A Quantum Jump in Computer Science

Gilles BRASSARD*

Université de Montréal, Département d'informatique et de recherche opérationnelle
C. P. 6128, Succursale Centre–Ville, Montréal (Québec), CANADA H3C 3J7
email: brassard@iro.umontreal.ca

Abstract. Classical and quantum information are very different. Together they can perform feats that neither could achieve alone. These include quantum computing, quantum cryptography and quantum teleportation. This paper surveys some of the most striking new applications of quantum mechanics to computer science. Some of these applications are still theoretical but others have been implemented.

1 Introduction: the Qubit

Classical and quantum information are very different. Classical information can be read, copied and transcribed into any medium; it can be transmitted and broadcast, but it cannot travel faster than the speed of light. Quantum information cannot be read or copied without disturbance, but in some instances it appears to propagate instantaneously or even backward in time. In principle—and sometimes also in practice—the two kinds of information can perform together several feats that neither could achieve alone. These include quantum computing, quantum cryptography and quantum teleportation.

Quantum mechanics has the potential to bring about a spectacular revolution in computer science. Even though current-day computers use quantum-mechanical effects in their operation, for example through the use of transistors, they are still very much classical computing devices. A supercomputer is not fundamentally different from a purely mechanical computer that could be built around simple relays: their operation can be described in terms of classical physics and they can simulate one another in a straightforward manner, given sufficient storage. All deterministic classical single-processor models of computation are computationally equivalent to the classical Turing machine up to polynomial factors. In contrast, computers could in principle be built to profit from genuine quantum phenomena that have no classical analogue, sometimes providing exponential speed-up compared to classical computers even if the classical computer is allowed to be probabilistic and to make occasional mistakes. Quantum information is also at the core of other phenomena that would be impossible to achieve in a purely classical world, such as unconditionally secure distribution of secret cryptographic material.

* This essay was written while the author was on sabbatical at the University of Wollongong, Australia; Research supported in part by NSERC and FCAR.

At the heart of it all is the quantum bit, or *qubit*, a word coined by Schumacher. A classical bit can take either value 0 or value 1. Nothing in between is allowed. According to quantum mechanics, a qubit can be in linear *superposition* of the two classical states. If we denote the classical states by $|0\rangle$ and $|1\rangle$, then a qubit can be in state

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

for arbitrary complex numbers α and β subject to $|\alpha|^2 + |\beta|^2 = 1$. The coefficients α and β are called the *amplitudes* of $|0\rangle$ and $|1\rangle$, respectively. A qubit is best visualized as a point on the surface of a unit sphere whose North and South poles correspond to the classical values. This is not at all the same as taking a value *between* 0 and 1 as in classical analogue computing.

In general, qubits cannot be measured reliably. The simplest measurement on $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$ consists of asking it to choose between the two classical states. In this case, the qubit *collapses* onto $|0\rangle$ with probability $|\alpha|^2$ and onto $|1\rangle$ with complementary probability $|\beta|^2$, and all additional information about the original superposition is lost irreversibly. Other measurements are also possible, but the more information you obtain about $|\Psi\rangle$, the more you risk spoiling it. Even though there is a continuum of possible quantum states of a single qubit, these states cannot all be distinguished reliably from each other: no measurement can extract more than one expected bit of information from any given qubit and even this much is not always possible. It is this manifestation of Heisenberg's uncertainty principle that makes quantum computing challenging but it also makes quantum cryptography possible.

2 Two Qubits Are More Than Twice Better Than One

We begin with a simple example that illustrates the difference between classical and quantum information. A *trine* is a qubit that can be either in state $|0\rangle$, $\frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}|1\rangle$ or $\frac{1}{2}|0\rangle - \frac{\sqrt{3}}{2}|1\rangle$, each possibility being equally likely [1]. From the point of view of classical information theory, a trine encodes $\log_2 3 \approx 1.585$ bits of information. However, the best possible measurement on a trine can do no better than rule out one of the three possibilities, leaving one bit of uncertainty about the original state of the qubit. Thus, it is not possible to extract more than $(\log_2 3) - 1 \approx 0.585$ bit of classical information from any given trine. Suppose now that you are presented with two trines that are guaranteed to be identical. Peres and Wootters have discovered that there is a measurement that can extract $\frac{\sqrt{2}}{3}\left(1 + \log_2(17 + 12\sqrt{2})\right) - \frac{3}{2} \approx 1.369 > 2 \times 0.585$ bits of information from both trines combined [31]. Thus it is possible in some cases to extract more than twice as much information from two *identical* qubits than from either one alone! This has no counterpart in classical information theory.

[1] If the qubit is implemented by the polarization of a photon such that $|0\rangle$ and $|1\rangle$ are encoded by vertical and horizontal polarizations, respectively, the three trine states correspond to a polarization that is either vertical, $+30°$ or $-30°$ from the horizontal.

3 Quantum Computing

Can computers simulate physics? This question was the topic of Feynman's keynote address at the 1981 Workshop on Physics of Computation [23]. This would hardly have been an original question had he not asked more specifically if physics can be simulated *efficiently* by *classical* computers. He gave reasons to believe the answer is negative. Today, we know that Feynman's intuition was correct—provided factoring is genuinely difficult for classical computers. Specifically, Shor [32], building on work of Deutsch and Jozsa [17], Bernstein and Vazirani [9] and Simon [33], has discovered how to factor arbitrarily large integers in polynomial time on a quantum computer. Thus, any efficient simulation of physics on a classical computer would yield an efficient classical algorithm for factoring. See also Berthiaume and Brassard [10] for early oracle results.

The concept of quantum computing evolved from the work of Benioff [2], Feynman [23, 24] and Deutsch [16] in the first half of the 1980s. In a classical probabilistic calculation, it is possible to program the computer to select one of several possible computation paths according to the laws of probability. In any specific instance of the calculation, however, only one of the potential paths is actually taken, and what-could-have-happened-but-did-not has no influence whatsoever on the actual outcome. In a quantum computer, all potential computation paths are taken simultaneously. More precisely, just as a qubit can be in linear superposition of states $|0\rangle$ and $|1\rangle$, a quantum computer can be in linear superposition of having followed several different alternatives simultaneously. What makes quantum computation so powerful—and mysterious—is that the possible paths that would yield the same result *interfere*, a notion that has no analogue in classical computation. This interference can be constructive or destructive, much as in Young's celebrated double-slit experiment. This allows for the reinforcement of the probability of obtaining desired results while at the same time reducing or even annihilating the probability of spurious results. This happens because the amplitude of reaching any given state is the sum of the amplitudes of reaching this state by all possible computation paths. Because amplitudes can be negative (even complex), the amplitude of reaching an undesirable result from one path can be cancelled by the amplitude of reaching it from another path. In the words of Feynman, "somehow or other it appears as if the probabilities would have to go negative" [23].

As a simple illustration of the phenomenon of quantum interference, consider a one-qubit quantum register whose possible classical states are $|0\rangle$ and $|1\rangle$. We have seen that it can be in any state of the form $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where $|\alpha|^2 + |\beta|^2 = 1$. Consider the quantum algorithm given in Fig. 1. Note that this algorithm does *not* say that we begin by measuring state $|\Psi\rangle$ to decide classically which path to follow: we want the program to exhibit interesting behaviour when state $|\Psi\rangle$ is already in quantum superposition before the call.

If we start in state $|\Psi\rangle = |0\rangle$, a call on QCF produces the quantum superposition of $|0\rangle$ and $|1\rangle$ with amplitudes $+\sqrt{2}/2$ and $-\sqrt{2}/2$, respectively. In other words, the result is $|\Psi\rangle = \frac{\sqrt{2}}{2}(|0\rangle - |1\rangle)$. If we start in state $|\Psi\rangle = |1\rangle$, on the other

procedure QCF

$$\textbf{if } |\Psi\rangle = |0\rangle \textbf{ then } |\Psi\rangle \leftarrow \begin{cases} |0\rangle \text{ with amplitude } +\sqrt{2}/2 \\ |1\rangle \text{ with amplitude } -\sqrt{2}/2 \end{cases}$$

$$\textbf{if } |\Psi\rangle = |1\rangle \textbf{ then } |\Psi\rangle \leftarrow \begin{cases} |0\rangle \text{ with amplitude } +\sqrt{2}/2 \\ |1\rangle \text{ with amplitude } +\sqrt{2}/2 \end{cases}$$

Fig. 1. Quantum coin flipping algorithm

hand, a call on QCF results in $|\Psi\rangle = \frac{\sqrt{2}}{2}(|0\rangle + |1\rangle)$. In either case, an immediate measurement of $|\Psi\rangle$ would yield either $|0\rangle$ or $|1\rangle$ with probability 50% since $|\sqrt{2}/2|^2 = |-\sqrt{2}/2|^2 = {}^1/_2$. In other words, if we start the computer in either classical state, call QCF once and observe the result, we have implemented a fair coin toss. This is at best useful for classical probabilistic computing.

Consider now what happens if we call QCF twice *without observing the result between calls*. Say we start in state $|\Psi\rangle = |1\rangle$. The first call produces $|\Psi\rangle = \frac{\sqrt{2}}{2}(|0\rangle + |1\rangle)$, as we have just seen. The second call transforms the $|0\rangle$ component of this state into $\frac{\sqrt{2}}{2}(|0\rangle - |1\rangle)$ and the $|1\rangle$ component into $\frac{\sqrt{2}}{2}(|0\rangle + |1\rangle)$, so that the net result is

$$\begin{aligned} |\Psi\rangle &= \tfrac{\sqrt{2}}{2}\left(\tfrac{\sqrt{2}}{2}(|0\rangle - |1\rangle) + \tfrac{\sqrt{2}}{2}(|0\rangle + |1\rangle)\right) \\ &= \tfrac{1}{2}(|0\rangle - |1\rangle + |0\rangle + |1\rangle) \;=\; |0\rangle. \end{aligned}$$

Similarly, if we start with $|\Psi\rangle = |0\rangle$ and apply QCF twice, we end up in state $|\Psi\rangle = -|1\rangle$. Thus, if we observe the result of two successive calls on QCF starting with the computer in a classical state, we find that it has switched to the logical negation of its original state. Because the first call on QCF completely randomizes either classical input if the output is observed, yet the second call produces a deterministic result that depends on the initial input if no observation is made between calls, Hayes suggested the following analogy: "It is as if we had invented a machine that first scrambles eggs and then unscrambles them".[2] It is slightly more challenging to find a quantum algorithm SQNOT whose effect, when composed with itself, is to transform $|1\rangle$ into $|0\rangle$ and $|0\rangle$ into $|1\rangle$ (rather than $-|1\rangle$), but this can be done: it is known as the *square root of not*. Try it!

To see more clearly the quantum interference at work, consider again what happens if we start in state $|1\rangle$ and make two calls on QCF. Thinking classically, there are four possible computation paths of the form $|a\rangle \to |b\rangle \to |c\rangle$, where

[2] The magic disappears if we implement quantum bits of memory with the polarization of single photons using an arbitrary coding scheme such as representing $|0\rangle$ with horizontally polarized photons and $|1\rangle$ with vertically polarized photons: in this case, algorithm QCF performs a simple polarization rotation by 45 degrees. But don't worry: more subtle uses of quantum interference, such as those used for quantum factorization, cannot be explained away so easily with classical thinking.

Fig. 2. The computation tree of two calls on QCF starting with $|\Psi\rangle = |1\rangle$

$|a\rangle = |1\rangle$ is the classical state before the first call, $|b\rangle$ is the classical state between the two calls and $|c\rangle$ is the classical state after the second call: $|1\rangle \to |0\rangle \to |0\rangle$, $|1\rangle \to |0\rangle \to |1\rangle$, $|1\rangle \to |1\rangle \to |0\rangle$ and $|1\rangle \to |1\rangle \to |1\rangle$. See Fig. 2. The amplitude of path $|1\rangle \to |0\rangle \to |1\rangle$, for example, is $\frac{\sqrt{2}}{2} \times (-\frac{\sqrt{2}}{2}) = -1/2$ because the amplitude of going from $|1\rangle$ to $|0\rangle$ is $\sqrt{2}/2$ and the amplitude of going back from $|0\rangle$ to $|1\rangle$ is $-\sqrt{2}/2$. (The amplitude of any given path is the product of the amplitudes along that path, much like probabilities multiply along each path in probabilistic algorithms.) Similarly, the amplitude of the three other paths is $+1/2$.

If there were no interference, the probability of following any given path would be the square of the magnitude of its amplitude, which in this case is $|\pm 1/2|^2 = 1/4$ for each path. Thus each path would be followed with equal probability and the end result would be equally likely to be $|0\rangle$ or $|1\rangle$. But what really happens is that the two paths ending in $|0\rangle$ interfere constructively: their amplitudes add up to produce $1/2 + 1/2 = 1$. By contrast, the two paths ending in $|1\rangle$ interfere destructively: their amplitudes add up to produce $-1/2 + 1/2 = 0$. The net result is that $|0\rangle$ is observed with certainty. There were two paths leading to $|1\rangle$ and each would have nonzero probability of being taken if we thought classically, yet they cancel each other out and $|1\rangle$ is never observed.

Although the above example is exceedingly simple, it illustrates the basic principle behind the design of quantum algorithms. The trick is to program the computer so that the computation paths that yield undesirable results interfere destructively whereas the "good" computation paths interfere constructively. Of course, interesting algorithms for the quantum computer are not restricted to working on a single qubit of memory. We saw already that two identical trines are more than twice as useful as a single trine. Similarly, a quantum register made out of n qubits can be exponentially more powerful than a single qubit. This is because the qubits in a quantum register do not have to be independent from

one another: they can be *entangled*. In general, an n–qubit quantum register can be in an arbitrary linear superposition of all 2^n classical states:

$$|\Psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$$

with the restriction that $\sum_{x \in \{0,1\}^n} |\alpha_x|^2 = 1$. If this quantum register is measured in the sense that it is asked to assume a classical value, the formula is much the same as with a single qubit: the register collapses onto some classical n–bit string x, and each possible such x is obtained with probability $|\alpha_x|^2$.

To understand the notion of entanglement, consider a two-qubit quantum register. Suppose at first that each qubit is independently in state $\frac{\sqrt{2}}{2}|0\rangle + \frac{\sqrt{2}}{2}|1\rangle$. The register is then in a state $|\Psi\rangle$ that can be represented by the *product* of two such states. By this we mean that amplitudes are multiplied but the product of classical states denotes their concatenation: $|0\rangle|1\rangle$ is the same as $|01\rangle$, a classical state for a two-qubit quantum register. Continuing our example,

$$\begin{aligned} |\Psi\rangle &= \left(\tfrac{\sqrt{2}}{2}|0\rangle + \tfrac{\sqrt{2}}{2}|1\rangle\right)\left(\tfrac{\sqrt{2}}{2}|0\rangle + \tfrac{\sqrt{2}}{2}|1\rangle\right) \\ &= \tfrac{1}{2}(\,|00\rangle + |01\rangle + |10\rangle + |11\rangle)\,. \end{aligned}$$

This is not a very interesting state. If it is asked to assume a classical value, each of the four possible classical two-bit strings is obtained with equal probability $1/4 = |1/2|^2$. This is exactly the same as saying that each of the two qubits, when measured, yields an independent random bit: the two qubits are not entangled.

Consider now a more interesting state for our quantum register.

$$|\widehat{\Psi}\rangle = \tfrac{\sqrt{2}}{2}(\,|00\rangle + |11\rangle) \tag{1}$$

This time, if the register is asked to assume a classical value, we observe either $|00\rangle$ or $|11\rangle$, each with probability $|\sqrt{2}/2|^2 = 1/2$, but we never observe $|01\rangle$ nor $|10\rangle$. The two qubits are entangled: the register cannot be represented as a product of two independent qubits. To see why this is remarkable, consider what happens if we perform a *partial* measurement. Let us say that we measure only the first qubit of the register. The outcome is equally likely to be $|0\rangle$ or $|1\rangle$. Similarly, the second qubit would have been equally likely to yield $|0\rangle$ or $|1\rangle$ had we measured *it* instead. But now that we have measured the first qubit, the outcome of a subsequent measurement on the second qubit is determined: it will yield the same value as the first. In other words, measuring one qubit can have an instantaneous effect on the state of another. This is as if we had a pair of *magic coins*: the outcome of either coin is not determined until it *or its twin* has been flipped. This works no matter how far apart the two qubits (or coins) are, which troubled Einstein very much indeed when he discovered a similar quantum phenomenon with Podolsky and Rosen 60 years ago [22]: he later called it "spooky action at a distance" [3] (*spukhafte Fernwirkungen*) in a letter to Born. Such states were subsequently observed in the laboratory.

[3] Although qubits appear to communicate instantaneously with one another, this phenomenon cannot be used to send a signal faster than light.

The relevance of entangled states to computing is best illustrated by the following example, which is due to Vazirani and inspired by Simon. Consider a function $h : \{0,1\}^n \to \{0,1\}^m$ and a quantum computer equipped with $n+m$ qubits of quantum register. Program your quantum computer to transform $|x, \mathbf{0}\rangle$ into $|x, h(x)\rangle$, where $x \in \{0,1\}^n$ and $\mathbf{0}$ denotes the all-zero m–bit vector. Now, prepare your quantum register in the following initial state

$$|\Psi_0\rangle = \sum_{x \in \{0,1\}^n} 2^{-n/2} \, |x, \mathbf{0}\rangle \, .$$

If you run your quantum computer on that state, the resulting state will be

$$|\Psi_1\rangle = \sum_{x \in \{0,1\}^n} 2^{-n/2} \, |x, h(x)\rangle \, .$$

In general, this is an entangled state that cannot be expressed as a product of the states of $n+m$ qubits. The first thing to notice is that you have computed the value of $h(x)$ for exponentially many different values of x in the time it would have taken to compute it on a single input with a classical computer. This phenomenon, known as *quantum parallelism*, is at the core of all speed-ups offered by quantum computing. Unfortunately, quantum mechanics severely restricts the information you can extract from $|\Psi_1\rangle$. In particular, it would be pointless to measure it in the way we have seen so far: if you ask $|\Psi_1\rangle$ to assume a classical value, it will simply collapse onto a pair $|x, h(x)\rangle$ for a randomly chosen $x \in \{0,1\}^n$ and you will be none the wiser since you could just as easily have obtained this with a classical computer. Instead, measure only the last m qubits of the register. As they collapse to a random value y in the image of h (with probabilities that depend on the number of preimages), entanglement causes the first n qubits to collapse to an equally-weighted superposition of all preimages of y under h. The resulting state is

$$|\Psi_2\rangle = \sum_{x \in h^{-1}[y]} \frac{1}{\sqrt{\#h^{-1}[y]}} \, |x, y\rangle \, .$$

For the cryptographically-minded, this is a way to generate collisions[4] out of an arbitrary hashing function h. But don't get too excited because, once again, the collision thus obtained cannot be extracted in general from the state $|\Psi_2\rangle$ of the quantum register. If you measure the first n qubits of your register now, you obtain a random preimage x of the random image y you had previously observed. This is pointless because you could produce the same result by first choosing x randomly and then computing $y = f(x)$. Thus, state $|\Psi_2\rangle$ encodes precious information, but you cannot get at it by a direct measurement. How tantalizing!

This is where quantum *interference* enters the game. Under some conditions on the type of function h considered, Simon [33] has discovered how to massage state $|\Psi_2\rangle$ in a way that interference patterns allow us subsequently to measure useful information that efficiently yields a collision with high probability.

[4] A *collision* is a pair of distinct values $x_1, x_2 \in \{0,1\}^n$ such that $h(x_1) = h(x_2)$.

Although his condition on h is somewhat unnatural, Simon's result is of great theoretical importance because he has also proved (assuming h is given by a black box) that no classical computer could find collisions on such functions in less than exponential time with better than an exponentially small probability of success, even if the classical algorithm is probabilistic and not required to return a correct answer with certainty. This was the first convincing argument that quantum computers can be exponentially faster than classical probabilistic computers. (It was known previously that quantum computers can be exponentially faster than classical deterministic computers [17] and that they can be somewhat faster than classical probabilistic computers [9].)

Shor's algorithm for the efficient factorization of large integers on the quantum computer follows a similar line of attack, but with some Fourier transforms and some number theory thrown in [32]. The details are beyond the scope of this essay. Note however that Ekert and Jozsa have written a particularly lucid account of Shor's algorithm [21]. You should also read Coppersmith's algorithm for approximating Fourier transforms on the quantum computer [15].

Before you run off and start writing your own quantum algorithms, you must be warned about a crucial restriction on such algorithms that is imposed by the laws of quantum mechanics. Clearly we must require that if classical state $|x\rangle$ leads to superposition $\sum \alpha_i |x_i\rangle$ then $\sum |\alpha_i|^2 = 1$ so that the probabilities sum up to 1 if an observation is made. For example, $|\sqrt{2}/2|^2 + |\pm\sqrt{2}/2|^2 = 1$ in algorithm QCF. By analogy with classical probabilistic algorithms, you may think this is the only restriction on quantum algorithms. Wrong! Quantum mechanics requires the operation of the computer to be *unitary*. This means that if X is the set of all possible classical states of the computer so that a general quantum state $|\Psi\rangle$ is represented by a column vector of amplitudes $|\Psi\rangle = (\alpha_x)_{x\in X}$, and if matrix U describes the program in the sense that the state that follows $|\Psi\rangle$ after one step is $U|\Psi\rangle$, then quantum mechanics requires that the inverse of U be its own conjugate transpose.

In the case of a one-bit computer, we could represent $|0\rangle$ as $\binom{1}{0}$ and $|1\rangle$ as $\binom{0}{1}$. In general, $\alpha|0\rangle + \beta|1\rangle$ is represented as $\binom{\alpha}{\beta}$. With this notation, consider

$$U = \frac{\sqrt{2}}{2}\begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \text{ and } \widehat{U} = \frac{1}{2}\begin{bmatrix} 1-i & 1+i \\ 1+i & 1-i \end{bmatrix}.$$

The reader is invited to verify that both these matrices are unitary, that U corresponds to algorithm QCF from Fig. 1 and that $\widehat{U}$ solves the square root of not (SQNOT) riddle: $\widehat{U}^2|b\rangle = |\bar{b}\rangle$ for $b \in \{0,1\}$.

Unitarity is not an easy constraint to work with. In particular, unitary operations must be reversible, which forces us to rethink the way we program. Fortunately, it has been proven by Bennett [3] that reversibility is not a serious handicap when it comes to efficient classical computation, so that polynomial-time quantum computers are at least as powerful as their classical counterpart. As we have seen, there are good reasons to believe that they are in fact significantly more powerful.

This is all very well in theory, but how realistic are quantum computers from a technological point of view? Will they forever remain the stuff of dreams? How about a special-purpose quantum factorization device? Quantum computing offers formidable technological challenges. The most serious of those concerns the need to keep quantum *coherence* for a macroscopic amount of time. The problem is that coherent superpositions are very fragile: they spontaneously turn into incoherent statistical mixtures—which are unsuitable for quantum parallelism—because of residual interactions with the environment. Consequently, quantum information disappears gradually from the computer. Moreover, according to Unruh, the coherence time decreases exponentially with the number of qubits involved [34], which means that the larger the problem you want to solve, the less time you have to tackle it! Landauer has pointed out the additional problem that quantum computation may be susceptible to spontaneous reversals, which would unpredictably cause the computation to go backwards [27]. (Landauer's argument applies to a "pure" quantum computer of the kind advocated by Deutsch [16] but it may not be a serious consideration if a classical computer is used to "drive" a quantum register, which would be sufficient to implement Shor's quantum factoring algorithm.) Yet another difficulty is that of error correction despite early work by Peres [30] and more recent work by Berthiaume, Deutsch and Jozsa [11]. In classical digital computing, discrete states are continuously snapped back onto the proper voltages for 0 and 1, but no similar process is available for quantum computing even if the intended states have discrete amplitudes. Consequently, the computer may slowly drift away from its intended state.

Nevertheless, the prospect for experimental quantum computing became less remote when DiVincenzo discovered that universal quantum computers can be built of two-input quantum gates [18], which is likely to greatly simplify the technology required to build them. (It is well-known that two-input gates are universal for classical computation but the same was not at all obvious for quantum computation because we saw that it has to be unitary—and therefore reversible—and universal classical reversible gates need at least three inputs.) It was later discovered that the natural quantum extension of the classical exclusive-or gate (also known as "controlled-not") is universal for quantum computation, if supplemented by simple operations on single qubits. Not only are theoreticians such as Lloyd contemplating designs for "potentially realizable quantum computers" [28], but several experimental physicists are currently attempting to realize basic quantum computing gates. Domokos, Raimond, Brune and Haroche are working on a prototype of the quantum exclusive-or gate based on atomic interferometry and microwave cavities capable of trapping single photons [19]. Another promising approach, proposed by Cirac and Zoller and being investigated experimentally by Wineland and others, makes use of trapped ions. See also [1, 21] for implementation proposals.

4 Time For Another Paradigm Shift

In his seminal 1936 paper, Turing insisted that any reasonable model of computation should be physically realizable and he argued that such was the case with what was to become known as the Turing machine. In the early days of computational complexity, it was suggested that a problem should be considered *tractable* if and only if it can be solved in polynomial time on a deterministic Turing machine. This obeyed Turing's paradigm because Turing machines can be implemented in the real world. Moreover, the notion seemed robust because it does not depend on the specifics of the Turing machine: it is invariant on the type of computer used, *provided we stick to single-processor deterministic computers.* It was soon realized, however, that the restriction to deterministic computation was unacceptable: there is strong evidence that probabilistic machines can solve more problems in polynomial time, especially but not necessarily if we tolerate an arbitrarily small error probability. Since probabilistic machines have as much claim to physical realism as deterministic ones, it was in keeping with Turing's spirit to extend the notion of tractability accordingly.

The time is ripe for yet another paradigm shift. Because we live in a quantum-mechanical world, it is no longer possible to accept the thesis according to which the definition of tractability should take account only of computations performed on computers whose operation is restricted by the laws of classical physics, perhaps with the mere addition of true randomness. Quantum computation is an important formal notion, distinct from and with as much claim to physical realism as classical deterministic and probabilistic computation. The fact that no quantum computer has yet been built should not be considered more fundamental an issue than that no classical computers existed when Turing set forth his model of computation. I strongly believe that the class of problems that can be solved efficiently on a quantum computer with arbitrarily small error probability is the best concept yet invented by humankind to model the notion of what is computationally tractable by a physically realizable device.

When I expressed those views [13] in my response to Hartmanis's Turing Award paper (in which he had said that "There is a strong effort to define and prove the *feasible* limits of computation, but even here the basic model of computation is not questioned"), Hartmanis responded "The exploration [of quantum computing] deserves our best effort. At the same time, a theoretical model alone is not going to change the limits of feasible computations. [...] Its relevance still depends on the possibility of implementing it." [26] Although he is obviously right from a practical point of view—knowing Shor's algorithm won't help you factor large numbers if you can't build the device—I strongly disagree from a theoretical standpoint. In my opinion, the theoretical notion of feasible computation should be modelled on our understanding of the physical world, not on our technological abilities. After all, even the classical Turing machine is an idealization that cannot be built in practice: even not taking account of the unbounded tape, any real implementation of a Turing machine would have nonzero probability of making a mistake! Does this discount the model? I think not.

5 Quantum Cryptography

Shor's breakthrough has two spectacular consequences. We saw already that one is the answer to Feynman's question: physics cannot be simulated efficiently on classical computers assuming factoring is classically hard. The other consequence is that much of classical cryptography will crumble if and when a quantum computer is built. (Of course, "classical" is used here to mean non-quantum, and it includes secret-key and public-key cryptography on the same footing, just as "classical physics" lumps together Newton's mechanics with Einstein's relativity.) This is because the security of a significant proportion of cryptography in use today depends directly or indirectly on the presumed difficulty of either factoring or extracting discrete logarithms, and Shor has also discovered how to extract discrete logarithms in polynomial time on a quantum computer [32]. Secret key cryptosystems such as the DES have not yet been shown to be efficiently breakable by quantum computers, but they become vulnerable if their keys are exchanged with public-key techniques, which is often the case. A direct quantum cryptanalysis of the DES would be an interesting academic exercise even though its remaining lifetime is likely to be shorter than the time needed to build the first quantum computer. At the very least, it would be as much fun as writing a program to break Enigma on a contemporary classical supercomputer.

The security of information-theoretic based cryptosystems such as the one-time-pad would not be compromised by the availability of a quantum computer, provided the long keys they need to operate are transmitted securely. Thus quantum mechanics has the potential to be the nemesis of most of contemporary cryptography but ironically it also offers salvation thanks to quantum cryptography [4, 7, 12], which exploits the impossibility to measure quantum information reliably. When information is encoded with non-orthogonal quantum states, any attempt from an eavesdropper to access it necessarily entails a probability of spoiling it irreversibly, which can be detected by the legitimate users. Using protocols that I have designed with Bennett, building on earlier work of Wiesner, this phenomenon can be exploited to implement a key distribution system that is provably secure even against an eavesdropper with unlimited computing power. This is achieved by the exchange of very tenuous signals that consist on the average of one-tenth of one photon per pulse. Thanks to further work in collaboration with Crépeau and others [6, 14], quantum techniques may also assist in the achievement of subtler cryptographic goals, such as protecting private information while it is being used to reach public decisions. Ekert suggested the use of entanglement as an alternative approach to quantum cryptography [20]. Quantum cryptography provides a typical example of the thesis presented in the opening paragraph of this essay: it is provably impossible to achieve the benefits of quantum cryptography using only classical information if no restrictions are assumed on the eavesdropper's computing power.

In the early days of quantum cryptography, more than 15 years ago, the topic was thought of as science fiction even by the researchers directly involved in developing it. Today, several prototypes have been built, including one that is fully

operational over 30 kilometres of ordinary optical fibre [29], and encouraging experiments on quantum key distribution under Lake Geneva (23 kilometres) have just been completed [25]. The reason why quantum cryptography went from science fiction to science fact in such a short time is *not* due to a technological breakthrough: rather it is because we discovered how to make reasonable use of quantum mechanics for cryptographic purposes.

Perhaps the same fate awaits quantum computing. Perhaps quantum computing as we conceive it now will forever remain beyond available technology, yet perhaps we shall think of new, more reasonable ways to harness the strange properties of quantum mechanics for computing purposes.

6 Quantum Teleportation

Quantum information cannot be measured in general, but it can be *teleported* from one place to another [5]. This provides a striking (theoretical) application for spatially separated entangled qubits such as those described by Equation (1). Assume Alice and Bob share a pair of entangled qubits and let Alice be given a "mystery" particle that holds an unknown qubit. What should she do if she wishes Bob to have the mystery qubit instead of her? Surely, she cannot measure the particle to learn the information, which she could then transmit to Bob over a classical channel: any such attempt would be almost certain to spoil the information irreversibly. Instead, she makes it interact in the proper way with her share of the entangled pair. As a result, Bob's share instantaneously becomes a replica of the mystery qubit up to rotation. At the same time Alice's mystery particle loses its information but she obtains two purely random bits of classical information that tell her which rotation Bob should perform on his replica to match the original: she must communicate these two bits to Bob over a classical channel because Bob's replica remains indistinguishable from a purely random qubit until he has performed the required rotation.

This process cannot be used to transmit information faster than light between Alice and Bob (thank goodness!), but it can be argued that all the information that was present in the mystery particle is transmitted instantaneously—except for two (random!) bits that need to chug along at the speed of light! If needed, imperfect stores of spatially separated entangled qubit pairs can be *purified* by local transformations and exchange of classical information [8]. Quantum teleportation and purification provide an example that is even more convincing than quantum cryptography for the thesis according to which classical and quantum information can benefit from one another.

Although the teleportation of people is still very far away [;-)], experiments are under way to teleport the state of single qubits. In fact, a few quantum gates for the exclusive-or would be sufficient to implement teleportation. Conversely, quantum teleportation could be a useful tool to move information around in a quantum computer and to transmit quantum information for quantum cryptographic purposes. Thus we see that it all fits together.

Acknowledgements

Charles H. Bennett initiated me to the wonders of quantum mechanics many years ago. Since then, I have benefited from stimulating discussions with too many people to list them all here. Among them, André Berthiaume, Paul Bratley, Claude Crépeau, Artur Ekert, Nicolas Gisin, Sophie Laplante, Ueli Maurer, Pierre McKenzie, David Mermin, Asher Peres, Jean–Michel Raimond, Neil Stewart and William Wootters have contributed directly to the writing of this essay. I am also very grateful to Canada's NSERC for the freedom and funding that came with the E. W. R. Steacie Memorial Fellowship.

References

1. Barenco, A., D. Deutsch, A. Ekert and R. Jozsa, "Conditional quantum dynamics and logic gates", *Physical Review Letters*, Vol. 74, 1995, pp. 4083 – 4086.
2. Benioff, P., "Quantum mechanical Hamiltonian models of Turing machines", *Journal of Statistical Physics*, Vol. 29, no. 3, 1982, pp. 515 – 546.
3. Bennett, C. H., "Logical reversibility of computation", *IBM Journal of Research and Development*, Vol. 17, 1973, pp. 525 – 532.
4. Bennett, C. H., F. Bessette, G. Brassard, L. Salvail and J. Smolin, "Experimental quantum cryptography", *Journal of Cryptology*, Vol. 5, no. 1, 1992, pp. 3 – 28.
5. Bennett, C. H., 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", *Physical Review Letters*, Vol. 70, no. 13, 29 March 1993, pp. 1895 – 1899.
6. Bennett, C. H., G. Brassard, C. Crépeau and M.–H. Skubiszewska, "Practical quantum oblivious transfer", *Advances in Cryptology: Proceedings of Crypto '91*, Springer-Verlag, pp. 351 – 366.
7. Bennett, C. H., G. Brassard and A. Ekert, "Quantum cryptography", *Scientific American*, October 1992, pp. 50 – 57. Reprinted in *The Computer in the 21st Century*, Scientific American, Inc., 1995, pp. 164 – 171.
8. Bennett, C. H., G. Brassard, B. Schumacher, J. Smolin and W. K. Wootters, "Purification of noisy entanglement, and faithful teleportation via noisy channels", manuscript, 1995.
9. Bernstein, E. and U. Vazirani, "Quantum complexity theory", *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, 1993, pp. 11 – 20.
10. Berthiaume, A. and G. Brassard, "Oracle quantum computing", *Journal of Modern Optics*, Vol. 41, no. 12, 1994, pp. 2521 – 2535.
11. Berthiaume, A., D. Deutsch and R. Jozsa, "The stabilisation of quantum computations", *Proceedings of the Third Workshop on Physics and Computation — PhysComp '94*, Dallas, November 1994, IEEE Computer Society Press, 1994.
12. Brassard, G., "Cryptology column — Quantum cryptography: A bibliography", *Sigact News*, Vol. 24, no. 3, 1993, pp. 16 – 20.
13. Brassard, G., "Time for another paradigm shift", *Computing Surveys*, Vol. 27, no. 1, 1995, pp. 19 – 21.
14. Brassard, G., C. Crépeau, R. Jozsa and D. Langlois, "A quantum bit commitment scheme provably unbreakable by both parties", *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science*, 1993, pp. 362 – 371.

15. Coppersmith, D., "An approximate Fourier transform useful in quantum factoring", IBM T. J. Watson Research Report RC 19642, Yorktown Heights, NY 10598, USA, 12 July 1994.
16. Deutsch, D., "Quantum theory, the Church–Turing principle and the universal quantum computer", *Proceedings of the Royal Society*, London, Vol. A400, 1985, pp. 97 – 117.
17. Deutsch, D. and R. Jozsa, "Rapid solution of problems by quantum computation", *Proceedings of the Royal Society*, London, Vol. A439, 1992, pp. 553 – 558.
18. DiVincenzo, D. P., "Two-bit gates are universal for quantum computation", *Physical Review A*, Vol. 51, 1995, pp. 1015 – 1022.
19. Domokos, P., J.–M. Raimond, M. Brune and S. Haroche, "A simple cavity-QED two-bit universal quantum logic gate: principle and expected performances", submitted to *Physical Review A*, 1995.
20. Ekert, A., "Quantum cryptography based on Bell's theorem", *Physical Review Letters*, Vol. 67, no. 6, 5 August 1991, pp. 661 – 663.
21. Ekert, A. and R. Jozsa, "Shor's quantum algorithm for factorising numbers", *Review of Modern Physics*, to appear.
22. Einstein, A., B. Podolsky and N. Rosen, "Can quantum-mechanical description of physical reality be considered complete?", *Physical Review*, Vol. 47, 1935, pp. 777 – 780.
23. Feynman, R. P., "Simulating physics with computers", *International Journal of Theoretical Physics*, Vol. 21, nos. 6/7, 1982, pp. 467–488.
24. Feynman, R. P., "Quantum mechanical computers", *Optics News*, February 1985. Reprinted in *Foundations of Physics*, Vol. 16, no. 6, 1986, pp. 507 – 531.
25. Gisin, N., personal communication, 18 May 1995.
26. Hartmanis, J., "Response to the essays 'On computational complexity and the nature of computer science' " *Computing Surveys*, Vol. 27, no. 1, 1995, pp. 59 – 61.
27. Landauer, R., "Is quantum mechanically coherent computation useful?", in *Proceedings of the Drexel-4 Symposium on Quantum Nonintegrability—Quantum Classical Correspondence*, D. H. Feng and B.–L. Hu (eds.), International Press, 1995.
28. Lloyd, S., "A potentially realizable quantum computer", *Science*, Vol. 261, 17 September 1993, pp. 1569 – 1571.
29. Marand, C. and P. Townsend, "Quantum key distribution over distances up to 30 km", *Optics Letters*, to appear.
30. Peres, A., "Reversible logic and quantum computers", *Physical Review A*, Vol. 32, 1985, pp. 3266 – 3276.
31. Peres, A. and W. K. Wootters, "Optimal detection of quantum information", *Physical Review Letters*, Vol. 66, no. 9, 4 March 1991, pp. 1119 – 1122.
32. Shor, P. W., "Algorithms for quantum computation: Discrete logarithms and factoring", *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science*, 1994, pp. 124 – 134.
33. Simon, D., "On the power of quantum computation", *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science*, 1994, pp. 116 – 123.
34. Unruh, W. G., "Maintaining coherence in quantum computers", *Physical Review A*, Vol. 51, February 1995, pp. 992 – 997.

This article was processed using the LaTeX macro package with LLNCS style