We're sorry we don't have a quantum computer to show you, but here (from left) Gillian Pierce; Brock Beauchamp, a junior in electrical and computer engineering; his mentor, James Arvo, associate professor of computer science; and John Preskill, professor of theoretical physics and mentor to Jacob West, a junior in physics, pose with Caltech's Center for Advanced Computing Research's Exemplar, the biggest machine Hewlett-Packard has yet built.

Who is going to be the next Carl Sagan? The next Stephen Jay Gould? This year, the Institute added a new course to the core curriculum: Core 1ab, Science Writing. To quote from the course's Web site, "Communicating scientific ideas is one of the most fundamental tasks that a scientist or engineer undertakes, and nonscientific audiences provide one of the most challenging groups to write for." During the two-quarter course, students write (and rewrite!) a 3,000-word essay on any topic in science, broadly defined. Since it's a writing course, not a lab course, they do not have to write about their own research but about any subject that appeals to them. This year's topics ranged from the history of science to flaviviruses, Fermat's Last Theorem, and the origin of the universe. The essays, according to program coordinator and editor Gillian Pierce, are supposed to be comparable to an article in Scientific American *or our own* E&S*, several of whose past stories were posted on the Web site as models.*

Science writing has been taught at other colleges, but never with so much faculty involvement. Each student picks a faculty mentor who is responsible for critiquing the essay's science content, while Pierce works on improving the writing. The faculty input adds an element of peer review to the process, making the course a good exercise for those students who will actually go on to publish academic papers.

The course, which will be required of all undergrads next year, was offered this year as an option. Fifteen adventurous students signed up. All of their papers will be published in an on-line journal (http://www.its.caltech.edu/~sciwrite/ejournalhome.htm)*, but we thought you might like to see a couple of the best ones, as chosen by Pierce and the staff of* E&S*. What follows are two very different looks at a hot research topic. And for what some people at Caltech are doing,* *see the sidebar on page 29.*

The Dawn Of Quantum Computation

by Brock Beauchamp

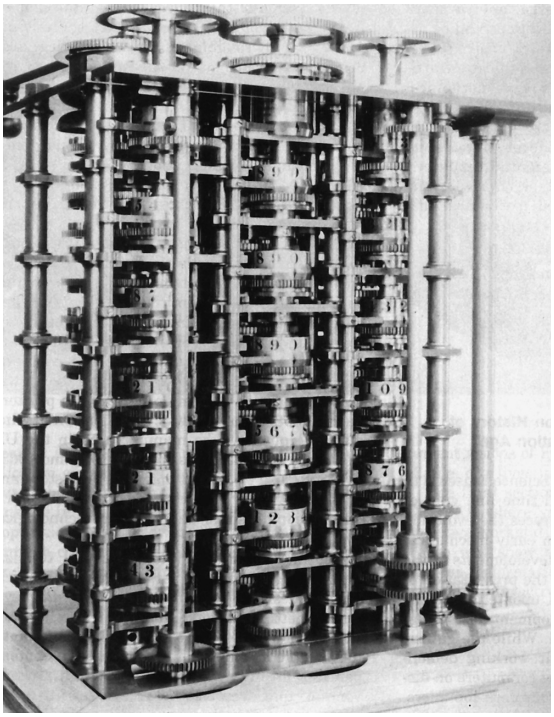

The largest surviving portion of Charles Babbage's Difference Engine, built in 1832, is in the Science Museum, London. Photo from "The Science Museum, London History of Computing and the 'Information Age,'" by Doron Swade, in the *Annals of the History of Computing*, volume 10, number 4, page 316, 1989. Copyright © 1989 IEEE.

In what now seems to be the dawn of time, around 500 B.C., the Babylonians invented a primitive "computer"—the lowly abacus. Over two thousand years later, in A.D. 1614, Scotsman John Napier, the inventor of the logarithm, renewed the interest in creating more advanced mechanical computers. The most famous of these improbable devices was the Babbage Difference Engine, which was drafted as a steam-powered apparatus that could solve one fixed problem, using thousands of gears and dials. It would have done these calculations with 20-decimal-place accuracy, but it was a costly and unwieldy feat of engineering that eventually lost funding. Such was the fate of most mechanical computers, which history remembers as little more than novelties, albeit novelties with foresight. Computation did not truly come of age until machines powered by vacuum tubes appeared on the scene in the early 20th century. When these behemoths were scaled down by the advent of the transistor in 1947, computational power that was once restricted to testing theories behind the H-bomb was available to the masses.

Today's computers are certainly faster than their predecessors, but they share many of the same inherent weaknesses. For example, they are stymied by the significant problem of factoring large numbers. Using the best algorithm to date, the number-field sieve, one can factor a 130-digit number in a little more than a month. However, factoring a 260-digit number, just twice the length, would require over a million years on the same computer! Clearly, an entirely different kind of tool is needed to solve such difficult problems, and many hope the quantum computer will be just that panacea.

The Challenge and the New Contender

In order to better appreciate these challenges, an understanding of computational complexity is helpful. To better systematize the difficulty of problems, they are often sorted into complexity classes. The gauge for complexity is how many steps it takes to solve a problem (the number of steps often being loosely called "time") with respect to the length of the input. Computer scientists are typically concerned with asymptotic complexity—that is, complexity as the size of the input grows very large. Using this criteria, many problems have been deemed intractable, meaning that any algorithm able to solve the problem has a prohibitive asymptotic complexity. (It is possible that there is some feasible way to approach "intractable" problems, but the evidence to date strongly suggests that the difficulty of these problems is unassailable.) In other words, making the problem just a little longer makes it considerably harder to solve. These "hard" problems are theoretically solvable on a computer, but quickly become impractical. For example, suppose that a company wanted to find the shortest route between all its regional offices. If there were 40 offices,

If entanglement gives the quantum computer its voice, it is quantum parallelism that gives it its muscle.

a computer would have to examine 40!, or $40 \times 39 \times \cdots \times 1$, which is approximately equal to 10^{45} different routes (that's a 1 followed by 45 zeroes), by first choosing one of the 40 offices, then one of the remaining 39, and so forth. Using current projections, the sun will supernova long before any computer could finish checking all of these possibilities! It seems as though there could not be a harder problem; however, there are well-formed problems that are uncomputable on any machine. The classic example is the halting problem: no program can be written that can tell whether or not any given program will eventually stop and return a value.

What new ammunition does quantum computation have to combat these difficulties? For one thing, quantum systems deal with information in an entirely different way. All information is represented in terms of an elementary unit called the qubit (short for "quantum bit," denoted in Dirac notation by "$|\phi\rangle$"). Qubits, which have no classical analog, exhibit a sort of quantum indeterminacy: the qubit is not in *any* state in particular until it is tested, after which it has a definite state. Because nature is ordered according to these quantum principles, each qubit is a complete representation of the system it represents, without any extraneous data. Information scientists are wont to describe such properties in the context of a fictitious conversation between Alice and Bob, so we will not break tradition here. In the classical scenario, Alice would look at her information and write, "Dear Bob, I have the state $|0\rangle$. Sincerely, Alice." Or, if it were a physical bit of information, she could simply make a copy and send it over to Bob. However, she cannot do this in a quantum information system. In the first case, Alice cannot simply measure her qubit and send the results as she did in the classical case. She might test her qubit $|\phi\rangle$, and in doing so force it into state $|0\rangle$, but she would not be sending *all* the information contained in the multiple states that were initially in $|\phi\rangle$. Second, it has been proven that it is physically impossible to clone a qubit while leaving the original untouched. This means that Alice cannot simply copy her qubit and send the copy to Bob. This leads to a very important result: the information contained in a qubit cannot be transmitted without sending the qubit itself.

Taking Advantage of Quantum Quirks

The inability to transmit qubits is no small problem—in order for quantum computers to be very useful, they need to be able to send information to other computers (in a network) and to the user (as output) without losing the copy they possess. The solution to this problem turns out to be the quirk known as quantum entanglement. It is a disturbing fact of modern physics that pairs of particles may be produced such that the measurement of one particle has an effect on the measurement of the other, even if they are separated by a great distance. At most, all Alice has to send to Bob is an explanation of what kind of measurements she performed on the "quantum twin" in her possession, which may be sent classically. The information that Bob gets is complete; his information perfectly reflects the state of Alice's qubit. However, because of the "no-cloning" theorem, Alice's qubit is destroyed in the process. Because of these properties, many refer to the process as quantum teleportation. According to Jeff Kimble, an expert in quantum optics at Caltech who demonstrated the first bona fide teleportation in 1998, "entanglement means if you tickle one, the other one laughs." Or, one could view entanglement like a pair of quantum dice that always add up to seven. Before one of the dice is rolled, neither die can be said to have a value. But when one is rolled, say as a three, that act determines the value of the other (to be a four, in this case). It's no sur-

prise that Einstein called this behavior "spooky action at a distance." While the mysteries of entanglement have stymied physicists for years, they are the keys to the quantum computer's ability to transfer and process information.

If entanglement gives the quantum computer its voice, it is quantum parallelism that gives it its muscle. Recall a fundamental property of the qubit: before it is tested, it is in many different states at the same time (technically speaking, a superposition of states). It is therefore possible that each one of these states could function like a separate computer, following a single computational path and coming up with a result. Each of these states then interferes with the others, like ripples on a pond, forming a peak that is interpreted as the final output. This is an important departure from the classical model, because it means that the right answer is only found with a certain probability. It will often take many trials before any degree of certainty can be established. Still, the ability to have so many parallel "computers" in one piece of hardware is what gives the quantum computer its unprecedented power.

The Birth of a Science

While there is no shortage of skepticism about quantum computation, there have been a number of early demonstrations of its promise. Like every other new technology, quantum computers began as a mere theoretical fascination, waiting in the wings for a practical application. In 1993, at the 35th Annual IEEE Symposium on the Foundations of Computer Science, Peter Shor delivered a groundbreaking paper that proved to be that "killer app." (Pronounced "eye-triple-E," IEEE stands for the Institute of Electrical and Electronics Engineers, a major clearinghouse for electrical standards and research.) More specifically, he presented an algorithm that can factor very large numbers, yet does so efficiently even as the input size grows bigger. Since many of the pieces that Shor incorporated into his algorithm have been known since 300 B.C., one may well wonder why his discovery was so remarkable. Although most of the methodology behind the algorithm is nothing new, Shor managed to use procedures from the classical realm that could benefit from quantum parallelism. This is particularly significant given that factorization is believed to be an intractable problem for classical computers. While it hasn't been proven to be one of those "hard" problems, it has thus far been such a Herculean feat that most cryptography depends on its difficulty. The connection is no mere coincidence—the ability of the quantum computer to make and break codes is what has driven most of the interest in the field. The prospect of a drastic increase in the speed of code-breaking algorithms was enough to make the scientific community, not to mention government agencies, stand up and take notice.

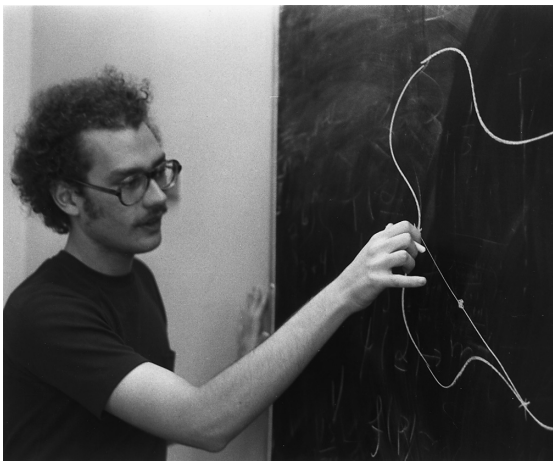

Peter Shor (BS '81, mathematics) won a national prize in the William Lowell Putnam Mathematics Competition as an undergraduate (see E&S, June/September 1981) and is now a big deal in quantum computing at AT&T Labs in Florham Park, New Jersey.

Though Shor's procedure is certainly the most famous quantum algorithm to date, there have been a number of other similar speedups. For example, in the field of computational chemistry, one of the most fundamental calculations is the determination of the thermal rate constant. In fact, some have suggested that the rate constant is "the single most important number characterizing chemical reactions." The rate constant is significant because it reveals how much energy a system must have for a reaction to proceed, as well as how quickly that reaction will take place. Recently, an algorithm has been published (Lidar and Wang, 1999) that computes the rate constant efficiently on a quantum computer. The resulting procedure drastically outperforms any exact classical calculation. A speedup has also been demonstrated for database searches in the field of information science. Searching a database is akin to looking for a forgotten client's telephone number in the phone book in order to find the client's full name. If there were N numbers in the phone book, one would have to flip through half the numbers on average before finding the right one. In 1996, L. K. Grover presented an algorithm that could perform the search in $\sqrt{N}$ steps on average. Although this is not a substantial speedup, it has been proven that the procedure is as fast as is possible, insofar as asymptotic complexity is concerned. Unfortunately, Grover's search algorithm is somewhat odd in that it is randomized, and therefore only gets the answer right about half the time. Its faults notwithstanding, it has the distinction of being the first quantum algorithm actually implemented (on an NMR-QC) that beats the classical analog. Furthermore, Grover's work has the potential to speed up a number of other seemingly unrelated problems.

In review, the power of the quantum computer is not the same across the board. Some problems get a modest speedup, like the search problem, while other problems get a drastic speedup, like factorization. Note, however, that the real power of this new breed of computer is an open avenue of investigation. Some scientists, such as Bennett et al., have argued persuasively that quantum com-

puters cannot put a dent in a very special class of intractable problems called "NP-Complete." If their assertion is in error, and NP-Complete problems are susceptible to quantum speedups, a vast array of very important problems could be solved efficiently. Bennett states that while his paper conclusively rules out the most straightforward approaches, it cannot make the categorical statement that *no* approach is possible. In truth, no one can yet say with certainty where the boundaries of complexity ought to fall. It does seem to be the case, however, that the realm of uncomputable problems is far beyond even the capacities of the quantum computer.

More Than a Speed Demon

In addition to their ability to speed up calculations, quantum computers bring much more to the table. Another significant feature they have to offer is error correction. This is important if quantum machines are to be able to communicate with one another, since every communication channel has some degree of unwanted noise. This is a well-established principle from classical communications, in which computer modems constantly check for errors that are caused by the noisy "static" on the phone lines. Also, if information is to be stored in any medium, there will necessarily be errors that arise and must be suppressed. These sources of error have been so thoroughly probed in the classical realm, it is currently unclear whether quantum algorithms will prove superior. In one sense, the new algorithms are inferior, in that up to nine qubits may need to be stored and updated for every qubit of data that is to be guarded from error. This requires much more storage than the classical algorithms use. There is, therefore, another very significant reason why these new forms of error correction are vital. Due to the sensitive state needed to create parallelism, quantum computers are highly susceptible both to minor flaws in their implementation and to undesirable interaction with the outside world. Both of these difficulties will be discussed later, but suffice it to say that without error-correcting codes, quantum computers could not do basic multiplication, let alone anything more complicated.

Finally, given the significant influence of cryptography in this budding science, any discussion would be remiss to exclude it. Equally noteworthy is the fact that many of these remarkable security protocols can be implemented with current technology. In 1995, H. Zbinden and his associates at the University of Geneva were able to use laser pulses to transmit qubits in a secure fashion. The pulses were sent across 23 kilometers of standard telecom fiber optics under Lake Geneva. The error rate, around three percent, was low enough to establish the viability of the protocol. Considering that an eavesdropper would be likely to introduce errors in approximately 25 percent of the qubits, the demonstrated error rate was sufficient to guarantee the privacy of the channel. Further enhancements with error-correcting codes would make the data all the more difficult to tamper with or intercept. Zbinden's experiment highlights an important advantage that quantum cryptography has over classical models: because qubits are changed when they are measured and cannot be cloned, a wiretapper cannot simply intercept them midstream without being noticed. However, as was demonstrated by C. A. Fuchs et al. in 1997, an eavesdropper can potentially take advantage of entanglement to glean partial information from a "secure" conversation. In conclusion, even though quantum cryptography is not yet foolproof, it promises to provide much greater security than any existing classical protocol.

Fallen Soufflés and Other Maladies

With all of these exciting new capabilities, one might expect to find quantum computers on the shelves sometime soon. However, there are a number of technical difficulties that some scientists think may never be resolved. Almost always, an underlying theory makes some assumptions that are very difficult to implement in practice. For example, the idealized quantum computer would have no internal flaws and no interaction with its environment. In reality, though, such complicating factors are always present, and they lead to the disruptive phenomenon called decoherence. Recall that in order for the computer to work properly, all the qubits have to be able to interfere in just the right way. Unfortunately, little flaws in the system upset the process (technically speaking, the system becomes "out of phase"). In addition, an even greater problem is that the system loses energy, and hence information, to its surroundings. These are no minor difficulties—information is lost 10 million times too fast to allow for the factorization of a 130-digit number! In that particular instance, it may well be easier to wait for classical computers to get faster than to try to compensate for such loss. Serge Haroche and Jean-Michel Raimond, two of the most outspoken pessimists about quantum computation, write that "the fundamental phenomenon of quantum decoherence, whose probability increases exponentially [i.e., very quickly] with the system size, will make it impossible to 'push back' … the quantum/classical boundary." Early experiments at least confirm the difficulty of the task: the ratio of speed to decoherence needs to be around one billion, in place of its empirical value of about 10. At this point, scientists are split; some believe that error correction will save the day, while others conclude that it would only make an unstable system all the more unwieldy.

How many customers would buy a calculator that couldn't be interrupted while it was working, failed to announce when it was done, and only got the right answer 50 percent of the time?

Sadly, decoherence is not the only substantial problem. There is another wrench in the works, one that might be called the problem of the "quantum soufflé." In today's electronic computers, one could (carefully) probe around in all sorts of circuits and measure voltages at a whim. However, quantum machines find that kind of prodding very rude, and they will refuse to give an answer. This is because testing the qubits collapses them into a single state, and the parallelism needed to solve the problem is lost in an instant. The tendency of the "quantum soufflé" to collapse is only half the explanation for its name. Everyone who has baked a soufflé (or at least seen Martha Stewart do so on television) knows that the oven needs to be set at just the right temperature and that the haute cuisine must be removed at precisely the right time if the final product is to be edible. It turns out that quantum computers are finicky in a similar fashion. Consider Grover's search algorithm, the one that had a 50/50 chance of coming up with the right answer after around $\sqrt{N}$ iterations. Of course, running through the procedure a few more times should give an even more accurate answer, right? Unfortunately, this is much like the temptation to crank the oven up a few degrees—it seems to make sense but doesn't help in the end. The probability of getting the right answer actually drops precipitously over the next few trials. The greatest difficulty in getting a quantum computer to market might well lie in writing the owner's manual.

The enthusiast would probably ask at this point, "Isn't it worth bearing with all these quirks to get a blazing fast computer?" The answer: not necessarily. It is important to realize that these speedups usually only outclass the classical computer on very large problems that require thousands of qubits and billions of logic gates. To make matters worse, it has been demonstrated that there are some problems that don't get any speedup from running on a quantum machine. As difficult as it is to build and operate a quantum computer, scientists would prefer to exploit alternatives whenever possible. After all, how many customers would buy a calculator that couldn't be interrupted while it was working, failed to announce when it was done, and only got the right answer 50 percent of the time?

Tomorrow and Beyond

Yes, there are a number of hurdles on the path to a large-scale quantum computer. However, this is to be expected in a field that has had most of its important questions posed within the last three or four years. Certainly, many of the questions are waiting to be asked in this realm of half magic, half science. At least for the foreseeable future, it appears that everyday silicon-and-wire computers will remain the standard. This conclusion is left tentative in hopes of avoiding the mistake of IBM chairman Thomas Watson, who forecast in 1943 that there would be "a world market for maybe five computers." After all, this nascent technology is already beginning to settle into its niche, poised to conquer problems previously thought to be invincible. Dawn has broken for the quantum computer, and it promises to be an exciting day.

A Limerick by Peter Shor

If the computers that you build are quantum,
Then spies everywhere will all want 'em.
Our codes will all fail,
And they'll read our e-mail,
Till we get crypto that's quantum,
  and daunt 'em. □

The Quantum Computer—An Introduction

by Jacob West

West (left) zips through a gnarly prime factorization problem with his quantum computer while Beauchamp (right) wrestles with his balky PC. Well, maybe someday…

What is a Quantum Computer?

Behold your computer. Your computer represents the culmination of years of technological advancements beginning with the early ideas of Charles Babbage (1791–1871) and the eventual creation of the first computer by German engineer Konrad Zuse in 1941. Surprisingly, however, the high-speed modern computer sitting in front of you is fundamentally the same as its gargantuan 30-ton ancestors, which were equipped with some 18,000 vacuum tubes and 500 miles of wiring! Although computers have become more compact and considerably faster in performing their task, the task remains the same: to manipulate and interpret an encoding of binary bits into a useful computational result. A bit is a fundamental unit of information, classically represented as a 0 or 1 in your digital computer. Each classical bit is physically realized through a macroscopic physical system, such as the magnetization on a hard disk or the charge on a capacitor. A document, for example, comprised of n characters stored on the hard drive of a typical computer is accordingly described by a string of $8n$ zeros and ones. Herein lies a key difference between your classical computer and a quantum computer. Where a classical computer obeys the well-understood laws of classical physics, a quantum computer is a device that harnesses physical phenomena unique to quantum mechanics (especially quantum interference) to realize a fundamentally new mode of information processing.

In a quantum computer, the fundamental unit of information (called a quantum bit, or qubit), is not binary but rather more quaternary in nature. This qubit property arises as a direct consequence of its adherence to the laws of quantum mechanics, which differ radically from the laws of classical physics. A qubit can exist not only in a state corresponding to the logical state 0 or 1 as in a classical bit, but also in states corresponding to a blend or superposition of these classical states. In other words, a qubit can exist as a zero, a one, or simultaneously as both 0 and 1, with a numerical coefficient representing the probability for each state. This may seem counterintuitive, because everyday phenomena are governed by classical physics, not quantum mechanics—which takes over at the atomic level. This rather difficult concept is perhaps best explained through an experiment. Consider the figures on the opposite page: In an experiment like that in figure a, where a photon is fired at a half-silvered mirror, it can be shown that the photon does not actually split by verifying that if one detector registers a signal, then no other detector does. With this piece of information, one might think that any given photon travels either vertically or horizontally, randomly choosing between the two paths. However, quantum mechanics predicts that the photon actually travels both paths simultaneously, collapsing down to one path only upon measurement. This effect, known as single-particle interference, can be better illustrated in a slightly more elaborate experiment, outlined in figure b. Figure b depicts an interesting experiment that demonstrates the phenomenon of single-particle interference. In this case, experiment shows that the photon *always* reaches detector A, *never* detector B! If

a single photon travels vertically and strikes the mirror, then, by comparison to the experiment in figure a, there should be an equal probability that the photon will strike either detector A *or* detector B. The same goes for a photon traveling down the horizontal path. However, the actual result is drastically different. The only conceivable conclusion is therefore that the photon somehow traveled both paths simultaneously, creating an interference at the point of intersection that destroyed the possibility of the signal reaching B. This is known as quantum interference and results from the superposition of the possible photon states, or potential paths. So although only a single photon is emitted, it appears as though an identical photon exists and travels the "path not taken," and is detectable only by the interference it causes with the original photon when their paths come together again. If, for example, either of the paths are blocked with an absorbing screen, then detector B begins registering hits again just as in the first experiment! This unique characteristic, among others, makes the current research in quantum computing not merely a continuation of today's idea of a computer, but rather an entirely new branch of thought. And it is because quantum computers harness these special characteristics that they have the potential to be incredibly powerful computational devices.

The Potential and Power of Quantum Computing

In a traditional computer, information is encoded in a series of bits, and these bits are manipulated via Boolean logic gates arranged in succession to produce an end result. Similarly, a quantum computer manipulates qubits by executing a series of quantum gates, each a unitary transformation acting on a single qubit or pair of qubits. In applying these gates in succession, a quantum computer can perform a complicated unitary transformation to a set of qubits in some initial state. The qubits can then be measured, with this measurement serving as the final computational result. This similarity in calculation between a classical and quantum computer affords that in theory, a classical computer can accurately simulate a quantum computer. In other words, a classical computer should be able to do anything a quantum computer can. So why bother with quantum computers? Although a classical computer can theoretically simulate a quantum computer, it is incredibly inefficient, so much so that a classical computer is effectively incapable of performing many tasks that a quantum computer could perform with ease. The simulation of a quantum computer on a classical one is a computationally hard problem because the correlations

Illustrations from "Quantum Computation" by David Deutsch and Artur Ekert, *Physics World*, March 1998, p. 47. See http://physicsWeb.org/toc/11/3 for related articles.

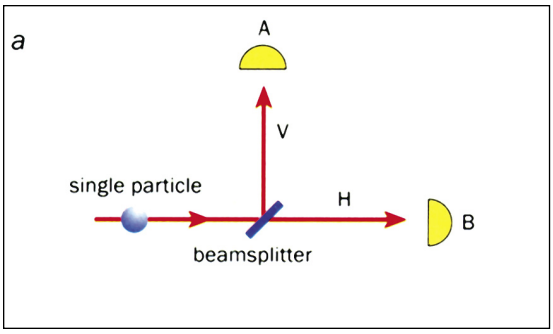

Here a light source emits a photon along a path toward a half-silvered mirror. This mirror splits the light, reflecting half vertically toward detector A and transmitting half toward detector B. A photon, however, is a single quantized packet of light and cannot be split, so it is detected with equal probability at either A or B. Intuition would say that the photon randomly leaves the mirror in either the vertical or horizontal direction. However, quantum mechanics predicts that the photon actually travels both paths simultaneously! This is more clearly demonstrated in figure b.

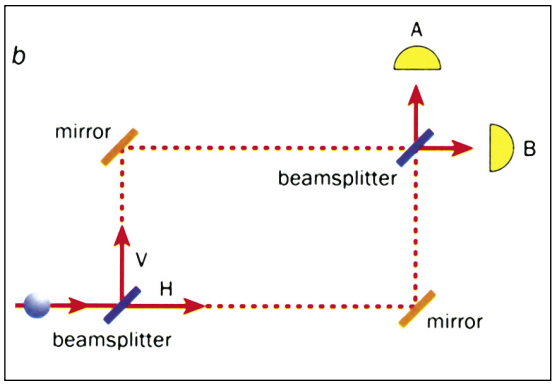

In this experiment, the photon first encounters a half-silvered mirror, then a fully silvered mirror, and finally another half-silvered mirror before reaching a detector; each half-silvered mirror introduces the probability of the photon traveling down one path or the other. Once a photon strikes the mirror along either of the two paths after the first beam splitter, the arrangement is identical to that in figure a, and so one might hypothesize that the photon will reach either detector A or detector B with equal probability. However, experiment shows that in reality this arrangement causes detector A to register 100 percent of the time, and detector B never! How can this be?

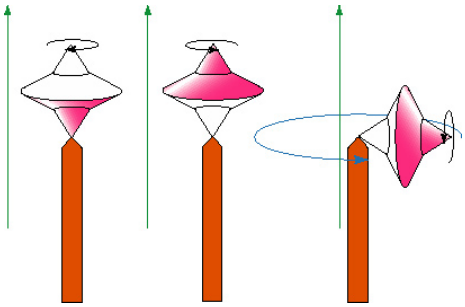

Above: Some atomic nuclei have a magnetic property that spins like a top. The spin axis prefers to align with an external magnetic field (green arrow), as shown at center. But a properly tuned radio pulse can tip the top—a 180-degree pulse (left) will flip it right over. And a 90-degree pulse (right) will knock it perpendicular to the field, causing it to precess like a gyroscope. This is the basis of nuclear magnetic resonance, or NMR. (After "Quantum Computing with Molecules," by Neil Gershenfeld and Isaac L. Chuang, Scientific American, June 1998.)

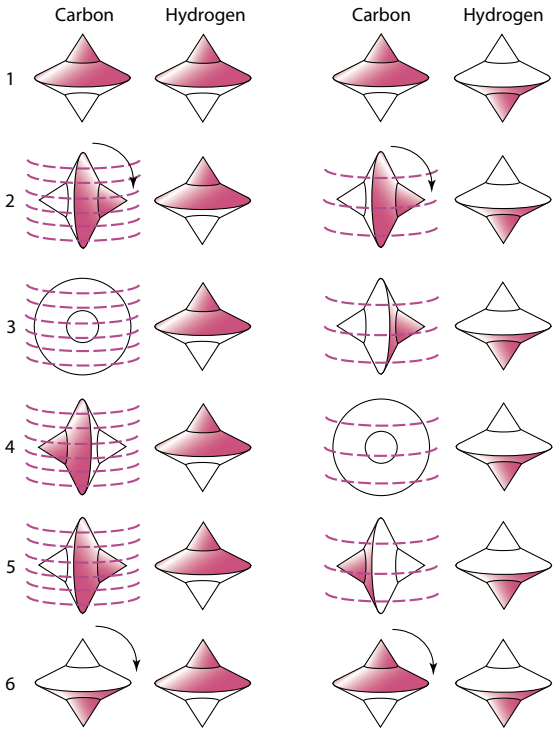

Above: A controlled-NOT gate inverts input A if and only if input B is 1. Gershenfeld and Chuang created a quantum controlled-NOT gate using chloroform molecules in an NMR machine. 1) The chloroform molecule contains a carbon-13 atom (input A) bound to a hydrogen atom (input B). 2) A 90-degree radio pulse tips both carbon nuclei perpendicular to the magnetic field (not shown). 3–5) The carbon nucleus precesses rapidly if the hydrogen nucleus is in state 1 (left), but more slowly if the hydrogen is in state 0 (right). 6) Applying another 90-degree pulse at just the right delay time inverts the carbon (left) or returns it to its original orientation (right).

among quantum bits are qualitatively different from correlations among classical bits, as first explained by John Bell. Take for example a system of only a few hundred qubits. This exists in a Hilbert space of approximately 10^{90} dimensions, which in simulation would require a classical computer to work with exponentially large matrices (to perform calculations on each individual state, which is also represented as a matrix), meaning it would take an exponentially longer time than even a primitive quantum computer.

Richard Feynman was among the first to recognize the potential in quantum superposition for solving such problems much much faster. For example, a system of 500 qubits, which is impossible to simulate classically, represents a quantum superposition of as many as 2^{500} states. Each state would be classically equivalent to a single list of 500 1's and 0's. Any quantum operation on that system—a particular pulse of radio waves, for instance, whose action might be to execute a controlled-NOT operation on the 100th and 101st qubits—would simultaneously operate on all 2^{500} states. Hence—with one fell swoop, one tick of the computer clock—a quantum operation could compute not just on one machine state, as serial computers do, but on 2^{500} machine states at once! Eventually, however, observing the system would cause it to collapse into a single quantum state corresponding to a single answer, a single list of 500 1's and 0's, as dictated by the measurement axiom of quantum mechanics. The reason this is an exciting result is because this answer, derived from the massive quantum parallelism achieved through superposition, is the equivalent of performing the same operation on a classical supercomputer with approximately 10^{150} separate processors (which is of course impossible)!

Early investigators in this field were naturally excited by the potential of such immense computing power, and soon the hunt was on to find something interesting for a quantum computer to do. Peter Shor, a research and computer scientist at AT&T Laboratories in New Jersey, provided such an application by devising the first quantum computer algorithm. Shor's algorithm harnesses the power of quantum superposition to rapidly factor very large numbers (on the order of 10^{200} digits and greater) in a matter of seconds. The premier application of a quantum computer capable of implementing this algorithm lies in the field of encryption, where one common (and best) encryption code, known as RSA, relies heavily on the difficulty of factoring very large composite numbers into their primes. A computer that could do this easily would naturally be of great interest to numerous government agencies that use RSA—previously considered to be "uncrackable"—and to anyone interested in electronic and financial privacy.

Encryption, however, is only one application of a quantum computer. In addition, Shor has

put together a toolbox of mathematical operations that can only be performed on a quantum computer, many of which he used in his factorization algorithm. Furthermore, Feynman asserted that a quantum computer could function as a kind of simulator for quantum physics, potentially opening the doors to many discoveries in that field. Currently the power and capability of a quantum computer is primarily theoretical speculation; the advent of the first fully functional quantum computer will undoubtedly bring many new and exciting applications.

A Brief History of Quantum Computing

The idea of a computational device based on quantum mechanics was first explored in the 1970s and early 1980s by physicists and computer scientists such as Charles Bennett of the IBM Thomas J. Watson Research Center, Paul Benioff of Argonne National Laboratory in Illinois, David Deutsch of the University of Oxford, and Feynman. The idea emerged when scientists were pondering the fundamental limits of computation. They understood that if technology continues to abide by Moore's Law, then the continually shrinking size of circuitry packed onto silicon chips will eventually reach a point where individual elements will be no larger than a few atoms. Here a problem arises, because at the atomic scale the physical laws that govern the behavior and properties of the circuit are inherently quantum mechanical in nature, not classical. This then raised the question of whether a new kind of computer could be devised based on the principles of quantum physics.

Feynman was among the first to attempt to provide an answer to this question by producing an abstract model in 1982 that showed how a quantum system could be used to do computations. He also explained how such a machine would be able to act as a simulator for quantum physics. In other words, a physicist would have the ability to carry out experiments in quantum physics inside a quantum-mechanical computer.

Later, in 1985, Deutsch realized that Feynman's assertion could eventually lead to a general-purpose quantum computer and published a crucial theoretical paper showing that any physical process, in principle, could be modeled perfectly by a quantum computer. Thus, a quantum computer would have capabilities far beyond those of any traditional classical computer. After Deutsch published this paper, the search began for interesting applications for such a machine.

Unfortunately, all that could be found were a few rather contrived mathematical problems, until Shor circulated in 1994 a preprint of a paper in which he set out a method for using quantum computers to crack an important problem in number theory, namely factorization. He showed how an ensemble of mathematical operations, designed specifically for a quantum computer, could be organized to enable such a machine to factor huge numbers extremely rapidly, much faster than is possible on conventional computers. With this breakthrough, quantum computing transformed from a mere academic curiosity directly into a national and world interest.

Obstacles and Research

The field of quantum information processing has made numerous promising advancements since its conception, including the building of two- and three-qubit quantum computers capable of some simple arithmetic and data sorting. However, a few potentially large obstacles still remain that prevent us from "just building one" or, more precisely, building a quantum computer that can rival today's modern digital computer. Among these difficulties, error correction, decoherence, and hardware architecture are probably the most formidable. Error correction is rather self-explanatory, but what errors need correction? The answer is primarily those errors that arise as a direct result of decoherence, or the tendency of a quantum computer to decay from a given quantum state into an incoherent state as it interacts, or entangles, with the state of the environment. These interactions between the environment and qubits are unavoidable, and induce the breakdown of information stored in the quantum computer, and thus errors in computation. Before any quantum computer will be capable of solving hard problems, research must devise a way to maintain decoherence and other potential sources of error at an acceptable level. Thanks to the theory (and now reality) of quantum error correction, first proposed in 1995 and continually developed since, small scale quantum computers have been built and the prospects of large quantum computers are looking up. Probably the most important idea in this field is the monitoring of phase coherence for error correction as a means to extract information and reduce error in a quantum system without actually measuring that system. In 1998, researchers at Los Alamos National Laboratory and MIT led by Raymond Laflamme managed to spread a single bit of quantum information (qubit) across three nuclear spins in each molecule of a liquid solution of molecules of alanine or trichloroethylene. They accomplished this using the techniques of nuclear magnetic resonance (NMR). This experiment is significant because spreading out the information actually made it harder to corrupt. Quantum mechanics tells us that directly measuring the state of a qubit invariably destroys the superposition of states in which it exists, forcing it to become either a 0 or 1. The technique of spreading out the information allows researchers to utilize the property of entanglement to study the interactions between states as an

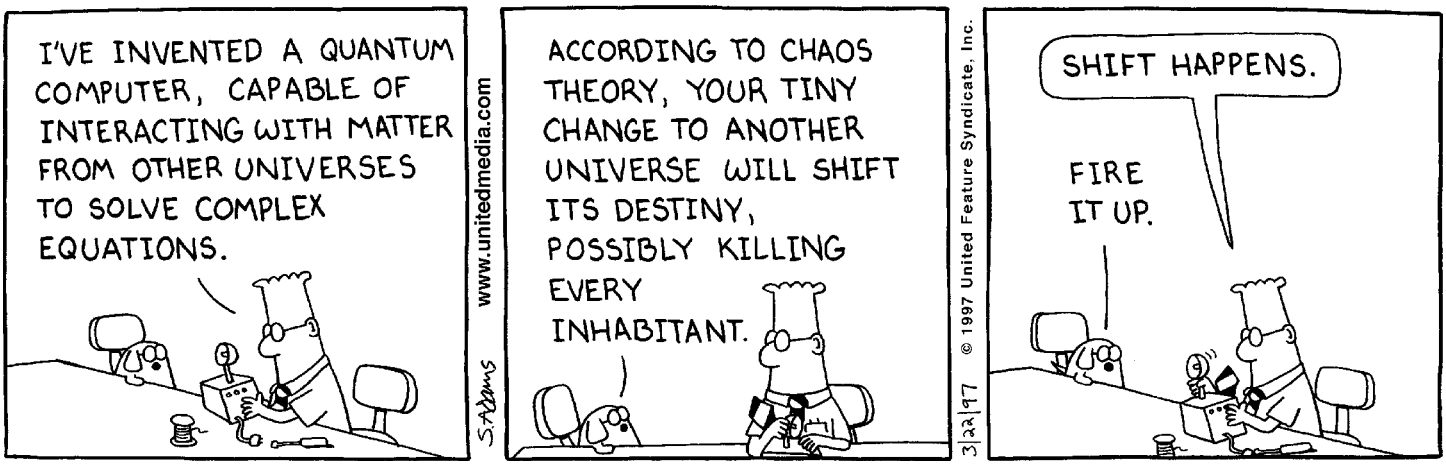

DILBERT reprinted by permission of United Feature Syndicate, Inc.

indirect method for analyzing the quantum information. Rather than a direct measurement, the group compared the spins to see if any new differences arose between them, without learning anything about the information itself. This technique gave them the ability to detect and fix errors in a qubit's phase coherence, and thus to maintain a higher level of coherence in the quantum system. This milestone has provided ammunition against skeptics and hope for believers. Currently, research in quantum error correction continues, with groups at Caltech (Preskill, Kimble), Microsoft, Los Alamos, and elsewhere.

At this point, only a few of the benefits of quantum computation and quantum computers are readily obvious, but before more possibilities are uncovered, theory must be put to the test. In order to do this, devices capable of quantum computation must be constructed. Quantum computing hardware is, however, still in its infancy. As a result of several significant experiments, NMR has become the most popular component in quantum hardware architecture. Only within the past year, a group from Los Alamos National Laboratory and MIT constructed the first experimental demonstrations of a quantum computer using NMR technology. Currently, research is under way to discover methods for battling the destructive effects of decoherence, to develop an optimal hardware architecture for designing and building a quantum computer, and to further uncover quantum algorithms to utilize the immense computing power available in these devices. Naturally this pursuit is intimately related to quantum error correction codes and quantum algorithms, so a number of groups are doing simultaneous research in a number of these fields. To date, designs have involved ion traps, cavity quantum electrodynamics (QED), and NMR. Though these devices have had mild success in performing interesting experiments, the technologies each have serious limitations. Ion-trap computers are limited in speed by the vibration frequency of the modes in the trap. NMR devices have an exponential attenuation of signal to noise as the number of qubits in a system increases. Cavity QED is slightly more promising; however, it still has only been demonstrated with a few qubits. Seth Lloyd of MIT is currently a prominent researcher in quantum hardware. The future of quantum computer hardware architecture is likely to be very different from what we know today; however, the current research has helped to provide insight as to what obstacles the future will hold for these devices.

Future Outlook

At present, quantum computers and quantum information technology remain in their pioneering stage. At this very moment obstacles are being surmounted that will provide the knowledge needed to thrust quantum computers up to their rightful position as the fastest computational machines in existence. Error correction has made promising progress to date, nearing a point now where we may have the tools required to build a computer robust enough to adequately withstand the effects of decoherence. Quantum hardware, on the other hand, remains an emerging field, but the work done thus far suggests that it will only be a matter of time before we have devices large enough to test Shor's and other quantum algorithms. Thereby, quantum computers will emerge as the superior computational devices at the very least, and perhaps one day make today's computers obsolete. Quantum computation has its origins in highly specialized fields of theoretical physics, but its future undoubtedly lies in the profound effects it will have on the lives of all humankind. □

Quantum computing at Caltech

While Gershenfeld and Chuang are tinkering with magnets, some folks at Caltech are playing with light. In this approach, photons carry information and atoms store it. All you need to do is design a gate that allows them to interact.

Valentine Professor and Professor of Physics Jeff Kimble has taken the first step in that direction. Kimble has been in the quantum-optics biz for over 20 years—see *E&S* Summer '93. This past February, his lab and collaborators in New Zealand successfully trapped a cesium atom, suspending it in a weak laser field in an "optical resonator"—a pair of mirrors, 10 microns apart, that are so highly reflective that a photon will bounce back and forth hundreds of thousands of times before escaping. The atom and the resonator share a quantum of excitation and could act as a gate.

Meanwhile, Professor of Theoretical Physics John Preskill has been thinking about error correction. In 1996, his grad student Daniel Gottesman (PhD '97) developed a systematic method for deriving quantum codes that could be used for fault-tolerant computation. Now Preskill is trying to design quantum fault-tolerance into the hardware. After all, that's what your hard disk does—the data is encoded in puddles of magnetic field that either point straight up or straight down. Oh, sure, an individual atom in the puddle might get zapped by a stray cosmic ray and flip its field the wrong way, but peer pressure from the surrounding atoms soon pushes it back into alignment. But qubits can "point" in any direction, and their errors are just wobbles of a degree or two. Fortunately, if you share the encoded information among many qubits, you only have to worry about errors that jiggle *all* of the qubits in exactly the same way.

One scheme Preskill is exploring exploits the Aharonov-Bohm effect, which is seen in an electron orbiting around a donut-shaped magnetic coil. As the electron moves, its wave function acquires a phase that depends only on the number of times per orbit that its path goes through the donut's hole. "It can take any path," Preskill explains. "As long as the number of windings is the same, the way the wave function changes is the same. So you use the particle's trajectory to store information that will be well protected." And unlike most things quantum, the bigger the system gets, the less likely it is to decohere. "You can pound on it with a hammer—bang! bang! bang!—and inflict a lot of local damage, but you can't damage nonlocal information unless many hammers conspire together. And the environment isn't smart enough to do that." An analogous optical system could be developed, he says.

How many qubits can happily coexist in one gate is not yet clear, but a real quantum computer will probably need an array of gates that will have to share information. Kimble's current setup consists of a forest of prisms, mirrors, beam splitters, and what have you that takes up about 50 square feet of benchtop. (And standard lab-model NMRs use powerful magnets that are bigger than washing machines and weigh over half a ton—not the sort of thing you'd want near your credit cards—and about another 1,000 pounds of radio-field generators and sundry electronic gear. Then there's that vial of funky liquid that they won't let you take on an airplane.) If quantum computing is ever going to go commercial, the apparatus clearly needs to become a lot more manageable.

So Assistant Professor of Physics Hideo Mabuchi (PhD '98), a former grad student of Kimble's, is beginning a collaboration with Professor of Physics Michael Roukes and Professor of Electrical Engineering, Applied Physics, and Physics Axel Scherer to build miniaturized solid-state optical systems. Roukes and Scherer are nanofabrication experts—makers of teeny-tiny machinery on computer chips. One of Roukes's specialties is micromagnets, and last year Scherer's lab, in collaboration with Summerfield Professor of Applied Physics Amnon Yariv and a group at USC, created a chip with an array of the world's smallest lasers, using quantum wells as light sources. The light is confined to an optical resonator that consists of a hexagonal array of tiny, carefully spaced holes drilled through a layer of atoms half a wavelength thick. The beam eventually emerges perpendicularly to the chip's surface, allowing optical communication with other components. But the lasing atom is embedded within the crystal, so any quantum entanglements would quickly decohere via the neighboring atoms. So the collaboration plans to drill a cavity in the center of Scherer's resonator. Then Roukes will lay down a couple of loops of nanowire that will electromagnetically trap a cesium atom in the cavity. It's Mabuchi's job to figure out how to entice the atom into the trap, and then verify that it's in there. Says Scherer, "Of all the approaches people are taking to create entangled states, this one, as ludicrous as it may seem, is probably the sanest. At least all the pieces work." Says Mabuchi, "One of the nice things about working with Axel and Mike is it gives us an understanding of how these devices were meant to be miniaturized and manufactured in the real world."

Kimble, Mabuchi, Preskill, Roukes, and Scherer have just launched a three- to five-year project funded by the Department of Defense's Multidisciplinary Component of the University Research Initiative (MURI). Their goal is to demonstrate quantum error-corrected communication over a 100-kilometer distance, incidentally developing technology that could later be used for quantum computing. We're still a long way from running Peter Shor's algorithm—just factoring 15 into 3×5 would require about 4,000 operations on four qubits, and it's anybody's guess how much effort it will take to get the system to hang together that long. But hey—it's a start. —*DS*