

The future of quantum information science

John Preskill
California Institute of Technology

*A talk at the NSF Workshop on Quantum Information Science
28 October 1999*

This is an especially opportune time to be holding this meeting, as we are approaching the close of the first century of quantum theory. We don't hear so much about it lately, but at the dawn of the 20th century, there was much anxiety about what became known as the Y(1.9)K problem. A crisis was averted when that problem was artfully resolved by this distinguished gentleman [Planck]. By launching quantum theory, Planck set the agenda for the physics of the early 20th century. Now, as Y2K approaches, we are again in need of visionary scientific leadership to face the challenges that confront 21st century science [Shor]. If you contemplate the faces of these two scientific leaders, you may conclude as I did that scientists are much more jolly now than they were 100 years ago. And why shouldn't we be happy – this is a great time to be a quantumist!

1 Three great ideas

What makes it great is that quantum information science has generated some great ideas in the latter part of the 20th century, and we are looking forward to finding out where these ideas will carry us.

The first great idea is quantum computation [Feynman, Deutsch, Shor]. We have learned that a computer that acts on quantum states rather than classical bits can perform tasks that are beyond the capability of any conceivable classical computer. (By highlighting these three milestones, I am slighting the important contributions of many people, in particular Vazirani and Yao and many others who have extended our understanding of quantum complexity since Shor.) But it is fair to say that the interest in quantum computation was greatly stimulated by Shor's factoring algorithm. Factoring this 130 digit number is a difficult problem for you and me and for our computers, but in principle it would be an easy problem for a quantum computer. The broader context though, is that the boundary between the easy problems and the hard problems, the problems that can be solved someday and the problems that we will never be able to solve, is different in our physical world than in the classical world modeled by Turing machines.

The second great idea is quantum key distribution [Bennett, Brassard]. (I received these photos of Charlie and Gilles over a noisy classical channel, but I expect that by applying a sophisticated error correction protocol you will be able to recognize them.) We have learned that, since eavesdropping on quantum information can be detected, key distribution via quantum communication is unconditionally secure. (Here I have slighted Wiesner, and the many people who have contributed to the proofs of security, such as Yao and Mayers.) Quantum key distribution is important because we all value our privacy, and we want to be able to talk to one another without worrying about who might be listening in. But the broader context is that there is an unavoidable tradeoff between gathering information about a quantum state and disturbing the quantum state, and we have far to go to appreciate all of the implications of that tradeoff.

The third great idea, and the most recent of the three, is quantum error correction [Shor, Steane]. We have learned that quantum information can be protected and processed fault tolerantly. (Many people have contributed to the development of this idea since it was first proposed, such as Gottesman, Knill, Laflamme, Aharonov, and Kitaev.) Suppose that we have a single qubit (shown here as a locked box, since we don't know what is inside) that we want to protect from error. We can encode that one qubit in correlations among 5 qubits. Then if one of the qubits is badly damaged, say by an interaction with its environment, it is possible by performing collective quantum operations on all 5 qubits to reverse the error and restore the original encoded state. Quantum error correction and fault tolerance are important because they will be crucial to operating quantum computers reliably so that they are likely to obtain correct answers. But the broader context is that we have now learned in principle how to prepare and manipulate very complicated quantum states, overcoming the debilitating effects of decoherence.

An idea can be called great if it can be expected to have broad implications that can't be easily anticipated in advance. By this standard, quantum computation, quantum key distribution, and quantum error correction are all worthy candidates for greatness. These three developments in quantum information science all pose great challenges for 21st century science.

2 Quantum information in the laboratory

One place where I expect quantum information to have substantial future impact is in the physics laboratory, by enabling new strategies for performing high-precision measurements. I have something in common with the National Science Foundation, because I am from Caltech, which means that we both worry about the LIGO project. (LIGO, the Laser-Interferometer Gravitational-Wave Observatory, is a joint Caltech-MIT project in which the NSF has invested heavily.) LIGO will be taking data soon, and then after a few years there will be a major upgrade of the detection system, to LIGO II. And a few years after that, around 2008, there will be a second upgrade. LIGO III, in its most sensitive frequency band, is supposed to achieve a sensitivity beyond the stan-

standard quantum limit for monitoring the position of a free mass. To reach that sensitivity, novel ideas will be needed.

This year I participated in a working group aimed at generating possible concepts for the LIGO III detection system, led by Kip Thorne, and to which Jeff Kimble made important contributions. And I think that the scope of the discussions that we had would not have been possible just a few years ago, when quantum information theory was in a much less mature state. (I should add that the real pioneer of this subject is Carl Caves, who recognized earlier than others that a deeper understanding of quantum information could suggest new ways to improve the precision of physics experiments.) So quantum information and Big Science will collide, and it may happen sooner than you might think.

But I expect that the more imminent and important impact of quantum information on experiment will occur on the tabletop. New strategies for performing interesting measurements are bound to arise that exploit quantum entanglement, quantum information processing, and quantum error correction.

How does an experimenter measure an unknown classical force (as LIGO does when it receives a gravitational wave signal)? An unknown force can be regarded as a Hamiltonian with unknown classical parameters, a “black box” that we need to probe. To do so, we prepare a suitable quantum state, allow it to evolve under the unknown Hamiltonian for a while, and then perform a measurement, obtaining classical data, the measurement outcome. From that data, we try to make an inference about what the unknown Hamiltonian is.

So in designing a measurement strategy, we seek a productive way to “query” the black box with quantum states, that will provide “optimal” information about what is inside. This sort of black box problem is one that has been much studied by the quantum complexity community, for example by Grover, who found an algorithm for searching an unsorted database for a distinguished state. We have learned from Grover’s algorithm that in an optimal procedure we need not just accept the box as given, but we can “drive” the box — by adding an additional controlled term to the unknown Hamiltonian. In a few settings, like that studied by Grover, we can precisely characterize the choice of driving term that allows us to extract optimal information. I envision that further study of such black box problems will produce many more such insights.

One important lesson that quantum information theory teaches us is that entangled strategies can gather more information than strategies that do not exploit entanglement. As one illustration, consider the case of measuring the state of a single qubit, the state of a spin-(1/2) object. The state can be envisioned as a spin vector pointing in some direction on the Bloch sphere, and we want to determine that direction. For example, we might want to measure a magnetic field which has caused the spin to precess, and we want the measurement to tell us something about how it has precessed, and hence about the direction and strength of the field.

Now suppose that we have two spins that both define the same direction, but contrast two cases: either the two spins are parallel or they are anti-parallel. In which case can we learn more about the direction? Of course if we were to measure the spins one at a time, it couldn’t possibly make any difference, either

way we would learn as much about how the spins point. But the best measurement does not measure the spins one at a time; it is a collective measurement on both spins at once. And indeed, as Gisin and Popescu recently noted, the anti-parallel case is actually better – from two antiparallel spins we can learn more. This is just one simple example of how the possibility of using entanglement makes it quite subtle to determine the best way to measure something, even in situations that seem very simple.

An instructive case study that demonstrates the synergy of quantum information science with science in the physics laboratory is the case of the ion trap quantum computer. In the wake of the discovery of Shor’s factoring algorithm, Artur Ekert, David DiVincenzo and others were saying that we needed to find a physical quantum system in which we could perform coherent conditional dynamics. Cirac and Zoller, who know as much as anyone about what’s possible in quantum optics, recognized that all the required ingredients are already available in ion traps. They observed that if a pulsed laser shines on an individual ion in a trap, then a transition can occur conditioned on the internal atomic state of the ion (indicated here by the colors red and green), and that if the transition occurs, a vibrational mode of the trap (a phonon) would be excited. Then we can shine a pulsed laser on another ion, so that a transition occurs conditioned on the presence of the phonon. This means that we can pick out two ions in the trap, and establish a nonlocal correlation (that is, entanglement) between those two ions. Whether or not you regard this development as a step toward the realization of a quantum computer, it is undeniably important, as it has enabled a breathtaking series of experiments by the NIST group headed by Dave Wineland. The original Cirac-Zoller proposal is now an important paradigm guiding the conception of new approaches to implementing coherent dynamics in other systems.

Now we are facing related experimental challenges on a variety of fronts, particularly in condensed matter systems. We’d like to make strong measurements at the level of single quanta, to induce well controlled coherent interactions among quanta, to fabricate new devices with sufficient precision to achieve that control. We already know something about how to do these things at the level of ions, neutral atoms, or photons, and we want to extend that ability to other systems, like the spins or electrons or nuclei, or phonons. These experimental developments greatly enhance the incentive to develop our understanding of quantum information, since we are approaching the era where we can create and maintain entangled quantum states – it is becoming increasingly urgent to ask which quantum states and operations are the most useful and/or important.

3 RUQ? The quantum-classical boundary

But I am exceeding my authority talking about experiment; that will be Jeff Kimble’s job. So let’s return to theory. An important and fundamental issue to a theorist is to understand in a more precise way the transition from classical to quantum behavior. There are many systems in which the strength of quantum

effects are governed by an adjustable parameter. Most of us think that we can recognize the difference between a system that is behaving very classically and one behaving very quantumly — but is there any sharp boundary between the two. And if so, where is the boundary?

There need not be a precise boundary. In many cases, there may be a smooth crossover from quantum to classical behavior, and in other cases, while there might be a well-defined boundary, it need not be very interesting. But perhaps the most interesting observation ever made about the difference between quantum and classical is that a classical system cannot efficiently simulate a quantum system. We can try to use this idea to establish a well defined boundary between quantum and classical behavior.

There are many facets to this issue of the quantum/classical boundary. I'll describe one setting in which the issue is easily visualized. Recall that a quantum state of a single qubit can be described as a point in the Bloch sphere. Imagine, an octahedron inscribed in the sphere, whose vertices are the six eigenstates of the three Pauli operators $\sigma_x, \sigma_y, \sigma_z$. An arbitrary single-qubit quantum gate rotates the Bloch sphere in an arbitrary fashion, but consider the special gates that map the octahedron to itself. And suppose that in addition, we can perform controlled-NOT gates, which, in a particular basis, flip a target qubit if and only if a control qubit takes the value $|1\rangle$. Finally, suppose that we have a supply of qubits that are eigenstates of σ_z and that we can measure σ_z of an arbitrary qubit.

These rules define an interesting computational model, adequate for implementation of quantum error correction. With these operations, we can encode a quantum error-correcting code, and can perform the recovery operations that reverse the errors. But this model does not suffice for universal quantum computation. In fact, a classical computer can simulate this model efficiently, as Knill was the first to observe explicitly.

We therefore ask, what needs to be added to this model to give it the computational power of a general quantum computer. An intriguing and appealing possible answer is offered by a conjecture due to Kitaev. We imagine adding to the model a reservoir of additional qubits, each prepared in the quantum state ρ . No additional power is added if ρ lies on or inside the inscribed octahedron, but according to Kitaev's conjecture, as long as the state ρ lies outside the octahedron, the model can no longer be efficiently simulated by a classical computer, and may indeed be capable of universal quantum computation. In this particular setting, then, the boundary between quantum and classical is easily characterized — the boundary *is* the octahedron. We could not imagine offering such an answer five years ago, or even making any sense of what it means.

This issue of the quantum/classical boundary is sufficiently important that we should try to approach it from a variety of directions. Dorit Aharonov has done some exciting work on using quantum complexity to define a phase boundary, which she will tell us about. I'd also like to remark that this issue has recently received attention in the context of liquid-state NMR computation. For a fixed number of spins, as we increase the temperature, we eventually enter the regime where the state of the spins in the thermal ensemble are unentangled.

This means that the state can be described as a probability distribution of classical spins, each oriented in a specified direction in space. Since the state is easily described classically, one might expect that the evolution of the spins could also be efficiently described in terms of a classical dynamical model. But in fact, for a range of temperature, there is no known way to describe efficiently the *evolution* of the spins under the entangling unitary transformations that can be applied in an NMR quantum computer. This failure to find a satisfactory classical model of the evolution of an essentially classical state indicates that there may be a hierarchy of computational models of varying strength, and hence a series of phases separating very classical behavior from very quantum behavior. Such a hierarchy in the context of ensemble computing was first suggested by Knill and Laflamme a few years ago.

4 Many-body quantum entanglement

Another topic that can be illuminated by quantum information theory is multi-particle entanglement. Actually, the discussion of NMR computing already relates to multi-particle entanglement, as a core issue is whether the thermal state of an ensemble of molecules is entangled or not. It is also of considerable interest to study the entanglement of a *pure* state of many particles. In the case of a quantum state of a system divided into two parts, one of the triumphs of quantum information theory is that the entanglement can be completely characterized. Given a bipartite pure state $|\psi\rangle_{AB}$, we imagine that many (M) copies of the state are shared by the two parties Alice and Bob. Then Alice and Bob each may perform local operations – unitary transformations and measurements – on their part of the state, and they can talk to one another about what they have done. When the dust settles, they have constructed some number N of EPR states – maximally entangled states of two qubits. If N is the largest possible yield of EPR pairs from such a procedure, then the asymptotic ratio N/M defines a measure of entanglement E . A beautiful result obtained by Bennett and collaborators is that this measure E is precisely the Von Neumann entropy $E = S(\rho_A) = S(\rho_B)$ of the density matrix that describes Alice’s or Bob’s half of the state.

Thus any form of two-party pure state entanglement can be converted to a standard currency – EPR pairs. Furthermore, this procedure is (asymptotically) reversible; we can convert copies of $|\psi\rangle_{AB}$ to EPR pairs and back again, losing a negligible number of copies if the initial number is large. Therefore, the EPR pair counting defines a universal exchange rate, whenever one wants to transform one type of bipartite pure state entanglement to another.

It would be nice to be able to define a standard currency for pure-state entanglement with more than two parts, but it is still unknown whether this is possible, even in the case of only three parties. Here is an example of an innocent looking open question. Imagine three parties, each with two qubits. These parties might, on the one hand, share three EPR pairs among them, or they might, on the other hand share two GHZ or three-particle “cat” states.

Can we locally and reversibly exchange many copies of one for many copies of the other? The answer is not known, although there is inspiring recent work on this question by the IBM group. The question is important, as it pertains to whether three-party cat states possess an entirely different kind of nonlocality than two-party EPR states.

I don't know if this will prove possible, but it would be satisfying if we could reversibly convert many copies of a given m -particle state to, say, a standard number of m -cats, $(m - 1)$ -cats, and so on down to 2-cats (EPR pairs). This accounting would then give us an unambiguous way to say "how entangled" the m -particle state is. I think that such a universal measure of many-particle pure-state entanglement would have many applications. It might enable us to identify new kinds of quantum critical phenomena at which the degree of entanglement of the ground state of a Hamiltonian changes discontinuously, or to characterize which kind of quantum simulation are hard to simulate on a classical computer.

5 From Fermilab to Feynmanlab: the fundamental physics of quantum information

Now, I am supposed to be talking about the future of quantum information science. And when the subject of the future of science comes up, especially if you are a card-carrying particle physicist as I am, one naturally wonders about our future understanding of the fundamental particles and their interactions. We have this citadel on the cornfield [Fermilab] with potential for important discovery extending into the early 21st century. But we are often told that the really interesting questions concern the fundamental physics at the scale 10^{19} GeV, not the TeV scale. So will we have to wait for the Planckatron? That means boosting the energy of the Tevatron by a factor of 10^{16} , and to get a reasonable event rate, boosting the luminosity by a factor of 10^{32} ; it is bound to be more than a factor of 10^{16} more expensive! It seems unlikely that we will be building the Planckatron on earth in the 21st century.

But perhaps we can carry on the quest for the physics of the Planck scale not at Fermilab, but at *Feynmanlab* – our future national quantum computing facility. Is that a ridiculous prospect? Well, I will appeal to authority. You might recognize this fellow [’t Hooft]. He was in the news recently, and will be traveling to Sweden later this year. I have the greatest respect for Gerard ’t Hooft – I know of no deeper and more original thinker. He is so original that sometimes he may seem a little flaky, but trust me he’s not. In a paper earlier this year, ’t Hooft proposes a model in which dissipative classical dynamics underlies quantum physics. In his model, the physics at the Planck scale *can* be efficiently simulated by a classical computer, which leads him to these fighting words:

It will never be possible to construct a ‘quantum computer’ that can factor a large number faster, and within a smaller region of space,

than a classical machine would do, if the latter could be built out of parts at least as large and as slow as the Planckian dimension.

The statement is laced with a good dose of 't Hooftian inscrutability. It is at least clear that 't Hooft agrees with Feynman that there is plenty of room at the bottom. But the remarkable thing is that 't Hooft believes that his model of physics at the Planck scale predicts that efficient factoring of numbers thousands of digits long will never be possible. A serious attempt to implement Shor's algorithm will probe deeply into the physics of quantum gravity.

More authority: Here is another fellow who has had a few good ideas in his career [Weinberg]. In the early 90's, Weinberg attempted to formulate testable alternatives to quantum mechanics, but failed rather dismally. He puts some positive spin on this outcome by saying:

This theoretical failure to find a plausible alternative to quantum mechanics suggests to me that quantum mechanics is the way it is because any small changes in quantum mechanics would lead to absurdities.

This is also an extraordinary statement. If indeed there are no "smooth deformations" of quantum mechanics that make any sense, we are compelled to understand why this is so. And if it is not so, then we are obligated to formulate small deviations from quantum mechanics that can be reconciled with its great successes down to the distance scales that we have so far probed experimentally.

I feel that ideas generated by the recent work on quantum fault tolerance might provide a useful framework for formulating intrinsic deviations from quantum mechanics that have so far escaped experimental detection. Consider, for example, the concept of a concatenated quantum code. Imagine that a single qubit is encoded in quantum correlations among five qubits, as described earlier. But each of those five, when inspected more carefully at higher resolution, is not actually a single raw qubit, but another block of five encoded by the same method as before. And each of those five, inspected at higher resolution, is again a block of five. And so on.

Now suppose that the fundamental qubits are subjected to errors, for example decoherence arising from interactions with the environment. Then as we observe a code block (coarser resolution) we find that the rate of decoherence is longer; the evolution is closer to unitary. And at still coarser resolution, even closer to unitary. Now we could imagine that the decoherence at the fundamental level is really intrinsic, that it arises from microscopic violations of quantum mechanics. (Stephen Hawking, for one, has argued that such violations should be expected in quantum gravity.) Now the picture arising from concatenated coding motivates a scenario in which, as we proceed to longer and longer distances, these violations of quantum mechanics become less and less apparent. We are led to wonder: are the laws of physics *attracted* to quantum mechanics in the infrared? Or in other words: is Nature fault tolerant.

Until just a few years ago, we lacked the tools to productively investigate these questions, but now quantum information science is providing such tools. I

am anticipating a much broader interface between quantum information science and fundamental physics in the future.

6 The second quantum century

So now as we close the first quantum century, a natural question that arises is, what happened *after* 1926. Indeed, the way I teach quantum mechanics to the Caltech sophomores today is not much different than I could have done it in the late 20's. We *have* learned some things since then, though. For one, we know more about the right Hamiltonian of the world – we have a standard model that describes physics with exquisite accuracy down to about 10^{-16} centimeters. We have developed powerful new tools to probe the consequences of the Hamiltonian: the renormalization group and broken symmetries are two of my personal favorites. And .. what else? Well, surprisingly late in the century, we have just begun to appreciate the implications of the tensor product structure of Hilbert space. It sounds pretty simple: the Hilbert space of a two part system is the tensor product of the Hilbert spaces of the parts. But there is much more to it than we would suspect on the surface. Indeed, the natural tensor product decomposition of Hilbert space is at the core of how quantum algorithms achieve a speedup over classical algorithms, and of how quantum error correction works.

So please know that quantum information science is much more than a faster way to factor large numbers. It has first of all earned an enduring place at the center of computer science, by revising our notions of cryptography, computational complexity, communication complexity, error correction, and fault tolerance. And it has given birth to some great ideas that are destined for wider application. I've indicated a few of the possibilities – that quantum information science will lead us to new strategies for performing high-precision measurements, that it will deepen our understanding of the quantum-classical boundary, that it will provide quantitative measures of multi-particle entanglement, and that it will offer a fresh perspective on the unsolved problems of fundamental physics.

It should be clear that in response to these new scientific challenges, a uniquely interdisciplinary community has arisen. This community deserves to be nurtured, and the National Science Foundation can play an important, perhaps even essential, role in assuring that it develops and prospers.