

Richard Doyle
Jet Propulsion Lab
rdoyle@jpl.nasa.gov

Quantum computing: The final frontier?

Richard J. Hughes, Los Alamos National Laboratory
Colin P. Williams, Jet Propulsion Laboratory

As NASA spacecraft explore deeper into the cosmos, speed-of-light-limited signal delays make it increasingly impractical to command missions from Earth. Future spacecraft will need greater onboard computing capacity to mimic human-level intelligence and autonomy. Unfortunately, computer manufacturers

will have difficulty providing the vastly increased computing power the space-exploration community will need.

The solution might well come from quantum computers, which offer properties of size, power, and robustness that are ideally suited to the space environment.

The potential of quantum technologies goes far beyond enhanced computing capacity. Future space missions will involve direct participation of non-NASA scientists. This will necessitate allowing more open access to spacecraft systems via free-space communication links. Quantum cryptography would allow such channels to be made absolutely secure and invulnerable to attack by malevolent hackers. To explore these

possibilities, this article describes the progress to date in understanding how quantum computers and related quantum information-processing devices might advance space exploration.

Smaller, faster, rad harder

Computing capacity conventionally grows through increased processor speed. Over the past three decades, for instance, the microprocessor industry has shrunk computer hardware by a factor of two roughly every 18 months, while maintaining robustness and mass-production capability. Smaller hardware means smaller distances for signals to travel inside microchips, so processor speed has risen dramati-

cally. However, as we enter the 21st century, chip manufacturers have begun to encounter problems in fabricating smaller computer hardware. Worse still, they are realizing that smaller computers don't obey the same rules as larger ones. They must, therefore, turn to alternative physical models of computation for designing tomorrow's computers.

Far from being bad news, the need to question computing's limits provides a new and unexpected opportunity for a quantum leap in computing capacity. By extrapolating the trend in miniaturization, we can conceive of computer hardware that uses single atoms to implement bits—the nuggets of classical information—by about 2020. At such scales, the dominant physics is quantum physics rather than classical physics.

These quantum rules admit subtle and counterintuitive physical effects, which can be harnessed to perform extraordinary operations on data stored in quantum bits, or *qubits*, that are impossible to realize on any classical computer—no matter how advanced it might be. Machines designed to exploit these extraordinary quantum effects are called quantum computers.^{1,2} (See the “Fundamentals of quantum computing” sidebar for a discussion of the physics involved.)

The expansion in the fundamental operations available to quantum computers lets them achieve an enhanced computing capacity, not simply by being smaller, but by running qualitatively new types of algorithms that cannot run as efficiently on any classical computer. Indeed, we already know that quantum computers can break codes using a factoring attack and simulate complex quantum mechanical systems exponentially faster than their classical counterparts.^{3,4}

In addition to a possible algorithmic

Quantum computing

Quantum computing is an exciting area from a computer science viewpoint. Not only is there the possibility of exponential speedup on some classes of problems, but quantum computing also offers a fresh take on some fundamental concepts such as the bit and the algorithm. The properties of quantum systems are non-intuitive, but they may be exploitable in ways that transform looming device physics constraints relating to computation into new opportunities.

NASA's interest in quantum computing is based in part on potential efficiency gains for the kinds of computational problems associated with onboard autonomy capabilities. But equally relevant, quantum computing devices may be energy efficient in unprecedented ways, and quantum information theory may offer new concepts for secure communications.

The authors are engaged in cutting-edge research in quantum communications and quantum algorithms, and they offer you their view of the potential of this new field and its relevance to NASA.

—Richard Doyle

Fundamentals of quantum computing

The fundamental nugget of classical information is the bit—a 0 or a 1. To understand a quantum computer, start by thinking about what happens as the physical structures implementing bits become smaller. Bits nowadays are so commonplace that we no longer give much thought to the properties that we expect them to possess. But as our computing and information-processing artifacts descend to ever-smaller scales, we must question whether bits behave as expected. For example, we assume that a bit always has some definite value, that this value is either 0 or 1, that we can make a perfect copy of a bit, that we can read a bit without affecting its value, and that reading one bit has no effect on the value of another, unread, bit. At sufficiently small scales, when using individual quantum systems to encode bits, all of these assumptions turn out to be wrong, because, in the words of Richard Feynman, “Nature isn’t classical, dammit!” As we understand the rules for manipulating quantum bits and can exploit them to conceive of new kinds of algorithms, we are creating the new field of quantum computing.

The qubit is the classical bit’s quantum analog. We can use almost any 2-state quantum system, such as an electron’s spin, a photon’s polarization orientation, or an ion’s internal energy levels, to encode a qubit. We can represent the two classical bit values as states labeled $|0\rangle$ and $|1\rangle$. Whereas a classical bit must be either 0 or 1, a qubit can be an arbitrary superposition of $|0\rangle$ and $|1\rangle$ simultaneously—a state such as $c_0|0\rangle + c_1|1\rangle$ with c_0 and c_1 complex numbers such that $|c_0|^2 + |c_1|^2 = 1$. We can visualize a qubit as a vector of unit length pointing from the origin to a point on the surface of a bounding sphere (see Figure A). The classical bit values correspond to the north and south poles, and superposition states correspond to the other points on the surface.

When measuring a superposition state $c_0|0\rangle + c_1|1\rangle$, it is generally impossible to predict with absolute certainty the bit value that the measurement will obtain. We only know that the probability of finding the qubit to be 0 is $|c_0|^2$ and that of finding it to be 1 is $|c_1|^2$. Moreover, the act of making the measurement appears to project the qubit into state $|0\rangle$ or state $|1\rangle$ consistent with the measurement outcome.

Any nontrivial computation requires a quantum memory register of n qubits. Fortunately, n need only be around 50 before we can imagine a quantum computer that can perform specialized computations that outstrip the capabilities of any classical rival. The state of n qubits lies in a 2^n -dimensional space spanned by the vectors

$$\underbrace{|00\dots 0\rangle, |00\dots 1\rangle, \dots, |11\dots 1\rangle}_n$$

If we measure the bit value of all n qubits, the possible states in which the complete memory register can be found correspond to the 2^n classically allowed bit strings that n bits can represent—that is, “00...0,” “00...1,” ..., or “11...1.” However, between measurements, the n qubits can exist in superposition of all the 2^n classical states. Quantum computers thus have tremendous capacity to work on several different computations at once. Alas, quantum mechanics forbids us from reading the answer to each of these superposed computations individually. We can only make some measurement that reveals a collective property of all the answers, which is often good enough to perform useful computation.

Quantum memory registers can exhibit quintessentially quantum behaviors that have no classical analogs. For example, suppose we had a 2-qubit quantum memory register in the state $(|00\rangle + |11\rangle)/\sqrt{2}$. This quantum state is rather strange: initially neither qubit is in a definite state, each could be found to be 0 or 1 with equal probability. But as soon as we measure one of the qubits, the bit value of the other (unmeasured) qubit will become definite. We say such a state is entangled, because we cannot think of the register’s state as being composed of a definite state for each qubit it contains. Instead, measurements made on one set of qubits have a side effect on another set of unmeasured qubits. The correlations between the bit values or entangled qubits can be much greater than anything classical bits can achieve. It is these excess correlations that ultimately account for much of quantum computing’s power.

So much for quantum memory registers and readout. How do we compute? That is, how do we make the qubits perform a purposeful computation? Surprisingly, the answer was discovered in 1926, long before quantum computers were invented! That’s because a quantum computer is first and foremost a quantum mechanical system, and Erwin Schrödinger discovered a formula that describes how all isolated quantum systems evolve in time. If the quantum computer’s memory register is described at time t by the state $|\psi(t)\rangle$, it must evolve in accordance with Schrödinger’s equation:

$$i\hbar \frac{\partial |\psi(t)\rangle}{\partial t} = H |\psi(t)\rangle$$

where H is an operator, called the Hamiltonian, related to the total energy of the system. This equation’s solution is

$$|\psi(t)\rangle = \exp(-i H t / \hbar) |\psi(0)\rangle = U |\psi(0)\rangle$$

where U is some unitary operator. (An operator is unitary if its inverse equals its conjugate transpose. Unitary operators are linear: given an input state we can compute a unique output state and vice versa. The unitarity property turns out to have an important physical implication. It means that a perfect quantum computer is necessarily a reversible computer.) This means that the computer’s final state comes from acting on the initial state with some unitary evolution operator U .

To interpret this physical evolution as a computation, we make the following correspondences. At time t , the state $|\psi(t)\rangle$ represents the memory register’s content; the algorithm that the quantum computer is executing is the unitary operator U ; the initial data on which the computer will act is encoded in the initial state $|\psi(0)\rangle$. The outcome from the quantum computation results from a measurement made on some or all of the qubits in state $|\psi(t)\rangle$. Thus, quantum algorithms are, in reality, just unitary transformations of some initial state vector $|\psi(0)\rangle$ into some final state vector $|\psi(t)\rangle$ followed by a measurement.

Once we know a unitary operator that implements a desired quantum computation, to make a practical quantum computer for performing this computation, we must break down this operator into a sequence of quantum logic gates that act on single qubits or pairs of qubits at a time. This defines a quantum circuit that performs the computation.

Quantum algorithms are now available for such problems as factoring large composite integers; computing discrete logarithms; estimating eigenvalues; determining means, medians, and maxima of functions; simulating stochastic processes; evaluating high-dimensional integrals; and finding collisions in functions. A few quantum algorithms are exponentially faster, but the majority are polynomially faster, than their classical counterparts.




Figure A. Picture a qubit as a vector contained in a sphere.




Figure 1. A spacecraft might have to replan its actions under extreme constraints on power, maneuverability, time, and computational resources.

advantage, quantum computers might also be more robust in the space environment than classical computing devices of comparable complexity. Quantum computers store information in exceedingly small physical structures, such as in the orientation of nuclear spins. These objects present far smaller cross sections to incoming radiation than conventional computer circuitry, making it less likely for radiation to damage a comparably powerful quantum device.

Furthermore, in principle, quantum computers are reversible: they can recover all of

the energy expended during computation. The only time energy must truly dissipate is when information is erased, so quantum computers could be more energy efficient than today's classical irreversible computers of comparable computing capacity.

Quantum algorithms for NASA applications

In certain mission-critical applications, computational processing speed is of paramount importance. For example, a remote spacecraft might be about to encounter a rare

astronomical body or perform a part of the mission for which there is only one opportunity for success, such as insertion into orbit or landing on a planetary surface. There will often be the need for onboard capability to process information, make decisions, and sometimes, replan elements of the mission in real time (see Figure 1). Unfortunately, planning is an NP-Hard problem, so fine-grained replanning quickly consumes available computational resources.

Currently, no known algorithm (classical or quantum) can solve NP-Hard problems in better than exponential time in the worst case. But it appears that quantum algorithms can be faster than classical ones by a significant factor. Consider solving the planning problem by reducing it to the k -SAT, or propositional satisfiability, problem, which is the canonical NP-Hard problem. The best classical tree-search algorithm for solving k -SAT is ResolveSAT.⁵ Other types of classical algorithms—local search algorithms such as Walk-SAT and Simulated Annealing—do better at solving k -SAT in practice, but their complexity is harder to pin down. Certain randomized algorithms do reasonably well, but invariably at the price of trading correctness for efficiency: some are guaranteed to terminate with correct results but have uncertain running times, whereas others have certain running times but are not guaranteed to terminate with correct solutions.

The state of the art in quantum algorithms for NP-Hard problems is a nested quantum tree-search algorithm invented by Nicolas Cerf, Lov Grover, and Colin Williams.⁶ This algorithm builds upon another quantum algorithm, called Grover's algorithm, for performing unstructured quantum search.⁷




Figure 2. Unstructured quantum search.

Unstructured quantum search. We can intuitively interpret the unstructured search problem as follows. Suppose we have a telephone directory and must find the name of the person who has a particular telephone number. Because the telephone directory is sorted by name, rather than by number, we can do no better than to pick a random starting point and examine one entry after another. If the directory contains N entries, in the worst case we must look at N items before finding the desired name. Consequently, unstructured classical search's complexity would appear to be $(O(N))$ and there would seem to be no way of improving on it.




Figure 3. A simple constraint-satisfaction problem. A seven-node graph is to be colored using at most three colors such that no two nodes connected directly by an edge have the same color. On the left is the graph and on the right a satisfying coloring of the graph. Graph coloring is an NP-Hard problem.

However, a quantum computer has new operations on data that lets it solve unstructured quantum search more efficiently. The idea behind the unstructured quantum search algorithm is to begin with an equally weighted superposition of all possible candidate telephone numbers (see Figure 2). Suppose there is an oracle (implemented as some secret quantum circuit) that, if given the index of a telephone directory entry, will tell us whether the index is or is not that of the sought-after item. Because the oracle is “quantum,” we can hand it a superposition of indices, rather than just a single index, and the oracle will respond with a superposition of replies indicating which indices do or do not correspond to solutions.

We can use the oracle to build an amplitude-amplification operator \hat{Q} whose repeated application tends to pump probability amplitude from the nonsolution states into the solution ones. The quantum computer does not read the contents of its memory register during this amplitude-amplification process, which allows it to remain in a quantum superposition state. The quantum computer makes only a single final measurement after a prescribed number of amplitude-amplification applications—just enough, in fact, to boost the probability of a measurement yielding the solution to close to 1. By amplitude amplifying $(\pi/4)\sqrt{N}$ times, the true solution will bubble to the surface and hence be detected.

Thus, for a task requiring $O(N)$ operations on a classical computer and that seems to be impossible to speed up, we can solve it in just $O\{(\pi/4)\sqrt{N}\}$ steps on a quantum computer.

Quantum algorithms for NP-Hard problems. In principle, we could use Grover’s unstructured quantum search algorithm to solve an NP-Hard problem such as con-

straint satisfaction, planning, scheduling, combinatorial optimization, propositional satisfiability, or diagnosis. For example, consider a constraint-satisfaction problem (CSP) consisting of μ variables that each can take on b possible values, together with a set of constraints that restrict the values that tuples of variables can assume simultaneously (see Figure 3).

Currently, the best known tree-search algorithm for solving the kind of CSP described above is ResolveSAT.⁵ ResolveSAT’s complexity is $O(b^{0.446\mu})$ (for a problem having μ variables and b values per variable). In principle, we could use a slightly modified version of Grover’s unstructured search algorithm to solve an NP-Hard problem by replacing the oracle with a quantum circuit that can decide whether a candidate solution is in fact a solution in polynomial time, then searching among the $N = b^\mu$ assignments of values to variables in a time complexity of roughly $O(b^{0.5\mu})$. However, this is less efficient than ResolveSAT, so the naïve use of Grover’s algorithm does not offer any advantage.

However, as NP-Hard problems have internal structure, a quantum computer could do better than to run Grover’s algorithm. Our structured quantum search uses a Grover-like quantum search at an intermediate level of the search tree to condition a subsequent quantum search in the leaves, which produces a nested quantum-search algorithm overall (see Figure 4).

Whereas ResolveSAT’s complexity is $O(b^{0.446\mu})$, the nested quantum-search algorithm has a complexity of $O(b^{0.333\mu})$ with one cut level, improving with more cuts. This is an exciting development for computer science because the recent improvements in SAT-solver algorithms, such as ResolveSAT, have only nibbled away at the exponent in the exponential cost function.

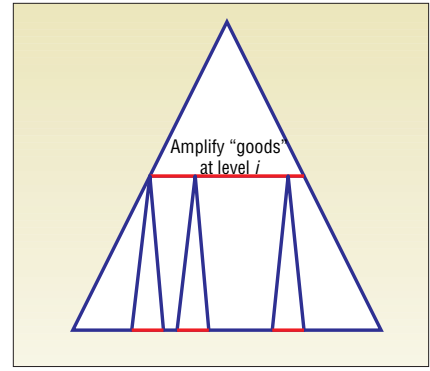


Figure 4. Nested quantum-search embeds one quantum search algorithm within another. There are N candidate solutions in the fringe, but only those that are extensions of a “good” at level i need be examined.

Quantum computing promises to take a big bite out of the exponent.

Could quantum computers do even better? Unfortunately, we now know that no quantum algorithm can do better than a square root speedup on unstructured quantum-search problems.⁸ Minor improvements beyond this are possible on average using parallel, punctuated, quantum searches.⁹ Because there is some probability of obtaining the correct answer before amplitude amplification completes, a society of k independent quantum searches performing partial amplitude amplification followed by premature measurement might find a solution earlier than expected. For a single search agent, this strategy results in 13% additional speedup on average for unstructured quantum search. However, the ultimate speedup attainable for structured quantum-search problems, such as NP-Hard problems, remains an open question.

Distributed quantum algorithms. Some future space missions will involve constellations of spacecraft acting in unison (see Figure 5). Such constellations permit simultaneous monitoring of an entire planet, rather than the periodic monitoring allowed by individual orbiting spacecraft. However, managing constellations of spacecraft presents new computational challenges. How can we fuse the data from different spacecraft to create a snapshot of the entire planet at a given instant? How much communication must take place between spacecraft to perform some collective task? These questions naturally lead us into thinking about whether we can combine quantum communications and quantum-computing techniques to provide new solutions to these kinds of problems.

For the simplest example of a NASA-relevant distributed computation, consider




Figure 5. A constellation of spacecraft must be synchronized to perform data fusion tasks.

synchronizing a pair of clocks on two different spacecraft. For fusion of time-dependent data to be possible, two spacecraft must agree on the time. Generally speaking, the more accurately clocks can be synchronized the better the resolution attainable from distributed sensors. Conventionally, synchronization occurs through an operational line-of-sight exchange of light pulses between two observers, say Alice and Bob, who are co-located with their clocks. While such a scheme is practical, it is not immune to errors. In the GPS satellite constellation, for example, the ability of the spaceborne atomic clocks to synchronize with a ground-based master atomic clock is limited by the fluctuating refractive index of the atmosphere, which causes the speed of light to vary randomly, limiting our ability to establish absolute distance and the resultant timing information with high accuracy.

Could quantum communications and distributed quantum computation help? We might think so because entangled particles have some very peculiar properties that let us avoid sending timing information through the atmosphere. For example, if Alice and Bob each hold one member of a pair of entangled particles in the state $(|01\rangle - |10\rangle)/\sqrt{2}$, when Alice measures her particle, she will instantaneously

cause the state of Bob's particle to become definite as well, regardless of the distance between them and the nature of the intervening medium. Similarly, if Alice and Bob start out by sharing pairs of entangled particles, certain distributed quantum computations, such as a joint appointment scheduling, can be performed using less communication overhead than ought to be necessary classically.¹⁰ Likewise, if Alice and Bob share entangled particles, Alice can send Bob two bits of classical information by sending only one qubit between them.¹¹ So, it appears that *shared prior entanglement* is a powerful computational and communications resource.

Assuming such shared prior entanglement, it appears possible to synchronize a pair of clocks without sending timing information through the atmosphere.¹² Alice and Bob start out having corresponding members of entangled pairs of atoms, each in the state $(|01\rangle - |10\rangle)/\sqrt{2}$. The atoms are indexed so Alice and Bob know which atom in Alice's ensemble corresponds to which (entangled) atom in Bob's ensemble. This shared ensemble is effectively a pre-clock from which Alice and Bob will later distill a pair of synchronized clocks. The pre-clock is idling because each atom in it is in the $(|01\rangle - |10\rangle)/\sqrt{2}$ state that does not evolve in time under any

symmetric operations by Alice and Bob.

To make a clock, Alice simultaneously measures each particle in her ensemble in the

$$\left\{ \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \right\}$$

basis. For each atom, Alice will either find it in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ state (Type I) or the $(|0\rangle - |1\rangle)/\sqrt{2}$ state (Type II). Alice's measurement does two things: it starts both Alice's atoms and Bob's atom "ticking" and tells Alice the indices of the atoms in both her Type I and Type II subensembles. Alice can choose to use either Type I or Type II atoms to make a clock. Bob's clock has complementary Type I and Type II subensembles of ticking atoms, but cannot yet tell which is which. Instead he must wait for a classical communication from Alice telling him the indices of the atoms that Alice used to make her clock. If Bob selects atoms with the complementary set of indices, he will have a clock that is ticking in synchrony with Alice's.

Thus Alice and Bob are left with synchronized clocks without any timing information having been transmitted through the atmosphere between them. This is really just a thought experiment at this stage, predicated on the assumption of shared prior entanglement between Alice and Bob. If this can be made to work, Alice and Bob will be able to synchronize their clocks to a greater precision than possible with line-of-sight optical light pulses. This, in turn, would let us improve the resolution attainable with distributed sensors—such as very long baseline interferometers—for imaging planets orbiting distant stars (see Figure 6).

It is not yet clear whether establishing shared prior entanglement is really any different from distributing a pair of clocks that start at a common location. Nevertheless, thinking about distributed quantum computations in space has stimulated new and exciting questions relating to relativistic quantum information theory.

See the "Solid-state quantum computing hardware" sidebar for a discussion of the hardware side of the equation.

Secure Earth-to-space quantum communications

On 27 April 1986, a hacker known as Captain Midnight briefly took over a satellite television broadcast to the US's East Coast. This celebrated incident highlighted

Solid-state quantum computing hardware

Much progress has been made in designing quantum computer hardware. The early work on ion traps, cavity quantum electrodynamics, and nuclear magnetic resonance has inspired several schemes for implementing quantum computer hardware in solid-state quantum electronics. The move to solid-state quantum electronics is appealing because it builds upon decades of experience in microprocessor fabrication technology. These schemes use electric charge, magnetic flux, superconducting phase, electron spin, or nuclear spin as the information bearing degree of freedom.¹⁻⁵ Although each scheme has various advantages and disadvantages, the approach based on harnessing quantized charge is especially appealing because we can fabricate the necessary superconducting circuitry for such a qubit using present-day e-beam lithography equipment. Also, quantum coherence, essential for creating superposed and entangled states, has been demonstrated experimentally.⁶ These qualities make the quantized charge-based qubit a strong contender for the basic element of a proof-of-concept quantum computer. Figure B shows a schematic diagram for a charge-based qubit with ancillary readout circuitry.

The basic idea is that in a superconductor the electrons team up in pairs, called Cooper pairs. By varying voltages and magnetic fluxes, we can make several Cooper pairs hop onto an island in a superconductor. If there are N Cooper pairs, we say the island is in state $|0\rangle$. If there are $N + 1$ Cooper pairs, the island is in state $|1\rangle$. As Cooper pairs are quantum objects, we can create superpositions of N and $N + 1$ Cooper pairs on the island simultaneously, thereby creating arbitrary superposition states of the form $c_0|0\rangle + c_1|1\rangle$. This is a qubit, the elemental building block of a quantum memory register.

So how do we act on such qubits to make them change their state? To make a practical design for a quantum computer, we must specify how to decompose any valid quantum computation into a sequence of elementary 1- and 2-qubit quantum gates which can be realized in physi-

cal hardware that is feasible to fabricate. The set of these 1- and 2-qubit gates is arbitrary provided it is universal—capable of achieving any valid quantum computation from a quantum circuit comprising only gates from this set. Traditionally, we have taken the set of universal gates to be the set of all 1-qubit quantum gates in conjunction with a single 2-qubit gate called controlled-NOT. However, many equally good universal gate sets exist,⁷ and there might be an advantage in using a nonstandard universal gate set if certain gate designs happen to be easier to realize in one hardware context than another. Certainly, we have known for some time that the square root of the 2-qubit exchange-interaction (the SWAP gate) is as powerful as the better-known controlled-NOT gate (*CNOT*) as far as computational universality is concerned. It makes sense, therefore, to see what gates are easy to make and then extend them into a universal set. We pursued this strategy at the Jet Propulsion Lab. In particular, we showed, in the context of




Figure B. Schematic diagram of a single charge-based qubit with an adjoining RF SET readout.

the importance of controlling and avoiding denial-of-service attacks on orbital assets. In addition to maintaining the security of uplinked command paths, ensuring the security of downlinked data—which is susceptible to passive eavesdropping—can be equally important. Quantum key distribution (QKD) is an emerging technology based on single photon communications that will provide satellite communications with greater security and convenience than present-day methods.

How cryptographic methods work.

Cryptographic methods can assure the security of data and command paths to orbital assets as follows. We can encrypt message (“plaintext”) P according to some algorithm E before transmission to produce a “ciphertext,” $C = E_K(P)$, where K is a secret parameter known as a cryptographic key (a random binary number sequence, typically a few hundred bits in length). On receiving the ciphertext, the intended recipient can invert the encryption process using the decryption algorithm D to recover the original message $P = D_K(C)$, provided the secret key K is known. Conversely, although the encryption and decryption

algorithms E and D might be publicly known, an eavesdropper passively monitoring the transmission C could not discern the underlying message P because of the

randomization the encryption process introduced, provided the key K remains secret. Secret key material therefore is a very valuable resource that must be avail-




Figure 6. Very long baseline interferometers can be used for imaging of planets orbiting distant stars.




Figure C. A single qubit, implemented as a quantized charge device fabricated in aluminum using e-beam lithography. In a superconductor, electrons pair up as Cooper pairs. If there are N Cooper pairs on the center island, the island contains a $|0\rangle$ qubit. If there are $N+1$ Cooper pairs on the island, the island contains a $|1\rangle$ qubit. As electrons are quantum, we can have a superposition of N and $N+1$ Cooper pairs on the island simultaneously.

charge-based qubits, that any 1-qubit operation and a 2-qubit operation called the square root of complex SWAP (or $\sqrt{i\text{SWAP}}$) is universal for quantum computation.⁸ Such a quantum gate has been designed and fabricated at JPL. Figure C shows a photograph of the qubit with associated connections needed for performing gate operations.

It is not yet clear which approach to quantum computing hardware will prove to be the most feasible or cost effective. Moreover, deploying quantum computer hardware in space poses unique challenges for quantum hardware designers. Power supplies must be of limited duration and be nonrenewable, and radiation levels can be significantly higher than on Earth.

Quantum computers might help to overcome these problems. As we've noted, quantum computers are naturally reversible computers—the energy expended during computation is, in principle, recoverable. Though possibly surprising, that's a well-known thermodynamic result.⁹ In practice, quantum computers are unlikely to realize this reversible thermodynamic ideal. Nevertheless, it is sensible to start with a computer that is reversible, in principle, if we are to conserve precious energy resources. Moreover, if quantum computers can

answer certain computational questions by running a quantum algorithm that requires fewer steps than any classical algorithm for the same problem, this too could result in a net energy saving.

Quantum computers might also be better suited for operation in a radiation environment. Quantum-scale structures present smaller cross sections to incoming radiation than do conventional computer memory components. In addition, certain quantum systems, such as nuclear spins, are relatively immune to ionizing radiation that might damage conventional computer hardware.

References

1. D. V. Averin, "Adiabatic Quantum Computation with Cooper Pairs," *Solid State Communications*, Vol. 105, 1998, pp. 659–664.
2. L. Tian et al., "Decoherence of the Superconducting Persistent Current Qubit," 1999; preprint available at <http://xxx.lanl.gov/abs/cond-mat/9910062>.
3. G. Blatter, V. Geshkenbein, and L. Ioffe, "Engineering Superconducting Phase Qubits," 1999; <http://xxx.lanl.gov/abs/cond-mat/9912163>.
4. R. Vrijen et al., "Electron Spin Resonance Transistors for Quantum Computing in Silicon-Germanium Heterostructures," 1999; <http://xxx.lanl.gov/abs/quant-ph/9905096>.
5. B. Kane, "Silicon-Based Quantum Computation," 2000; <http://xxx.lanl.gov/abs/quant-ph/0003031>, submitted to *Fortschritte der Physik* Special Issue on Experimental Proposals for Quantum Computation.
6. Y. Nakamura, A. Pashkin, and J. S. Tsai, "Coherent Control of Macroscopic Quantum States in a Single-Cooper-Pair Box," *Nature*, Vol. 398, 1999, pp. 785–787.
7. D. Di Vincenzo, "Two Bit Gates are Universal for Quantum Computation," *Physical Review A*, Vol. 51, No. 2, Feb. 1995, pp. 1015–1022.
8. P. Echehnach et al., "Universal Quantum Gates for Single Cooper Pair Box Based Quantum Computing," *Proc. Progress in Electromagnetics Research Symp., Session on Solid-State Quantum Computing Hardware*, PIERS, Washington, D.C., 2000.
9. C. H. Bennett, "Logical Reversibility of Computation," *IBM J. Research and Development*, Vol. 17 (1973) pp. 525–532.

able on the satellite and at the ground station. Moreover, frequent key changes are necessary to ensure security. So, if the key material supplied at launch should be used up during normal operations, methods for secure key distribution to satellites on-orbit are of paramount importance to ensure that third parties cannot acquire even partial knowledge of the new key.

It is obviously impractical to send a courier to a satellite, and key transmissions themselves must be assumed susceptible to passive eavesdropping. Public-key encryption methods, which derive their security from the perceived difficulty of certain mathematical problems such as factoring large integers, can be conveniently used to securely distribute new keys by broadcast, but are subject to increasing challenges, including unanticipated advances in computational techniques, technology, and algorithms. For example, there have recently been rapid advances in the size of integers that have been factored with Internet collab-

orations,¹³ and new proposals for special purpose, high-speed factoring hardware have emerged.¹⁴ Furthermore, because it is notoriously difficult to assess an adversary's future computational capabilities accurately, today's passively monitored public-key transmissions could become retroactively vulnerable well before their intended security lifetime elapses.¹⁵ If large-scale quantum computation becomes feasible, essentially all public-key methods will become vulnerable.^{3,16}

QKD is unconditionally secure, no matter what present or future technology an adversary might possess; its security is assured by laws of Nature.¹⁷ A QKD procedure starts with the sender Alice generating a secret random binary number sequence R_A . For each bit in the sequence, Alice prepares and transmits a single photon over a quantum channel (a low-loss, faithful transmission medium) to the recipient Bob (see Figure 7). He measures each arriving photon and attempts to identify the bit

value Alice has transmitted. Alice's photon state preparations and Bob's measurements are chosen from sets of nonorthogonal possibilities. For example, using the so-called B92 protocol¹⁸ Alice agrees with Bob (through public discussion) that she will transmit a vertically polarized single-photon state for each 0 in her sequence, and a 45 degree-polarized single-photon state for each 1 in her sequence.

Bob agrees with Alice to test the polarization of each arriving photon for horizontal polarization to reveal 1s, or negative 45-degree polarization to reveal 0s. Bob's choices of polarization are set by random bit values from a secret random binary sequence R_B , which he generates and proceeds through in synchronization with Alice. In this scheme, Bob will never detect a photon for which he and Alice have used a preparation and measurement pair that corresponds to different bit values, such as horizontal and vertical polarizers, which happens for 50% of the bits in Alice's sequence.

However, for the other 50% of Alice's bits, the preparations and measurements use nonorthogonal polarizations, such as vertical and negative 45 degrees, resulting in a quantum-mechanically random 50% detection probability for Bob on this portion. Thus, by detecting single photons, Bob identifies a random 25% portion of the bits in Alice's random sequence, assuming she transmits exactly one photon for each bit and there are no bit losses in transmission or detection. This 25% efficiency factor is the price that Alice and Bob must pay for secrecy. (In practice there will be additional losses, but photons that fail to reach the receiver merely reduce the key rate without leaking any information to adversaries.) Bob and Alice reconcile their common bits through discussion over a public channel with Bob revealing the locations, but not the bit values, in the sequence where he detected photons; Alice retains only those detected bits from her initial sequence.

The resulting detected bit sequences comprise the raw key material from which a pure key is distilled using classical error detection techniques. Because the key does not exist until after the quantum transmissions are complete there is no prior record of a key that could be compromised by insiders.

Because a photon is an indivisible elementary particle, the QKD transmissions cannot be passively tapped in the conventional sense, so adversaries would need to undertake far more risky active attacks. However, Heisenberg's Uncertainty Principle ensures that any active attack will not permit an adversary to faithfully read the key transmissions. Moreover, such an attempt will inevitably disturb the transmissions, letting the intended users detect the attempted eavesdropping, as well as put a rigorous upper bound on the information that might have been leaked. With this bound, Alice and Bob can apply the technique of *privacy amplification* in which they agree (by public discussion) to produce a new bit sequence formed from the parities of suitably chosen random subsets of their reconciled bit sequences, which ensures that any adversary could do no better than guess the resulting bit sequence. The keys produced at the end of this procedure will be secret. They could be used to initialize conventional cryptographic hardware, typically requiring a few hundred bits.

Transmitting secure signals. For successful ground-to-satellite key generation with

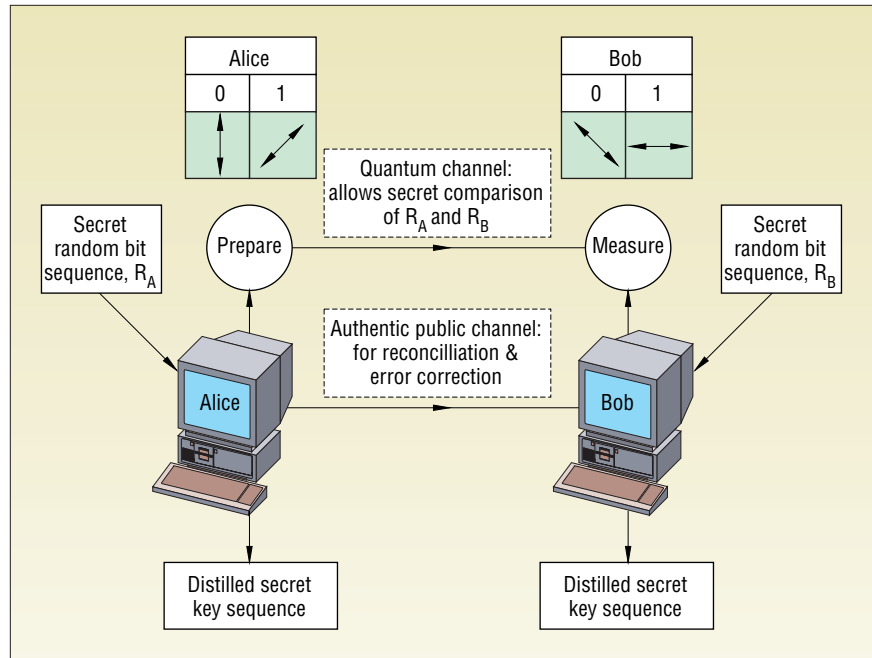


Figure 7. The B92 quantum-key distribution protocol.

QKD, we will need reliable single-photon transmission and detection through a turbulent atmosphere against a high background. For good atmospheric transmission, free-space QKD can operate at a wavelength near 770 nm where the transmission from surface to space can be as high as 80%, and the polarized QKD photons will faithfully transmit because the depolarizing effects of Faraday rotation in the ionosphere and of atmospheric turbulence are negligible. We can readily produce photons at this wavelength with rugged, low-power semiconductor lasers, control their properties, and detect them with efficiencies as high as 65% with off-the-shelf components.

Atmospheric turbulence will cause beam wander, which we can overcome using the optical beam-control techniques developed in free-space laser communications for high-bandwidth terrestrial, surface-to-satellite, satellite-to-satellite, and (potentially) deep-space communications. QKD is compatible with and can take advantage of the optical techniques developed for this new communications infrastructure.

At Los Alamos, ground-based experiments have shown that the single optical photons of QKD can faithfully transmit through a turbulent atmosphere and be reliably detected even against a daylight background over a point-to-point 1.6-km path.²⁰ This distance was limited only by the lengths of the available ranges. Because most of the optical effects of atmospheric turbulence on a surface-to-space path occur within 2 km of the ground, these results demonstrate that

QKD could effectively serve to securely rekey a satellite on-orbit from a ground station (or for satellite-to-satellite key generation).²¹ For illustration, we can estimate the key generation capability of QKD between a ground station and a low Earth orbit (LEO) satellite (~600 km altitude) in one overhead pass. We will assume that the QKD transmitter (Alice) is at the ground station and the receiver (Bob) is on the satellite.

To estimate the key generation rate, we can assume 20-cm diameter optics at both the transmitter and satellite receiver. Beam-wander from atmospheric turbulence at night at a typical optical communications ground station can be 1 to 5 arc seconds, but for this analysis we assume a worst case "seeing" of ~10 times the diffraction limit (10 arc seconds of wander). With a laser pulse rate of 10 MHz, one photon-per-pulse on average, and atmospheric transmission and detector efficiencies as above, a key generation rate of ~500 Hz should be feasible. Higher key rates would be possible under more typical seeing conditions. Also, with a beam fine-pointing control system, as used in laser communications systems, the beam could be locked onto the satellite, increasing the key rate to ~40 kHz.

It would also be possible to place the QKD transmitter on the satellite and the receiver on the ground. Because most of the optical influence of atmospheric turbulence would occur in the final ~2 km of the beam path, a higher key rate would be possible. In either case, the bit error rate (BER) from background photons would be



Richard Hughes is a Laboratory Fellow and Quantum Information Science team leader in the Physics Division at Los Alamos National Laboratory. He is principal investigator of several projects in quantum computation and quantum cryptography. Richard obtained his PhD in theoretical elementary particle physics from the University of Liverpool, England, and has held research positions at Oxford University and The Queen's College, Oxford; the California Institute of Technology; and CERN, Geneva, Switzerland. In 1996 and 1998 he was awarded Los Alamos Distinguished Performance Awards for his quantum cryptography research, and in 1997 he was awarded the Los Alamos Fellows' Prize for his work on quantum information science. He became a Fellow of the American Physical Society in 1999. In his spare time he competes in ultra running events in excess of 100km. Contact him at hughes@lanl.gov.




Colin P. Williams is a principal scientist at the Jet Propulsion Laboratory and an acting associate professor of computer science at Stanford University. He is also chief scientific officer of Quantum Confidential Inc., a company commercializing quantum technologies. His current research interests include quantum computing, quantum cryptography, quantum lithography, computational phase transitions, and AI. He holds a PhD in artificial intelligence from the University of Edinburgh, an MSc in atmospheric physics from Imperial College, University of London, and a BSc in mathematical physics from the University of Nottingham. He was a former research assistant to Stephen Hawking and research scientist at Xerox PARC. Contact him at the Jet Propulsion Lab., California Inst. of Technology, Pasadena, CA 91109; colin.p.williams@jpl.nasa.gov.

no worse than the few percent level seen in the Los Alamos ground-based experiments.²⁰ Although such BERs are very much larger than those acceptable in conventional communications, they are tolerated in QKD because of the ensured secrecy of the transmitted bits. Interactive error-correction methods for removing all such errors are available for use in QKD.

From these simple analyses, we see that during the several minutes that a satellite would be in view of the ground station, there would be adequate time to acquire the satellite, perform the QKD transmissions for ~1 minute, and produce a minimum of ~10,000 raw bits, from which a shorter error-free key stream of several thousand bits would be produced after error correction and privacy amplification. Under more typical seeing conditions or with beam fine pointing implemented, we could produce up to 10^5 secret key bits in a 1-minute QKD transmission. So, multiple new cryptographic keys could be generated between a ground station and a LEO satellite in one overhead pass using available technology. Satellite QKD could also serve to provide secure key distribution to two ground-based users (Alice and Bob): they could each generate independent quantum keys with the same satellite, which would then transmit the XOR of the two keys to Bob. Bob would then XOR this bit string with his key to produce a key that agrees with Alice's. Alice and Bob could then use their shared key for encrypted communications over any convenient channel.

Based on these feasibility arguments, QKD will be capable of providing the ulti-

mate level of security for future generations of satellites.

 **QUANTUM COMPUTING AND COMMUNICATIONS** are two aspects of the new field of quantum information theory. This field is still in its infancy but is already promising to have a major impact on space exploration. Quantum computing offers an alternative and relatively unexplored route to increased computing capacity for onboard autonomy and machine intelligence. Quantum algorithms let us solve old problems in new ways, sometimes exponentially faster than possible with conventional computers. Quantum communications could enable unconditionally secure communications with manned spacecraft and let us encrypt sensitive data en route to and from satellites. New and unforeseen applications of quantum technologies are being discovered every year. ■

References

1. C.P. Williams and S. Clearwater, *Ultimate Zero & One: Computing at the Quantum Frontier*, Copernicus Books, New York, 1999.
2. C.P. Williams and S. Clearwater, *Explorations in Quantum Computing*, Springer-Verlag, New York, 1998.
3. P.W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Algorithms on a Quantum Computer," *SIAM J. Computing*, Vol. 26, No. 1484–1509, 1997.

4. D.S. Abrams and S. Lloyd, "A Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors," *Physical Review Letters*, Vol. 83, 1999, pp. 5162–5165.
5. R. Paturi et al. "An Improved Exponential-Time Algorithm for k-SAT," *Proc. IEEE 39th Symp. Foundations of Computer Science*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1998, pp. 628–637.
6. N.J. Cerf, L.K. Grover, and C.P. Williams "Nested Quantum Search and Structured Problems," *Physical Rev. A*, Vol. 61, 2000.
7. L.K. Grover, "A Fast Quantum Mechanical Algorithm for Database Search," *Proc. 28th ACM Ann. Symp. Theory of Computing*, ACM Press, New York, 1996, pp. 212–219.
8. C. Zalka, "Grover's Quantum Searching Algorithm is Optimal," *Physical Rev. A*, Vol. 60, 1999, pp. 2746–2751.
9. R. Gingrich, C.P. Williams, and N.J. Cerf, "Generalized Quantum Search with Parallelism," *Physical Review A*, Feb. 2000.
10. H. Buhrman et al., "Multipartite Quantum Communication Complexity," 1997; <http://xxx.lanl.gov/abs/quant-ph/9710054>.
11. C.H. Bennett and S.J. Wiesner, *Physical Rev. Letters*, Vol. 69, 1992, p. 2881.
12. R. Jozsa et al., "Quantum Clock Synchronization Based on Shared Prior Entanglement," *Physical Review Letters*, Vol. 85, No. 9, 2000.
13. S. Cavallar et al., "Factorization of a 512-Bits RSA Modulus," *LNCS*, Vol. 1807, 1999, p. 1.
14. A. K. Lenstra and A. Shamir, "Analysis and Optimization of the TwinkleFactoring Device," *LNCS*, Vol. 1807, 1999, p. 35.
15. D. Atkins et al., "The Magic Words are Squeamish Ossifrage" *LNCS*, Vol. 917, 1994, p. 265.
16. R.J. Hughes, "Cryptography, Quantum Computation and Trapped Ions," *Philosophical Trans. Royal Soc.*, Vol. A356, p. 1853, 1998.
17. R.J. Hughes et al., "Quantum Cryptography," *Contemporary Physics*, Vol. 36, 1995, p. 149.
18. R. J. Hughes and J. E. Nordholt, "Quantum Cryptography Takes to the Air," *Physics World*, May 1999, p. 31.
19. W.T. Buttler et al., "Daylight Quantum Key Distribution Over 16 km," *Physical Rev. Letters*, Vol. 84, 2000, p. 5652.
20. R.J. Hughes et al., "Free-Space Quantum Key Distribution in Daylight," *J. Modern Optics*, Vol. 47, 2000, p. 549.