

Beyond Bits: The Future of Quantum Information Processing



Two of the 20th century's most powerful ideas, quantum mechanics and computer science, are uniting into a yet more powerful body of knowledge, giving birth to new technologies and applications in a wide variety of industries.

Andrew M. Steane
Oxford
University

Eleanor G. Rieffel
FX Palo Alto
Laboratory

Today's computers, for all their marvels, operate on the same fundamental principle as the mechanical devices dreamed up by Charles Babbage in the 19th century and later formalized by Alan Turing: One stable state of the machine represents one number. Even seemingly non-standard computation models, such as the one based on DNA, share this basic principle. What else could they do? Yet physicists have shown that the laws describing the natural world are not the simple laws of classical mechanics—they are the subtler laws of quantum physics, and they invite us to think differently about computing. The computational principles that have guided us so well until now stem from classical physics and thus, we can be certain, are only partly right.

Recently, physicists and computer scientists have realized that not only do our ideas about computing rest on only partly accurate principles, but they miss out on a whole class of computation. Quantum physics offers powerful methods of encoding and manipulating information that are not possible within a classical framework. The potential applications of these quantum information processing methods include provably secure key distribution for cryptography, rapid integer factoring, and quantum simulation.

Information theory and quantum theory were among the most significant conceptual revolutions of the 20th century. Understanding of these theories led to the century's major technological advances. In the 21st century, we expect to see these theories unite to form an even more powerful force for advancement: quantum information theory.

QUANTUM COMPUTING

Throughout the history of computing, the bit has

remained the basic computational unit of information. Quantum mechanics enables the encoding of information in quantum bits (qubits). Unlike a classical bit, which can store only a single value—either 0 or 1—a qubit can store both 0 and 1 at the same time. Furthermore, a quantum register of 64 qubits can store 2^{64} values at once. Quantum computers can perform computations on all these values at the same time. However, extracting the results of these massive parallel computations has proved tricky, limiting the number of applications that have shown significant speed increases over classical computing. Classical parallelism can also increase the number of values handled simultaneously, but long before reaching the amount of parallelism achievable by a quantum computer, a classical system runs out of space. For classical systems, the amount of parallelism increases in direct proportion to system size; for quantum systems, it increases exponentially with size, as illustrated in Figure 1.

Quantum systems can operate in entangled states—states in which parts of the system correlate in ways that have no classical analog. Entangled states, such as the EPR pairs we discuss shortly, are responsible for most of the parallelism quantum systems achieve. Thus, computation using quantum parallelism is often called entanglement-enhanced information processing.

Any attempt to extract information from a state requires measurement. Unfortunately, in quantum computing, any measurement disturbs the state, thus destroying the quantum parallelism. Essentially, we can ask one, and only one, question about results generated by quantum parallelism before having to redo the computation. Moreover, the sort of question we can ask is restricted and is a subject of active research.



Millennium Speculations

We invited a few international experts to speculate informally on the future of quantum information processing. The first four responses come from theoreticians in quantum physics and computer science; the last comes from an experimentalist in quantum and atomic physics. (The questions we asked are labeled “AS/ER.”)

Lov K. Grover, Lucent Technologies

Let me start with a visionary quotation: “Where a calculator on the Eniac is equipped with 18,000 vacuum tubes and weighs 30 tons, it is possible that the computer of the future may have only 1,000 vacuum tubes and weigh only 1.5 tons” (*Popular Mechanics*, Mar. 1949). The idea is that due to the intrinsic unpredictability of innovation when we are talking of revolutionary technologies, I (or anyone else) could be way off the mark.

AS/ER: When will a 10-bit quantum computer be built? A 100-bit?

Grover: A 10-bit, special-purpose in two years, a 10-bit, general-purpose in 10 years, 100 bits in 100 years.

AS/ER: Will a quantum computer large enough to factor 1,000-digit numbers ever be built?

Grover: Yes, provided we do not find anything wrong with quantum mechanics.

Continued in next sidebars

Peter Shor¹ found a single question to ask in attacking the factoring problem, but researchers have found such a question for only a few problems.

For factoring, quantum parallelism provides a speed increase so immense it turns impossible computations into practical ones. Quantum computers are also exponentially better than classical computers at calculating the properties of quantum systems. Such calculations might appear to concern only a handful of physicists, but in fact they will impact a broad range of industries. For example, quantum physics is essential to the manufacture of increasingly small or complex devices, and it directly underlies chemistry. Suppose we wish to microfabricate a complicated, highly precise nanoscale device. We will need to understand a host of subtle quantum effects even to design the device, and time spent on a quantum computer to gain this understanding will be invaluable. Or consider pharmaceuticals. Among the biological molecules harnessed by evolution, we can expect to find a few that take advantage of subtle quantum effects that classical computing will have difficulty unraveling. Again, quantum computer time will be crucial.

The fact that measurement disturbs the quantum state turns out to be a benefit in other situations. Suppose we wish to communicate secretly. If we use quantum bits, a spy cannot learn anything without disturbing them—a disturbance we will notice. In fact—and this illustrates how subtle quantum physics is—the only type of message we know how to share in this quantum-secure manner is a completely random string of bits! However, a random string is the perfect key on which to base standard (classical) cryptography schemes. Using a communication system whose key transmissions are guaranteed by the laws of nature, the bank-account holder and the military commander need never again have a false sense of security.

QUANTUM PHYSICS

Quantum physics is the description of the principles underlying everything physical (the “laws of nature”) as scientists currently understand them. But quantum physics concepts such as entanglement and the measurement problem have no analog in everyday experience and are therefore difficult to understand. For the past century, physicists have carried on intense and unsatisfying debates about how to interpret some aspects of quantum mechanics, and these debates continue today. Nevertheless, the theory’s mathematical basis is thoroughly understood and enormously successful—of all physical theories, quantum physics makes the most precise predictions. This detailed understanding of nature’s workings has enabled the development of a wide variety of applications, including transistors, lasers, and medical-imaging techniques. Thus, quantum phenomena, although in some ways inexplicable, are usable and useful.




Figure 1. Classical versus quantum parallelism: To achieve the same degree of parallelism as (a) 300 quantum processors ($n = 300$), we would need (b) 2^{300} classical processors. Since 2^{300} is more than the number of particles in the universe, to say that quantum computing enables an astronomical increase in parallelism is obviously an understatement.

Qubits

In a famous experiment, light from a single source passes through two slits, creating an interference pattern on a screen. Even when the light source emits only one photon at a time, interference patterns appear. Standard quantum theory postulates that each photon travels both paths at once. Thus, a particle can be in two places at the same time. In such situations, we say that the particle’s position is in a superposition of two states.

The two paths that a particle travels can represent the two states of a bit, 0 and 1. In quantum mechanics, whenever a system has two or more possible

Millennium Speculations

Seth Lloyd, MIT



Quantum information processing is likely to have a large future impact on physics—we may eventually understand the fundamental dynamics of the world in terms of information rather than Lagrangians—and on engineering—more and more technologies are gravitating toward the quantum realm. It will also continue to change applied mathematical research in computational complexity.

A 10-bit quantum computer will probably be built in two years. I predict the number of bits will double every one to four years. Eventually, a quantum computer large enough to factor 1,000-digit numbers ought to be built, barring unforeseen global catastrophe. It will be expensive and will require 10,000 bits or more for error correction. This should take at least 30 years. If funding runs dry, it will take longer.

Quantum cryptography could become profitable, but only in limited contexts for at least 30 years.

Moore's law is already slowing down. In terms of ultimate physical limits to computation, it cannot continue for more than another 300 years, at which point we would have computers that are 40 orders of magnitude faster than current machines. The memory version of Moore's law is likely to last considerably longer than the speed version and get closer to the ultimate physical limits of approximately Avogadro's number of bits per cubic centimeter at room temperature.



Figure 2. Strange quantum correlations: When Alice and Bob play with their coins individually, the coins appear to behave normally—tossing each coin gives perfectly random results. However, we soon notice that every time Alice's coin lands heads, so does Bob's, and vice versa. Since there is no possible means of communication between the coins, this must be magic, right? No, it's quantum mechanics.

avenues, it can explore them simultaneously. Any two-state system, like the photon's paths, can represent a qubit. In a quantum computer, we might instead use two orbits of an electron in an atom to represent our qubit. The atom can exist in a superposition of 0 and 1, just as a struck bell can vibrate at two different frequencies simultaneously.

Quantum parallelism

Quantum computers operate on both values stored in any qubit at the same time. Moreover, n qubits, each in a superposition of 0 and 1, encode 2^n values, and a quantum computer can compute on all these values at once. This massive parallelism, exponential

in the number of particles used in the computation, is called quantum parallelism. Any classical circuit has a corresponding quantum circuit.² So a quantum computer can make any calculation on all values in roughly the same time an ordinary computer would take to make a calculation on a single value.

Entangled particles

At the heart of quantum parallelism lies the fact that quantum superpositions of multiple particles admit strange correlations with no classical analogs. To get an idea of these correlations, imagine that Alice and Bob are each given a coin (which we'll say behaves like a particle). When Alice and Bob play with their coins individually, the coins appear to behave normally; in particular, tossing each coin gives perfectly random results. However, we soon notice that every time Alice's coin lands heads, so does Bob's, and vice versa. There is no possible means of communication between the coins. Magic? No, quantum mechanics (see Figure 2).

Although physicists have not seen such correlations between large, separate objects, including coins, they have created such correlations between individual atoms. We call objects correlated in this way *quantum mechanically entangled*. Pairs of objects exhibiting such correlations are called EPR pairs (after Einstein, Podolsky, and Rosen, who first discussed them). Most possible quantum states contain quantum correlations, and this fact is responsible for the exponential nature of quantum parallelism and the success of quantum algorithms. It was Richard Feynman's observation that classical computers could not efficiently simulate certain types of entanglement that promoted early interest in quantum computing.³

The measurement problem

Accessing results obtained through quantum parallelism requires measuring the final state of the qubits. Any measurement instrument registers a single result, even though the quantum computer may be storing a superposition, possibly huge, of different values. Nature resolves this paradox by destroying the other values in the superposition whenever a measurement is made, transforming the state from a possibly complex superposition to a simple state consisting of the single value read. For quantum computing, this difficulty in accessing values is a severe limitation, requiring highly unconventional programming techniques to finesse the problem, such as those of Peter Shor and Lov Grover.^{1,4}

QUANTUM ALGORITHMS

The trick to getting the information you want when you measure is to make quantum transformations to the superposition of all values. Two methods currently exist. Shor's algorithm measures a common property



Millennium Speculations

David Deutsch, Oxford University

AS/ER: What kind of impact can we expect from research into quantum information processing?

Deutsch: Quite sweeping. For the next few decades, however, it will be almost entirely a philosophical impact. After that, a technological one too.

AS/ER: Will a quantum computer large enough to factor 1,000-digit numbers ever be built? If so, when?

Deutsch: Of course, in a few decades.

AS/ER: Will quantum cryptography become profitable, and, if so, when?

Deutsch: I'm amazed that it hasn't already.

AS/ER: Will RSA be broken? If so, when and how?

Deutsch: Of course. Five minutes after the first sufficiently large quantum factorization engine is built. Note that for purposes of long-term security, it is already broken by the mere theory of quantum computers.

of all the output values. Grover's algorithm amplifies the results of interest.

The best classical algorithm for factoring large integers takes an amount of time that increases superpolynomially with the number's length, and is thus unworkable for large numbers. A variety of encryption schemes, including the widely used RSA algorithm (named for its inventors, Rivest, Shamir, and Adelman), rely on this fact. Shor, building on previous workers' insights, surprised the computing world and spurred a flurry of interest in quantum computing when he found a quantum algorithm that can factor numbers in an amount of time that increases only with the cube of the number of digits in the number. This elegant algorithm first computes all the values of a certain function using quantum parallelism and entanglement and then uses a quantum analog of the Fourier transform. Measurement

then gives a value from which one can extract the function's period and use it to factor the number in question. (See the sidebar "Factoring, RSA, and Shor's Factoring Algorithm.")

Factoring, RSA, and Shor's Factoring Algorithm

The long-held belief that factoring large integers is computationally infeasible on conventional computers underlies the most robust encryption schemes in use today, including RSA.

Public-key encryption

RSA is a public-key encryption scheme; that is, it enables a person to post a public key by which others can encrypt messages that only the person who posted the key can decrypt. Public-key encryption relies on trapdoor functions—functions that are hard to compute but become easy once a piece of information is known.

RSA establishes a public key using two large primes p and q , and a number e that shares no factors with $p - 1$ and $q - 1$. Finding such numbers is a relatively easy computational task. A user then posts e and n , the product of p and q , while keeping p and q secret. To encode a message M , the user computes $M^e \bmod n$, a computation that takes little time.

To decrypt the message, you need to find a d such that $M^{ed} \bmod n = M$ for any message M . If you know just e and n , finding such a d is difficult, but once you know p and q , finding d becomes easy. Thus, finding d in terms of e and n is a trapdoor function, provided it is hard to find n 's factors p and q . But if factoring becomes computationally easy, so does deciphering RSA-encrypted messages.

Shor's algorithm

The most dramatic example of the power of quantum computing to date is Shor's bounded-probability, polynomial-time factoring algorithm.¹ Shor attacked the factoring problem by combining knowledge from classical and quantum computer science. Suppose a number $n = pq$ has two unknown factors p and q . A simple classical algorithm can deduce p and q from two other numbers. The first, a , can be chosen randomly, while the second, r , is the period of the function $f(x) = a^x \bmod n$, which must be calculated.

Shor's algorithm describes how to determine r efficiently with a quantum computer. A superposition of all numbers x between zero and n^2 is placed in a register of quantum bits. All the values of $f(x) = a^x \bmod n$ are calculated simultaneously using quantum parallelism and stored in a second register. Each result remains entangled with the x that generated it; if the first register were measured and found to be x , a measurement of the second register would yield $a^x \bmod n$, just as in the EPR correlations. However, such measurements must not be carried out! They would destroy the superposition. Rather, we want to know how frequently the sequence of remainders repeats itself, since that is the sought-after r . So we perform a quantum Fourier transform on the first register, the most subtle step in the algorithm. The entanglement between registers allows the peri-

odicity of the second register to be reflected in the first.

Although the quantum Fourier transform is based on the classical fast Fourier transform, it is efficient on a quantum computer, another remarkable piece of quantum parallelism. The relatively few operations of the quantum Fourier transform result in the complex overlapping of vast numbers of computational strands. In the end, the first register stores a superposition that is essentially only multiples of n^2/r . We measure the register to learn one of these multiples, from which, with high probability, we can deduce r .

Shor's algorithm also breaks generalizations of the RSA scheme that use the discrete-logarithm problem. Since large quantum computers have yet to be built, RSA-encrypted data remains secure for the time being. Nevertheless, even today, people should worry about data that must stay secure for decades into the future. Although quantum cryptography provides a provably secure key distribution scheme, it loses the advantages of public keys, and there is no known quantum method to regain them. No one yet knows how to create a public-key distribution scheme based on quantum methods.

Reference

1. P. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *SIAM J. Computing*, Oct. 1997, pp. 1,484-1,509.

Millennium Speculations

Artur Ekert, Oxford University



Physicists will use the language of information processing to rephrase the fundamental theories in physics. In fact, I believe that in the next decade courses in quantum mechanics will start with a notion of a qubit rather than with particles in silly potential wells. Computer science will be recognized as a branch of physics. Logicians and philosophers will be digesting issues such as the physics of mathematical proof and where the certainty of our mathematical knowledge comes from. Then quantum technology will follow—I do hope not to live long enough to see the launch of Quantum Windows.

AS/ER: Will a quantum computer large enough to factor 1,000-digit numbers ever be built?

Ekert: Of course, in a few decades.

AS/ER: Will quantum cryptography become profitable?

Ekert: Within this decade.

AS/ER: Will RSA be broken?

Ekert: Within this decade. Not necessarily via factorization or quantum means.

Grover's quantum algorithm for unstructured search provides a quadratic speedup over what is possible classically. Although this speedup is not as dramatic as that achieved by Shor for the factoring algorithm, it is provably better than any possible classical algorithm. Shor's is simply better than any

currently discovered classical algorithm (although most researchers conjecture that it is better than any possible classical algorithm). Grover's algorithm slowly changes the quantum state so that the state being searched for is more and more likely to be the one measured. A curious aspect of this algorithm, unlike most classical probabilistic algorithms, is that you have to know when to stop. The longer you run most classical probabilistic algorithms, the better your results. But if you run Grover's algorithm too long, the chance of obtaining the desired result decreases.

Although a quadratic speedup can have significant practical impact (the fast Fourier transform gives a quadratic speedup over the Fourier transform), we would hope for more. For unstructured search, however, more is not possible; Grover's algorithm is known to be optimal. Clearly, even great feats of imagination will not enable quantum parallelism to exponentially speed up all problems.

THE FUTURE

The future poses two large questions: What quantum information-processing devices can we build? And what can such devices do?

Building quantum key distribution devices

Research labs have realized quantum key distribution over a standard fiber-optic network at distances of up to

Quantum Key Distribution

Quantum key distribution was the first development in quantum information theory likely to have a commercially viable application. Researchers have proposed a variety of distribution schemes and implementations. As an example, we have chosen a scheme developed by Charles Bennett and Gilles Brassard, who built on the ideas of Stephen Wiesner.¹ The scheme aims to establish a random secret key, consisting of a sequence of 0s and 1s, known only to the two people who wish to communicate.

One of the parties, Alice, generates a random sequence of 0s and 1s, a subsequence of which will become the secret key. She sends Bob one photon for each number in the sequence, randomly and secretly choosing {0,1} encoded as either {horizontal polarization, vertical polarization} or {-45° polarization, +45° polarization} for each photon. Bob makes a random choice between vertical and diag-

onal orientation of his polarization detector for each photon he receives. Depending on its orientation, a polarization filter can distinguish either vertical from horizontal polarization or diagonal +45° from diagonal -45°. However, a vertical detector, attempting to measure diagonal light, randomly forces the photon state to either vertical or horizontal. Similarly, a diagonal detector forces the photon state to +45° or -45°. Bob's choice of orientation will match Alice's half the time. When they match, Bob's result corresponds to what Alice transmitted; the other half of the time, he gets random (thus useless) results. To find out which are his useful results, he and Alice tell each other publicly which orientation they used for every photon sent but not the precise states transmitted and received. They keep the bits whose orientations matched and throw away the rest, obtaining a shared random key of roughly $n/2$ bits, assuming nobody tampered with the transmission.

So far the scheme has gained nothing by using quantum states. However, if Eve, an eavesdropper, intercepts the photons en route, she must introduce a polarization detector to learn anything; the detector will disturb many of the photons it measures. In fact, Eve randomizes the polarization whenever she doesn't by luck choose the same orientation as Alice. To test for eavesdroppers, Alice and Bob publicly compare some of their good bits. If the bits match, no randomization of these bits occurred. If enough bits match, Bob and Alice can be highly confident that no eavesdropper was present and can use the rest of the key to communicate privately.

Reference

1. C.H. Bennett, G. Brassard, and A.K. Ekert, "Quantum Cryptography," *Scientific American*, Oct. 1992, pp. 50-57.

48 km and are testing distribution by laser beam to satellites. Currently, the data rate is low, and security loopholes still exist. However, once the technology a spy would need to exploit such loopholes becomes available, it can also be used to close them. Thus, quantum key distribution is already close to commercial viability. The “Quantum Key Distribution” sidebar describes a proposed scheme. Researchers are working to combat the noise problem, extend the distribution distance, and develop single-photon light sources to enhance the data rate. Finding other ways quantum effects can improve communication and encryption of both quantum and classical data is also an active research area.

Building quantum computers

The largest quantum computers to date have achieved about 100 logic operations on two qubits or 10 operations on seven qubits. At this scale, a device does not perform valuable computational tasks but does teach us about the principles of quantum computing and their physical realization. The seven-qubit system, a nuclear magnetic resonance (NMR) spectrometer, demonstrates the superposition principle but not entanglement. So far, experimentalists have been able to efficiently create entanglement only in a laser-cooled ion trap, in which lasers drove two charged atoms into an entangled state with 90 percent reliability. Figure 3 shows a schematic diagram of an ion trap. Figure 4 is a photograph of an actual ion trap from Andrew Steane’s lab. Many researchers have proposed larger quantum computers using ion trap, NMR, and other methods. So far, all these proposals suffer from noise and scaling problems severe enough that no one sees a way to implement more than about 100 qubits.

On the other hand, no one sees a fundamental physical barrier to building large quantum computers. Thus, we can be optimistic that a new proposal, or a significant variant of a current proposal, will lead to such computers. It seems certain, however, that a quantum computer large enough to break current RSA standards will not be built within the next 20 years. How long it takes quantum computers to develop depends not only on how difficult they are to implement but also on how much incentive exists to build them. That incentive in turn depends on how broad a class of algorithms researchers find and how quickly miniaturization reaches its fundamental physical limits.

The end of Moore’s law

More than 25 years ago, Gordon Moore predicted that the computing power of a single chip would continue to double every 18 months—a prediction so accurate that it is known as Moore’s law. This increase in computing speed results mainly from the increasing miniaturization of components.




Figure 3. Schematic diagram of a 5-qubit ion trap, a rudimentary quantum computer. The five calcium ions trapped by electrodes represent five quantum bits. Laser beams can perform quantum transformations on the ion strings and are also used to cool and stabilize the individual ions.




Figure 4. A photograph of an actual ion trap used in Andrew Steane’s laboratory. The four central electrodes (blue)—only about 1 mm in width and 3 cm long—trap the calcium ions to be laser-cooled and used in experimental manipulation of quantum information.

Sometime in the next 10 to 20 years, component technology will reach various strict physical limits that will mean the end of Moore’s law. For example, if miniaturization were to continue at its current rate, by 2010 or so it would reach the limit of transistors acting on single electrons. Depending on the success of nanotechnology, miniaturization will increase computing power by somewhere between 1,000 and 100,000 times the current speed—and that is all. Achieving any speed increases beyond these bounds



Millennium Speculations

David Wineland, National Institute of Standards and Technology

Significant attention is being paid to possible condensed-matter implementations of quantum computation. This is natural, given the immense success of classical computers based on this technology. However, the requirements are drastically different for quantum and classical computation, so all possibilities should continue to be explored.

It seems that a useful quantum computer is a way off. The two killer applications seem to be factoring and searching—we need more! Nevertheless, quantum communication and quantum cryptography seem to be nearly in hand, and we could see practical implementation relatively soon.

AS/ER: When will a 10-bit quantum computer be built?

Wineland: Notwithstanding the arguments about the “quantumness” of NMR information processors, it seems likely that some aspects of computation with 10 bits will be achieved in NMR in the next few years. A competitor for this goal might be trapped ions if solutions to specific problems can be found.

AS/ER: A 100-bit quantum computer?

Wineland: Tough, but I’d say from 10 years to never.

AS/ER: Will a quantum computer large enough to factor 1,000-digit numbers ever be built?

Wineland: I’d say if a 100-bit quantum computer can be built, a much larger one could also be built, thus allowing big-number factorization.

will require different approaches. Quantum information processing has a dual role to play here. Learning how to carefully control and manipulate quantum information will enable us to get closer to these fundamental physical limits. And, in certain cases, quantum computers will allow us to bypass the limits.

Quantum-computer applications

Most computationally hard practical problems, like nondeterministic polynomial-time (NP) problems, have some structure, so the optimality limit of Grover’s algorithm does not apply. But these problems do not have as regular a structure as the factoring problem. In the five years since the announcement of Shor’s algorithm, many researchers have tried unsuccessfully to find quantum algorithms to solve NP-complete problems. Another approach is possible. Classical heuristic algorithms for NP-complete problems are used frequently in industry with reasonable success. Although these algorithms are slower than desirable and sometimes fail completely, they are the best way known today for solving this wide class of practical problems.

The average efficiency of both classical and quantum heuristics is hard to analyze. We test classical heuristics by running them on a large random sample of problems. Before we can evaluate the usefulness of heuristic quantum algorithms, we may have to wait for large quantum computers. If an efficient quantum algorithm that generally solves NP-complete problems

continues to elude researchers—as seems likely—any quantum heuristics that improve on classical heuristics could prove significant.

Nevertheless, quantum computers will not replace classical computers. In many situations, quantum computers, because of constraints on the way they must be built, will perform much more slowly than classical computers. Only special situations requiring the great efficiency that quantum parallelism provides will benefit from quantum computers.

Working with quantum states

With the steady push toward smaller and smaller appliances, some devices will be so small that classical physics will no longer suffice, and knowledge of quantum physics will become essential. Already, atom wave devices based on quantum atom interference rival the most sensitive inertial sensors, such as the laser gyros and gravity detectors used in spacecraft. The implications of better understanding and control of quantum phenomena are by no means restricted to computing. The development of even relatively small quantum computers would enable the efficient simulation of quantum systems, which in turn would lead to an even better understanding of quantum phenomena and their uses.

Already our potential to control quantum phenomena has increased greatly with a development straight from quantum information science: quantum error correction. Active error correction is essential to any reliable complex system, biological or man-made. Think of the feedback loop involved in such a simple activity as picking up a glass, for example, or the governor on a steam engine and its modern cousins in the form of control loops engineered into any complex industrial process. But for many years quantum error correction was believed impossible. Classical error correction relies on amplification to detect the error and an irreversible damping process. Quantum theorists knew well that quantum information cannot be amplified, and any damping disturbs the computation just as measurement does. The only hope, it seemed, was the dim possibility of building the computer so well that it doesn’t generate errors. It is almost certain, however, that a computer large enough to run Shor’s algorithm on interesting cases would generate errors. Quantum systems are notoriously fragile. Interaction with the external environment easily results in unintended measurements that corrupt the quantum states.

Fortunately, quantum entanglement proved such a fascinating subject in its own right that researchers sought new ways to think about it. Inspired by classical information theory, Robert Calderbank and Shor and, independently, Andrew Steane, saw how to perform quantum error correction. To correct an

error, we only need to understand what quantum transformation of the state occurred—and then undo it. We can set up mechanisms that determine what transformation occurred without gaining any information about the actual state before or after the error. If we avoid gaining information about the state, we avoid the problems described in the previous paragraph. Quantum error correction is now one of the most highly developed areas of quantum information theory, and it will have applications wherever people need to control and manipulate quantum systems. (Several sources provide more details on quantum information processing.⁵⁻⁸)

Quantum information theory seeks to unite some of the most influential ideas of 20th-century science: quantum mechanics, computer science, and information theory. The development of quantum information theory has only begun, and only a few applications are known, mostly in quantum system control and data security. Where exactly the theory will lead is hard to predict, but it seems poised to contribute to some of the most exciting ideas of the 21st century. Quantum information theory gives us an ideal framework for developing a better understanding of how nature works and what it will let us do. Such advancements in knowledge led to new technologies and applications in the past and surely will do so again—to those we have suggested and to those yet undreamed of. ❖

References

1. P. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *SIAM J. Computing*, Oct. 1997, pp. 1,484-1,509.
2. D. Deutsch, "Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer," *Proc. Royal Society of London, Series A*, A400, 1985, pp. 97-117.
3. R.P. Feynman, "Quantum Mechanical Computers," *Lectures on Computation*, A.J.G. Hey and R.W. Alice, eds., Addison-Wesley, Reading, Mass., 1996, pp. 185-211.
4. L.K. Grover, "A Fast Quantum Mechanical Algorithm for Database Search," *Proc. 28th Ann. ACM Symp. Theory of Computing*, ACM Press, New York, 1996, pp. 212-219.
5. A. Hey, ed., *Feynman and Computation*, Perseus Books, Reading, Mass., 1999.
6. J. Preskill and A. Kitaev "Course Information for Physics 229 Advanced Mathematical Methods of Physics," <http://www.theory.caltech.edu/people/preskill/ph229/index.html#lecture>.
7. E. Rieffel and W. Polak, "An Introduction to Quantum Computing for Non-Physicists," to be published in *ACM Computing Surveys*, June 2000; <http://xxx.lanl.gov/abs/quant-ph/9809016>.



Millennium Speculations

We finished our survey with a light-hearted bet, based on one informally suggested by Peter Shor some years ago: "The first factorization of a 1,000-digit number will be done on a quantum, not a classical, computer." (We are assuming a product of suitably chosen, approximately 500-digit primes.) We asked our experts, "Which side of this bet would you choose—quantum or classical—and how many bottles of wine would you venture on it?" The answers:

Grover: The classical—Moore's law will win out.

Deutsch: Quantum. I don't want to win wine. Would you take chocolate?

Lloyd: I suspect that Peter is right. I'd bet a few bottles of good wine.

Ekert: By some sort of a quantum computer or a hybrid quantum/classical device. I would bet a couple of bottles of Medoc Premier Cru.

Wineland: I'd bet on classical. I'm not really a betting person, so I'd bet just one bottle of wine, which I probably wouldn't receive or have to deliver in my lifetime anyway.

And ourselves?

Steane: Quantum, one good bottle.

Rieffel: Quantum. Ten pounds of the best chocolate.

8. A. Steane, "Quantum Computing," *Reports on Progress in Physics*, Vol. 61, 1998, pp. 117-173; <http://xxx.lanl.gov/abs/quant-ph/9708022>.

Andrew M. Steane is a university lecturer and fellow of Exeter College, Oxford University, where he is developing an experiment to investigate quantum information theory in laser-cooled trapped ions. From 1995 to 1999, he held a Royal Society University Research Fellowship at Clarendon Laboratory, Oxford. Through his study of the theory of quantum interference involving many particles, he established the basic principles of quantum error correction theory. Earlier, he was a junior research fellow at Merton College, Oxford, and then a European research fellow at the École Normale Supérieure, Paris, where he developed experiments in laser cooling and control of atoms, which led to the interferometry of atomic de Broglie waves. He received the Maxwell Medal of the Institute of Physics for this work. Steane has a PhD in physics from Oxford University. He is a member of the Institute of Physics. Contact him at a.steane1@physics.ox.ac.uk.

Eleanor G. Rieffel is a research scientist at FX Palo Alto Laboratory, where she has been working since 1996. Currently, she is conducting research in both quantum information theory and adaptive systems. She received a PhD in mathematics from UCLA in 1993 and spent the next three years as Busemann Assistant Professor in the Department of Mathematics at USC. She is a member of the IEEE, the ACM, and the AMS. Contact her at rieffel@pal.xerox.com.