements were performed on them, what the outcomes of the initial measurements were, and what subsequent unitary gates were applied.
4. While the *purpose* of the state of n Cbits is an anthropocentric add-on to its intrinsic character (what, for example, is the purpose of the velocity of a classical particle?), one would not bother to keep track of the state of n Qbits were it not that this information about their past history has a specific purpose: It enables us to determine the correlations between the initial and final measurement outcomes after any intermediate sequence of applications of unitary gates.

All of these features resonate strongly with the constellation of ideas known as the Copenhagen interpretation. The quantum state of a system is not an objective property of that system. It merely provides an algorithm enabling one to infer from the initial set of measurements and outcomes ("state preparation") the probabilities of the results of a final set of measurements after a specified intermediate time evolution. We ourselves have direct access to nothing beyond the outcomes of such measurements.

Why did my bare-bones, no-nonsense, pedagogically motivated, minimalist introduction to quantum mechanics come out sounding so Copenhagen? I think there are several reasons:

a. Proponents of the Copenhagen interpretation (notably Heisenberg and Peierls) have always maintained that the quantum-mechanical formalism does not describe "the system" but "our knowledge of the system."

[3] The art of quantum computation, of course, is to construct unitary transformations leading to final states in which only informative values of x are associated with appreciable probabilities $|a_x|^2$.

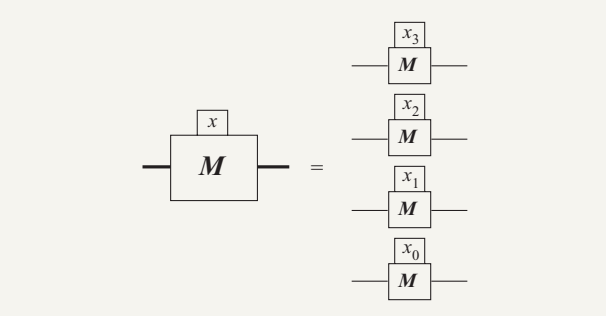

Figure 12

Construction of a 4-Qbit measurement gate from four 1-Qbit measurement gates. The integer x has the binary expansion $x_3x_2x_1x_0$.

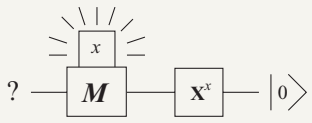

Figure 13

Use of a 1-Qbit measurement gate to prepare an off-the-shelf Qbit so that its associated state is $|0\rangle$.

Quantum computation is the first application of quantum mechanics that does not use it to further our understanding of the physical world. On the contrary, quantum computation exploits the known quantum-mechanical character of the physical world to expedite the processing of knowledge, as represented symbolically by constituents (Qbits) of that world. It is therefore not surprising that the Copenhagen interpretation should provide a congenial setting for the exposition of quantum computation.

b. A computation uses a finite set of Qbits. It has an unambiguous beginning and end. There is always a world external to the computation. If there were not an outside world, there would be no point in doing the computation because there would be nobody or nothing to take advantage of the output. Nobody (well, at this stage practically nobody) wants to view the entire universe (single or multiple) as one colossal quantum computer, sufficient unto itself. The Copenhagen interpretation is characterized by a similar modesty of scope. Physics is a tool for relating some aspects of our experience to other aspects. Every application of physics begins and ends with an appeal to experience.

c. The pedagogical device of restricting "measurement" to measurement in the computational basis, treating measurement in other bases as computational-basis measurement preceded by an appropriate unitary transformation, resonates with the Copenhagen notion of the primacy of the classical world. The computational basis states are (actually, one should replace "are" with "are isomorphic to") the states that describe ordinary classical Cbits. By restricting "measurement" to the computational basis, I have automatically arranged for the input and output of every quantum computation to be describable in the ordinary old-fashioned language of classical computation—numbers flashed on a classical display. Bohr always insisted that our knowledge of the world had to be formulated in ordinary language, or we could not communicate it to anybody else.
d. My computer science students know very little physics. They are therefore immune to any temptation to reify the states of Qbits into properties of the associated physical systems. If you think you too are immune from such temptation, ask yourself whether you do or do not believe that a horizontally polarized photon is *intrinsically* different from a vertically polarized photon. If you do, you are a victim of that very temptation.

I conclude by translating the possibly obscure quotation at the head of this section into the quite straightforward form it assumes in the context of quantum computation:

In our description of a quantum computation the purpose is not to disclose the real essence of the Qbits but only to track down statistical relations between initial and final measurement outcomes.

Acknowledgment

This work was supported by the National Science Foundation under Grant No. 0098429.

References and note

1. D. DiVincenzo, P. Hayden, and B. Terhal, "Hiding Quantum Data"; see *http://arxiv.org/quant-ph/0207147/*.
2. C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, "Teleporting an Unknown Quantum State via Dual Classical and Einstein–Podolsky–Rosen Channels," *Phys. Rev. Lett.* **70,** 1895–1899 (1993).
3. C. A. Fuchs, "Notes on a Paulian Idea: Foundational, Historical, Anecdotal and Forward-Looking Thoughts on the Quantum"; see *http://arxiv.org/quant-ph/0105039/*.
4. C. H. Bennett and S. J. Wiesner, "Communication via One- and Two-Particle Operators on Einstein–Podolsky–Rosen States," *Phys. Rev. Lett.* **69,** 2881–2884 (1992).
5. N. D. Mermin, "Deconstructing Dense Coding" *Phys. Rev. A* **66,** 032308 (2002).
6. N. Bohr, *Collected Works*, Vol. 6, J. Kalckar, Ed., North-Holland Publishing Co., Amsterdam, 1985, p. 296.
7. The pedagogical approach to quantum mechanics summarized here is described in more detail in N. D. Mermin, "From Cbits to Qbits: Teaching Computer Scientists Quantum Mechanics," *Amer. J. Phys.* **71,** 23–30 (2003). Lecture notes for the course itself are regularly updated, revised, and reposted at *http://www.ccmr.cornell.edu/~mermin/qcomp/CS483.html*.

Received May 19, 2003; accepted for publication May 27, 2003

N. David Mermin *Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853 (ndm4@cornell.edu).* Dr. Mermin is Horace White Professor of Physics in the Cornell University Physics Department. Harvard University awarded him an A.B. degree in mathematics in 1956 and a Ph.D. in physics in 1961. He is the author of "Space and Time in Special Relativity," "Solid State Physics" (with Neil W. Ashcroft), and "Boojums All the Way Through." Over the years he has done theoretical research in solid-state physics, statistical physics, low-temperature physics, crystallography, and the foundations of quantum mechanics. From 1984 to 1990, Dr. Mermin was Director of the Laboratory of Atomic and Solid State Physics at Cornell, and since 2000 has been a founding member of the Cornell Faculty of Computation and Information. He is a member of the American Academy of Arts and Sciences and the U.S. National Academy of Sciences.