Bulk Spin-Resonance Quantum Computation

Neil A. Gershenfeld and Isaac L. Chuang*

Quantum computation remains an enormously appealing but elusive goal. It is appealing because of its potential to perform superfast algorithms, such as finding prime factors in polynomial time, but also elusive because of the difficulty of simultaneously manipulating quantum degrees of freedom while preventing environmentally induced decoherence. A new approach to quantum computing is introduced based on the use of multiple-pulse resonance techniques to manipulate the small deviation from equilibrium of the density matrix of a macroscopic ensemble so that it appears to be the density matrix of a much lower dimensional pure state. A complete prescription for quantum computing is given for such a system.

The information-processing capability of quantum systems has long been of theoretical interest but has recently become of great practical importance because of a series of remarkable results leading up to Shor's algorithm for finding prime factors in polynomial time instead of exponential time (*1*–*4*). This development has led many groups to try to realize a quantum computer, using systems such as trapped ions (*5, 6*), quantum dots (*7*), and cavity quantum-electrodynamics (*8*–*10*). The experimental challenge is to find a system that has the nonlinear interactions that are required for computation and that simultaneously can be influenced externally in order to control it but that does not couple to the environment so strongly that the quantum coherence is rapidly lost (*11, 12*). The most successful realization to date is a two quantum bit ("qubit") computer that uses a cooled single ion of beryllium (*6*). Because of the enormous experimental effort required in such an experiment to isolate a small number of quantum degrees of freedom and cool them to their ground state, these approaches appear unlikely to provide a practically accessible means for computing with a larger number of qubits.

In the quest to build a quantum computer, one particularly attractive physical system has been the nuclear spin because of its extremely good isolation from electronic and vibrational mechanisms that can lead to decoherence. In fact, in Bloch's original 1946 paper on nuclear induction (*13*), he pointed out that relaxation times can be inconveniently long (for observation) and suggested that paramagnetic ions be introduced to shorten them. Of course, we are interested here in long relaxation times; today, in chemical applications of nuclear magnetic resonance (NMR), coherence times on the order of thousands of seconds can be observed (*14*), and protons in ordinary water can remain coherent for tens of seconds. Furthermore, the routine use of complex sequences of hundreds of radio frequency (rf) pulses for manipulating spins in modern NMR spectroscopy (*15*) points to the promise of nuclear spins for representing quantum information.

However, there is a problem: Quantum computation requires the preparation, manipulation, coherent evolution, and measurement of pure quantum states. In an ensemble of systems prepared in a distribution of states, such as is the case for nuclear spins in thermal equilibrium at room temperature, there is a statistical mixture of pure states. If computation were naïvely attempted with such a mixture, the resulting destructive interference between different states would eliminate the coherence needed for quantum algorithms.

These pitfalls and potentials for NMR quantum computation have been noted (*16, 17*). One suggestion was to create an apparatus to address single nuclear spins one at a time, in a spirit similar to manipulating single ions in a trap. This approach would avoid the thermal problem because by definition a single system is always in a pure quantum state. Unfortunately, the most promising experimental approach for this, that of using an atomic-resolution magnetic scanning force probe (*18*), is very difficult to realize because of the high sensitivity required to see the extremely small magnetic induction signal created by a single nucleus and the complications introduced by the presence of a scanning probe.

This article presents a new approach to quantum computing based on using bulk samples rather than isolated degrees of freedom. The problem, of course, is that such samples microscopically are in a thermal distribution of states, and it is impractical to hope to cool macroscopic materials to their ground state; furthermore, bulk samples are macroscopic ensembles whose members cannot be addressed individually for readout. We present here solutions to these problems. First, we report on a procedure to take advantage of the structure present in thermal equilibrium to introduce into the system's large density matrix a perturbation that acts exactly like a much smaller dimensional effective pure state. We then show how quantum computation can be performed using this ensemble system in such a way that the result is deterministic and can be read out efficiently. One great advantage of this approach is that, because of the massive redundancy provided by having a large ensemble of identical copies of the system, environmental interactions or intentional measurements only weakly perturb the computer's state. Thus, quantum computation becomes experimentally accessible in many naturally existing materials.

From the perspective of the NMR chemist, our scheme is unusual in a fundamental respect. In NMR spectroscopy, the primary purpose is to elucidate molecular structure and chemical dynamics, and great efforts are made to enhance the desired signal and render the detected spectra into a form that reflects properties of the system under study. Our purpose here is very different—we view each molecule as a single computer, whose state is determined by the orientations of its spins. Sequences of rf pulses, which manipulate spin orientations and couplings, constitute quantum logic gates and perform unitary transformations on the state. We begin with molecules with known configurations and spin couplings and deliberately suppress most of the NMR signal in order to produce an output that is independent of the molecular structure and is due only to the quantum computation that was performed.

A different perspective may be taken by the computer scientist. Our machine can be visualized as a collection of $O(10^{23})$ redundant N-bit computers. Because the N spins in each molecule may be in entangled quantum superposition states, our computer can be a quantum one. Therefore, although the clock cycle time of each computer may be slow [each rf rotation pulse can take O(100) microseconds], there can be a significant speed-up over a comparable classical computer because we may execute exponentially

N. A. Gershenfeld is at the Physics and Media Group, MIT Media Lab, Cambridge, MA 02139, USA. I. L. Chuang is with the Institute for Theoretical Physics, University of California Santa Barbara, Santa Barbara, CA 93106, USA. E-mail: ike@isl.stanford.edu

*Present address: Ginzton Laboratory, Stanford University, Stanford, CA 94305.

faster quantum algorithms. The challenge is to find a global way to use the ensemble of independent quantum computers without needing to address them individually.

Our solution operates by rearranging the states of *N*-spin molecules (which are initially in thermal equilibrium) such that a portion forms a uniform background that does not contribute to the measured signal, and the remainder forms a deviation that does. This deviation will behave as a pure state just as a vacancy in an electronic conduction band in a semiconductor behaves exactly like a quantum particle (a hole). The procedure utilizes traditional pulsed NMR techniques but requires novel excitation sequences. Once the system is prepared in this fashion, further pulses can effect the computation and cause the result to be read out. First, we will show how an effective pure state can be distilled from a thermal ensemble, and we will then turn to the more familiar question of how to apply arbitrary unitary transformations to such pure states and hence compute (*17*, *19*).

Pure states from a thermal ensemble. In a conventional pulsed NMR experiment, a transverse rf magnetic field is applied to a sample in a strong longitudinal magnetic field B_0. The strong field induces a Zeeman splitting of the levels of the spin system. In thermal equilibrium, for a two-level system (such as uncoupled protons) the ratio of the populations of the higher energy and lower energy levels differs from unity by $\sim 10^{-6}$. This small difference gives rise to a macroscopic magnetization, which can be detected by its precession in a pickup coil.

Consider, for example, a typical homonuclear two-spin molecular system such as (2,3)-dibromothiophene. In a 4.7-T field,

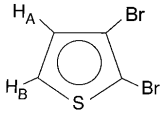

the precession frequencies of the two hydrogen nuclei in this molecule are both ~200 MHz, but are slightly different by their chemical shift that arises from their unique local chemical environments. Typical chemical shifts are a few parts per million, which translates to a few kilohertz in strong magnets. Single-spin rotations may be accomplished by applying semiselective transverse rf pulses, which are resonant at one of the proton frequencies, and cause the addressed nuclear spin to rotate around the transverse axis while still precessing around B_0. Spins on the same molecule may also interact with each other through dipolar- or electron-mediated interactions; these nonlinearities are used to accomplish logic operations, as will be discussed below. For simplicity of discussion, we shall consider only spin-1/2 systems and scalar-coupling Hamiltonians here.

At high temperature, the state of a system of *N*-spin molecules is well described by the density matrix (see Box 1 for a derivation):

$$\hat{\rho} = \frac{\hat{I}}{2^N} + \hat{\rho}_\Delta \tag{1}$$

The two terms describe an equilibrium part that is proportional to the identity $\hat{I}$, and a remaining traceless deviation part $\hat{\rho}_\Delta$, which we shall term the "deviation density matrix" (*20*). Under the action of a unitary transformation (for example, free evolution or imposed pulse sequences), the identity part of the density matrix will not change:

$$\hat{U}\hat{\rho}\hat{U}^\dagger = \hat{U}\left[\frac{\hat{I}}{2^N} + \hat{\rho}_\Delta\right]\hat{U}^\dagger = \frac{\hat{I}}{2^N} + \hat{U}\hat{\rho}_\Delta\hat{U}^\dagger \tag{2}$$

Therefore, the dynamics of the ensemble of noninteracting molecules can further be approximated by using just the density matrix

Box 1. The thermal equilibrium density matrix. Let us assume that the state of the entire system is well described as an ensemble of non-interacting molecules. For example, this is a good approximation in a liquid sample, where the rapid tumbling averages out intermolecular interactions (*15*). The ensemble averaged state is described by a density matrix, which is a tensor product of the density matrices for each molecule:

$$\hat{\rho} = \hat{\rho}_{\text{molecule 1}} \otimes \hat{\rho}_{\text{molecule 2}} \otimes \hat{\rho}_{\text{molecule 3}} \otimes \cdots \tag{14}$$

Because in thermal equilibrium these density matrices are all identical, and because during their further evolution they do not interact, it is sufficient to consider the evolution of a single molecular density matrix to represent the whole sample. The overall system of $O(10^{23})$ *N*-spin molecules has a huge number of degrees of freedom, but because it is not possible to address the molecules individually, the system acts as though it has only *N* degrees of freedom. This means that we obtain only a small number of useful quantum bits from the enormous number of underlying degrees of freedom in the sample.

In thermal equilibrium, the *N* spins of each molecule are arranged in some distribution of energy eigenstates (that is, aligned with or against B_0). These states are described by a $2^N \times 2^N$ diagonal density matrix whose elements give the average populations of the 2^N eigenstates. For example, the equilibrium density matrix for a single spin is:

$$\hat{\rho}_i^{\text{equilibrium}} = \begin{bmatrix} p\downarrow & 0 \\ 0 & p\uparrow \end{bmatrix} \tag{15}$$

where $p\downarrow$ and $p\uparrow$ denote the population probabilities for the two energy levels. In terms of the Boltzmann factor $\alpha_i = \hbar\omega_i/2kT$:

$$\hat{\rho}_i^{\text{equilibrium}} \approx \frac{1}{2}\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \frac{1}{2}\begin{bmatrix} \alpha_i & 0 \\ 0 & -\alpha_i \end{bmatrix} \tag{16}$$

where w_i is the resonant frequency of the i^{th} spin. For protons at room temperature, $\alpha \approx 4 \times 10^{-6}$ times B_0 in tesla. Although this is a very small number, this small deviation is what almost all NMR experiments measure.

For a molecule with *N* spins, in thermal equilibrium the density matrix is approximately given by the tensor product of the states of of the individual spins;

$$\hat{\rho} = \hat{\rho}_i \otimes \hat{\rho}_2 \otimes \cdots \otimes \hat{\rho}_N \tag{17}$$

For two spins, we find that:

$$\hat{\rho} = \frac{1}{4}\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} + \frac{1}{4}\begin{bmatrix} \alpha_1+\alpha_2 & 0 & 0 & 0 \\ 0 & \alpha_1-\alpha_2 & 0 & 0 \\ 0 & 0 & -\alpha_1+\alpha_2 & 0 \\ 0 & 0 & 0 & -\alpha_1-\alpha_2 \end{bmatrix} \tag{18}$$

written in the energy eigenstate basis $|\downarrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\uparrow\uparrow\rangle$.

For the $N = 4$ case, the diagonal elements are approximately:

$$\text{diag}(\hat{\rho}_\Delta) = \frac{\alpha}{16}\begin{bmatrix} \downarrow\downarrow\downarrow\downarrow & \downarrow\downarrow\downarrow\uparrow & \downarrow\downarrow\uparrow\downarrow & \downarrow\downarrow\uparrow\uparrow & \downarrow\uparrow\downarrow\downarrow & \downarrow\uparrow\downarrow\uparrow & \downarrow\uparrow\uparrow\downarrow & \downarrow\uparrow\uparrow\uparrow & \uparrow\downarrow\downarrow\downarrow & \uparrow\downarrow\downarrow\uparrow & \uparrow\downarrow\uparrow\downarrow & \uparrow\downarrow\uparrow\uparrow & \uparrow\uparrow\downarrow\downarrow & \uparrow\uparrow\downarrow\uparrow & \uparrow\uparrow\uparrow\downarrow & \uparrow\uparrow\uparrow\uparrow \\ 4 & 2 & 2 & 0 & 2 & 0 & 0 & -2 & 2 & 0 & 0 & -2 & 0 & -2 & -2 & -4 \end{bmatrix} \tag{19}$$

where the 16 spin states have been explicitly labeled. Now, a unitary transform on the spin states can be used to selectively exchange populations among the different energy levels; mathematically, this amounts to relabeling the states. Experimentally, this may be accomplished with logic gate pulse sequences (shown below) or by selective rf pulses. Either way, we can produce the final state:

$$\text{diag}(\hat{\rho}'_\Delta) = \frac{\alpha}{16}\begin{bmatrix} \downarrow\downarrow\downarrow\downarrow & \downarrow\downarrow\downarrow\uparrow & \downarrow\downarrow\uparrow\downarrow & \downarrow\downarrow\uparrow\uparrow & \downarrow\uparrow\downarrow\downarrow & \downarrow\uparrow\downarrow\uparrow & \downarrow\uparrow\uparrow\downarrow & \downarrow\uparrow\uparrow\uparrow & \uparrow\downarrow\downarrow\downarrow & \uparrow\downarrow\downarrow\uparrow & \uparrow\downarrow\uparrow\downarrow & \uparrow\downarrow\uparrow\uparrow & \uparrow\uparrow\downarrow\downarrow & \uparrow\uparrow\downarrow\uparrow & \uparrow\uparrow\uparrow\downarrow & \uparrow\uparrow\uparrow\uparrow \\ 4 & 0 & 0 & 0 & -2 & 2 & 2 & 2 & 2 & -2 & -2 & -2 & -4 & 0 & 0 & 0 \end{bmatrix} \tag{20}$$

for the macroscopic deviation from identity. Because this is governed by quantum evolution equations, the macroscopic signal reflects quantum dynamics (within the decoherence time scale, which can allow for thousands of rf pulses). This interesting observation is the usual starting point in NMR theory.

Now, let us consider the physical meaning of the terms in ρ_Δ at thermal equilibrium, where the diagonal terms give the state populations and the off-diagonal terms are zero. For example, suppose that we have an $N = 2$ spin system that has the deviation density matrix:

$$\hat{\rho}_\Delta = \alpha \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (3)$$

This representation means that of the 10^{23} molecules, $(\frac{1}{4} + 3\alpha) \times 10^{23}$ are in the $|\downarrow\downarrow\rangle$ state, while $(\frac{1}{4} - \alpha) \times 10^{23}$ are in each of the $|\downarrow\uparrow\rangle$, $|\uparrow\downarrow\rangle$, and $|\uparrow\uparrow\rangle$ states. The system is clearly a statistical mixture, meaning that it is an ensemble of molecules in a variety of the four possible states. Nevertheless, the system behaves like a pure state.

Why? Suppose that we tip one of the spins (in all of the molecules) into the transverse plane so that its precession generates a signal in our pickup coil. If the $|\downarrow\downarrow\rangle$ state had population $\frac{1}{4} - \alpha$ just like the other states, then because of the balanced distribution, for every spin pointed along $\hat{x}$, there would be another in the $-\hat{x}$ direction, and the net magnetization would be zero. However, the $|\downarrow\downarrow\rangle$ state has excess relative population, and thus the only net signal arises from $4\alpha \times 10^{23}$ $|\downarrow\downarrow\rangle$ spins, which are all in the same pure quantum state. The lesson is that an excess (or deficient) single state spin population among a uniformly populated background of levels behaves like a pure state.

We can use this result as follows. Let us make the reasonable assumption that the chemical shifts of the N spins are small compared to their average frequency ω, that is, that $|\omega_i - \omega_j| \ll \omega$. In this case, for $N = 2$ we have that:

$$\hat{\rho}_\Delta = \frac{2\alpha}{4} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (4)$$

Let us think of the second spin as a label; it tells us which of two pure states, $|\downarrow\rangle$ or $|\uparrow\rangle$, the first spin was prepared in. Mathematically, this is represented by the block-diagonal structure of $\hat{\rho}_\Delta$, which indicates that the state of the first spin is a pure state when conditioned on a measurement result of the second spin. Thus, if we arrange for unitary transforms to occur to the first spin alone (and for nothing to happen to the second spin), then the state shown in Eq. 4 becomes:

$$\left[\begin{array}{c|c} U & 0 \\ \hline 0 & U \end{array}\right] \left[\begin{array}{c|c} |\downarrow\rangle\langle\downarrow| & 0 \\ \hline 0 & -|\uparrow\rangle\langle\uparrow| \end{array}\right] \left[\begin{array}{c|c} U & 0 \\ \hline 0 & U \end{array}\right]^\dagger = \left[\begin{array}{c|c} U|\downarrow\rangle\langle\downarrow|U^\dagger & 0 \\ \hline 0 & -U|\uparrow\rangle\langle\uparrow|U^\dagger \end{array}\right] \quad (5)$$

Here, U represents the unitary transform effected by a series of pulse sequences applied to the spin system; it may be some one-bit operation carried out by a quantum computation. The meaning of this result is that the output is either $U|\downarrow\rangle\langle\downarrow|U^\dagger$ or $U|\uparrow\rangle\langle\uparrow|U^\dagger$, and moreover, these two are distinguished by the state of the second spin. Experimentally, this difference is determined in the final readout of the state. Thus, we may distill $k = 1$ qubits from $N = 2$ thermal spins.

This concept is easily generalized to $N > 2$ spins. The fundamental idea is to identify equally populated states that naturally exist in a thermal ensemble; unitary transforms are then used to group together these states to form a uniform background against which a differently populated single pure-state can express itself. We label this group using some other spins and in this manner produce a state in which a few of the spins are in pure states when conditioned on the state of the others.

As a final example, consider the $N = 4$ case (see Box 1). The final state can be represented by four independent pairs of pure state spins. That is, we may write:

$$\rho'_\Delta \doteq \frac{4\alpha}{16} \left[\begin{array}{c|c|c|c} |\downarrow\downarrow\rangle\langle\downarrow\downarrow| & 0 & 0 & 0 \\ \hline 0 & |\downarrow\downarrow\rangle\langle\downarrow\downarrow| & 0 & 0 \\ \hline 0 & 0 & |\downarrow\downarrow\rangle\langle\downarrow\downarrow| & 0 \\ \hline 0 & 0 & 0 & |\downarrow\downarrow\rangle\langle\downarrow\downarrow| \end{array}\right] \quad (6)$$

where "$\doteq$" indicates equality up to a sign and constant offset factor in each block. The four blocks are four separate two-spin pure states (let us call these our qubit spins) when conditioned on the remaining spins (the two "ancilla" spins). Note that the state of the qubits alone would not give a pure state; it is necessary to know the state of the ancilla spins as well. In this sense, the logical qubits we have are actually collective states of the four spin system. By freezing evolution of the ancilla spins (experimentally feasible with a technique known as decoupling), we prohibit the four block states from interacting with each other, and thus we need not concern ourselves with the dynamics of the ancilla.

In a similar manner, we may distill a three qubit pure state from six spins, a six qubit pure state from eight spins, and so forth. If the resulting block-diagonal density matrix is written as $\hat{\rho} = \hat{\rho}_1 \circ \hat{\rho}_2 \circ \cdots \circ \hat{\rho}_M$, then a unitary transformation that acts on just one set of spin degrees of freedom distributes over the blocks permitting an arbitrary transform to be applied to the qubit block:

$$\begin{aligned} &(\hat{U}_1 \otimes \hat{I}_2 \otimes \cdots \otimes \hat{I}_M)(\hat{\rho}_1 \circ \hat{\rho}_2 \circ \cdots \circ \hat{\rho}_M) \\ &\quad \cdot (\hat{U}_1 \otimes \hat{I}_2 \otimes \cdots \otimes \hat{I}_M)^\dagger \\ &\quad = \hat{U}_1 \hat{\rho}_1 \hat{U}_1^\dagger \circ \hat{U}_1 \hat{\rho}_2 \hat{U}_1^\dagger \circ \cdots \circ \hat{U}_1 \hat{\rho}_M \hat{U}_1^\dagger \end{aligned} \quad (7)$$

The $N = 2$ and $N = 4$ cases are special, in that we are able to use all of the eigenstates as either background or pure states. In general, this will not be the case, and some states of the ancilla spins will correspond to "garbage blocks" rather than to pure states. Alternatively, additional smaller pure spin systems may be fit into these blocks. The blocks may be thought of as being parallel computers whose outputs we may distinguish by using the state of the ancilla.

We now indicate how our procedure is tolerant of errors. Perturbations in the preparation process, for example, due to the $|\omega_i - \omega_j|$ never being exactly zero, will not pose a problem for the above scenario for the following reason. Suppose there is some small error in preparing the initial condition, $\hat{\rho} = 2^{-N}\hat{I} + \hat{\rho}_\Delta + \epsilon\hat{\rho}_{\text{error}}$. By the same argument used in Eq. 2, the result of a complete computation (which can be expressed as the unitary transform $\hat{U}_{\text{qc}}$) is $2^{-N}\hat{I} + \hat{U}_{\text{qc}}\hat{\rho}_\Delta\hat{U}_{\text{qc}}^\dagger + \epsilon\hat{U}_{\text{qc}}\hat{\rho}_{\text{error}}\hat{U}_{\text{qc}}^\dagger$. The magnitude of the initial error is unchanged, and thus for small ϵ, standard statistical techniques can be used to extract the computational result efficiently. Errors that do perturb the evolution of the computation can be reduced through quantum-error correction (*21, 22*).

The principle behind our idea is that there may exist significant structure in the state distribution of a molecule's spins, even in thermal equilibrium. This structure corresponds to regularity from which we may distill something close to being a pure state (that has very low von-Neumann entropy, Tr [$\hat{\rho}$ log $\hat{\rho}$]), for a subset of the spins. Our procedure is a kind of compression, which makes good use of the negative entropy available in the initial state of the system. In Eq. 6, the excess entropy is contained in the constant offsets not shown between the qubit blocks. Alternatively, excess entropy can be rejected to the "garbage" blocks,

which are distinguished from those of interest by the ancilla spins. Thus, a subset of the available spins inside each molecule are used to label the remainder, creating conditionally pure states.

Finally, how well does this procedure scale? The density matrix for a pure state of k quantum bits (qubits) must have $2^k - 1$ zero eigenvalues. Under the small chemical shift assumption, the initial ρ_Δ for N spins will have $N!/[(N/2)!]^2 \approx 2^N/(N\pi/2)^{1/2}$ approximately zero eigenvalues (which come from states with equal spin up and spin down populations). Unitary transforms leave the spectrum of eigenvalues unchanged but may be used to permute the elements to arrange for a convenient block-matrix structure, as previously demonstrated. The necessary condition for preparing a k-qubit effective pure state is that:

$$k \leq \log_2\left[\frac{N!}{[(N/2)!]^2} + 1\right] \approx N - O(\log_2 N) \quad (8)$$

For large N, we may, therefore, fit at least one $k \approx N$ qubit computer in a molecule with N spins because, for example, logic gate sequences can be used to assemble together into one block the set of all states with equal numbers of up and down spins (these states have equal Boltzmann populations) and the single all spin-down state (which is highly populated in comparison).

However, this does not come for free—the price we pay is a decrease in the observed signal strength. Fundamentally, this is because we have not changed the overall temperature of the system; only a very small conditional subset has "cold" spins. In the high-temperature approximation we use here, the observed signal strength S of the deviation density matrix scales as:

$$S \propto nV\frac{N\alpha}{2^N} \quad (9)$$

where n is the molecular density, V is the sample volume, and $N\alpha/2^N$ is the probability of finding a particular N-spin configuration. For α fixed at room temperature, performing a $k = 10$ qubit calculation with an $N = 15$ spin molecule decreases S by ≈ 0.01 compared to a single-spin signal from the sample. Modern phase-cycling and subtraction techniques make it possible to see such a signal. Such a 10-qubit quantum computer explores a state space of size 2^{10}; no other (non-NMR) concept has come close to realizing such a large system.

To scale beyond 10 spins, improvements can be made by using more clever entropy packing techniques (*23*) that take advantage of additional structure in the initial thermal state, but these will only give additional polynomial factors of N. Gradient fields (*24*) or coupling to other degrees of freedom could also be used to perform nonunitary operations that can make use of more of the initial thermal state. Initial spin cooling may be used (such as by coherence transfer, optical pumping, or conventional refrigeration), which can increase the signal strength exponentially. It is also possible to design quantum algorithms that produce answers one bit at a time, trading off a logarithmic increase in the number of times a calculation needs to be run for greater signal-to-noise.

Single-spin operations. We have demonstrated that a pure state initial condition can be prepared from a thermal bulk spin ensemble. To complete the task of showing that this system may perform quantum computation, we need to show that two more things are possible: any unitary transformation (quantum logic gate), and final state readout. Taking advantage of the fact that any quantum logic gate can be constructed from combining arbitrary single qubit rotations and controlled-NOT gates (*25*), we concentrate here on showing how the traditional tools of pulsed NMR allow us to perform these two functions.

To start, we assume that we have uncoupled identical spins in a strong magnetic field $\mathbf{B} = B_0\hat{\mathbf{z}}$. The Hamiltonian for the spin degree of freedom is $H = -\boldsymbol{\mu}\cdot\mathbf{B} = -\gamma\hbar B_0\hat{I}_z$, where γ is the gyromagnetic ratio for the spin and $\hat{I}_z$ is the operator for the z component of the angular momentum. The time-advance operator $\exp(iHt/\hbar) = \exp(-i\gamma B_0 t\hat{I}_z) = \hat{R}_z(\theta = \gamma B_0 t)$ thus shows the time evolution to be a rotation by an angle θ about the $\hat{z}$ axis.

If a nearly resonant oscillating rf field is now applied in the $\hat{\mathbf{x}}$ direction, then in the rotating-wave approximation, we find that in the rotating frame the spin evolves under an effective field $\{B_1\cos(\varphi)\hat{\mathbf{x}}, B_1\sin(\varphi)\hat{\mathbf{y}}, (B_0 - \omega/\gamma)\hat{\mathbf{z}}\}$ (where φ is the rf phase). For spins close to resonance, even a moderately large B_1 can dominate the small deviation of ω from γB_0 due to chemical shifts, and so it is a good approximation to ignore the B_0 term and take all of the nearby spins to rotate around the effective field B_1 by the same amount. Because it is possible to independently rotate around each axis, it is possible to compose these operators to generate an arbitrary rotation. We shall use $R_\phi(\theta)$ to denote a single-spin rotation around the axis $\phi = \{x, y, z\}$.

Arbitrary single-spin rotations are still possible with interacting spins. For example, in a two-spin system such as (2,3)-dibromothiophene, if the interaction strength is weak compared to the Zeeman energy, then the Hamiltonian can be taken to have the scalar coupling form $H = \hbar\omega_A\hat{I}_{zA} + \hbar\omega_B\hat{I}_{zB} + \hbar\omega_{AB}\hat{I}_{zA}\hat{I}_{zB}$ (for strong coupling, the full $\hat{I}_A\cdot\hat{I}_B$ term must be used). These three terms lead to three rotation operators that commute and can be applied in any order. The first two give the time-evolution operators $R_{zA}(\omega_A t)$ and $R_{zB}(\omega_B t)$, rotations of the A and B spin degrees of freedom around $\hat{z}$. The third term is:

$$\hat{R}_{zAB}(\omega_{AB}t) = e^{i\omega_{AB}t\hat{I}_{zA}\hat{I}_{zB}} = \cos(\omega_{AB}t/2)\cdot\hat{I} + i\sin(\omega_{AB}t/2)\begin{bmatrix}1&0&0&0\\0&-1&0&0\\0&0&-1&0\\0&0&0&1\end{bmatrix} \quad (10)$$

The effect of this term is a coupled two-spin rotation.

We now assume that the chemical shift $\omega_A - \omega_B$ is sufficiently large so that the spins can be addressed individually. Then a 180° pulse on A about any axis ϕ has the property that it may be used to reverse the evolution of terms containing $\hat{I}_{zA}$ in the Hamiltonian. Namely, $\hat{R}_{\phi A}(180)\hat{R}_{zA}(\omega_A t) = \hat{R}_{zA}(-\omega_A t)\hat{R}_{\phi A}(180)$, and $\hat{R}_{\phi A}(180)\hat{R}_{zAB}(\omega_{AB}t) = \hat{R}_{zAB}(-\omega_{AB}t)\hat{R}_{\phi A}(180)$ (*26*). This mechanism is known as refocusing and is useful to apply selected terms in the Hamiltonian and to remove reversible broadening effects (such as spin interactions and magnetic-field inhomogeneity). Repeated fast refocusing, known as decoupling, is also useful because it completely stops the dynamics of the affected terms.

Thus, one way to apply a $\hat{x}$ rotation to one spin alone in the coupled two-spin system is by a refocusing pulse between free evolutions, by using $R_{xB}(180)\exp[iHt/\hbar]R_{xB}(180)\exp[iHt/\hbar] = R_{zA}(2\omega_A t)$. Of course, $R_z(\theta) = R_y(-90)R_x(\theta)R_y(90)$ can also be performed. Refocusing is more generally useful because of its selectivity. For example, if refocusing pulses are applied to both spins, then the only active term is $R_{zAB}(2\omega_{AB}t)$ (refocusing it twice has no net effect). If the two coupled spins cannot be addressed separately, standard NMR techniques can be used to effect selective and semiselective excitations by using longer nonselective pulse sequences (*15*).

The controlled-NOT. Given the ability to perform arbitrary single-bit operations, the next element required for quantum logic gates is a nonlinear interaction between spins, such as the controlled-NOT (CNOT) operation, which conditionally flips one spin based on the value of the other (*17*). The interaction terms in the Hamiltonian provide the required nonlinearity. For the two-spin case with scalar coupling, a CNOT can be implemented as a controlled phase shift preceded and fol-

lowed by rotations, given by the sequence $C_{AB} = R_{yA}(-90)R_{zB}(-90)R_{zA}(270 = -90)R_{zAB}(180)R_{yA}(90)$ (Fig. 1). Multiplying the rotation matrices corresponding to the pulses sequence shown in Fig. 1 yields the operator corresponding to

$$C_{AB} = \frac{1}{2^{5/2}}\begin{bmatrix}1 & -1 & 0 & 0\\ 1 & 1 & 0 & 0\\ 0 & 0 & 1 & -1\\ 0 & 0 & 1 & 1\end{bmatrix} \cdot \begin{bmatrix}1-i & 0 & 0 & 0\\ 0 & 1-i & 0 & 0\\ 0 & 0 & 1+i & 0\\ 0 & 0 & 0 & 1+i\end{bmatrix} \times \begin{bmatrix}1-i & 0 & 0 & 0\\ 0 & 1+i & 0 & 0\\ 0 & 0 & 1-i & 0\\ 0 & 0 & 0 & 1+i\end{bmatrix} \cdot \begin{bmatrix}1+i & 0 & 0 & 0\\ 0 & 1-i & 0 & 0\\ 0 & 0 & 1-i & 0\\ 0 & 0 & 0 & 1+i\end{bmatrix} \times \begin{bmatrix}1 & 1 & 0 & 0\\ -1 & 1 & 0 & 0\\ 0 & 0 & 1 & 1\\ 0 & 0 & -1 & 1\end{bmatrix} = \sqrt{-i}\begin{bmatrix}1 & 0 & 0 & 0\\ 0 & 1 & 0 & 0\\ 0 & 0 & 0 & 1\\ 0 & 0 & 1 & 0\end{bmatrix} \quad (11)$$

This is the CNOT operation (up to an irrelevant overall phase, which could be removed by extra rotations). In terms of the CNOT operation, we may now give the pulse sequence necessary to transform Eq. 19 to Eq. 20 in Box 1: $C_{AC}C_{BC}C_{CB}C_{CA}$, where the four spins are labeled alphabetically from right to left (note that no transform needs to occur to spin D).

An experimental demonstration of the feasibility of this CNOT sequence (albeit in a heteronuclear system) is routinely provided by the common NMR sequence INEPT (insensitive nuclei enhancement by polarization transfer) (*27*). This similar sequence, $R_{yA}(90)R_{zAB}(90)R_{xA}(90)$, is typically used to transfer the population from the abundant proton to the less-responsive carbon nucleus. The corresponding unitary operator has the same form as the CNOT matrix in Eq. 11, but there are phase differences among the non-zero elements.

Molecular structure determinations can involve sequences of hundreds of pulses that probe coupling in networks of hundreds of spins. To construct systems with many qubits, it is possible to use just the two-spin interaction terms because the process of refocusing or decoupling turns on or off interaction terms in the Hamiltonian based on the parity of the number of times that the affected spin appears in each term, and such a parity matrix is a complete basis. Furthermore, couplings are not required between all pairs of nuclei; it is only necessary to have local interactions because universal computations can be performed with a quantum cellular automata (*16*). It is also possible to directly use higher order terms to save steps [for example, six-quantum transitions have been observed in benzene (*28*)]; an open question is how the scaling of such higher order terms can be used to reduce the number of operations over an algorithm based solely on two-spin terms. More powerful techniques, such as average Hamiltonian theory (*15*), may also be useful for quantum computation.

Readout. Finally, we show how logical-state readout may be accomplished. A general quantum computation produces as output some pure state $|\psi\rangle = \hat{U}_{qc}|\psi_0\rangle$. The result is usually found from a measurement of the probabilities $|\langle\psi|\psi\rangle|^2$. Can these be obtained from an NMR quantum computer? Two issues must be considered: in an NMR system, only certain spin states contribute to observable experimental quantities. Furthermore, only ensemble averages are accessible.

Recall that our computational basis was chosen to be the longitudinal orientation of each spin. However, in an NMR apparatus, these terms do not contribute a signal. Rather, what is experimentally detected by the pickup coil is the net transverse magnetization:

$$nV\langle\mu_x + i\mu_y\rangle = nV\gamma\hbar\mathrm{Tr}[\hat{\rho}_\Delta(\hat{I}_x + i\hat{I}_y)] \quad (12)$$

In other words, if the density matrix is expanded in a basis of products of angular momentum operators (*15*), then the signal in the pickup coil gives the coefficients of the $\hat{I}_x$ and $\hat{I}_y$ terms. An arbitrary computation can be designed to move the result to these terms. Furthermore, the readout sequence can be designed such that only those blocks in the density matrix that correspond to a valid computation will output, either by conditioning the readout flip pulses on the ancilla state or by using phase cycling techniques. Therefore, measurement of the free induction decay can identify the coherence terms that contain the computation result (by the precession frequency) and read out the answer (by determining the amplitude and phase).

The second issue arises because the usual output from a quantum computation is nondeterministic, and thus averages out in an ensemble averaged measurement. However, this need not be the case, as we show here. Consider, for example, Shor's quantum factoring algorithm, whose output is a random number known to be close to some integer ratio c/r, where c is unknown (and nearly uniformly distributed), and r is the desired result (the answer to a discrete logarithm problem, which can be used to factor a composite number). If only one quantum computer were involved, projective measurement of each final qubit state would retrieve the random ratio, from which r could be determined efficiently with some classical calculation. However, when an entire ensemble of quantum computers exists and only ensemble averaged measurement results are available, this algorithm fails because each computer ends up in a final state with different c. That is, $\langle c/r\rangle$ is not necessarily rational, and thus r cannot be determined from it.

This problem can be resolved by a small modification of Shor's algorithm. In this case, the i^{th} quantum computer in the ensemble will obtain a result $x_i = c_i/r$, where c_i is a random (unknown) integer, which is almost certainly co-prime to r. Physically, x_i could be represented in each quantum computer as the state of some spins inside the

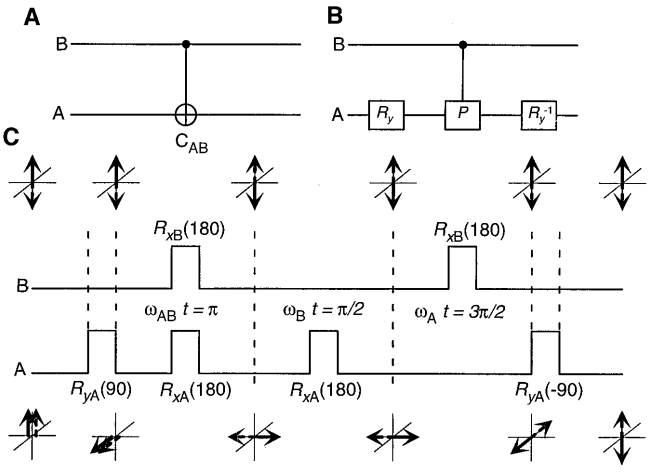

Fig. 1. (**A**) A controlled-NOT gate acting on two qubits, (**B**) the controlled-NOT gates implemented by a controlled phase shift gate (specified by a unitary matrix with diagonal elements {1, 1, 1, −1}) preceded and followed by π/2 rotations, and (**C**) the pulse sequence corresponding to the components in (B).

molecule. Now, a classical computation may also be performed on a quantum computer (for a spin computer, this need only be carried out within the T_1 time, and not the T_2 time, because phase damping is inconsequential for classical states), and thus using a continued fraction expansion algorithm, each quantum computer may determine c_i from x_i and thus obtain r (*4*, *29*). This algorithm involves the following steps to determine p and q from $x \approx p/q$: let $\lfloor x \rfloor$ denote the greatest integer less than or equal to x. A series of numbers $a_j = \lfloor 1/\xi_j \rfloor$ are calculated, where $\xi_0 = 1/x$ and $\xi_j = 1/\xi_{j-1} - a_{j-1}$. These then give a series of convergents (p_j, q_j) from the recurrence relation $p_0 = a_0$, $q_0 = 1$, $p_1 = a_1 a_0 + 1$, $q_1 = a_1$, $p_j = a_j p_{j-1} + p_{j-2}$, $q_j = a_j q_{j-1} + q_{j-2}$, where the desired (p, q) is one of the convergents. r is found to be the q_j that satisfies the discrete logarithm. All of this can be accomplished in polynomial time.

The freedom to append this continued fraction step to Shor's algorithm has been pointed out elsewhere (*30*). What was not recognized previously is that this is useful because it allows an ensemble to provide as an output a deterministic estimate of r. This output r will be randomly distributed [say, according to $p(r)$] about some mean $\langle r \rangle = \sum_r p(r)\, r$, whose value we wish to determine. Our apparatus can measure this result bit-by-bit in the following manner: Let r_m be the m^{th} bit of r. Then it follows that

$$\langle r \rangle = \sum_m 2^m \left[\sum_r p(r) r_m \right] \qquad (13)$$

where the bracketed term is experimentally observable. Therefore, in one experiment, the N quantities $\langle r_m \rangle$ are determined as described above, efficiently giving the desired result $\langle r \rangle$. For other known quantum algorithms (*31*–*33*), similar modifications can be used to make the algorithm give a deterministic output.

We have described one strategy for performing bitwise readout of the computation result. To emphasize the variety of possible strategies, consider an alternative and possibly more efficient approach: one may encode the state of multiple qubits into the phase of the signal output from each molecule. For example, quadrature phases can be used to encode the state of two qubits, and so on. The microscopic phases are averaged to become a detectable macroscopic quantity, and by virtue of the quantum algorithm, phases corresponding to the desired output signal will add coherently. However, because high phase resolution requires more observation time, a trade-off will exist between such an approach and the alternative of multiple repeated experiments. The general idea is that there exist many classical coding strategies that may improve the efficiency with which a bulk spin quantum computer can communicate the answer through the ensemble measurement channel.

A final practical but central issue is the problem of decoherence. Slow decoherence is one of the primary attractions of an NMR quantum computer. A global limit is set by the thermal relaxation time scale T_1, which can be up to thousands of seconds, and a more stringent limit by T_2, the phase damping time, which is shorter but still may be a few seconds long. The critical figure of merit is the ratio of the decoherence and logic gate time scales. Depending on the spin system chosen, single spin operations using semiselective pulses may require roughly 0.1 to 10 ms, whereas conditional dynamics through coupled evolution may require roughly 1 to 100 ms. Thus, a few thousand primitive logic operations may be feasible within the decoherence time.

The viability of bulk-spin quantum computing arises because on one hand, nuclear systems are very well decoupled from their environment, while on the other hand, we can obtain a strong signal from these systems because of the large ensemble. These two apparently contradictory attributes are compatible here because most of the Hilbert space of the total system is not used for quantum computation. The validity of this scenario is demonstrated by any multidimensional NMR experiment [such as correlated spectroscopy, or COSY (*26*)] in which off-diagonal peaks occur because of quantum coherences among spins. Study of these peak widths has led to a detailed understanding of the relevant decoherence processes (*15*), with the surprising conclusion that many multidimensional spectra may be adequately predicted by ignoring decoherence entirely. Long coherence times are the final requisite for practical NMR quantum computation.

Conclusions. With these primitives of preparing the initial state, performing arbitrary unitary transformations, and reading out the state, it is straightforward to program more complex operations. For example, the maximally entangled GHZ state (*34*) (which has yet to be experimentally realized) can be reached from three qubits with the sequence $C_{AC}\, C_{AB}\, R_{yA}(90)$. As a final programming example, we give the pulse sequence for a three-spin system to implement a Toffoli gate using controlled NOT operations (*25*): $R_{yC}(-45)\ C_{BC}\ R_{yC}(-45)\ C_{AC}\ R_{yC}(45)\ C_{BC}\ R_{yC}(45)$. This is a "controlled-controlled-NOT" gate in that when both A and B are logical one, then the state of C is inverted. Although this unitary transform matches the classical truth table, internally there are entangled and superposition states, and so the corresponding classical network of CNOT gates would not give the correct answer. This sequence (preceded by the appropriate initialization and followed by readout) can therefore serve as a proof-of-principle test for bulk-spin quantum computing (*35*).

A more ambitious proof-of-principle is a demonstration of Shor's factoring algorithm by using a 6-qubit system (which can be prepared by using $N = 9$ spins, with relative signal strength $S \approx 0.1$) to "factor" the number 15 with the procedure described in (*36*). Assuming perfect timing, a few tens of rf pulses would be needed to prepare the initial state, ~40 more to perform the computation, and a final few tens of pulses to initiate the readout. This experiment is a daunting challenge for the current ion-trap quantum computer but should be feasible for many NMR systems (by using more complex molecules than the simple introductory examples shown here).

We have introduced to the study of quantum computation some familiar notions from multidimensional pulsed NMR (the ability to use decoupling and refocusing to selectively manipulate a Hamiltonian) and some new ones (the use of extra spins to prepare an ensemble density matrix deviation from thermal equilibrium that has collective states that act exactly like low-dimensional pure states, and the usefulness of quantum algorithms with deterministic outputs). We believe that our observations have significant practical implications: Complex experimental apparatus to isolate a small number of spins can be replaced by simple naturally occurring materials, with the massive redundancy providing macroscopically detectable signals, so that experimental quantum computing may become a widely accessible field. Although extensions to handle the problems of scaling beyond ~10 qubits may be possible, even a system of that size offers an unprecedented opportunity to observe and explore superfast quantum algorithms, quantum error correction, and the evolution of decoherence and entanglement.

REFERENCES AND NOTES

1. P. Shor, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE Computer Society, Los Alamitos, CA, 1994).
2. D. P. Divincenzo, *Science* **270**, 255 (1995).
3. S. Lloyd, *Sci. Am.* **273**, 44 (1995).
4. A. Ekert and R. Jozsa, *Rev. Mod. Phys.* **68**, 1 (1996).
5. J. Cirac and P. Zoller, *Phys. Rev. Lett.* **74**, 4091 (1995).
6. C. Monroe *et al*., *ibid.* **75**, 4714 (1995).
7. S. Bandyopadhyay and V. Roychowdhury, *Jpn. J. Appl. Phys.* **35**, 3350 (1996).
8. P. Domokos, J. Raimond, M. Brune, S. Haroche, *Phys. Rev. Lett.* **52**, 3554 (1995).
9. I. L. Chuang and Y. Yamamoto, *Phys. Rev. A* **52**, 3489 (1995).
10. Q. A. Turchette *et al*., *Phys. Rev. Lett.* **75**, 4710 (1995).
11. W. G. Unruh, *Phys. Rev. A* **51**, 992 (1995).
12. I. L. Chuang, R. Laflamme, P. W. Shor, W. H. Zurek, *Science* **270**, 1633 (1995).

13. F. Bloch, *Phys. Rev.* **70**, 460 (1946).
14. T. E. Chupp, E. R. Oteiza, J. M. Richardson, T. R. White, *Phys. Rev. A* **38**, 3998 (1988).
15. R. R. Ernst, G. Bodenhausen, A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Oxford Univ. Press, Oxford, 1994).
16. S. Lloyd, *Science* **261**, 1569 (1993).
17. D. P. DiVincenzo, *Phys. Rev. A* **50**, 1015 (1995).
18. K. Wago *et al*., *J. Vac. Sci. Technol. B* **14**, 1197 (1996).
19. S. Lloyd, *Phys. Rev. Lett.* **75**, 346 (1995).
20. This matrix is often called the "reduced density matrix" in the NMR literature [see, for example, (*15*)]. However, we shall not use that terminology here because it may be confused with a different meaning of the same term in quantum-information theory, where it refers to the density operator obtained by a partial trace over a subsystem.
21. A. Steane, *Phys. Rev. Lett.* **77**, 793 (1996).
22. P. Shor, *Phys. Rev. A* **52**, 2493 (1995).
23. B. Schumacher, *ibid.* **51**, 2738 (1995).
24. Y. Zhang, W. Mass, D. Cory, *Mol. Phys.* **86**, 347 (1995); see also D. Cory, A. Fahmy, T. Havel, *Proc. Natl. Acad. Sci. U.S.A.*, in press.
25. A. Barenco *et al*., *Phys. Rev. A* **52**, 3457 (1995).
26. G. Mateescu and A. Valeriu, *2D NMR: Density Matrix and Product Operator Treatment* (Prentice Hall, Englewood Cliffs, NJ, 1993).
27. G. Morris and R. Freeman, *J. Am. Chem. Soc.* **101**, 760 (1979).
28. W. Warren, S. Sinton, D. Weitekamp, A. Pines, *Phys. Rev. Lett.* **43**, 1791 (1979).
29. G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers* (Oxford Univ. Press, London, ed. 4, 1960).
30. P. Shor, *SIAM J. Comput*., in press.
31. D. Deutsch and R. Jozsa, *Proc. R. Soc. London Ser. A* **439**, 553 (1992).
32. D. Simon, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE Computer Society, Los Alamitos, CA, 1994), pp. 116–123.
33. A. Y. Kitaev, *Los Alamos Natl. Lab. E-print* quant-ph/9511026 (1995).
34. D. Greenberger, M. Horne, A. Shimony, A. Zeilinger, *Am. J. Phys.* **58**, 1131 (1990).
35. D. Beckman, A. N. Chari, S. Devabhaktuni, J. Preskill, *Los Alamos Natl. Lab. e-print* quant-ph/9602016 (1996).
36. Details of the pulse sequence are available at http://physics.www.media.mit.edu/projects/spins/spins.html which introduces spin computing and is linked to http://physics.www.media.mit.edu/projects/spins/pulses/pulses.html.
37. Supported by the MIT Media Lab's Things That Think consortium. I.L.C. acknowledges financial support by the Fannie and John Hertz Foundation. We are grateful for the hospitality of the Institute for Theoretical Physics (NSF grant PHY94-07194), and the Institute for Scientific Interchange (with support from Elsag-Bailey). Many people have contributed to the coherent evolution of these ideas, including S. Lloyd, D. P. DiVincenzo, P. W. Shor, C. Bennett, J. Smith, J. Jacobson, T. Toffoli, N. Margolus, and W. H. Zurek.

10 September 1996; accepted 2 December 1996

Alterations in Synaptic Strength Preceding Axon Withdrawal

H. Colman, J. Nabekura,* J. W. Lichtman†

Permanent removal of axonal input to postsynaptic cells helps shape the pattern of neuronal connections in response to experience, but the process is poorly understood. Intracellular recording from newborn and adult mouse muscle fibers temporarily innervated by two axons showed an increasing disparity in the synaptic strengths of the two inputs before one was eliminated. The connection that survived gained strength by increasing the amount of neurotransmitter released (quantal content), whereas the input that was subsequently removed became progressively weaker, because of a reduction in quantal content and a reduction in quantal efficacy associated with reduced postsynaptic receptor density. Once the synaptic strengths of two inputs began to diverge, complete axonal withdrawal of the weaker input occurred within 1 to 2 days. These experiments provide a link between experience-driven changes in synaptic strength and long-term changes in connectivity in the mammalian nervous system.

The ability of the nervous system to respond to experience in an enduring way may depend on alterations in the structure or function of synaptic connections. An indelible synaptic alteration induced by experience early in postnatal life is the loss of some of the axonal inputs that converge on a target cell (*1*). A large body of work also indicates that the strength of existing synapses can be potentiated or depressed in response to activity (*2*), and several investigations have suggested that such changes in synaptic efficacy ultimately lead to structural plasticity (*3*). However, the relation between alterations in synaptic strength and permanent structural changes in synaptic connectivity is not well understood.

The neuromuscular junction is a simple and accessible place to examine the relation between functional and structural synaptic changes, especially during early postnatal life when each muscle fiber undergoes a transition from polyneuronal to single innervation (*4*). Earlier attempts with the use of techniques with relatively low sensitivity failed to detect functional correlates of synapse loss at the neuromuscular junction, and led to the hypothesis that the loss of synaptic transmission during synapse elimination must be abrupt and must result from the sudden degeneration of the eliminated axonal branch and all of its synapses (*5*). Subsequent anatomical studies, however, provided no evidence of degeneration (*6*) and later studies demonstrated a progressive loss of synaptic area before axon withdrawal (*7*). Utilizing more sensitive physiological techniques, we have now reexamined whether alterations in synaptic strength occur prior to axonal withdrawal.

Intracellular recording from muscle fibers (n = 600) in the mouse trapezius muscle, chosen because of its favorable anatomical features, showed that most fibers underwent the transition from multiple to single axonal innervation during the first two postnatal weeks with some muscle fibers achieving single innervation substantially earlier than others (*8*). On postnatal day 2, approximately three-fourths of muscle fibers were multiply innervated (>95 percent by two axons), whereas about one-third were multiply innervated on day 6 and less than one-tenth remained multiply innervated at day 10. The progressive loss of polyneuronal innervation in the trapezius indicated that synapse elimination here (as elsewhere) was not occurring synchronously on each postsynaptic cell. Rather, some target cells achieved single innervation much sooner than others. Thus, if there were functional changes in synaptic strength associated with the elimination of synapses, at any one time, different muscle fibers should be at different stages in this process.

In order to independently activate two axons converging at the same junction, suction electrodes were applied to two nearby nerve branches projecting to the same region of the neonatal mouse trapezius (Fig. 1A). Muscle fibers innervated by two axons (Fig. 1B), one input traveling through each of these nerve branches, were selected for study by intracellular recording. High magnesium (10 to 17 mM) recording solution was used to reduce the size of the evoked endplate potentials (EPPs) (Fig. 1C) (*9*), so that the quantal content of each input could be measured by repetitive stimulation (mean = 770) with the method of failures (*10*).

Changes in quantal content. During the first 10 days after birth, the average quantal contents of the inputs to individual fibers diverged (Fig. 2A). At young ages (P1 to

The authors are in the Department of Anatomy and Neurobiology, Box 8108, Washington University School of Medicine, 660 South Euclid, St. Louis, MO 63110, USA.

*Present address: Department of Physiology, Kyushu University Faculty of Medicine, 3-1-1 Maidashi Higashiku, Fukuoka, 812-82 Japan.
†To whom correspondence should be addressed.