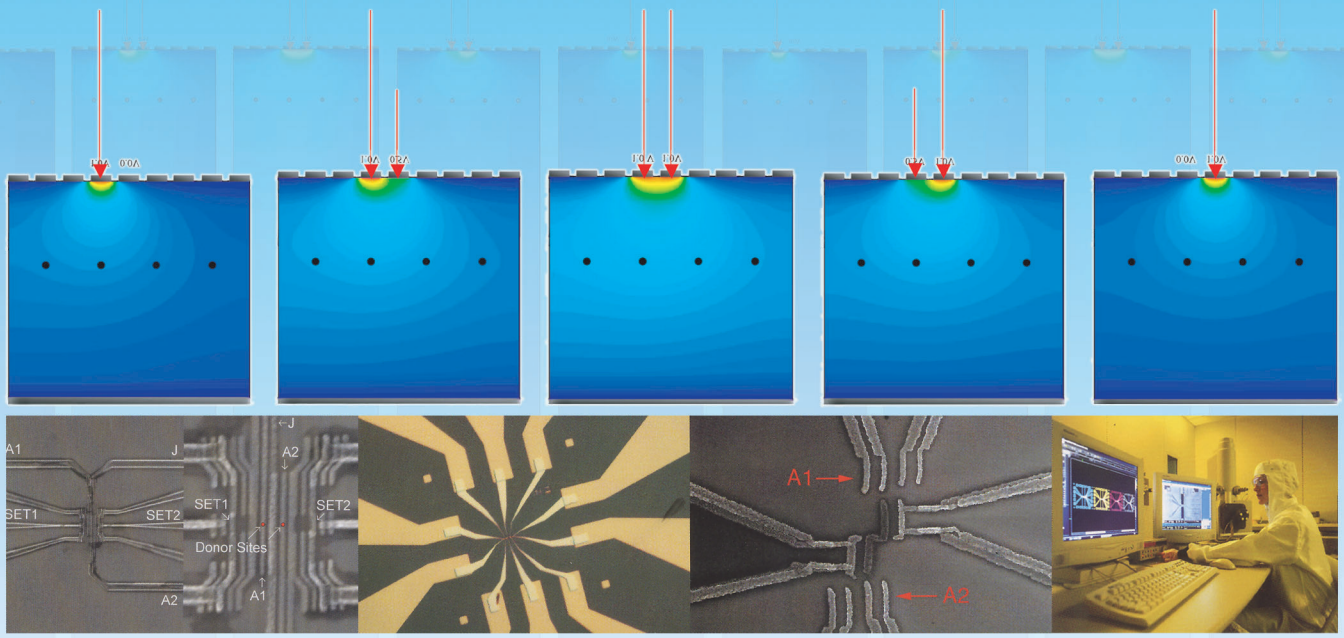

Toward a Silicon-Based Nuclear-Spin Quantum Computer

Developing the technology for a scalable, solid-state quantum computer

Robert G. Clark, P. Chris Hammel, Andrew Dzurak, Alexander Hamilton, Lloyd Hollenberg, David Jamieson, and Christopher Pakes as told to Jay Schecker

One of the major challenges in quantum computing is to identify a system that can be scaled up to the number of qubits needed to execute nontrivial quantum algorithms. Peter Shor's algorithm for finding the prime factors of numbers used in public-encryption systems (numbers that typically consist of more than a hundred digits) would likely require a quantum computer with several thousand qubits. Depending on the error correction scheme appropriate to the particular computer, the required number could be much larger. Although ion-trap or nuclear-magnetic-resonance (NMR) quantum "computers" containing a few (less than 10) qubits have been successfully demonstrated, it is not clear that those systems can be scaled up.

Solid-state quantum computers may be more likely candidates. As a result, a number of solid-state schemes have been proposed, most of which fall into two categories: The physical qubits are spin states of individual electrons or nuclei, or they are charge or phase states of superconducting structures.

A particularly promising scheme is the silicon-based nuclear-spin computer, proposed a few years ago by Bruce Kane (1998), then of the University of New South Wales in Sydney, Australia, and now of the University of Maryland in College Park, Maryland. Shown in Figure 1, the Kane computer architecture consists of a linear array of phosphorus atoms embedded beneath the surface of a pure silicon wafer. Each phosphorous atom serves as a qubit, and the linear array forms the computer's quantum register.[1] The entire wafer is placed in a strong, static magnetic field $\mathbf{B}_0$ at sub-Kelvin temperatures, and alignment of the phosphorous atom's nuclear spin as parallel or antiparallel to $\mathbf{B}_0$ corresponds to the $|0\rangle$ and $|1\rangle$ states of the qubit, respectively. (Throughout this article, we will follow the convention of Kane

[1] Two-dimensional qubit arrays are also possible but require complex electrode geometries and additional insulators.

and use the notation $|0\rangle$, $|1\rangle$ to designate both qubit and nuclear spin states. We will use arrows, $|\downarrow\rangle$ or $|\uparrow\rangle$, to designate electron spin states.)

In order to execute a quantum algorithm, we need to rotate individual qubits in Hilbert space and couple two qubits together. We accomplish both operations with an array of gate electrodes[2] that lies on top of the wafer but is isolated from the pure silicon by a thin insulating layer of silicon dioxide (SiO_2). Referring to Figure 1, notice that each A-gate sits precisely above a phosphorous atom and each J-gate lies between adjacent atoms. As discussed later, a small positive voltage applied to the A-gate gives independent control of the qubit directly under the gate, while voltage applied to the J-gate allows coupling two qubits together through an electron-mediated interaction.

The Center for Quantum Computer Technology (CQCT), headquartered in Sydney, Australia, and Los Alamos National Laboratory are trying to fabricate Kane's silicon-based quantum computer. Although we can call upon the technology, techniques, and collective experience of the huge silicon semiconductor industry, building the computer is still a daunting enterprise. We need to manipulate individual phosphorus atoms and place them precisely within a defect-free, isotopically pure silicon matrix. We need to create metal gates on the nanoscale, lay them within a few atoms of each other, and then ensure that each gate is aligned properly over the buried qubits. Simply put, the ability to do this level of nanofabrication does not exist at this time.

Employing a "see-what-works-best" strategy, we have initiated parallel research approaches for nearly every fabrication stage. If one approach fails, we still have a backup. Our current focus is on developing a prototype

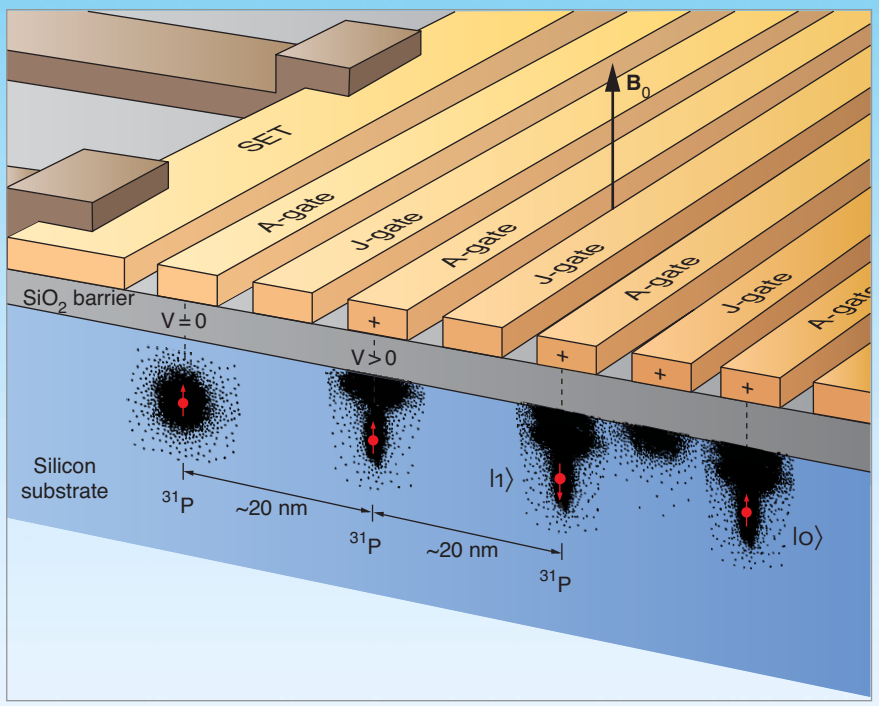

Figure 1. Schematic of the Silicon-Based Quantum Computer

The architecture of Kane's solid-state quantum computer has a linear array of phosphorous donor atoms buried into a pure silicon wafer, with an interdonor spacing of about 20 nm. In the presence of a large magnetic field and at sub-Kelvin temperatures, the nuclear spins of the donor atoms can be aligned either parallel or antiparallel with the field, corresponding to the $|0\rangle$ and $|1\rangle$ qubit states, respectively. An array of metal gates lies on the surface of the wafer, isolated from the silicon by a barrier layer of SiO_2. The A-gates lie directly above the donor atoms and enable individual control of single qubits. The J-gates lie between the donors and regulate an electron-mediated coupling between adjacent nuclear spins, thus allowing for two-qubit operations. Readout of the qubit is performed with either a single electron transistor, as shown, or with a magnetic-resonance force microscope (MRFM, not shown). The electron clouds shown here are schematic—the actual situation is more complicated.

two-qubit device. By stretching many technologies beyond their heretofore-assumed limits, we have come tantalizingly close to achieving that goal. In the sections that follow, we describe the computer and some critical technologies in greater detail, and we also outline our progress in building a prototype.

Design Features of the Silicon-Based Computer

In our solid-state architecture, individual phosphorus atoms are embedded in a sea of silicon. These two elements were chosen for several reasons, the first and foremost being that phosphorous is the standard dopant for conventional silicon-based semiconductor devices and there is tremendous working knowledge of phosphorus and silicon from the conventional computing industry.

The second reason stems from the need to control the spin environment. We require our qubits to have nuclear spin $I = 1/2$, but we also want the surrounding environment to be spin free. Otherwise, unwanted spin-spin interactions between the qubit and any nearby nuclear spins would compromise the coherent states needed for

[2] In this context, a "gate," like a transistor gate, is a device used for controlling charge motion. It does not refer to a logical operation such as a **cnot** gate.

quantum computation. The only stable phosphorous isotope, phosphorous-31, is a spin-1/2 nucleus, so nature has automatically satisfied our qubit criterion. Creating a spin-free environment, however, is a little more difficult. Natural silicon contains a mixture of three isotopes: silicon-28, -29, and -30. Whereas the even-numbered isotopes are spin free, silicon-29 has a spin of $I = 1/2$. As a result, we estimate that to do quantum computation, we will need to reduce the silicon-29 content in our silicon wafer to about one part in 10^5. Those stringent isotopic purity levels can be reached with current technology.

Finally, the nuclear spin of a phosphorous atom in a silicon host is stable. One way to infer the stability is to examine the spin-lattice relaxation time T_1, which characterizes the time it takes for a spin system with a net alignment (a net magnetization) to return to its thermally equilibrated magnetization. At temperatures of about 1 kelvin, the nuclear-spin relaxation time $T_{1,n}$ (where the subscript "n" is for the nucleus) for phosphorus in silicon is expected to be longer than 10 hours (Feher 1959).

The nuclear spin qubits, however, interact with the environment through their donor electrons; as a result, the electron spin stability is also important, particularly for qubit readout (see later discussion). The electron-spin relaxation time $T_{1,e}$ (where the subscript "e" is for the electron) is approximately 1 hour at about 1 kelvin (Honig and Stupp 1960). The phase decoherence time of the electron spin, as measured by the spin-spin relaxation time $T_{2,e}$, is shorter still. Although never measured for an isolated electron system such as our qubit scheme, the $T_{2,e}$ for an ensemble of electrons was measured to be approximately 0.5 millisecond (Gordon and Bowers 1958). A recent theoretical study, appropriate for a single phosphorus donor atom in silicon, indicates a $T_{2,e}$ of the order of 1 second (Mozyrsky et al. 2002). This value for $T_{2,e}$ should be long enough for us to perform a quantum computation and read out the result.

The Spin Hamiltonian and Single-Qubit Operations. To understand the physics underlying the operation of the silicon-based computer, recall that phosphorus has one more electron in its outer electron shell than silicon. When it is placed into a silicon crystal lattice, phosphorus fulfills its role as a surrogate silicon atom and still has one electron left over. At very low temperatures, that "donor" electron remains bound—albeit rather loosely—to the phosphorus nucleus. The electron "talks" to the nucleus primarily through the charge (Coulomb) interaction and to a lesser degree through the hyperfine interaction, which is between the electron spin and the nuclear spin.

We exploit the hyperfine interaction to individually address single qubits. The effective low-energy, low-temperature Hamiltonian describing the spin interactions for the electron spin and the nuclear spin of an atom in the presence of a static magnetic field $\mathbf{B}_0$ is given by

$$H = \mu_B \mathbf{B}_0 \sigma_z^e - g_n \mu_n \mathbf{B}_0 \sigma_z^n + A\sigma^e \cdot \sigma^n \ , \quad (1)$$

where σ_z^e and σ_z^n are Pauli spin matrices, μ_B and μ_n are the Bohr and nuclear magnetons, respectively, and g_n is the nuclear g-factor. The hyperfine interaction is described by the term $A\sigma^e \cdot \sigma^n$.

For large values of $\mathbf{B}_0$, the Hamiltonian in Equation (1) leads to a set of energy levels that correspond to the four hyperfine levels, $|0\downarrow\rangle$, $|1\downarrow\rangle$, $|0\uparrow\rangle$, and $|1\uparrow\rangle$. At the sub-Kelvin operating temperature of the computer, however, the electron spins are completely spin-polarized in the lower-energy state $|\downarrow\rangle$. Thus, for one-qubit operations, we may ignore the electron spin polarization to a good approximation and consider only the two nuclear states $|0\rangle$ and $|1\rangle$ (Goan and Milburn 2000). The energy difference between those states is

$$\Delta E_0 = \hbar\omega_0 \cong 2g_n \mu_n \mathbf{B}_0 + 2A + 2A^2/\mu_B \mathbf{B}_0 \ , \quad (2)$$

where ω_0 is called the nuclear resonance frequency. The resonance frequency, which is typically in the radio-frequency (rf) range, is equal to the Larmor frequency, or the rate at which the nuclear spins precess about the magnetic-field lines.

Suppose $\mathbf{B}_0$ is aligned along the z-axis, and the nuclear spin is initially in the $|0\rangle$ state. If we subject the spins to a secondary magnetic field that is oscillating in the x-direction at the nuclear resonance frequency, that is, a field $\mathbf{B}_1 = B_1 \cos(\omega_0 t)\, \hat{x}$, then the nuclear spins will begin to rotate toward the $(-z)$-axis, or from the parallel to the antiparallel alignment (see the box "Spin Manipulation with Magnetic Resonance" on page 288). The spin rotation is equivalent to rotating a qubit in Hilbert space from the $|0\rangle$ state to some superposition of the $|0\rangle$ and $|1\rangle$ states.

As described, the $\mathbf{B}_1$ field will affect all spins simultaneously. To address a particular spin, we use the A-gate directly above it and modify that donor atom's hyperfine coupling. The parameter A in Equation (1) is proportional to the magnitude of the electron probability density at the site of the nucleus, $\Psi_e(0)$:

$$A = 8/3\pi\, \mu_B\, g_n\, \mu_n |\Psi_e(0)|^2 \ . \quad (3)$$

As seen in Figure 2, placing a positive voltage on the A-gate above the phosphorous atom attracts the atom's electron cloud toward the surface and away from the nucleus, thereby reducing the magnitude of $\Psi_e(0)$. The hyperfine

energy levels of that one atom change slightly, and the resonance frequency needed to rotate the nuclear spin is reduced from, say, ω_0 to $\omega\_$. If the frequency of the $\mathbf{B}_1$ field is set to $\omega\_$, that is, $\mathbf{B}_1 = B_1 \cos(\omega\_ t)\,\hat{x}$, then only the spin directly under the A-gate will be in resonance and will begin to rotate. Removing the voltage on the A-gate halts the rotation.

A one-qubit gate operation is therefore implemented if the silicon wafer is subjected to a transverse oscillating B-field of frequency $\omega\_$ and if the A-gate above a qubit is pulsed for a defined period. Throughout the duration of the pulse, the qubit is in resonance with the secondary magnetic field and rotates through some angle in Hilbert space. When the voltage is removed at the end of the pulse, the qubit is left in the desired superposition of the $|0\rangle$ and $|1\rangle$ states.

Two-Qubit Operations. To select adjacent pairs of qubits for two-qubit operations, we apply a positive voltage to the J-gate between them. As seen in Figure 3(a), this procedure causes the electron wave functions of the two donor atoms to overlap, and the electron spins couple together through the electron-spin exchange interaction. Because each electron is coupled to its nucleus through the hyperfine interaction, turning on the electron-spin exchange interaction also couples the nuclear spins together.

The coupled nuclear-electron spin system is fairly complex, but we can gain insight into it by looking at the effective Hamiltonian for the system:

$$H_{\text{coupled}} = H^1 + H^2 + J\sigma^{1e}\cdot\sigma^{2e} \quad . \qquad (4)$$

This Hamiltonian is valid at energy scales that are small compared with the electron-binding energies of the donor atoms. The first two terms correspond to the hyperfine Hamiltonian—Equation (1)—of each donor, respectively, and the last term accounts for the spin exchange

Continued on page 290

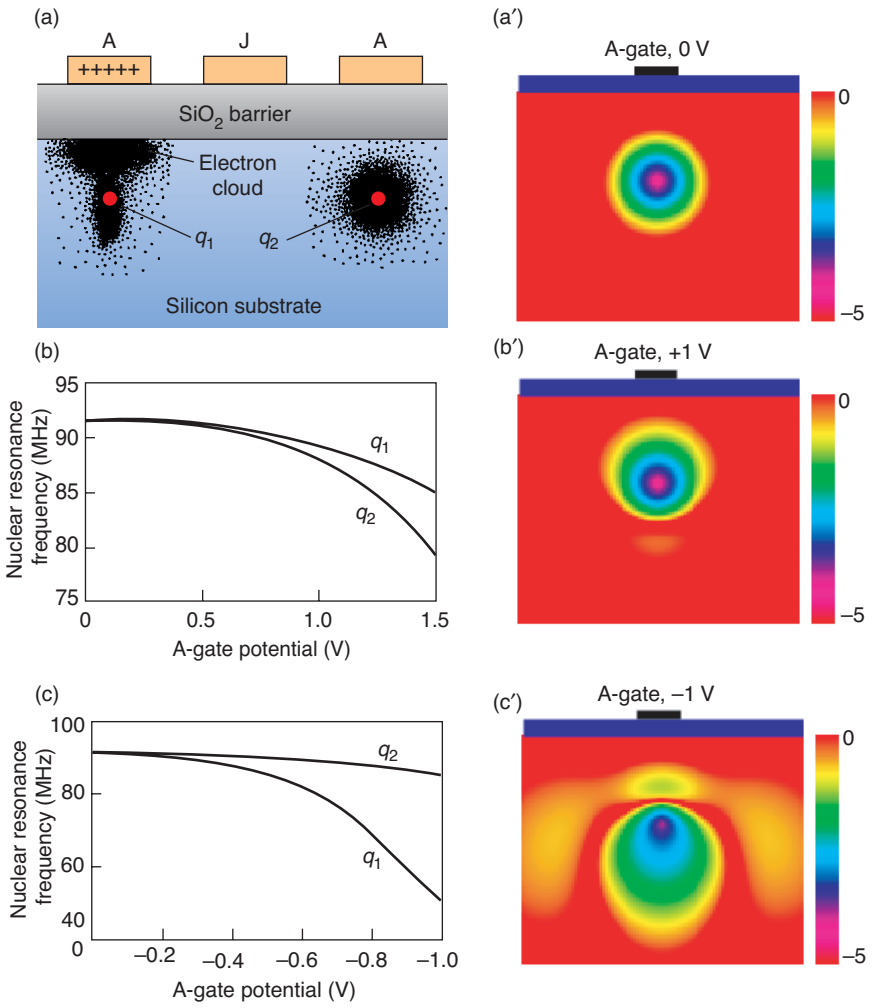

Figure 2. A-Gate Control of One Qubit

We use magnetic resonance techniques to rotate nuclear spins and place them in arbitrary superpositions of $|0\rangle$ and $|1\rangle$ qubit states. (a) In this cartoon, a small voltage is applied to the A-gate directly above a qubit. The donor electron moves away from the ^{31}P nucleus. (a′) This plot of the electron probability density surrounding a donor atom with no voltage on the A-gate was obtained by solving the Schrödinger equation nonperturbatively in an isotropic effective-mass hydrogenic basis. The plot is a cross section through the nucleus, with the color variations on a logarithmic scale. The atom is buried 20 nm below the Si/SiO$_2$ interface. (b) The graph shows the variation of the nuclear transition frequency as a function of A-gate voltage. (b′) The color plot shows the electron probability density. A positive voltage on the A-gate pulls the electron away from the nucleus, thus reducing the hyperfine coupling—the parameter A in Equation (1) in the text. In a B-field of about 2 T, the resonance frequency of a phosphorous nucleus q_1 is $\nu_0 = \omega_0/2\pi \approx 93$ MHz. With a gate voltage of 1 V, the resonance frequency of q_1 reduces to about $\nu\_ = \omega\_/2\pi \approx 90$ MHz, while a neighboring nucleus q_2 is in resonance at about 87 MHz. (The proximity of the oxide barrier has a fairly large effect on the qubits, and the positive gate voltage affects q_2 more than q_1.) Subjecting the silicon wafer to a transverse, oscillatory magnetic field of frequency $\nu\_$ would cause only q_1 to respond. (c)–(c′) Initial calculations indicate that the electron probability density is more responsive to a negative gate bias, which results in better frequency discrimination between adjacent qubits.

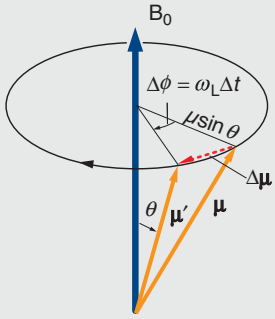

$$|\boldsymbol{\tau}| = |\boldsymbol{\mu} \times \mathbf{B}_0| = \mu B_0 \sin\theta$$
$$\boldsymbol{\tau} \cong \Delta\mathbf{J}/\Delta t = \gamma^{-1}\Delta\boldsymbol{\mu}/\Delta t$$
$$|\boldsymbol{\tau}| = \gamma^{-1}(\mu\sin\theta\ \omega_L \Delta t)/\Delta t$$
$$= \gamma^{-1}\mu\sin\theta\ \omega_L$$
$$\gamma^{-1}\mu\sin\theta\ \omega_L = \mu B_0 \sin\theta$$
$$\Rightarrow \omega_L = \gamma B_0$$

Figure A. Larmor Precession
Magnetic moments precess around magnetic field lines at the Larmor precession frequency ω_L, which is derived above.

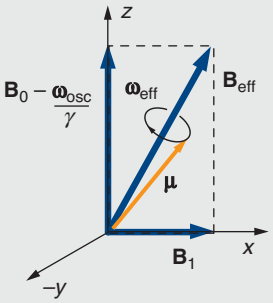

Figure B. Effective Magnetic Field in the Rotating Frame
The motion of the moment about $\mathbf{B}_{eff}$ is easier to describe in the frame rotating about the *z*-axis at the same frequency ω_{osc} as the oscillating field $\mathbf{B}_1$, since then $\mathbf{B}_1$ is static. $\mathbf{B}_0$ is reduced by the amount ω_{osc}/γ.

Spin Manipulation with Magnetic Resonance

Magnetic resonance is the traditional technique for detecting and manipulating any particle, such as an electron, atom, or nucleus, that has a magnetic moment $\boldsymbol{\mu}$. The manipulation is controlled by a combination of static and oscillating magnetic fields. Classically, a particle's magnetic moment is proportional to its angular momentum $\mathbf{J}$ through the relation

$$\boldsymbol{\mu} = q/2m\ \mathbf{J}\ , \tag{1}$$

where q is the charge of the particle and m is its mass. Remarkably, a similar relation holds true in quantum mechanics, although we must also take into account the particle's intrinsic spin angular momentum. In general, we can write

$$\boldsymbol{\mu} = \gamma\mathbf{J}\ , \tag{2}$$

where the parameter γ is known as the gyromagnetic ratio. It is related to the constants in equation (1) by a dimensionless constant known as the g-factor,

$$\gamma = \mathrm{g}\ (q/2m)\ . \tag{3}$$

The magnitude and sign of the g-factor depend on the specific atom or nucleus, but are always approximately 1.

In the classical picture of a randomly oriented moment in a magnetic field $\mathbf{B}_0 = B_0\hat{z}$, the moment would like to lower its energy by aligning itself parallel to the applied field. But the magnetic field can only produce a torque on the moment, $\tau = \boldsymbol{\mu} \times \mathbf{B}_0$. Because the torque is directed perpendicular to the plane defined by the field, the moment does not align with the field, but like a spinning gyroscope that resists the force of gravity, precesses around the magnetic field line. By using the fact that the torque is equal to the rate of change of the angular momentum, we can derive the angular precession frequency of the moment (see Figure A):

$$\omega_L = \gamma B_0\ , \tag{4}$$

where ω_L is called the Larmor frequency and is measured in radians per second. Equation (4) is the single most important equation of magnetic resonance. It says that the frequency of precession about a magnetic-field line is proportional to both the magnitude of the magnetic field and the gyromagnetic ratio. Interestingly, as derived in the equations accompanying the figure, the frequency is independent of the angle θ that specifies the orientation of the magnetic moment. The Larmor frequency enables us to identify the particle because the gyromagnetic ratio is distinct for electrons and different nuclei. The Larmor frequency $(\omega_0/2\pi)$ for an electron is about 28 gigahertz per tesla (MHz/T) and for a proton, roughly 45 MHz/T.

The moments precessing in the applied field can be manipulated in several ways. One common method is pulsed magnetic resonance. For a short period, we apply an oscillating magnetic field along the x-axis, $\mathbf{B}_1 = B_1\cos(\omega_{osc}t)\,\hat{x}$, where $B_1 << B_0$. The moment will begin to precess around a time-dependent total magnetic field consisting of $\mathbf{B}_0$ plus $\mathbf{B}_1$. This complicated motion can be better understood by examining the moment in a reference frame that rotates around the z-axis with frequency ω_{osc}, the same frequency as $\mathbf{B}_1$. In the rotating frame, $\mathbf{B}_1$ becomes a static field and the precession frequency about the z-axis is reduced: $\omega_L \rightarrow \omega_L - \omega_{osc}$.

Phenomenologically speaking, in the rotating frame the magnetic moment "sees" an effective field of magnitude

$$\mathbf{B}_{\text{eff}} = \left(B_0 - \frac{\omega_{\text{osc}}}{\gamma} \right) \hat{z} + B_1 \hat{x}$$
$$= \frac{1}{\gamma} \left(\omega_{\text{L}} - \omega_{\text{osc}} \right) \hat{z} + B_1 \hat{x} \quad , \qquad (5)$$

which is illustrated in Figure B. Equation (5) tells us that, when the frequency of the $\mathbf{B}_1$ equals the Larmor frequency, namely, at the resonance condition $\omega_{\text{L}} = \omega_{\text{osc}}$, the effective field has no z-component. Only the $\mathbf{B}_1$ field remains, and the moment will precess around the x-axis at an angular frequency ω_1 set by the magnitude of $\mathbf{B}_1$, namely, $\omega_1 = \gamma B_1$. Thus in the laboratory, we can rotate a moment about the x-axis by setting the frequency of $\mathbf{B}_1$ to the Larmor frequency. We control the rate of rotation by adjusting the field strength and the amount of rotation by restricting the length of time that the $\mathbf{B}_1$ field is applied.

Pulsed magnetic resonance can be used to manipulate a qubit. Suppose a qubit state is defined by the nuclear spin orientation such that the spin aligned parallel to $\mathbf{B}_0$ represents the state $|0\rangle$ whereas the spin aligned antiparallel to the field represents the state $|1\rangle$. We send a current pulse through an inductive coil to create the field B_1. If the pulse is timed to last for one-half of a precession period, or $t = \pi/\omega_1$, then the spins will rotate around the x-axis for π radians, or by 180°. If the qubit was initially in the $|0\rangle$ state, it would now be in the $|1\rangle$ state. Similarly, we can pulse the current for a time $t = \pi/(2\omega_1)$—a so-called $\pi/2$ pulse—and rotate the qubit into an equal superposition of the $|0\rangle$ and $|1\rangle$ states, namely the state $1/\sqrt{2}\ (|0\rangle + |1\rangle)$. (See Figure C.)

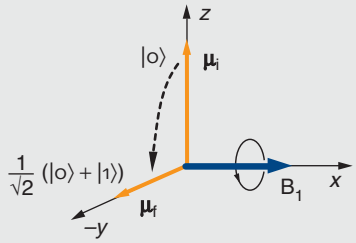

Figure C. Pulsed Magnetic Resonance

When ω_{osc} is made equal to ω_{L}, a moment will begin to rotate about the *x*-axis. We place a qubit into an equal superposition of logical states by rotating the moment through 90° with a $\pi/2$ pulse, in which B_1 is turned on for a time $t = \pi/(2\omega_1)$, where $\omega_1 = \gamma B_1$.

We can also make moments rotate continuously about the x-axis. In a process known as cyclic adiabatic inversion, we sweep ω_{osc} through a range that includes the Larmor frequency. When we start the sweep, $\omega_{\text{osc}} << \omega_{\text{L}}$. According to Equations (5), there is little cancellation of the static field $\mathbf{B}_0$, and $\mathbf{B}_{\text{eff}}$ will lie substantially along the z-axis. As the frequency approaches ω_{L}, there is more cancellation, and $\mathbf{B}_{\text{eff}}$ begins to rotate toward the x-axis. When $\omega_{\text{osc}} = \omega_{\text{L}}$, $\mathbf{B}_{\text{eff}}$ points along the x-axis. Continuing to sweep the frequency to $\omega_{\text{osc}} >> \omega_{\text{L}}$ will eventually cause $\mathbf{B}_{\text{eff}}$ to point along the $(-z)$-axis. If ω_{osc} is swept slowly enough (the adiabatic condition), the moments will continue to precess around $\mathbf{B}_{\text{eff}}$ and will follow its rotation in the x-z plane from $+z$ to $-z$ (See Figure D). Reversing the sweep will cause $\mathbf{B}_{\text{eff}}$ to rotate backwards. By continuously sweeping ω_{osc} back and forth through the resonance frequency, we effectively make the spins rotate continuously around the y-axis.

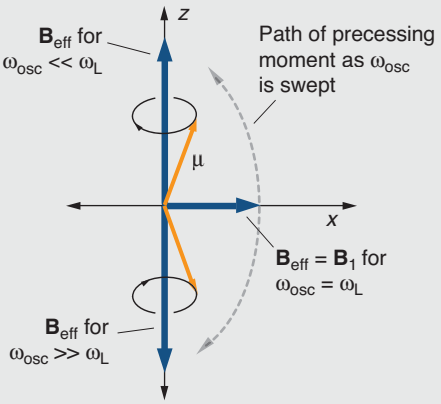

Figure D. Cyclic Adiabatic Inversion

$\mathbf{B}_{\text{eff}}$ rotates about the *y*-axis when ω_{osc} is swept through the resonance frequency ω_{L}. If ω_{osc} changes slowly, the moment continues to precess about $\mathbf{B}_{\text{eff}}$ and we can rotate the moment about the *y*-axis.

Cyclic adiabatic inversion provides one of the mechanisms by which we detect electron moments with a magnetic resonance force microscope (MRFM). A small number of moments are in resonance with $\mathbf{B}_0$, $\mathbf{B}_1$, and the gradient field produced by the magnetic tip at the end of the MRFM cantilever. We use cyclic adiabatic inversion to selectively rotate those moments, thus producing a tiny oscillating magnetization within the sample that in turn produces an oscillating force on the MRFM cantilever. By adjusting the rate at which we sweep ω_{osc}, we can match the forcing frequency to the cantilever's resonant frequency, and even a small number of moments can drive the cantilever into a detectable oscillation.

Continued from page 287

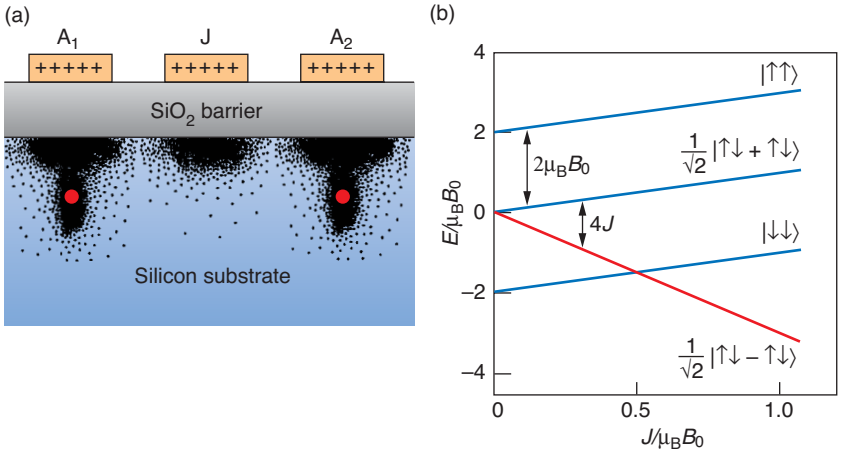

Figure 3. Coupling Two Qubits with a J-Gate
(a) When the gates A_1 and A_2 are appropriately biased, application of a small positive voltage to the J-gate lowers the potential barrier between adjacent donor sites and turns on an electron-spin exchange interaction between qubits, as seen in this cartoon. The electrons interact with the nuclei through the hyperfine interaction; hence, the two nuclear spins become coupled to each other. (b) The graph shows energy levels of the coupled electron-spin system as a function of the electron-spin exchange coefficient *J*, which can be modified by voltage applied to the J-gate. The electrons couple their spins to form three states with $S = 1$ (shown in blue) and one with $S = 0$ (shown in red). For $J/\mu_B B_0 < 0.5$, the electrons occupy the lowest energy $S = 1$ state $|\downarrow\downarrow\rangle$. Two-qubit operations are performed in this low *J* regime.

interaction. The exchange coupling coefficient J is proportional to the overlap between the wave functions of the two donor electrons, so its strength is a function of the J-gate voltage.

We first examine the coupled electron-spin states by ignoring (for a moment) the contribution of the nuclear spins to H^1 and H^2 in Equation (4). The effect of the spin exchange interaction is to create coupled electron-spin states, three with total spin $S = 1$ and one with total spin $S = 0$. The respective wave functions are

$$S = 1 \quad |\uparrow\uparrow\rangle \ , \quad 1/\sqrt{2}\,|\uparrow\downarrow + \downarrow\uparrow\rangle \ , \text{ and } \quad |\downarrow\downarrow\rangle \ ,$$

$$S = 0 \quad |S\rangle = 1/\sqrt{2}\,|\uparrow\downarrow - \downarrow\uparrow\rangle \ .$$

In the absence of a magnetic field, the energy difference between the states with $S = 1$ and $S = 0$ is $4J$, an amount known as the exchange energy. In the presence of the magnetic field $\mathbf{B}_0$ that permeates the quantum computer, the $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ states are Zeeman-split around zero by an amount $\pm 2\mu_B B_0$, and the energies of the coupled electron-spin states vary as a function of J, as seen in Figure 3(b). Notice that the lowest-energy $S = 1$ state and the $S = 0$ state cross when the exchange energy becomes equal to the Zeeman splitting, that is, when $4J = 2\mu_B B_0$. We exploit that crossing in a qubit readout scheme discussed later.

We now consider the nuclear spin states. Conceptually, for every coupled electron-spin state, there are four possible orientations of the two nuclear spins, corresponding to the uncoupled ($J = 0$) nuclear states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. Thus, there are sixteen nuclear spin states in all. Formally, we must use Equation (4) to find the energies and eigenfunctions of all sixteen.[3] If we focus only on those states associated with the electron ground state $|\downarrow\downarrow\rangle$ and assume $4J < 2\mu_B B_0$, then in order of decreasing energy, the coupled nuclear-spin states are the following:

$$|11\rangle \ ,$$
$$|\Phi_+\rangle = 1/\sqrt{2}\,|10 + 01\rangle \ ,$$
$$|\Phi_-\rangle = 1/\sqrt{2}\,|10 - 01\rangle \ , \text{ and}$$
$$|00\rangle \ . \qquad (5)$$

The electron-spin exchange interaction shifts the energy of the $|\Phi_-\rangle$ state below that of $|\Phi_+\rangle$ by an amount

$$\Delta E_J = \hbar\omega_J = 2A^2\left(\frac{1}{\mu_B B_0 - 2J} - \frac{1}{\mu_B B_0}\right) \ , \qquad (6)$$

where ω_J is the nuclear exchange frequency. For $B_0 = 2T$ and $4J = 0.124$ milli-electron-volt (meV), ω_J has a value of about $(2\pi)75$ kilohertz, a frequency that approximates the rate at which binary operations can be performed on the computer.

The spin exchange interaction causes the wave functions of Equation (5) to evolve and rotate between spin states. One possible result is that the nuclear spins undergo a coordinated swapping of states: $|q_1 q_2\rangle \rightarrow |q_2 q_1\rangle$ (see the box "The Swap" on the facing page). Thus, the spin exchange interaction should automatically implement the logical two-qubit **swap** gate.

Of more interest is the **cnot** gate, which along with single-qubit operations, forms a universal set of gates from which any quantum algorithm can be executed. In the Kane system, the **cnot** corresponds to the conditional rotation of a target spin by 180°, provided the control spin is in the state $|1\rangle$. In principle, it can be realized by subjecting the wafer to a transverse

[3] The energy differences between the four nuclear states associated with each electron state are very small. A graph of the nuclear-electron energy levels would look identical to Figure 3(b), except that under high magnification one would see that each line consists of four closely spaced lines.

magnetic field $\mathbf{B}_1$ and applying voltages to the A- and J-gates (Goan and Milburn 2000).

Suppose that the two electron spins are initially uncoupled ($J = 0$) and that the hyperfine coupling constants A_1 and A_2 of the two donor atoms are equal ($A_1 = A_2$). In that case, biasing the A-gates such that $A_1 > A_2$ distinguishes the control qubit from the target. We then turn on the spin exchange interaction ($J > 0$) and slowly make A_1 equal to A_2. The result would be that the uncoupled qubit state $|10\rangle$ evolves adiabatically into the state $|\Phi_+\rangle = 1/\sqrt{2}\,|10 + 01\rangle$, and $|01\rangle$ evolves into $|\Phi_-\rangle = 1/\sqrt{2}\,|10 - 01\rangle$. When A_1 equals A_2, the energy splitting between the two states is given by Equation (6). The states $|11\rangle$ and $|00\rangle$ are unaffected by the procedure.

We next apply a linearly polarized oscillating field B_1 with frequency that is resonant with the $|11\rangle - |\Phi_+\rangle$ energy difference. The field is left on until the $|11\rangle$ state has rotated into the $|\Phi_+\rangle$ state and vice versa. By executing a reverse of the sequence of steps performed at the beginning of the operation, we adiabatically transform the $|\Phi_+\rangle$ and $|\Phi_-\rangle$ states back into $|10\rangle$ and $|01\rangle$, respectively. We effect the change

$$\begin{aligned} |00\rangle &\rightarrow |00\rangle\ , \\ |01\rangle &\rightarrow |01\rangle\ , \\ |10\rangle &\rightarrow |11\rangle\ , \text{ and} \\ |11\rangle &\rightarrow |10\rangle\ . \end{aligned}$$

That is, the only qubits that are flipped are those in which the control qubit is in the state $|1\rangle$. Thus, we expect to be able to perform the **cnot** operation.

Approaches to Readout

One can evaluate the result of a quantum computation only by reading the final state, $|0\rangle$ or $|1\rangle$, of a qubit. Likewise, the ability to determine the state of a given qubit is critical to initializing the quantum register. Ideally, we would read the qubit state directly by measuring the donor atom's nuclear-spin state. But direct detection of a single nuclear spin is currently impossible and may forever be out of our grasp. (The strength of the coupling between a magnetic field and the nuclear spin is set by the magnitude of the nuclear magneton μ_n, which is very small.) We are therefore forced to find some other means of reading out the qubit state.

The potential solution is to correlate the nuclear spin states of a target atom with the electron spin and to find some way of determining the electron spin state. We are currently pursuing two distinct detection schemes, one involving a single electron transistor (SET) and the other, a magnetic resonance force microscope (MRFM). Both approaches require that we push the respective technologies beyond the current state of the art.

The Swap

Before the J-gate is turned on, the two nuclear spins are uncoupled, and each is described by the following energy eigenstates: $|\Psi_1\rangle = |00\rangle$, $|\Psi_2\rangle = |01\rangle$, $|\Psi_3\rangle = |10\rangle$, and $|\Psi_4\rangle = |11\rangle$. Once the J-gate is turned on, the coupled eigenstates are $|00\rangle$, $|\Phi_-\rangle = 1/\sqrt{2}\,|10 - 01\rangle$, $|\Phi_+\rangle = 1/\sqrt{2}\,|10 + 01\rangle$, and $|11\rangle$.

Suppose the uncoupled nuclear spins were originally in the state $|\Psi_2\rangle = |01\rangle$, and then voltage was applied quickly to the J-gate. In terms of the eigenstates of the coupled system, the system finds itself in the state

$$|\Psi_2\rangle = 1/\sqrt{2}\,(|\Phi_+\rangle - |\Phi_-\rangle)\ . \qquad (1)$$

The time evolution of this wave function (up to an overall phase) is given by

$$|\Psi_2(t)\rangle = 1/\sqrt{2}\,(|\Phi_+\rangle - e^{-i\omega_J t}\,|\Phi_-\rangle)\ , \qquad (2)$$

where ω_J is the nuclear exchange frequency. After a time $t = \pi/\omega_J$, the wave function will evolve into

$$|\Psi_2(\pi/\omega_J)\rangle = 1/\sqrt{2}\,(|\Phi_+\rangle + |\Phi_-\rangle) = |\Psi_3\rangle\ . \qquad (3)$$

That is, the system will have evolved from the state $|01\rangle$ to the state $|10\rangle$. The spins will have swapped. If we quickly remove the voltage from the J-gate, the two-spin system will remain in the state $|10\rangle$.

Single-Charge Measurement Using SETs. The idea behind this technique, first described by Kane (1998), is to couple the target qubit q_t to a readout qubit q_r by a J-gate, and then infer the state of q_t by monitoring the donor electrons of the coupled system. If q_t is in the state $|0\rangle$, we can cause both electrons to become localized around the readout atom (they would occupy the so-called D^- state). If q_t is in the state $|1\rangle$, each donor electron would remain bound to its respective atom. An SET would be used as an ultrasensitive electrometer to determine whether one or two electrons were bound to the readout atom.

The procedure can be understood with reference to Figure 4(a), which shows the coupled nuclear-spin states in the vicinity of the electron spin crossing. As discussed in the previous section, for $J/\mu_B B_0 < 0.5$, the lowest-energy electron spin state is the S = 1 state $|T\rangle = |\downarrow\downarrow\rangle$, but for $J/\mu_B B_0 > 0.5$,

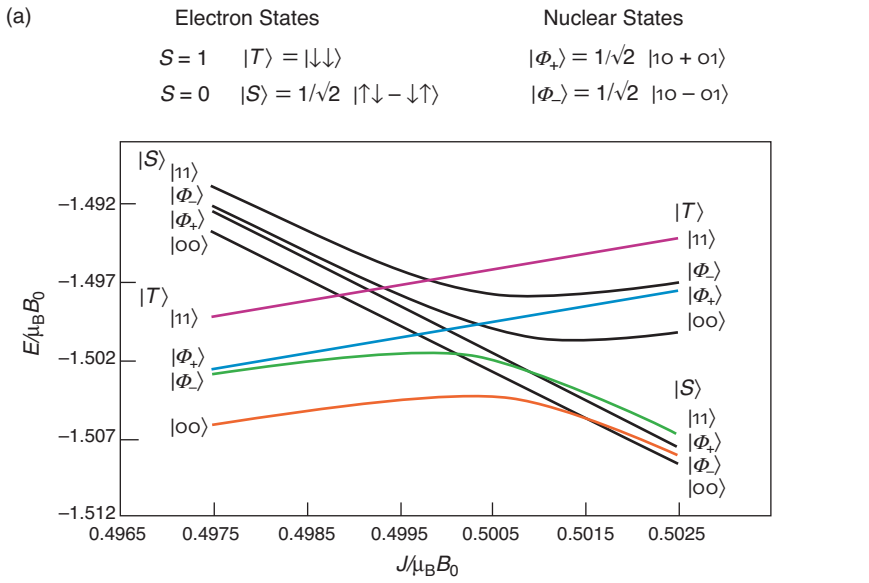

(b) States Adiabatically Evolved through the Crossing

	$J/\mu_B B_0 < 0.5$, $A_t - A_r \gg \hbar\omega_J$		$J/\mu_B B_0 < 0.5$, $A_t = A_r$		$J/\mu_B B_0 > 0.5$, $A_t = A_r$
q_t initially in state $\|1\rangle$	$\|T\rangle\|11\rangle$	→	$\|T\rangle\|11\rangle$	→	$\|T\rangle\|11\rangle$
	$\|T\rangle\|10\rangle$	→	$\|T\rangle\|\Phi_+\rangle$	→	$\|T\rangle\|\Phi_+\rangle$
q_t initially in state $\|0\rangle$	$\|T\rangle\|01\rangle$	→	$\|T\rangle\|\Phi_-\rangle$	→	$\|S\rangle\|11\rangle$
	$\|T\rangle\|00\rangle$	→	$\|T\rangle\|00\rangle$	→	$\|S\rangle\|\Phi_-\rangle$

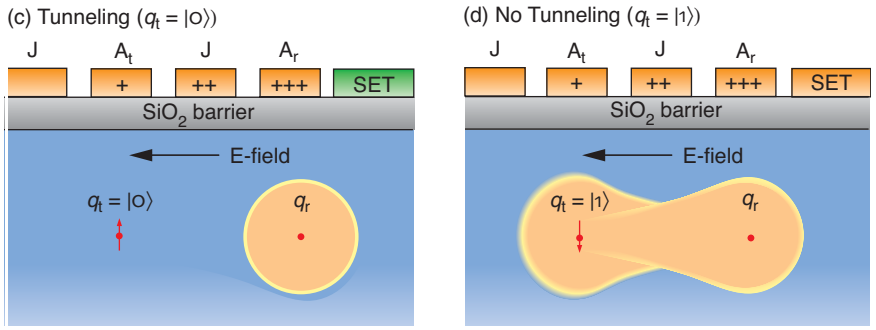

Figure 4. Single-Electron Transistor (SET) Readout Scheme

(a) The graph shows the eight lowest-energy nuclear-spin states for the coupled target and readout qubits $|q_t, q_r\rangle$ in the region where the $S = 0$ and the lowest energy $S = 1$ electron-spin states cross. (b) We can adiabatically evolve the nuclear-electron states by biasing the J- and A-gates, as seen in this (partial) sequence of steps. The electrons are initially in the $S = 1$ state $|T\rangle$. If q_t was initially in the $|1\rangle$ state, then the electrons will remain in $|T\rangle$ regardless of the state of q_r. If initially $q_t = |0\rangle$, then at the end of the sequence, the electrons will be in the $S = 0$ state $|S\rangle$. (c) Only the two electrons in the $|S\rangle$ state can bind to a single phosphorous atom in silicon. Given a suitable biasing of the gate electrodes, we can try to induce an electron to tunnel to a readout qubit q_r. If the tunneling is successful, the electrons were in the $|S\rangle$ state, and $q_t = |0\rangle$. The tunneling current would be detected by an SET located near q_r. (d) If no tunneling occurs, the two electrons were in the $|T\rangle$ state, and hence $q_t = |1\rangle$.

the $S = 0$ state $|S\rangle = 1/\sqrt{2}\ |\uparrow\downarrow - \downarrow\uparrow\rangle$ assumes the lower energy.

Figure 4(a) shows what happens to the eight lowest-energy nuclear-spin states as the electron-spin states cross. Focusing on the four states initially associated with $|T\rangle$, we see that after the crossing, the two higher-energy nuclear states $|11\rangle$ and $|\Phi_+\rangle$ remain coupled to $|T\rangle$, while the two lower-energy states $|00\rangle$ and $|\Phi_-\rangle$ get coupled to $|S\rangle$. In other words, as we increase J, we can adiabatically evolve both the nuclear- and electron-spin systems. If the target qubit was originally in the state $|0\rangle$, then regardless of the state of the readout qubit, the electrons will evolve into the $S = 0$ spin state. If q_t is originally in the state $|1\rangle$, the electrons will remain in the lowest energy $S = 1$ spin state. The sequence of steps, similar to those used to implement the **cnot** gate, is outlined in Figure 4(b).

We next use the fact that the only two-electron bound state of a phosphorous atom in silicon is the D^- state with total spin $S = 0$. As seen in Figure 4(c), we would bias the A- and J-gates to create an electric field between the two donor atoms. If the electrons are in the $S = 0$ state, the target electron can transfer to the readout atom, and we would know that the target atom was initially in the state $|0\rangle$.

An SET would be used to detect the presence of the second donor electron about the readout atom. In many ways, an SET is like an ordinary transistor, in that a gate electrode moderates the current flowing between a source and drain electrodes. The difference is that between the SET's source and drain lies an extremely small metallic island, which is isolated from each electrode by small patches of insulating material. The insulator acts as a tunnel junction. For current to flow, electrons must tunnel from the source to the island and then from the island to the drain. The tunneling current is greatly affected by the

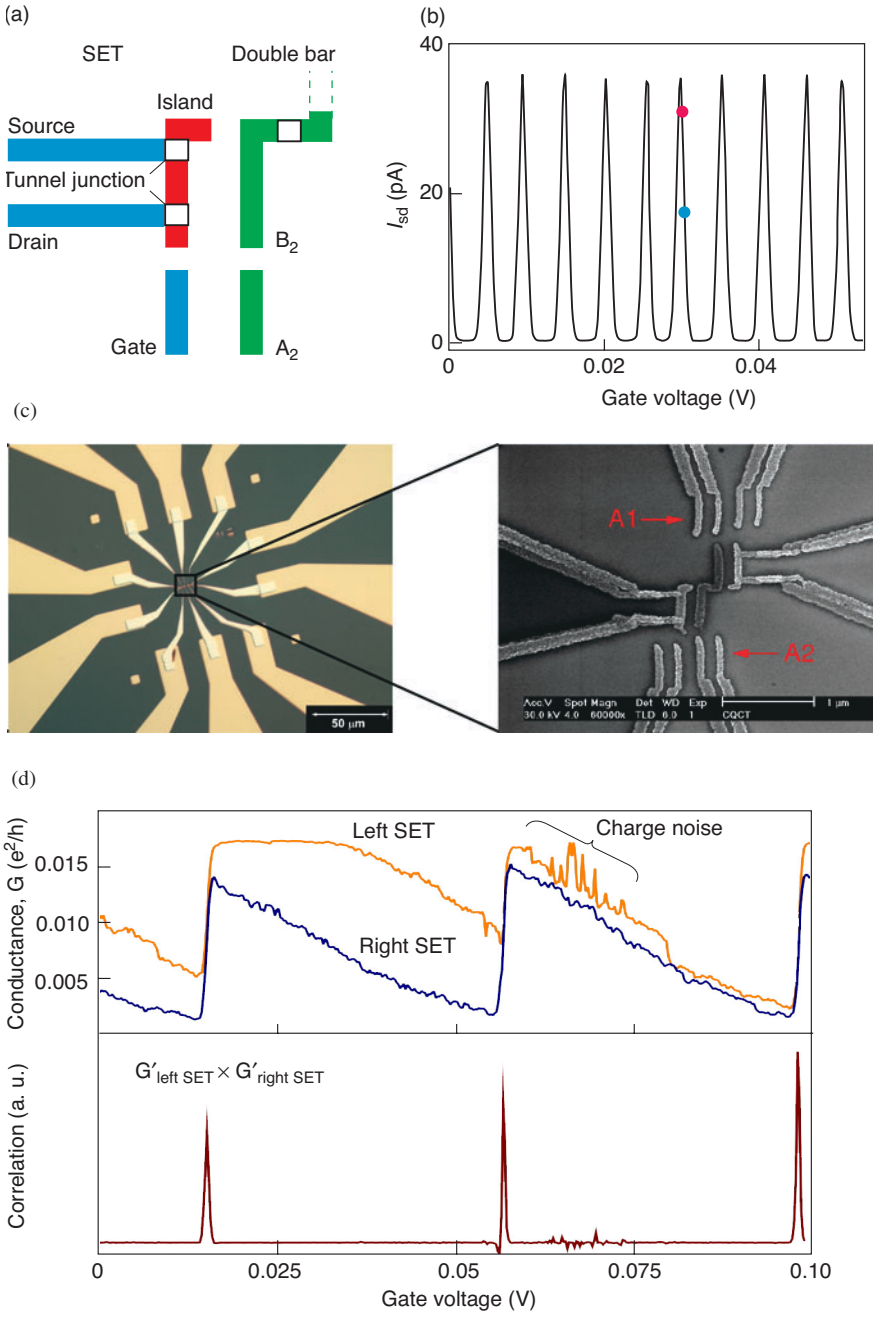

Figure 5. Twin SET Device for Readout Simulation

(a) The figure shows one half of a twin SET/ double-bar test device. The SET consists of a small metal island connected to source and drain leads by tunnel junctions and a gate electrode that is capacitively coupled to the island. Electrons can pass from source to drain only by tunneling through both junctions. The SET is located next to a metal bar B_2, which is isolated from the other bar B_1 by a tunnel junction. (b) The tunneling current I_{sd} in the SET is strongly influenced by a change in the local charge distribution. If the gate voltage is originally biased at V_0, (blue dot), then a change in the local charge distribution effectively modifies it to $V_0 - \delta$, and the source-drain current will change dramatically (red dot). (c) This is an image of the twin-SET test device obtained with a scanning electron microscope. The image to the right is a magnified version of the central region. The twin-SET device is fabricated by a double-angle evaporation process, which replicates each of the features. Unequal voltage on A_1 and A_2 causes an electron to tunnel from one bar to the next. (d) The movement of charge is detected as a change in the source/drain conductance in both SETs simultaneously. The two signals can be correlated to discriminate the charge transfer signal from reproducible charge noise or from random noise events.

capacitive coupling between the gate and the island. This means that for particular voltage biases on the gate, source, and drain, current flow through the SET becomes exquisitely sensitive to minute changes in the charge distribution of the local environment. The presence of a single additional electron is readily detectable as a change in the SET's source/drain conductance.

We have developed several readout simulation devices to test the properties of our SETs built in house. In the device seen in Figure 5, two thin metal bars, isolated from each other by a tunnel junction, substitute for the phosphorous atoms. Control gates are used to electrostatically "push" single electrons from one bar to the next. The two SETs are then used to detect the change in the charge distribution due to the discrete, single-electron tunneling events. Those events cause the output of both SETs to change abruptly. In contrast, signals due to unwanted charge noise (reproducible fluctuations in the conductance versus voltage curve) tend not to affect both SETs simultaneously. By correlating the outputs of the two SETs, we are able to clearly identify the single-charge transfer events and reject

spurious signals that would interfere with the readout.

Other factors, however, also need to be considered before we use an SET in a qubit readout scheme. Suppose the target qubit q_t is initially in the $|1\rangle$ state. Then, for high values of J, the coupled electrons will remain in the higher-energy state $|T\rangle$ (refer to Figure 4). This means that the coupled atomic system could lower its energy if one of the polarized electrons "relaxed" and flipped to form the state $|S\rangle$. The electron would then transfer to the readout qubit, and we would erroneously deduce that q_t was initially in the $|0\rangle$ state! Recent results suggest that the spin relaxation time is of the order of milliseconds. We must therefore pull information out of the SET on an even shorter time scale. We must be able to determine that a change occurred in the SET conductance at a time t_0, rather than a few milliseconds after t_0.

Unfortunately, that is difficult to do with an ordinary SET. The measurements are made at liquid helium temperatures, and the SET, sitting in a cryostat, must somehow be connected to the outside world. The capacitance of the connecting cables is fairly large, and when combined with the intrinsic resistance of the SET, produces a resistance-capacitance (or R-C) time constant for the device that is longer than the spin relaxation time. Information about the SET conductance takes too long to propagate to the outside world.

The solution to this problem is to develop a fast readout SET (Schoelkopf et al. 1998). Known as an rf SET, it is basically an ordinary SET coupled to an impedance-matching circuit that negates the effects of the external capacitance. We have recently developed a very sensitive reflection-mode rf SET that operates at 430 megahertz. It can detect the movement of a single electron in the device shown in Figure 5 in about 1 microsecond. For a system containing discrete phosphorous atoms, the readout time would likely increase to about 100 microseconds, but that is still sufficient for the readout approach discussed in this section.

The MRFM. The second approach to readout is to perform a direct measurement of the spin state of the electron surrounding the qubit and thereby infer the qubit state. To do so, we are developing an MRFM, which combines the versatility and chemical specificity of magnetic resonance with the high-resolution and ultrahigh sensitivity of an atomic force microscope (AFM). The key feature of the MRFM is that only spins contained within a defined area in the sample—the so-called sensitive slice—contribute to the detected signal. Because the location and size of that slice can be controlled, there is selective sensitivity to deeply buried structures.

Our MRFM, developed at Los Alamos in collaboration with Michael Roukes of Caltech, is illustrated in Figure 6. The microscopic, sharp-pointed magnetic tip is bonded to the end of a tiny cantilever. As in an ordinary AFM, the tip is scanned over a sample, and signals are recorded at every point. In our instrument, however, the magnetic field from the tip $\mathbf{B(r)}$ interacts with all the electron spins in the substrate through the magnetic gradient force, $\mathbf{F(r)} = (\mathbf{m}\cdot\nabla)\mathbf{B(r)}$, where $\mathbf{m}$ is the net magnetization of the spins. Depending on the spin orientation, the force on the tip is either repulsive or attractive. The net orientation of the electron spins in the sample, therefore, causes a tiny deflection of the cantilever.

We interact with only a subset of the spins through magnetic resonance. The sample is immersed in a static magnetic field $\mathbf{B}_0 = B_0\hat{z}$, so the precession frequency of the spins around the magnetic-field lines is proportional to $\mathbf{B}_0$ plus the z-component of $\mathbf{B(r)}$, that is, the total magnetic field in the z-direction**.** Magnetic resonance comes into play when we subject the spins to an oscillating magnetic field $\mathbf{B}_1$ that is aligned in the x-direction. Because the magnitude of $\mathbf{B(r)}$ decreases rapidly with distance, only those spins that are at the correct distance from the tip are in resonance with the $\mathbf{B}_1$ field. Spins that are too close to the tip will have a higher resonant frequency; those that are farther away, a lower frequency. Thus, for a given field gradient and a fixed cantilever position, a resonance frequency becomes correlated with positions inside the sample. With reference to Figure 6, all spins that lie within a small, hemispherical shell beneath the tip (the sensitive slice) have the same resonance frequency.

To detect those spins, we use the technique of cyclic adiabatic inversion, discussed in the box "Spin Manipulation with Magnetic Resonance" on page 288. In essence, we continuously rotate the selected spins at the resonant frequency of the cantilever. The continuous up and down reorientation of the spins creates an oscillating force on the tip that amplifies the cantilever's natural up-down motion. The situation is analogous to pushing a child's swing at its natural frequency of oscillation: with each push, the amplitude of the motion becomes larger. After thousands of spin rotations, the amplitude of the cantilever's up-down motion has increased by about an angstrom, which is large enough to be detected with a fiberoptic interferometer. The fiber sits slightly above the back of the cantilever, and laser light sent down the fiber interferes with itself as it reflects from both the cantilever and the fiber's end. By monitoring changes in the interference pattern, we can detect the oscillations.

The orientation of the nuclear spins can be inferred from the frequency at which the electron spin resonance

occurs. Because of the hyperfine interaction, the resonance frequency of an electron spin flip depends on the nuclear-spin state. Considering the hyperfine states of a single qubit, the $|0\downarrow\rangle$ to $|0\uparrow\rangle$ transition has a different energy than the $|1\downarrow\rangle$ to $|1\uparrow\rangle$ transition, and thus there are two resonance frequencies for an electron spin transition. Measurement of, say, the higher resonance frequency would correspond to the nuclei in the sample being aligned with the $\mathbf{B}_0$ field.

The discussion so far has centered on detecting many nuclear spins, but to read out the result of a quantum computation, we need to measure a single nuclear spin. That such measurement is at all possible is due to the exceedingly high spatial resolution of the MRFM, which is determined by the thickness Δz of the hemispherical shell. The thickness is inversely proportional to the magnitude of the field gradient:

$$\Delta z \cong \frac{\Delta\omega_r}{\gamma\left|\nabla_z \mathbf{B}(\mathbf{r})\right|} \quad , \qquad (7)$$

where γ is the gyromagnetic ratio and $\Delta\omega_r$ is the linewidth of the resonant electron-spin transition that is being driven by the MRFM. For phosphorus atoms in silicon, $\Delta\omega_r/\gamma$ is on the order of 1 milligauss. A field gradient of about 10^5 tesla per meter (T/m) will then produce a thickness that is much less than 1 angstrom, even when the hemispherical shell extends several hundred angstroms beneath the substrate surface. In that case, the sensitive slice would be so thin that only a single donor electron would be in resonance with the MRFM probe.

We have conducted numerous experiments to measure the field gradient of our specialized magnetic tips. With the tip about 2 micrometers from the surface, we have measured a field gradient approaching 10^4 T/m (see Figure 7). From this value, we estimated Δz and the volume of our hemispherical shell. Then, knowing the spin density of the sample, we estimated the number of spins that contribute to the signal. For the data shown in Figure 7, the number is between one thousand and ten thousand electron spins.

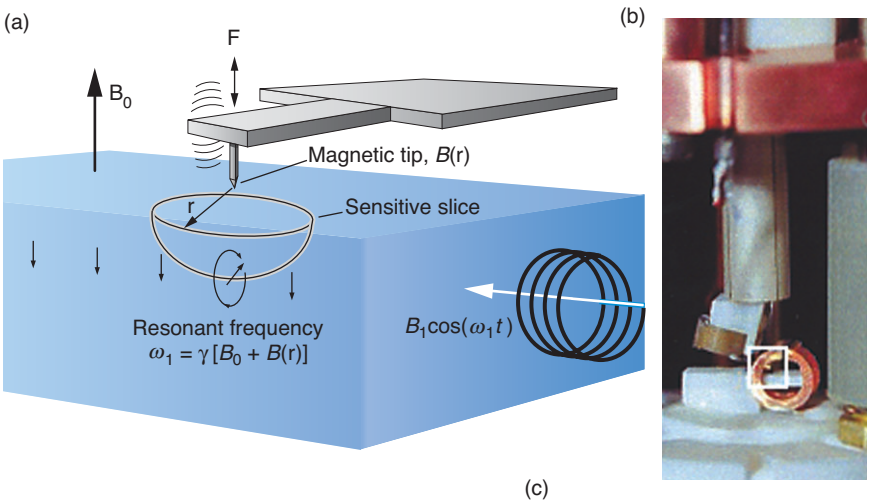

Figure 6. The Magnetic Resonance Force Microscope (MRFM)
(a) The MRFM can, in principle, deduce the state (up or down) of a single nuclear spin. The entire device sits at approximately liquid helium temperatures in a static magnetic field B_0. The tip at the end of the cantilever is coated with a magnetic material that generates a magnetic field B(*r*) that changes rapidly with the distance *r*. The interaction between the electron spins in the sample and the magnetic field gradient due to B(*r*) produces a force that deflects the cantilever. We interact with only a small subset of spins, located within a hemispherical shell of radius r_1, by subjecting the sample to an oscillating magnetic field $B_1\cos(\omega_1 t)$, where $\omega_1 = \gamma[B_0+B(r_1)]$. By using the technique of cyclic adiabatic inversion, we can cause the spins to oscillate between the up and down states at the cantilever resonance frequency, thus driving the cantilever into measurable oscillation. We detect the oscillation with an optical device. The electron-resonance frequency can then be correlated with a nuclear spin orientation. (b) The MRFM tip assembly and sample mount are shown in this photo. The vertical tube is a piezo scanning tube, which moves the tip over the sample, while the circular feature is the induction coil that produces B_1. The white box highlights the magnetically coated tip, shown under high magnification in (c).

Because the field gradient increases nearer to the tip, sensitivity should be greater if the tip is closer to the surface. But mechanical and thermal noise also deflect the tip and cantilever. As we begin to interact with fewer spins, the "signal" force due to spins eventually becomes less than the "noise" force due to unwanted sources. By equating expressions for the signal force to the noise force, we can derive an expression for the minimum detectable magnetic moment, $m_{\min}$, needed to give a signal to noise of 1:

$$m_{\min} = \frac{1}{\left|\nabla_z \mathbf{B}(\mathbf{r})\right|}\sqrt{\frac{2kk_B T\Delta\nu}{\pi Q f_c}} \quad . \qquad (8)$$

In Equation (8), k_B is the Boltzmann constant; T, the temperature; and $\Delta\nu$, the detection bandwidth. The other parameters describe the cantilever: its force constant k, resonant frequency f, and quality factor Q. The three key

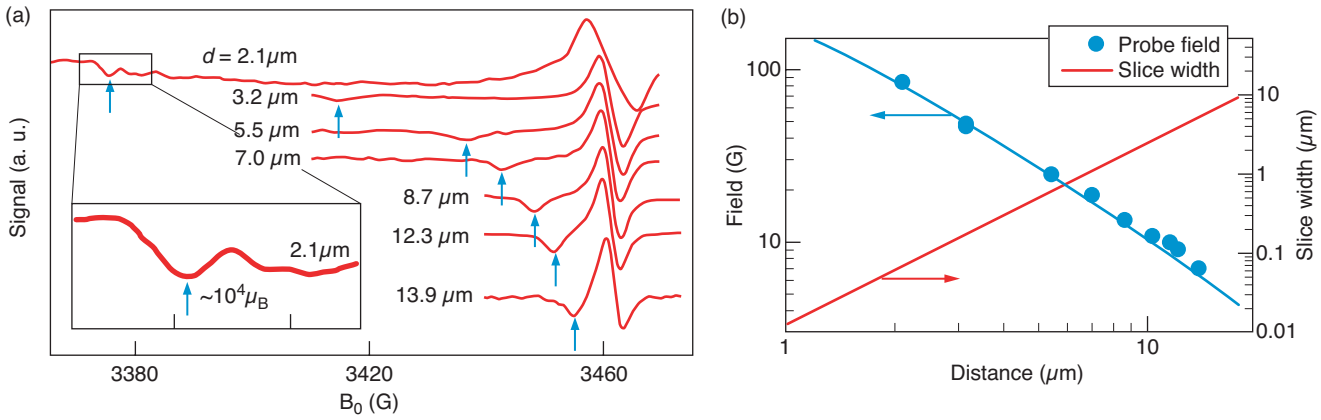

Figure 7. Sensitivity of the MRFM

(a) If the magnetic tip is kept at a fixed distance *d* from the sample, then lowering the value of the magnetic field corresponds to sweeping the sensitive slice upwards, toward the surface. At some point, the slice leaves the sample, and the resonance condition changes dramatically. That change is seen in the derivative of the MRFM signal as a dip, indicated by the arrows. (b) By following the dip as a function of tip height, we can measure the tip's magnetic field. From the field gradient, we then calculate the width of the sensitive slice using Equation (7) in the text. Knowing the spin density within the sample, we use the slice width to deduce how many spins produced the signal and thereby infer the MRFM sensitivity.

parameters that we can optimize are the field gradient, temperature, and quality factor.

We believe that the sensitivity of the MRFM is currently limited by surface contamination on the sample. As the tip approaches the surface, the contamination acts like a viscous force that damps the oscillatory motion—that is, it lowers the Q in Equation (8). To solve this problem, we are upgrading the equipment so that the sample be transferred from a surface preparation chamber into the microscope without leaving the ultrahigh vacuum environment. The system will also be cooled to temperatures between 250 and 300 millikelvins in a helium-3 dilution refrigerator, a technique that is compatible with maintaining the sample under ultrahigh vacuum.

Detection of a single electron moment requires that

$$m_{\text{min}} = 1\ \mu_B \cong 10^{-23}\ \text{joule/tesla} \quad . \quad (9)$$

Given a field gradient of 10^5 T/m, the signal force on the cantilever is approximately 10^{-18} newton (the weight of approximately two million phosphorus atoms). We believe that an upgraded, low-temperature microscope will allow us to observe the magnetic resonance signal of a single electron spin.

SSQC Fabrication Progress

Implementing our quantum-computing scheme requires that we produce a very regular array of phosphorus atoms in pure silicon, in which each donor is located precisely beneath a metal A-gate on the surface. The spacing between adjacent phosphorus donors is chosen to ensure that the electron-spin exchange interaction is minimal when there is no voltage on the J-gate lying between the donors. We want the two electron wave functions to overlap, but only slightly. Calculations (Goan and Milburn 2000) indicate that a separation of 10 to 20 nanometers between donors is required.

A nominal donor spacing of 20 nanometers translates into gate structures that are less than 10 nanometers in width. Fabricating a highly regular metal array on that scale, even with state-of-the-art techniques, is at the limits of the electron-beam techniques used in making conventional electronics. That problem, however, pales when compared with the difficulties we face in making a precisely aligned array of phosphorous donors that is buried under layers of silicon. The difficulties have led us to pursue two different fabrication strategies, known as the top-down and bottom-up approaches.

In the top-down approach, phosphorous atoms will be implanted by ion bombardment into specific sites on the silicon wafer. Because the ion scatters as it slows down in the silicon, we will not know the exact location of the donors, only that they will lie within close range of the defined implantation area. The top-down approach provides a rapid means to demonstrate proof of principle and allows us to fabricate a two-donor device that can be used to test readout

strategies and, possibly, quantum operations. Scaling this approach to large numbers of qubits will be challenging, because of the irregular spacing of the donor array.

In contrast, the bottom-up approach will use a scanning tunneling microscope (STM) with which to place phosphorous atoms on a clean silicon surface in a precisely arranged array. The array will then be overgrown with silicon, and the gate structures will be laid down by electron beam lithography (EBL). This approach, although more difficult to implement, could in principle allow us to build a Kane-type computer with the required precision. The bottom-up approach is not discussed here but is described in detail in the article "Fabricating a Qubit Array with a Scanning Tunneling Microscope" on page 302.

Top-Down Approach for Creating a Two-Donor Device. A host of issues surrounds the operation and readout of the nuclear-spin quantum computer. A key concern involves the transfer of the electron from the target to the readout atom during readout. The two electrons in the D^- state are not bound very strongly to the phosphorous atom, and the electron may be lost during the transfer. The initial phosphorus-phosphorus (P-P) system would transform into a $P\text{-}P^+$ system. If that is the case, we may need to use extra electrodes in order to create a deeper potential that will confine the electron or to employ a different readout atom (such as tellurium) that has a more strongly bound two-electron state.

Our current goal for the top-down approach is to produce a device that can be used to study the controlled electron transfer between two donors. We intend to ionize one of the two phosphorus atoms and then study the coherent transfer of the remaining electron between the two donor atoms in the $P\text{-}P^+$ system. For this purpose, we have relaxed the stringent constraints of the Kane computer architecture and designed the simple device shown in Figure 8. It consists of two A-gates separated by a single J-gate. Two phosphorous atoms will be implanted between the A- and J-gates, an arrangement that is sufficient for charge transfer experiments.

Device fabrication starts with a wafer that is already topped with a barrier layer of SiO_2. To deposit metal A- and J-gates on the surface

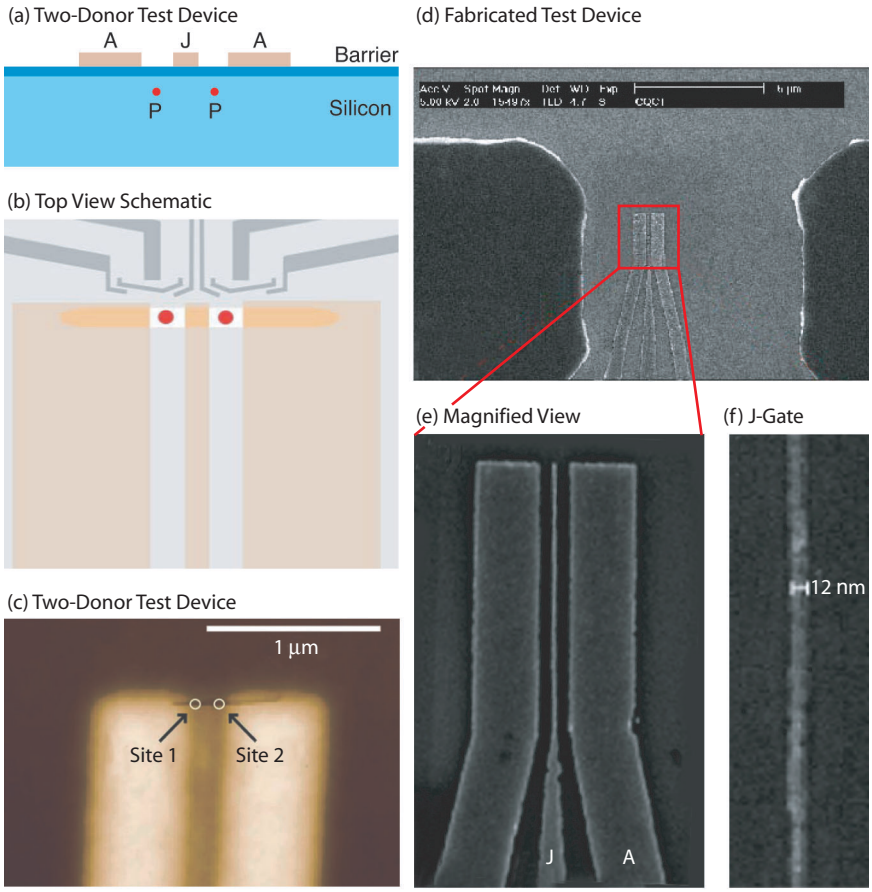

Figure 8. Fabrication of a Two-Donor Device

(a) For proof-of-principle experiments involving the transfer of an electron between two phosphorus donors, it is not necessary to configure A-gates above the donors. Instead, we are configuring the three-gate device shown in side view in this illustration. (b) The top view schematic shows that the donors will be implanted between the A- and J-gates (red circles). The single horizontal line at the top of the gates represents an opening in a polymer resist layer. Ions can enter the substrate only where the line crosses the gates. The schematic also shows representative SET devices located near the implantation sites. (c) AFM image of an actual device prior to implantation. The narrow horizontal line near the end of the gates is a 20-nm-wide opening in the resist. (d) An SEM image of a fabricated metal-gate array is shown. The large metal structures on either side of the gates are aluminum electrodes used to detect the impact of ions during implantation. (e) This magnified view shows the central A-, J-, and A-gates. We have fabricated gate arrays with J-gate widths of less than 15 nm and gate separations down to 30 nm. The image in (f) shows a J-gate made from a titanium/gold alloy that is only 12 nm wide.

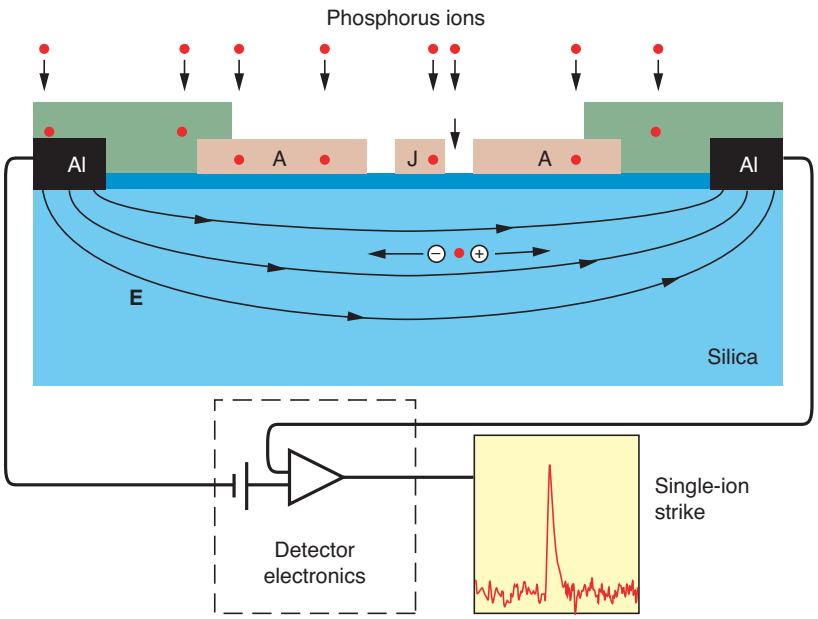

Figure 9. Detection of a Single-Ion Impact

We can detect in situ the impact of a single phosphorous ion in the silicon wafer. Aluminum electrodes are deposited on either side of the implantation site. An electric field applied between the two electrodes is used to separate electron-hole pairs that are created by ion impact. Each type of charge carrier migrates to its respective electrode and produces a current pulse that is detected by an external circuit. The implantation is halted after two such pulses have been detected.

before ion implantation, we use EBL techniques. A new layer of resist is created that covers the entire surface, including the gates. A second EBL exposure then patterns a thin line across the gates. This pattern is developed so that two tiny channels, each approximately 15 × 30 nanometers, are created on each side of the J-gate. The channels extend down to clean SiO_2 and define the implantation sites.

Next, we bombard the wafer with phosphorous ions. Although most of the ions are stopped in the mask, some go through the channels, strike the wafer, and get implanted about 10 nanometers below the Si/SiO_2 interface. After implantation, the device is heated to between 900°C and 950°C to anneal any damage to the silicon lattice. As a final step, we lay down the SETs. Creating the SETs after the anneal (instead of making them in the same step as the control gates) protects their fragile tunnel barriers, which would likely be degraded should they be submitted to temperatures above 900°C.

Because we want only one phosphorus atom per implantation site, the key to this entire process is the ability to detect a single ion after it has struck the silicon. And it is the properties of the silicon itself that help us fulfill this task. The energetic phosphorous ion produces a cascade of electron-hole pairs as it slows down and comes to a stop in the silicon matrix. Those charge species can be separated by an applied electric field, accelerated, and detected as a current pulse in an external circuit (see Figure 9). Voltage applied to surface electrodes straddling the implantation sites produces the field and transmits the pulses. The intrinsic silicon substrate makes this in situ particle-detection system highly efficient because the accelerating electric field extends fully between the two electrodes. We have demonstrated detection efficiencies of over 99 percent. Unfortunately, we cannot tell where the ion falls, and as there are two holes, there is only a 50 percent chance of creating a two-donor device with one donor in each hole. Although there are ways to improve those odds (by masking all but one hole with a specialized AFM cantilever, for example), in the short term, a 50 percent success rate is acceptable.

So far, no one has come close to seeing the transfer of a single electron between two precisely located, nanofabricated donor atoms. We hope to do so with a device similar to the one described above by September or October of 2002. We would then be in a position to study the coherent transfer of the electron between two donor atoms and possibly obtain information on decoherence mechanisms of relevance to spin readout.

Unlike the simple test device shown in Figure 8, an ideally configured device would have the A-gates directly above the phosphorous donors. We have designed a process to fabricate such a device (McKinnon et al. 2001). A multilayer electron-beam-sensitive resist is deposited on top of a SiO_2-coated silicon wafer, the resist is partially developed, and a linear array of ion-implantation channels is patterned in the resist with EBL techniques. The wafer is bombarded with phosphorous ions. The resist is then fully developed, and triple-angle metal evaporation is used to deposit the metal gate array and the SETs on the SiO_2 surface. Because only one mask is used to define the location of both the ion implantation sites and the gates, the A-gates are automatically registered over the qubits.

Figure 10 shows a six-donor test

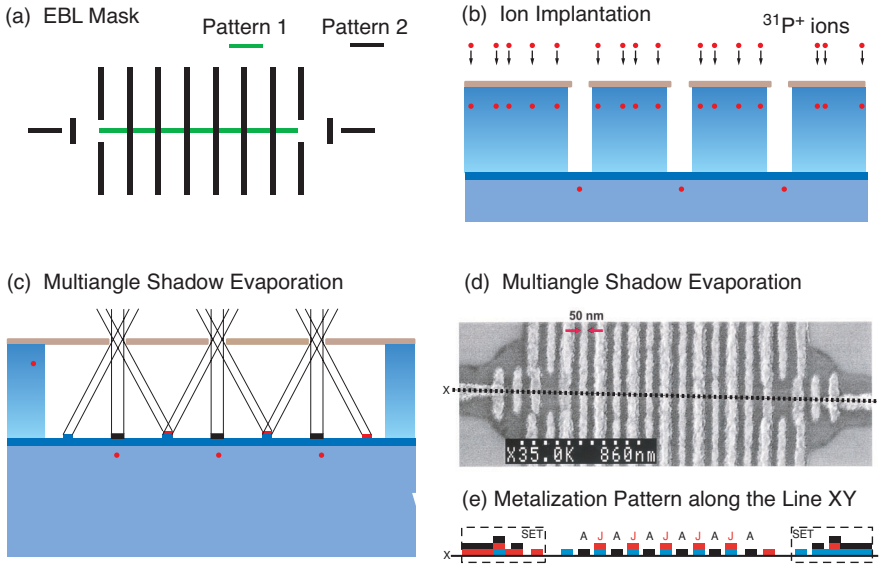

Figure 10. Creating a Multidonor Device

We have developed a process for creating gates directly over implanted ions. (a) In separate steps, patterns 1 and 2 of this EBL mask are partially developed in a trilayer resist that coats the substrate. The resist then sustains a series of lines with tiny channels that extend down to the substrate, where patterns 1 and 2 intersect. (b) A cross section of the resist after partial development shows the channels. The wafer is now bombarded with phosphorus ions. Some ions make it to the silicon surface and are implanted 5–10 nm below the Si/SiO_2 interface. (c) The resist is fully developed and material is removed, leaving behind a large cavity between the SiO_2 surface and the self-supporting top layer. Triple-angle shadow evaporation is then used to lay down an array of metal gates on the SiO_2 surface. (d) This photo shows a potential six-donor device with two readout SETs. No ions were implanted into this device. (e) A schematic side view of the device reveals the metalization pattern that results from the triple-angle shadow evaporation process.

device made by triple-angle evaporation. No ions were implanted into this device, and as can be seen in the figure, the triple-angle process does not yet result in gates that are sufficiently narrow to allow implementing the Kane scheme. We also have to address the problem of maintaining the integrity of the SET through the high-temperature anneal. However, the fundamental idea is robust.

The precision of the donor arrays produced by the top-down approach will be limited by straggling, which is inherent to ion implantation, and by the diffusion of dopants during the annealing step. Recent calculations indicate that small irregularities in the ion array could impair the operation of the quantum computer. That is why we are also pursuing the bottom-up fabrication approach, which might lead to a device with a very regular, well-characterized donor array.

Concluding Remarks

Phosphorous in silicon is a very clean, well-understood solid-state system. In its turn, NMR is a very well understood nuclear-spin manipulation technique. Performing NMR on a silicon chip implanted with phosphorus can therefore make for a very powerful quantum computer.

But the creation of a silicon-based solid-state computer presents such an enormous technical challenge that we must explore several strategies for building and implementing almost every aspect of the device. Hence, we investigate both SETs and magnetic resonance force microscopy as a means to read out the qubit state. Similarly, we have pursued two complementary fabrication strategies: the top-down process, which uses industrial production techniques, such as ion implantation and EBL, to produce a few-qubit device, and the bottom-up process, which involves advanced STM techniques and conventional molecular-beam epitaxy. Although the bottom-up approach is less suited to high-throughput production, it has the potential of leading to large, highly regular qubit arrays. We have made significant progress along all these parallel development paths.

Currently, scaling up a solid-state computer to over a million qubits is a goal that appears so distant as to be nearly out of sight. Yet less than fifty years ago, computer companies attempting to reduce the size of their machines were just becoming aware of a new strategy known as integrated circuits. Those early chips were crude and contained but a few transistors, but from them, evolved the modern, densely packed integrated circuits of today. Like those early chips, the quantum devices developed so far are rudimentary. No doubt, the challenges we face in building a real silicon-based quantum computer are significant, but our initial results offer hope that large-scale quantum computing may one day be realized. ■

Acknowledgments

Building a silicon-based solid-state quantum computer is a difficult and challenging project that cannot be achieved without the efforts of a large research team. At the University of New South Wales, Tilo Buehler,

Rolf Brenner, David Reilly, Bob Starrett, and Dave Barber have worked in the development of SET and rf SET readout schemes. Similarly, important contributions to the top-down fabrication approach have been made by Fay Stanley, Eric Gauja, Nancy Lumpkin, Rita McKinnon, Mladen Mitic, Linda Macks, and Victor Chan. At the University of Melbourne, Cameron Wellard has supported the device modeling; Changyi Yang, Paul Spizzirri, and Robert Short have worked on single-ion implantation technologies; and Jeff McCallum and Steven Prawer have contributed in a range of areas. The Magnetic Resonance Force Microscope research would not have been possible without the contributions of Denis Pelekhov, Michael Roukes, Andreas Suter, and Zhenyong Zhang.

Further Reading

Feher, G. 1959. Electron Spin Resonance Experiments on Donors in Silicon. *Phys. Rev.* **114** (5): 1219.

Goan, H., and G. Millburn. 2000 (February). Silicon-Based Electron-Mediated Nuclear Spin Quantum Computer. CQCT, University of Queensland internal report.

Gordon, J. P., and K. D., Bowers. 1958. Microwave Spin Echoes from Donor Electrons in Silicon. *Phys. Rev. Lett.* **1**: 368.

Honig, A., and E. Stupp. 1960. Electron Spin-Lattice Relaxation in Phosphorus-Doped Silicon. *Phys*. *Rev*. **117** (1): 69.

Kane, B. E. 1998. A Silicon-Based Nuclear Spin Quantum Computer. *Nature* **393**: 133.

McKinnon, R. P., F. E. Stanley, T. M. Buehler, E. Gauja, K. Peceros, L. D. Macks, M. Mitic et al. 2002. Nanofabrication Processes for Single-Ion Implantation of Silicon Quantum Computer Devices. (To be published in *Smart Mater. Structures*.)

Mozyrsky, D, Sh. Kogan, and G. P. Berman. 2001. Time Scales of Phonon Induced Decoherence of Semiconductor Spin Qubits. [Online]: http://eprints.lanl.gov (cond-mat/0112135).

Schoelkopf, R. J., P. Wahlgren, A. A. Kozhevnikov, P. Delsing, and D. E. Prober, 1998. The Radio-Frequency Single-Electron Transistor (RF-SET): A Fast and Ultra-sensitive Electrometer. *Science*, **280**: 1238.

Robert G. Clark received Bachelor of Science and Ph.D. degrees in physics from the University of New South Wales in Sydney, Australia, in 1973 and 1983, respectively. Between 1984 and 1990, Bob was a lecturer in physics at the University of Oxford, England, and research group head at Clarendon Laboratory at Oxford. During this period, he received a number of awards and distinctions, including the Wolfson Award for prestigious research (1988). In 1990, Bob became professor of experimental physics at the University of New South Wales. He is now heading the Center for Quantum Computer Technology in Sydney, Australia, and is leading the Australian effort for developing a silicon-based solid-state quantum computer.

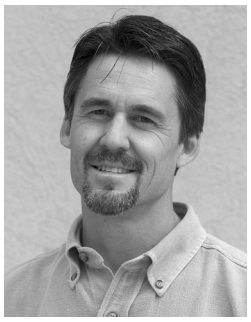

P. Chris Hammel received his B.A. from the University of California at San Diego and his Ph.D. from Cornell University (1984). Following a postdoctoral appointment at the Massachusetts Institute of Technology, Chris came to Los Alamos National Laboratory on a J. Robert Oppenheimer Fellowship. He became a staff member at Los Alamos in 1989. He is a Fellow of Los Alamos National Laboratory and was awarded the Los Alamos Fellows Prize in 1995. He is also a Fellow of the American Physical Society. In June, 2002, Chris joined the Physics Department at the Ohio State University as an Ohio Eminent Scholar. In addition to his current research, which is focused on the development of an ultrasensitive magnetic resonance force microscope, he is actively engaged in magnetic resonance studies of correlated electron systems including high-temperature cuprates and heavy fermions.

Andrew Dzurak is a program manager at the Centre for Quantum Computer Technology and Associate Professor within the School of Electrical Engineering at the University of New South Wales, Sydney. He has collaborated with researchers at Los Alamos since 1996. Andrew has 15 years of experience in the fabrication and cryogenic electrical measurement of quantum effects in semiconductor nanostructures, beginning with a Ph.D. at the University of Cambridge (the United Kingdom). Andrew moved to UNSW to take up an Australian Postdoctoral Fellowship and has worked closely with Robert Clark to establish the Semiconductor Nanofabrication Facility, of which he has been Deputy Director since 1995. Within the Centre, Andrew manages research in integrated devices for the control and readout of phosphorous qubits in silicon.

Alex Hamilton received his Ph.D. from the University of Cambridge, the United Kingdom, in 1993. He is a senior lecturer in physics in the School of Physics at the University of New South Wales, as well as manager for the Quantum Measurement Program at the Centre for Quantum Computer Technology. He is interested in the study of the fundamental properties of submicron electronic devices, with particular interest in the regime in which strong interactions between adjacent devices influence the behavior of the complete system.

Llloyd Hollenberg is a senior lecturer at the University of Melbourne, and since 2001 has been manager of the Device Modeling Program at the Centre for Quantum Computer Technology. He received his Ph.D. from the University of Melbourne in 1989. His interests include the general aspects of quantum many body systems, and in particular field theoretic and spin systems. At present, his work focuses on the physics of the Kane quantum computer and on quantum algorithms.

David N. Jamieson is an associate professor and reader at the University of Melbourne, Australia. He received his Ph.D. from the University of Melbourne, and has been Director of the Microanalytical Research Centre (MARC) of the School of Physics, University of Melbourne since 1996. In 2000, he joined the Centre for Quantum Computer Technology as the manager of the Ion Beam Program. He aims to adapt the technology of single ion detection and control to the construction of phosphorous arrays in silicon.

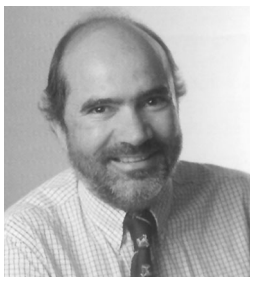

Christopher I. Pakes received his Ph.D. in physics in 1999 from the College of Physics, Birmingham, the United Kingdom. He has studied metrological applications of nanotechnology and spin-polarized properties of oxide functional materials, using ultrahigh vacuum, low temperature STM and dc SQUIDs. In 2000, Chris joined the Center for Quantum Computer Technology and was appointed manager of the Atomic Level Manipulation and Imaging Program in 2001.