

Nuclear Magnetic Resonance:
a quantum technology for computation and
spectroscopy

H. K. Cummins*
Oxford Centre for Quantum Computation
Clarendon Laboratory
Parks Road
Oxford OX1 3PU
UK
holly.cummins@qubit.org

J. A. Jones
Oxford Centre for Quantum Computation
and Oxford Centre for Molecular Sciences
New Chemistry Laboratory
South Parks Road
Oxford OX1 3QT
UK
jonathan.jones@qubit.org

August 3, 2000

Abstract

In this article we consider nuclear magnetic resonance (NMR) as an example of a quantum technology; we consider in particular detail the implementation of quantum computers using NMR. We begin by outlining the physical principles underlying NMR, and give an introduction to the quantum mechanics involved. We next discuss the general characteristics of quantum technologies and the ways and extent to which these characteristics are expressed in NMR. We then give an introduction to the subject of quantum computation and its implementation using NMR. Finally, we describe some spectroscopy techniques which also exploit the quantum nature of NMR.

*To whom correspondence may be addressed.

1 Introduction

Quantum physicists become accustomed to the glamour of their field. Their subject is the very small, the very strange, the theoretically forbidden, and, usually, the very difficult. It is easy to forget, after wrestling with the more exotic aspects of the microscopic world all day, that quantum physics might also include aspects which are concrete, wet, and, worst of all, occasionally organic. One of these disconcertingly substantial facets of quantum mechanics is Nuclear Magnetic Resonance spectroscopy (NMR). While physicists generally leave NMR to chemists, biochemists, and doctors, NMR does in fact have many interesting quantum aspects; liquid-state NMR is an unusual and often overlooked example of a quantum technology. NMR has recently been used to perform quantum computations which use quantum effects to achieve non-classical increases in computational efficiency. Many spectroscopic techniques also exploit quantumness to extract otherwise inaccessible geometric information about a molecule or to probe uncooperative atoms by giving them the attributes of more cooperative atoms.

2 Basics of NMR spectroscopy

Nuclear magnetic resonance was first observed in 1946 by two groups who independently probed the behaviour of the hydrogen nuclei (in water and paraffin) in strong magnetic fields. The physicists who discovered it anticipated it would be an ideal method for measuring the magnetic moments of various atomic nuclei, but further investigation revealed that the resonance frequency of a nucleus was influenced by its chemical environment. This unfortunate dependence on chemical details made the technique useless for the simple physical experiments envisioned but simultaneously transformed it into a valuable spectroscopic tool [1]. NMR is now an important analytic method in many fields, including chemistry, biology, medicine, materials science, and geology.

In essence, NMR is the study of transitions between the Zeeman levels of an atomic nucleus in a magnetic field. For experimental and theoretical convenience, attention is often restricted to the spin-half nuclei. In the presence of a magnetic field B_z , directed along the z -axis, the degeneracy of the two spin states of a spin-half nucleus is lifted by the Zeeman Hamiltonian,

$$H = -\frac{1}{2}\hbar\gamma\sigma_z B_z \quad (1)$$

where γ (the gyromagnetic ratio) is a constant characteristic of the nucleus

and σ_z is one of the Pauli spin matrices. The allowed energy levels are multiples of the eigenvalues of σ_z ; that is,

$$E = \pm \frac{1}{2} \hbar \gamma B_z \quad (2)$$

In order to make E a reasonable size, B_z is made as large as possible by using sophisticated superconducting electromagnets. The most advanced of these have fields around 20 T, which is approximately 400 000 times the earth's magnetic field, while typical laboratories might have magnets whose field is about 5–15 T. These extraordinarily strong magnets are also elaborately engineered to set up fields which are homogeneous across the sample to within one part in 10^{10} .

What justifies this level of effort in manufacturing strong homogeneous magnets and makes equation (2) interesting is that the precise B_z experienced by a nucleus in a magnetic field, while largely determined by the magnet, is also affected by tiny fields set up by any nearby electrons, so that the same nucleus in different chemical environments will have slightly different energy levels. The presence of such a set of energy levels can be detected by spectral absorption as long as we can control some interaction that causes transitions between levels; in the case of NMR this is provided by the coupling between the nucleus and a rotating (or alternating) magnetic field applied perpendicular to the main field with an angular frequency such that

$$\hbar \omega = \Delta E \quad (3)$$

where ΔE is the spacing between energy levels, and ω lies in the radiofrequency region of the spectrum for currently achievable magnetic fields.

Although what we have described is already a useful tool for chemical analysis, it doesn't incorporate any of the spookier features of quantum mechanics. To do this, we need to turn our attention to molecules which contain more than one spin-half nucleus, since the second important interaction in liquid-state NMR, responsible for the creation of non-classical correlations between spins, arises from spin-spin interactions in these systems. Perhaps surprisingly, the most important spin-spin coupling is not dipole-dipole coupling: the nuclei do interact with one another as dipoles, but these magnetic interactions are averaged out in liquids by rapid molecular tumbling. There is a second interaction between nuclei, related to the Fermi contact interaction and mediated by shared valence electrons, which is not averaged out completely. This coupling is known as the scalar or J-coupling. It is directly

observable in NMR spectra as a splitting in the NMR signals corresponding to each nucleus, as shown in figure 1.



Figure 1: Cytosine and its NMR spectrum when dissolved in deuterated water (D_2O). The two sets of peaks each correspond to a signal from a hydrogen atom. For chemical reasons, only the two hydrogen atoms (the white spheres) bonded to carbon (black spheres) give rise to NMR signals. The other three undergo rapid exchange with deuterium in the solvent (light grey spheres).

A substantial scalar coupling is only seen between two nuclei if they share significant electron density, implying they are close together in the same molecule. Close means *logically* close, not *spatially* close: the two nuclei must be connected by a small number of chemical bonds. When the coupling between two nuclei is small compared with the difference between their NMR frequencies (weak coupling) the coupling Hamiltonian takes the simple form

$$H = J_{12} (\sigma_{z1} \otimes \sigma_{z2}) \quad (4)$$

where J_{12} , the spin-spin coupling constant, depends on the details of the molecular structure, but is typically in the range 1 Hz to 1 kHz.

Note that the J-coupling Hamiltonian is not applied to the system from outside. Instead, it is an inherent interaction which will be manifested any time the system is allowed to evolve naturally ('free precession'). Taking into account the Zeeman effect and the spin-spin coupling, the total Hamiltonian of a two spin system is

$$H = \frac{1}{2} (\omega_1 (\sigma_{z1} \otimes \mathbb{I}_2) + \omega_2 (\mathbb{I}_1 \otimes \sigma_{z2}) + \frac{1}{2} J_{12} (\sigma_{z1} \otimes \sigma_{z2})) \quad (5)$$

The notation σ_{z1} simply means 'an ordinary σ_z matrix that is here being used in reference to the first spin', and \mathbb{I}_2 similarly means 'an ordinary

identity matrix that pertains to the second spin'. The identity matrices are necessary so that the matrix sums make sense. For a system with n spins, the Hamiltonian will have n Zeeman terms and $\frac{1}{2}n(n-1)$ coupling terms, although some coupling terms may be so small that they can be safely neglected.

2.1 Representations of quantum states

Having determined the Hamiltonian and the allowed energy levels of the system, we now turn to its allowed states. The simplest possible quantum mechanical system consists of one particle whose state is described by the state vector $|\psi\rangle$. Because of quantisation, the particle will have a fixed number of discrete energy levels; we are considering a spin-half system, so we restrict it to two levels which we call $|0\rangle$ and $|1\rangle$. The spin is a quantum object and can therefore exist in a superposition of these levels: a general description of its state is $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are both complex numbers and $\alpha\alpha^* + \beta\beta^* = 1$. Written another way,

$$|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad (6)$$

The density matrix representation of this same state is

$$\rho = |\psi\rangle\langle\psi| = \begin{bmatrix} \alpha\alpha^* & \alpha\beta^* \\ \beta\alpha^* & \beta\beta^* \end{bmatrix} \quad (7)$$

For example, the state $(|0\rangle + |1\rangle)/\sqrt{2}$ can be represented as

$$|\psi\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \text{or} \quad \rho = |\psi\rangle\langle\psi| = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad (8)$$

We can always express ρ in the form $\rho = \frac{1}{2}(\mathbb{I} + s_x\sigma_x + s_y\sigma_y + s_z\sigma_z)$, a linear combination of the Pauli spin matrices, subject to the constraints that s_x , s_y , and s_z are all real and $s_x^2 + s_y^2 + s_z^2 = 1$. Once we have written it in this form, it is natural to picture ρ as a unit vector called the Bloch vector. Figure 2 shows the Bloch representation of $(|0\rangle + |1\rangle)/\sqrt{2}$.

What about multiple particles? One obvious way of representing larger systems is to use tensor products of representations of smaller systems: $\rho = \sigma_1 \otimes \sigma_2 \otimes \cdots \otimes \sigma_n$. For example, a system with two particles, one in an

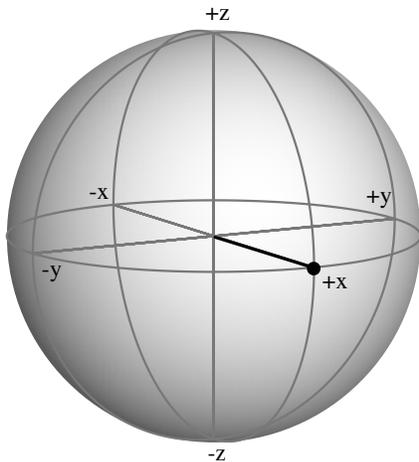


Figure 2: Bloch vector representation of the state $(|0\rangle + |1\rangle)/\sqrt{2}$.

equally weighted superposition and one in the state $|0\rangle$, is represented by

$$\rho = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (9)$$

However, quantum mechanics also allows many multi-particle quantum states which cannot be decomposed into products of constituent states and therefore cannot be represented in this way. These states are known as *inseparable* or *entangled*. For example, consider the two particle state

$$\rho = \frac{1}{2} (|01\rangle + |10\rangle)(\langle 01| + \langle 10|) = \frac{1}{2} \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (10)$$

for which no decomposition into single particle subsystems exists. (The reader can easily verify this by trying to factor the density matrix.) This representational difficulty hints at an underlying physical difficulty. Inseparability implies that the information about the state of any particular particle in an inseparable system is not ‘owned’ by the particle, but is rather shared between two or more particles in the system. Einstein, Podolsky, and Rosen suggested that this implied the particles were ‘communicating’

with one another faster than light, a violation of special relativity [2]. They argued this apparent paradox could be only be resolved by the introduction of hidden variables to restore determinism and locality behind the scenes. In 1964 Bell proved that that a theory with local hidden variables could be experimentally distinguished from one without them and subsequent experiments seem to have ruled out the possibility of local hidden variables [3, 4].

The importance of entanglement to quantum mechanics cannot be underestimated. Schrödinger said ‘he would not call that *one* but rather *the* characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought’ (cited in [5]). Entanglement and the associated paradoxes of non-locality have been dealt with extensively in the literature, but see references [6, 7, 8] for an introductory discussion.

Is there a way of visualising multiparticle states that is analogous to the Bloch representation for a single state? The answer is, unfortunately, no. Separable states can be represented as a set of Bloch spheres, but the Bloch representation breaks down for more general states and no satisfactory alternative is known. Although this is inconvenient, it should not come as too great a surprise: the correlations between entangled states are so fundamentally removed from our everyday understanding of nature that it would be surprising if there *were* a means of representing them with a picture or simple classical analogy.

A second representational issue arises when we try to represent parts of larger systems. Say we would like to represent, however inadequately, the state of the first particle in the system described by equation (10). We can do this by averaging (‘tracing’) over the possible states of the second particle, which gives us

$$\rho = \frac{1}{2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (11)$$

Since equation (10) describes an inseparable state, we know this representation is to some extent wanting, and we would like a way of quantifying the fact that important information is missing. Notice that, while ρ is ordinarily defined to be $|\psi\rangle\langle\psi|$, there is no way of writing equation (11) as a product of state vectors. This is a general property of fragments of entangled systems and we call density matrices of this type *mixed*. The state described by (10) is in fact the maximally mixed state. Formally, a density matrix is mixed if

it cannot be written in the form

$$\rho = |\phi_i\rangle \langle \phi_i| \tag{12}$$

but must instead be represented as a sum of several such outer products. For example, the matrix in (11) can be represented as $\frac{1}{2} |0\rangle \langle 0| + \frac{1}{2} |1\rangle \langle 1|$.

We can visually represent one spin of a multi-spin state in a similarly incomplete way as a mixed Bloch vector. Reconsider figure 1, the cytosine spectrum. The spectral line from each hydrogen nucleus is split, and this suggests that their Bloch vectors are in some sense split as well. The two sub-vectors correspond to the two possible states of the *other*, coupled, hydrogen nucleus. They precess at slightly different frequencies, which is why they have slightly different locations in the spectrum. Since the sub-vectors don't always point in exactly the same direction, their mean length may be less than unity, and so mixed vectors are shorter than their pure counterparts; the most mixed state possible is just a dot at the origin of the Bloch sphere. Figure 3 shows an example set of sub-vectors and mean vector.

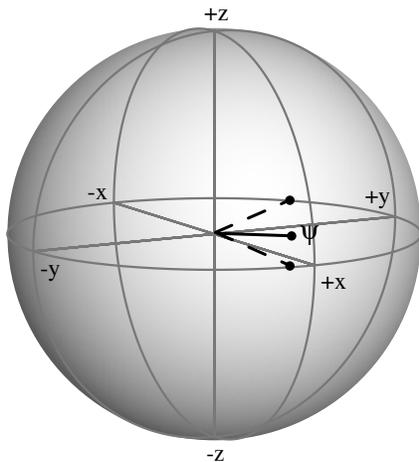


Figure 3: Two ways of representing a mixed state Bloch vector. The mean of the two sub-components (dashed lines) is a vector of less than unit length (solid line).

We have so far described the states of one spin, or of a small assembly of spins, but there will be an astronomical number of spins present in the sample probed in an NMR experiment. A macroscopic volume of liquid is a

complicated and superficially unquantum object, and we would prefer to ignore the bulk of the liquid and investigate a single molecule whose quantum behaviour we can analyse mathematically. Could we do this by using spatial localisation techniques, like those used in magnetic resonance imaging, to pick out individual molecules? Although this technique would be conceptually tidy, it is not practical for several reasons. The signal from any individual molecule is extremely weak, around $2\mu\text{eV}$, so making measurements with reasonable signal to noise would be extremely difficult, although perhaps not absolutely impossible. More seriously, molecules are tiny, and even the spacing between them is very small, so achieving sufficient spatial resolution would be tricky. Even were the spatial resolution fine enough, however, the rapid motion of molecules in solution means that the targeted molecule would soon diffuse out of the selected region.

We are therefore forced to detect the combined signal from all the molecules. Fortunately, however, we can almost entirely neglect the macroscopic nature of the observed system. The mathematical representation of the apparent state of an NMR system is considerably simpler than the representation of its actual state. Taking the number of spins in the system to be P , the full density matrix of the system has 2^P elements, making detailed quantum mechanical calculations hopeless. However, rotational averaging reduces dipole–dipole coupling between spins in neighbouring molecules to a second-order effect, and so the system can be considered in terms of an easily-manipulated reduced density matrix of size 2^n , where n is the number of spins in any one molecule. Instead of being one monstrously complicated system, the molecules in solution behave like an ensemble of small discrete systems. They are isolated from one another by dipole averaging, and from the external environment by simple surface to volume considerations of the liquid. Nonetheless, the macroscopic and ensemble aspects of the system are not entirely without implications; we will discuss some of the more interesting ones later.

2.2 Representations of quantum operators

Like quantum states, quantum operations are conveniently represented as matrices, given by

$$\mathbf{U} = \exp\left(\int_0^\tau -i\mathbf{H} dt\right) \quad (13)$$

where \mathbf{U} is the propagator, \mathbf{H} is the operation’s Hamiltonian, and \hbar is set to unity. Although in principle each operation has a unique matrix repre-

sentation, global phase factors are meaningless in NMR. The propagators of Hamiltonians which differ only by a multiple of the identity matrix differ only by a global phase factor, and so these Hamiltonians are in practice experimentally indistinguishable. All propagators obey one constraint, which is that they must be *unitary*. This is an expression of the quantum mechanical axiom that all quantum operations are reversible.

Any unitary operation on a spin can be described a rotation about some axis. In the case of RF pulses, this correspondence is particularly natural because a given rotation angle corresponds in practice to a pulse duration, $\theta = \tau\omega$, the angle of the axis in the xy -plane corresponds to the phase of the RF field, and the declination of the axis from the xy -plane (if any) is determined by the field's detuning. For example, we would write the operator given by application of the Hamiltonian $H = \omega\sigma_x$ for a period of $\tau = \pi/\omega$ as

$$\mathbf{U} = \exp\left(\int_0^{\frac{\pi}{\omega}} -\frac{1}{2}i\omega\sigma_x dt\right) = \begin{bmatrix} 0 & -i \\ -i & 0 \end{bmatrix} = 180_x \quad (14)$$

Operators that arise from the J-coupling affect the relations between spins and so they cannot be conveniently described as rotations; instead, they are simply described in terms of time periods. The operator for the J-coupling has the general form

$$\mathbf{U} = \begin{bmatrix} \exp\left(-\frac{1}{4}iJt\right) & 0 & 0 & 0 \\ 0 & \exp\left(\frac{1}{4}iJt\right) & 0 & 0 \\ 0 & 0 & \exp\left(\frac{1}{4}iJt\right) & 0 \\ 0 & 0 & 0 & \exp\left(-\frac{1}{4}iJt\right) \end{bmatrix} \quad (15)$$

for a coupling of strength J over a time t . Since we generally want Jt to be some dimensionless fraction, we describe the operator as a time in units of $1/J$, for example, $\pi/(4J)$.

3 Basics of quantum technologies

Quantum effects are manifested in a wide array of physical systems, but not every system in which quantumness is observable ought to be considered an example of a quantum technology. Technologies are sets of tools, techniques, and concepts which together provide a means of manipulating a system into performing a task. This task cannot simply be the elucidation of the techniques and concepts involved—while a pencil can be described as

a simple technology for writing, it cannot be sensibly described as a technology for learning about pencils, or, worse yet, as a technology for being a pencil. Similarly, a quantum technology must achieve some goal beyond the manipulation and observation of the system's own quantum nature. Despite this qualification, manipulation and observation of quantumness are important technical pre-requisites of any quantum technology. For the class of technologies we are considering, those which relate to quantum information processing, two kinds of manipulation are required: the ability to control the evolution of the system (Hamiltonian sculpting), and the ability to control the starting state of the system (initialization). Observation is a necessary final step, so that we can read out the answer of our computation or the transmitted information. Both readout and initialization are complicated by the ensemble nature of NMR in ways that are often unfortunate but also unexpectedly interesting.

3.1 Ensemble quantum technologies

It is surprising that a system so tangibly macroscopic as liquid in a test tube should be a viable quantum technology. The bulk nature of NMR does in fact have several repercussions at the quantum level, some of which are more fundamental than others. These implications have experimental interest in themselves, and they also force us to think deeply about some of the finer points of quantum mechanics. The study of the ensemble aspects of NMR is controversial and has inspired a number of theoretical investigations into previously unexplored aspects of quantum mechanics. For a general discussion of the issues involved, see reference [9].

We instinctively associate quantum mechanical behaviour with the very small. Quantum mechanics was not developed until technology advanced to the point where scientists were able to investigate the tiniest constituents of matter directly. The very small, it seemed, had a variety of baffling and essentially disturbing behaviours that weren't manifested by the larger objects that had previously been used as the benchmark of normality. However, once scientists had adjusted to quantum jumps, superpositions, and apparent action at a distance, they stopped asking why small things didn't act like big things, and started wondering why big things didn't act like small things. The general conclusion was that macroscopic systems decohere rapidly, and that this dilution of the quantumness is responsible for the apparently 'normal' behaviour.

Decoherence is an important concept in the quantum mechanics of the macroscopic. A system becomes decohered if it interacts with its environ-

ment, since this interaction causes information about the system to be spread through the environment; when the environment is traced out this information is lost. Having decohered corresponds to being in a mixed state; a completely decohered system is in the maximally mixed state—that is, the density matrix is a multiple of the identity matrix.

When we consider a test tube of liquid in a spectrometer, we mentally divide the state of each molecule into two parts: a large undetectable maximally mixed portion, and a small detectable (‘effective’) portion, ρ_0 :

$$\rho = (1 - \epsilon)\mathbb{I}/N + \epsilon\rho_0. \tag{16}$$

where $N = 2^n$ and n is the number of qubits in the molecule. Coherent states can be temporarily excited only in the ρ_0 part. These excited states will themselves decohere (‘relax’) by interacting with one another and the environment. There are a number of factors that determine how long this relaxation takes, including the previously neglected dipole–dipole interaction. Luckily, however, decoherence of the effective state is not a critical consideration in NMR since the relaxation times are long enough that we can conduct complex experiments without crippling losses to decoherence. The mental division into effective state and undetectable maximally mixed state allows us to largely neglect the ensemble nature of NMR; nonetheless, the presence of many copies of ρ_0 does present problems, particularly in the initialization and readout of quantum states.

3.2 Observation

Even the small, undecohered, part of an NMR experiment contains multiple copies of the effective state, and this ensemble nature has an immediate practical implication. It is an axiom of quantum mechanics that a measurement of a state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ will return, at random, either 0 with probability $|\alpha|^2$ or 1 with probability $|\beta|^2$ (assuming the measurement was made in the computational basis). The measurement will project $|\psi\rangle$ into one of $|0\rangle$ or $|1\rangle$; if we got an answer of 0, $|\psi\rangle = |0\rangle$ after the measurement; if the measurement gave 1, $|\psi\rangle$ will be turned into $|\psi\rangle = |1\rangle$. The measurement returns only partial information about the true state and collapses it in the process. In an NMR measurement, on the other hand, a measurement of an effective state $|\psi\rangle$ is in fact a measurement of a whole set of $|\psi\rangle$ s; it can only return mean values of α and β , and the superposition will not be destroyed by the measurement. Being able to measure the average popula-

tion non-destructively seems like a wonderful advantage, but many quantum information processing techniques rely on projective measurement.

These difficulties aside, the mechanism of observation in an NMR system is very simple. Spins in a superposition will precess around the z -axis, while spins with no component in the xy -plane are in an eigenstate and will not precess. Ensembles of precessing spins give rise to a precessing net magnetic moment which, by Faraday's law, induces a voltage in a receiver coil. We can measure the magnitude and phase of the detected signal and draw conclusions about the declination of the net magnetic moment from the z -axis and its phase. The magnitude of the observed signal will decrease as the system relaxes to equilibrium, and so this measurement process is called 'free-induction decay'. It is usually necessary to apply a 90° pulse to the system before measurement, since this has the effect of knocking z -magnetisation (an eigenstate) down into the xy -plane where it can be observed.

3.3 Hamiltonian sculpting

The power and utility of NMR as a quantum technology stems from the malleability of the NMR Hamiltonian. Terms can be modified or suppressed entirely, almost at will. This allows the experimenter precise control over the interactions within the system and the response of the system to external probing. The terms in the NMR Hamiltonian, given in equation (5), are inherent to the system and cannot be changed, but the average Hamiltonian over some period depends on the inherent Hamiltonian equation (5) and also whatever RF field Hamiltonians may have been applied. Its form may be considerably simpler than that of the inherent Hamiltonian. The literature dealing with the various possible manipulations of the average Hamiltonian is vast, but we will briefly describe some of the more important techniques.

The most basic and useful manipulation is the spin echo. Spin echoes can, among other things, undo the effects of spatial inhomogeneities in the magnetic or RF field, suppress specific interactions, and refocus the effects of chemical shifts. Spin echoes can be understood by analogy with a budget tape-deck whose fast-forward and rewind buttons have both broken. In order to return the tape to its initial position after playing it for some time τ , it is necessary to take the tape out, flip it over, and then play it again for a time τ . The important idea is that the action of time can be in some sense be reversed by flipping the system. The spin echo is also analogous to a colony of lemurs who leave their nest in a group to go forage in the morning. Some leap quickly through the trees, while others scurry more slowly on the

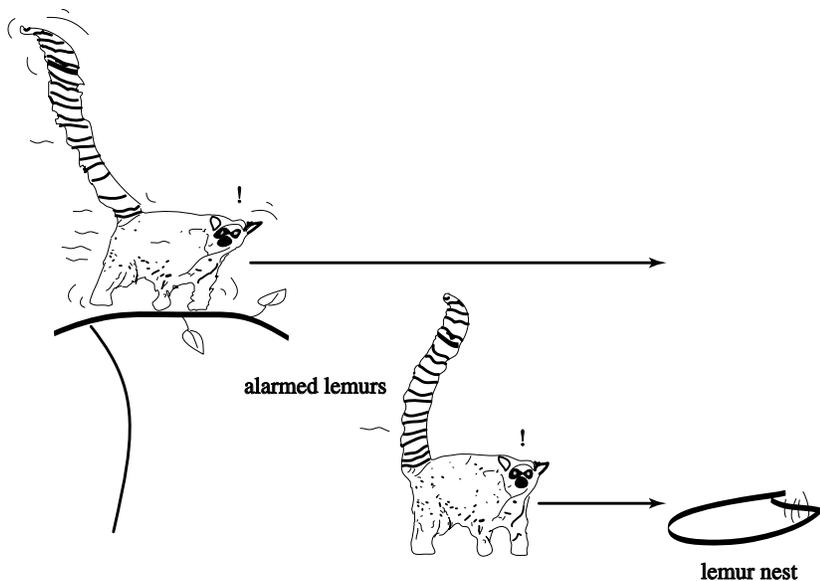


Figure 4: When lemurs spot a fossa in the distance, they return to their nest. Lemurs in the trees are further from the nest but are able to jump rapidly from branch to branch, returning home at the exact same time as the lemurs on the ground.

ground, so the lemurs are soon dispersed through the forest. When one of the lemurs spots a predatory fossa in the distance and raises an alarm, all the lemurs turn around in dismay and head back to the nest. The lemurs in the trees make rapid progress, but they're further out from the nest than the slow-moving ones on the ground, and so the lemurs necessarily all arrive back at the nest in a clump at the same time. This process is illustrated in figure 4. Viewed thermodynamically, this restoration of the dispersed and disordered lemur population into a neat group might seem like an unphysical violation of the second law of thermodynamics, but this is not the case; the unitarity of the dispersal process means that there is no increase in entropy associated with it, and there is therefore no decrease in entropy associated with its reversal. While they are most commonly used in NMR, spin echoes are by no means unique to NMR, and they may find wide application in future quantum technologies.

The normal course of events in an NMR experiment is that an RF pulse

is applied to the sample to excite it, signal appears, and then the signal gradually disappears as the sample decoheres back to thermal equilibrium. It is certainly not normal for signal to appear when no pulse has been applied, but this is exactly what happens in a spin echo experiment; a 180° pulse is applied after the signal has disappeared, nothing happens for some time, and then signal spontaneously reappears. This resurrection of a decayed signal was initially quite startling, not least to the discoverer of the spin echo, Erwin Hahn. The reason the spin echo works is that much of what appears to be decoherence in NMR is in fact dephasing across the sample. Dephasing is unitary and can therefore be reversed with an appropriate pulse sequence.

Consider the two-spin Hamiltonian, (5), again. In addition to reversing dephasing error terms, spin echoes can be used to suppress any two of its three terms, or, in longer combinations, any one of its terms. Selectively applying a 180° pulse to the first spin undoes its evolution and the effects of the J-coupling but allows the second spin to evolve, selectively refocussing the second spin similarly kills the $\omega_2\sigma_{z2}$ term and the $J_{12}(\sigma_{z1} \otimes \sigma_{z2})$ term but keeps the $\omega_1\sigma_{z1}$ term, while simultaneously refocussing both spins deletes the $\omega_1\sigma_{z1}$ and $\omega_2\sigma_{z2}$ terms but keeps the J-coupling term. Refocussing the first and second spins at different times removes the J-coupling term but merely scales down both evolution terms. For systems with more spins, the pulse sequences for creating arbitrary scaled down Hamiltonians are longer but not unreasonably so.

A second, less versatile, sculpting technique is spin decoupling. J-couplings to a specific nucleus can be eliminated by continuously applying a field at the resonant frequency of the offending nucleus. This causes rapid transitions between its $|0\rangle$ and $|1\rangle$ states, which causes rapid flips between the sub-vectors of the coupled nuclei. If these flips are sufficiently rapid, it will be impossible to distinguish the two components of the Bloch vector and no splitting will be observed in the spectrum. As a side-effect, all signal from the target nucleus is obliterated, which is sometimes undesirable.

3.4 Controlling the starting state

Hamiltonian sculpting allows us to precisely control the evolution of an effective quantum state, and so it seems that driving a system into an effective state of our choosing ought to be a trivial problem. Unfortunately, the operations we can achieve by Hamiltonian sculpting are all reversible, and this severely limits the states we can achieve. As initialization requires that the system be placed in some state $|\psi\rangle$, independent of its state before the be-

gining of the initialization process, it manifestly cannot be achieved by any reversible process. Initialization schemes must therefore be quite different from ordinary operations.

In many quantum technologies, initialization is achieved by cooling to the ground state. This is not a practical approach for NMR, as the energy gaps involved are tiny compared with the Boltzmann factor at any reasonable temperature. Instead, three alternate experimentally accessible processes with non-unitary effects are commonly exploited in NMR. These are relaxation, magnetic field gradients, and phase cycling. The first of these is the simplest, since it requires no special equipment or post-processing, and is, in fact, inevitable. Relaxation brings the system into a reproducible thermal equilibrium which is an adequate starting state for many purposes. The second technique, gradient crushing, relies on the fact that the sample forms a macroscopic ensemble; by applying Hamiltonians which vary over the sample the *ensemble averaged* evolution can be non-unitary. This is most commonly achieved by momentarily destroying the spatial homogeneity of the main magnetic field, but similar effects can be achieved using spatially inhomogeneous RF fields. The last approach relies on combining the results of several subtly different NMR experiments by post-processing; as this is done by classical methods, such processing is not confined to unitary transformations. In conventional NMR experiments this is referred to as phase cycling, and plays a central role in many sequences, although for many purposes it has now been replaced by the use of gradient techniques. In practice spectroscopic techniques normally rely on relaxation to initialize the system, while quantum computation usually requires the more sophisticated gradient crushes or phase cycling.

Because initialization by relaxation is so effortless, it is easy to underestimate its importance as an initialization technique. The reproducibility of NMR spectroscopy relies on the well-defined nature of the relaxed state. More critically, relaxation into thermal equilibrium is necessary before a net magnetic moment arises. The earliest attempts at an NMR experiment, by Dutch physicist C. J. Gorter, famously failed because he happened to choose a sample with an unusually long relaxation time, which meant the system did not achieve its initialized state before measurements were attempted.

While all of the initialization processes described allow for the manipulation of the effective state, none, even the most non-unitary, allows for the manipulation or elimination of the maximally mixed portion of the NMR state. The silent presence of this large decohered volume of solution presents some fundamental issues. A widely used criterion for quantumness (in the sense of ‘provable non-classicalness’) is separability. A non-separable state

clearly has some entanglement (the converse assertion is mathematically irrefutable but nonetheless more philosophically controversial). By definition, the effective state of an NMR system exhibiting entanglement is non-separable, but the state of the entire volume of solution has been shown to be separable for low polarisation, implying a ‘classical’ description of the state exists [10]. It has not been shown, however, that transformations between these ‘classical’ states can be described classically [11]. This forces a reconsideration of the notion that entanglement is the defining feature of quantum mechanics, since non-entangled systems can apparently display non-classical behaviour. NMR has made it clear that entanglement is as yet poorly understood within the quantum mechanics community, and mixed state entanglement is currently a very active research area [10, 11, 12].

4 Quantum computation

Hamiltonian sculpting gives us the tools needed to manipulate quantum spins and the relationships between them. Currently, the most intriguing use of this power is as a means of implementing quantum computers. Quantum computers are information processing devices which operate by and exploit the laws of quantum mechanics, potentially allowing them to solve problems which are intractable using classical computers [13]. They can implement new types of quantum mechanical algorithms which are far more computationally efficient than their classical equivalents. The first of these quantum algorithms to be described was Deutsch’s algorithm [14, 15], while more recent algorithms include Grover’s search algorithm [16] and Shor’s factoring algorithm [17]. Shor’s factoring algorithm is particularly relevant because its implementation would compromise the security of widely-used public-key cryptosystems such as the RSA system [18] developed by Rivest, Shamir, and Adleman.

Quantum computers offer huge potential, but present technology is largely incapable of realising this potential because quantum computers require assemblies of quantum spins and these are difficult to prepare, isolate, manipulate, and observe. Although there have been many proposals for systems which could be used to implement quantum computers, technical difficulties have so far prevented much experimental progress. To date NMR quantum computers provide the only working means of implementing full quantum algorithms. Although it is widely accepted that NMR in its current form is not a viable means of implementing large quantum information processing devices, many interesting small-scale quantum algorithms and protocols

have been implemented to date.

4.1 Quantum algorithms

Several two- and three-qubit algorithms have been implemented using NMR in the past few years, including Deutsch's algorithm [19, 20] for exposing double-sided coins in a single glance and algorithms [21, 22, 23] for finding and counting needles in haystacks. (It goes without saying that the coins and haystacks are strictly metaphorical.) Some interesting quantum protocols have also been demonstrated, including quantum state teleportation [24] and a two bit phase error correction [25] and detection [26] codes.

The first quantum algorithm to be demonstrated was a solution of Deutsch's problem. In its simplest form, this problem concerns the analysis of single bit binary functions. There are four possible functions which take one bit as input and return a single bit as a result. Two of these return the same value (zero or one) regardless of the input (constant functions), while one returns the input and one flips the input (balanced functions). Given an unknown example of one of these four functions, we would like to know if the function is one of the constant ones or one of the balanced ones. Ordinarily, we must completely characterise the function by trying it out for both inputs in order to categorise it. This is a frustrating exercise because we have to evaluate the function twice and get two bits of information out, but we only actually want one bit of information. We would prefer to evaluate the function once and get out just the information we want. The problem is analogous to deciding if a tossed coin is legitimate. If someone gives us a coin for a game of chance, we would like to establish first of all that the coin is not double-sided. Our opponent resents this officiousness, and so they charge us ten pounds for each side of the coin we look at. Obviously, in the classical case we must inspect first one side of the coin and then the other side and pay twenty pounds, but that gives us much more information than we wanted because we then know whether the coin has two heads, two tails, a head and a tail, or a tail and a head (although we'd be hard-pressed to distinguish those last two cases in a normal coin). Deutsch worked out a quantum algorithm which works by feeding the function a superposition of zero and one and then interfering the output with itself to get the answer.

Another algorithm that has been implemented in two and three qubits is Grover's search algorithm. Grover's algorithm has less historical importance than Deutsch's problem, but it is potentially much more useful. Grover's algorithm allows a state satisfying some condition (for example, 'is this a needle?') to be quickly picked out of a large sample (for example, a

haystack). Like Deutsch's algorithm, Grover's search algorithm achieves a non-classical efficiency boost by simultaneously processing a superposition of states. We would like to find a state that meets some criteria, so we alternately feed a superposition to the function to label the satisfying state and apply a series of gates to drive the superposition into the labelled state. A classical search over a domain of size N for one of k satisfying elements would require about $\frac{1}{2}(N/k)$ function evaluations, whereas the quantum search requires only $O(\sqrt{N/k})$.

The quadratic speedup in searching is exciting and has several significant applications, but the speedup and potential utility pale when Grover's algorithm is compared to Shor's factoring algorithm, which is exponentially faster than the best known classical algorithm. This exponential speedup is significant because public-key cryptosystems like RSA and composite systems based on it, like PGP (Pretty Good Privacy), rely on the presumed exponential difficulty of factoring for their security. A related exponentially-faster algorithm exists for solving the discrete logarithm problem, which is the basis of the DSS (Digital Signature Standard) cryptosystem. Public-key cryptosystems are absolutely ubiquitous in modern communication, and so were these cracking algorithms to be realised on large computers, the economic impact would be enormous. Although one element of Shor's algorithm has been demonstrated on an NMR quantum computer [27], it seems that any full implementation is out of reach for the moment.

Richard Feynman is often credited with first highlighting the power of quantum information processing in 1982 by observing that quantum systems were impossible to simulate efficiently using classical means [28]. One corollary of this is the well-known one that quantum systems might have remarkable information processing capabilities, which has led to the proliferation of quantum algorithms. A second and less well-studied corollary is that quantum systems could efficiently simulate other quantum systems [29]. This second corollary was recently demonstrated by a group who performed an efficient simulation of a truncated quantum harmonic oscillator using NMR techniques [30].

Like classical computers, quantum computers have four basic elements: (qu)bits, logic gates, a means of initialization, and a means of reading out the answer. Once these elements are in place, the implementation of any algorithm is relatively straightforward. We have already touched on initialization and measurement in NMR, but we will revisit them in a quantum computation context as gates, as well as discussing the implementation of qubits and ordinary logic gates.

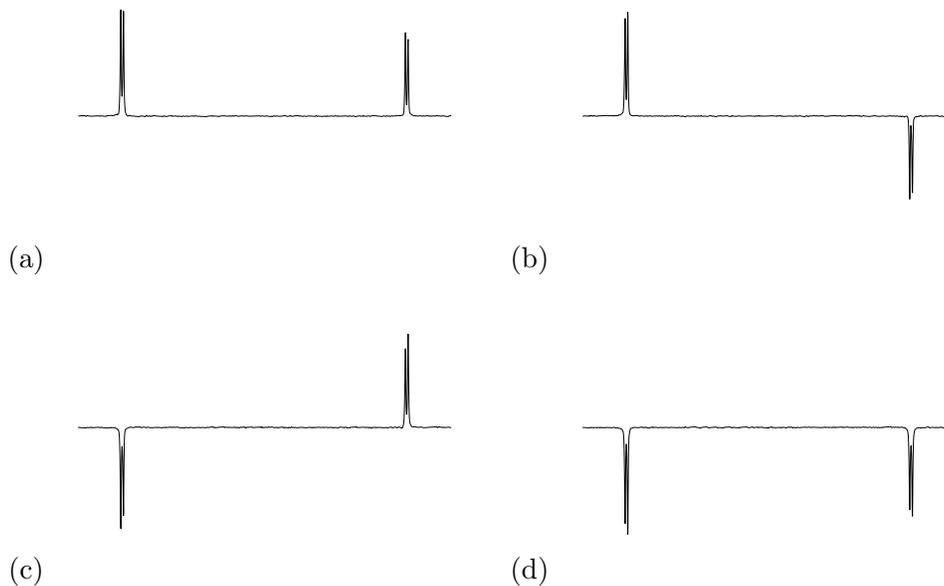


Figure 5: Spectra of cytosine with the qubits in states (a) $|00\rangle$, (b) $|01\rangle$, (c) $|10\rangle$, and (d) $|11\rangle$. Note that the state $|00\rangle$ is similar but not identical to the thermal equilibrium shown in figure 1; the asymmetric peaks here are due to experimental imperfections.

4.2 Bits and qubits

Qubits are the quantum analogues of bits. Unlike bits, qubits can exist in superpositions of states or entangled with one another. Therefore, while classical computers are limited to performing operations on one n -bit state at a time, a quantum computer which handles n qubits can process a simultaneous superposition of 2^n possible states. Qubits are generally individual spin $\frac{1}{2}$ particles like photons or atoms. Consider figure 1, the cytosine spectrum, again. Each group of peaks corresponds to a qubit. In the spectrum shown, the qubits are in thermal equilibrium; figure 5 shows spectra of the same molecule with the qubits in the states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$.

Two-qubit computers are theoretically interesting, but we would naturally like to build much larger computers. In this we are limited by several factors [31], of which two are major issues; the first is the difficulty of engineering molecules with a sufficient number of spin-half nuclei appropriately coupled to one another, and the second is the difficulty of resolving the sig-

nals from many qubits sharing a limited bandwidth. For example, the range of Larmor frequencies found for ^1H nuclei in simple organic compounds, and working at a ^1H frequency of about 500 MHz, is only about 5000 Hz, and this limited frequency range can soon ‘fill up’. Extra bandwidth can be found by using a variety of spin half nuclei, since the NMR transition frequencies of different nuclei are very different, but this approach cannot be continued indefinitely, as the number of suitable nuclei is small: the only obvious candidates are ^1H , ^{13}C , ^{15}N , ^{19}F and ^{31}P .

4.3 Logic gates

A logic gate is a controlled interaction between a targeted spin and its environment. A one qubit gate is an interaction between the spin and the outside environment, and a two qubit gate is an interaction between a spin, its environment, and another spin [32]. A one-qubit gate can be represented as a simple rotation of the Bloch vector, while a two-qubit gate affects the relationships between spins and so its effects cannot be quite so easily visualised. In the case of NMR, an applied RF field mediates the interactions responsible for one qubit gates while the combined effects of the RF field and J-coupling are responsible for implementing two qubit gates such as the controlled-NOT gate.

Although there are many possible logic gates, any algorithm can be efficiently implemented using a limited repertoire of gates. In the classical case, any logic gate can be built efficiently from the two-bit NAND gate (along with implicit SWAP and CLONE gates which interchange and duplicate bits, respectively), so any algorithm can be implemented using only NANDs. In the quantum case, almost any gate is universal, although it is often more convenient to use a larger set of gates. A convenient set usually quoted is one qubit rotations plus the two qubit controlled-NOT gate, but an equivalent adequate set is one-qubit rotations and a two qubit controlled phase shift.

4.3.1 One qubit gates

We describe below the operation of some of the more common one and two qubit gates. The simplest non-trivial gate is the NOT gate,

$$\text{---} \bigoplus \text{---} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad (17)$$

Up to an (irrelevant) global phase factor, this is equivalent to the matrix given in (14), and so we can implement NOT gates using 180_x pulses. The NOT gate can be understood classically, but there are several commonly used quantum gates which have no classical equivalents. One of these is the square-root-of-NOT gate,

$$\text{---} \bigcirc \sqrt{\quad} \text{---} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -i \\ -i & 1 \end{bmatrix} \quad (18)$$

The square-root-of-NOT gate has the property

$$\text{---} \bigcirc \sqrt{\quad} \text{---} \text{---} \bigcirc \sqrt{\quad} \text{---} = \text{---} \bigoplus \text{---} \quad (19)$$

It is easy to see that such a gate can be implemented using a 90_x RF pulse. Another very common gate is the Hadamard gate, which, like the square-root-of-NOT, takes eigenstates to superpositions, but is also self inverse. It is described by

$$\text{---} \boxed{\text{H}} \text{---} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (20)$$

The Hadamard gate can be implemented in NMR using a pulse which is deliberately applied off-resonance to give $\exp(-i\pi(\sigma_x + \sigma_z)/2\sqrt{2})$ or a two pulse pair, $90_y \cdot 180_x$. It is often experimentally simpler to use the pseudo-Hadamard gate, 90_y ,

$$\text{---} \boxed{\text{h}} \text{---} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \quad (21)$$

which takes eigenstates to uniform superpositions, but is not self-inverse. The inverse pseudo-Hadamard gate is implemented by the rotation 90_{-y} . Another useful operation that has no classical analogue is the phase shift operation,

$$\text{---} \bigcirc \frac{\pi}{2} z \text{---} = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} \quad (22)$$

This gate rotates the Bloch vector about the z -axis. It is rarely used in theoretical descriptions of algorithms but it can be useful in their adaptation to NMR.

4.3.2 Two qubit gates

Although there are in principle a wide variety of two-bit gates, all interesting two bit gates implement conditional dynamics where the state of one qubit is used to control changes made to the state of a second qubit. The controlled-NOT is a particularly useful two-qubit gate because it is comparable to a classical exclusive-OR gate and, along with a complete set of one-qubit gates, it comprises a set of adequate gates. In matrix-form, the controlled-NOT gate is given by

$$\begin{array}{c} \bullet \\ | \\ \oplus \end{array} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (23)$$

For the purposes of NMR computation, it is simplest to replace the controlled-NOT with Hadamard transforms and a symmetric controlled phase shift, and then replace the Hadamard transforms by pseudo-Hadamard gates, which gives the following circuit:

$$\begin{array}{c} \bullet \\ | \\ \pi_z \\ | \\ \bullet \\ \boxed{h^{-1}} \quad \boxed{h} \end{array} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (24)$$

Here

$$\begin{array}{c} \bullet \\ | \\ \pi_z \\ | \\ \bullet \end{array} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (25)$$

The controlled phase shift is less analogous to classical gates like the exclusive-OR than the controlled-NOT is, but it is a more natural algorithmic building block in an NMR context. The operator of the controlled phase shift is equivalent (within a global phase shift) to $\exp(-i\frac{\pi}{4}(\sigma_z^1 + \sigma_z^2 - \sigma_z^1\sigma_z^2))$. The transformation described by equation (25) can be implemented (in a two-qubit system) by allowing the system to evolve under the influence of

its Hamiltonian, with some tweaking, or it can be implemented in a more controlled way by using the identities $\theta_z = 90_{-x} \cdot \theta_y \cdot 90_x$ to replace the z -rotations by operations which can be implemented using RF pulses.

Controlled phase shifts and controlled-NOTs are important gates because they spread information about the control-bit through a larger system, potentially creating entanglement. They can also be used as building blocks of routines which move information around in a system without creating entanglement. For example, consider the SWAP network:

$$(26)$$

Swapping becomes important in networks where not every spin is coupled to all the other spins; in this case it is sometimes necessary to swap distant qubits through the network to bring them near each other before a controlled-NOT can be applied between them. Note the distinction here between spins (or nuclei) and qubits (or states); while we loosely speak of spins and qubits as though they were the same thing, spins are physical objects and qubits are abstract quanta of information. Therefore we can move qubits through a molecule—without ever changing the positions of the actual spins—by swapping information between spins and then mentally adjusting which qubits we map onto which spins. Several techniques in conventional NMR spectroscopy are related to the swap sequence.

4.3.3 Measurement gates

The ensemble nature of NMR means that there are actually two types of measurements we can make in the course of implementing an algorithm. The first, more obvious, measurement is a standard NMR ensemble measurement. This usually takes places at the end of an algorithm, and it tells us the answer (or, technically speaking, the average state of the measured qubit). We can represent such a measurement in a logic network as a period of observable free precession,

$$(27)$$

There are many circumstances, however, under which we want to conduct a measurement in the middle of the logic network. Some of the most exciting algorithms, such as Grover’s fast search algorithm and Shor’s factoring algorithm, as well as popular protocols like teleportation and quantum error-correction, require measurements midway through. What happens in these algorithms and protocols is that we measure the state of the system, an eigenvalue of the measurement operator is returned, and we take some important algorithmic action based on the result of the measurement. But what exactly is a measurement, in this context?

Ordinarily, a measurement is the extraction of some information about the microscopic system out to a macroscopic level. This process can be described quantum mechanically as the entanglement of the system to its environment. The quantum mechanical description does not specify the size of the environment, and so we are perfectly entitled to entangle our state to a one qubit ‘environment’, thereby completely bypassing the macroscopic part of the process. We achieve this entanglement by using a controlled-NOT gate. Therefore, the following gate, previously described as a controlled-NOT, is also, in fact, the ‘measurement gate’:



Under normal circumstances, the measurement would be followed by some action based on the result, and these conditional dynamics can also be built directly into the algorithm. But wait—while we might have achieved the entanglement-and-control part of the measurement process within our algorithm, we certainly have not managed to collapse the wave function to an eigenstate. As it happens, this is absolutely unimportant, because quantum mechanics is linear. If our algorithm dictates we take action a if a measurement returns $|0\rangle$ and action b if a measurement returns $|1\rangle$, we can build this decision process right into the logic circuit so that the appropriate action is taken on the relevant part of a superposition. Processes analogous to macroscopic measurement can also be effected using decoherence to destroy off-diagonal elements in the density matrix [24].

4.3.4 Input gates

Initialization is the process of placing a quantum computer in some well defined initial state. As was already discussed, this can be difficult in en-

semble systems. Although the thermal equilibrium state is adequate for many spectroscopic applications, a quantum computer really requires its qubits to be in the pure state $|\mathbf{0}\rangle = |000\dots 0\rangle$ prior to beginning the computation. For this reason conventional liquid state NMR was long ruled out as a practical technology for quantum computation. In 1996, however, Cory, Fahmy, and Havel pointed out that an effective (in the sense used in section 3.1) pure state was sufficient, and described a procedure for making effective, or pseudo-pure, ground states [33]. For the simplest possible quantum computer, comprising a single qubit, this process is trivial, as the thermal equilibrium density matrix has exactly the desired form, but with larger systems the situation is more complicated. For a system of n spin-half nuclei, assumed to all be of the same nuclear species (a homonuclear spin system), the 2^n eigenstates will be distributed across an evenly spaced ladder of $n + 1$ groups of energy levels, with the number of (nearly degenerate) states within each group given by Pascal's triangle, and the population of each state determined by the Boltzmann equation. Transforming this complex state into the pseudo-pure ground state requires a non-unitary process. The original approach of Cory *et al.* is based on the use of magnetic field gradients and is in many ways the most satisfying, but an alternative 'temporal averaging' scheme based on post processing [34] has also proved extremely popular.

Although in general gradient crushes or phase cycling are required to produce acceptable ground states for quantum computation, there does exist one elegant initialization scheme which requires only relaxation processes [35]. In this process a pure state is produced by *logically* selecting an appropriate subset of the thermal equilibrium state. While the thermal equilibrium spin density matrix for an n spin system ($n > 1$) is not unitarily related to the state $|\mathbf{0}\rangle = |000\dots 0\rangle$, subsets of energy levels can be chosen which *do* have the correct pattern of populations; the computation is then performed within this subset of states. Although theoretically pleasing this scheme appears to be more complex than the other approaches, and only two experimental implementations have been reported [36, 37]. Regardless of the technique used to achieve it, a pseudo-pure ground state input is represented in a logic network as

$$|\mathbf{0}\rangle \text{ ————— } \quad (29)$$

The difficulty of initialization in NMR has sparked theoretical research into the previously uncontested assumption that useful computations must take as their starting state a pure (or at the very least, pseudo-pure) state

that has been initialized to some value. In classical computation the only available states are pure, and initialization is simple, and so there has never been any reason to doubt the truth of the assumption. In NMR, on the other hand, mixed states are much more accessible than pure states. It is tempting to ask if they can themselves be a computational resource even though the question seems ridiculous, since maximally mixed states cannot be detected and are impervious to unitary transformations. Surprisingly, it has been shown recently that even maximally mixed states, in combination with one pure qubit, actually *can* be used to perform quantum algorithms with better-than-classical power [12, 38]. One of the dogmas of quantum information processing, first stated by Ekert and Josza, is that its power stems from the presence of entanglement [39], but the research into mixed states prompted by NMR quantum computation is forcing a serious reconsideration of the Ekert–Josza dogma.

5 Quantum spectroscopy

One reason why the implementation of quantum computers using NMR has been able to outstrip implementations using any other technology is that many of the techniques required were already well known in the NMR spectroscopy community. In fact, there are two major groups of conventional NMR techniques which exploit quantum effects. Both use similar routines based on partial swaps but with different aims. The first group, polarisation transfer sequences like INEPT (Insensitive Nuclei Enhanced by Polarisation Transfer) and variants, involve creating correlations between two spins in order to extract information about the first spin by exploiting the second, more measurement-friendly, spin. The second class, COSY (CORrelation SpectroscopY) sequences and variants, create correlations between a large number of spins and then use the presence or absence of these correlations to infer structural information about a molecule. (Both of these techniques highlight the fondness for whimsical acronyms in the NMR community.) INEPT sequences are useful if the nucleus to be measured has a small gyromagnetic ratio, since these nuclei are less sensitive. COSY sequences, on the other hand, are primarily useful in untangling complex spectra.

5.1 Coherences

Although entanglement is well-known in the NMR spectroscopy community, it is rarely discussed in the terms of Einstein, Podolsky, and Rosen. Instead it is described using a technical language whose major lexical ingredient is

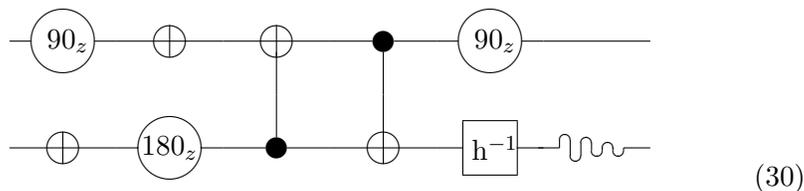
‘orders of coherence’. The reason a macroscopic system such as liquid in a test tube is able to display quantum behaviour is that there is coherence between the various pseudo-pure states. What this means in practical terms is that their Bloch vectors point in the same direction; were this not the case, some (or all) components would go to zero in any measurement of the average state and so significant information would be lost. While this coherence is both a prerequisite for any useful manipulation of the system and a fairly natural consequence of the fact that the applied RF field is phase-homogenous across the sample, it is still a special enough condition that it is given the name quantum coherence. Technically, unentangled qubits lying in the xy -plane (in the state $|0\rangle + \exp(i\phi)|1\rangle$) exhibit single quantum coherence. Two-party entangled states of the form $(|01\rangle + |10\rangle)/\sqrt{2}$, which would be called an Eistein–Podolsky–Rosen (EPR) pair by physicists, are known as zero-quantum coherences. A double quantum coherence describes another entangled state very similar to the two-party ‘cat-state’ $(|00\rangle + |11\rangle)/\sqrt{2}$ (named after the Schrödinger cat), while a triple quantum-coherence describes the three-party cat-state $(|000\rangle + |111\rangle)/\sqrt{2}$, and so on for higher orders of coherence.

5.2 Polarisation transfer

One of the uses of higher order coherences in NMR spectroscopy is to increase the sensitivity of the spectroscopy of spins with low gyromagnetic ratios. These are issues particularly with nuclei like ^{13}C or ^{15}N ; their low gyromagnetic ratio means the energy levels are closely spaced, which in turn means the excess population in the ground state is quite small. We can increase the polarisation, and consequently the observable signal, by transferring polarisation from relatively highly-polarised ^1H spins. This gives us a potential polarisation gain of γ_H/γ_I , where γ_I is the gyromagnetic ratio of the weakly polarised (insensitive) species.

The most basic polarisation transfer sequence that exploits quantum coupling is the INEPT sequence, but a neater variation known as refocused-INEPT is also used. These sequences are essentially partial swaps, in which the state of the hydrogen is transferred to the insensitive nucleus without transferring its state back to the hydrogen. We can drop the transfer-back part of the process because only the spectrum of the insensitive nucleus is of interest.

Written as a logic circuit, refocused-INEPT is



The upper qubit is the highly-polarised one, and the lower qubit is the poorly-polarised one. At the end of the network, we measure only the state of the lower qubit, since only it is of interest. The refocused-INEPT network is not particularly intuitive in form, but the corresponding NMR pulse sequence is very simple to implement. The two controlled-NOT gates are responsible for the partial swap character of refocused-INEPT.

The net action of the refocused-INEPT sequence on a typical starting state of $(\sigma_{z1}) \otimes \mathbb{I}_2$ (neglecting multiples of the identity matrix) is

$$\sigma_{z1} \otimes \mathbb{I}_2 \xrightarrow{\text{R-INEPT}} \mathbb{I}_1 \otimes \sigma_{x2} \quad (31)$$

It uses the coupling between the spins to turn what was a maximally mixed state of the second spin into a superposition which can then be manipulated and observed. The process is somewhat analogous to using one of a pair of coupled pendula to set the other one swinging.

5.3 Correlation spectroscopy

The object of correlation spectroscopy is to arrive at spectral assignments by determining the network of spin-spin couplings in a molecule. It does this by extending the spectral information over two dimensions and looking for cross-peaks between spectral lines. For example, consider the spectrum of the sugar derivative shown in figure 6. All the peaks must be associated with nuclei, but it's not clear which are self-contained peaks and which arise from splittings, nor is it at all obvious what peaks might belong to what nuclei. It would be useful to know what couplings existed between the peaks, since this would allow us to begin piecing out a relation between the observed spectrum and the geometry of the molecule.

We can do this by taking a two-dimensional spectrum in which cross-peaks exist only between atoms whose scalar coupling is relatively large, since we can simply read off the correlation network from such a spectrum. A correlation spectrum of the same sugar derivative is shown in figure 7. The conventional spectrum can be seen running diagonally down the figure.

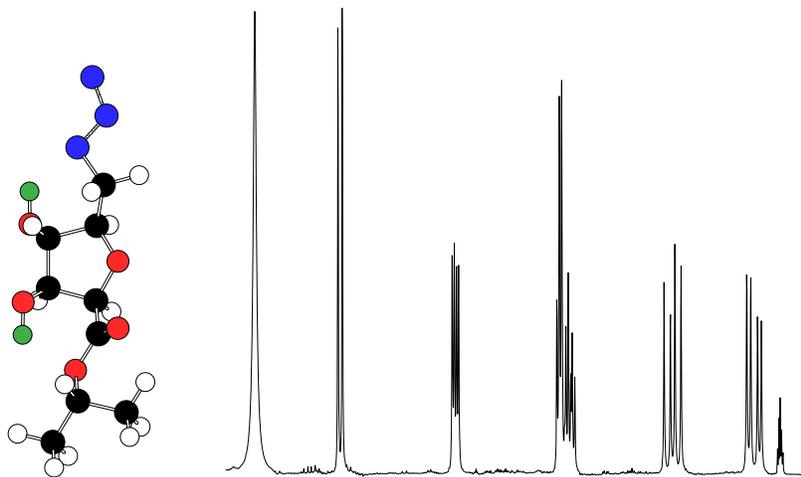


Figure 6: A sugar derivative and part of its NMR spectrum. The spectrum is sufficiently complicated that it becomes difficult to make assignments or infer geometric information about the molecule. Signal arises from the hydrogen atoms (white spheres).

Although the presentation of the data in figure 7 gives a better impression of the overall form of the COSY data, it is somewhat unusual; in practice it is much easier to interpret the spectra using flat contour plots, as shown in figure 8. Only part of the spectra and the corresponding part of the sugar molecule are shown.

5.3.1 Evolution and measurement—again

In order to understand where the second dimension in two-dimensional spectroscopy comes from, it's necessary to return once more to the subject of measurement and evolution. Ordinarily, we consider the two processes to be quite distinct—we allow spins to evolve over the course of a pulse sequence, and then we measure them at the end. Physically, however, the processes are almost identical. A spin is normally probed in NMR spectroscopy by using an RF pulse to knock the Bloch vector onto the xy -plane, where the precession of the superposition state about the z -axis induces a voltage in the receiver coil. In quantum computation, we generally want to know the present state of the spin, and so we look to see if the spin is inducing a current in the receiver coil. In both cases the measurement process involves

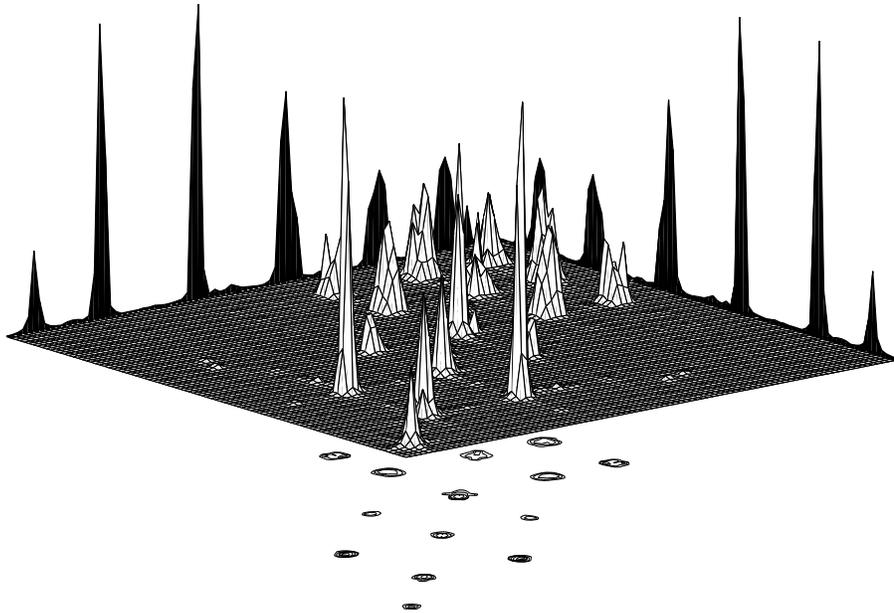


Figure 7: A COSY spectrum of the system shown in figure 6. The central region shows a three dimensional plot of the COSY data, the side panels show the (low resolution) one-dimensional spectrum, and a contour plot is shown on the ‘floor’. Peaks on the main diagonal which share off-diagonal cross-peaks correspond to atoms which are physically close to each other.

making a record of the changing state of the receiver coil with time and then taking a Fourier transform of this time-domain signal to extract the Larmor frequencies of the various spins in the system. During this time, the spins will continue to evolve under the influence of the system’s Hamiltonian. Not only is this natural, it is an essential part of the measurement process, because it is the changing state of the system that causes the changing state of the receiver coil. For example, the sub-components of an entangled spin will evolve at different frequencies under the influence of a J-coupling, and this frequency difference is manifested as a splitting in the spectral line. This frequency difference also means that the sub-components will generally end up in different places from one another at the end of the measurement, and non-classical correlations might have arisen in the system. Ordinarily, we don’t care, because the measurement was the final stage of the experiment.

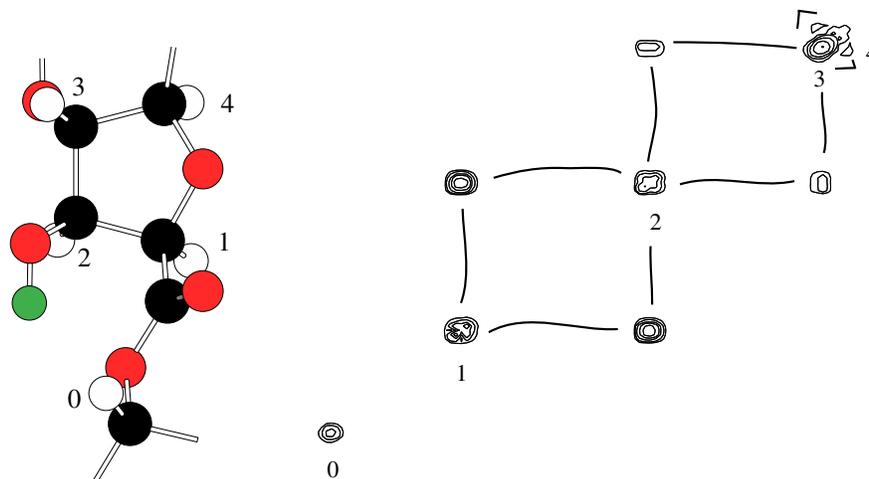
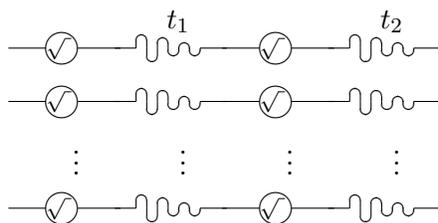


Figure 8: A portion of the sugar from figure 6 and the COSY spectrum shown in figure 7, with cross-peak connections drawn in. (Some cross-peaks lie outside the region shown.) Looking at the connections and observing that hydrogen 1 must be close to hydrogen 2, which must be close to hydrogen 3, which must be close to hydrogen 4, allows those hydrogens to be identified as part of the carbon chain in the sugar ring. Hydrogen 0 doesn't share any cross-peaks with these hydrogens, and so it must lie some distance from the central ring—examination of the rest of the COSY spectrum allows its location to be determined more exactly. Note that it's purely fortuitous that the spectral frequency happens to be correlated with the position in the molecule—if this were always the case, spectroscopy would be much simpler than it is.

What happens if we introduce a second period of evolution, t_1 , in the middle of the experiment? An evolution step like this is the basis of the controlled-NOT gate, and so it ought not to have very extraordinary consequences beyond the possible introduction of correlations into the system. But now what happens if we observe the variation of the final signal with t_1 ? Suddenly the extra evolution is both a measurement *and* an evolution step. Assuming (sensibly enough) that we also measure the system at the end of the experiment, we now have two time parameters, and so we can perform a two-dimensional Fourier transform on our data set. The way this is done in practice is that the sequence will be repeated many times for different values of t_1 , giving a series of free-induction decays in t_2 which can

be Fourier-transformed into a series of spectra. This series of spectra is then Fourier-transformed with respect to t_1 to give a two-dimensional spectrum.

While there are a variety of two-dimensional techniques, figure 7 was taken using the oldest and most basic of these, known as a COSY sequence. The COSY logic circuit is shown below (in general there will be many spins in the system):



The first RF pulse flips the spins down onto the y -axis so that they can be ‘measured’, the time delay allows correlations to build up in the system, the second RF pulse converts the correlations into a form which will be rendered observable by further free precession, and the final time delay then converts this antiphase magnetisation into an observable superposition of the second spin. (Observation of the signal also takes place during the second time period.) By working out the evolution of the system mathematically, one finds that the normal spectrum will be observed on the diagonal, and that cross-peaks will be observed between spectral lines if there exists a substantial J-coupling between them—exactly as we see in figures 7 and 8. In other words, if atoms are close together, they will couple and become correlated, whereas if they are spatially separated they won’t. We can extend this technique to three and four dimensions to extract even more information about coupling networks from large biomolecules.

6 Conclusions

NMR is a seductive subject for many reasons, but two of them are basic: it allows us to probe the material character of the world, and it allows us to probe the *possibilities* of the world. Because of its quantum character, NMR gives us a versatile and sensitive spectrographic tool for chemical analysis; in other words, it is a means of working out what things are made of and how molecules are put together. This directly extends our understanding of our surroundings but it is not the only worthwhile application of NMR; NMR also extends our understanding of what we can *do* with our

surroundings. We want to know what restrictions nature places on its exploitation, even if achieving this understanding means running into these restrictions. The limits on what can be done with classical computers and the more generous limits that exist for quantum computers make implementing a quantum computer using NMR an extraordinarily interesting problem. Even if NMR quantum computers never get beyond a primitive level, we still have a desire to build these primitive prototypes. Part of the reason is of course the fact that tricks and insights physicists gain from collaboration with chemists may be applicable to other implementations of a quantum computer or even to other branches of physics. NMR quantum computation has already prompted a flurry of investigations into the foundations of quantum mechanics and its connection to the deeper aspects of computer science. The real experimental attraction of an NMR quantum computer is more subtle, however; we want to try just to see if we *can*, even if our own best predictions tell us that we ultimately *can't*.

Acknowledgements

HKC thanks NSERC (Canada) and the TMR programme (EU) for their financial assistance. JAJ is a Royal Society University Research Fellow. The OCMS is supported by the UK EPSRC, BBSRC, and MRC. We are grateful to Tim Claridge (University of Oxford) for providing the data shown in figures 6, 7, and 8.

Holly Cummins took her first degree at McMaster University, Canada. She received early exposure to NMR during summer projects working on medical magnetic resonance imaging, and is now relieved to get away from biological squishy bits. She is currently studying for a DPhil at the Clarendon Laboratory in Oxford. Her primary research interest is the NMR implementation of quantum computers, which has led to a secondary interest in graphical depictions of quantum phenomena.

Jonathan Jones studied chemistry for his first degree at Oxford University, and then continued with a DPhil in NMR. His research interests in NMR have kept him largely in Oxford, but have led him to Departments as varied as Biochemistry and Physics. He now holds a Royal Society University Research Fellowship and works at both the Clarendon Laboratory and the Oxford Centre for Molecular Sciences. His research interests include NMR quantum computation and the development of NMR techniques to study protein folding.

The Oxford Centre for Quantum Computation can be found at www.qubit.org.

References

- [1] R. Freeman. *Spin Choreography: Basic Steps in High Resolution NMR*. Spektrum, 1997.
- [2] A. Einstein, N. Rosen, and B. Podolsky. *Physical Review*, 47:777, 1935.
- [3] J. S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195, 1964.
- [4] A. Aspect, J. Dalibard, and G. Roger. Experimental test of bell inequalities using time-varying analyzers. *Physical Review Letters*, 49:1804, 1982.
- [5] J. S. Bell. *Speakable and unspeakable in quantum mechanics*. Cambridge University Press, 1993.
- [6] L. Hardy. Spooky action at a distance in quantum mechanics. *Contemporary Physics*, 39(6):419, 1998.
- [7] N. D. Mermin. *Boojums all the way through*. Cambridge University Press, 1990.
- [8] M. B. Plenio and V. Vedral. Teleportation, entanglement and thermodynamics in the quantum world. *Contemporary Physics*, 39(6):431, 1998.
- [9] R. Fitzgerald. What really gives a quantum computer its power? *Physics Today*, January:20, 2000.
- [10] S. L. Braunstein, C. M. Caves, R. Jozsa, N. Linden, S. Popescu, and R. Schack. Separability of very noisy mixed states and implications for NMR quantum computing. *Physical Review Letters*, 83(5):1054, 1999.
- [11] R. Schack and C. M. Caves. Classical model for bulk-ensemble NMR quantum computation. *Physical Review A*, 60(6):4354, 1999.
- [12] E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25):5672, 1998.
- [13] C.H. Bennett and D.P. DiVincenzo. Quantum information and computation. *Nature*, 404(6775):247, 2000.
- [14] D. Deutsch. In R. Penrose and C. J. Isham, editors, *Quantum Concepts in Space and Time*. Clarendon, Oxford, 1986.

- [15] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London, Series A*, 439(1907):553, 1992.
- [16] L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79:325, 1997.
- [17] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Review*, 41:303, 1999.
- [18] B. Schneier. *Applied cryptography : protocols, algorithms, and source code in C*. Wiley, New York, second edition, 1996.
- [19] J. A. Jones and M. Mosca. Implementation of a quantum algorithm on a nuclear magnetic resonance quantum computer. *Journal of Chemical Physics*, 109:1648, 1998.
- [20] I. L. Chuang, L. M. K. Vandersypen, X. Zhou, D.W. Leung, and S. Lloyd. *Nature*, 393:143, 1998.
- [21] I. L. Chuang, N. Gershenfeld, and M. Kubinec. Experimental implementation of fast quantum searching. *Physical Review Letters*, 80:3408, 1998.
- [22] J. A. Jones, M. Mosca, and R. H. Hansen. Implementation of a quantum search algorithm on a quantum computer. *Nature*, 393:344, 1998.
- [23] L. M. K. Vandersypen, M. Steffen, M. H. Sherwood, C. S. Yannoni, G. Breyta, and I. L. Chuang. Implementation of a three-quantum-bit search algorithm. *Applied Physics Letters*, 76(5):646, 2000.
- [24] M.A. Nielsen, E. Knill, and R. Laflamme. Complete quantum teleportation using nuclear magnetic resonances. *Nature*, 396(6706):52, 1998.
- [25] D. G. Cory, M. D. Price, W. Maas, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo. Experimental quantum error correction. *Physical Review Letters*, 81(10):2152, 1998.
- [26] D. Leung, L. Vandersypen, X. L. Zhou, M. Sherwood, C. Yannoni, M. Kubinec, and I. Chuang. Experimental realization of a two-bit phase damping quantum code. *Physical Review A*, 60(3):1924, 1999.
- [27] Y. S. Weinstein, S. Lloyd, and D. G. Cory. Implementation of the quantum fourier transform. 1999. www.arxiv.org/quant-ph/9906059.

- [28] R. P. Feynman. *International Journal of Theoretical Physics*, 21, 1982.
- [29] S. Lloyd. Universal quantum simulators. *Science*, 273(5278):1073, 1999.
- [30] S. Somaroo, C. H. Tseng, T. F. Havel, R. Laflamme, and D. G. Cory. Quantum simulations on a quantum computer. *Physical Review Letters*, 82(26):5381, 1999.
- [31] J. A. Jones. NMR quantum computation:a critical evaluation. *Fortschritte der Physik*, 2000. www.arxiv.org/quant-ph/0002085, in press.
- [32] A. Barenco. Quantum physics and computers. *Contemporary Physics*, 37(5):375, 1996.
- [33] D.G. Cory, A. F. Fahmy, and T. F. Havel. Nuclear magnetic resonance spectroscopy: an experimentally accessible paradigm for quantum computing. In T. Toffoli, M. Biafore, and J. Leão, editors, *PhysComp96*, page 87. New England Complex Systems Institute, 1996.
- [34] E. Knill, I. Chuang, and R. Laflamme. Effective pure states for bulk quantum computation. *Physical Review A*, 57(5):3348, 1998.
- [35] N. Gershenfeld and I. Chuang. The usefulness of NMR quantum computing - response. *Science*, 277(5332):1689, 1997.
- [36] L. M. K. Vandersypen, C. S. Yannoni, M. Steffen, M. H. Sherwood, and I. L. Chuang. Realization of logically labeled effective pure states for bulk quantum computation. *Physical Review Letters*, 83(15):3085, 1999.
- [37] K. Dorai, Arvind, and A.Kumar. *Physical Review A*, page 042306, 2000.
- [38] E. Knill and R. Laflamme. Quantum computation and quadratically signed weight enumerators. 1999. www.arxiv.org/quant-ph/9909094.
- [39] A. Ekert and R. Jozsa. Quantum algorithms: entanglement-enhanced information processing. *Proceedings of the Royal Society of London, Series A*, 356(1743):1769, 1998.