

Toward a table-top quantum computer

by Y. Maguire
E. Boyden
N. Gershenfeld

To investigate the scaling issues for quantum computation using thermalized ensembles of spins, we are building a rudimentary quantum computer using nuclear magnetic resonance (NMR). We discuss the experimental issues facing the scalability of this technique including the number of qubits, issues of classical versus quantum behavior, and polarization. Due to the weak signals measured in NMR, the most significant challenge to practical devices is polarization enhancement. We present the architecture for a simple NMR spectrometer that is table-top, low-cost, and software-radio driven using commodity electronics and unconventional permanent magnet designs. Optimizations specific to using this machine as a computational device, such as probe design, compilation, and on-line control, are discussed.

With the quest for faster microcomputer chips in the semiconductor industry, exponential increases in speed must come at the exponentially increasing cost of some other resource. Advances in the semiconductor industry have capitalized on two resources: space, by shrinking the size of integrated circuitry in semiconductor fabrication; and money, by the ever increasing investment in fabrication plants. It is likely that we will, in the not too distant future, reach fundamental physical limits: transistors will be atomic in scale and memory will consist of only a handful of electrons. Time does not offer much help as a resource, because light takes 100 picoseconds (ps) to travel the 3 centimeters (cm) of a typical chip. Exponentially large physical space or energy are not realistically available resources. People have proposed DNA computers¹ as a type of massive parallelism to solve computational problems. This

solution does not scale either, because a computationally hard problem requires exponential amounts of DNA, which can quickly require samples more massive than the earth itself.¹ The dream of quantum computing is that quantum mechanics provides an untapped, scalable exponential resource. Quantum computers can, in theory, factor numbers in exponentially fewer steps than classical computers (thus cracking modern cryptographic protocols), simulate quantum systems with exponentially fewer resources than now possible, and search databases in $\mathcal{O}(\sqrt{N})$ queries rather than $\mathcal{O}(N)$ on an ordinary machine.

The power of quantum computing begins with a profound observation. Whereas a classical bit can be in one of two states, 0 or 1, a quantum bit, or “qubit” (e.g., a spin-1/2 particle) can be in a weighted superposition of both states. For example,

$$a|0\rangle + b|1\rangle$$

where a and b are complex numbers, and $|a|^2 + |b|^2 = 1$. When measured with a readout operator, the qubit appears to collapse to state $|0\rangle$ with probability $|a|^2$, and to state $|1\rangle$ with probability $|b|^2$. The state of two qubits can be written as

$$|\Psi\rangle = \sum_i c_i |i\rangle = a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle$$

©Copyright 2000 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

where $|a|^2 + |b|^2 + |c|^2 + |d|^2 = 1$, and the probability of measuring the amplitude of each state is given by the magnitude of its squared coefficient. In general, the state of n qubits is specified by $2^{n+1} - 1$ real numbers—an exponential amount of information, relative to the number of physical particles required. Most of these states are entangled—to create them requires some kind of interaction between the qubits, and the qubits cannot be treated entirely independently from one another.

Designing quantum algorithms is made difficult by the fact that it is impossible to extract the entire exponential amount of information encoded in the state of a quantum system: only a polynomial amount of information can be read out, since a measurement collapses the exponential enormity of the state into a handful of classical values, with probabilities of occurrence given by the amplitudes as described in the previous paragraph. Furthermore, the only computational primitives allowed are unitary operators (which evolve the state of a system under Hamiltonian dynamics) that are reversible and therefore conserve probability. Nevertheless, a collection of wonderful algorithms have been discovered. In 1994, Peter Shor² showed quantum computers could factor large numbers into prime factors in polynomial steps, compared to exponential steps³ on classical computers. Prime factorization is an essential part of modern public key cryptographic protocols, paramount to privacy in the electronic world. In 1997, Lov Grover⁴ demonstrated a search algorithm that could search an unsorted list for a single element quadratically faster on a quantum computer than on a classical computer.

Introduction to NMR

Liquid state nuclear magnetic resonance (NMR) is the manipulation of bulk degrees of freedom to obtain information about the underlying nuclear spins in an applied magnetic field. In order to understand how to build a quantum computer using NMR, we begin with the general requirements for a quantum computer.⁵

1. It starts in a known state (a pure state such as the ground state).
2. It is able to efficiently represent any unitary mapping from one state to another from a basic set of elementary unitary mappings.
3. It is able to concatenate these mappings to form a circuit, that is maintain quantum states long enough to represent any unitary mapping.

4. It performs projective measurement in order to read out the computation results.

To implement any of the above requires that we understand the dynamics of nuclear spins when they act as an ensemble. Quantum statistical mechanics allows us to compute the equilibrium thermal state of the liquid. Moreover, if we know the dynamics of a single molecule, we can easily calculate the dynamics of the entire liquid. To begin, the dynamics of quantum systems are described by operators that map one quantum state to another, $|\Psi\rangle \rightarrow |\Phi\rangle$. For the nuclear spins of our quantum computer, operators are finite dimensional matrices. The Hamiltonian \mathcal{H} is a special operator that describes the complete dynamics of a molecule under time evolution and also determines the energy eigenstates of the system. The Hamiltonian of a rapidly tumbling molecule in a liquid containing n spin-1/2 nuclei with a strong magnetic flux density $\vec{B} = B_0 \hat{z}$ is given by

$$\mathcal{H} = \sum_{i=1}^n \hbar \gamma_i (B_0 + \delta B_i) I_{iz} + \sum_{i < j} 2\pi \hbar J_{ij} I_{iz} I_{jz}$$

where I_{iz} represents the z -axis spin operator for nuclei i , γ_i is the gyromagnetic ratio of spin i , δB_i describes a local magnetic flux density adjustment for each spin i due to the local chemical environment, and J_{ij} describes the coupling, in Hertz, between spins i and j . We assume for the discussion of the dynamics below that we are in the weak coupling limit where $|\gamma_i \delta B_i - \gamma_j \delta B_j| \gg J_{ij}$. In equilibrium, the Hamiltonian dictates that spin i precesses quickly about the z -axis at a frequency $\gamma_i B_0 / 2\pi$ adjusted by the local term, $\gamma \delta B_i / 2\pi$, the so-called chemical shift. This precession frequency is called the Larmor frequency. This fast oscillation (usually at several MHz) is slowly modulated by the J-coupling (which usually proceeds at a few Hz). These precession states are the equilibrium states of the system, called the eigenstates. With the energy eigenstates, we are now ready to begin to understand the spin dynamics. Given a $(2^n - 1)$ -dimensional basis $|a\rangle$ of single-molecule eigenstates, we can write the state of a thermal liquid ensemble of molecules as a density matrix ρ , where the matrix element ρ_{ab} is given by

$$\rho_{ab} = \overline{c_a c_b^*}$$

where the average is taken over the thermal ensemble, and $c_a c_a^*$ is the probability of being in the state $|a\rangle$. Given a time-evolution operator

$$U(t_1, t_0) = e^{(i\mathcal{H}(t_1-t_0))/\hbar}$$

where \mathcal{H} is the Hamiltonian of the quantum system during the period from time t_0 to time t_1 , the density matrix evolves as

$$\rho(t_1) = U\rho(t_0)U^{-1}$$

To know the dynamics of the system for all future times, we just need to know the initial state of the system. Quantum statistical mechanics provides the answer with the thermalized (maximum entropy) density matrix

$$\rho(0) = e^{-\mathcal{H}/k_B T} / \mathcal{Z} = \frac{1}{\mathcal{Z}} \left(\mathbf{1} + \sum_{i=1}^n \frac{\hbar \gamma_i B}{kT} I_{iz} + \dots \right)$$

The partition function \mathcal{Z} is the normalization factor such that $\text{tr}(\rho(0)) = 1$. Because the Zeeman energies $\gamma_i B_0$ are usually about 10^{-5} times smaller compared to $k_B T$, the first two terms of the Taylor expansion of the exponential serve as a very good approximation, with the second term being about 10^{-5} times the magnitude of the first, constant term. If our measurement apparatus is a coil wrapped around the sample detecting the classical magnetization signal, it turns out that the constant term is not observable, so we ignore it and concentrate on the remainder (called the deviation density matrix). The only time-evolution operators of interest in this simple model of NMR (where we neglect relaxation mechanisms and higher order spin-spin interactions), will be the single-qubit rotations ($\theta = \gamma_i B_1 t$),

$$X_i(\theta) = e^{i\theta I_x}$$

which corresponds to a magnetic flux density $\vec{B}_1 = 2B_1 \cos(\gamma_i(B_0 + \delta B_i)t)\hat{x}$, resonant with spin i , for a time t . If we think of the spin dynamics in the frame rotating about the z-axis at its Larmor frequency, we can decompose this linearly polarized field into two counter-rotating circularly polarized fields, one of which rotates with this “rotating” frame, while the other rotates oppositely at $2\gamma_i(B_0 + \delta B_i)$ relative to the nuclear magnetization. We can ignore the latter term and concentrate on the former. Then B_1 becomes a static field and B_0 becomes zero, a much easier system to analyze. The \vec{B}_1 field simply rotates the spin in the Y - Z plane in this rotating frame. We can create this field with the same measurement coil and just control the amplitude, phase, and time duration of our radio-frequency (RF) field. If we time the pulse such that

$\gamma_i B_1 t = \pi/2$ for the RF field, this is known in NMR literature as a $(\pi/2)_x$ pulse. Similarly,

$$Y_i(\theta) = e^{i\theta I_y}$$

correspond to performing similar operations along the y-axis. The double-qubit rotation

$$ZZ_{ij}(\theta) = e^{-i\pi J_{iz} I_{jz}}$$

corresponds to evolving a pair of J-coupled spins for a time t .

It turns out the X_i , Y_i , and ZZ_{ij} time-evolution operators are computationally complete, in that together they can approximate any unitary operator with any desired accuracy. Using our coil and an arbitrary waveform RF generator, we can satisfy the second of the quantum computing requirements. This coil actually allows us to satisfy the last requirement as well since we use it to perform projective measurement. One subtle point is that some quantum algorithms have been designed to give out probabilistic results from single quantum systems. Fortunately, it has been shown⁶ that they can be modified to give deterministic results for ensemble systems.

Decoherence is the process by which a quantum system loses quantum coherence to its environment and becomes classical. In NMR, decoherence is measured by how long a spin on one molecule stays in phase coherence with another molecule. This interval is called the spin-spin relaxation time, T_2 . The beauty of NMR for quantum computing is that T_2 is much larger than the gate speed. This allows hundreds, even thousands of quantum gates to be implemented, allowing the third quantum computing requirement to be satisfied. One can see why this is possible: the nucleus is a perfect, frictionless trap, surrounded by a cloud of electrons that shield most external decoherence mechanisms; we are spreading relatively few degrees of freedom over the $\mathcal{O}(2^{10^{23}})$ -sized Hilbert space of a typical liquid. No other system to date has shown such long coherence times and, paradoxically, the ability to control the nuclear states on fast time scales. Quantum error correction is a wonderful theory that has shown that one can compute indefinitely if the decoherence rate per gate is less than a threshold,⁷⁻¹⁰ which has been estimated to be between $\mathcal{O}(10^{-4})$ to $\mathcal{O}(10^{-6})$ per gate.¹¹ At the cost of many more ancilla spins, NMR is the closest demonstrated technology ($\mathcal{O}(10^{-3})$) to reaching this threshold. By modify-

ing the chemistry of the liquids, researchers at IBM¹² are improving this threshold for NMR.

We have not satisfied the first quantum computing requirement, the creation of a pure state. A thermalized liquid only gives us the mixed state provided by the almost maximum entropy mixed state. One might think this is impossible, since to go from a high entropy state to a zero entropy state requires non-unitary operators. The key insight from a couple of groups is that there are degrees of freedom for rejecting this entropy.^{6,13,14} Logical labeling uses extra spins to label the remaining spins as ground state spins; spatial labeling uses external gradient fields to erase terms in the density matrix by creating a spatial superposition of quantum computers and then mixing them together; temporal labeling creates a series of experiments that are temporally averaged to give an initial ground state and the corresponding evolution. For prototyping our quantum computer, we use an inefficient but conceptually simple method, exhaustive averaging. For two qubits, one might run three experiments, for example, using preparatory gates to create deviation density matrices

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.6 & 0 & 0 \\ 0 & 0 & -0.6 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -0.6 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0.6 \end{bmatrix},$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0.6 & 0 \\ 0 & 0 & 0 & -0.6 \end{bmatrix}$$

performing the desired computation on each, then averaging the three experimental results—which gives the same result as if one had performed the computation on a state equal to the average of the three initial density matrices,

$$\begin{bmatrix} 1.33 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$+ \begin{bmatrix} -0.33 & 0 & 0 & 0 \\ 0 & -0.33 & 0 & 0 \\ 0 & 0 & -0.33 & 0 \\ 0 & 0 & 0 & -0.33 \end{bmatrix}$$

The second term is unobservable and the first term is a pure state. This works because quantum mechanics is linear, and just as a Cartesian vector can be resolved into x , y , and z terms, we can resolve our density matrix into more easily manipulated components.¹⁴

Issues for thermal nuclear resonance

In this section we discuss issues facing the scaling of liquid nuclear resonance to nontrivial computational complexity. We begin with spin polarization and its essential role in the scalability of quantum computers, then address the both quantum and classical nature of NMR in thermal liquids, and finally discuss how to increase the number of qubits in the system.

Polarization. It seems counterintuitive to be able to make the observable term of the thermalized density matrix behave exactly as if it were in the ground state of the system. Delving deeper, consider the largest signal that can be observed by logical labeling to be the difference between the most probable and least probable terms of the density matrix:^{15,6}

$$S = ((\rho_{eq})_{max} - (\rho_{eq})_{min}) \gamma \hbar (I_z)_{max}$$


Solving this using the equation for the thermalized density matrix $\rho(0)$ given earlier in this section yields

$$S = \frac{N}{2^N} \gamma \hbar \frac{\sinh(N\beta\gamma\hbar B_0/2)}{(\cosh(\beta\gamma\hbar B_0/2))^N}$$

The plot of this function is shown in Figure 1A, with a baseline set as spin noise, where quantum information is lost to the environment via stimulated emission with the coil.¹⁶ It is apparent that between 20 and 50 qubits, no current technique can measure the quantum computing signal. Looking at the problem in the domain of polarization, Figure 1B shows the polarization needed to extract a specific number of qubits with a signal-to-noise ratio of 3 decibels (dB) relative to spin noise. This plot shows the spin polarization must approach unity to make the system perform better than classical systems. Both plots show that the Boltzmann limit is not scalable for quantum computation.

It would be nice to dramatically increase the spin polarization but keep rotational and translational degrees of freedom thermalized. Optical pumping¹⁷ is a technique potentially able to accomplish this. Each quantum of circularly polarized light carries one unit

Figure 1 (A) shows the intrinsic magnetization versus temperature as a function of the number of qubits. Plot (B) shows the polarization versus the maximum number of qubits with a 3 dB signal-to-noise ratio.



of angular momentum that can put an atom into a highly polarized electronic state. The electronic polarization can then polarize the nuclear spins through relaxation mechanisms. In noble gases that do not interact and have simple electronic orbitals, close to 50 percent polarizations have been obtained. Happer et al.¹⁸ did pioneering work on optical pumping of rubidium gases that transfer polarization to either xenon-129 or helium-3. Transferring this polarization to nuclear states efficiently is an open, difficult question, but one that needs to be addressed. Another reason to have optical photons interact with nuclear spins is for quantum error correction. To implement any error correction algorithm, one needs a way to reject errors in the calculation in a way that does not influence the other spins that carry valuable information. A photon channel could be a means to do this—if a spin is polarized by optical pumping its previous state is erased. In summary, polarization is the most serious challenge facing practical quantum computation using ensembles of thermalized nuclear systems. If it can be done, remaining challenges should be surmountable, as will be shown below.

Is it really quantum? Even before we begin to address issues of polarization, it is paramount to determine if a mixed state system, in which the quantum system is represented by an ensemble of

quantum systems, actually has the requirements to perform quantum computing. This was a question in the original paper by Gershenfeld and Chuang.⁶ The paper showed how to do state preparation, algorithms, and readout on an NMR system, but it was not clear that quantum information processing was actually being carried out in the system. Shack and Caves¹⁹ have almost answered this question. They show that for a number of qubits less than 16, at a modest spin polarization of $\varepsilon = 2 \times 10^{-6}$, the density matrix during the entire experiment is fully separable. Being separable means that the statistics of any measurement of the density matrix can be fully described by classical probabilities. In the high temperature limit the quantum computing signal can be simplified to

$$S = \frac{N \hbar \gamma B_0}{2^N 2k_B T} = \frac{N}{2^{N+1}} \varepsilon$$

The density matrix of the spins in NMR can be described classically¹⁹ if

$$S \leq \eta \equiv \frac{1}{1 + 2^{2N-1}}$$

This means that for N smaller than a certain number, entanglement cannot exist. One could rephrase

the calculations and ask “What is the minimum polarization needed so that a bulk NMR system with a number of qubits greater than one has a density matrix that is not separable?” The minimum polarization is 44.4 percent. Fortunately, this is not the end of the story: quantum computing is not just about the density matrix, it is also about unitary evolution. Shack and Caves¹⁹ show that if a unitary transform can be written as an outer product of matrices

$$V = V_1 \otimes V_2 \dots \otimes V_N$$

then a classical description with positive transition probabilities also can fully describe the system.

However, when they tried to include entangling unitaries, operators that cannot be written as in the equation above, they found a classical system could describe the system with classical transition probabilities only at the expense of an exponential decrease ε in signal-to-noise ratio for each step. They further showed that an exponential signal decrease per step can be avoided for K entangling unitaries by adding N ancillary qubits. K is the largest integer such that

$$\varepsilon \leq \eta^{K+1}$$

For the classical system to be able to use a single entangling unitary with no signal reduction at a polarization of 2×10^{-6} , this corresponds to $K = 6$ qubits. One could pose the question, “What is the minimum polarization such that for $N \geq 2$, no entangling unitaries can be simulated by the classical system without exponential signal decrease?” For 2 qubits, this corresponds to 3.9×10^{-3} , above the range of current spectrometers, but within the realm of possibility of even conventional polarization enhancement techniques such as the Nuclear Overhauser Effect²⁰—dynamic nuclear polarization with electrons. The model used by the authors for avoiding exponential signal decrease is only one of many possibilities: for a system to perform genuine information processing, all possible classical models cannot account for the entangling unitary evolutions without an exponential cost. Most believe that although the initial state of thermal NMR is separable, the dynamics will not be.

The question of whether NMR is quantum or not is an academic question—unity polarization, as addressed in the previous section, is paramount for scalable quantum computation. With unity polarization

both the initial states and the dynamics will surely be quantum mechanical.

Number of qubits. The last major barrier for thermal NMR quantum computing is the number of qubits. To be able to implement error correction on a quantum computer may require somewhere between 10 and 100 qubits to describe a single qubit.¹⁰ To represent thousands of fully coupled qubits with separate resonances is impossible with current techniques. Fortunately, Seth Lloyd²¹ provided a solution. If one has a polymer chain of the type $D-A-B-C-A-B-C \dots$, where each letter corresponds to a unique resonance of a nuclear species, then this system is capable of performing any quantum computation. This architecture suffers a slight slowdown to pass information up and down the chain, but the huge gains that quantum algorithms offer are still preserved. This architecture is called a Single Instruction Multiple Data (SIMD) cellular automata. Recently, it has been shown²² that only $\dots A-B-A-B \dots$ is necessary, which will make molecular synthesis considerably easier.

Hardware for the table-top quantum computer

We have begun the development of instrumentation to address the issues raised in the previous section. An NMR quantum computer starts with a collection of molecules in a liquid, where each molecule is a redundant quantum computer. In the presence of our homogeneous magnet field, a small percentage of these spins align with the magnetic field and begin to precess. By applying amplitude, frequency, phase and time controlled magnetic energy, we can rotate the spins around an arbitrary axis in the Hilbert space. To do this, we wrap a coil around the liquid sample, and add reactive elements to make it resonant at several frequencies. We use a digitally controlled transmitter that generates these waveforms. With sensitive electronics and a switch, we can detect the magnetization of the liquid and perform a measurement of the quantum state. A general system diagram is shown in Figure 2.

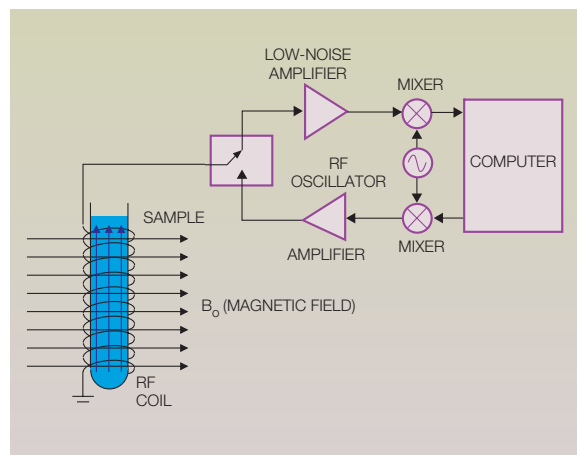
Computational optimizations. Current commercial NMR spectrometers are incredible machines capable of doing a range of tasks, such as imaging and quantitative chemical structure determination. With the instrumentation outlined below, we hope to match current spectrometers in performance but end up with a device that can be treated more as a computer. Since we are trying to make a computer, there

are optimizations that we can do to improve our system performance and size over commercial NMR spectrometers. The first has to do with the sample—commercial systems need to routinely analyze a number of samples. As a result, these systems are engineered for cylindrical symmetry for easy access. The ends contribute an inhomogeneity in the RF field, leading to errors in the computational gates. The major constraint for this configuration is that the active region of a sample cannot be larger than about 1.5 cm in length and 1 cm in diameter. For our computer, we need only a single sample and can engineer our system for spherical symmetry with no end effects. A spherical probe would have higher RF homogeneity and allow much smaller magnet designs.

Since we are interested in a computational result rather than a chemical spectrum, we can use a single qubit to serially produce an answer. This poses an enormous advantage since we can make a very high Q electrical,²³ optical,²⁴ or mechanical readout system (decoupled from the transmitter in the last two cases) that could dramatically improve our sensitivity. Another technique is to use an electron channel to measure an answer. This is an old technique²⁰ called ENDOR, or Electron-Nuclear Double Resonance. Going through the electron channel will optimistically give a gyromagnetic ratio of improvement. The gyromagnetic ratio is inversely proportional to rest mass, and thus approximately three orders of magnitude improvement in signal can be obtained. Species-dependent relaxation mechanisms may make this approach insubstantial, but is worth investigating further for candidate quantum computing molecules.

Magnet design. To make a table-top system, one needs a compact, homogeneous magnetic field. We are investigating permanent magnet structures to solve this problem. There are a number of constraints for the system: small enough to fit on a table-top; as high a magnetic flux density as possible; a homogeneity of the flux density across the sample of about one part per billion. The latter two points require some justification. The basic signal-to-noise ratio for a coil is proportional to $B_0^{7/4}$.²⁵ Given NMR signals are extremely small, it is important to maximize this signal-to-noise ratio. If there are spatial field variations, there will be spins in the sample that will resonate slightly faster or slower than others. With the spins out of phase, this is an artificial means of introducing coherence that shortens the maximum possible coherence time of the liquid. Empirically, it has been found that for a homogeneity of one part per

Figure 2 General system schematic for an NMR quantum computer



billion, all experimentally induced decoherence will be smaller than the natural decoherence processes. Due to manufacturing tolerances, this lofty goal cannot be reached in practice solely from the magnet alone. The field must be corrected in a controlled manner with slight perturbations, a process called shimming. With shimming, the magnet tolerance required is about 1 to 10 parts per million (ppm). This will be explained in more depth in the next section.

We have done finite element modeling of two basic magnet geometries. The first is based on an arrayed magnet geometry invented by Klaus Halbach.²⁶ It can be shown²⁷ that for a magnetization distribution within a cylindrical volume as shown in Figure 3, one can obtain a perfectly homogeneous solution within the volume of field strength

$$|B| = B_r \ln(r_2/r_1)$$

B_r is the remanent flux density, that is the flux density output in a zero magnetic field. This is quite surprising since one is able to obtain flux densities inside the volume that are much larger than this remanence. To construct this magnet, Halbach developed an approximation using segments of magnets that achieve a high order approximation to this field. The finite element model is shown in Figure 4. The design goal was to achieve 1 ppm homogeneity over a 1 cm diameter sphere in the middle of the magnet. This was accomplished with $r_1 = 2$ cm, $r_2 = 6.5$ cm, $|\vec{B}| = 1.513$ T.

Figure 3 Magnetization distribution within the magnet that yields a perfectly homogeneous field inside. This is true for both cylinders and spheres.





Figure 4 Finite element model of the Halbach array with a discretized magnetization approximation that yields a 1 ppm homogeneity inside



This is very encouraging, but from a practical point of view has a couple of limitations. Manufacturing tolerances on modern rare earth magnets, usually made by sintering, pressure bonding, or extruding

are about 1 to 10 thousandths of an inch and the error on the magnetization angle can approach two degrees. These two facts make the best possible homogeneity on the order of one part per million. Two insights lead to alternatives. The first is to construct the magnet out of many tiny magnets. With enough magnets, the statistics of the distribution of errors for the ensemble of magnets average out high spatial frequencies in the magnetic field. The second method is more practical and simply requires judicious placement of high permeability materials. From Laplace's equation, taking the Fourier transform in space yields

$$\Phi(\mathbf{k}) = \frac{1F(\mathbf{k})}{(\mathbf{k} \cdot \mathbf{k})\mu}$$

where $F(\mathbf{k})$ is the spatial Fourier transform of $\nabla \cdot \mathbf{M}$, the magnetization. The frequency component $F(\mathbf{k})d\mathbf{k}$ will be attenuated by a factor of $(\mathbf{k} \cdot \mathbf{k})\mu$. For ferromagnetic materials with μ_r approaching 100000, the high permeability enhances the filtering of Laplace's equation by a factor μ_r over using permanent magnet materials alone. Fermilab²⁸ took advantage of this point in the design of their anti-proton recycler ring, where they required small, homogeneous magnets. They iteratively shaped the pole pieces of a traditional H-loop design.

Shimming. Shimming, or the modification of the bias magnetic field by precisely controlled electromagnets (or, in some cases, permanent magnets, μ -metal, and chunks of steel), can increase the homogeneity of the bias magnetic field by several orders of magnitude. This is essential for NMR due to the fact that nonidealities in the magnetic field can greatly reduce the lifetime of quantum coherences and blur features in the NMR spectrum. The observed signal will be a sum of the signals of many different isochromats, or collections of spins at the same frequency. A broadening of a waveform in the frequency-domain is associated with a shrinking of the length of the signal in the time-domain, since as the signal components at various frequencies dephase and become random with respect to one another, their sum tends toward zero. To perform a quantum computation, the signal must not decay below the noise level until the computation is complete. Thus we require that the homogeneity of our magnetic field be within one part in 10^9 .

In a region without currents, a magnetic field obeys Laplace's equation, $\nabla^2 \vec{B} \sim \nabla^2 B = 0$, which can be

solved in terms of spherical harmonics. (We set the x - and y -components of the magnetic field to zero, and consider only the z -component, which we refer to as B , in the following discussion.) Spherical harmonics are enumerated by indices (l, m) , which are called the *order* and the *degree*, respectively. For $m = 0$, the *zonal* harmonics, one has cylindrical symmetry (i.e., the polar coordinate ϕ is irrelevant), which is appropriate for many magnets, where the dominant inhomogeneity is due to cylindrical end effects. The equation for B at the point (r, θ, ϕ) is then of the form

$$B(r, \theta) = \sum C_n(r/a)^n P_n(\cos\theta)$$

where P_n is the n th Legendre polynomial, and C_n and a are parameters. In general, the complete solution for all (l, m) (including the *tesseral* terms, for which $m \neq 0$) is, for the point (r, θ, ϕ) ,

$$B = \sum_{l=0}^{\infty} \sum_{m=0}^l C_n(r/a)^l P_{lm}(\cos\theta) \cos(m(\phi - \psi_{lm}))$$

where P_{lm} is called the *associated Legendre polynomial*, and the last ϕ dependent cosine term sinusoidally modulates the magnetic field as one traces a circle in the xy plane. All the harmonics except for $l = m = 0$ contribute to the inhomogeneities, since a homogeneous field, by definition, has a magnitude which does not depend on any coordinates. Since the higher-order terms rapidly decrease in amplitude for most realistic solutions of Laplace's equation, adjusting a few terms should suffice to remove most inhomogeneities: this is the motivation for including between 10 and 30 shim coils with every modern NMR spectrometer.

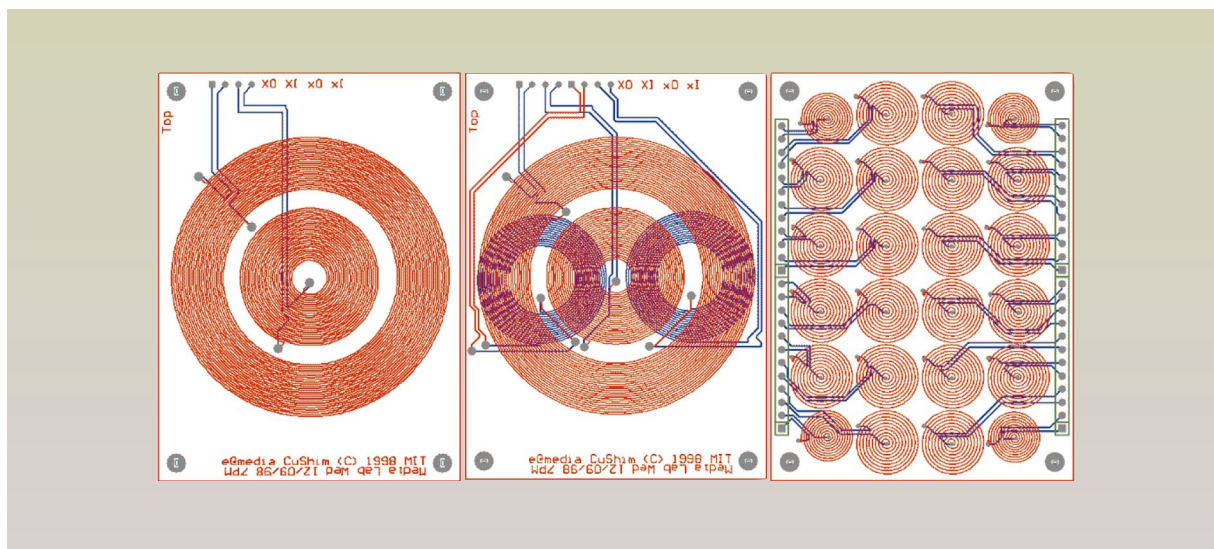
Spherical harmonics are orthogonal, so in theory one should be able to tune each coil just once. In reality they are not precisely orthogonal, since a small translation of one coil with respect to another leads to high-order interferences between the various orders. Thus shimming with orthogonal coils is a slightly non-linear problem, and for a spectrometer with n shim coils, one must conduct a search over an n -dimensional space of coil currents, trying to maximize the height and narrowness of the peaks in the spectrum. Many cumulative weeks of human effort are usually required for the shimming operation for the lifetime of a spectrometer, but small corrections are usually automatable. We are very interested in making shimming an on-line process, where a classical computer

is continuously correcting the shim currents for maximum resolution with no human intervention. The purpose is to begin to abstract away the spectrometer aspect of the machine and focus on the computational aspects. To do so will require modern non-linear search methods, as will be described in the software section.

For a 1 tesla (T) magnetic flux density with part-per-100000 homogeneity, in order to get part-per-billion homogeneity, the error in these small fields must be on the order of 10^{-5} Gauss, roughly 10000 times smaller than the Earth's magnetic field. The physical coils we developed are planar, as shown in Figure 5. They are simply designed and fabricated using normal techniques for printed circuit boards. This is an old technique, as planar, orthogonal shimming was developed²⁹ in the 1950s to shim electromagnets, but fell out of favor when superconducting magnets became adopted. We have built electronics to address any number of coils: each coil has a bipolar 16-bit digital-to-analog converter (DAC) capable of outputting between -100 and 100 milliamperes (mA) of current, so that the least significant bit corresponds to $305 \mu\text{A}$ of current. For a 25-turn coil of diameter 5 cm, this creates a magnetic field with magnitude ~ 1 gauss (G) (1 cm away from the coil center), so that the least significant bit corresponds to about 10^{-5} G. Thus the specified hardware is, at least in theory, capable of creating the necessary fields. The actual performance depends on many things, including the stability of the shimming system's power supply, the susceptibility of the electronics to external noise, and unavoidable noise imposed by the laws of thermodynamics.

Transmitter. To generate arbitrary time series RF signals without prohibitive amounts of analog hardware, we have investigated minimal software radio approaches. With permanent magnet systems, resonance frequencies will reach only about 80 megahertz (MHz). This is well within the domain of the best digital-to-analog converters. In making a robust transmitter for NMR, one must be very careful with transmitter noise, phase noise of the oscillator, and long-term stability. To begin development, we began with a VXI-based, 40 million sample per second DAC made by Hewlett-Packard. The HPE1445 has 13 bits of digital resolution and generates arbitrary waveforms of any duration, and the clock is derived from a very stable oven-controlled crystal oscillator (OCXO). Initially, we have been interested in creating a system that can address two heteronuclear species such as protons and carbon-13 (many more

Figure 5 Planar shim coils to replace complex orthogonal coils in commercial spectrometers



qubits can exist in each heteronuclear subspace). We constructed a small analog board that takes the DAC output and, with filtering, creates simultaneous arbitrary waveforms for both species. A photograph of the board is shown in Figure 6. The signal from the DAC is split into two paths and is filtered, then converted by an RF mixer to each respective Larmor frequency. The local oscillators, also made by Hewlett-Packard, are switched between transmit and receive mode by gallium arsenide switches made by Microwave Associates (M/A-COM). These switches were chosen for their unusually high isolation of about 70 dB between channels and fast switching speed. High isolation is important because any unnecessary noise leakage to the receiver will degrade receiver performance and allow high power transmit pulses to be demodulated in the receiver. Control of the transmitter switches and other functions on this board and the receiver was accomplished with a Hitachi 32-bit processor.

Receiver. The signal magnitude produced by a coil interacting with the spin magnetization after a $(\pi/2)$ pulse is given³⁰ by

$$S \propto N_s \beta (\gamma B_0)^2 (L V_s \eta)^{1/2},$$


where N_s is the number of spins, β is a parameter to indicate the inhomogeneity of the field B_1 , L is the coil inductance, V_s is the sample volume, and η

is the ratio of the volume of the active sample region to the coil. The noise term, due to fluctuation-dissipation, is the traditional Johnson type (ignoring equipartition and spin-noise¹⁶):

$$N = \sqrt{F 4 k_B T R \Delta f},$$

where F is the noise figure of the receiver, $k_B T$ is Boltzmann's constant times the temperature, R is the resistance of the coil, and Δf is the input receiver bandwidth. For designing a receiver, we see it is important to minimize both the input bandwidth and the noise figure. Since signals in NMR are on the order of -150 to -130 dBm (dB relative to 1 mW), careful attention to these details is paramount for detecting signals. For the receiver, we started with a VXI-based, 20-million sample per second analog-to-digital converter (ADC) made by Hewlett-Packard. The HPE1437 is an amazing device capable of representing 23 bits of noise-free digital resolution. The maximum dynamic range of an ADC without processing gain is $1.76 + 6.02N$ dB, where N is the nominal bit resolution. Thus, the theoretical dynamic range of this device is 140.22 dB. However, since we routinely detect signals that are -130 dBm, analog electronics are still required to get the maximum digital resolution. With this, we developed a minimal receiver that performed some basic amplification and filtering. The signals from the protons and carbon-13

Figure 6 Photograph of the transmitter board with a 25-cent coin to give scale




are amplified by an Avantek 105 amplifier, a 1.6 dB noise figure amplifier with 27 dB of gain. After a wide, bandpass filter, the signal is downconverted, filtered, and then amplified with an integrated device, the RF2612 made by RF Micro Devices. The signal is then buffered with an op amp and then combined for the HPE1437. The overall RF specifications for the board are a 1.7 dB NF, with 60–104 dB of gain. A photograph of the board is shown in Figure 7. The ADC digitally downsamples the signal to obtain the amplitude and phase of the signal and then applies digital filtering and downsampling to reduce the output data rate. What we hope we have shown is not that this system is superior to commercial systems, but is a simple, inexpensive, and software-driven one.

FPGA version. The system that we have developed should be capable of doing all the tasks necessary to implement a quantum computer, but it is hardly table-top in size or function. With so many VXI modules and separate pieces, we sought a minimal implementation of all these functions. Field Programmable Gate Arrays (FPGAs) are software programmable devices that give one the ability to develop small microprocessors, memory controllers, waveform generators, and timers usually imple-

mented in application-specific integrated circuits (ASICs). We are developing a couple of boards around these devices that we expect will be able to amply do all the required tasks of our current system. The design that is currently being implemented is shown in Figure 8. Control of the board is implemented using the Internet Protocol (IP) over Ethernet. An FPGA serves as a controller that interfaces to a small Ethernet physical interface and Memory Access Controller (MAC) using the User Datagram Protocol (UDP) over IP. A three-wire serial bus is optoisolated and carries data over a bus to the transmitter and receiver cards. The transmitter card has an FPGA controlling 512 kilobytes (kB) of static RAM and an Analog Devices AD9856, a 200 MHz complementary metal oxide semiconductor (CMOS) digital upconverter. Over Ethernet, one supplies the frequency, amplitude, and phase of a signal digitally and this $3'' \times 2''$ card creates the waveform with submicrosecond time resolution. The same analog receiver design described in the previous section is implemented here except that the signal is mixed down to a standard intermediate frequency of 455 kHz. A third board containing an FPGA with 512 kB of SRAM and the AD9260 samples the signal at 10 MSPS (mega samples per second) and 16-bit resolution. On the FPGA, digital filters and downsamplers are applied

Figure 7 Photograph of the receiver board



and the signal is digitally downconverted with a quadrature local oscillator to recover amplitude and phase. Over the three-wire serial bus, the data are transmitted to the FPGA Ethernet controller and can be sent to any device on our network to be visualized and processed. Timing control for this whole system is controlled by a small oven-controlled crystal oscillator, which has excellent phase noise characteristics and temperature/time stability. To add other spins, one simply needs to add a stack of the three aforementioned boards, a high power switch, and RF amplifier.

Software

In this section we discuss the pieces of software needed to run on our quantum computer, including low-level programming to address our hardware and optimization software for shimming.

Pulse language and system design. We designed a programming language,³¹ q , to control the hardware presented in the last section. A q program is parsed by an interpreter, QOS (or Quantum Operating System), which runs on a PC. The job of QOS is to coordinate the timing and configuration of the hardware, compile pulse sequences into waveforms for broadcast, and process the signals that return. It is

the digital core of the software-radio architecture of this NMR system. A q program is essentially assembly microcode of the basic operations for an NMR quantum computer. A sample q program that performs one of the three pulse sequences required for executing Grover's algorithm, which performs a search on N unsorted items with only $\mathcal{O}(\sqrt{N})$ database accesses, is briefly explained in the accompanying sidebar.

Optimization for shimming. We built a machine-learning environment for QOS. Our initial priority was to continuously automate the tedious shimming process, but it quickly became apparent that with appropriate control software, the NMR spectrometer could automatically perform very complicated tasks. For example, one might then run a multidimensional NMR experiment like COSY (CORrelated Spectroscopy), extracting the various J -couplings between spins, as well as any other interactions that might be apparent. This is a very difficult task, requiring excellent hardware and software, since most real molecules in natural environments are difficult to differentiate from their surroundings.

All nonlinear search optimization techniques require a function to be optimized. Usually evaluating this function is an expensive thing to do—in our case,

Declare the program type

```
% This is a comment
% A name for this program – Grover's algorithm to discover the
% |11> state
Qname('grover|11>')
Qsample('13CHC13')      % Molecule – chloroform
Qversion('1.0 esb')     % Program version
Qdecouplingstyle('minimal') % Whether decoupling is on/off
```

Declare the molecular parameters

```
Qnumqbs(2) % number of qubits
% Proton(H) on chloroform is first qubit – give NMR specific information
% for qubit
Qspecies(1,'1H', 26.7510e7, 42574931.880108, 35000000, 1, 0.001, 5)
% Carbon(13C) is 2nd qubit
Qspecies(2,'13C', 6.7263e7, 10705040.871934, 8000000, 0.0159, 0.001, 5)
```

Prepare the state (for temporal averaging purposes)

```
% A square (hard) pulse (as opposed to a soft gradient or sinc pulse)
% applied about the y-axis for 90 degrees
Qy(1,90,square)
% wait for a specified time in decimal units
Qwait(0.002325581395349, dec, 0)
Qy(2, -90, square)
% An x-pulse
Qx(2, -90, square)
Qy(2, 90, square)
Qy(1, -90, square)
Qx(1, -90, square)
Qy(2, 90, square)
Qwait(0.002325581395349, dec, 0)
Qy(1, -90, square)
Qx(1, -90, square)
Qy(1, 90, square)
Qy(2, -90, square)
Qx(2, -90, square)
```

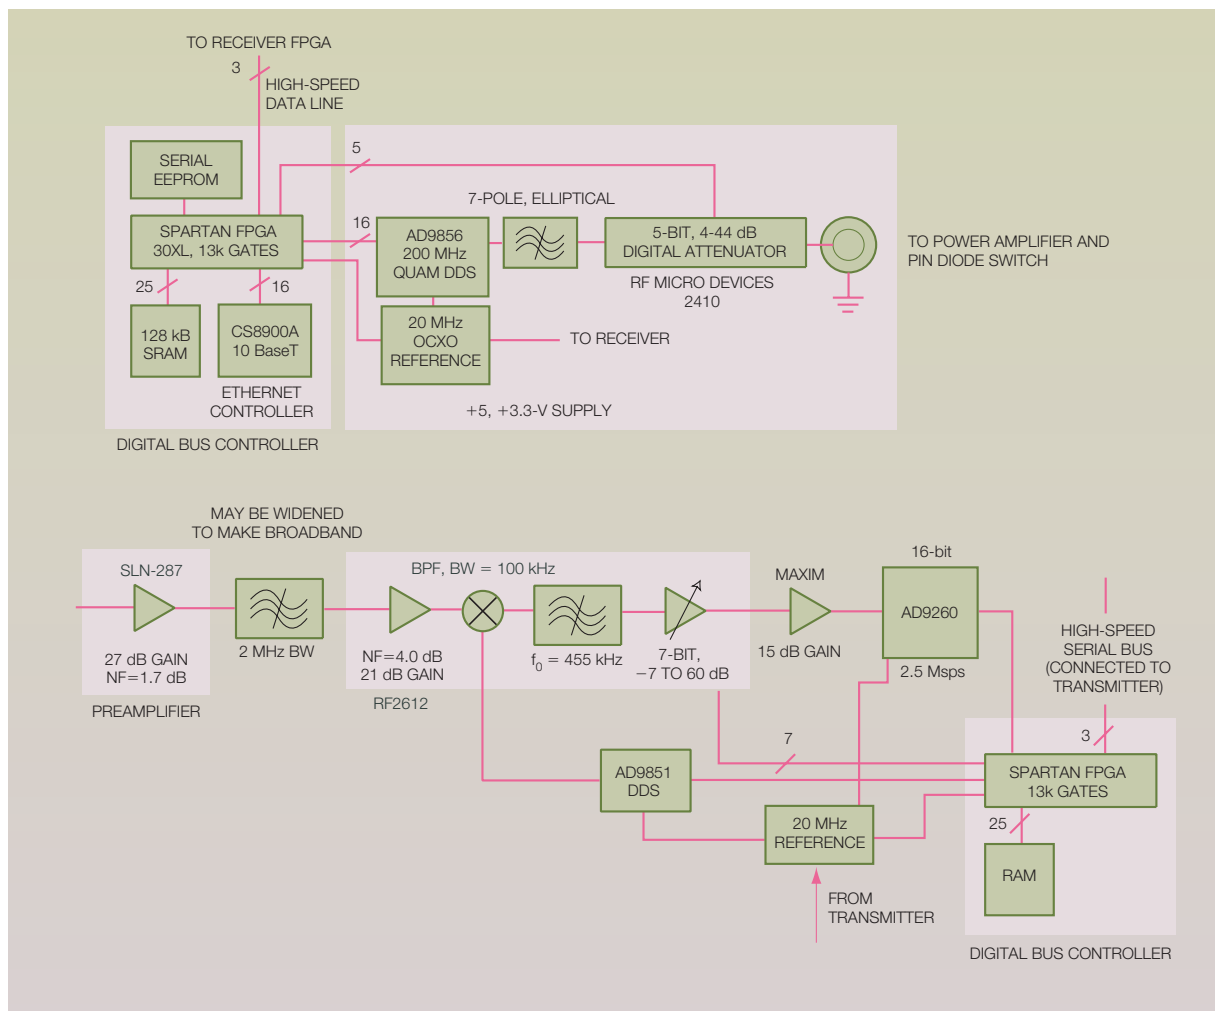
Perform Grover's algorithm including readout

```
Qx(1,180,square)
Qy(1,90,square)
Qx(2,180,square)
Qy(2,90,square)
Qwait(0.002325581395349, dec, 0)
Qy(2,-90,square)
Qx(2,90,square)
Qy(1,-90,square)
Qx(1,90,square)
Qwait(0.002325581395349, dec, 0)
Qy(2,-90,square)
Qx(2,-90,square)
Qy(1,-90,square)
Qx(1,-90,square)
```

requiring many seconds per execution. Since we want to maximize the dephasing time (thus improving the homogeneity of the magnetic field), a suitable function is the sharpness of a peak in the 1-D NMR spectrum. Thus a simple protocol for calculating this

sharpness is: run a $\pi/2$ pulse program with specific values for the parameters being optimized, and collect the data from the NMR spectrometer. Then compute the power spectrum of the signal, and fit a Lorentzian (as described earlier) to the lineshape,

Figure 8 Overall schematic for the transmitter and receiver electronics for an FPGA-based software radio



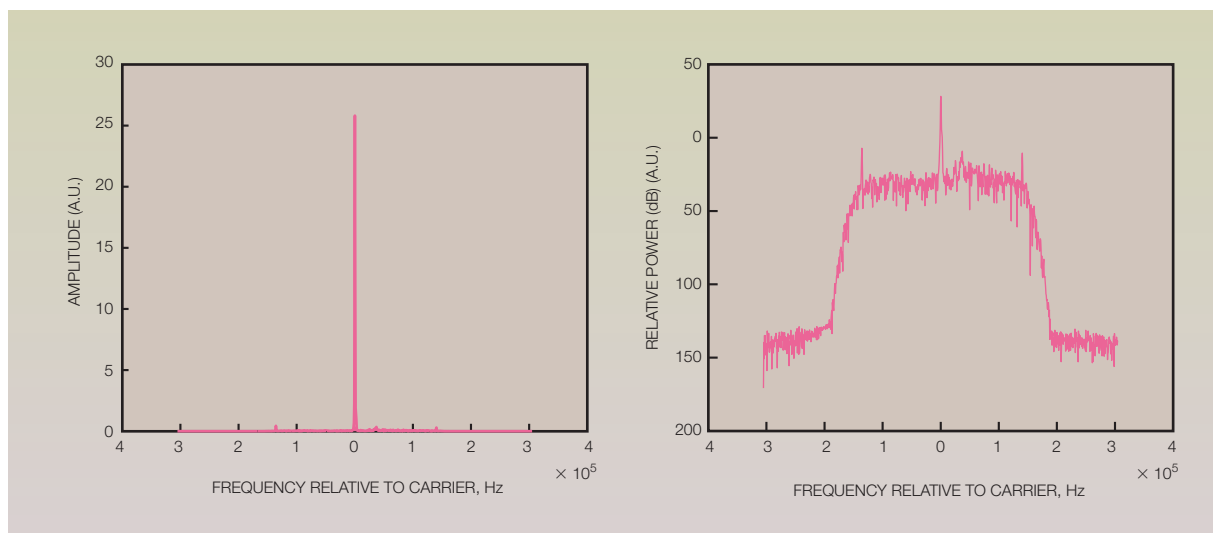
extracting the sharpness ω_0/a . The sharpness can be maximized using any of a number of methods—the Nelder-Mead downhill simplex method, Powell’s direction set method, or the conjugate gradient method (using finite differences to calculate the gradient).

Results

Our results are preliminary but indicative of things to come. The plots outlined below are relatively trivial on commercial spectrometers, but our results come from a machine that is significantly less in cost and complexity. The plot in Figure 9 shows the in-

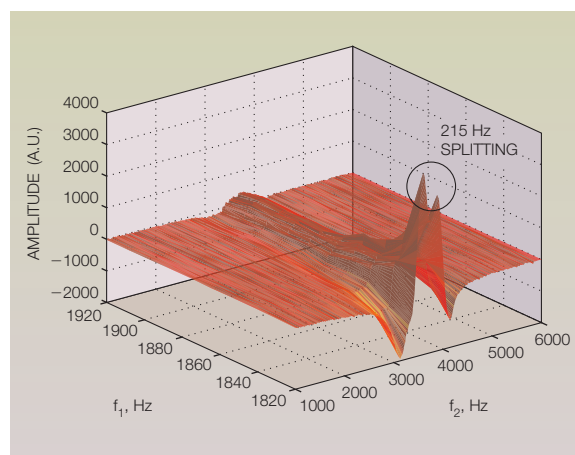
trinsic sensitivity of our spectrometer on a 0.5 cm^3 50 percent glycerine in water sample. The sample was prepared in a 5 millimeter NMR tube and the data were collected for 3.28 milliseconds. This is short compared to the intrinsic coherence time of the protons in the liquid, but the sample has been inhomogeneously broadened by our low homogeneity test electromagnet. With over 50 dB of signal to noise over a bandwidth of approximately 1 kHz, our system is competitive with commercial systems, and with computational optimizations we are pursuing, we hope to develop our system beyond what is done commercially.

Figure 9 Plot showing the sensitivity of the proton spectrum of glycerine, showing an excellent signal-to-noise ratio is in excess of 40 dB. The linear amplitude plot is on the left with a power plot on the right. The data were collected for 3.28 ms, and the bandwidth of the signal is approximately 1 kHz.



The next plot is much more important. It is a correlation spectroscopy plot of the correlation of one spin to another. In this case, we are showing the correlation of a carbon atom and a hydrogen atom, proof that we can in principle create logical gates with our system. Richard Ernst³² invented this technique as a means to show coupling patterns between spins, that is to identify $\pi J_{ij} I_{zi} I_{zj}$ terms of the Hamiltonian from ωI_i , terms that were ambiguous in previous efforts. Without delving into the details of the rich subject of 2-D correlation spectroscopy, we have demonstrated an experiment that isolates the coupling term to show that the nonlinear term of the Hamiltonian is present. The experiment goes as follows: we apply a $\pi/2$ pulse on the hydrogen spin that creates a $(1/\sqrt{2})(|1\rangle + |0\rangle)$ state, then we let the system evolve for a variable time τ_1 that is changed from experiment to experiment. During this time, the $|0\rangle$ state of the hydrogen spin will interact with the $|0\rangle$ state of the carbon spin to make this state resonate slightly faster during this time. The other $|1\rangle$ state of the hydrogen spin will interact with the $|0\rangle$ state of the carbon spins to resonate slightly slower. The difference in resonance frequency is the J coupling frequency. Our sample was 1 milliliter (mL) of C^{13} labeled chloroform with a coupling frequency of 215 Hz. A plot is shown in Figure 10, highlighting a splitting of the hydrogen resonance by 215 Hz. This is

Figure 10 Three-dimensional plot showing the resolution of J -coupling in chloroform, required to implement gates. This proton spectrum shows a 215 Hz splitting via the interaction with C^{13} .



direct evidence of the coupling between spins, which with shimming will be used to implement gates. We could not go further because of the inhomogeneity of our magnet. With our shimming system, we should be able to implement a simple quantum algorithm.

Conclusions

We have demonstrated the feasibility of a table-top NMR apparatus performing quantum computations, approaching the relevant performance of commercial spectrometers, and potentially scaling to non-trivial sizes through optimizations for computation rather than spectroscopy. Success beyond engineering diligence almost wholly depends on the success of unity polarization enhancement for thermal ensembles of nuclear spins. Optical pumping techniques show promise but still fall short of what is required for this technique to be useful. However, if this major obstacle can be overcome, we will have a machine that will have unprecedented computational degrees of freedom. This could potentially be very important not only for implementing quantum algorithms but also for controlling and manipulating large coherent quantum systems. Such a system could one day lend itself to simulating problems that physicists and semiconductor engineers only dream of. The important realization is that nature itself is a very powerful computational architecture. We just needed to figure out how to talk to it.

Acknowledgments

We gratefully acknowledge the MIT Media Lab's Things That Think consortium and the support of the Defense Advanced Research Projects Agency (DARPA) under the Nuclear Magnetic Resonance Quantum Computer (NMROC) initiative. We thank our collaborators at Stanford, Berkeley, MIT, and IBM. Yael Maguire acknowledges financial support from IBM and the National Science Research Council of Canada (NSERC).

Cited references and note

1. L. Adleman, "Molecular Computation of Solutions to Combinatorial Problems," *Science* **266**, 1021–1023 (1994).
2. P. W. Shor, "Algorithms for Quantum Computation: Discrete Logarithms and Factoring," *Proceedings, 35th Annual Symposium on Foundations of Computer Science*, IEEE, New York (1994), pp. 124–134.
3. It has not been ruled out that a better classical algorithm exists that would make the difference between classical and quantum factoring algorithms insubstantial.
4. L. K. Grover, "Quantum Mechanics Helps in Searching for a Needle in a Haystack," *Physical Review Letters* **79**, 325–328 (1997).
5. I. L. Chuang et al., "Bulk Quantum Computation with Nuclear Magnetic Resonance: Theory and Experiment," *Proceedings of the Royal Society of London* **454** (1998), pp. 447–467.
6. N. A. Gershenfeld and I. L. Chuang, "Bulk Spin-Resonance Quantum Computation," *Science* **275**, 350–356 (1997).
7. D. Aharonov and M. Ben-Or, "Fault Tolerant Quantum Computation with Constant Error," <http://xxx.lanl.gov/abs/quant-ph/9611025>.
8. A. Yu Kitaev, "Fault-Tolerant Quantum Computation by Anyons," xxx.lanl.gov/abs/quant-ph/9707021.
9. D. Gottesman, "Theory of Fault-Tolerant Quantum Computation," *Physical Review A* **57**, No. 1, 127–137 (January 1998).
10. J. Preskill, "Reliable Quantum Computers," *Proceedings of the Royal Society of London A* **454**, No. 1969 (1998), pp. 385–410.
11. A. Steane, "Quantum Computing," *Reports on Progress in Physics* **61**, No. 2, 117–173 (1998).
12. C. Yannoni, M. Sherwood, L. Vandersypen, D. Miller, M. Kubinec, and I. Chuang, "Nuclear Magnetic Resonance Quantum Computing Using Liquid Crystal Solvents," *Applied Physics Letters* **22**, No. 75, 3563–3565 (1999).
13. D. Cory, A. Fahmy, and T. Havel, "Ensemble Quantum Computing by NMR Spectroscopy," *Proceedings of the National Academy of Science, USA*, **94** (1997), p. 1634.
14. E. Knill, I. Chuang, and R. Laflamme, "Effective Pure States for Bulk Quantum Computation," *Physical Review A* **57**, No. 5, 3348–3363 (1998).
15. W. S. Warren, "The Usefulness of NMR Quantum Computing," *Science* **277**, 1688–1689 (1997).
16. T. Sleator et al., "Nuclear-Spin Noise and Spontaneous Emission," *Physical Review B* **36**, No. 4, 1969–1980 (1987).
17. A. Kastler, "Optical Methods for Studying Hertzian Resonances," *Science* **158**, 214–221 (1967).
18. W. Happer, E. Miron, S. Schaefer, D. Schreiber, W. A. V. Wijngaarden, and X. Zeng, "Polarization of the Nuclear Spins of Noble-Gas Atoms by Spin Exchange with Optically Pumped Alkali-Metal Atoms," *Physical Review A* **29**, 3092–3110 (1984).
19. R. Shack and C. M. Caves, "Explicit Product Ensembles for Separable Quantum States," *Journal of Modern Optics* **47**, No. 2–3, 387–399 (2000).
20. C. P. Slichter, "Principle of Magnetic Resonance," Springer-Verlag, Inc., New York (1996).
21. S. Lloyd, "A Potentially Realizable Quantum Computer," *Science* **261**, 1569–1571 (1993).
22. S. C. Benjamin, "Schemes for Parallel Quantum Computation Without Local Control of Qubits," *Physical Review A* **61**, No. 2, 020301/1–4 (2000).
23. P. Broekaert and J. Jeener, "Suppression of Radiation Damping in NMR in Liquids by Active Electronic Feedback," *Journal of Magnetic Resonance Series A* **113**, No. 1, 60–64 (1995).
24. J. M. Kikkawa and D. D. Awschalom, "All-Optical Magnetic Resonance in Semiconductors," *Science* **287**, 473–476 (2000).
25. Y. Maguire, *Towards a Table-Top Quantum Computer*, M.Sc. thesis, MIT, Cambridge, MA (1999).
26. K. Halbach, "Permanent Magnets for Production and Use of High Energy Particle Beams," *Proceedings of the Eighth International Workshop on Rare Earth Cobalt Permanent Magnets and Their Applications*, Dayton, Ohio (1985), p. 103.
27. M. G. Abele, *Structures of Permanent Magnets: Generation of Uniform Fields*, John Wiley & Sons, Inc., New York (1993), pp. 1–67.
28. "Permanent Magnets and the Antiproton Recycler Ring: A Step in the Right Direction," *Cern Courier* **37**, No. 4 (May 1997).
29. J. E. Golay, "Field Homogenizing Coils for Nuclear Spin Resonance Instrumentation," *Review of Scientific Instruments* **29**, No. 4 (1958).
30. D. I. Hoult and R. E. Richards, "The Signal-to-Noise Ratio

- of the Nuclear Magnetic Resonance Experiment,” *Journal of Magnetic Resonance* **24**, 71–85 (1976).
31. E. Boyden, *Quantum Computation: Theory and Implementation*, M.Sc. thesis, MIT, Cambridge, MA (1999).
 32. R. E. Ernst, G. Bodenhausen, and A. Wokaun, “Principles of Nuclear Magnetic Resonance in One and Two Dimensions,” Clarendon Press, Oxford (1987).

Accepted for publication July 5, 2000.

Yael Maguire *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: yael@media.mit.edu).* Mr. Maguire is a Ph.D. candidate at the MIT Media Laboratory. He is an IBM Student Fellow holding a master’s degree in media arts and sciences from MIT. Mr. Maguire received a degree in engineering physics with Highest Honors from Queen’s University in Canada. He is currently working on building a table-top device for quantum computing, which holds the promise of computing faster than traditional, semiconductor-based computers. He currently holds an NSERC fellowship from Canada and is generally interested in the fundamental ties between information processing and physics.

Edward S. Boyden III *Stanford Neurosciences Program, Beckman Center B103, Stanford, California 94305 (electronic mail: eboyden3@stanford.edu).* Mr. Boyden is a doctoral student in neurosciences at Stanford University, where he works on molecular analyses of small neural networks using genetic, electrophysiological, and optical techniques. He received his master’s degree in electrical engineering and computer science, and his bachelor’s degree in physics, both from MIT. He worked on many of the projects in the MIT Physics and Media Group throughout his undergraduate years.

Neil Gershenfeld *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307.* Dr. Gershenfeld leads the Physics and Media Group at the MIT Media Laboratory and directs the Things That Think research consortium. His laboratory investigates the relationship between the content of information and its physical representation, from developing molecular quantum computers, to smart furniture, to virtuosic musical instruments. Author of the books *When Things Start to Think*, *The Nature of Mathematical Modeling*, and *The Physics of Information Technology*, Dr. Gershenfeld has a B.A. degree in physics with High Honors from Swarthmore College, a Ph.D. from Cornell University, was a Junior Fellow of the Harvard University Society of Fellows, and a member of the research staff at AT&T Bell Laboratories.