

Quantum Information and Quantum Computation

John Preskill
15 January 1997

Abstract

Information is something that can be encoded in the state of a physical system, and a *computation* is a task that can be performed with a physically realizable device. Therefore, since the physical world is fundamentally quantum mechanical, the foundations of information theory and computer science should be sought in quantum physics. In fact, quantum information has weird properties that contrast sharply with the familiar properties of *classical* information. A quantum computer – a new type of machine that exploits the quantum properties of information – could perform certain types of calculations far more efficiently than any foreseeable classical computer. To build a functional quantum computer will be an enormous technical challenge. As a first step, new methods for *quantum error correction* are being developed that can help to prevent a quantum computer from crashing.

I suspect that many of you here tonight are aware that this is 1997. We are approaching the dawn of a new millennium. [2001 theme.] And as we look back at the 20th century, one of the most notable achievements of our civilization has been the development of our information technology. I have a notebook computer up here with me tonight, and I think it's a pretty hot machine. But surely by, say, the end of the 21st century, the information technology that impresses us so much today will have been far surpassed by new technology that we cannot even imagine today. Even so, I intend to speculate tonight about the future of information technology.

This sort of projection of the future of technology is a task fraught with danger. Here is one cautionary tale. This past year, we celebrated an important anniversary, the 50th anniversary of the ENIAC, which many people regard as the first electronic digital computer. It is interesting to see what people were saying back in the 40's about the future of electronic computing. Here is a quote from Popular Mechanics that appeared in 1949: "Where a calculator on the ENIAC is equipped with 18,000 vacuum tubes and weighs 30 tons, computers in the future may have only 1,000 vacuum tubes and perhaps only weigh one and a half tons." Well, in fact, the computing power of the ENIAC is roughly equivalent to what is in a digital watch, and we have had digital watches since the 70's. So the visionary who said this evidently was not thinking big enough, or small enough.

So no one can accurately predict the future of technology; that's a given. And aside from that, I am particularly ill equipped for this task. I am a theoretical physicist, not an engineer, and I am sure that there are many people in this auditorium who know more than I do about how computers work. But when I look back at the 20th century, I know without hesitating that the crowning intellectual achievement of the century has been the

discovery of the quantum theory. So it is natural for me to wonder how the discoveries concerning quantum theory in the 20th century will impact the technology of the 21st century.

Where information is concerned, a physicist knows that, whatever information might be, it is something that can be encoded and stored in the state of some physical system, like the pages of a book, or the sectors of a hard disk. But we also know that all physical systems are fundamentally quantum mechanical systems. So information is something that can be encoded in a quantum state. The question we will address, our theme for the evening, is, "Can the computers of the future better exploit the quantum properties of information, to perform tasks that are beyond what can conceivably be achieved with conventional silicon-based information technology?"

To get started, we'll need to recall some basic facts about information. Probably many of you have played the game 20 questions. It's fun. One player thinks of an object, and the other player has to guess what it is. That player asks questions to which the answer is either yes or no. Well, it has to be a clear question...Anyway, if you ask the right questions, you can figure out what the object is, just from the yes or no answers. A highly skilled player, by asking 20 questions, could in principle distinguish about 1,000,000 different objects. And if we are willing to allow more questions, in principle any number of objects could be distinguished. So we say that any amount of information can be encoded in the yes/no answers. All information can be reduced to elementary units, what we call bits. Each bit is a yes or a no, or sometimes we represent it as the number 0 or the number 1.

I like to visualize a bit as an object, let's say a ball, that can be either one of two colors, let's say either red or green. Bits are valuable, so we should take good care of them. So we get a bit box. We can open the lid of the box, seal the ball up inside for safekeeping, and then if we open the box later on, the ball pops out. Whatever color ball we put in the box is what comes out, so we can recover our bit and read it.

But in quantum theory, the elementary unit of information is something rather different from the classical bit — I'll call it a quantum bit, or a "qubit" for short. We may think of a quantum bit as a box with a ball stored inside, but in this case, we can open the box through either one of two doors, door 1 or door 2. To an experimental physicist, the two doors correspond to two different ways to measure the quantum state of an atom, or of a particle of light, but let's not worry about that, we'll just think of it as a box with two doors.

If we put a red or a green ball into the box, either through door 1 or door 2, close the door, and then open the same door again, then the color of the ball that comes out is the same as the color that we put in; it is just like a classical bit. But suppose that we put the ball into door number 1, and then we open door number 2. Then the color of the ball that comes out doesn't have anything to do with what we put in, the color is completely random. 50% of the time it will be red and 50% of the time it will be green. So if we open the wrong door, we can't read the information that was put into the box. Let's try it again. This time we put the balls into door number 2, and if we reopen door number 2, the color of the ball that comes out is the same as what we put in. But when we put the

balls into door number 2 and then open door number 1, the color of the ball that comes out is random; half the time it is what we put in, and half the time it is not.

So we have seen one way in which quantum information, information encoded in qubits, is different than classical information, information encoded in bits. You can't read quantum information unless you open the right door. But there is a more interesting difference between classical and quantum information, and to appreciate it, we'll need to suppose that we have two boxes. The boxes can be far apart. One of them is here at Caltech in Pasadena, and the other is in the custody of a friend of mine who lives in the Andromeda galaxy. These boxes have some peculiar properties. First of all, if I open my box here in Pasadena, either through door 1 or door 2, the color of the ball that comes out is completely random. And the same is true for my friend in Andromeda. So when either one of us opens his box, through either door, we don't get any information; we don't find out anything about what is inside. But that's funny . . . we have two boxes so we should have been able to store two bits of information. How is that information encoded? Where is it hiding?

The answer is that the information is contained in correlations between what happens when I open a box in Pasadena and what happens when my friend opens a box in Andromeda. If I open my box through door number 1, I might find a red ball or I might find a green ball. But if I find a green ball, then if my friend also opens door number 1 or his box, he finds a green ball, too. And if I find a red ball, he always finds a red ball. Same thing if I open door number 2 and he opens door number 2 — we are guaranteed to find balls of the same color. So what he finds is perfectly correlated with what I find if we open the same door. Well, there are several different ways in which what happens when we open a box in Pasadena can be correlated with what happens when we open a box in Andromeda, and we have chosen one of those ways — that's information. But there is no way to get access to any of the information, no way to read it, just by making observations in Pasadena or Andromeda. Instead, the information is spread out in a very nonlocal way, shared equally, in a sense, between the box in Pasadena and the box in Andromeda. This property of quantum information, that it can be encoded in this nonlocal way, is what we call quantum entanglement. It is the crucial way in which quantum information is different than classical information, and it is what underlies much of what I want to talk about tonight.

[Very interesting.] Yes, thank you. But not everyone is so impressed by these correlations. Correlations are not really so exotic; we encounter them all the time in everyday life. For example, I have a friend who, on any given day, decides at random to wear either red socks or green socks. (You may think my friend is rather eccentric, but among physicists he is considered quite normal.) Anyway, I can trust my friend to always wear two socks of the same color; he is very fastidious about that. So his socks are perfectly correlated. That means that, as soon as I see one of his feet, and notice that he is wearing a red sock on that foot, I know for sure before I even look that there is a red sock on the other foot. And if I see a green sock on one foot, I know for sure that there is a green sock on the other foot, even without looking. So correlations are not at all unusual. On the one hand my friend always wears two socks of the same color, and on the other hand when we open

our two quantum boxes, each through the same door, we always find two balls of the same color. Is there really any fundamental difference between these two things? Aren't the boxes just like the soxes?

In fact, I want to argue that there is a profound difference between the boxes and the soxes. To explain why, it will be helpful to consider an even more peculiar friend of mine, one with three feet. This fellow also decides at random every day how to wear his socks, but he always wears an even number of red socks (either 0 or 2) and an odd number of green socks (1 or 3). I know him well, and I trust my friend; he never wears one red sock and never wears three red socks. That means that once I have seen two of his feet, I know with certainty, before I even look, what color sock is on the third foot. If I see a red sock and a green sock, the third sock must be red. If I see two green socks or two red socks, the third sock must be green.

Now I want to consider an analogous situation with quantum boxes instead of socks. We have three boxes. And suppose that we decide to open door number 2 of one of the boxes, and door number 1 of the other two boxes. Every time we try this, we find that the number of red balls is even (0 or 2); we never find one red ball or three red balls. I have tried this a million times, so I am sure that it's true. Trust me. You don't trust me? Okay, I'll show you ... Here I am opening door 2 on the third box, and door 1 on the other boxes. And we find two red balls. Now I'm opening door 2 on the second box and door 1 on the others. This time 0 red balls; all the balls are green. This time I'm opening door 2 on the first box ... two red balls. Door 2 on the second box ... two red balls. See, it always works.

This is interesting, because suppose that I want to know what will happen when I open the third box, either through door 1 or door 2. I can find out ahead of time, before I open that box, by opening the first two boxes. Let's say I am interested in knowing what will happen when I open door number 1 on the third box. I can open door number 1 of the first box and door 2 of the second box. I found one green ball and one red ball. So now I know for sure what will happen when I open door 1 of the third box. What will happen? Shout it out. Okay, let's see ... yep, it's red, like we thought. Now this time I would like to know what will happen when I open door number 2 on the third box. So I open door number 1 on the first two boxes. I found two green balls. So that means that when we open door number 2 of the third box we'll find ... what? Yes, it works, it is green.

It could be that these boxes are very far apart. Maybe one is in New York, one in Chicago, and one in Pasadena. Well, it seems obvious that opening boxes in New York and Chicago cannot have any influence on what happens when we open a box in Pasadena. So a reasonable person would say that the boxes are like socks. When I look at two of my friends feet, and see a red sock and a green sock, I don't think that I made his third sock turn red by looking at the first two socks. I just think that I found enough about what socks he is wearing today to know that the third sock is red. He had a red sock on that foot all along, but I didn't know it until I looked at the first two feet. So naturally, we assume that it is the same for the boxes. Opening the first two boxes did not change anything inside the third box; it just gave us enough information to figure out what was in the third box all along.

All right, now let's try something new. What will happen if we open door number 2 on all three boxes? We haven't tried this before, so we don't really know. But let's try to use some theory. Let's see if we can make a prediction about what will happen before we open the boxes. That will make the experiment more fun.

Well, how are we going to make a prediction? I'll reason this way ... Another thing we haven't tried is opening door number 1 on all three boxes, so I don't know what would happen if we did that. But let's make an assumption. Let's suppose that if we open door number 1 on all three boxes, we'll find three red balls. It's just a hypothesis, but let's assume this. Well, if that's the case, I can tell you for sure what will happen if we open door number 2 of any of the boxes. You remember, we know for sure that if we open door number 2 on one box, and door number 1 on the other two, we always find an even number of red balls. There is no doubt about that; we have checked it a million times. So, if we make our assumption about what happens when we open door 1 on all the boxes, what can we say about what happens when we open door 2 on, say, the first box? Well, since we find two red balls when we open door 1 of the second and third boxes, we know for sure that we'll have to find a green ball when we open door 2 of the first box. Similarly, we have to find a green ball if we open door number 2 of the second box, because we find two red balls when we open door 1 on the first and third boxes. By the same argument, we have to find a green ball when we open door 2 of the third box. So from what we already know about the boxes, we can deduce that if we find three red balls when we open door 1 of all three boxes, we must find three green balls when we open door 2 of all three boxes.

Is the reasoning clear to everyone? Look, this is the hardest part of my lecture. Just pay close attention for the next few minutes. I promise that after that it will get easier.

Okay. The only thing is that we don't really know what will happen when we open door number 1 of all three boxes; we just assumed we would find three red balls, and maybe that's not true. But we can make a list of all of the things that we could conceivably find were we to open door 1 on all three boxes. There are altogether eight possibilities, and here is the list. Now, for each of these eight possibilities, we can use the same reasoning that I just described to infer what will happen if we open door number 2 on all the boxes. I won't go through the reasoning for every case, but it is a good exercise; you can do it yourself when you get home tonight, if you're not too sleepy. And we find something interesting. For every one of the eight possibilities for what could happen when we open door number 1 on all the boxes, we conclude that when we open door number 2 on all of them the number of red balls must be even — it can be 0 or 2, but it is never 1 or 3. So, it turns out that we really can make a prediction.

Let's review what we've done. We have three boxes. And we know for sure that if we open door number 2 on one box and open door 1 on the other two, we always find an even number of red balls. No doubt about it; we've checked it a million times. Now have been able to use simple logic to infer what must happen if we open door number 2 on all three boxes. We concluded that we must find an even number of red balls. I'm anxious to find out if the theory works. Are you all ready to do the experiment? Okay, here we go!

[Drum roll.] One red ball? What the ... something seems to have gone wrong!

[Scream.] Yes, I'm very upset about it too. But .. oh .. I bet I know what it must be; it's experimental error! These experiments with the boxes and balls are very difficult to do. We'll try it again, and this time I'll be extra careful. [Drum roll.] Three red balls? This doesn't make any sense! We better try it once more. [Drum roll.] One red ball! I guess we have to accept that the experiment is correct. Every time we try it we find an odd number of red balls. But I just don't understand ... [Twilight Zone theme.]

If we review our reasoning to try to figure out where we went wrong, there seems to be only one place where we could have tripped up. We made the eminently reasonable assumption that the act of opening a box in New York and opening a box in Chicago could have no influence on the contents of a box in Pasadena. That's certainly how it works for socks. But our experiment has shown that it doesn't work that way for boxes. The boxes are not like the soxes. The experimental evidence has left us with no choice but to conclude that opening two of the boxes doesn't just tell us what was in the third box all along; opening those boxes actually shapes what we will find when we open the third box. Hence there is a strange kind of nonlocality built into the foundations of quantum physics.

Not everyone is happy about this. One who was unhappy was Albert Einstein. He derided these nonlocal correlations between boxes as "spooky action at a distance" (except he said it in German). Einstein argued from this syllogism: A. I am Albert Einstein. B. I do not understand quantum mechanics. Therefore, C. Quantum mechanics is wrong. Actually, it is a strong argument. But we have been living with quantum mechanics for over seventy years now, and we still can't find anything wrong with it. So it seems that it doesn't really matter whether Einstein liked it or not, we have to accept that this weird nonlocality is an essential part of the description of Nature.

The human mind does not seem to be well equipped to grasp this aspect of Nature, and so we speak of the weirdness of quantum theory. Some people think that weirdness is ugly, but I don't really think so. If Nature is weird, so be it, and let's try to get used to it. But we can also go a step further, and see if we can put the weirdness to work. Does quantum weirdness enable us to perform tasks that would be impossible in a less weird world?

As we search for ways to exploit the weird properties of quantum information, a good place to start is to think about copying information. What would it mean to build a quantum copy machine? Well, I suppose it would mean this. We have a quantum box, and I happen to have put a ball inside through door 2. The copier looks at the box and builds a second box. Now if we open both boxes, the original and the copy, through door number 2, we will find balls of the same color. And if I have put a ball in the original box through door number 1, then when I open both the original and the copy through door number 1, I will find balls of the same color. But there is no such machine. The trouble is that to copy what is inside, the copier needs to open the box. But it has no way of knowing whether I put my ball in door 1 or door 2. It might guess right, and open the correct door, and then it can make a good copy. But it guesses wrong and opens the wrong door, it will damage the information that I stored in the box, and it won't be able to copy it faithfully. So quantum information cannot be copied.

[Scream.] Yes, it is disconcerting. Sometimes it is very useful to be able to copy information. For example, I would have been uncomfortable if I could not have made a copy of this electronic presentation. But sometimes it might be a good feature if information cannot be copied. For example, it might be good if we could carry quantum dollar bills — then we would not have to worry about counterfeiters. Well, I don't know what kind of money people will be carrying around in their pockets one hundred years from now; I don't even know if they'll have pockets. But one thing I am sure about. Even at the end of the 21st century, humans are going to want to keep secrets. That we can't copy quantum information might be a good thing, if it enables us to build a quantum telephone that cannot be tapped.

You see, I have a friend named Alice, and Alice is very anxious to place a very private phone call to a man named Bob. But Alice has a very nosy friend named Eve, and Eve is habitually listening in to Alice's calls on the extension. Alice needs to find a way of calling Bob where she can be sure that Eve is not eavesdropping. Obviously, Alice and Bob should use some sort of code. But it happens that Eve, in addition to being very nosy, is also extremely clever. And she likes nothing better than breaking codes — she absolutely revels in code-breaking. Alice is not at all confident that she is smart enough to think up a code that Eve cannot break.

But Alice knows that if she only had a string of random numbers, and she could send those random numbers to Bob, then Alice could use the random number as a key to code her message, and Bob could use the key for decoding. Then if Eve were to intercept the message, it would be completely random junk, and there would be no way for Eve to decode it. That is, there would be no way for Eve to decode the message unless she had the key. So the problem that Alice and Bob need to solve is this. How can Alice send the random key to Bob, and be assured that there is not way for Eve to have intercepted the key?

Let's see how quantum information can be used to solve this problem. Alice first assembles some random bits — balls that are either red or green. Then she gets some boxes, and decides at random to put each ball into either door 1 or door 2 of the corresponding box. She seals the boxes shut, and then sends them to Bob. Bob does not know which door Alice used on each of the boxes. But he decides at random to open each box through door 1 or door 2. Now it will happen by chance that about half the time, Bob will open the same door that Alice used, and in those cases the ball that Bob finds will be of the same color as the ball that Alice put in the box. But about half the time, Bob will open the wrong door, and then he won't find out anything about what Alice put in the box.

But now that Bob has safely recovered and opened the boxes, so it is too late for Eve to do anything about it, Alice can make a public announcement telling which door she used on each of the boxes; it doesn't matter if Eve finds out at this point. With this information, Bob now knows for which boxes he opened the right door, and he can share that information publicly with Alice. Now Bob and Alice share a set of random bits; Alice has successfully sent the key to Bob.

But what about Eve? She will surely try to intercept some of the boxes while they

are being shipped from Alice to Bob. But to find out what is inside, she will need to open the boxes. If she happens to open the right door (the door that Alice used), then she can copy the information, and Bob will never be the wiser. But sometimes she will open the wrong door, and so might change the color of the ball. Now Bob and Alice can conduct a test to see what Eve has been up to. They can publicly compare a small portion of their key, to make sure that Bob really received what Alice sent. If they agree, then they can be highly confident that Eve didn't open any boxes (or at least that she opened very few boxes). But if they disagree, then they know that Eve has been up to no good.

We can't stop Eve from trying to eavesdrop; that's what Eves do. But because quantum information cannot be copied without disturbing the information, we can detect Eve's activity. If the test is unsuccessful, then Alice won't use that key to encode a message. But if the test is successful, Alice can converse with Bob secure in the knowledge that Eve cannot listen in. So Eve is out of business.

So we have found one feature of quantum information, that it cannot be copied, that can be put to use: it can be used for private communication. But there are deeper properties of quantum information that I think have far greater technological potential. To understand why, let's return to the differences between classical bits and quantum bits. Let's suppose that we have 10 boxes containing classical bits. There are a lot of possible arrangements of those classical bits—lots of ways to put red and green balls in the boxes. But any one arrangement is very simple to describe; I just have to tell you whether each ball is green or red. With quantum bits things are different. Suppose we have 10 quantum boxes. Now it is quite complicated to describe even one typical arrangement of the 10 boxes. In this case, it is not correct to say that each ball is either red or green. Typically, each ball has the potential to be red and the potential to be green, depending on which door we open. Furthermore, the boxes are correlated. Opening any one of the 10 boxes has an influence on what happens when we open the other nine, so our description must include a characterization of those influences. It turns out that to give a complete mathematical description of a typical configuration of 10 quantum boxes, I would have to write down about 1,000 numbers.

And the complexity of the description rapidly escalates as I add more boxes. With 20 quantum boxes, we need about 1,000,000 numbers to give a complete descriptions of all the influences of each box on the others. With 30 boxes, we need 1,000,000,000 numbers. It turns out that for a relatively modest number of boxes (about 300), to write down a complete description of a typical configuration would require more numbers than the number of atoms in the visible universe. It is clear that no such description could ever be written down, even in principle.

So there is no hope of even describing the typical state of a few hundred qubits, no way to write down the description using ordinary classical bits. This feature of quantum information seemed very intriguing to this gentleman — Richard Feynman. Feynman was led to ask (in 1982) a very interesting question, a question he repeated each time he taught a course on computation at Caltech, the course on which this book is based. Feynman wondered, might it be possible that a computer that operates on qubits (rather than classical bits) would be capable of performing tasks that would be inconceivable

using conventional silicon-based digital technology? Feynman's idea was that there may be problems that are very hard to solve using ordinary computers that would become easy to solve if we used a quantum computer instead.

To get a better feel for what this idea means, let's consider an example of a problem that is hard for ordinary computers to solve. You may know what a prime number is — we say that a whole number is prime if it cannot be divided evenly by any whole number aside from itself and 1. Here is the list of all the prime numbers less than 100. Now, with prime numbers, we can play an interesting game, called Find the Factors. I give you a number that can be expressed as a product of two prime numbers, and you have to tell me what those prime factors are. (It's like Jeopardy — I give you the answer, the product of the two numbers — and you have to tell me the question: what two numbers I multiplied together to get the product.

Okay, let's start with this one. 15 can be written as a product of two prime numbers. Who knows what they are? Right, $15 = 3 \times 5$. How about this one: what are the prime factors of 91? Right, it's $91 = 7 \times 13$. But as the numbers get bigger, the game gets a lot harder. Can anyone do this one? [Jeopardy theme.] Are there any lightning calculators out there? Sorry time's up. But if I had given you a piece of paper and a pencil and a few minutes, you probably could have told me that $2537 = 43 \times 59$. As you can see, as the size of the number that we are trying to factor increases, the difficulty of the game escalates very rapidly. Until we get to this: This is a 130 digit number that can be expressed as a product of two 65 digit prime factors. [Mission: Impossible theme.] It's hard, but it is not quite impossible. This number holds the world's record; this is the hardest factoring problem that has ever been solved by a computer. It was done last year, and it is interesting how it was done — the computation involved a network of hundreds of powerful workstations collaborating and communicating over the internet. The computation took several months. But as we add digits to the number to be factored, the time required to do the computation grows so explosively that, say, factoring a 200 digit number is still far beyond what existing computers can accomplish. So perhaps this is a good context to consider Feynman's challenge. Classical computers will never be able to factor very large numbers. Could a quantum computer do better?

Why is factoring so hard? Searching for the prime factors of our number is like trying to unlock a padlock; if we can find the right key (the right prime factors) the key will open the lock. But there are many, many keys to try, many possible prime numbers that might divide our number. We can solve the problem by trying one key after another, until we finally find the key that opens the lock, but because there are so many keys, this takes a very long time. With a quantum computer we can do much better — we can try many, many keys in many, many locks all at the same time. As we have seen, a collection of a modest number of qubits (just a few hundred) can in a sense encode an enormous amount of information. By performing our computation only once, but on qubits rather than ordinary bits, we can achieve the same effect as if we had performed the computation with ordinary bits over and over and over again.

The secret of the quantum computer is that we can invoke a kind of massive parallelism, we can do a very large number of computations all at the same time. Designers

of conventional computers often speak of parallelism, of computers with many processors working together on a problem all at the same time. But a quantum computer can achieve a level of parallelism that we could never dream of with a conventional machine — with only hundreds of qubits, we can perform simultaneously a number of computations that exceeds the number of atoms in the visible universe. We'll never build a conventional computer with that many processors.

What might this mean in practice? With conventional computers, we can now factor a 130 digit number in a few months, let's say one month. But if we take into account how the difficulty of the computation grows as we add digits, we can estimate that that same network of computers would be able to factor a 400 digit number in about 10 billion years, about the age of the universe. So factoring a 400 digit number really is Mission: Impossible. Even with vast advances in computing power, we won't be factoring 400 digit numbers anytime soon. But suppose we had a quantum computer that could also factor a 130 digit number in one month. (That's a very big assumption, but let's make it anyway.) Because of the massive parallelism that a quantum computer can employ, the time it takes to do the computation grows at a much more modest rate. We can estimate that it would take a few years to factor the 400 digit number, which would be feasible. Because of this much more favorable scaling of the computation time with the size of the problem, quantum computers will always have a huge advantage over classical computers for sufficiently complex problems. That a quantum computer could be an efficient factoring machine was pointed out by an exceptionally clever computer scientist named Peter Shor. Shor is a former Caltech undergraduate who is now at AT&T research.

So I have tried to persuade you that a quantum computer would be a wonderful thing to have if we could only build one. But how will we build one? What sort of hardware will a quantum computer have? If we want to be able to manipulate quantum bits, one way to do that (perhaps not the only way, but one way) would be to encode and process the information at the level of single atoms. The technology now exists to suspend an array of individual atoms in a vacuum using electromagnetic fields, and to store the trapped atoms for a long time. Each atom can be in either one of two possible quantum states, so we can still represent them as red or green balls. Since they are individual atoms, such tiny little fellows, you might think it would be hard to see whether each atom is red or green. But in fact that is not very hard. We can shine a laser on each of the atoms, and if the color of the laser light is chosen just right, then all of the red atoms will scatter the light so they will glow visibly, and the green atoms won't interact with the light at all, so they will remain dark. So we easily see which of the atoms are red and which are green. Here is an actual photograph of atoms illuminated by laser light that was taken in a laboratory in the National Institute of Standards and Technology. You can clearly see each atom (a mercury atom in this case) as a diffuse ball of light. This photograph is not actual size; the atoms in the picture are actually about 10 microns apart.

Of course, we want to do a lot more than just look to see if the atoms are red or green; we want to process the information in the atoms and build up a complex and interesting quantum computation. And in particular, if the quantum computer is to realize its potential to perform tasks beyond what classical computers can do, it must prepare and

manipulate configurations of the atoms in which they have complex nonlocal correlations. I will briefly sketch how this might be done. First we shine a laser on one of the atoms. If we choose the color of the laser light just right, then the light will not interact with the atom at all if the atom is red. But if it is green and we leave the laser on for just the right amount of time, then the atom will change from green to red, and at the same time the laser will stimulate that atom, and all the atoms in the trap, to begin vibrating back and forth. If we now direct a laser at another of the atoms, if we choose the proper color for the laser light, then this laser will not interact with the atom if it is not vibrating, but if it is vibrating, and we leave the laser on just long enough, that atom will change color, and the vibration of all the atoms will cease.

So look at what we have achieved. If the first atom is red, no laser ever interacts with any atom and nothing ever happens. But if the first atom is green, then two atoms change color. This operation therefore induces a correlation between the colors of two atoms in the trap. By performing many such operations in succession, we can build up a complex and interesting quantum computation.

It is currently possible to do experiments like this involving one or two qubits and one or two operations. But for a quantum computer to be able to do computations that compete with the best that digital computers can achieve, we will need to scale this up enormously. We'll need machines with hundreds or thousands of qubits (not necessarily atoms in a trap, but qubits of some kind) capable of performing millions or billions of operations. Clearly, the technology has a long way to go before a quantum computers can fulfill their destiny to become the world's fastest machines. But some of us at Caltech are quite excited about the potential of quantum computation, and we have established an institute that is dedicated to exploring whether quantum computers can really work.

[Snare drum.] We call ourselves QUIC, the center for Quantum Information and Computation. There are three Caltech professors involved. Jeff Kimble has to do the experiments, so we let him hold the cabbage. Steve Koonin thinks about quantum computation on those frequent occasions when he has nothing to do as Provost. Colleagues at USC and MIT are also involved — Al Despain, a computer scientist at USC, and Seth Lloyd, a former Caltech postdoc who is now on the MIT faculty. I have had fun for the past year or two working with a talented group of Caltech students and postdocs on problems relating to quantum computation.

[Rocky theme.] This is the team. And here you see us at work, opening some boxes and checking out the balls that are inside. You hear a lot of talk about quantum weirdness, but to really understand what it means you have to get to know these guys. [Put me in, coach.]

Well, what do we really do all day with our balls and our boxes. One thing that we have thought about a lot is: what if a quantum computer makes a mistake? How do you prevent it from making errors? Errors can be a problem even with classical information. We all have bits that we cherish, because information can be very valuable to us, but everywhere there are dragons lurking who delight in tampering with our bits. [Jaws theme; fiendish laughter.] With classical information there are ways to protect ourselves against

the dragons. Say we have a ball that is supposed to be red. Then we can store three copies of the red ball. Then, once in a while a dragon may appear and paint one of our balls green. But there is a busy little beaver who checks the balls periodically, and whenever he sees that one of the balls is a different color than the others, he changes the color of that ball so that it matches the color of the other two. So we see that redundancy (having three red balls instead of one) can protect us from errors. If the busy beaver is quick enough, he can prevent the dragon from damaging our bits. [Applause.]

But what can we do to protect a quantum bit from the dragon? Here too we can try to use redundancy for protection, but we can't do it in quite the same way as with classical bits. We can't replace our box by three identical boxes, the original plus two replicas, because we have already seen that quantum information cannot be copied. Furthermore, when the dragon comes along he might open door number 1 of the box, change the color of the ball, and reclose the box, or he might open door 2, change the color, and close the box. And the beaver needs to be able to fix the error without knowing whether the dragon opened door 1 or door 2.

Well, it turns out that it really is possible to protect quantum information from errors. In this case, it isn't enough to replace a box by three boxes, we actually need five boxes. And the boxes are not all identical replicas of the information that we want to safeguard. Instead, the information to be protected is encoded in correlations involving all five of the boxes, like the nonlocal correlations between Pasadena and Andromeda. That way, there isn't any information in any one of the boxes, instead the information is shared by all the boxes. That means that if the dragon damages one of the boxes, the information still remains intact, because it wasn't in that box anyway. Now the beaver can come along and tell which box the dragon has messed with, and reset that box to its original state. So it seems that redundancy can be used to protect quantum information as it can protect classical information. [Applause.] But the redundancy works in a quite different way; information is protected by storing it in correlations involving many boxes.

We have discussed a lot of properties of quantum information tonight. We saw that the fundamental unit of quantum information is the qubit, which we may envision as a box with a ball inside that is either red or green, such that we can open the box to see what is inside through either one of two doors. The qubits can have peculiar correlations that we cannot reconcile with our usual classical notion of a correlation: the boxes are not like the soxes. Quantum information cannot be copied, and we may therefore use it for private communication. The mathematical description of even a modest number of qubits is exceedingly complex, and we may therefore use qubits to perform massively parallel computations, and achieve an enormous speed up compared to the time required to do a computation on a conventional computer. We can safeguard quantum information from errors by encoding the information in correlations involving many boxes. And I have told you that the first experiments that process quantum information have been carried out in the last year of two.

Clearly the technology must progress a long way before quantum computers are ready to fulfill their destiny as the world's fastest machines. There is a long road ahead. But somewhere at the end of that road, high on a hill, stands a shining castle. We are just

beginning to venture down that road. [Mission: Impossible theme.] Oh, I know, it is hard ... no one said it would be easy. But I really believe that we will eventually arrive at the gates of that castle. When will that be? I really have no idea. But we have come a long way in the 50 years since the ENIAC, and it seems reasonable to me that in another 50 years quantum computer will be in widespread use to solve very hard problems. I could be completely wrong. Maybe quantum computers will never be in widespread use. [Hallelujah!] Or perhaps I am being way too conservative, like Popular Mechanics in 1949.

And what of the shorter-term prospects for putting the weirdness of quantum theory to work for fun and profit? Well, the technology for quantum communication is much more mature than that for quantum computation. Prototype key exchange devices have already been build and tested. These might conceivably see commercial use in just a few years, though at first they would be only for the most paranoid users requiring the utmost in privacy. Ideas generated by recent work on quantum computation are leading experimentalists like Jeff Kimble to develop new methods for preparing exotic quantum states, and for performing new types of measurements that we could not even conceive of a few years ago. And recent theoretical developments are deepening our understanding of quantum information and the ways it differs from classical information. Particularly significant, I think, is the finding that quantum information can be protected from errors with suitable coding methods; I expect that development to have broad ramifications throughout experimental physics.

So the road to quantum computation may be a long one, and there is no telling for sure how long, but it certainly has been and will continue to be a fascinating voyage. And that is all I have to say. [Hallelujah! Mission: Impossible theme.]