

Liquid State NMR Quantum Computing

Lieven M. K. Vandersypen, Costantino S. Yannoni, and Isaac L. Chuang

Volume 9, pp 687–697

in

Encyclopedia of Nuclear Magnetic Resonance

Volume 9: Advances in NMR

(ISBN 0471 49082 2)

Edited by

David M. Grant and Robin K. Harris

© John Wiley & Sons, Ltd, Chichester, 2002

Liquid State NMR Quantum Computing

Lieven M. K. Vandersypen, Costantino S. Yannoni, and
Isaac L. Chuang

Department of Applied Physics, Selft University of Ledmology, Selft,
The Netherlands

1	Introduction	1
2	Quantum Computation	1
3	NMR Quantum Computers	4
4	Summary and Conclusions	9
5	References	9

1 INTRODUCTION

Since its invention, NMR spectroscopy has developed from a technique for studying physical phenomena such as magnetism into a tool for acquiring information about molecules in chemistry and biology. Furthermore, it was pointed out early on in 1955, almost as an anecdote, that nuclear spins could also be used for storing information using spin echoes.¹

This insight beautifully illustrated a notion that was developed in a very different context: information is physical and cannot exist without a physical representation.² In recent decades, the relationship between physics and information has been revisited from a new perspective: could the laws of physics play a role in *how* information is processed? The answer appears to be yes. If information is represented by systems such as nuclear spins governed by the laws of quantum mechanics, an entirely new way of doing computation, quantum computation (QC), becomes possible. Quantum computing is not just different or new; it offers an extraordinary promise, the capability of solving certain problems which are beyond the reach of any machine relying on the classical laws of physics.

The practical realization of quantum computers is still in its infancy. Interestingly, over 40 years after the initial suggestion of using spins to represent (classical) information, NMR spectroscopy has become the workhorse for experimental explorations of quantum information processing.

In this article, we first explain how quantum computers work and why they could solve certain problems so much faster than any classical computer. Next, we describe how quantum computers can be implemented using nuclear magnetic resonance (NMR) techniques and what is involved in designing and implementing QC pulse sequences, preparing a suitable initial state and interpreting the output spectra. We close with an overview of the state of the art and the prospects for NMRQC and other QC implementations.

Good reviews of quantum computation and information can be found in Refs.^{3–5} A specialized review of and a good introduction to NMRQC are given in Refs.^{6,7} respectively.

2 QUANTUM COMPUTATION

2.1 INEPT as a Computation

We start from a familiar place for many NMR spectroscopists, the INEPT pulse sequence (see **INEPT**, Volume 4). This sequence was designed to transfer polarization from a high γ nucleus to a low γ nucleus. However, it can also be viewed as a *logic gate* (Figure 1) which flips one spin conditioned upon the orientation of a neighboring spin.

If we arbitrarily assign “0” to a spin up and “1” to a spin down, we can think of spin-1/2 nuclei as bits in a digital computer. We remind the reader that bits (“0” or “1”) can be used to represent numbers in a similar way to decimal numbers (“0” through “9”). In decimal representation, 503 means $3 \times 10^0 + 0 \times 10^1 + 5 \times 10^2$; similarly, in binary representation, 1101 means $1 \times 2^0 + 0 \times 2^1 + 1 \times 2^2 + 1 \times 2^3$ (which corresponds to the number 13 in decimal representation). All computers today process information in binary representation, and are able to rapidly perform multiplication and division (repeated addition and subtraction) as well as the most complex computations, via a sequence of very simple, elementary operations, called logic gates, acting on just one or two bits at a time.

In this binary framework, the INEPT sequence corresponds to an elementary two-bit operation known as the controlled-NOT or for short CNOT gate within the phase corrections discussed in Section 3.1. The CNOT gate performs a NOT operation on a second bit, flipping it from “0” to “1” or from “1” to “0”, if and only if the value of a first bit is “1”. The input to the logic gate is the initial state of the spins, and the output is the final state of the spins. The four possible input values and the corresponding output values are tabulated in Figure 2.

The CNOT combined with single-spin rotations provides for a *universal* set of logic gates. This means that *any* computational task can be implemented using a sufficiently large number of nuclear spins simply by concatenating CNOTs and single-spin rotations in the proper way.^{8–10} In summary, spin-1/2 nuclei in a molecule can serve as bits in a computer, while pulses and delay times provide universal logic gates.

2.2 Quantum Parallelism

Everything we have discussed so far has been purely classical. The Bloch sphere picture of Figure 1 reinforces this




Figure 1 The evolution of one of two coupled heteronuclear spins during an INEPT type pulse sequence, when the other spin is up (solid line) or down (dashed line). The rotating frame is set on resonance with the first spin so there is no need to refocus chemical shift. The usual read-out pulse is left out. The same pulse sequence can be applied to two homonuclear spins using spin-selective pulses




Figure 2 (a) Input and output states for the INEPT pulse sequence and (b) for the corresponding CNOT gate

classical view of the spins; however, nuclear spins are quantum objects. In Dirac notation, the state of a spin can be denoted $|0\rangle$ for a spin in the ground state (along z), and $|1\rangle$ for a spin in the excited state (along $-z$), corresponding to the two classical values for a bit (“0” and “1”). Now, a spin along the x axis is in reality in a coherent superposition state of spin up and spin down, written as $(|0\rangle + |1\rangle)/\sqrt{2}$, while a spin along the y axis is in the state $(|0\rangle + i|1\rangle)/\sqrt{2}$, etc. A spin-1/2 particle is thus more complex than a classical bit. Any two-level quantum system, such as a spin-1/2 particle, can serve as a quantum bit (qubit).

The difference between the quantum and classical descriptions becomes clear as soon as more than one quantum particle is considered. For example, it is well known that the state of n interacting spin-1/2 cannot be described simply by n sets of coordinates on the Bloch sphere. In order to include phenomena such as multiple-quantum coherence, we need recourse to $4^n - 1$ real numbers in the product operator expansion or equivalently to density matrices of dimension $2^n \times 2^n$. Furthermore, the evolution of a closed system of n spins can only be described by $2^n \times 2^n$ unitary matrices (see **Liouville Equation of Motion**, Volume 4). The number of degrees of freedom that need to be specified in a classical description of the state and dynamics of n coupled spins thus increases exponentially with the number of spins.

Richard Feynman proposed in 1982 that the exponential complexity of quantum systems might be put to good use in simulating the dynamics of another quantum system,¹¹ a task which requires exponential effort on a classical computer. David Deutsch extended and formalized this idea in 1985, and introduced the notion of “quantum parallelism”.¹²

Consider a (classical) logic gate which implements a function f with one input bit x and one output bit $f(x)$. If $x = 0$, the gate will output $f(0)$; if $x = 1$, the output will be $f(1)$. The analogous quantum logic gate is described by a unitary operation which transforms a qubit as

$$|0\rangle \mapsto |f(0)\rangle \quad \text{and} \quad |1\rangle \mapsto |f(1)\rangle$$

However, due to the possibility of preparing coherent superposition states and to the linearity of quantum mechanics, the same gate also performs the transformation

$$\frac{|0\rangle + |1\rangle}{\sqrt{2}} \mapsto \frac{|f(0)\rangle + |f(1)\rangle}{\sqrt{2}}$$

In this sense, it is possible to evaluate $f(x)$ for both input values in one step. Next consider a different logic gate which implements a function $g(x)$ with two input bits. We can prepare each qubit in a superposition of “0” and “1”. Formally, the state of the joint system is then written as $(|0\rangle + |1\rangle) \otimes (|0\rangle + |1\rangle)/2$. Leaving out the tensor product symbol as well as any normalization factors, the state can be written as $(|0\rangle + |1\rangle)(|0\rangle + |1\rangle)$, or $|0\rangle|0\rangle + |0\rangle|1\rangle + |1\rangle|0\rangle + |1\rangle|1\rangle$, which is further abbreviated to $|00\rangle + |01\rangle + |10\rangle + |11\rangle$. Therefore, a set of two spins can be in a superposition of the four states “00”, “01”, “10” and “11”. A quantum logic gate implementing $g(x)$ then transforms this state as

$$|00\rangle + |01\rangle + |10\rangle + |11\rangle \mapsto |g(00)\rangle + |g(01)\rangle + |g(10)\rangle + |g(11)\rangle$$

Thus the function has been evaluated for the four possible input values in parallel. In general, a function of n qubits implemented on a quantum computer can be evaluated for all 2^n input values in parallel. In contrast to classical computers, for which the number of parallel function evaluations increases at best linearly with their size, the number of parallel function evaluations grows exponentially with the size of a quantum computer (the number of qubits).

Obviously, this is true only so long as the coherent superposition states are preserved throughout the computation. This means that the computation should be completed before quantum coherence is lost due to “decoherence” processes (in NMR, spin–spin and spin-lattice relaxation; see **Relaxation: An Introduction**, Volume 6). Since some degree of decoherence is unavoidable, practical quantum computers appeared virtually impossible to build, until quantum error correction was conceived, as discussed in Section 2.4.

Even if the coherence time is long compared to the duration of a typical logic gate and quantum error correction is employed, can we really access the exponential power exhibited by quantum systems? The postulates of quantum mechanics dictate that a measurement – which induces instantaneous and complete decoherence – of a qubit in a superposition state $|f(0)\rangle + |f(1)\rangle$ will give either “ $f(0)$ ” or “ $f(1)$ ”, with equal probabilities. Similarly, after doing 2^n computations all at once, resulting in a superposition of 2^n output values, measurement of the quantum bits randomly returns a single output value. A more clever approach is thus needed exploiting quantum parallelism requires the use of quantum algorithms.

2.3 Quantum Algorithms

Special quantum algorithms allow one to take advantage of quantum parallelism in order to solve certain problems in far fewer steps than is possible classically. When comparing the capability of two computers to solve a certain type of problem, the relevant criterion is not so much what resources (time, size, signal-to-noise ratio, ...) are required to solve the problem but rather how quickly the required resources grow with the problem size.

A particularly important criterion is whether the required resources increase exponentially or polynomially with the problem size. Exponentially difficult problems are considered intractable; they become impossible to solve when the problem size is large. In contrast, polynomially difficult problems are considered tractable. The interest in quantum computing is based on the fact that certain problems which

appear intractable (resources grow exponentially with problem size) on any classical computer are tractable on a quantum computer.

This was shown in 1994 by Peter Shor, almost 10 years after Deutsch introduced quantum parallelism. Shor's quantum algorithm¹³ allows one to find the period of a function exponentially faster than any classical algorithm. The importance of period-finding lies in that it can be translated, using some results from number theory, to finding the prime factors of integer numbers, and thus also to breaking widely used cryptographic codes. These codes are based precisely on the fact that no *efficient* classical algorithm is known for period-finding or factoring, i.e., the effort required to factor a number on classical computers increases exponentially with the number of digits of the integer to be factored. In contrast, Shor's algorithm is efficient: on a quantum computer, the effort to factor an integer increases only polynomially (to be precise, the third power) with the number of digits of the integer. As a result, while factoring a 1000-digit number is believed to be beyond the reach of any machine relying on the classical laws of physics, such a feat could be accomplished on a quantum computer.

The first quantum algorithm was invented by Deutsch and Jozsa¹⁴ in 1992. This algorithm allows a quantum computer to solve with certainty an artificial mathematical problem known as Deutsch's problem. Even though this algorithm does not have much application, it is historically significant as it provided the first steps towards Shor's algorithm, and because it is a simple quantum algorithm that can be experimentally tested.

Another class of quantum algorithms was discovered in 1996 by Lov Grover. These algorithms¹⁵ allow quadratic speed-ups of certain search problems, for which there is no better approach classically than to try all L candidate solutions one at a time. A quantum computer using Grover's algorithm needs to make only \sqrt{L} such trials. Even though this speed-up is only quadratic rather than exponential, it is still significant.

The last currently known application of quantum computers lies in simulating other quantum systems,¹⁶ as Feynman conjectured. Even a computer consisting of no more than a few dozen quantum bits could outperform the fastest classical computers in solving important physics problems, such as calculating the energy levels of an atom.¹⁷

In the remainder of this section, we briefly review the structure of Shor's algorithm, because it is so important and at the same time gives good insight into how quantum computing works (for a more detailed explanation, see Refs.^{4,13}). The crucial step in Shor's factoring algorithm is the use of the quantum Fourier transform (QFT) to find the period r of the function $f(x) = a^x \bmod M$, which means $f(x)$ is the remainder after division of a^x by M , where M is the integer to be factored, and a is an integer which is more or less randomly chosen.^{4,13}

The QFT performs the same transformation as the (classical) fast Fourier transform (FFT), but can be computed exponentially faster. As always, we do not have access to all the individual output values; the QFT merely allows us to *sample* the FFT but we will see that this suffices for period-finding. The FFT_N takes as input a string of N complex numbers x_j and produces as output another string of N complex numbers

y_k , such that

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j e^{2\pi ijk/N} \quad (1)$$

For an input string with numbers which repeat themselves with period r , the FFT_N produces an output string with period N/r , as illustrated in the following four examples for $N = 8$ (the output strings have been rescaled for clarity)

r	input string		output string	N/r	
8	1 0 0 0 0 0 0 0	\mapsto	1 1 1 1 1 1 1 1	1	(a)
4	1 0 0 0 1 0 0 0	\mapsto	1 0 1 0 1 0 1 0	2	(b)
2	1 0 1 0 1 0 1 0	\mapsto	1 0 0 0 1 0 0 0	4	(c)
1	1 1 1 1 1 1 1 1	\mapsto	1 0 0 0 0 0 0 0	8	(d)

If r does not divide into N , the inversion of the period is approximate. Furthermore, the FFT turns shifts in the locations of the numbers in the input string into phase factors in front of the numbers in the output string:

1 0 0 0 1 0 0 0	\mapsto	1 0	1 0	1 0	1 0	(e)			
0 1 0 0 0 1 0 0	\mapsto	1 0	$-i$	0	-1	0	i	0	(f)
0 0 1 0 0 0 1 0	\mapsto	1 0	-1	0	1 0	-1	0	(g)	
0 0 0 1 0 0 0 1	\mapsto	1 0	i	0	-1	0	$-i$	0	(h)

The QFT performs exactly the same transformation, but differs from the FFT in that the complex numbers are stored in the amplitude and phase of the terms in a superposition state. The amplitude of the $|000\rangle$ term represents the first complex number, the amplitude of the $|001\rangle$ term the second number and so forth. For clarity, we will label the states $|000\rangle, |001\rangle, \dots, |111\rangle$ as $|0\rangle, |1\rangle, \dots, |7\rangle$. As an example, we see from (f) that the QFT transforms the state $|1\rangle + |5\rangle$ to the state $|0\rangle - i|2\rangle - |4\rangle + i|6\rangle$.

The QFT is incorporated in actual quantum algorithms as outlined in Figure 3. A first register (group of qubits) is prepared in a superposition of all its possible states. A second register is initialized to the ground state (for factoring a number M , the size of the second register must generally be at least $\log_2 M$ and the first register must be at least twice as large). For example, if register 1 has three qubits and register 2 has two qubits, the state of the system is prepared in

$$(|0\rangle + |1\rangle + |2\rangle + |3\rangle + |4\rangle + |5\rangle + |6\rangle + |7\rangle)|0\rangle \quad (2)$$




Figure 3 Schematic diagram of the main steps in quantum algorithms for period-finding

Then the function $f(x)$ is evaluated (in NMR by applying a pulse sequence, as discussed in Sections 2.1 and 3.1), where x is the value of the first register, and the output value $f(x)$ is stored in the second register. Since the first register is in an equal superposition of all $|x\rangle$, the function is evaluated for all values of x from 0 to 7 in parallel. For example, let $f(x) = 3$ for even x and $f(x) = 1$ for odd x , which means the period r is 2 (in real applications, we have a description of f but do not know r in advance). Evaluation of $f(x)$ then transforms the state of equation (2) to the state

$$\begin{aligned} & |0\rangle|3\rangle + |1\rangle|1\rangle + |2\rangle|3\rangle + |3\rangle|1\rangle \\ & + |4\rangle|3\rangle + |5\rangle|1\rangle + |6\rangle|3\rangle + |7\rangle|1\rangle \quad (3) \\ & = (|0\rangle + |2\rangle + |4\rangle + |6\rangle)|3\rangle + (|1\rangle + |3\rangle + |5\rangle + |7\rangle)|1\rangle \quad (4) \end{aligned}$$

We pause to point out that this state is *entangled*, which means that it cannot be written as a product of single-qubit states. The state $|00\rangle + |01\rangle + |10\rangle + |11\rangle$ is an example of an unentangled state, because it can be written as $(|0\rangle + |1\rangle)(|0\rangle + |1\rangle)$, a product of single-qubit states. In contrast, $|00\rangle + |11\rangle$ is a simple example of an entangled state. Entanglement has no classical analogue and is believed to lie at the heart of the exponential speed-up offered by quantum computation.

In order to appreciate the role of the QFT, suppose we now measure the second register in equation (4) (this measurement can be left out but simplifies the explanation). The state of the first register will collapse to either

$$|0\rangle + |2\rangle + |4\rangle + |6\rangle \quad \text{or} \quad |1\rangle + |3\rangle + |5\rangle + |7\rangle \quad (5)$$

depending on whether the measurement of register 2 gave “3” or “1”. We see that all eight possible outcomes are still equally likely, so the result of a measurement does not give us any useful information. But if we apply the QFT, the first register will be transformed to

$$|0\rangle + |4\rangle \quad \text{or} \quad |0\rangle - |4\rangle \quad (6)$$

Now a measurement of the first register does give useful information, because only multiples of N/r are possible outcomes, in this example “0” and “4”.

This concludes the quantum part of the computation. From the measurement result, a classical computer can efficiently calculate the inverted period N/r , and thus also r , with high probability of success using results from number theory. Now that r is known, the factors of the integer M can be quickly computed from number theory as well, with high probability (the probability of success can be further increased by repeating the whole procedure a few times).

We close with two final remarks on quantum algorithms: (1) quantum computing cannot offer any speed-up for many common tasks, such as adding up two numbers or word processing, which can already be completed efficiently on a classical computer; and (2) there are many problems which take exponential effort on both quantum and classical computers. It would be somewhat disappointing from a practical viewpoint if no other applications were found; however, our understanding of the connection between physics and information and computation has already changed dramatically.

2.4 Quantum Error Correction

Any quantum computation must be completed within the coherence time, T_2 and T_1 in NMR, as pointed out in Section 2.2. T_1 and T_2 processes alter the state of the qubits and are therefore a source of errors. For many years, this requirement led to widespread pessimism about the practicality of quantum computers. In 1995, however, Peter Shor and Andrew Steane independently discovered quantum error correction^{46,47} and showed that it is possible to correct for truly random errors caused by decoherence.

This came as quite a surprise, because quantum error correction had to overcome three important obstacles: (1) the no-cloning theorem, which states that it is not possible to copy unknown quantum states;⁵ (2) measuring a quantum system affects its state; and (3) errors on qubits can be arbitrary rotations in Hilbert space, compared with simple bit flips for classical bits. Quantum error correction requires many extra operations and extra qubits (ancillae), however, which might introduce more errors than are corrected, especially because the effect of decoherence increases exponentially with the number of entangled qubits, in much the same manner that multiple quantum coherences decay exponentially faster than single quantum coherences. Therefore, a second surprising result⁴⁸ was that provided the error rate (probability of error per elementary operation) is below a certain threshold, and given a fresh supply of ancilla qubits in the ground state, it is possible to perform arbitrarily long quantum computations.

The threshold error rate is currently estimated⁴⁸ to be about 0.001%. The actual error rate in NMRQC is approximated by $1/2JT_2$, where $2JT_2$ is roughly the number of operations that can be computed within the coherence time. For small molecules the error rate is typically on the order of 0.1–1%, two to three orders of magnitude too high. A remarkable implication of quantum error correction, however, is that if (1) a molecule is found which achieves the threshold error rate; and (2) select spins can be fully polarized, both T_1 and T_2 could in principle be infinitely lengthened by applying an error correction pulse sequence.

3 NMR QUANTUM COMPUTERS

3.1 Pulse Sequence Design

At first sight, the translation of abstract quantum algorithms or function evaluations into actual pulse sequences may appear obscure: However, systematic techniques^{18,19} exist to make pulse sequence design relatively straightforward. The starting point is that: each quantum algorithm can be described by a sequence of transformations under unitary operators.

Such unitary transformations represent rotations in Hilbert space (a multidimensional extension of the Bloch sphere). Examples of unitary transformations are evolution during RF pulses and free evolution under the system Hamiltonian; relaxation processes give rise to nonunitary transformations. Once the desired unitary operators have been identified, arbitrary unitary operators can be translated into sequences of single-qubit rotations and CNOT gates.

These building blocks can be readily implemented in NMR (see Section 3.2). Decompositions into other sets of elementary gates are also possible, and can be helpful in simplifying the pulse sequences.^{20,38} In any case, it is crucial that the duration

of the pulse sequence design process as well as the length of the resulting pulse sequence do not increase exponentially with the problem size.

We now point out two important distinctions between QC and conventional pulse sequences. On the one hand, QC sequences must be more general: QC sequences must perform the desired transformation for arbitrary input states. In contrast, conventional sequences are often designed assuming a particular input state. As a first example of this difference, the sequence of Figure 1 assumes that both spins are in Zeeman states, i.e., aligned along $\pm z$. It implements the unitary operator

$$\hat{U}_{\text{INEPT}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -i & 0 \end{pmatrix} \quad (7)$$

which is similar to but different from the unitary operator for the CNOT gate, defined as

$$\hat{U}_{\text{CNOT}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (8)$$

implemented, for example, by $90_z^a 90_{-z}^b 90_x^b 1/2 J_{ab} 90_{-y}^b$.

As a second example, consider the so-called Hadamard gate, defined as

$$\hat{U}_{\text{Had}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (9)$$

This gate creates a superposition state starting from a basis state: it transforms $|0\rangle$ to $|0\rangle + |1\rangle$ (z to x in the Bloch sphere) and $|1\rangle$ to $|0\rangle - |1\rangle$ ($-z$ to $-x$). At first sight, this transformation could be done simply via a 90_y pulse. However, the unitary operator for 90_y

$$\hat{U}_{90_y} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \quad (10)$$

is different from \hat{U}_{Had} ; e.g., applying \hat{U}_{Had} twice has no net effect, but applying \hat{U}_{90_y} twice produces U_{180_x} . A possible sequence to implement \hat{U}_{Had} exactly is $90_y 180_x$.

On the other hand, QC sequences can be more specific: QC sequences can be specialized for a specific molecule using full knowledge of its spectral properties. In contrast, conventional sequences must work for any molecule, because the spectral properties of the molecule are usually not known in advance. Exact knowledge of the chemical shifts and J -coupling constants allows one not only to greatly simplify the pulse sequences, but also to achieve much more accurate unitary transformations than would otherwise be possible.

Finally, while systematic procedures exist to design a pulse sequence, there is a need to develop tools to find the pulse sequence of the shortest duration and with the fewest RF pulses. Even small-scale quantum computations easily involve tens to hundreds of gates acting on multiple spins and precise control of the spin dynamics is difficult to maintain throughout such long sequences of operations, as shown in the next section.

3.2 Implementation of Computations

The implementation of quantum computations with NMR can be based on single-spin rotations and CNOT gates, since any quantum algorithm can be translated into these building blocks. Although these elementary operations appear quite easy to implement, the requirements for precision in QC experiments are unusually high, due to the large number of pulses and the quantitative nature of the information contained in the output spectra.

Implementation of accurate *single-spin rotations* about an axis in the x - y plane is relatively easy in heteronuclear molecules; yet, it can be very demanding for homonuclear spin systems because spin selectivity requires longer pulses resulting in considerable coupled evolution of the spins during the pulses.²¹ Clearly, some degree of homonuclear dynamics is unavoidable when more than a handful of qubits are involved.

We therefore begin by reviewing the requirements for pulse shaping²² (see **Selective Pulses**, Volume 7; **Shaped Pulses**, Volume 7). First, the magnetization corresponding to each of the lines in a multiplet must be rotated about exactly the same axis and over exactly the same angle, i.e., off-resonance effects due to line splitting must be removed. This requires self-refocusing shaped pulses or tailored composite pulses²³ (see **Composite Pulses**, Volume 2). Second, the effect of J -coupling between unselected spins must be removed, either during the pulse or later in the pulse sequence. Third, all pulses must be universal rotors, i.e., the rotation must be independent of the initial state of the spin. Fourth, the unselected spins must not be affected by the RF irradiation. This last requirement is difficult to satisfy because of transient Bloch–Siegert effects²⁴ which can result in substantial (tens of degrees) phase shifts of nearby unselected spins. However, it is possible to estimate and compensate for the Bloch–Siegert shift.^{25,26} Finally, simultaneous (as opposed to consecutive) pulses at two or more nearby frequencies are desirable in order to keep pulse sequences short, but transient Bloch–Siegert shifts greatly deteriorate such simultaneous rotations.^{27,28} Accurate simultaneous rotations at nearby frequencies can be achieved using a special correction technique.²⁹ Even then, simultaneous pulses on well-coupled spins may excite multiple quantum coherences.³⁰

There are a number of hardware requirements for successful execution of QC experiments. Good RF coil homogeneity is crucial in avoiding excessive signal attenuation and related errors. Furthermore, it is desirable that one frequency source and transmitter board per qubit be available. If there are more qubits than spectrometer channels, the carrier frequency must be jumped to the appropriate frequencies throughout the pulse sequence, or phase ramping techniques must be employed.³¹ A dedicated frequency source for each qubit also makes it easy to keep track of the rotating frame of each spin and apply all the pulses on any given spin with the correct relative phase. This removes the need for extra pulses to refocus chemical shift evolution.²² Alternatively, software rotating frames can be created by detailed bookkeeping of the time elapsed since the beginning of the pulse sequence such that the evolution of the rotating frame of any given spin with respect to the carrier reference frame can be calculated. The phases of the pulses throughout the pulse sequence, as well as the receiver phase, can then be adjusted accordingly.^{25,36} The same technique can be used to implicitly realize single-spin rotations about z .

Alternatively, z -rotations can be implemented using resonance offsets or composite pulses.²²

Two strategies exist for implementing CNOT gates (both assume first order spectra). If all the spins are mutually coupled, CNOT's can be realized via line-selective pulses which invert specific lines within a multiplet.¹⁰ In practice, it is usually more convenient to use pulse sequences such as the one demonstrated in Figure 1.^{8,10} For molecules with several coupled spins, the sequence of Figure 1 must be expanded with extra pulses to refocus undesired J -couplings; systematic methods exist to design efficient refocusing schemes.^{28,33,34} A CNOT between two uncoupled spins can be realized by swapping qubit states.^{32,35} For example, for a CNOT between two spins a and c which are not mutually coupled but which are both coupled to a third spin b , the procedure is as follows: apply a pulse sequence which swaps the state of a and b (via $\text{CNOT}_{ab} \text{CNOT}_{ba} \text{CNOT}_{ab}$), then perform a CNOT between b and c , and then swap a and b again. The net result is CNOT_{ac} ; spin b is unaffected. It is thus not necessary that all spins be pairwise coupled as long as the network of couplings includes all n spins.

An alternative to imposing the correct evolution on all the spins at all times, both for RF pulses and for CNOT-type gates, is to allow erroneous evolutions which will later be reversed. Such techniques have been highly successful in certain standard pulse sequences (see **Decoupling Methods**, Volume 3), but are much harder to develop for the nonintuitive and nontransparent QC pulse sequences. Nevertheless, it has been shown experimentally that a large degree of cancellation of erroneous evolutions is possible even in QC experiments: about 300 two-qubit gates involving over 1350 90° RF pulses have been successfully concatenated.³⁶ A general methodology designed to take advantage of this possibility has yet to be developed.

From this discussion, it will be clear that the selection of a suitable molecule is crucial for NMRQC. The desired properties are (1) sufficiently large chemical shifts for good addressability, (2) large coupling constants, while maintaining first-order spectra, for fast two-qubit gates (or a coupling network that matches the pattern of connectivities needed for the algorithm), and (3) long T_2 and T_1 in order to allow time to execute many logic gates. Furthermore, high γ nuclei are desirable for good sensitivity. More mundane but equally important requirements are that the molecule be stable, synthesizable, soluble and safe.

3.3 State Initialization

Apart from experiments designed to produce nonthermal spin polarizations, setting up a proper initial state for the nuclear spins is a concept worth revisiting for NMR spectroscopists. Since it is such a crucial step in quantum computing, this entire section is therefore devoted to state initialization.

Most quantum computations require a *pure* initial state, for example a set of fully polarized spins, in the state $|0\dots 0\rangle$. However, nuclear spins in thermal equilibrium at room temperature are in an almost totally random state: for typical magnetic field strengths, the ground ($|0\rangle$) and excited ($|1\rangle$) state probabilities differ by only about 1 part in 10^5 . The spins are then said to be in a *mixed* (nonpure) state. The polarization could be increased using hyperpolarization techniques (see **Optically Enhanced Magnetic Resonance**, Volume 5) but the state of the art is still very far from cooling nuclear spins into the ground state.

The conceptual breakthrough which made NMR quantum computation possible at room temperature was the concept of *effective pure* or *pseudo-pure* states:⁸⁻¹⁰ effective pure states are mixed states which produce the same signal as a pure state to within a scaling factor.

The signature of an effective pure state for n spins is that all but one of the 2^n populations are equal, and that no coherences are present. The density matrix then consists of an identity component and a pure state component. The identity density matrix is not observable in NMR since only population differences are observed, and furthermore does not transform under unitary evolutions ($UIU^\dagger = I$). Thus the visible signal is produced solely by the one distinct population, corresponding to a pure state. In product operator notation,³⁷ the effective pure ground state is proportional to $\widehat{I}_z + \widehat{S}_z + 2\widehat{I}_z\widehat{S}_z$ for two spins, to $\widehat{I}_z + \widehat{S}_z + \widehat{R}_z + 2\widehat{I}_z\widehat{S}_z + 2\widehat{I}_z\widehat{R}_z + 2\widehat{S}_z\widehat{R}_z + 4\widehat{I}_z\widehat{S}_z\widehat{R}_z$ for three spins, and so forth. Figure 4 shows how the effective pure state preparation is manifest in the spectrum of one of five coupled spins. A characteristic of effective pure (basis) states is that only one line survives in each multiplet.

Three methods are known for preparing effective pure states starting from thermal equilibrium.

- (1) *Logical labeling*^{8,38} consists of applying a pulse sequence which rearranges the thermal populations such that a subset of the spins is in an effective pure state, conditioned upon the state of the remaining spins. Then the computation is carried out within this embedded subsystem.³⁹




Figure 4 (a) Spectrum of pentafluorobutadienyl cyclopentadienyldicarbonyliron complex in thermal equilibrium. (b) The same spectrum after preparing an effective pure state $|0000\rangle$. (Reproduced from Ref.²⁶ by permission of the American Physical Society)

For example, the Boltzman populations for the states $\{|000\rangle, |001\rangle, |010\rangle, |011\rangle, |100\rangle, |101\rangle, |110\rangle, |111\rangle\}$ for a homonuclear three-spin system deviate from the uniform background by $\{3a, a, a, -a, a, -a, -a, -3a\}$ respectively, where

$$a = \frac{1}{2^n} \frac{\hbar\omega}{2k_B T} \ll 1$$

After rearranging the populations for the eight spin states as $\{3a, -a, -a, -a, a, a, a, -3a\}$, the last two qubits are in an effective pure state conditioned upon the first qubit being $|0\rangle$. As the total number of qubits n in the molecule increases, the relative fraction of effective pure qubits goes to 1, but the preparation sequence becomes complex quite rapidly for large n and the signal strength scales as $n/2^n$.

- (2) *Temporal averaging*⁴⁰ is similar to phase cycling (see **Phase Cycling**, Volume 6), since it consists of adding up the spectra of multiple experiments. However, instead of changing just the phase of some pulses, each experiment starts off with a different state preparation sequence which permutes the populations. For two heteronuclear spins, adding together three experiments which yield respective population deviations $\{a, b, -b, -a\}$, $\{a, -b, -a, b\}$ and $\{a, -a, b, -b\}$ is equivalent to performing an experiment with population deviations $\{3a, -a, -a, -a\}$. For arbitrary n , at least $(2^n - 1)/n$ experiments are needed,²⁶ since the effective pure state is made up of $2^n - 1$ product operator terms and the starting state, thermal equilibrium, contains n terms.
- (3) *Spatial averaging*¹⁰ uses a pulse sequence containing magnetic field gradients (see **Field Gradients & Their Application**, Volume 3) to equalize all the populations but the ground state population. Only one experiment is involved, but the preparation sequence quickly becomes unwieldy for large spin systems and the signal strength decreases exponentially with n .

To date, temporal and spatial averaging have been the most popular choices for preparing effective pure states. Several hybrid schemes^{25,40} have also been developed which trade off complexity of the preparation steps for the number of experiments. Nonetheless, all these state preparation schemes have in common that creating effective pure states incurs an exponential cost either in the signal strength or in the number of experiments involved.

Such an exponential overhead is of course not tolerable for quantum computations. The reason for this cost is that effective state preparation techniques simply select out the signal from the ground state population present in thermal equilibrium, and the fraction of molecules in the ground state is proportional to $n/2^n$.

A significant breakthrough by Schulman and Vazirani resulted in a method to cool a subset of the spins in a molecule down to the ground state without any exponential overhead.^{41,42} The idea is to redistribute the entropy of the spins so that some have zero entropy (pure state) while the entropy of the remaining spins increases. In a sense, this method is simply a polarization transfer scheme. Surprisingly, both the length of the cooling pulse sequence and the number of spins that must be sacrificed increase only about linearly with the number of pure spins desired. Furthermore, the cooling algorithm approaches the entropic bound in the limit

of large numbers of spins. However, with initial polarization $\alpha \ll 1$, a molecule with at least k/α^2 spins is needed to obtain k pure spins. This method is therefore impractical when starting from equilibrium at room temperature, with $\alpha \approx 10^{-5}$. Nevertheless, a combination of hyperpolarization techniques and the Schulman–Vazirani scheme may some day become practical. In any case, despite the exponential cost incurred when preparing effective pure states, the highly random initial state represents *no fundamental* obstacle to scalable quantum computation.

3.4 Read-out

In NMR spectroscopy, only one nucleus is traditionally observed. In QC, however, the concept of a single observe channel and one or more decoupler channels does not apply: the output of a quantum computation is the final state of one or several spins. The final states of each of the output spins must thus be read out.

If each of the output spins ends up in $|0\rangle$ or $|1\rangle$ (or in reality in the effective pure state corresponding to $|0\rangle$ or $|1\rangle$), the answer can be read-out directly by applying a pulse which rotates the spin from $\pm z$ to $\pm x$. With properly referenced receiver phase settings, the spectrum for each output spin then consists of either absorption or emission lines, indicating whether the output value of the corresponding bit is “0” or “1” (Figure 5).

If the output state is a superposition state, the situation is a bit more complicated. For a single quantum computer subject to a “hard” measurement (assumed in Section 2.2), the superposition “collapses” to one of the terms in the superposition, with probabilities given by the square of the amplitude of each term. In contrast, NMR spectra result from averaging over a large number of independently operating quantum computers (molecules), and for each spin the area under the spectral lines is proportional to the $|0\rangle - |1\rangle$ probability difference.

The output state of equation (6) serves as an example: half of the molecules in the ensemble collapse to $|0\rangle$ ($|000\rangle$), while the other half collapses to $|4\rangle$ ($|100\rangle$). In other words, spins 2 and 3 always end up in “0” so their spectral lines are absorptive; in contrast, the signal of spin 1 averages to zero because there are equally many molecules in which spin 1 ends up in “0” as in “1”. It is not clear that such bit-wise averages of probabilistic output states are generally sufficient to solve the problem of interest. For Shor’s period-finding algorithm, this problem can be circumvented⁸ by performing the classical post-processing steps (Section 2.3) on the quantum computer using ancillae qubits; any classical computation can also be done on a quantum computer.⁵ In this way, the output state becomes the period r for all the molecules in the ensemble (as opposed to an average over all the multiples of N/r) and the measurement result becomes deterministic.

Instead of recording the signal of each of the output spins, it is sometimes possible to use the extra information provided by the line splittings due to J couplings to derive the output state of several qubits from the spectrum of a single spin. Since each of the lines in the multiplet can be identified with specific states of the other spins (as in Figure 4), the presence or absence of each line in the multiplet gives information about the state of the other spins (Figure 5).

Finally, while the spectra of a few select spins suffice to obtain the answer to a computation, the full density matrix

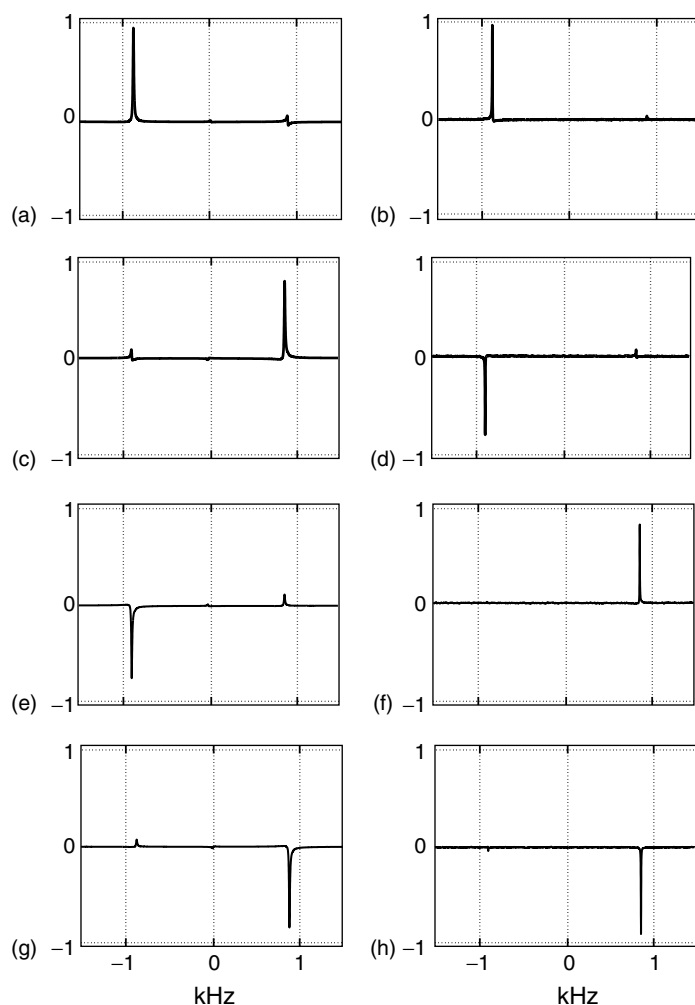


Figure 5 Output spectra of the proton (a) and carbon (b) spins of $^{13}\text{CHCl}_3$ (dissolved in a liquid crystal solution) for four different executions of Grover's search algorithm. Only the real part of the spectra is shown, and frequencies are relative to ν_{H} and ν_{C} . A positive or negative line in the spectrum indicates that the corresponding spin was in $|0\rangle$ or $|1\rangle$ before the read-out pulse. Furthermore, the position of the line in the spectrum of one spin also reveals the state of the other spin. For example, if the ^1H line is at $\nu_{\text{H}} - J_{\text{CH}}/2$, the ^{13}C spin is in $|0\rangle$; a ^1H line at $\nu_{\text{H}} + J_{\text{CH}}/2$ indicates the ^{13}C spin is in $|1\rangle$. Thus, the state of the two qubits for each of the four cases is (from top to bottom) $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$. (Reproduced from Ref.⁶¹ by permission of the American Institute of Physics)

conveys much more information. This extra information can be used to expose the presence of errors such as multiple quantum coherences not visible in the single output spectra and furthermore is a useful tool for debugging pulse sequences. The procedure for reconstructing the density matrix is called quantum state tomography.^{43–45} It consists of repeating the computation multiple times, each time looking at the final state of the spins after applying different sets of read-out pulses which rotate different elements of the density matrix to observable positions. However, since this procedure involves on the order of 4^n experiments, it is practical only for experiments involving a few spins.

3.5 State of the art and Outlook

To date, only very simple demonstrations of quantum algorithms, all using liquid-state NMR techniques, have been carried out. Variations of Grover's algorithm have been demonstrated with two^{44,49,50} and three qubits,³⁶ the Deutsch–Jozsa algorithm with two,^{43,51} three⁵² and five⁵³ qubits, and the

period-finding algorithm with five qubits.²⁶ Most recently, a seven-qubit quantum computer has been used to factor the number 15 using Shor's algorithm.⁶⁷ A quantum simulation has been implemented with two⁵⁴ and three⁵⁵ qubits, and quantum error detection and correction with two,⁵⁶ three⁵⁷ and five⁶⁸ qubits. Finally, a 7-spin coherence has been created and observed²⁵ (see Ref.⁶ for additional references). While these experiments indeed demonstrate the principles of quantum information processing, they all involve far fewer qubits than would be needed to solve a problem beyond the reach of classical machines.

Despite the rapid progress in recent years, scaling liquid state NMRQC to tens or hundreds of qubits may be impractical for several reasons, although none of them appear fundamental. In particular, as the number of qubits increases: (1) the strength of the signal selected with current state initialization techniques decreases exponentially,^{8,9,58} but Schulman–Vazirani cooling⁴¹ does not suffer any such exponential overhead (Section 3.2); (2) the chemical shift separations unavoidably become smaller,

but Lloyd³² showed that for universal quantum computation it suffices to have a linear chain $d - abc - abc - \dots - abc$, in which only nearest neighbour spins are coupled and with only four distinct chemical shifts δ_a , δ_b , δ_c and δ_d ; (3) J -couplings become smaller or even unresolved, but this can be circumvented by swapping spin states (Section 3.2).³² While these obstacles are not fundamental, the solutions all make the pulse sequences, very likely impractically, much longer. This would require increasingly longer T_2 and T_1 for larger molecules, while in practice the T_2 and T_1 tend to become shorter (see **Relaxation: An Introduction**, Volume 6).

NMRQC has also brought up new theoretical issues: (1) since only ensemble averaged results are available because of the large number of molecules in a sample tube, some information is lost that would be available in an idealized quantum computer such as a single molecule at zero degrees Kelvin. For the known quantum algorithms, this information can be retrieved by performing classical post-processing steps on the quantum computer (Section 2.3); and (2) since the density matrix of nuclear spins at room temperature is very close to the identity matrix, it is not possible to produce genuinely entangled states between the nuclear spins in small thermally polarized molecules in liquid solution.⁵⁹ This observation has sparked a stimulating debate about the “quantumness” of NMR, because it implies that each of the states produced in NMRQC experiments so far is classical. However, all attempts to describe the dynamics of a set of coupled spins by an efficient classical model have been unsuccessful. It is thus conjectured that even though the states show no entanglement, the dynamics of the spins is truly quantum mechanical,⁶⁰ a proposition which will appear obvious to most NMR spectroscopists. In fact, it has also been the starting point of this introduction to NMR quantum computing.

4 SUMMARY AND CONCLUSIONS

In many respects, liquid state NMR provides an ideal test bed for elementary quantum computations. The degree of control over the evolution of multiple coupled qubits – the result of 50 years of technology development – the long relaxation times of nuclear spins and a set of new insights^{8,9} made it possible to perform certain computations in *fewer steps* than is possible using any classical machine. This is in itself a remarkable achievement.

It is unlikely that liquid state NMR could ever be used to solve problems *faster* than any classical machine, but it has already inspired many other NMR based proposals for quantum computing. Liquid-crystal solvents have been used to partly reintroduce dipole-dipole couplings (see **Liquid Crystals: General Considerations**, Volume 4) to speed up the gate time and increase the number of gates possible within the coherence time.⁶¹ New molecular architectures based on liquid crystal solutions are now being investigated. Solid-state NMR near 0K could circumvent the state initialization problem, but also poses new questions in terms of addressability and coherence times (see **Internal Spin Interactions & Rotations in Solids**, Volume 4). Several approaches to solve these issues have been proposed.^{62,63} Another proposal which received much attention consists of doing NMR on impurity atoms placed in a linear or two-dimensional array, with chemical

shifts and couplings controlled by electrodes placed on top of and in between the impurity atoms.⁶⁴

Furthermore, there is a plethora of very different experimental approaches to building quantum computers (for an extensive review, see Ref.⁶⁵). Four trapped ions have recently been entangled⁶⁶ and a single quantum logic gate has been implemented between two photons coupled with each other via interactions with an atom in a cavity.⁶⁹ In the long run, the most scalable approaches may be those based on solid-state technology, such as electron spins in quantum dots or magnetic fluxes in SQUIDS.⁶⁵

It is clear that none of these proposals will be easy to implement since they all require substantial and innovative development of technology. The success of any approach will depend on the ratio of the coherence time to the gate duration, i.e., how many gates can be completed within the coherence time, and on the achievable degree of quantum control.

Many of the problems in quantum control are similar for different experimental systems, and we therefore expect that other quantum computer implementations will benefit from the ideas, concepts and solutions which arise from liquid state NMR experiments. In addition, we hope that some of the techniques developed within the context of quantum computation may find more general application in NMR.

The possible payoff for successful quantum computing is tremendous: to solve problems beyond the reach of any classical computer. It is not clear at this point whether quantum computers will fulfill this promise, but in any case quantum computing has already provided an exciting new perspective on NMR and, more broadly, on the connection between physics, information and computation.

5 REFERENCES

1. A. G. Anderson, R. L. Garwin, E. L. Hahn, J. W. Horton, G. L. Tucker, and R. M. Walker, *J. Appl Phys.*, 1955, **26**, 1324.
2. R. Landauer, *Phys. Today*, 1991, **44**, 5, 22.
3. C. Bennett and D. P. DiVincenzo, *Nature*, 2000, **404**, 247.
4. A. Ekert and R. Jozsa, *Rev. Mod. Phys.*, 1996, **68**, 733.
5. M. A. Nielsen and I. L. Chuang, ‘Quantum Computation and Quantum Information’, Cambridge University Press: Cambridge, 2000.
6. J. A. Jones, *Fortschr. Phys.*, 2000, **48**, 909.
7. N. Gershenfeld and I. L. Chuang, *Sci. Am.*, 1998, June.
8. N. Gershenfeld and I. L. Chuang, *Science*, 1997, **275**, 350.
9. D. G. Cory, M. D. Price, and T. F. Havel, *Physica D*, 1998, **81**, 2152.
10. D. G. Cory, A. F. Fahmy, and T. F. Havel, *Proc. Natl. Acad. Sci. USA*, 1997, **94**, 1634.
11. R. P. Feynman, *Int. J. Theor. Phys.*, 1982, **21**, 467.
12. D. Deutsch, *Proc. R. Soc. Lond. A*, 1885, **400**, 97.
13. P. Shor, in ‘Proc. 35th Annual Symposium on the Foundations of Computer Science’, IEEE Comp. Soc. Press: Los Alamitos, CA, 1994, 124.
14. D. Deutsch and R. Jozsa, *Proc. R. Soc. Lond. A*, 1992, **439**, 553.
15. L. Grover, *Phys. Rev. Lett.*, 1997, **79**, 4709.
16. S. Lloyd, *Science*, 1996, **273**, 1073.
17. D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.*, 1999, **83**, 5162.
18. A. Barenco, C. Bennett, R. Cleve, D. P. DiVincenzo, N. H. Margolus, P. W. Shor, T. Sleator, J. A. Smolin, and M. Weinfurter, *Phys. Rev. A*, 1995, **52**, 3457.
19. M. D. Price, S. S. Somaroo, A. E. Dunlop, T. F. Havel, and D. G. Cory, *Phys. Rev. A*, 1999, **60**, 2777.

20. J. A. Jones, R. H. Hansen, and M. Mosca, *J. Magn. Reson.*, 1998, **135**, 353.
21. N. Linden, H. Barjat, R. J. Carbajo, and R. Freeman, *Chem. Phys. Lett.*, 1999, **305**, 28.
22. R. Freeman, 'Spin Choreography', Spektrum: Oxford, 1997.
23. H. K. Cummins and J. A. Jones, *New J. Phys.*, 2000, **2**, 6.
24. L. Emsley and G. Bodenhausen, *Chem. Phys. Lett.*, 1990, **168**, 297.
25. E. Knill, R. Laflamme, R. Martinez, and C.-H. Tseng, *Nature*, 2000, **404**, 368.
26. L. Vandersypen, M. Steffen, G. Breyta, C. Yannoni, R. Cleve, and I. Chuang, *Phys. Rev. Lett.*, 2000, **85**, 5452.
27. Ě. Kupče and R. Freeman, *J. Magn. Reson. A*, 1995, **112**, 261.
28. N. Linden, Ě. Kupče, and R. Freeman, *Chem. Phys. Lett.*, 1999, **311**, 321.
29. M. Steffen, L. M. K. Vandersypen, and I. L. Chuang, *J. Magn. Reson.*, 2000, **146**, 369.
30. Ě. Kupče, J. M. Nuzillard, V. S. Dimitrov, and R. Freeman, *J. Magn. Reson. A*, 1994, **107**, 246.
31. S. L. Patt, *J. Magn. Reson.*, 1992, **96**, 94.
32. S. Lloyd, *Science*, 1993, **261**, 1569.
33. D. W. Leung, I. L. Chuang, F. Yamaguchi, and Y. Yamamoto, *Phys. Rev. A*, 2000, **61**, 042310.
34. J. A. Jones and E. Knill, *J. Magn. Reson.*, 1999, **141**, 322.
35. D. Collins, K. W. Kim, W. C. Holton, H. Sierzputowska-Gracz, and E. O. Stejskal, *Phys. Rev. A*, 2000, **62**, 022304.
36. L. M. K. Vandersypen, M. Steffen, M. H. Sherwood, C. S. Yannoni, G. Breyta, and I. L. Chuang, *Appl. Phys. Lett.*, 2000, **76**, 648.
37. O. W. Sorenson, G. W. Eich, M. H. Levitt, J. Bodenhausen, and R. R. Ernst, *Prog. NMR Spectrosc.*, 1983, **16**, 163.
38. L. M. K. Vandersypen, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, *Phys. Rev. Lett.*, 1999, **83**, 3085.
39. D. Suter, A. Pines, and M. Mehring, *Phys. Rev. Lett.*, 1986, **57**, 242.
40. E. Knill, I. L. Chuang, and R. Laflamme, *Phys. Rev. A*, 1998, **81**, 5672.
41. L. J. Schulman and U. Vazirani, 'Proc. 31st Ann. Symp. on Theory of Comp. Science', 1998, p. 322.
42. D. Chang, L. M. K. Vandersypen, and M. Steffen, see <http://xxx.lanl.gov/abs/quant-ph/0011055>.
43. I. L. Chuang, N. Gershenfeld, M. Kubinec, and D. Leung, *Proc. R. Soc. Lond. A*, 1998, **454**, 447.
44. I. L. Chuang, N. Gershenfeld, and M. Kubinec, *Phys. Rev. Lett.*, 1998, **80**, 3408.
45. I. L. Chuang, L. M. K. Vandersypen, X. Zhou, D. W. Leung, and S. Lloyd, *Nature*, 1998, **393**, 143.
46. P. W. Shor, *Phys. Rev. A*, 1995, **52**, 2493.
47. A. Steane, *Phys. Rev. Lett.*, 1996, **77**, 793.
48. D. Aharonov and M. Ben-Or, see <http://xxx.lanl.gov/abs/quant-ph/9906129>.
49. J. A. Jones, M. Mosca, and R. H. Hansen, *Nature*, 1998, **393**, 344.
50. J. A. Jones and M. Mosca, *Phys. Rev. Lett.*, 1999, **83**, 1050.
51. J. A. Jones and M. Mosca, *J. Chem. Phys.*, 1998, **109**, 1648.
52. N. Linden, H. Barjat, and R. Freeman, *Chem. Phys. Lett.*, 1998, **296**, 61.
53. R. Marx, A. F. Fahmy, J. M. Myers, W. Bermel, and S. J. Glaser, *Phys. Rev. A*, 2000, **62**, 012310.
54. S. Somaroo, C.-H. Tseng, T. F. Havel, R. Laflamme, and D. G. Cory, *Phys. Rev. Lett.*, 1999, **82**, 5381.
55. C.-H. Tseng, S. Somaroo, Y. Sharf, E. Knill, R. Laflamme, T. F. Havel, and D. G. Cory, *Phys. Rev. A*, 2000, **61**, 012302.
56. D. Leung, L. Vandersypen, X. Zhou, M. Sherwood, C. Yannoni, M. Kubinec, and I. Chuang, *Phys. Rev. A*, 1999, **60**, 1924.
57. D. G. Cory, W. Maas, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo, *Phys. Rev. Lett.*, 1998, **81**, 2152.
58. W. S. Warren, *Science*, 1997, **277**, 1688.
59. S. L. Braunstein, C. M. Caves, R. Jozsa, N. Linden, S. Popescu, and R. Schack, *Phys. Rev. Lett.*, 1999, **83**, 1054.
60. R. Schack and C. Caves, *Phys. Rev. A*, 1999, **60**, 4354.
61. C. S. Yannoni, M. H. Sherwood, D. C. Miller, I. L. Chuang, L. M. K. Vandersypen, and M. G. Kubinec, *Appl. Phys. Lett.*, 1999, **75**, 3563.
62. D. G. Cory, R. Laflamme, E. Knill, L. Viola, T. F. Havel, N. Boulant, G. Boutis, E. Fortunato, S. Lloyd, R. Martinez, C. Negrevergne, M. Pravia, Y. Sharf, G. Teklemariam, Y. S. Weinstein, and W. H. Zurek, *Fortschr. Phys.*, 2000, **48**, 875.
63. F. Yamaguchi and Y. Yamamoto, *Appl. Phys. A*, 1999, **68**, 1.
64. B. E. Kane, *Nature*, 1998, **393**, 133.
65. Special Issue, *Fortschr. Phys.*, 2000, **48**, 9–11.
66. C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe, *Nature*, 2000, **404**, 256.
67. L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, *Nature*, 2001, **414**, 883.
68. E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, *Phys. Rev. Lett.*, 2001, **86**, 5811.
69. Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, *Phys. Rev. Lett.*, 1995, **75**, 4710.