Ensemble Quantum Computing by Nuclear Magnetic Resonance Spectroscopy

David G. Cory[1]
Nuclear Engineering, M.I.T., Cambridge, MA 02139
Amr F. Fahmy[2]
Timothy F. Havel[2]
B.C.M.P., Harvard Medical School, Boston, MA 02115

TR-10-96 (December 18, 1996; revised December 30, 1996)

A Technical Report of the
Center for Research in Computing Technology
Harvard University
Cambridge, Massachusetts

A shortened version of this report was accepted for publication in the Proceedings of the National Academy of Sciences by Richard M. Karp on December 16, 1996.

---

1. Supported by NSF / DMR / 9357603.
2. Supported by NSF / MCB / 9527181.

Ensemble Quantum Computing by Nuclear Magnetic Resonance Spectroscopy

David G. Cory (dcory@mit.edu)
Dept. of Nuclear Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139

Amr F. Fahmy (amr@deas.harvard.edu)
Biological Chemistry and Molecular Pharmacology
Harvard Medical School, Boston, MA 02115

Timothy F. Havel (havel@menelaus.med.harvard.edu)
Biological Chemistry and Molecular Pharmacology
Harvard Medical School, Boston, MA 02115

Abstract

A *quantum computer* is a machine which can operate in parallel on all its possible inputs at once, and can be programmed much like a conventional digital computer. It is based on the fact that the number of degrees of freedom in a collection of coupled two-state quantum systems grows exponentially with the number of systems. A major problem with quantum computers is that the amount of information that can be extracted from a quantum system is extremely limited, because the act of measurement irreversibly alters the state of the system. This phenomenon, known as *wave function collapse*, severely restricts the range of problems that can be solved on a quantum computer. We present a new computational model, which differs from a quantum computer only in that the result of measuring an observable is the *expectation value* of the observable, rather than a random eigenvalue thereof. We show that such an *expectation value quantum computer* can solve NP-complete problems in polynomial time, by showing that it can solve the satisfiability problem in linear time. This abstract computational model can be realized, to a certain extent, by nuclear magnetic resonance spectroscopy on macroscopic ensembles of quantum spins, namely molecules in a test tube. This is made possible by identifying a manifold of statistical spin states, called *pseudo-pure states*, whose mathematical description is isomorphic to that of an isolated spin system. The result is a novel *NMR computer* which can be programmed much like a quantum computer, but in other respects more closely resembles a DNA computer. Most notably, when applied to intractable combinatorial problems an NMR computer can use an amount of sample, rather than time, which grows exponentially with the size of the problem. Although NMR computers will be limited by current technology to exhaustive searches over only 15 to 20 bits, searches over as much as 50 bits are in principle possible, and more advanced algorithms could greatly extend the range of applicability of such machines.

1. Introduction

Several physical implementations of computational models other than the standard von Neumann model have recently been proposed, which in principle scale better on certain types of computational problems. Most notably, Adleman has solved a travelling salesman problem by DNA computing (1), and Shor has shown theoretically that a quantum computer should be able to factorize integers in polynomial time (2). The term “NP-complete” refers to a class of computationally intractable problems of particular interest, both because they are “polynomially equivalent” to one another, and because they are encountered in many important applications (3, 4, 5). The travelling salesman problem includes a set of NP-complete instances, which indicates that DNA computing may someday be useful in solving NP-complete problems. In contrast, it is widely believed that a quantum computer cannot solve NP-complete problems in polynomial time (6). Up to now, however, only very small problems have been solved by DNA computing (7), and no one has yet succeeded in building a quantum computer able to handle more than two bits of information (8, 9).

In this paper, we consider another physical mechanism that is capable of computation, namely *nuclear magnetic resonance* (NMR) spectroscopy. This approach is based on the fact that the spins in each molecule of a liquid sample are largely isolated from the spins in all other molecules. As a result, each molecule is effectively an independent quantum computer. Via radio-frequency electromagnetic pulses, it is straightforward to manipulate each of the spins in every molecule identically across the sample, as in a synchronous, single-instruction parallel machine. Similarly, what one observes in such an experiment is the sum of certain observables over all the molecules in the sample, which is proportional to the ensemble-average expectation value of the observable. Thus, such a machine performs a calculation using quantum parallelism at the molecular level, and then amplifies its results to the macroscopic level via a form of classical parallelism. We call a computer based on this principle an *ensemble quantum computer.*

Other researchers have proposed implementing an atomic-scale quantum computer by NMR and analogous physical mechanisms, e.g. (8, 9, 10). The reason a *macroscopic* version has not previously been pursued lies in the fact that the magnetic moments of the individual spins are very small, so that even with superconducting magnets it is not possible to perfectly align the spins at temperatures above ca. one milli-Kelvin. This in turn makes it impossible to prepare a liquid sample in a *pure state*, in which the spin states of all the molecules are identical. We have solved this problem by introducing a new concept into NMR spectroscopy, called *pseudo-pure states*, whose behavior is similar in many respects to a pure state. In particular, they can be described by a type of spinor, which evolves via unitary transformations under the Hamiltonians of NMR, and whose expectation values are easily obtained from the corresponding ensemble-average expectation values. We call these spinors *pseudo-spinors*, to emphasize the fact that their physical interpretation differs from that of the spinors which describe isolated spin systems.

Such an *NMR computer* can be programmed electronically much like a quantum computer, but can be implemented using macroscopic liquid samples at room temperature and pressure like a DNA computer. In principle, it can also trade an exponential increase in the amount of time required to solve intractable combinatorial problems, including NP-complete problems, for an exponential increase in the size (physical dimensions) of the sample on which it operates. Preliminary experiments which validate this novel approach to computing have already been performed (11). Although scaling it to problems beyond the reach of conventional computers promises to be very difficult, the theoretical implications of such a machine are intriguing, and its general principles may prove more broadly applicable. Because the overall approach combines results from several very different fields, an elementary exposition of its main features is needed to provide a common knowledge base for future interdisciplinary research. That is the purpose of the present paper.

2. Implicit Parallelism in Linear Analogue Computers

The purpose of this section is to show that the implicit parallelism which is available in quantum computers via coherent superpositions can actually be achieved by an entirely classical device. Consider an electronic system with a linear response function, 2^n binary inputs, and 1 output. Suppose further that if all the inputs are zero except for the k-th, which is $+1$ volts, then the output is $f(k) \in \{0, 1\}$ volts, where f is a boolean function of the bits in the binary representation of the integer k. By reduction to satisfiability (see section 4), any NP-complete problem can be formulated as proving the existence of a solution to a boolean equation of the form $f(k) = 1$, and hence solved by such a device. We shall call these devices *linear analogue computers*, or LAC's. Each LAC computes just one function f, so building it and programming it are essentially the same thing. The input can be prepared by a binary decoder, which is the only nonlinear component needed.

Being linear, a LAC satisfies the superposition principle. If we let $\mathbf{v}_k$ be the k-th elementary unit vector, we may regard f as a scalar-valued function on a vector space of dimension 2^n, such that $f(\mathbf{v}_k) = f(k)$ for all k. Linearity then implies that any linear combination of the inputs results in an output equal to the same linear combination of all the constituent outputs, i.e.

$$f\left(\sum_{k=0}^{2^n-1} c_k \mathbf{v}_k\right) = \sum_{k=0}^{2^n-1} c_k f(\mathbf{v}_k) \qquad [1]$$

for any scalar coefficients c_k. Therefore, if we set $c_k = 2^{-n}$ for all k, the output will be equal to the probability that a randomly chosen input k is a solution to the equation $f(k) = 1$, and hence nonzero if and only if such a solution exists. In principle, this can be done in constant time, once the LAC in question has been built. Note we have not needed to assume that the vectors are orthogonal!

Of course, the size and complexity, and hence the time needed to build, a LAC increases exponentially with the size n of the problem it is designed to solve. Therefore, this approach does not lead to a polynomial-time solution to NP-complete problems. It has nevertheless been shown that a device with many of the same capabilities can in principle be built using the properties of quantum systems, whose size grows only linearly with the size of the problem it can be applied to. Because of its importance here, we shall present a brief account of how such a *quantum computer* (QC) works in the next section.

3. Quantum Computing and Exponential State Spaces

A distinctive property of multiparticle quantum systems is that the dimensionality of their state space grows exponentially with the number of particles (12). This is because the Hilbert space for a system of distinguishable particles must be taken as the tensor product of the Hilbert spaces of the individual particles, in order to model the correlations among them. In particular, if one considers the spin dynamics of a system of n spin $1/2$ particles, whose Hilbert spaces are all of dimension 2, the dimension of the Hilbert space of the combined system is 2^n. It is for this reason that many researchers have proposed constructing a QC from a system of spin $1/2$ particles, for example in an ion trap (13).

The vectors (or "wave functions") in the Hilbert space of a system of spins are called *spinors*. The standard basis consists of the joint eigenvectors of the total and z-component of the spin angular momentum. The encoding used in quantum computing maps each integer k in the range $[0, 2^n - 1]$ to the k-th basis element $\mathbf{v}_k$ versus a particular ordering of the basis. For a single spin $1/2$ particle, the basis consists of the spin "down" (antiparallel to the z-axis) state, which is represented by the vector $[1, 0]$ and is denoted by the "bra" $\langle 0|$, together with the spin "up" state, which is represented by the vector $[0, 1]$ and is denoted as $\langle 1|$; the corresponding column vectors are denoted by the "kets" $|0\rangle$ and $|1\rangle$, respectively. The basis vectors of an n-spin system are formed by taking the tensor or *Kronecker products* of the basis vectors of its constituent spins in some arbitrary but fixed order. The general definition of the Kronecker product "$\otimes$" of an $M \times N$ matrix $\mathbf{A}$ with an $M' \times N'$ matrix $\mathbf{B}$ is the $MM' \times NN'$ matrix given by

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & \dots & a_{1N}\mathbf{B} \\ \dots & \dots & \dots \\ a_{M1}\mathbf{B} & \dots & a_{MN}\mathbf{B} \end{bmatrix} \qquad \text{where} \qquad \mathbf{A} = \begin{bmatrix} a_{11} & \dots & a_{1N} \\ \dots & \dots & \dots \\ a_{M1} & \dots & a_{MN} \end{bmatrix}. \qquad [2]$$

Thus the basis vectors of a two-spin system are given by:

$$\begin{aligned} |00\rangle &\equiv \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \begin{bmatrix} 1 \\ 0 \end{bmatrix} \equiv \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \qquad |01\rangle \equiv \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \begin{bmatrix} 0 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \\ |10\rangle &\equiv \begin{bmatrix} 0 \\ 1 \end{bmatrix} \otimes \begin{bmatrix} 1 \\ 0 \end{bmatrix} \equiv \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \qquad |11\rangle \equiv \begin{bmatrix} 0 \\ 1 \end{bmatrix} \otimes \begin{bmatrix} 0 \\ 1 \end{bmatrix} \equiv \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \end{aligned}, \qquad [3]$$

A nontrivial complex linear combination of the basis vectors is called a *coherent superposition*.

In a QC, any logical operation on a register in the basis state k must transform it to another basis state l. Because energy dissipation rapidly destroys microscopic order and hence the ability to store information, a QC must be completely isolated from its surroundings. Isolated quantum systems naturally evolve by *unitary* transformations, which are necessarily *reversible*. This prevents one from implementing conventional logic gates such as the AND gate, or indeed any gate whose number of output bits is less than the number of input bits. Fortunately, any gate can be made reversible by copying some of the input bits to the output bits. For example, the quantum XOR operation, which copies the first input bit and overwrites its second bit with the output, has the following truth table:

$$
\begin{aligned}
[1, 0, 0, 0] \equiv \langle 00| &\rightarrow \langle 00| \equiv [1, 0, 0, 0] \\
[0, 1, 0, 0] \equiv \langle 01| &\rightarrow \langle 01| \equiv [0, 1, 0, 0] \\
[0, 0, 1, 0] \equiv \langle 10| &\rightarrow \langle 11| \equiv [0, 0, 0, 1] \\
[0, 0, 0, 1] \equiv \langle 11| &\rightarrow \langle 10| \equiv [0, 0, 1, 0]
\end{aligned}
\qquad [4]
$$

The unitary matrix for this quantum XOR gate (which is sometimes called the "conditional NOT") is simply

$$
\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \qquad [5]
$$

For further discussion and examples, the reader is referred to (9, 10, 14, 15).

It may be seen that a quantum computer essentially implements a *reversible* LAC, but there are some significant differences. First, since the coefficients in a coherent superposition are generally complex, a quantum computer exhibits interference effects not present in the simple LAC described above (unless it used alternating voltages as its inputs!). Second, the input / output relation of a quantum computer is fully programmable, so a single such device can solve many problems. Third, and most importantly, the size of the system needed to achieve a given computation grows only *linearly* with the size of the input.

There are however some serious problems with this approach. The first is that it promises to be very difficult to build a QC of any significant size, because it requires the precise assembly, isolation, control, and measurement of an atomic-scale system. A more fundamental problem lies in the fact that the amount of information that can be extracted from even perfectly precise measurements on quantum systems is extremely limited, because the act of observation irreversibly alters the system.

For example, if one places a two-spin system in the superposition

$$c_0|00\rangle + c_1|01\rangle + c_2|10\rangle + c_3|11\rangle = \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \end{bmatrix}, \qquad [6]$$

and measures the z-component of the spins, one will observe one of the four basis states $k = 0, \ldots, 3$, each with probability $|c_k|^2 = c_k c_k^*$. The system is forced into the observed state by the act of observation, and all subsequent measurements of the same observable will therefore yield the same result. This phenomenon is sometimes called *wave function collapse*. Thus, it is not possible to completely determine the state of the system (i.e. the coefficients c_k) from a finite number of measurements on identically prepared copies of a given quantum system.

4. The Expectation Value Quantum Computer

Let us imagine an abstract computer that functions like a quantum computer in every respect but one: instead of a random eigenvalue, the result of a measurement of an observable on the system is the *expectation value* of that observable. This computational model further assumes that the expectation value can be measured to arbitrarily high precision, and in time proportional to the number of digits in the result. We shall call our abstract computer an *expectation value quantum computer*, or EVQC.

The following result serves as the motivation for much of the rest of this paper: *an expectation value quantum computer is capable of solving NP-complete problems in polynomial time*.

We shall prove this by showing that an EVQC can solve a well-known NP-complete problem, namely the *satisfiability* (SAT) problem (5), in linear time. In this problem, one is given a boolean function $f\colon \{0, 1\}^n \to \{0, 1\}$ in conjunctive normal form. This is a logical conjunction of clauses, where each clause is a disjunction of boolean variables or their negations. The variables are assigned values from the set $\{0, 1\}$, and such an assignment is said to satisfy f if the value of f on this assignment is 1. The SAT problem is then to determine if there exists a satisfying assignment for the given function f.

If the state w of the system is given by the ket $|w\rangle$ versus a fixed basis, and $\mathbf{K}$ is the Hermitian matrix which represents the observable K versus the same basis, elementary quantum mechanics tells us that the expectation value of K is the value of the quadratic form $\langle w|\mathbf{K}|w\rangle$. Given a system of $n+1$ spin $1/2$ particles, we may represent each assignment $x \in \{0, 1\}^n$, together with the state of the $(n+1)$-th output bit $y \in \{0, 1\}$, by an elementary basis vector $|x, y\rangle$ in a Hilbert space of dimension 2^{n+1}. Using results from quantum computing (14), it is possible to convert the boolean operations in f into a sequence of unitary transformations whose product $\mathbf{U}_f$, applied to any single basis state $|x, 0\rangle$, yields $\mathbf{U}_f|x, 0\rangle = |x, f(x)\rangle$. If we now prepare the input in the superposition of all possible inputs given by

$$\sum_{x \in \{0, 1\}^n} c_{(x, 0)}\, |x, 0\rangle \ , \qquad [7]$$

where $c_{(x,\,0)} = 2^{-n/2}$ and $c_{(x,\,1)} = 0$ for all x, then the result of performing the computation on this superposition is

$$\begin{aligned} \mathbf{U}_f \sum_{x \in \{0,1\}^n} c_{(x,\,0)} |x, 0\rangle &= \sum_{x \in \{0,1\}^n} c_{(x,\,0)} \, |x, f(x)\rangle \\ &\equiv \sum_{x \in \{0,1\}^n} d_{(x,\,f(x))} \, |x, f(x)\rangle \equiv z_f \ . \end{aligned} \qquad [8]$$

Now consider the observable S whose Hermitian matrix $\mathbf{S}$ is

$$\mathbf{S} = \mathbf{1} \otimes \mathbf{1} \otimes \ldots \otimes \mathbf{1} \otimes |1\rangle\langle 1| = \mathbf{Diag}(0, 1, 0, 1, \ldots, 0, 1) \ , \qquad [9]$$

where $\mathbf{1}$ is the 2×2 identity matrix and $\mathbf{Diag}$ is the diagonal matrix of its arguments. Since each nonzero output coefficient $d_{(x,\,1)} = c_{(x,\,0)} = 2^{-n/2}$, its expectation value on the above superposition is

$$\begin{aligned} \langle z_f | \mathbf{S} | z_f \rangle &= \sum_x \sum_{x'} c_{(x,\,0)} c^*_{(x',\,0)} \langle x, f(x) | \mathbf{S} | x', f(x') \rangle \\ &= \sum_x \sum_{x'} c_{(x,\,0)} c^*_{(x',\,0)} \langle x | x' \rangle \langle f(x) | 1 \rangle \langle 1 | f(x') \rangle \\ &= \sum_x \left| c_{(x,\,0)} \right|^2 |\langle f(x) | 1 \rangle|^2 = \sum_x \left| d_{(x,\,1)} \right|^2 \\ &= 2^{-n} \left| \{ x \in \{0,1\}^n | \, f(x) = 1 \} \right| \end{aligned} \qquad [10]$$

This is just the fraction of all assignments which satisfy the given function f. In the case $n = 2$, we have explicitly

$$\mathbf{S} = \begin{array}{c c} \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} & \begin{array}{c} |000\rangle \\ |001\rangle \\ |010\rangle \\ |011\rangle \\ |100\rangle \\ |101\rangle \\ |110\rangle \\ |111\rangle \end{array} \end{array} , \qquad [11]$$

so that $\langle z_f | \mathbf{S} | z_f \rangle = |d_{001}|^2 + |d_{011}|^2 + |d_{101}|^2 + |d_{111}|^2$.

The time required to program $\mathbf{U}_f$, to prepare the superposition, and (in our computational model) to measure the expectation value of S to the requisite precision of 2^{-n}, are all $O(n)$, thus showing that an EVQC can solve SAT in linear time. Because all NP-complete problems are polynomially equivalent by Cook's theorem (3), this proves that an EVQC can solve NP-complete problems in polynomial time. Since our EVQC essentially counts the number of satisfying assignments, we have further shown that it is capable of solving #P-complete problems in polynomial time (5).

It should be noted that the operator $S \equiv S_{n+1}$, whose expectation value tells us the fraction of states in which the $(n+1)$-th spin is 1, can easily be altered to an operator S_k whose expectation value provides the same information about the k-th spin. Thus given any computable function $g : \{0,1\}^m \to \{0,1\}^n$ and a unitary transformation $\mathbf{U}_g$ such that $\mathbf{U}_g |x, 0\rangle = |x, g(x)\rangle$, we can find the value of $g(x)$ on any single input state $x \in \{0,1\}^m$ by simply measuring the expectation values of the n operators S_{m+1} through S_{m+n}. The implication is that any Turing computable function can be computed by an EVQC, and hence we can use implicit parallelism in any such computation by operating on a superposition.

5. Ensemble Quantum Computing and NMR Spectroscopy

An EVQC cannot be realized by any known finite physical system, but it can be approximated by an *ensemble* (large collection) of independent and identical QC's, over which one can measure the sum of the observable evaluated on each QC in the ensemble. Such an *ensemble quantum computer*, or EQC, performs a calculation using quantum parallelism, and then uses a form of classical parallelism to estimate an expectation value. Although an EQC differs from a single QC in significant respects, it also depends on quantum parallelism in an essential way, and cannot be equated with an ensemble of conventional computers. In particular, each QC in the ensemble operates on the *same* superposition over all possible inputs, rather than different (or random) inputs as in an ordinary massively parallel computer. Even more importantly, whereas conventional computers do not appear to occur naturally on the molecular scale, quantum computers do! In this section, we describe how this fact makes it possible to construct an EQC with $\sim 10^{23}$ QC's in it, using nuclear magnetic resonance (NMR) spectroscopy. In the next section, we shall consider the physical limits on what such an *NMR computer* can do.

All atomic nuclei with intrinsic spin behave like tiny magnets (16). Therefore, when they are placed in a magnetic field, their magnetic dipoles tend to align themselves parallel to the field. The Hamiltonian which gives the energy difference between the parallel and antiparallel states is called the *Zeeman Hamiltonian*. Even with superconducting magnets, this energy difference is very small compared to the thermal energy $k_B T$. Hence when a macroscopic ensemble of identical spins is placed in a magnetic field, the net alignment of their magnetic moments with the field is likewise very small at room temperature. If one could reduce the temperature sufficiently close to absolute zero, however, near perfect alignments would be obtained. Under these conditions, all the spins are in the same quantum state, which is described by saying that the ensemble as a whole is in a *pure state*.

In such a system, the magnetization of the sample is the sum of the magnetizations due to the individual spins. Each of the microscopic observations is a random variable, but when one takes the sum over a macroscopic number of observations, the result is a *deterministic* quantity equal to the expectation value of the observable times the number of spins. If all the spins in the ensemble were identical, we would only have a one-bit EQC. If all the spins were distinguishable, on the other hand, the result would be a massive QC, but without the ability to measure well-defined expectation values. What is needed is a partition of the spins into distinguishable equivalence classes of indistinguishable spins. Fortunately, such a partitioning occurs very naturally in chemistry, where the spin-active nuclei in a molecule are distinguished by their electronic environments, but in a macroscopic sample of identical molecules each type of spin is indistinguishable between molecules.

A standard technique in NMR spectroscopy is to use pulses of radio-frequency (RF) radiation to transform the state of the spins by unitary transformations. Since the distinguishable spins in a molecule generally have distinct resonance frequencies, the frequency range of these pulses can be made selective for single spins. For example, a selective pulse of sufficient duration and intensity to rotate the net magnetization of the k-th spin by $\pi/2$ and in-phase with the imaginary component of the carrier corresponds to the unitary matrix

$$\mathbf{1} \otimes \ldots \otimes \mathbf{1} \otimes \mathbf{U}_{\pi/2} \otimes \mathbf{1} \otimes \ldots \otimes \mathbf{1} \qquad \left(\mathbf{U}_{\pi/2} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \right), \qquad [12]$$

where the matrix $\mathbf{U}_{\pi/2}$ occurs as the k-th factor of the Kronecker product.

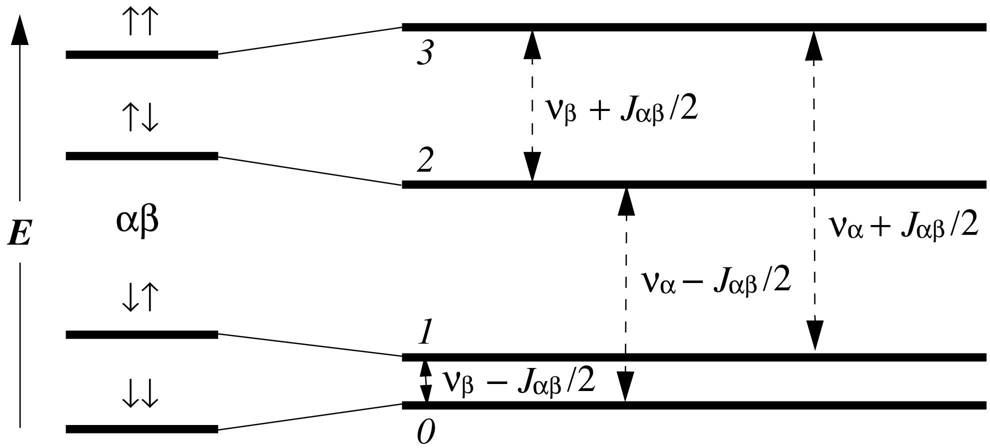

Figure 1. The four energy levels associated with two spins α and β, whose resonance frequencies are ν_α and ν_β Hertz. The energy levels when there is no coupling between the spins are shown on the left, and those with a coupling of $J_{\alpha\beta}$ Hertz on the right. The allowed transitions between the energy levels are indicated with two-headed dashed arrows.

For a two-spin system, we have the energy level diagram shown in Fig. 1. The four dashed two-headed arrows are the transitions allowed by the selection rules for angular momenta, and connect pairs of states that differ by only single spin flips (i.e. whose Hamming distance is one). The spins are generally also coupled to one another, either by through-space dipole-dipole interactions, or by a through-bond effect called *scalar coupling*. This coupling causes the energy differences associated with the various transitions to be generically distinct. Tipping both spins into the xy-plane with a nonselective $\pi/2$ pulse $\mathbf{U}_{\pi/2} \otimes \mathbf{U}_{\pi/2}$ causes them to precess in phase, generating a detectable (macroscopic) rotating magnetic moment. The real part of the Fourier transform of the resulting signal gives an

NMR spectrum with two pairs of peaks, as shown in Fig. 2. The intensity of each pair is proportional to the population difference between states in which the corresponding spin is "up" and those in which it is "down".

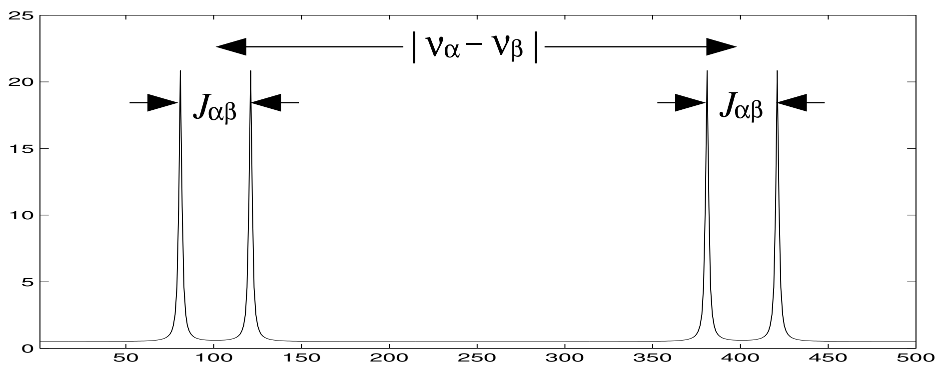

Figure 2. The simulated NMR spectrum of a two-spin system, with $\nu_\alpha = 100$, $\nu_\beta = 400$ and $J_{\alpha\beta} = 40$ Hertz. From left to right, the four peaks in this spectrum correspond to the transitions $0 \leftrightarrow 1$, $2 \leftrightarrow 3$, $0 \leftrightarrow 2$ and $1 \leftrightarrow 3$ indicated in Fig. 1 (be forewarned, however, that NMR spectroscopists traditionally plot their spectra with frequency increasing from right to left). In practice, this spectrum would be obtained by applying a nonselective $\pi/2$ pulse to the equilibrium state and Fourier transforming the resulting signal. The peak intensities are proportional to the population differences between pairs of states, and their widths to the inverse decoherence time (known as T_2 in NMR).

The quantum XOR gate introduced above flips one of these spins, given that the other spin is "up", which in turn corresponds to the transitions $1 - 3$ and $2 - 3$ for an output on the first and second spins, respectively. Thus, one can implement the quantum XOR gate by a single RF pulse whose frequency range spans only the peak of one of these transitions, and which is of sufficient length and intensity to invert the populations of the corresponding pair of states. This is an example of a *Pound-Overhauser double resonance* experiment (16). More generally, we have shown that one can implement the universal Toffoli gate via NMR pulse sequences, and hence effect any desired unitary transformation by these means (11). In particular, the unitary transformation $\mathbf{U}_f$ needed to compute the boolean function f in SAT can be performed by an appropriate pulse sequence.

The expectation value of the operator S in the SAT algorithm above corresponds to the sum of the intensities of the 2^n peaks that are obtained by flipping the $(n+1)$-th spin. For a two-spin system, the matrix $\mathbf{S} = \mathbf{S}_2$ can be written in terms of the operators

$$\mathbf{P}_{01} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{P}_{23} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \qquad [13]$$

whose expectation values give the population differences between the subscript states (and hence the intensities of the corresponding peaks), as

$$\mathbf{S}_2 = (\mathbf{1} \otimes \mathbf{1} - \mathbf{P}_{01} - \mathbf{P}_{23})/2 = (\mathbf{1} \otimes \mathbf{1} - \mathbf{1} \otimes \mathbf{I}_z)/2 , \qquad [14]$$

where

$$\mathbf{I}_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \qquad [15]$$

is the matrix of the operator for the z-component of the spin. Finally, a superposition in which all the basis states are equiprobable is easily obtained by applying a

$\pi/2$ pulse whose frequency range spans all the peaks due to flips of the input spins, but which misses those peaks due to the output spin. Thus, the SAT problem is easily and naturally solved by NMR spectroscopy ... at least in small instances.

6. The Physical Limitations of NMR Computing

Although straightforward in principle, there are several theoretical and practical considerations which limit the sizes of problems that can be solved by an NMR implementation of an EQC. In this section we shall consider the most important such obstacles and, where possible, establish bounds on the corresponding limits. In the following, we denote the number of spins in a molecule by n, the number of molecules present in the sample by M, and the total number of spins by $N = nM$.

The mean-square error in an estimated expectation value is σ^2/M, where the variance σ^2 depends on the problem instance. In order to reliably solve e.g. a SAT problem on n variables, this error must be less than 2^{-2n}, but the resulting bound on n depends on σ^2. To avoid this, we instead require only that at least one molecule samples at least one solution with at least some fixed probability $p_{\min}$. This probability is given by $p = 1 - e^{-r}$, where $r = M/2^n$ is the average number of molecules per state. The result turns out to depend on $p_{\min}$ only through the term $\log(\log((1 - p_{\min})^{-1}))$, and hence we shall simply set $p_{\min} = 1 - e^{-1}$, i.e. $r > 1$. We typically have $N \sim 10^{23} \approx 2^{76}$ in an NMR sample, so if we wish to solve a problem of size n, we would have $M \sim 2^{76}/n$, or $r \sim 2^{76}/(n2^n)$. This leads to the approximate bound:

$$(n2^n < 2^{76}) \quad \Rightarrow \quad (\log_2(n) + n < 76) \quad \Rightarrow \quad (n < 70) \quad . \qquad [16]$$

This bound on the number of bits over which we can perform a reliable search for a single desired state can be increased by seven or eight by time-averaging the signals, and perhaps as much as ten by using very large samples. Because the number of bits that can be handled depends logarithmically on the number of repetitions and/or size, the absolute limit on what any practical NMR computer can handle remains well below 100 bits. In practice, the finite signal-to-noise ratio of the receiver will reduce these estimates substantially. These estimates, however, apply only to solving combinatorial problems by exhaustive search. With the use of more sophisticated search procedures, it may be possible to reduce the rate of growth in the sample size required and so go considerably larger. It is also conceivable that algorithms for some problems exist which can be implemented more efficiently (in terms of time) on an NMR computer than they can be on conventional or quantum computers, using only a fixed sample size.

Several other obstacles to scaling an NMR computer to large problems also exist, but it is more difficult to establish hard bounds on the size of the problem that these obstacles permit us to solve. For example, it is difficult to find molecules with more than about ten spins in them and with a large coupling constant between every pair of spins (particularly if one is relying upon only scalar coupling). In addition, the number of peaks that can be independently resolved in the NMR spectrum of a molecule is limited by both the separation as well as the width of the peaks. Nevertheless, we need only find one suitable molecule (out of the myriads available in chemistry) in order to obtain a general EQC on that number of spins. As a result, we do not expect this obstacle to be the ultimate bottleneck.

The width of the peaks is also important because it is equal to the inverse *spin-spin relaxation time* T_2, otherwise known as the decoherence time in quantum computing. Since decoherence introduces errors into the intermediate results, it limits the time available for computation. The difficulty of isolating microscopic systems from their environment well enough to attain long decoherence times has proven to be one of the chief obstacles to implementing a true QC. In contrast, the nuclear spins in a molecule are generally quite well isolated from its motional and electronic degrees of freedom.

In solids, the dominant mechanism of spin-spin relaxation is the dipole-dipole interaction, which typically results in decoherence times of a small fraction of a millisecond. In liquids, on the other hand, the dipole-dipole interaction is averaged to zero by the rapid rotational diffusion of nearby molecules, leading to decoherence times that can be of order many seconds. The intramolecular scalar couplings between spins, however, are not averaged to zero, and can exceed one hundred Hertz. The number of selective pulses that can be used in one experiment, which is of order the ratio of the coupling constants to the decoherence rate, can therefore be over a thousand. It is possible that error correcting schemes like those used to control decoherence and other errors in quantum computation (see e.g. (17)) can be developed for NMR computing, further increasing the time available for computation.

More importantly, this dipolar averaging effect largely isolates the spins in different molecules from one another, thus making each molecule into an independent QC. It follows that liquid NMR samples are an extremely suitable means of implementing an EQC. Unfortunately, the small energies associated with the Zeeman Hamiltonian make it impossible to prepare a pure state at temperatures above about 10^{-3} Kelvin, but molecules do not remain in the liquid state as one approaches the absolute zero of temperature. For this reason we have developed a novel means of simulating pure states at room temperature.

7. Pseudo-Pure States

A *mixed state* is a statistical mixture of independent quantum systems which are not all in the same quantum state. The mixed states of spin systems are described by a generalization of spinors, known as the *density matrix*, which enables us to compute their statistical properties (18). The density matrix Ψ of a pure state is obtained from the corresponding spinor ψ simply by taking a dyadic product, i.e. $\Psi = |\psi\rangle\langle\psi|$. Such a density matrix is necessarily an idempotent projection operator whose only nonzero eigenvalue has the value one. In the case of a two-spin system in the state $|00\rangle$, for example, we have

$$\Psi = |00\rangle\langle 00| = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} . \qquad [17]$$

The density matrix of a mixed state is obtained by taking the average of the density matrix of the spin system over an ensemble of pure states, i.e.

$$\Psi = \int_{\{\psi\}} p(\psi)\, |\psi\rangle\langle\psi|\, d\psi = \overline{|\psi\rangle\langle\psi|}\ , \qquad [18]$$

where $p(\psi)$ is the probability density of the pure state described by the spinor ψ and $\{\psi\}$ denotes the set of all unit norm spinors. Such a density matrix can be an arbitrary positive semi-definite Hermitian matrix normalized to unit trace. The diagonal elements are the relative *populations* of the various energy levels, while the off-diagonal elements represent *coherences*, i.e. correlations in the phases of precessing spins in pairs of energy levels across the sample.

The density matrix evolves in time according to the *Liouville-von Neumann equation*

$$\frac{\mathrm{d}\boldsymbol{\Psi}}{\mathrm{d}t} = i\,[\boldsymbol{\Psi}, \mathbf{H}] = i\,(\boldsymbol{\Psi}\mathbf{H} - \mathbf{H}\boldsymbol{\Psi}) \ , \qquad [19]$$

where $\mathbf{H}$ is a matrix representation of the spin Hamiltonian of the system, and $[\boldsymbol{\Psi}, \mathbf{H}]$ denotes the matrix commutator. This has the general solution

$$\boldsymbol{\Psi}(t) = \mathbf{U}^{\dagger}(t)\boldsymbol{\Psi}(0)\mathbf{U}(t), \qquad [20]$$

where $\mathbf{U}(t)$ is a time-dependent unitary matrix. For a time-independent Hamiltonian, this matrix has the form $\mathbf{U}(t) = \mathbf{Exp}(it\mathbf{H})$. Finally, the *ensemble average* of the expectation value of any observable K is obtained as the trace product $\mathrm{tr}(\mathbf{K}\boldsymbol{\Psi})$ of the corresponding matrix $\mathbf{K}$ with the density matrix.

Eq. [18] is usually regarded as a Gibbs ensemble average, which is a purely thought construction used to compute time averages. As described above, however, the spins in the M different molecules of a liquid are essentially independent of one another. Hence the total density matrix $\tilde{\boldsymbol{\Psi}}$ factorizes into a Kronecker product of M *reduced* density matrices, each of size $2^n \times 2^n$, as follows:

$$\tilde{\Psi} = \overline{|\psi_1 \ldots \psi_M\rangle\langle\psi_1 \ldots \psi_M|} = \overline{(|\psi_1\rangle \otimes \ldots \otimes |\psi_M\rangle)(\langle\psi_1| \otimes \ldots \otimes \langle\psi_M|)}$$
$$= \overline{|\psi_1\rangle\langle\psi_1| \otimes \ldots \otimes |\psi_M\rangle\langle\psi_M|} = \overline{|\psi_1\rangle\langle\psi_1|} \otimes \ldots \otimes \overline{|\psi_M\rangle\langle\psi_M|} \qquad [21]$$
$$= \Psi_1 \otimes \ldots \otimes \Psi_M \qquad \text{where} \qquad \Psi_k \equiv \overline{|\psi_k\rangle\langle\psi_k|} \, .$$

Moreover, since all the molecules are identical, all the reduced density matrices $\Psi_k = \Psi$ are equal.

A *coherent* excitation of the system acts identically on all the molecules at once by definition, and hence likewise corresponds to a Kronecker product of identical unitary matrices:

$$(\mathbf{U} \otimes \ldots \otimes \mathbf{U})^{\dagger}(\Psi \otimes \ldots \otimes \Psi)(\mathbf{U} \otimes \ldots \otimes \mathbf{U})$$
$$= (\mathbf{U}^{\dagger}\Psi\mathbf{U}) \otimes \ldots \otimes (\mathbf{U}^{\dagger}\Psi\mathbf{U}) \qquad [22]$$

In addition, the ensemble average of any macroscopic observable that is a sum of identical quantum observables K over all the molecules has a matrix of the form $\sum(\mathbf{1} \otimes \ldots \otimes \mathbf{K} \otimes \ldots \otimes \mathbf{1})$ (sum over $\mathbf{K}$ in each place), and so has the value:

$$\mathrm{tr}[\textstyle\sum(\mathbf{1} \otimes \ldots \otimes \mathbf{K} \otimes \ldots \otimes \mathbf{1})(\Psi \otimes \ldots \otimes \Psi)]$$
$$= \textstyle\sum[\mathrm{tr}(\Psi) \ldots \mathrm{tr}(\mathbf{K}\,\Psi) \ldots \mathrm{tr}(\Psi)] = M\,\mathrm{tr}(\mathbf{K}\,\Psi) \qquad [23]$$

Since the Hamiltonian is such an observable, and since (19)

$$\mathbf{Exp}(\,it \textstyle\sum(\mathbf{1} \otimes \ldots \otimes \mathbf{H} \otimes \ldots \otimes \mathbf{1})) = \mathbf{Exp}(it\mathbf{H}) \otimes \ldots \otimes \mathbf{Exp}(it\mathbf{H}) \;, \qquad [24]$$

it follows that the reduced density matrix also satisfies the Liouville-von Neumann equation. The net result is that a large number of copies of a single type of molecule in a liquid sample constitutes a excellent physical approximation to a Gibbs ensemble, at least for the spin degrees of freedom.

We define a *pseudo-pure state* to be one whose density matrix can be shifted by adding an appropriate multiple of the unit matrix to it so as to obtain a scalar multiple of the density matrix of a pure state. To obtain a physically intuitive expression for such a density matrix, suppose the Hamiltonian has the form $\mathbf{H} = \gamma H_0[|\psi\rangle\langle\psi| - \hat{\mathbf{1}}]$, where $\hat{\mathbf{1}} \equiv \mathbf{1}/2^n$, ψ is a unit spinor, γ is the gyromagnetic ratio, and H_0 is the magnetic field strength. If β is the inverse of the thermal energy in units of Hertz, the corresponding equilibrium density matrix can be written in closed form as (19)

$$\Psi = \frac{\mathbf{Exp}(-\beta\gamma H_0[|\psi\rangle\langle\psi| - \hat{\mathbf{1}}])}{\mathrm{tr}[\mathbf{Exp}(-\beta\gamma H_0[|\psi\rangle\langle\psi| - \hat{\mathbf{1}}])]} = \frac{\mathbf{1} + (e^{-\beta\gamma H_0} - 1)|\psi\rangle\langle\psi|}{2^n + e^{-\beta\gamma H_0} - 1}. \qquad [25]$$

If we define the *polarization* $\alpha = (e^{-\beta\gamma H_0} - 1)/(e^{-\beta\gamma H_0} + 1)$, this becomes

$$\Psi = \frac{(1-\alpha)\mathbf{1} + 2\alpha|\psi\rangle\langle\psi|}{(1-\alpha)2^n + 2\alpha} \qquad (-1 \le \alpha \le 1), \qquad [26]$$

which we will also take as our general definition of a pseudo-pure state, whether or not it is at equilibrium. In other words, the density matrix of a pseudo-pure state is an arbitrary unit trace matrix of size $2^n \times 2^n$ whose eigenvalues lie in the interval $[0, 1]$, with $2^n - 1$ of them degenerate. The spinor ψ in Eq. [26] will be called a *pseudo-spinor*, to emphasize the fact that its physical interpretation differs from that of the spinor for an isolated spin system.

Pseudo-pure states promise to be of great utility in implementing an EQC for the following reasons:

1. Like true pure states, each pseudo-spinor determines a unique pseudo-pure density matrix via Eq. [26] above, and each pseudo-pure density matrix determines a pseudo-spinor that is unique up to an overall phase factor (assuming that α is known).
2. Like true pure states, when the density matrix of a pseudo-pure state (as in Eq. [26]) is transformed by a given unitary transformation (as in Eq. [20]), the corresponding pseudo-spinor is transformed by the *same* unitary matrix, since $\mathbf{U}\Psi\mathbf{U}^\dagger \propto (1-\alpha)\mathbf{1} + 2\alpha(\mathbf{U}|\psi\rangle)(\mathbf{U}|\psi\rangle)^\dagger$.
3. Like true pure states, the ensemble average expectation value of an observable K versus a pseudo-pure density matrix (as in Eq. [26]) yields the *ordinary* expectation value of the observable versus the corresponding pseudo-spinor, since $\mathrm{tr}(\mathbf{K}\Psi) \propto (1-\alpha)\mathrm{tr}(\mathbf{K}) + 2\alpha\langle\psi|\mathbf{K}|\psi\rangle$ where $\mathrm{tr}(\mathbf{K})$ is a known constant.
4. Unlike true pure states, one can prepare pseudo-pure states in an NMR sample at room temperature, for example by an appropriate sequence of RF and gradient pulses applied to the equilibrium state (11).

The molecules of a sample in a pseudo-pure state constitute a statistical mixture of quantum states. Nevertheless, there is a preponderance of a particular state present, which manifests itself when we add up the magnetizations of all the molecules, in effect making each molecule cast its “vote” for the final spectrum. The net result is that we can emulate a QC by NMR spectroscopy on macroscopic liquid samples in open test tubes at room temperature and pressure. Moreover, the number of operations required to do this is identical to the number of operations executed by the QC! The only real difference is that one can determine the state of the system in terms of its expectation values without wave function collapse which, as we have seen, is actually a significant advantage for most computational purposes. For an experimental demonstration of the fact that one can apply quantum logic gates to coherent superpositions of pseudo-spinors, the reader is referred to Ref. (11).

8. Coherent versus Incoherent Superpositions

The fact that it is possible to efficiently emulate a QC with an EQC may seem surprising, since a QC derives its power from wave function interference and collapse. In order to better understand how it can nevertheless be done, note that the set of $2^n \times 2^n$ Hermitian matrices can be regarded as a 2^{2n}-dimensional vector space over the *real* numbers, whose scalar multiplication and addition are given by the corresponding operations on the matrices. The trace of the product of two matrices defines a Euclidean inner product on this space. The positive semi-definite Hermitian matrices form a *convex cone*, whose *extremal rays* are generated by the rank one density matrices of pure states.

The space of all Hermitian matrices inherits a (nonorthogonal) basis of pure states from the standard spinor basis. For a two-spin system, this standard basis is:

$$
\begin{array}{cc}
|00\rangle\langle 00| \quad |01\rangle\langle 01| & |10\rangle\langle 10| \quad |11\rangle\langle 11| \\
(|00\rangle + |01\rangle)(\langle 00| + \langle 01|) & (|00\rangle - i\,|01\rangle)(\langle 00| + i\,\langle 01|) \\
(|00\rangle + |10\rangle)(\langle 00| + \langle 10|) & (|00\rangle - i\,|10\rangle)(\langle 00| + i\,\langle 10|) \\
(|00\rangle + |11\rangle)(\langle 00| + \langle 11|) & (|00\rangle - i\,|11\rangle)(\langle 00| + i\,\langle 11|) \\
(|01\rangle + |10\rangle)(\langle 01| + \langle 10|) & (|01\rangle - i\,|10\rangle)(\langle 01| + i\,\langle 10|) \\
(|01\rangle + |11\rangle)(\langle 01| + \langle 11|) & (|01\rangle - i\,|11\rangle)(\langle 01| + i\,\langle 11|) \\
(|10\rangle + |11\rangle)(\langle 10| + \langle 11|) & (|10\rangle - i\,|11\rangle)(\langle 10| + i\,\langle 11|)
\end{array}
\qquad [27]
$$

We shall call real linear combinations of these standard basis matrices *incoherent superpositions*, to contrast them with coherent superpositions of spinors. Observe that coherent superpositions of spinors give rise, via the dyadic product, to incoherent superpositions of density matrices versus this basis of pure states.

The extremal rays of the cone are an algebraic surface in this space, which is homogeneously parametrized by spinors via the quadratic mapping given by the dyadic product. The trace one matrices constitute an affine hyperplane perpendicular to the identity matrix, which may be regarded as the center of the cone. The density matrices of pseudo-pure states therefore lie on line segments which connect the extremal rays to the center of the cone in this affine hyperplane. The two-sided unitary transformations act transitively on the set of extremal rays, and hence also on the points of the line segments given by a fixed polarization α. A general pseudo-pure state can be uniquely expanded as a linear combination of pseudo-pure states with the same polarization, one on each of the line segments ending in one of the standard basis states. The coefficients in such a linear combination are not independent of one another, but are quadratic functions of the standard spinor coordinates. Nevertheless, the Liouville-von Neumann equation is first-order in time just like the Schrödinger equation, and hence the superposition principle holds. As we saw by studying the LAC earlier this, together with a state space whose dimension increases exponentially with the number of bits, is all that is needed in order to obtain "quantum parallelism".

The existence of an EQC based on NMR shows that quantum parallelism is logically and physically distinct from the probabilistic aspects of quantum computing, and in particular from wave function collapse. Wave function collapse nevertheless provides a powerful filtering and amplification mechanism in quantum computing, which is not available in ensemble quantum computing. This is the reason why an NMR computer may require an exponentially increasing sample size to efficiently solve intractable problems. On the other hand, wave function collapse is by its very nature uncontrollable, and hence considerable ingenuity is required to develop algorithms which can take advantage of it. Thus far, no one has found an efficient QC algorithm for NP-complete problems, which we have shown is in principle straightforward on an EQC. Via pseudo-pure states, one can also implement a small EQC much more easily than a true QC of equal size.

The price one pays for using pseudo-pure states in NMR samples at room temperature is the loss of a factor of about one million in the effective number of molecules per state, since even with superconducting magnets the net polarization of the spins is only about one part in a million. Since this factor figures only logarithmically in the physical limits (Eq. [16]), we can in principle still solve SAT problems on up to about 50 bits with an NMR computer that operates on pseudo-pure states at room temperatures. Of course, a number of practical considerations will limit us to substantially smaller problems, at least using the presently available materials and technologies. Prominent among these is the signal-to-noise ratio that can be attained with conventional electronics at room temperatures, which will probably make exhaustive searches over all inputs on more than about 15 to 20 bits impractical in the foreseeable future. This should nevertheless be sufficient to motivate considerable research into alternative materials, technologies and perhaps algorithms that will enable us to push these limits back yet further.

9. Conclusions

Researchers today are divided concerning whether or not it will ever be possible to build a quantum computer of any significant size. We have explained how a macroscopic analogue of a quantum computer can be implemented today, using commercially available NMR spectrometers and ordinary liquids at room temperature and pressure. Because it uses the parallelism inherent in macroscopic ensembles instead of the filtering and amplification concomitant upon wave function collapse, an NMR computer may require an exponentially increasing sample size to obtain a detectable signal. This in turn will limit the size of the problems that can be reliably solved by NMR. The experimental issues involved in NMR computing, and in particular general methods for the preparation of pseudo-pure states, will be dealt with in subsequent publications (see also (11)).

Despite its acknowledged limitations, it is possible that an NMR computer (or possibly some other type of EQC) will someday be built that can (with suitable algorithms) solve problems beyond the reach of conventional computers. Besides NP-complete problems, an NMR computer should be able to factorize integers by either a direct search procedure, or by procedures based on Miller's approach which scale better (2, 10). Another likely application would be to simulate the statistical behavior of open quantum systems, as originally proposed for closed quantum systems using quantum computers by R. P. Feynman (see also (20)). These theoretical questions will likewise be considered in subsequent publications.

Acknowledgments

Thanks are due to G. Wagner (B.C.M.P., Harvard Medical School) and R. Griffin (Francis Bitter Magnet Lab, M.I.T.) for their support and encouragement. This work was funded by N.S.F. grant MCB 9527181 to T.F.H. and N.S.F. grant DMR 9357603 to D.G.C. Several of the key ideas in this paper were present in the extended abstract submitted by the same authors on May 12, 1996 to the Fourth Workshop in Physics and Computation, and subsequently presented at Boston University on November 24, 1996.

References

1. Adleman, L. (1994) “Molecular computation of solutions to combinatorial problems”, *Science* **266**, 1021-4.

2. Shor, P. W. (1994) “Algorithms for quantum computation: discrete logarithm and factoring”, *Proc. 35th Symp. Found. Computer Sci.* (IEEE Press), 124-3.

3. Cook, S. A. (1971) “The complexity of theorem-proving procedures”, *Proc. 3rd Ann. ACM Symp. on Theory of Computing* (Assoc. for Computing Machinery, New York, NY), 151-158.

4. Karp, R. M. (1972) “Reducibility among combinatorial problems”, in R. E. Miller and J. W. Thatcher (eds.), *Complexity of Computer Computations* (Plenum Press, New York, NY), 85-103.

5. Garey, M. R. and Johnson, D. S. (1979) *Computers and Intractability* (W. H. Freeman, San Francisco).

6. Bennett, C. H., Bernstein, E., Brassard, G. & Vazirani, U. (1996) “Strengths and weaknesses of quantum computing”, *SIAM J. Comput.*, to appear.

7. Lipton, R. J. (1995) “DNA solution of hard computational problems”, *Science* **268**, 542-5.

8. Lloyd, S. (1993) “A potentially realizable quantum computer”, *Science* **261**, 1569-71 (see also commentary by Lloyd, S. (1994) “Envisioning a quantum supercomputer”, *Science* **263**, 695-5).

9. DiVincenzo, D. P. (1995) “Quantum computation”, *Science* **270**, 255-61.

10. Ekert, A. and Jozsa, R. (1996) “Shor’s quantum algorithm for factorizing numbers”, *Rev. Mod. Phys.* **68**, 733-53.

11. Cory, D. G., Fahmy, A. and Havel, T. F. (1996) Extended abstract, presented

at the Fourth Workshop on Physics and Computation (Nov. 22-4, 1996), and published by the New England Complex Systems Institute (http://dynamics.bu.edu/InterJournal/papers/[33]_113096200052_.html).

12. Cohen-Tannoudji, C., Diu, B. and Laloë, F. (1977) *Quantum Mechanics* (J. Wiley & Sons, New York, NY), vols. 1 & 2.

13. Cirac, J. I. and Zoller, P. (1995) "Quantum computation with cold trapped ions", *Phys. Rev. Lett.* **74**, 4091-4.

14. Deutsch, D. (1989) "Quantum computational networks", *Proc. R. Soc. Lond.* A **425**, 73-90.

15. Brassard, G. (1995) "A quantum jump in computer science", in van Leeuwen, J. (ed.), *Computer Science Today* (Lect. Notes Comput. Sci. 1000, Springer-Verlag), 1-14.

16. Slichter, C. P. (1990) *Principles of Magnetic Resonance* (3rd ed., Springer-Verlag).

17. P. W. Shor, (1996) "Fault-tolerant quantum computation", in 37th Annual Symposium on the Foundations of Computer Science, Burlington, Vermont, 14-16 October 1996.

18. Blum, K. (1981) *Density Matrix Theory and Applications* (Plenum Press, New York, NY).

19. Najfeld, I. and Havel, T. F. (1994) "Derivatives of the matrix exponential and their computation", *Aiken Labs* TR-33-94 (partially published in *Adv. Appl. Math.*, **16**, 321-75, 1995).

20. Lloyd, S. (1996) "Universal quantum simulators", *Science* **273**, 1073-8.