QUANTUM MECHANICS

Error-Correcting Code Keeps Quantum Computers on Track

Quantum computing is a computer scientist's dream. By exploiting the ability of a quantum system such as an array of atoms to be in many different energy states at once, a quantum computer can, in theory, perform vast numbers of computations at the same time, tackling problems that would overwhelm conventional machines. But so far, this fantasy has remained just that: a fantasy. No one has built a quantum computer, much less programmed one to calculate anything. And one bite of reality keeping computer scientists from realizing this fantasy has been the notorious fragility of quantum states, which makes quantum systems vulnerable to errors.

Some recent mathematical discoveries, however, have given computer scientists cause for optimism. To ensure that information remains intact, classical computers rely on error-correcting codes, which include duplicate bits that serve as "quality control" for the rest of the data. That strategy could not work for quantum computers, it seemed, because in quantum mechanics it's impossible simply to duplicate a quantum state; read or copy the state, and you will have altered it. But in a paper published last year in *Physical Review A*, mathematician Peter Shor of AT&T Bell Labs (now AT&T Research) has shown how—at least in theory—to nudge a quantum system back into line without looking at it directly. Shor's theoretical feat is now triggering a flurry of error-correction schemes based on his method.

Such efforts are "extremely important if one wants to do extended quantum computations," says Seth Lloyd of the Massachusetts Institute of Technology, a quantum-computing pioneer. That's because quantum computers are unlikely ever to operate as reliably as their conventional counterparts. "Quantum information has a tendency to degrade more quickly than classical information," he explains. Among the difficulties is that as quantum systems interact with their environment, they tend to drift away from whatever state they're supposed to be in.

The result is an information leak, because these states encode information, albeit not in the usual way. Each bit of information—a qubit—is not simply a binary 0 or a 1; instead it's a combination of the two, with coefficients that are associated with the probability that an observer will find the qubit in a particular state. For example, classically an atom can be in a ground state or in a higher energy state, but quantum-mechanically it can be in a superposition of the two. More generally, a system of n qubits—residing, say, in the ground or higher energy levels of n atoms in a row—consists of a combination of 2^n classical states, each accompanied by a coefficient representing the probability that the quantum system will "collapse" into that state when the qubits are measured.

Quantum computation actually takes place on these coefficients as, for example, a burst of laser light alters the mixture of states in a row of atoms. But there's no way to peep into the computation process, because observing a qubit forces it to "choose" between its two states, collapsing many threads of the computation into one. The only time the operator of the computer actually examines a qubit is at the end of the computation, to learn the outcome. As a result, a qubit cannot be duplicated, except by repeating the computation that produced it.

And that poses a problem for conventional error correction, which is based on redundancy. For example, the simplest way to ensure accurate storage or transmission of a conventional 0-1 bit of information is to create three copies of the bit and then take a "majority vote" among the copies whenever the bit needs to be read. Such a strategy reduces the probability of an error from, say one out of a million to a tiny bit less than three out of a trillion. More sophisticated codes can do even better, safeguarding long strings of "information" bits with just a few extra, error-correction bits. But the methods all make the sensible assumption that bits can be read and copied with impunity—that information doesn't disappear just because you look at it, as it does in a quantum computer.

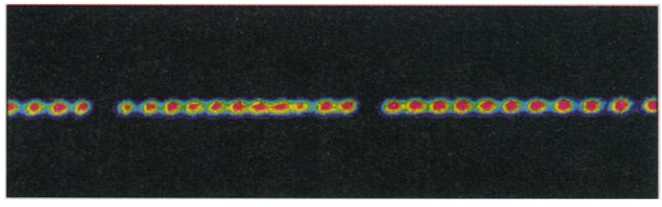

J. MILLER & J. BERGQUISTNIST

You can't always count on atoms. Computing with quantum systems like this row of mercury ions requires a means of error correction.

The key to getting around this barrier came last year, when Shor showed how to safeguard a single piece of quantum information by encoding it as a combination of states in a nine-qubit system—spreading the information content of one qubit across nine of them. Shor's code is constructed so that the original quantum information remains intact even if an error occurs in one of the nine qubits. While based on the conventional "majority vote" approach to error correction, the quantum code corrects errors without explicitly counting ballots. In effect the quantum computer just determines which, if any, vote differs from the others, and registers this information in some ancillary qubits. Counterintuitive as it seems, measuring the ancillary qubits—and thereby losing some of the information stored in them—restores the original nine qubits to their correct state.

The existence of quantum error-correcting codes "came as quite a surprise," says Lloyd. "Before Shor came up with this idea, nobody thought it was possible." Now that the barrier has been broken, however, quantum error-correcting schemes are proliferating.

Researchers at IBM and Los Alamos National Laboratory have streamlined Shor's method to one that embeds single qubits in five-qubit states that are impervious to single errors. Meanwhile, Shor and Rob Calderbank, also at AT&T Research, and, independently, Andrew Steane at Oxford University, have shown how to create quantum analogs of other, more powerful, codes that can correct multiple errors in long strings of bits. "All these other schemes," says Lloyd, "are refinements of Shor's original scheme."

Even with the surge of results, the theory of quantum error correction "is pretty much still getting started," says Shor. One problem, Lloyd notes, is that error correction, being a computation of its own, runs its own risk of making mistakes. "In order for quantum computation to work, you've got to have error correction that is insensitive to errors committed during the process of correction," he says. "There are some ideas floating around" on how to do this, "but nothing written up," he adds.

Then there's the ultimate problem: figuring out how to put these schemes to work in an actual quantum computation. "The problem with all these techniques is they're good for storing quantum bits, but they're not yet good for computing," notes Shor. No one has figured out a way to perform quantum computation directly on information that is distributed over multiple qubits, as these techniques require. Nevertheless, this also looks like a solvable problem, Lloyd says (he and Shor each report having found a likely solution). Quantum computers may be getting closer to reality, and this time it may byte, not bite.

–Barry Cipra