International Journal of Theoretical Physics, Vol. 42, No. 2, February 2003 (© 2003)

Quantum Information: Qubits and Quantum Error Correction

Charles H. Bennett[1]

Received September 10, 2002

This edited transcript is a lively, informal presentation of quantum computation and quantum error correction by one of the founders and major contributors to the field.

KEY WORDS: Quantum error correction; quantum information.

Charlie Bennett: So I will give a little introduction to quantum information with this very restricted kind of quantum mechanics, just discrete Hilbert spaces, no momenta and coordinates, nor even any harmonic oscillators, but in an effort to connect it with what we would like to mean by “information” or a “bit.”

Now here we have where the founding of this field—the physics of information—really took place (Fig. 1 [figures are from a slide presentation]). I don’t know what this dog is, or what he or she has been doing in the field since, but this is a really diverse group of people. I guess you have seen this picture already, and I have to update it.

David Finkelstein: Number 26 on the group photo is Stan—Stan Kugell is his name.

Charlie Bennett: Okay, so I am going to begin actually with some of the slides I use with business types at IBM. [Laughter.] So I tell them first about that information which we know about (Fig. 2).

You can tell I am going to make this part of something larger—you’ve probably figured that out already. (I actually used Micrographics Designer for this—and that is why PowerPoint instincts do not serve right.)

[1] IBM Research, Yorktown Heights, New York; e-mail: bennetc@watson.ibm.com

1	Freeman Dyson
2	Gregory Chaitin
3	James Crutchfield
4	Norman Packard
5	Panos Ligomenides
6	Jerome Rothstein
7	Carl Hewitt
8	Norman Hardy
9	Edward Fredkin
10	Tom Toffoli
11	Rolf Landauer
12	John Wheeler
13	Frederick Kantor
14	David Leinweber
15	Konrad Zuse
16	Bernard Zeigler
17	Carl Adam Petri
18	Anatol Holt
19	Roland Vollmar
20	Hans Bremerman
21	Donald Greenspan
22	Markus Buettiker
23	Otto Floberth
24	Robert Lewis
25	Robert Suaya
26	Stan Kugell
27	Bill Gosper
28	Lutz Priese
39	Madhu Gupta
30	Paul Benioff
31	Hans Moravec
32	Ian Richards
33	Marian Pour-E
34	Danny Hillis
35	Arthur Burks
36	John Cocke
37	George Michaels
38	Richard Feynman
39	Laurie Lingham
40	Thiagarajan
42	Gerard Vichniac
43	Leonid Levin
44	Lev Levitin
45	Peter Gacs
46	Dan Greenberger

Fig. 1. MIT Endicott House "Physics and Computation" meeting, May 6–8, 1981.

So this information is part of a larger subject which is quantum information. I use this older-looking font (Fig. 3) to indicate it's older than classical information in the sense that the classical information in computing theory was really formalized in the middle of the 20th century—it has been applied throughout physics, chemistry, and engineering with enormous success, but until recently was not applied directly to information processing—whereas quantum physics was formalized in the first third of it. But people didn't realize that it was—really what they were studying was something about information rather than something about

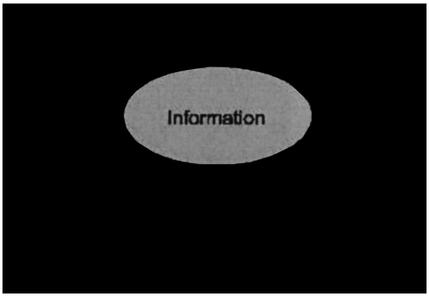

Information

Fig. 2. Information, to begin with. You'll suspect this is part of a larger subject.

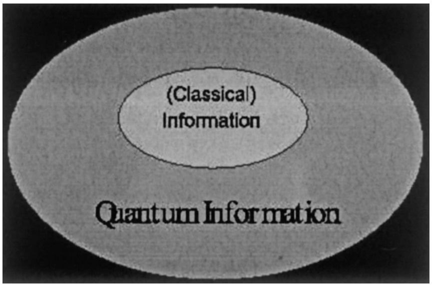

Fig. 3. Classical and quantum information.

physics, not both information and physics. And I'll try to argue why the latter should be true.

So information really is a very useful abstraction. It is the notion of distinguishability abstracted away from what we are distinguishing, or from the carrier of information (Fig. 4). The digital revolution is based on the fact that information can be expressed in bits, and any needed transformation can be accomplished by simple logic operations ("gates") acting on the bits two at a time. The fact that information is independent of its physical embodiment makes possible Moore's law. Making bits ten times smaller and cheaper increases their usefulness, unlike making shoes or cars ten times smaller and cheaper.

We usually take it for granted that information can be read or copied without disturbing it. At it of course cannot travel faster than light or backwards in time.

And, that these bits and gates are fungible, which is of great practical importance because it means that we can make them a thousand times smaller and cheaper and they remain useful for their original purposes, unlike a car that's a thousand times smaller.

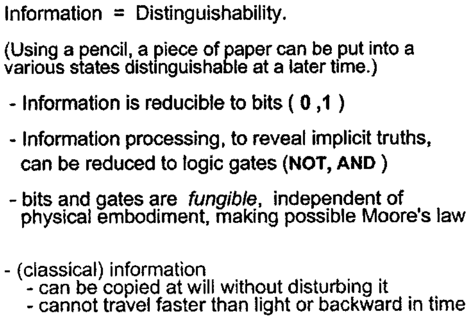

Information = Distinguishability.

(Using a pencil, a piece of paper can be put into a various states distinguishable at a later time.)

- Information is reducible to bits (**0** ,**1**)
- Information processing, to reveal implicit truths, can be reduced to logic gates (**NOT, AND**)
- bits and gates are *fungible*, independent of physical embodiment, making possible Moore's law

- (classical) information
  - can be copied at will without disturbing it
  - cannot travel faster than light or backward in time

Fig. 4. Characterizing properties of information.

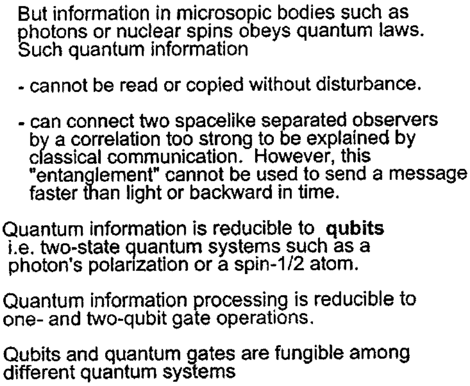

But information in microsopic bodies such as photons or nuclear spins obeys quantum laws. Such quantum information

- cannot be read or copied without disturbance.
- can connect two spacelike separated observers by a correlation too strong to be explained by classical communication. However, this "entanglement" cannot be used to send a message faster than light or backward in time.

Quantum information is reducible to **qubits** i.e. two-state quantum systems such as a photon's polarization or a spin-1/2 atom.

Quantum information processing is reducible to one- and two-qubit gate operations.

Qubits and quantum gates are fungible among different quantum systems

Fig. 5. How quantum information is different from, and similar to, classical information.

And of course people had taken for granted these residual physical properties that you can copy it without disturbing it, and it can’t travel faster than light or backwards in time.

But really, information is not quite like that. We know that in microscopic bodies it obeys slightly different laws (Fig. 5). You can’t always read it or copy it without disturbing it, and we have the phenomenon of entanglement, which is a correlation too strong to be explained by classical communication, which is nevertheless useless for giving yourself advice in hindsight.

But there are some similarities to classical information. That is, that any transformation you might want to make of a quantum state can be reduced to operations that act on one and two qubits at a time. And the qubit is a state of a two-state quantum system.

Gerry Sussman: Charlie?

Charlie Bennett: Yes.

Gerry Sussman: I’m trying to get the sense of this. In the second point you have here with the bullet, it says: “However, this entanglement cannot be used to send a message”—that’s the same as for classical.

Charlie Bennett: Yes, that’s the same.

Gerry Sussman: Why is there a ”however” there?

Charlie Bennett: Well, because the entanglement exhibits phenomena which are not explainable by classical communication—

Gerry Sussman: Oh, I see.

Charlie Bennett: —because it involves parts that are supposedly not interacting anymore. So this creates enormous confusion, and you have to make this point.

Gerry Sussman: Okay. It's not a distinction.

Charlie Bennett: Yes.

Gerry Sussman: There's no distinctions of course.

Charlie Bennett: Yes.

Gerry Sussman: Okay.

Charlie Bennett: This is a similarity.

Gerry Sussman: Yes.

Charlie Bennett: So the idea is, the program here is that all this stuff is really mathematics rather than physics. And we ought to develop a theory of information which generalizes the theory of distinguishability to include these quantum properties, of course as well as using quantum mechanics for physics purposes.

Well, now a lot of this has to do with the theory of reversible computation, which I will say a little bit more about, but one of the illustrations of that is I think it is very important to define "bit" and "qubit."

Well, I think you can find "bit" in the ordinary dictionaries. This qubit definition (Fig. 6) I take from the Random House Unabridged Dictionary of 2006. [Laughter.] But is useful for people like us to think about these important concepts

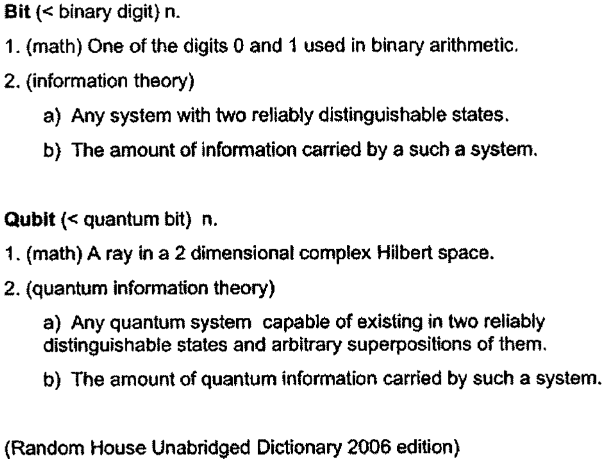

Bit (< binary digit) n.

1. (math) One of the digits 0 and 1 used in binary arithmetic.

2. (information theory)

   a) Any system with two reliably distinguishable states.

   b) The amount of information carried by a such a system.

Qubit (< quantum bit) n.

1. (math) A ray in a 2 dimensional complex Hilbert space.

2. (quantum information theory)

   a) Any quantum system capable of existing in two reliably distinguishable states and arbitrary superpositions of them.

   b) The amount of quantum information carried by such a system.

(Random House Unabridged Dictionary 2006 edition)

Fig. 6. To see what essential about the notions of bit, and qubit, think of how a dictionary defines "bit," and how it might define "qubit."

and how we ought to define them in a way that a dictionary-maker would do. And so I would like to solicit comments and criticisms and suggestions for—see, they're coming up already. Let's hear it. Yes, yes.

David Finkelstein: Well saying a "qubit" is a ray is like saying an electron is a ray. An electron might be described somehow by ray, but the ray is a mathematical object. A qubit is a physical embodiment of information.

Charlie Bennett: No, no, this is a definition of it as a mathematical entity.

David Finkelstein: Ah.

Charlie Bennett: So I think one of the reasons for using the word—

David Finkelstein: —call a state vector of a qubit.

Charlie Bennett: Yes.

David Finkelstein: If you call this a qubit, what do you call the state vector?

Charlie Bennett: I call it a spin or something. I would call it something physical.

David Finkelstein: A physical thing is a spin. Okay.

Charlie Bennett: So one of the justifications for having a new word like qubit is that you want to abstract the distinguishability properties away from a two-state quantum system—

David Finkelstein: Do you realize—

Charlie Bennett: —away from the particular system you're talking about.

David Finkelstein: —1 and 2 are in conflict. Namely, in 2 it is a quantum system, in 1 it is a mathematical object.

Charlie Bennett: Well, no, this isn't—in a dictionary words have different meanings in different contexts.

David Finkelstein: Okay.

Charlie Bennett: Like the word "simple" that you were using, that doesn't mean—I didn't even know what it meant when you first started using it. I knew I ought to, but . . . [Laughter.]

Charlie Bennett: So in other words, this is intended to be a—this is more or less what you might find in a good dictionary. That is, this is a very—I mean almost not very exciting because this is just talking about one of these two digits, and then this is a more profound idea, and actually two more profound ideas.

This is presumably a physical thing that has two distinguishable states, and this is the abstraction of the amount of information carried by such a system, which is a little circular because we haven't said what information is yet. But I am trying to do the same thing down here.

David Finkelstein: I accept that completely. Thank you.

Charlie Bennett: Well the fundamental principle of quantum mechanics is the superposition principle (Fig. 7), which is illustrated in—well, I guess I lean towards the idea that there is a reality and we're trying to understand what its actual nature is. It's a sort of, I guess you would say this is a Platonic view of Nature.

So these are some typical axioms for quantum mechanics (Figs. 8 and 9)

- That for every system there is a Hilbert space of dimensionality equal to its maximum number of reliably distinguishable states.
- Every ray in Hilbert space corresponds to a possible state.
- And spontaneous evolution is a unitary transformation.
- And then we have these other two things: The Hilbert space of a composite system is a tensor product of the Hilbert spaces of its parts, which you can probably derive that from the other axioms.
- And then this one, which if you believe in the Many Worlds' interpretation you don't need at all, having to do with measurement.

Brian Hayes: Did Moses bring the footnotes along with the Tablets?

Charlie Bennett: No, no. These were added by rabbis later on. That's the way they do these things. [Laughter.]

Charlie Bennett: And actually I think there is significant—when I talk about this, I give another analogy to the Ten Commandments, because these are supposedly the fundamental principles that everything obeys. You just look out the window and nothing really seems to be obeying them.

Tom Toffoli: Isn't that Moore's Law?

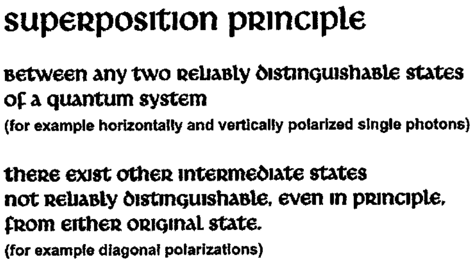

superposition principle

between any two reliably distinguishable states
of a quantum system
(for example horizontally and vertically polarized single photons)

there exist other intermediate states
not reliably distinguishable, even in principle,
from either original state.
(for example diagonal polarizations)

Fig. 7. The foundation of quantum mechanics is incomplete distinguishability, embodied in the superposition principle.

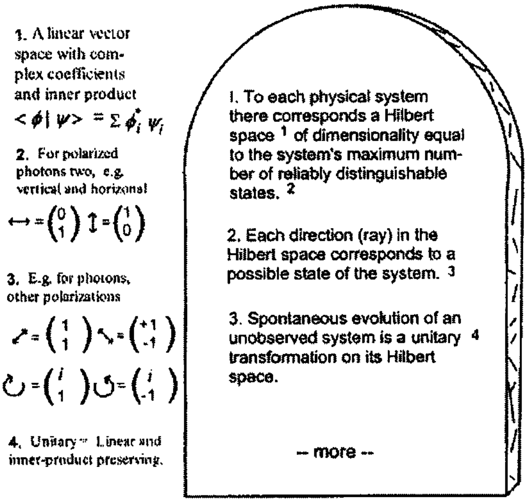

1. To each physical system there corresponds a Hilbert space [1] of dimensionality equal to the system's maximum number of reliably distinguishable states. [2]

2. Each direction (ray) in the Hilbert space corresponds to a possible state of the system. [3]

3. Spontaneous evolution of an unobserved system is a unitary [4] transformation on its Hilbert space.

-- more --

1. A linear vector space with complex coefficients and inner product $<\phi|\psi> = \Sigma \phi_i^* \psi_i$

2. For polarized photons two, e.g. vertical and horizonal

↔ $= \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ ↕ $= \begin{pmatrix} 1 \\ 0 \end{pmatrix}$

3. E.g. for photons, other polarizations

↗ $= \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ ↘ $= \begin{pmatrix} +1 \\ -1 \end{pmatrix}$

↻ $= \begin{pmatrix} i \\ 1 \end{pmatrix}$ ↺ $= \begin{pmatrix} i \\ -1 \end{pmatrix}$

4. Unitary = Linear and inner-product preserving.

Fig. 8. Typical axioms for quantum mechanics.

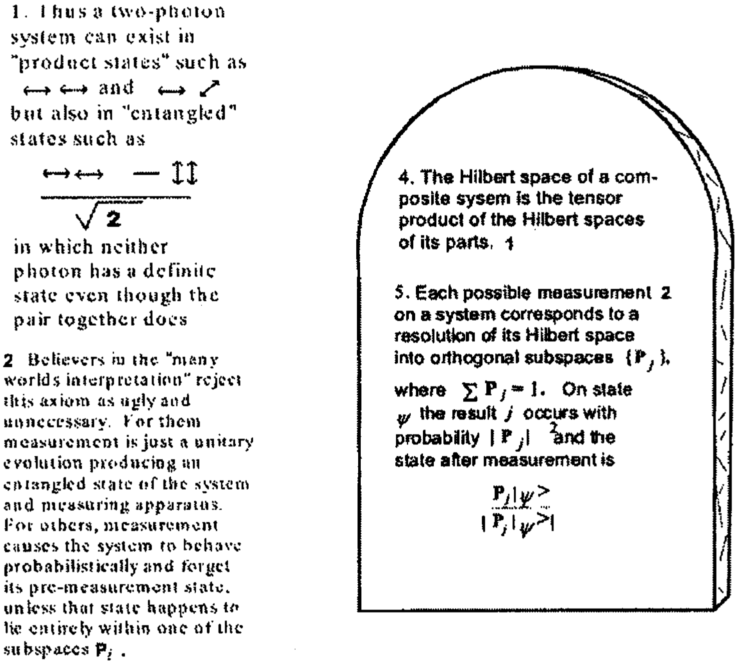

4. The Hilbert space of a composite sysem is the tensor product of the Hilbert spaces of its parts. [1]

5. Each possible measurement [2] on a system corresponds to a resolution of its Hilbert space into orthogonal subspaces $\{\mathbf{P}_j\}$, where $\sum \mathbf{P}_j = 1$. On state ψ the result j occurs with probability $|\mathbf{P}_j|^2$ and the state after measurement is

$$\frac{\mathbf{P}_j|\psi>}{|\mathbf{P}_j|\psi>|}$$

1. Thus a two-photon system can exist in "product states" such as ↔↔ and ↔↗ but also in "entangled" states such as

$$\frac{↔↔ \; - \; ↕↕}{\sqrt{2}}$$

in which neither photon has a definite state even though the pair together does

2 Believers in the "many worlds interpretation" reject this axiom as ugly and unnecessary. For them measurement is just a unitary evolution producing an entangled state of the system and measuring apparatus. For others, measurement causes the system to behave probabilistically and forget its pre-measurement state, unless that state happens to lie entirely within one of the subspaces $\mathbf{P}_j$.

Fig. 9. More axioms for quantum mechanics.

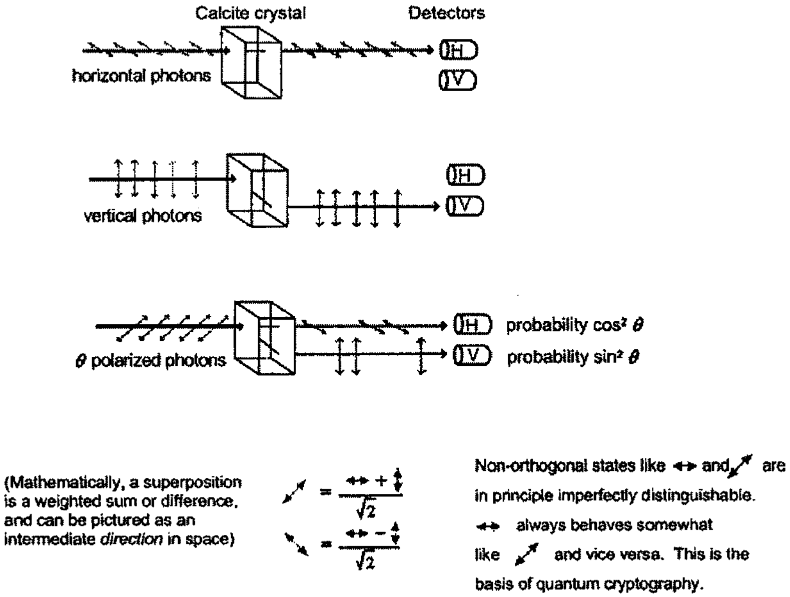
Calcite crystal
Detectors
horizontal photons
0H
0V
vertical photons
0H
0V
θ polarized photons
0H
0V
probability cos² θ
probability sin² θ
(Mathematically, a superposition is a weighted sum or difference, and can be pictured as an intermediate *direction* in space)
↗ = (↔ + ↕)/√2
↘ = (↔ − ↕)/√2
Non-orthogonal states like ↔ and ↗ are in principle imperfectly distinguishable.
↔ always behaves somewhat like ↗ and vice versa. This is the basis of quantum cryptography.

Fig. 10. Superposition principle illustrated by photons.

Charlie Bennett: It is missing an "o."

So here is an illustration of the superposition principle (Fig. 10). Horizontal and vertuical photons can be reliably distinguished, and can be used to carry one bit each. But an intermediate state in principle it is not reliably distinguishable from vertical and horizontal photons. In fact, a diagonal photon behaves sometimes like a horizontal photon and sometimes like a vertical one; it mathematically behaves like a linear combination of vertical and horizontal. Two polarized photons can be reliably distingushed if and only if their polarization directions are at roght angles to one another ("orthogonal").

And this remark of processing this quantum data can be viewed as gates acting on these qubits. The kind of gates that you need are only one- and two-bit gates. One-bit gates are just arbitrary rotations in the two-dimensional Hilbert space.

In the two-bit gates (Fig. 11), all you really need is an exclusive OR, or as they call it in the quantum information Controlled-NOT in which the first qubit controls whether the second one is flipped or not. And because it is a quantum system, if you give it a superposition of inputs, it gives you a superposition of outputs, and this generates an entangled state.

Quantum cryptography just involves preparing and measuring qubits, but a quantum computer allows them to interact, via quantum logic gates, to perform computations. The most famous 2-bit quantum gate is the Controlled-NOT,

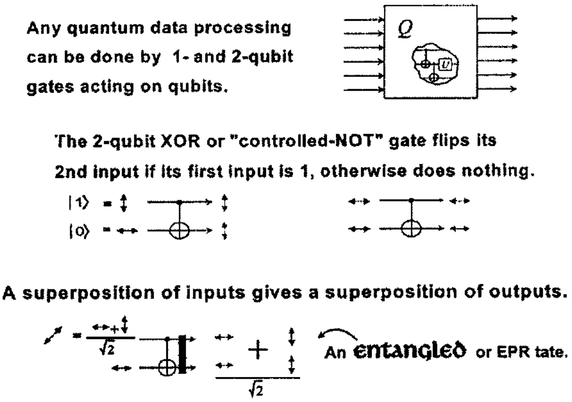

Any quantum data processing can be done by 1- and 2-qubit gates acting on qubits.

The 2-qubit XOR or "controlled-NOT" gate flips its 2nd input if its first input is 1, otherwise does nothing.

A superposition of inputs gives a superposition of outputs.

An entangled or EPR tate.

Fig. 11. Quantum gates.

the quantum version of the classical Exclusive-OR or XOR gate. It can be used to copy a vertical or horizontal qubit, but if one tries to copy a diagonal qubit, the copying attempt fails and a new kind of quantum state results instead, a so-called *entangled* or Einstein–Podolsky–Rosen state of the two qubits that have interacted.

It is probably unnecessary in this audience to explain what entanglement is and isn't, or try to explain it. It is a state of the whole that is not expressible in—I'm scared of that word, actually, "not expressible"—in terms of the states of its parts.

When two systems are entangled, they have a definite relation, even though neither has a state of its own. This is an idea actually that—in my youth around 1967 to be exact—this was an idea that was very easy to explain to people in California: I don't know what I think, and you don't know what you think, but we know that we think exactly the same thing. [Laughter.] But it's an idea that has an exact mathematical description (Fig. 12).

So the excitement of this field comes of course from the fact that if you build a quantum computer—that is, the computer in which the data in the intermediate states exists in the form of qubits that can be entangled, and you use these quantum gates on it, then some problems that appear to be very hard for classical computers, you just make them out of ANDs and NOTs, can be done much faster (Fig. 13).

Much of the interest in quantum computers stems from the fact that they could greatly speed the solution of some hard problems, the most famous of which is the factoring of large numbers.

Gerry Sussman: Can I object to that slide? We don't know that the classical theme can't be solved in polynomial time.

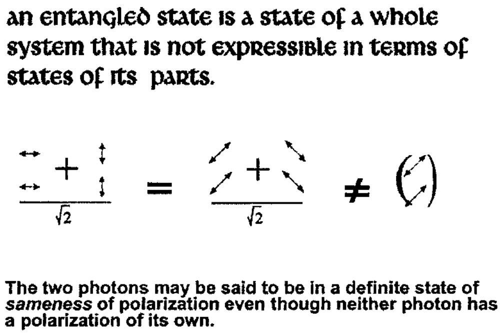

an entangled state is a state of a whole system that is not expressible in terms of states of its parts.

The two photons may be said to be in a definite state of *sameness* of polarization even though neither photon has a polarization of its own.

Fig. 12. Quantum entanglement.

Charlie Bennett: I said so.

Gerry Sussman: Okay, but it says there is exponential speedup.

Charlie Bennett: Okay, you should object to the slide, not what I'm saying about it. I said "which appear" to be harder.

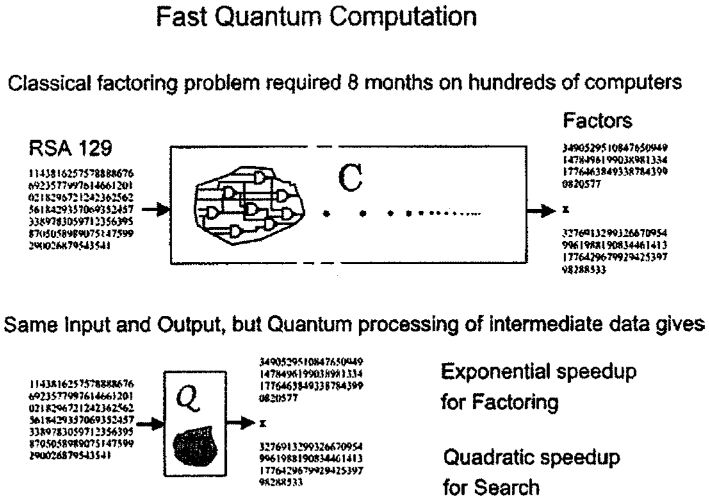

Fast Quantum Computation

Classical factoring problem required 8 months on hundreds of computers

RSA 129

114381625757888676
6923577997614661201
0218296721242362562
5618429357069352457
3389783059712356395
8705058989075147599
290026879543541

Factors

3490529510847650949
1478496199038981334
1776463849338784399
0820577
x
3276913299326670954
9961988190834461413
1776429679929425397
98288533

Same Input and Output, but Quantum processing of intermediate data gives

114381625757888676
6923577997614661201
0218296721242362562
5618429357069352457
3389783059712356395
8705058989075147599
290026879543541

3490529510847650949
1478496199038981334
1776463849338784399
0820577
x
3276913299326670954
9961988190834461413
1776429679929425397
98288533

Exponential speedup
for Factoring

Quadratic speedup
for Search

Fig. 13. Fast quantum computation. Top: a very large classical computation was required to factor the number RSA 129. Bottom: the same number could have been factores in a much smaller number of steps on a quantum computer (the shading insid the quantum computer indicates that during the computation the qubits on the different wires are entangled, even though the final answer is not). If one attempted to observe this intermediate data before the computation was done, the data would be disturbed and the computation would give the wrong answer.

Gerry Sussman: Right. Okay. It’s the slide.

Charlie Bennett: Yes, yes. That’s right.

Gerry Sussman: The computer should be at exponential speed.

Charlie Bennett: Yes.

Tom Toffoli: Why do you make the box, the Q box, smaller than the C box?

Charlie Bennett: Well because the duration of the computation is this extent here (indicating). The number of bits involved in it is the height of it. So this was really intended to be much longer.

So this is—I wanted to get into a feature of quantum information that I think is particularly worth thinking about in connection with the question of: Is it digital or is it analog?

What is the relation of quantum information in this limited arena, rather well-understood arena, to the question of discrete versus continuous?

In what sense is it discrete, and in what sense is it continuous?

That is what I really wanted to concentrate on. And this question is very much connected with the quantum error correction (Fig. 14).

And quantum error correction in turn is one of the main things that makes this whole idea worth thinking about practically, because if it were not for quantum error

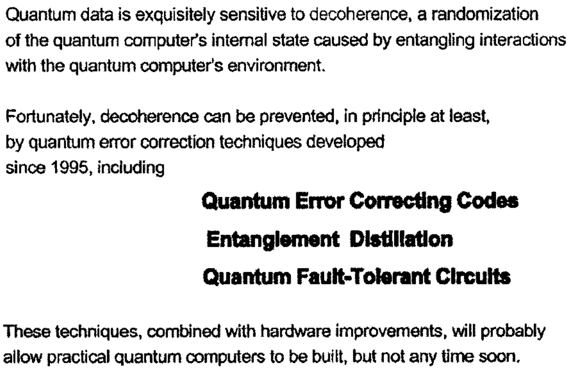

Quantum data is exquisitely sensitive to decoherence, a randomization of the quantum computer’s internal state caused by entangling interactions with the quantum computer's environment.

Fortunately, decoherence can be prevented, in principle at least, by quantum error correction techniques developed since 1995, including

Quantum Error Correcting Codes

Entanglement Distillation

Quantum Fault-Tolerant Circuits

These techniques, combined with hardware improvements, will probably allow practical quantum computers to be built, but not any time soon.

Fig. 14. Quantum error correction. Between 1993, when Shor dicovered his fast quantum factoring algorithm, and 1995, quantum computers were thought to be a fascinating theoretical idea but wildly impractical and unlikely ever to be built. This changed when error-correcting techniques were discovered. These techniques are the quantum version of the discovery by von Neumann that a reliable classica computer could be built out of unreliable parts, if the parts were connected together in a properly redundant fashion.

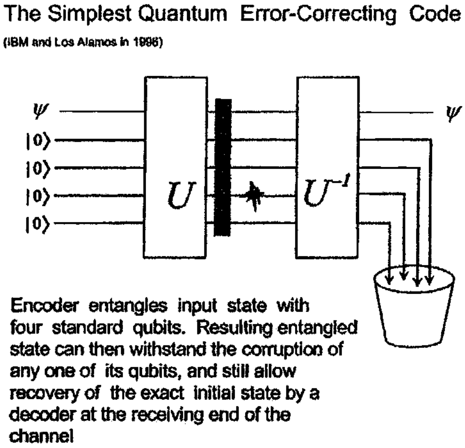
The Simplest Quantum Error-Correcting Code
(IBM and Los Alamos in 1996)
ψ
|0⟩
|0⟩
|0⟩
|0⟩
U
U^{-1}
ψ
Encoder entangles input state with four standard qubits. Resulting entangled state can then withstand the corruption of any one of its qubits, and still allow recovery of the exact initial state by a decoder at the receiving end of the channel

Fig. 15. Simple example of quantum error correction.

correction there would be very little hope of ever building a quantum computer. And now there is some hope. So I will say, you know, that we will probably have a quantum computer but not anytime soon . . . A couple of years ago I promised it to some journalist in the next millenium!

So there are—one of the main techniques of quantum error correction is the quantum error-correcting code. And this is an illustration of the simplest one that will correct an error in a single qubit, which has been expanded and encoded in an entangled state of five qubits (Fig. 15).

And the performance of the error-correcting code is that if you take any of these five qubits and do an arbitrary corruption of it, and then you undo this unitary operation that the encoder did, the bad effects of the noise of the corruption will be entirely syphoned off into these ancillary qubits and the original one will come out in its original state.

And this was an idea that was not obvious to the discoverers of quantum error-correcting codes with the first ones of Peter Shor and Andrew Stein, and they weren't obvious to other people because we were thinking about it in the wrong way (Fig. 16).

It has often been said that classical error correction is based on making multiple copies and then doing a measurement and doing majority voting. And both of those things sounded like something that you can't do with quantum information.

So indeed anything that you might do to this encoded and corrupted version of the qubit, it would seem that you would be learning something about this state, and because of the uncertainty principles essentially you can't learn about a state without disburting it, so if this was an unknown state it would seem like there was no way to repair the error without finding out something which would disturb the state.

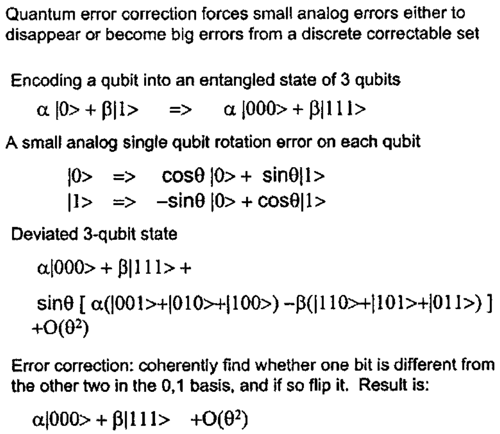

Quantum error correction forces small analog errors either to disappear or become big errors from a discrete correctable set

Encoding a qubit into an entangled state of 3 qubits

$\alpha\,|0\rangle + \beta|1\rangle \quad \Rightarrow \quad \alpha\,|000\rangle + \beta|111\rangle$

A small analog single qubit rotation error on each qubit

$|0\rangle \Rightarrow \cos\theta\,|0\rangle + \sin\theta|1\rangle$

$|1\rangle \Rightarrow -\sin\theta\,|0\rangle + \cos\theta|1\rangle$

Deviated 3-qubit state

$\alpha|000\rangle + \beta|111\rangle +$

$\sin\theta\,[\,\alpha(|001\rangle+|010\rangle+|100\rangle) - \beta(|110\rangle+|101\rangle+|011\rangle)\,]$
$+O(\theta^2)$

Error correction: coherently find whether one bit is different from the other two in the 0,1 basis, and if so flip it. Result is:

$\alpha|000\rangle + \beta|111\rangle \quad +O(\theta^2)$

Fig. 16. The right way to think of quantum error correction.

But this was all really very wrong thinking. And the reason it is wrong is that here is a good example of an even simpler error-correction code that just wraps into three bits (indicating). This is really just like the classical triple redundancy, simplest error-correcting code.

So what we will say here is: We map a zero into three zeroes, and a one into three ones. But an arbitrary combination of a superposition of zero and ones is not mapped into three copies of the superposition, because that would be cloning and you can't do it. You couldn't make the encoder.

Instead, you map it into a linear combination, the same linear combination of three zeroes and three ones. And that means something that you can do with a rather simple quantum circuit like this (indicating).

You just take your original qubit coming in here, and you have two zeroes, and you conditionally flip this one (indicating), and you conditionally flip this one (indicating). And now whatever X, whether X is zero or one or a linear superposition of them, that is the encoder that produces that state.

So there is no problem making the state, and clearly it is an entangled state. And one of the things people say about quantum errors which really is central to the question of continuum versus discrete is to say well how can a quantum error-correction process work?

Suppose each of my bits—this is actually a continuum quantity; it can be anywhere, these two numbers, these two complex numbers, can vary continuously. So suppose I introduce a very small error and I suppose this is zero, drifts off a little bit and just rotates like that and becomes partly one, and this one rotates in the orthogonal direction.

So that would be something like if I was thinking of these as polarizations, it just shifts a little bit like that. So how would I find an error like that and fix it? Because after all, any one of these rotated states should be a possible state of my qubit. It looks like a really hard error-correction process.

Well in fact because of this entangled encoding it is not especially hard at all. The zero, as I say the zero suffers this damage. The one suffers that damage. And if you work it out for the three qubit state, supposing they both—each of the three qubits independently suffers that same rotation? Then we get something like this. We have the original state, and then we have a linear term in this small angle. And then we have a quadratic term.

And now what we need to do to get rid of this linear term is a kind of oblivious error-correction. That is, I don't want to measure these bits because if I did I'd find out what they were. So I want to find out how they're wrong without finding out what they are.

So I make a quantum circuit that ascertains whether they're all three the same without telling me whether they are zeroes or ones. And then if they're not all three the same, it takes the one that's different and puts it back to agree with the other two.

Now I drew a circuit like that, or part of a circuit, to show you that it's not exceptionally complicated. I have my three bits here, and I'm trying to find out—in this case I'm trying to find out whether the A bit, whether they're all three the same, or whether the A bit is different and the B and C are the same.

So what I will do here is to test whether A and B are the same. If they are not the same, I flip this bit down here. This is one of my ancillary bits. Then I put these back the way they were. Then I test whether A and C are the same. And if not, I flip this.

And then I do a Toffoli gate, which says: Is A different from both B and C? And the only way it can be different from both B and C is for it to have one value while B and C have the other value.

And so if that is true, I remember that fact and then I undo all the calculations that allowed me to calculate that so A, B, and C are put back exactly as they were to start, including the error that may or may not be present in A. And these two wires are put back into the zero state [indicating], and this (indicating) wire has state which is zero if everything is okay, and it has a one if the A bit is wrong and disagrees with the other two.

And so then I just flip the A bit conditionally on this wire, and I have corrected the error without finding out what it is.

And the whole theory of quantum error-correcting codes, this is sort of a baby example because this one will not correct phase errors, but this illustrates the idea that if the state deviates a little bit the process of error correction is in a way like a quantum measurement. It forces the system to decide whether there is no error, or whether there is a big error. And if there is a big error, it corrects it.

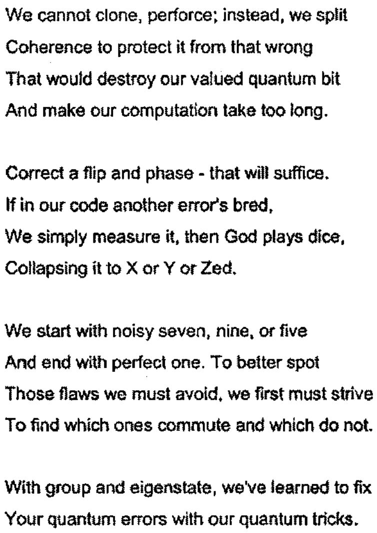

We cannot clone, perforce; instead, we split
Coherence to protect it from that wrong
That would destroy our valued quantum bit
And make our computation take too long.

Correct a flip and phase - that will suffice.
If in our code another error's bred,
We simply measure it, then God plays dice,
Collapsing it to X or Y or Zed.

We start with noisy seven, nine, or five
And end with perfect one. To better spot
Those flaws we must avoid, we first must strive
To find which ones commute and which do not.

With group and eigenstate, we've learned to fix
Your quantum errors with our quantum tricks.

Fig. 17. Dan Gottesmann's quantum error correction sonnet.

And this was such a beautiful idea that one of the discoverers of it, Dan Gottesman, wrote a sonnet about it, which I will now display (Fig. 17). Can everybody read that? [Pause]

So the process there that I was telling about where an arbitrary continuum error gets collapsed to one of a discrete set, all of which are correctable, is this line in here (indicating).

And this is more or less engineering. This just says how this kind of idea of developed in the manner of traditional fault-tolerant computation (Fig. 18), do it all over in the quantum way and you encode the data that you're working on.

Fault tolerant computation involves more that error-correcting codes. It must control errors that happen during the error-correction process, and the inevitable spread of errors when bad qubits interact with good ones. Fortunately, all this can be done, if the error rate of the individual gates and wires can be made low enough to begin with.

You do the computational transformations, which I'm not showing here. You have restoring circuits which take in clean qubits and suck out the noise. And they are robust in their construction enough that they can recover even from errors that are made during the error-correction process.

And so the theory of that—the rather complicated theory of that—has been pretty well developed for particular error models, and will probably be developed in concert with the kinds of errors that occur in whatever physical implementations people pursue for quantum computers.

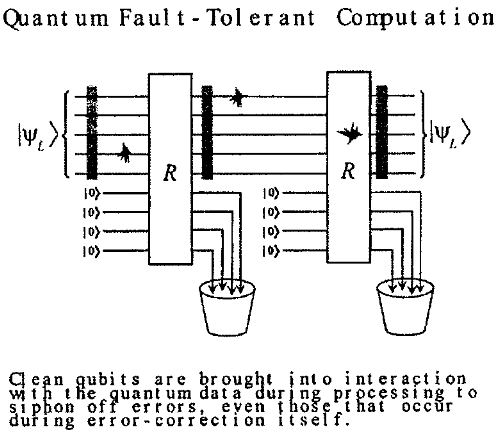
Quantum Fault-Tolerant Computation
$|\psi_L\rangle$
R
R
$|\psi_L\rangle$
$|0\rangle$
$|0\rangle$
$|0\rangle$
$|0\rangle$
$|0\rangle$
$|0\rangle$
$|0\rangle$
$|0\rangle$
Clean qubits are brought into interaction with the quantum data during processing to siphon off errors, even those that occur during error-correction itself.

Fig. 18. Quantum fault-tolerant computation.

Well now just a few other things you can do with quantum information that are difficult or impossible to do with classical information is cryptographic key distribution (Fig. 19).

In Quantum Key Distribution, users “Alice” and “Bob” communicate by a quantum channel (grey photons) and a classical channel (black bits). an eavesdropper (“Eve”) eavesdrop on all their classical messages, and can eavesdrop on the photons as much as she dares. but of course eavesdropping on the photons disturbs them. The dilemma means two possible outcomes. If Eve eavesdrop only a little, Alice and Bob will be able to agree on a secret key, a supply of bits known to them and no one else. If Eve eavesdrops too much, Alice and Bob will almost always detect the eavesdropping and abort the protocol. Except with negligible probability, Alice and Bob will not be tricked into agreeing on a key that is not secret.

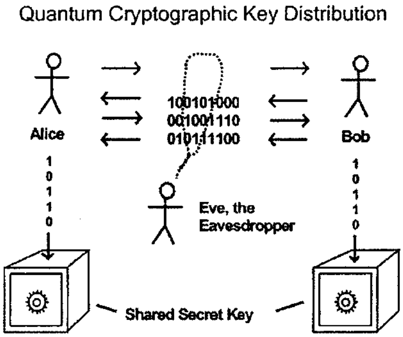

Fig. 19. Quantum key distribution.

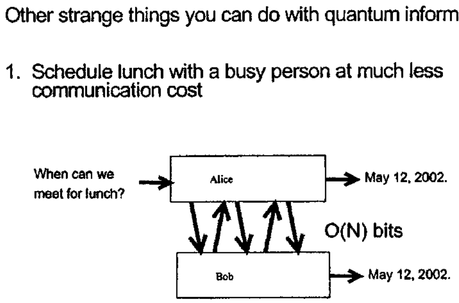

Other strange things you can do with quantum information

1. Schedule lunch with a busy person at much less communication cost

Fig. 20. Scheduling lunch with a busy person.

There also exist classical method of key distribution, in which all communiaction between Alice and Bob is classical, but these methods are insecure in principle. they can all be broken by Eve with a sufficiently powerful classical computer, and many in widespread use today could easily be broken on a quantum computer. But, unless the laws of quantum mechanics are incorrect, quantum key distribution cannot be broken by any amount of computing power, quantum or classical.

Aside from the question of how you might build a quantum computer, the question of what one might do with it is wide open. Quantum computers have beenshown to be capable of many other surprising tasks than fast factoring or fast search.

Scheduling lunch with a busy person (Fig. 20), you can prove that it takes an amount of communication equal to the smaller person's calendar to find out a day on which you are both free, or if there is no free day.

Gerry Sussman: The smaller calendar?

Charlie Bennett: The smaller calendar. So you have two people with calendars. You're trying to do an OR of ANDs. And if you are allowed a quantum

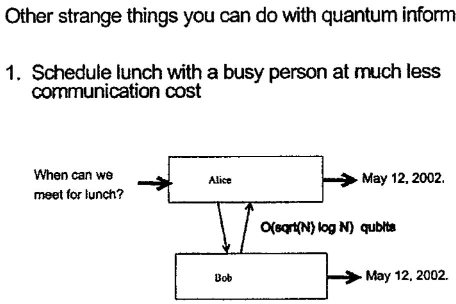

Other strange things you can do with quantum information

1. Schedule lunch with a busy person at much less communication cost

Fig. 21. Quantum speedup of the lunch appointment problem.

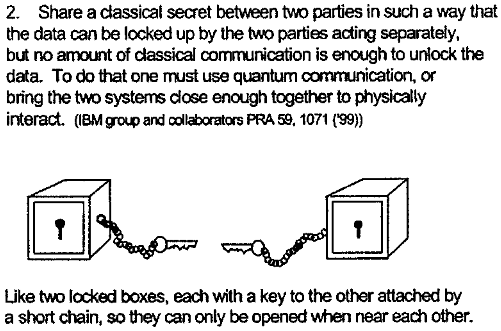

2. Share a classical secret between two parties in such a way that the data can be locked up by the two parties acting separately, but no amount of classical communication is enough to unlock the data. To do that one must use quantum communication, or bring the two systems close enough together to physically interact. (IBM group and collaborators PRA 59, 1071 ('99))

Like two locked boxes, each with a key to the other attached by a short chain, so they can only be opened when near each other.

Fig. 22. Sharing a classical secret.

communication, it can be reduced approximately to the square root of the amount of communication (Fig. 21).

And then you can share classical data between two parties in such a way that if you have classical hijackers they cannot force you to give up the data, and in quantum communication it is necessary to recover the secret (Fig. 22).

And then there is all this quite variety of phenomena connected with quantum channels and their capacities, and how you can enhance the capacity of the classical channel for sending—a quantum channel for sending classical information by giving entanglement (Fig. 23).

One promising theoretical field is entanglement-assisted communication. Here the two parties, Alice and Bob, share entangled particles beforehand, and one asks how the entanglement helps them perform various communication tasks, such as sending classical messages from Alice to Bob. Entanglement doubles

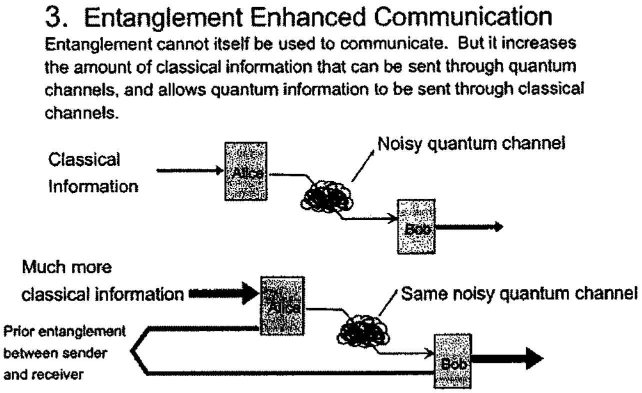

3. Entanglement Enhanced Communication

Entanglement cannot itself be used to communicate. But it increases the amount of classical information that can be sent through quantum channels, and allows quantum information to be sent through classical channels.

Classical Information
Alice
Noisy quantum channel
Bob
Much more classical information
Alice
Same noisy quantum channel
Prior entanglement between sender and receiver
Bob

Fig. 23. Entanglement enhanced communication.

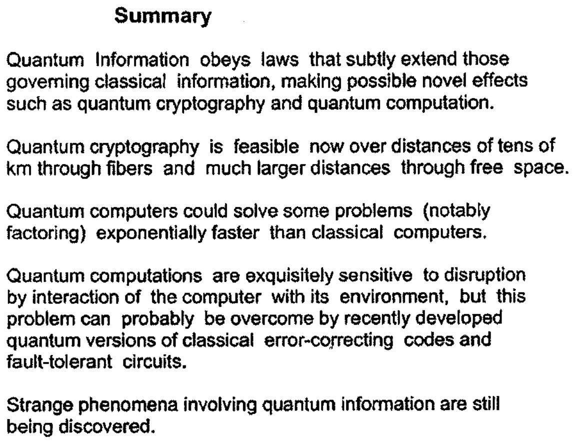

Summary

Quantum Information obeys laws that subtly extend those governing classical information, making possible novel effects such as quantum cryptography and quantum computation.

Quantum cryptography is feasible now over distances of tens of km through fibers and much larger distances through free space.

Quantum computers could solve some problems (notably factoring) exponentially faster than classical computers.

Quantum computations are exquisitely sensitive to disruption by interaction of the computer with its environment, but this problem can probably be overcome by recently developed quantum versions of classical error-correcting codes and fault-tolerant circuits.

Strange phenomena involving quantum information are still being discovered.

Fig. 24. Summary.

the rate at which classical information can be sent through a noiseless quantum channel. If the channel is noisy, the rate can be increased by an even larger factor.

Another thing that can be done with prior entanglement is to allow intact qubits to be sent through a classical channel.

But I won't go on with that because that's really separate from the main topic of this workshop. And this is probably a good place to stop and ask for arguments (Fig. 24).

Thank you [Applause].

1. QUESTIONS SESSION

Tom Toffoli: I have a question that could have been asked even in the 1930s. It's something similar to what you would discuss for quantum error correcting. Instead of digitizing analog information—I'm talking about classical information—just put your stream of classical information, plus two or three streams of constants, and then send them through gates and so on, essentially use the redundancy for correcting directly at the analog level instead of converting to bits and then using the redundancy present at the digital level. How come something like that hasn't been done when digital transistors were expensive? Do you understand the question?

Charlie Bennett: Yes, I think. You're saying why isn't a quantum error correction used for correcting classical analog data?

Gerry Sussman: Yes.

Tom Toffoli: Or even classical analog error correction, something similar to that, to correct classical analog data instead of doing B to A and A to B, and do the correction only in that B?

Charlie Bennett: Well I think the reason is that the structure of classical analog data, and the corrections of errors in it, are very different from quantum data. In quantum data you have two things that you don't have with classical data. One is entanglement, and the other is this ability to force an error to collapse into a discrete choice of whether it happened or didn't happen.

Tom Toffoli: Okay.

Charlie Bennett: And so I think that is why it doesn't have any direct application to analog error correction.

Gerard 't Hooft: You're saying that in some sense a quantum computer is an analog computer, because now we have continuous numbers to work with instead of discretized numbers of classical bits?

Charlie Bennett: I would say almost the opposite. I would say that the proper notion of digital data, of discrete data, is the quantum notion. And that a classical bit is like saying that a real number is a complex number that happens to have no imaginary part.

A classical bit is a qubit that happens to have one of two standard orthogonal values. A classical—I'm glad you asked this question because I meant to—oh, there it is—I mean to say, if we believe, as I certainly do, that quantum information is a generalization of classical information that includes everything about it that's good and more, and besides that it's more physically more realistic, I have to say what a classical bit is and what a classical wire is, and what a classical gate is.

I can all those things. I already said what a classical bit is. It is a quantum bit, a qubit that has one of two standard orthogonal values.

A classical wire is a noisy quantum channel (drawing on flipchart). If I send a zero or a one through this (indicating), and a zero here (indicating), the result will be a copy, a classical copy, and I throw part of it away into the waste basket, or lose it in the environment.

If this is any value other than zero or one, this becomes an entangled state. And when I throw away part of it, it looks like a probabilistic thing has happened to this.

So this is the definition, or is an adequate definition of a classical wire. It's a wire that carries zeroes and ones. But if you try to send a superposition through it, or send part of an entangled state through it, it spoils it.

And any kind of classical computer is just a quantum computer in which every wire has been substituted by this defective kind of wire.

Or to put it another way, a classical wire is a quantum wire with an evesdropper. So I think really the best way to think about it is the central notion of bit is qubit, and a classical bit is a special case. And it is discrete in the sense that it has a discrete set of distinguishable states, but it is continuous in its amplitudes.

It is very different from an analog quantity, from a position of a compass needle. I think the analogy that everybody makes that a spin is a little bit like a bar magnet that can point in any direction is leading people in the wrong direction.

James Baugh: A way to see it is that an analog signal can encode an infite number of classical bits. A qubit can encode—

Charlie Bennett: can encode only one classical bit. And that looks bad, but in fact it is good because it means that this error correction is a robust process, whereas trying to correct analog errors is really hopelessly nonrobust.

James Baugh: That’s right.

Charlie Bennett: Oh, this is a big spin. Well, let’s see. If I know as much as I think I know about how cathode-ray tubes work, if I hold this up to the screen it should make these things move around a little bit. [Laughter.]

Voice: Try holding it up to your disk drive. [Laughter.]

Charlie Bennett: And keep it away from my credit card and my disk card—

Ed Fredkin: And your watch.

Charlie Bennett: What does it do to a watch?

Ed Fredkin: I’m thinking of a mechanical watch; mechanical watches get screwed up.

Charlie Bennett: Who has a mechanical watch anymore? But you’re right—if it has a fair amount of magnetic parts, they can get magnetized and will stick to each other.

Norman Margolus: Okay, on the question of whether quantum computation is analog computation, there is another aspect of course—I don’t know whether or not it bothers you or not—but when you have a classical ensemble, the fact that, with a thousand coins, any particular pattern has a very low probability, that doesn’t seem very strange—you have to get some pattern, right? But if you have a superposition of states of spins, any individual one of which has a terribly low amplitude, and yet it’s the interference between them that produces the computation, somehow those states seem more real and you’re using all of them. They somehow are all implemented in the universe, and that somehow seems worse.

Charlie Bennett: Oh, to me it seems better but I guess that's just—

Norman Margolus: But it seems more analog in that sense.

Charlie Bennett: Well actually I think quantum analog information is a good term, if we apply it to things like the modes of an electromagnetic field and signals which come from an infinite-dimensional Hilbert space.

Charlie Bennett: So there are quantum analog computations which actually go over to resemble in large degree classical analog computations, because you have two of these things that live in an infinite-dimensional Hilbert space, and you know that in the real universe you can't approximate them beyond some finite number of dimensions. And then other than that, they will either not populate those states or they won't populate them accurately, or you won't be able to control them.

Norman Margolus: Yes. But it seems to me that in the quantum computation, when you're producing these states each of which has small amplitude, that somehow there should be a view that those amplitudes are a function of the apparatuses—that the informtion that you're thinking about as being in those amplitudes is somehow in the apparatus, and the qubits are really holding only a small amount of information. I don't know if there's a view like that, but there should be. Anyway, that's just a thought.

Gerry Sussman: What appears to be magical is that the classical error-correction is based on dissipating the error. Dissipation is crucial.

Charlie Bennett: So it is in quantum error correction.

Gerry Sussman: Okay, but if you go back to diagram you had on the sheet, what bit is being dissipated?

Charlie Bennett: That's a very good question.

Gerry Sussman: No, not that one. The previous one.

Charlie Bennett: The previous one. I'll tell you which one it is.

Gerry Sussman: Something has to do away.

Charlie Bennett: Yes. Exactly. Very good question. The bottom one. That's the answer. Norman has got it. This bit tells you whether that error occurred. What you cannot do is to make the error go away and look as if it hadn't occurred by a unitary process.

Gerry Sussman: Right.

Charlie Bennett: You find out whether it occurred. You correct it. And because of the setup of this code and the kind of error it was, you can correct it without finding out out anything about the data.

Gerry Sussman: Sure. But after this wire comes out, mustn't it actually somehow go to—

Charlie Bennett: You have to throw it in the trash and replace it by a fresh zero.

Gerry Sussman: Why?

Charlie Bennett: Because if you try to use this one again, it will put an error in instead of taking it out.

Gerry Sussman: So it's got to be heat bending somewhere?

Charlie Bennett: Yes.

Gerry Sussman: Okay. Now—okay, that's what I needed to see. Okay, thank you.

Brian Hayes: So this is not a reversible process?

Charlie Bennett: Well, yes and no. This is one of the nice things, I mean for somebody who has worked on reversible computing for many years. Most people didn't think it was except that almost all the people who thought it was interesting are in this room [Laughter]—there's a nice combination here. In order to keep the thing reversible—see, if I go this far, I have already found out whether this bit is a zero or a one. But I have now leaked a lot of extra data. I have measured some of these things. And I don't really want to know about the other stuff.