NMR techniques for quantum control and computation

L. M. K. Vandersypen*

Kavli Institute of NanoScience, Delft University of Technology, 2628 CJ Delft, The Netherlands

I. L. Chuang†

Center for Bits and Atoms and Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

(Published 12 January 2005)

Fifty years of developments in nuclear magnetic resonance (NMR) have resulted in an unrivaled degree of control of the dynamics of coupled two-level quantum systems. This coherent control of nuclear spin dynamics has recently been taken to a new level, motivated by the interest in quantum information processing. NMR has been the workhorse for the experimental implementation of quantum protocols, allowing exquisite control of systems up to seven qubits in size. This article surveys and summarizes a broad variety of pulse control and tomographic techniques which have been developed for, and used in, NMR quantum computation. Many of these will be useful in other quantum systems now being considered for the implementation of quantum information processing tasks.

CONTENTS

I. Introduction 1037
II. The NMR System 1039
A. The system Hamiltonian 1039
1. Single spins 1039
2. Interacting spins 1040
a. Direct coupling 1040
b. Indirect coupling 1040
B. The control Hamiltonian 1041
1. Radio-frequency fields 1041
2. The rotating frame 1042
C. Relaxation and decoherence 1043
III. Elementary Pulse Techniques 1043
A. Quantum control, quantum circuits, and pulses 1043
1. Quantum gates and circuits 1043
2. Implementation of single-qubit gates 1044
3. Implementation of two-qubit gates 1044
4. Refocusing: Turning off undesired $I_z^i I_z^j$ couplings 1045
5. Pulse sequence simplification 1047
6. Time-optimal pulse sequences 1048
B. Experimental limitations 1049
1. Cross-talk 1049
2. Coupled evolution 1050
3. Instrumental errors 1050
IV. Advanced Pulse Techniques 1051
A. Shaped pulses 1051
1. Amplitude profiles 1051
2. Phase profiles 1053
B. Composite pulses 1054
1. Analytical approach 1054
2. Numerical optimization 1056
C. Average-Hamiltonian theory 1057
1. The Magnus expansion 1057
2. Multiple-pulse decoupling 1058
3. Reversing errors due to decoherence 1059
V. Evaluation of Quantum Control 1059
A. Standard experiments 1059
1. Coherent oscillations driven by a resonant field 1059
2. Coherent oscillations initiated by a kick 1060
3. Ramsey interferometry 1060
4. Measurement of T_2 1060
5. Measurement of T_1 1061
6. Measurement of $T_{1\rho}$ 1061
B. Measurement of quantum states and gates 1062
1. Quantum state tomography 1062
2. Quantum process tomography 1063
C. Fidelity of quantum states and gates 1064
1. Quantum state fidelity 1064
2. Quantum gate fidelity 1065
D. Evaluating scalability 1065
VI. Discussion and Conclusions 1065
References 1067

I. INTRODUCTION

Precise and complete control of multiple coupled quantum systems is expected to lead to profound insights in physics as well as to novel applications, such as quantum computation (Bennett and DiVincenzo, 2000; Nielsen and Chuang, 2000; Galindo and Martin-Delgado, 2002). Such coherent control is a major goal in atomic physics (Wieman *et al.*, 1999; Osborne and Coontz, 2002; Leibfried *et al.*, 2003), quantum optics (Zeilinger, 1999; Osborne and Coontz, 2002) and condensed-matter research (Clark, 2001; Maklin *et al.*, 2001; Osborne and Coontz, 2002; Zutic *et al.*, 2004), but

*Electronic address: lieven@qt.tn.tudelft.nl
†Electronic address: ichuang@mit.edu

 ©2004 The American Physical Society

surprisingly, many of the leading experimental results are coming from one of the oldest areas of quantum physics: nuclear magnetic resonance (NMR).

The development of NMR control techniques originated in a strong demand for precise spectroscopy of complex molecules: NMR is the premier tool for protein structure determination, and in modern NMR spectroscopy, often thousands of precisely sequenced and phase-controlled pulses are applied to molecules containing hundreds of nuclear spins. More recently, over the past seven years, a wide variety of complex quantum information processing tasks have been realized using NMR, on systems ranging from two to seven quantum bits (qubits) in size, on molecules in liquid (Chuang, Vandersypen, *et al.*, 1998; Jones *et al.*, 1998; Nielsen *et al.*, 1998; Somaroo *et al.*, 1999; Knill *et al.*, 2000; Vandersypen *et al.*, 2001), liquid crystal (Yannoni *et al.*, 1999), and solid-state samples (Zhang and Cory, 1998; Leskowitz *et al.*, 2003). These demonstrations have been made possible by application of a menagerie of new and previously existing control techniques, such as simultaneous and shaped pulses, composite pulses, refocusing schemes, and effective Hamiltonians. These techniques allow control and compensation for a variety of imperfections and experimental artifacts invariably present in real physical systems, such as pulse imperfections, Bloch-Siegert shifts, undesired multiple-spin couplings, field inhomogeneities, and imprecise system Hamiltonians.

The problem of control of multiple coupled quantum systems is a signature topic for NMR and can be summarized as follows: given a system with Hamiltonian $\mathcal{H}=\mathcal{H}_{\text{sys}}+\mathcal{H}_{\text{control}}$, where $\mathcal{H}_{\text{sys}}$ is the Hamiltonian in the absence of any active control, and $\mathcal{H}_{\text{control}}$ describes terms that are under external control, how can a desired unitary transformation U be implemented, in the presence of imperfections, and using minimal resources? Similar to other scenarios in which quantum control is a well-developed idea, such as in laser excitation of chemical reactions (Walmsley and Rabitz, 2003), $\mathcal{H}_{\text{control}}$ arises from precisely timed sequences of multiple pulses of electromagnetic radiation, applied phase-coherently, with different pulse widths, frequencies, phases, and amplitudes. However, importantly, in contrast to other areas of quantum control, in NMR $\mathcal{H}_{\text{sys}}$ is composed from multiple distinct physical pieces, i.e., the individual nuclear spins, providing the tensor product Hilbert-space structure vital to quantum computation. Furthermore, the NMR systems employed in quantum computation are better approximated as being closed, as opposed to open, quantum systems.

Nuclear spins and NMR provide a wonderful model and inspiration for the advance of coherent control over other coupled quantum systems, as many of the challenges and solutions are similar across the world of atomic, molecular, optical, and solid-state systems (see, for example, Steffen, 2003). Here, we review the control techniques employed in the field of NMR quantum computation, focusing on methods that are robust under experimental implementation, and including experimental prescriptions for evaluation of the efficacy of the techniques. In contrast to other reviews of NMR quantum computation which have appeared in the literature (Cory *et al.*, 2000; Jones, 2000; Vandersypen, 2001), and introductions to the subject (Gershenfeld and Chuang, 1998; Jones, 2001; Steffen *et al.*, 2001; Vandersypen *et al.*, 2002), we do not assume prior knowledge of, or give specialized descriptions of quantum computation algorithms, nor do we review NMR quantum computing experiments. And although we do not assume prior detailed knowledge of NMR, a self-contained treatment of several advanced topics, such as composite pulses, and refocusing, is included. Finally, because the primary purpose of this article is to elucidate control techniques which may generalize beyond NMR, we also assume a regime of operation in which relaxation and decoherence mechanisms are simple to treat and physical evolution is dominated by closed-system dynamics.

The organization of this article is as follows. In Sec. II, we briefly review the physics of NMR, using a Hamiltonian description of single and interacting nuclear spins 1/2 placed in a static magnetic field, controlled by radio-frequency fields. This establishes a foundation for the first major part of this review, Sec. III, which discusses the ways in which the control Hamiltonian can be used to construct all the elementary quantum gates, and the limitations that arise from the given system and control Hamiltonian, as well as from instrumental imperfections. The second major part of this review, Sec. IV, presents three classes of advanced techniques for tailoring the control Hamiltonian, which permit accurate quantum control despite the existing limitations: the methods of amplitude and frequency shaped pulses, composite pulses, and average Hamiltonian theory. Finally, in Sec. V, we describe a set of standard experiments, derived from quantum computation, which demonstrate coherent qubit control and can be used to characterize decoherence. These include procedures for quantum state and process tomography, as well as methods for evaluating the fidelity of quantum states and gates.

For further reading on NMR, we recommend the textbooks of Abragam (1962), Ernst, Bodenhausen, and Wokaun (1987) and Slichter (1996) for their rigorous discussions of the nuclear-spin Hamiltonian and standard pulse sequences; Freeman (1997) for an intuitive explanation of advanced techniques for control of the spin evolution; and Levitt (2001) for an intuitive understanding of the physics underlying the spin dynamics. Many useful reviews on specific NMR techniques are compiled in the Encyclopedia of NMR (Grant and Harris, 2001).

For additional reading on quantum computation, we recommend the book by Nielsen and Chuang (2000) for the basic theory of quantum information and computation; Bennett and DiVincenzo (2000); and Braunstein and Lo (2000) for reviews of the state of the art in experimental quantum information processing; and Lloyd (1995), for a simple introduction to quantum computation. Excellent presentations of quantum algorithms are given by Ekert and Jozsa (1996) and Steane (1998).

The original papers introducing NMR quantum computing are those of Cory *et al.* (1996, 1997; Cory, Price,

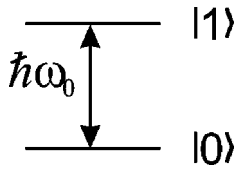

FIG. 1. Energy diagram for a single spin-1/2 particle.

and Havel, 1998), and Gershenfeld and Chuang (1997). Gershenfeld and Chuang (1998) and Steffen *et al.* (2001) give elementary introductions to NMR quantum computing, while introductions geared towards NMR spectroscopists are presented by Jones (2001) and Vandersypen *et al.* (2002). Summaries of NMR quantum computing experiments and techniques are given by Cory *et al.* (2000), Jones (2000), and Vandersypen (2001).

II. THE NMR SYSTEM

We begin with a description of the NMR system, based on its system Hamiltonian and the control Hamiltonian. The system Hamiltonian gives the energy of single and coupled spins in a static magnetic field, and the control Hamiltonian arises from the application of radio-frequency pulses to the system at, or near, its resonant frequencies. A rotating reference frame is employed, providing a very convenient description.

A. The system Hamiltonian

1. Single spins

The time evolution of a spin-1/2 particle (we shall not consider higher-order spins in this paper) in a magnetic field $\vec{B}_0$ along $\hat{z}$ is governed by the Hamiltonian

$$\mathcal{H}_0 = -\hbar\gamma B_0 I_z = -\hbar\omega_0 I_z = \begin{bmatrix} -\hbar\omega_0/2 & 0 \\ 0 & \hbar\omega_0/2 \end{bmatrix}, \quad (1)$$

where γ is the gyromagnetic ratio of the nucleus, $\omega_0/2\pi$ is the Larmor frequency,[1] and I_z is the angular momentum operator in the $\hat{z}$ direction. I_z, I_x, and I_y relate to the well-known Pauli matrices as

$$\sigma_x = 2I_x, \quad \sigma_y = 2I_y, \quad \sigma_z = 2I_z, \quad (2)$$

where, in matrix notation,

$$\sigma_x \equiv \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}; \quad \sigma_y \equiv \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}; \quad \sigma_z \equiv \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (3)$$

The interpretation of Eq. (1) is that the $|0\rangle$ or $|\uparrow\rangle$ energy (given by $\langle 0|\mathcal{H}|0\rangle$, the upper left element of $\mathcal{H}$) is lower than the $|1\rangle$ or $|\downarrow\rangle$ energy ($\langle 1|\mathcal{H}|1\rangle$) by an amount $\hbar\omega_0$, as illustrated in the energy diagram of Fig. 1. The energy splitting is known as the Zeeman splitting.

We can pictorially understand the time evolution $U = e^{-i\mathcal{H}t/\hbar}$ under the Hamiltonian of Eq. (1) as a precessing motion of the Bloch vector about $\vec{B}_0$, as shown in Fig. 2. As is conventional, we define the $\hat{z}$ axis of the Bloch sphere as the quantization axis of the Hamiltonian, with $|0\rangle$ along $+\hat{z}$ and $|1\rangle$ along $-\hat{z}$.

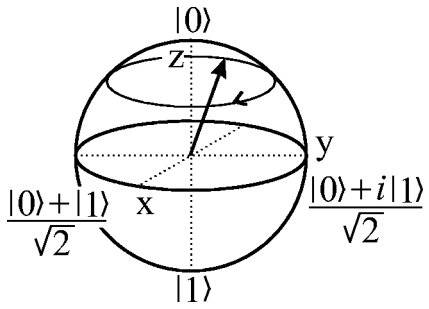

FIG. 2. Precession of a spin-1/2 particle about the axis of a static magnetic field.

For the case of liquid-state NMR, which we shall largely restrict ourselves to in this article, typical values of B_0 are 5–15 T, resulting in precession frequencies ω_0 of a few hundred MHz, the radio-frequency range.

Spins of different nuclear species (*heteronuclear spins*) can be easily distinguished spectrally, as they have very distinct values of γ and thus also very different Larmor frequencies (Table I). Spins of the same nuclear species (*homonuclear spins*) which are part of the same molecule can also have distinct frequencies, by amounts known as their *chemical shifts* $\tilde{\sigma}_i$.

The nuclear-spin Hamiltonian for a molecule with n uncoupled nuclei is thus given by

$$\mathcal{H}_0 = -\sum_{i=1}^{n} \hbar(1-\tilde{\sigma}_i)\gamma_i B_0 I_z^i = -\sum_{i=1}^{n} \hbar\omega_0^i I_z^i, \quad (4)$$

where the i superscripts label the nuclei.

The chemical shifts arise from partial shielding of the externally applied magnetic field by the electron cloud surrounding the nuclei. The amount of shielding depends on the electronic environment of each nucleus, so like nuclei with inequivalent electronic environments have different chemical shifts. Pronounced asymmetries in the molecular structure generally promote strong chemical shifts. The range of typical chemical shifts $\tilde{\sigma}_i$ varies from nucleus to nucleus, e.g., ≈10 parts per million (ppm) for ^{1}H, ≈200 ppm for ^{19}F, and ≈200 ppm for ^{13}C. At B_0=10 T, this corresponds to a few kHz to tens of kHz (compared to ω_0's of several hundred MHz). As an example, Fig. 3 shows an experimentally measured spectrum of a molecule containing five fluorine spins with inequivalent chemical environments.

TABLE I. Larmor frequencies (MHz) for some relevant nuclei, at 11.74 T.

Nucleus	^{1}H	^{2}H	^{13}C	^{15}N	^{9}F	^{31}P
$\omega_0/2\pi$	500	77	126	−51	470	202

[1]We shall sometimes leave the factor of 2π implicit and call ω_0 the Larmor frequency.

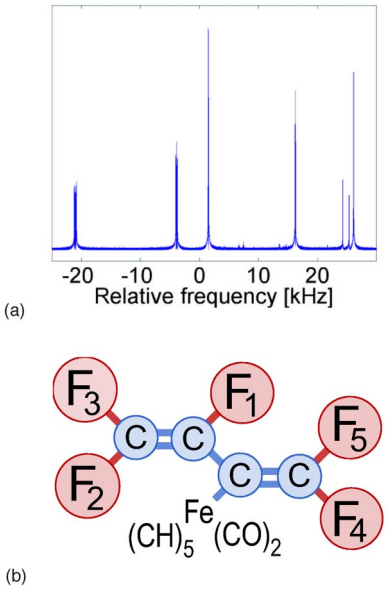

FIG. 3. (Color in online edition) Fluorine NMR spectrum (absolute value) centered around ≈470 MHz of a specially designed molecule, shown in (b). The five main lines in the spectrum correspond to the five fluorine nuclei in the molecule. The two small lines derive from impurities in the sample. The NMR spectra were acquired by recording the oscillating magnetic field produced by a large ensemble of precessing spins and by taking the Fourier transform of this time-domain signal. The precession motion of the spins is started by applying a radio-frequency pulse (Sec. II.B.1), which tips the spins from their equilibrium position along the $\hat{z}$ axis into the $\hat{x}$-$\hat{y}$ plane. (b) From Vandersypen, Steffen, Breyta, Yannoni, Cleve, and Chuang, 2000.

In general, the chemical shift can be spatially anisotropic and must be described by a tensor. In liquid solution, this anisotropy averages out due to rapid tumbling of the molecules. In solids, the anisotropy means that the chemical shifts depend on the orientation of the molecule with respect to $\vec{B}_0$.

2. Interacting spins

For nuclear spins in molecules, nature provides two distinct interaction mechanisms which we now describe, the direct dipole-dipole interaction, and the electron-mediated Fermi contract interaction known as J coupling.

a. Direct coupling

The *magnetic dipole-dipole interaction* is similar to the interaction between two bar magnets in each other's vicinity. It takes place purely through space—no medium is required for this interaction—and depends on the internuclear vector $\vec{r}_{ij}$ connecting the two nuclei i and j, as described by the Hamiltonian

$$\mathcal{H}_D = \sum_{i<j} \frac{\mu_0 \gamma_i \gamma_j \hbar}{4\pi |\vec{r}_{ij}|^3} \left[\vec{I}^i \cdot \vec{I}^j - \frac{3}{|\vec{r}_{ij}|^2} (\vec{I}^i \cdot \vec{r}_{ij})(\vec{I}^j \cdot \vec{r}_{ij}) \right], \tag{5}$$

where μ_0 is the usual magnetic permeability of free space and $\vec{I}^i$ is the magnetic moment vector of spin i. This expression can be progressively simplified as various conditions are met. These simplifications rest on averaging effects and can be explained within the general framework of average-Hamiltonian theory (Sec. IV.C).

For large $\omega_0^i = \gamma^i B_0$ (i.e., at high B_0), $\mathcal{H}_D$ can be approximated as

$$\mathcal{H}_D = \sum_{i<j} \frac{\mu_0 \gamma_i \gamma_j \hbar}{8\pi |\vec{r}_{ij}|^3} (1 - 3\cos^2\theta_{ij}) [3 I_z^i I_z^j - \vec{I}^i \cdot \vec{I}^j], \tag{6}$$

where θ_{ij} is the angle between B_0 and $\vec{r}_{ij}$. When $|\omega_0^i - \omega_0^j|$ is much larger than the coupling strength, the transverse coupling terms can be dropped, so $\mathcal{H}_D$ simplifies further to

$$\mathcal{H}_D = \sum_{i<j} \frac{\mu_0 \gamma_i \gamma_j \hbar}{4\pi |\vec{r}_{ij}|^3} (1 - 3\cos^2\theta_{ij}) I_z^i I_z^j, \tag{7}$$

which has the same form as the J coupling we describe next [Eq. (9)].

For molecules in liquid solution, both intramolecular dipolar couplings (between spins in the same molecule) and intermolecular dipolar couplings (between spins in different molecules) are averaged away due to rapid tumbling. This is the case we shall focus on in this article. In solids, similarly simple Hamiltonians can be obtained by applying multiple-pulse sequences which average out undesired coupling terms (Haeberlen and Waugh, 1968), or by physically spinning the sample at an angle of $\arccos(1/\sqrt{3})$ (the "magic angle") with respect to the magnetic field.

b. Indirect coupling

The second interaction mechanism between nuclear spins in a molecule is the *J coupling* or *scalar coupling*. This interaction is mediated by the electrons shared in the chemical bonds between the atoms and due to the overlap of the shared electron wave function with the two coupled nuclei, a Fermi contact interaction. The *through-bond coupling strength J* depends on the respective nuclear species and decreases with the number of chemical bonds separating the nuclei. Typical values for J are up to a few hundred Hz for one-bond couplings and down to only a few Hz for three- or four-bond couplings. The Hamiltonian is

$$\mathcal{H}_J = \hbar \sum_{i<j} 2\pi J_{ij} \vec{I}^i \cdot \vec{I}^j = \hbar \sum_{i<j} 2\pi J_{ij} (I_x^i I_x^j + I_y^i I_y^j + I_z^i I_z^j), \tag{8}$$

where J_{ij} is the coupling strength between spins i and j. Similar to the case of dipolar coupling, Eq. (8) simplifies to

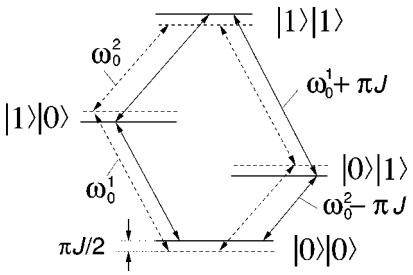

FIG. 4. Energy-level diagram for (dashed lines) two uncoupled spins and (solid lines) two spins coupled by a Hamiltonian of the form of Eq. (7) or Eq. (9) in units of $\hbar$.

$$\mathcal{H}_J = \hbar \sum_{i<j}^{n} 2\pi J_{ij} I_z^i I_z^j, \tag{9}$$

when $|\omega_i - \omega_j| \gg 2\pi|J_{ij}|$, a condition easily satisfied for heteronuclear spins and which can also be satisfied for small homonuclear molecules.

The interpretation of the scalar coupling term of Eq. (9) is that a spin "feels" a static magnetic field along $\pm\hat{z}$ produced by the neighboring spins, in addition to the externally applied $\vec{B}_0$ field. This additional field shifts the energy levels as in Fig. 4. As a result, the Larmor frequency of spin i shifts by $-J_{ij}/2$ if spin j is in $|0\rangle$ and by $+J_{ij}/2$ if spin j is in $|1\rangle$.

In a system of two coupled spins, the frequency spectrum of spin i therefore actually consists of two lines separated by J_{ij} and centered around ω_0^i, each of which can be associated with the state of spin j, $|0\rangle$ or $|1\rangle$. For three pairwise coupled spins, the spectrum of each spin contains four lines. For every additional spin, the number of lines per multiplet doubles, provided all the couplings are resolved and different lines do not lie on top of each other. This is illustrated for a five-spin system in Fig. 5.

The magnitude of all the pairwise couplings can be found by looking for common splittings in the multiplets of different spins. The relative signs of the J couplings can be determined via appropriate spin-selective two-pulse sequences, known in NMR as two-dimensional correlation (soft-COSY) experiments (Brüschweiler *et al.*, 1987) or via line-selective continuous irradiation; both approaches are related to the CNOT gate (Sec. III.A.3). The signs cannot be obtained from just the simple spectra.

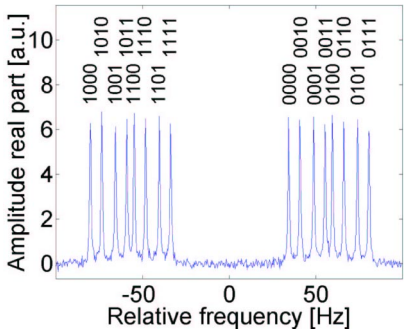

FIG. 5. (Color in online edition) The spectrum of spin F_1 in the molecule of Fig. 3. This is an expanded view of the left line in the spectrum of Fig. 3. Frequencies are given with respect to ω_0^1. The state of the remaining spins is as indicated, based on $J_{12}<0$ and $J_{13},J_{14},J_{15}>0$; furthermore, $|J_{12}|>|J_{13}|>|J_{15}|>|J_{14}|$. From Vandersypen, Steffen, Breyta, Yannoni, Cleve, and Chuang, 2000.

In summary, the simplest form of the Hamiltonian for a system of n coupled nuclear spins is thus [from Eqs. (4) and (9)]

$$\mathcal{H}_{\rm sys} = -\sum_i \hbar\omega_0^i I_z^i + \hbar \sum_{i<j} 2\pi J_{ij} I_z^i I_z^j. \tag{10}$$

In almost all NMR quantum computing experiments performed to date, the system is well described by a Hamiltonian of this form.

B. The control Hamiltonian

1. Radio-frequency fields

We turn now to physical mechanisms for controlling the NMR system. The state of a spin-1/2 particle in a static magnetic field $\vec{B}_0$ along $\hat{z}$ can be manipulated by applying an electromagnetic field $\vec{B}_1(t)$ which rotates in the $\hat{x}$-$\hat{y}$ plane at ω_{rf}, at or near the spin precession frequency ω_0. The single-spin Hamiltonian corresponding to the radio-frequency (RF) field is, analogous to Eq. (1) for the static field B_0,

$$\mathcal{H}_{rf} = -\hbar\gamma B_1[\cos(\omega_{rf}t+\phi)I_x - \sin(\omega_{rf}t+\phi)I_y], \tag{11}$$

where ϕ is the phase of the RF field, and B_1 its amplitude (the minus sign in front of the sine term makes the RF field evolve in the same sense as the spin evolution under $\mathcal{H}_0$). Typical values for $\omega_1=\gamma B_1$ are up to $\approx$50 kHz in liquid NMR and up to a few hundred kHz in solid NMR experiments. For n spins, we have

$$\mathcal{H}_{rf} = -\sum_i^n \hbar\gamma_i B_1[\cos(\omega_{rf}t+\phi)I_x^i - \sin(\omega_{rf}t+\phi)I_y^i]. \tag{12}$$

In practice, a magnetic field is applied which oscillates along a fixed axis in the laboratory, perpendicular to the static magnetic field. This oscillating field can be decomposed into two counter-rotating fields, one of which rotates at ω_{rf} in the same direction as the spin and so can be set on or near resonance with the spin. The other component rotates in the opposite direction and is thus very far off-resonance (by about $2\omega_0$). As we shall see, its only effect is a negligible shift in the Larmor frequency, called the Bloch-Siegert shift (Bloch and Siegert, 1940).

Note that both the amplitude B_1 and phase ϕ of the

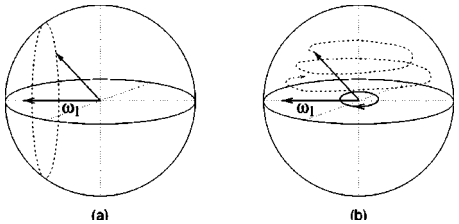

FIG. 6. Nutation of a spin subject to a transverse RF field (a) observed in the rotating frame and (b) observed in the lab frame.

RF field can be varied with time,[2] unlike the Larmor precession and the coupling terms. As we shall shortly see, it is the control of the RF field phases, amplitudes, and frequencies which lies at the heart of quantum control of NMR systems.

2. The rotating frame

The motion of a single nuclear spin subject to both a static and a rotating magnetic field is rather complex when described in the usual laboratory coordinate system (the *lab frame*). It is much simplified, however, by describing the motion in a coordinate system rotating about $\hat{z}$ at ω_{rf} (the *rotating frame*):

$$|\psi\rangle^{rot} = \exp(-i\omega_{rf}tI_z)|\psi\rangle. \tag{13}$$

Substitution of $|\psi\rangle$ in the Schrödinger equation $i\hbar(d|\psi\rangle/dt) = \mathcal{H}|\psi\rangle$ with

$$\mathcal{H} = -\hbar\omega_0 I_z - \hbar\omega_1[\cos(\omega_{rf}t + \phi)I_x - \sin(\omega_{rf}t + \phi)I_y] \tag{14}$$

gives $i\hbar(d|\psi\rangle^{rot}/dt) = \mathcal{H}^{rot}|\psi\rangle^{rot}$, where

$$\mathcal{H}^{rot} = -\hbar(\omega_0 - \omega_{rf})I_z - \hbar\omega_1[\cos\phi I_x - \sin\phi I_y]. \tag{15}$$

Naturally, the RF field lies along a fixed axis in the frame rotating at ω_{rf}. Furthermore, if $\omega_{rf} = \omega_0$, the first term in Eq. (15) vanishes. In this case, an observer in the rotating frame will see the spin simply precess about $\vec{B}_1$ [Fig. 6(a)], a motion called *nutation*. The choice of ϕ controls the nutation axis. An observer in the lab frame sees the spin spiral down over the surface of the Bloch sphere [Fig. 6(b)].

If the RF field is off-resonance with respect to the spin frequency by $\Delta\omega = \omega_0 - \omega_{rf}$, the spin precesses in the rotating frame about an axis tilted away from the $\hat{z}$ axis by an angle

$$\alpha = \arctan(\omega_1/\Delta\omega), \tag{16}$$

and with frequency

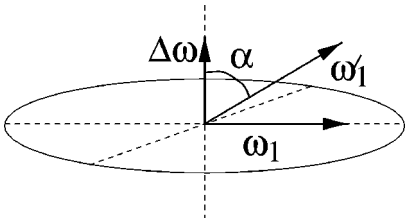

FIG. 7. Axis of rotation (in the rotating frame) during an off-resonant radio-frequency pulse.

$$\omega_1' = \sqrt{\Delta\omega^2 + \omega_1^2}, \tag{17}$$

as illustrated in Fig. 7.

It follows that the RF field has virtually no effect on spins that are far off resonance, since α is very small when $|\Delta\omega| \gg \omega_1$ (see Fig. 8). If all spins have well-separated Larmor frequencies, we can thus in principle selectively rotate any one qubit without rotating the other spins.

Moderately off-resonance pulses ($|\Delta\omega| \approx \omega_1$) do rotate the spin, but due to the tilted rotation axis, a single such pulse cannot, for instance, flip a spin from $|0\rangle$ to $|1\rangle$ (see again Fig. 8). Of course, off-resonance pulses can also be useful, for instance, for direct implementation of rotations about an axis outside the $\hat{x}$-$\hat{y}$ plane.

We could also choose to work in a frame rotating at ω_0 (instead of ω_{rf}), where

$$\mathcal{H}^{rot} = -\hbar\omega_1\{\cos[(\omega_{rf} - \omega_0)t + \phi]I_x - \sin[(\omega_{rf} - \omega_0)t + \phi]I_y\}. \tag{18}$$

This transformation does not give a convenient time-independent RF Hamiltonian (unless $\omega_{rf} = \omega_0$), as was the case for $\mathcal{H}^{rot}$ in Eq. (15). However, it is a natural starting point for the extension to the case of multiple

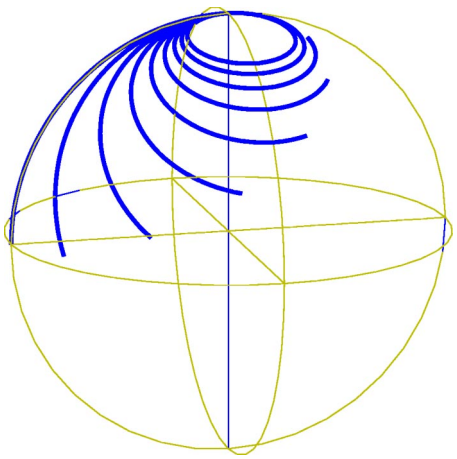

FIG. 8. (Color in online edition) Trajectory in the Bloch sphere described by a qubit initially in $|0\rangle$ (along $+\hat{z}$), after a 250-μs pulse of strength $\omega_1 = 1$ kHz is applied off-resonance by $0, 0.5, 1, \ldots, 4$ kHz. On-resonance, the pulse produces a 90° rotation. Far off-resonance, the qubit is hardly rotated away from $|0\rangle$.

[2]For example, the Varian Instruments Unity Inova 500 NMR spectrometer achieves a phase resolution of 0.5° and 4095 linear steps of amplitude control, with a time base of 50 ns. Additional attenuation of the amplitude can be done on a logarithmic scale over a range of about 80 dB, albeit with a slower time base.

spins, where a separate rotating frame can be introduced for each spin:

$$|\psi\rangle^{rot} = \left[\prod_i \exp(-i\omega_0^i t I_z^i)\right]|\psi\rangle. \tag{19}$$

In the presence of multiple RF fields indexed r, the RF Hamiltonian in this multiply rotating frame is

$$\mathcal{H}^{rot} = \sum_{i,r} -\hbar\omega_1^r\{\cos[(\omega_{rf}^r - \omega_0^i)t + \phi^r]I_x^i - \sin[(\omega_{rf}^r - \omega_0^i)t + \phi^r]I_y^i\}, \tag{20}$$

where the amplitudes ω_1^r and phases ϕ^r are under user control.

The system Hamiltonian of Eq. (10) is simplified, in the rotating frame of Eq. (19); the I_z^i terms drop out, leaving just the $J_{ij}I_z^iI_z^j$ couplings, which remain invariant. Note that coupling terms of the form $\vec{I}^i\cdot\vec{I}^j$ do not transform cleanly under Eq. (19).

Summarizing, in the multiply rotating frame, the NMR Hamiltonian $\mathcal{H}=\mathcal{H}_{\rm sys}+\mathcal{H}_{\rm control}$ takes the form

$$\mathcal{H}_{\rm sys} = \hbar\sum_{i<j} 2\pi J_{ij}I_z^iI_z^j, \tag{21}$$

$$\mathcal{H}_{\rm control} = \sum_{i,r} -\hbar\omega_1^r\{\cos[(\omega_{rf}^r - \omega_0^i)t + \phi^r]I_x^i - \sin[(\omega_{rf}^r - \omega_0^i)t + \phi^r]I_y^i\}. \tag{22}$$

C. Relaxation and decoherence

One of the strengths of nuclear spins as quantum bits is precisely the fact that the system is very well isolated from the environment, allowing coherence times to be long compared with the dynamical time scales of the system. Thus our discussion here focuses on closed-system dynamics, and it is important to be aware of the limits of this approximation.

The coupling of the NMR system to the environment may be described by an additional Hamiltonian term $\mathcal{H}_{\rm env}$, whose magnitude is small compared to that of $\mathcal{H}_{\rm sys}$ or $\mathcal{H}_{\rm control}$. It is this coupling which leads to decoherence, the loss of quantum information, which is traditionally parametrized by two rates: T_1, the energy relaxation rate, and T_2, the phase randomization rate (see also Secs. V.A.4 and V.A.5).

T_2 originates from spin-spin couplings which are imperfectly averaged away, or unaccounted for in the system Hamiltonian. For example, in molecules in liquid solution, spins on one molecule may have a long-range, weak interaction with spins on another molecule. Fluctuating magnetic fields, caused by spatial anisotropy of the chemical shift, local paramagnetic ions, or unstable laboratory fields, also contribute to T_2. Nevertheless, in well-prepared samples and in a good experimental apparatus at reasonably high magnetic fields, the T_2 for molecules in solution is easily on the order of 1 s or more. This decoherence mechanism can be identified with elastic scattering in other physical systems; it does not lead to loss of energy from the system.

T_1 originates from couplings between the spins and the "lattice," that is, excitation modes that can carry away energy quanta on the scale of the Larmor frequency. For example, these may be vibrational quanta, paramagnetic ions, chemical reactions such as ions exchanging with the solvent, or spins with higher-order magnetic moments (such as ^{2}H, ^{17}Cl, or ^{35}Br), which relax quickly due to their quadrupolar moment's interacting with electric field gradients. In well-chosen molecules and liquid samples with good solvents, T_1 can easily be tens of seconds, while isolated nuclei embedded in solid samples with a spin-zero host crystal matrix (such as ^{31}P in ^{28}Si) can have T_1 times of days. This mechanism is analogous to inelastic scattering in other physical systems.

The description of relaxation in terms of only two parameters is known to be an oversimplification of reality, particularly for coupled spin systems, in which coupled relaxation mechanisms appear (Redfield, 1957; Jeener, 1982). Nevertheless, the independent spin decoherence model is useful for its simplicity and because it can capture well the main effects of decoherence on simple NMR quantum computations (Vandersypen *et al.*, 2001), which are typically designed as pulse sequences shorter in time than T_2.

III. ELEMENTARY PULSE TECHNIQUES

This section begins our discussion of the main subject of this article, a review of the control techniques developed in NMR quantum computation for coupled two-level quantum systems. We begin with a quick overview of the language of quantum circuits and its important universality theorems, then connect this with the language of pulse sequences as used in NMR, and indicate how pulse sequences can be simplified. The main approximations employed in this section are that pulses can be strong compared with the system Hamiltonian while selectively addressing only one qubit at a time, and can be perfectly implemented. The limits of these approximations are discussed in the last part of the section.

A. Quantum control, quantum circuits, and pulses

The goal of quantum control, in the context of quantum computation, is the implementation of a unitary transformation U, specified in terms of a sequence $U = U_kU_{k-1}\cdots U_2U_1$ of standard "quantum gates" U_i, which act locally (usually on one or two qubits) and are simple to implement. As is conventional for unitary operations, the U_i are ordered in time from right to left.

1. Quantum gates and circuits

The basic single-qubit quantum gates are rotations, defined as

$$R_{\hat{n}}(\theta) = \exp\left[-\frac{i\theta\hat{n}\cdot\vec{\sigma}}{2}\right], \tag{23}$$

where $\hat{n}$ is a (three-dimensional) vector specifying the axis of the rotation, θ is the angle of rotation, and $\vec{\sigma} = \sigma_x\hat{x} + \sigma_y\hat{y} + \sigma_z\hat{z}$ is a vector of Pauli matrices. It is also convenient to define the Pauli matrices [see Eq. (3)] themselves as logic gates, in terms of which σ_x can be understood as being analogous to the classical NOT gate, which flips $|0\rangle$ to $|1\rangle$ and vice versa. In addition, the HADAMARD gate H and $\pi/8$ gate T

$$H = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad T = \begin{bmatrix} 1 & 0 \\ 0 & \exp(i\pi/4) \end{bmatrix} \tag{24}$$

are useful and widely employed. These and any other single-qubit transformation U can be realized using a sequence of rotations about just two axes, according to Bloch's theorem: for any single-qubit U, there exist real numbers α, β, γ, and δ such that

$$U = e^{i\alpha}R_x(\beta)R_y(\gamma)R_x(\delta). \tag{25}$$

The basic two-qubit quantum gate is a controlled-NOT (CNOT) gate,

$$U_{\mathrm{CNOT}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \tag{26}$$

where the basis elements in this notation are $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$ from left to right and top to bottom. U_{CNOT} flips the second qubit (the target) if and only if the first qubit (the control) is $|1\rangle$. This gate is the analog of the classical exclusive-OR gate, since $U_{\mathrm{CNOT}}|x,y\rangle = |x, x\oplus y\rangle$, for $x,y \in \{0,1\}$ and where $\oplus$ denotes addition modulo two.

A basic theorem of quantum computation is that up to an irrelevant overall phase, any U acting on n qubits can be composed from U_{CNOT} and $R_{\hat{n}}(\theta)$ gates (Nielsen and Chuang, 2000). Thus the problem of quantum control can be reduced to implementing U_{CNOT} and single-qubit rotations, where at least two nontrivial rotations are required. Other such sets of universal gates are known, but this is the one that has been employed in NMR.

These gates and sequences of such gates may be conveniently represented using quantum circuit diagrams, employing standard symbols. We shall use a notation commonly employed in the literature (Nielsen and Chuang, 2000) in this article.

2. Implementation of single-qubit gates

Rotations on single qubits may be implemented directly in the rotating frame using RF pulses. From the control Hamiltonian, Eq. (22), it follows that when an RF field of amplitude ω_1 is applied to a single-spin system at $\omega_{rf=\omega_0}$, the spin evolves under the transformation

$$U = \exp[i\omega_1(\cos\phi I_x - \sin\phi I_y)t_{pw}], \tag{27}$$

where t_{pw} is the *pulse width* (or pulse length), the time duration of the RF pulse. U describes a rotation in the Bloch sphere over an angle θ proportional to the product of t_{pw} and $\omega_1 = \gamma B_1$, and about an axis in the $\hat{x}$-$\hat{y}$ plane determined by the phase ϕ.

Thus a pulse with phase $\phi = \pi$ and $\omega_1 t_{pw} = \pi/2$ will perform $R_x(90)$ [see Eq. (23)], which is a 90° rotation about $\hat{x}$, denoted for short as X. A similar pulse but twice as long realizes a $R_x(180)$ rotation, written for short as X^2. By changing the phase of the RF pulse to $\phi = \pi/2$, Y and Y^2 pulses can similarly be implemented. For $\phi = 0$, a negative rotation about $\hat{x}$, denoted $R_x(-90)$ or $\bar{X}$, is obtained, and similarly $\phi = -\pi/2$ gives $\bar{Y}$. For multiqubit systems, subscripts are used to indicate on which qubit the operation acts, e.g., $\bar{Z}_3^2$ is a 180° rotation of qubit 3 about $-\hat{z}$.

It is thus not necessary to apply the RF field along different spatial axes in the lab frame to perform $\hat{x}$ and $\hat{y}$ rotations. Rather, the phase of the RF field determines the nutation axis in the rotating frame. Furthermore, note that only the *relative* phase between pulses applied to the same spin matters. The absolute phase of the first pulse on any given spin does not matter in itself. It just establishes a phase reference against which the phases of all subsequent pulses on that same spin, as well as the read-out of that spin, should be compared.

We noted earlier that the ability to implement arbitrary rotations about $\hat{x}$ and $\hat{y}$ is sufficient for performing arbitrary single-qubit rotations [Eq. (25)]. Since $\hat{z}$ rotations are very common, two useful explicit decompositions of $R_z(\theta)$ in terms of $\hat{x}$ and $\hat{y}$ rotations are

$$R_z(\theta) = XR_y(\theta)\bar{X} = YR_x(-\theta)\bar{Y}. \tag{28}$$

3. Implementation of two-qubit gates

The most natural two-qubit gate is the one generated directly by the spin-spin coupling Hamiltonian. For nuclear spins in a molecule in liquid solution, the coupling Hamiltonian is given by Eq. (9) (in the lab frame as well as in the rotating frame), from which we obtain the time evolution operator $U_J(t) = \exp[-i2\pi J I_z^1 I_z^2 t]$, or in matrix form

$$U_J(t) = \begin{bmatrix} e^{-i\pi Jt/2} & 0 & 0 & 0 \\ 0 & e^{+i\pi Jt/2} & 0 & 0 \\ 0 & 0 & e^{+i\pi Jt/2} & 0 \\ 0 & 0 & 0 & e^{-i\pi Jt/2} \end{bmatrix}. \tag{29}$$

Allowing this evolution to occur for time $t = 1/2J$ gives a transformation known as the *controlled phase gate*, up to a 90° phase shift on each qubit and an overall (and thus irrelevant) phase

$$U_{\mathrm{CPHASE}} = \sqrt{-i}\bar{Z}_1\bar{Z}_2U_J(1/2J) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}. \tag{30}$$

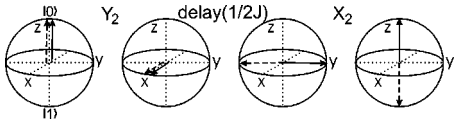

FIG. 9. Bloch-sphere representation of the operation of the CNOT_{12} gate between two qubits 1 and 2 coupled by $\hbar 2\pi J I_z^1 I_z^2$. Here, qubit 2 starts off in $|0\rangle$ (along $\hat{z}$) and is depicted in a reference frame rotating about $\hat{z}$ at $\omega_0^2/2\pi$. Solid and dashed arrows correspond to the case where qubit 1 is $|0\rangle$ and $|1\rangle$, respectively. Adapted from Gershenfeld and Chuang, 1997.

This gate is equivalent to the well-known CNOT gate up to a basis change of the target qubit and a phase shift on the control qubit

$$\begin{aligned} U_{\mathrm{CNOT}} &= iZ_1^2 \bar{Y}_2 U_{\mathrm{CPHASE}} Y_2 \\ &= iZ_1^2 \bar{Y}_2 [\sqrt{-i}\bar{Z}_1 \bar{Z}_2 U_J(1/2J)] Y_2 \\ &= \sqrt{i} Z_1 \bar{Z}_2 X_2 U_J(1/2J) Y_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \end{aligned} \quad (31)$$

The core of this sequence, $X_2 U_J(1/2J) Y_2$, can be graphically understood via Fig. 9 (Gershenfeld and Chuang, 1997), assuming the spins start along $\pm\hat{z}$. First, a spin-selective pulse on spin 2 about $\hat{y}$ (an rf pulse centered at $\omega_0^2/2\pi$ and of a spectral bandwidth such that it covers the frequency range $\omega_0^2/2\pi \pm J_{12}/2$ but not $\omega_0^1/2\pi \pm J_{12}/2$) rotates spin 2 from $\hat{z}$ to $\hat{x}$. Next, the spin system is allowed to freely evolve for a duration of $1/2J_{12}$ seconds. Because the precession frequency of spin 2 is shifted by $\pm J_{12}/2$ depending on whether spin 1 is in $|1\rangle$ or $|0\rangle$ (see Fig. 4), spin 2 will arrive in $1/2J$ seconds at either $+\hat{y}$ or $-\hat{y}$, depending on the state of spin 1. Finally, a 90° pulse on spin 2 about the $\hat{x}$ axis rotates spin 2 back to $+\hat{z}$ if spin 1 is $|0\rangle$, or to $-\hat{z}$ if spin 1 is in $|1\rangle$.

The net result is that spin 2 is flipped if and only if spin 1 is in $|1\rangle$, which corresponds exactly to the classical truth table for the CNOT. The extra $\hat{z}$ rotations in Eq. (31) are needed to give all elements in U_{CNOT} the same phase, so the sequence works also for superposition input states.

An alternative implementation of the CNOT gate, up to a relative phase factor, consists of applying a line-selective 180° pulse at $\omega_0^2 + J_{12}/2$ (see Fig. 4). This pulse inverts spin 2 (the target qubit) if and only if spin 1 (the control) is $|1\rangle$ (Cory, Price, and Havel, 1998). In general, if a spin is coupled to more than one other spin, half the lines in the multiplet must be selectively inverted in order to realize a CNOT. Extensions to doubly controlled NOT's are straightforward: in a three-qubit system, for example, this can be realized through inversion of one out of the eight lines (Freeman, 1998). As long as all the lines are resolved, it is in principle possible to invert any subset of the lines. Demonstrations using very long multifrequency pulses have been performed with up to five qubits (Khitrin *et al.*, 2002). However, this approach cannot be used whenever the relevant lines in the multiplet fall on top of each other.

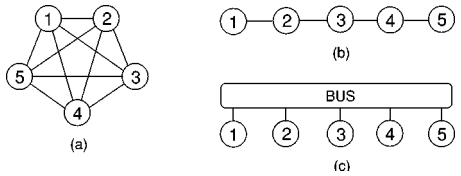

FIG. 10. Three possible coupling networks between five qubits. (a) A full coupling network. Such networks will in practice always be limited in size, as physical interactions tend to decrease with distance. (b) A nearest-neighbor coupling network. Such linear chains with nearest-neighbor couplings or two-dimensional variants are used in many solid-state proposals. (c) Coupling via a "bus." This network is used in ion-trap schemes, for example. As in case (a), the bus degree of freedom will in reality couple well to only a finite number of qubits.

If the spin-spin interaction Hamiltonian is not of the form $I_z^i I_z^j$ but contains also transverse components [as in Eqs. (5), (6), and (8)], other sequences of pulses are needed to perform the CPHASE and CNOT gates. These sequences are somewhat more complicated (Bremner *et al.*, 2002).

If two spins are not directly coupled to each other, it is still possible to perform a CNOT gate between them, as long as there exists a network of couplings that connects the two qubits. For example, suppose we want to perform a CNOT gate with qubit 1 as the control and qubit 3 as the target, CNOT_{13}, but 1 and 3 are not coupled to each other. If both are coupled to qubit 2, as in the coupling network of Fig. 10(b), we can first swap the states of qubits 1 and 2 (via the sequence CNOT_{12} CNOT_{21} CNOT_{12}), then perform a CNOT_{23}, and finally swap qubits 1 and 2 again (or relabel the qubits without swapping back). The net effect is CNOT_{13}. By extension, at most $O(n)$ SWAP operations are required to perform a CNOT between any pair of qubits in a chain of n spins with just nearest-neighbor couplings [Fig. 10(b)]. SWAP operations can also be used to perform two-qubit gates between any two qubits that are coupled to a common "bus" qubit [Fig. 10(c)].

Conversely, if a qubit is coupled to many other qubits [Fig. 10(a)] and we want to perform a CNOT between just two of them, we must remove the effect of the remaining couplings. This can accomplished using the technique of refocusing, which has been widely adopted in a variety of NMR experiments.

4. Refocusing: Turning off undesired $I_z^i I_z^j$ couplings

The effect of coupling terms during a time interval of free evolution can be removed via so-called "refocusing" techniques. For coupling Hamiltonians of the form $I_z^i I_z^j$, as is often the case in liquid NMR experiments [see Eq. (9)], the refocusing mechanism can be understood at a very intuitive level. Reversal of the effect of coupling Hamiltonians of other forms, such as in Eqs. (5), (6), and (8), is less intuitive, but can be understood within the

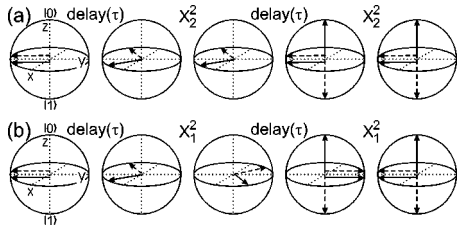

FIG. 11. Bloch-sphere representation of the operation of two simple schemes to refocus the coupling between two coupled qubits. The diagram shows the evolution of qubit 1 (in the rotating frame) initially along $-\hat{y}$, when qubit 2 is in $|0\rangle$ (solid) or in $|1\rangle$ (dashed). The refocusing pulse can be applied to either (a) qubit 2 or (b) qubit 1.

framework of average Hamiltonian theory (Sec. IV.C).

Let us first look at two ways of undoing $I_z^i I_z^j$ in a two-qubit system. In Fig. 11(a), the evolution of qubit 1 in the first time interval τ is reversed in the second time interval, due to the 180° pulse on qubit 2. In Fig. 11(b), qubit 1 continues to evolve in the same direction all the time, but the first 180° pulse causes the two components of qubit 1 to be refocused by the end of the second time interval. The second 180° pulse ensures that both qubits always return to their initial state.

Mathematically, we can see how refocusing of J couplings works using the fact that for all τ

$$X_1^2 U_J(\tau) X_1^2 = U_J(-\tau) = X_2^2 U_J(\tau) X_2^2, \tag{32}$$

which leads to

$$X_1^2 U_J(\tau) X_1^2 U_J(\tau) = I = X_2^2 U_J(\tau) X_2^2 U_J(\tau). \tag{33}$$

Replacing all X_i^2 with Y_i^2, the sequence works just the same. However, if we sometimes use X_i^2 and sometimes Y_i^2, we get the identity matrix only up to some phase shifts. Also, if we applied pulses on both qubits simultaneously, e.g., $X_1^2 X_2^2\ U_J(\tau) X_1^2 X_2^2\ U_J(\tau)$, the coupling would not be removed.

Figure 12 gives insight into refocusing techniques in a multiqubit system. Specifically, this scheme preserves the effect of J_{12}, while effectively inactivating all the other couplings. The underlying idea is that a coupling between spins i and j acts "forward" during intervals where both spins have the same sign in the diagram, and acts "in reverse" whenever the spins have opposite signs. Whenever a coupling acts forward and in reverse for the same duration, it has no net effect.

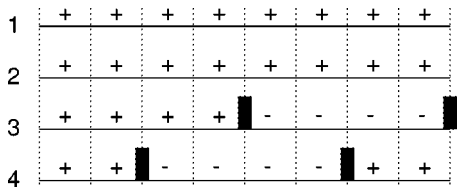

FIG. 12. Refocusing scheme for a four-spin system, designed to leave J_{12} active the whole time but to neutralize the effect of the other J_{ij}. The interval is divided into slices of equal duration, and the "+" and "−" signs indicate whether a spin is still in its original position, or upside down. The black rectangles represent 180° pulses, which flip the corresponding spin.

Systematic methods for designing refocusing schemes for multiqubit systems have been developed specifically for the purpose of quantum computing. The most compact scheme is based on Hadamard matrices (Jones and Knill, 1999; Leung *et al.*, 2000). A Hadamard matrix of order n, denoted by $H(n)$, is an $n \times n$ matrix with entries ± 1 such that

$$H(n) H(n)^T = nI. \tag{34}$$

The rows are thus pairwise orthogonal, and any two rows agree in exactly half of the entries. Identifying $+1$ and -1 with $+$ and $-$ as in the diagram of Fig. 12, we see that $H(n)$ gives a valid decoupling scheme for n spins using only n time intervals. An example of $H(12)$ is

$$\begin{bmatrix}
+ & + & + & + & + & + & + & + & + & + & + & + \\
+ & + & + & - & - & + & - & - & + & - & - & + \\
+ & + & + & + & - & - & - & + & - & + & - & - \\
+ & - & + & + & + & - & - & - & + & - & + & - \\
+ & - & - & + & + & + & - & - & - & + & - & + \\
+ & + & - & - & + & + & - & + & - & - & + & - \\
+ & - & - & - & - & - & - & + & + & + & + & + \\
+ & - & + & - & - & + & + & - & - & + & + & - \\
+ & + & - & + & - & - & + & - & - & - & + & + \\
+ & - & + & - & + & - & + & + & - & - & - & + \\
+ & - & - & + & - & + & + & + & + & - & - & - \\
+ & + & - & - & + & - & + & - & + & + & - & -
\end{bmatrix}. \tag{35}$$

If we want the coupling between one pair of qubits to remain active while removing the effect of all other couplings, we can simply use the same row of $H(n)$ for those two qubits.

$H(n)$ does not exist for all n, but we can always find a decoupling sequence for n qubits by taking the first n rows of $H(\bar{n})$, with $\bar{n}$ the smallest integer that satisfies $\bar{n} \geqslant n$ with known $H(\bar{n})$. From the properties of Hadamard matrices, we can show that $\bar{n}/n$ is always close to 1 (Leung *et al.*, 2000). So decoupling schemes for n spins require $\bar{n}$ time intervals and no more than $n\bar{n}$ 180° pulses.

Another systematic approach to refocusing sequences is illustrated via the following four-qubit scheme (Linden, Barjat, *et al.*, 1999):

$$\begin{bmatrix}
+ & + & + & + & + & + & + & + \\
+ & + & + & + & - & - & - & - \\
+ & + & - & - & - & - & + & + \\
+ & - & - & + & + & - & - & +
\end{bmatrix}. \tag{36}$$

For every additional qubit, the number of time intervals is doubled, and 180° pulses are applied to this qubit after the first, third, fifth, … time interval. The advantage of this scheme over schemes based on Hadamard matrices

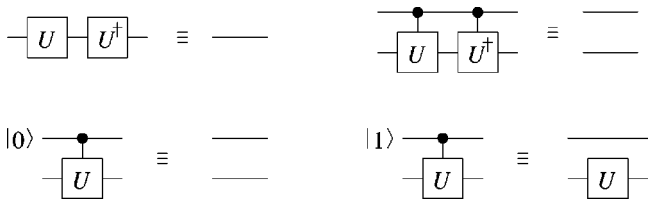

FIG. 13. Simplification rules for quantum circuits, drawn using standard quantum gate symbols, where time goes from left to right, each wire represents a qubit, boxes represent simple gates, and solid black dots indicate control terminals.

is that it does not require simultaneous rotations of multiple qubits. The main drawback is that the number of time intervals increases exponentially.

We end this subsection with three additional remarks. First, each qubit will generally be coupled to no more than a fixed number of other qubits, since coupling strengths tend to decrease with distance. In this case, all refocusing schemes can be greatly simplified (Jones and Knill, 1999; Linden, Barjat, *et al.*, 1999; Leung *et al.*, 2000).

Second, if the forward and reverse evolutions under J_{ij} are not equal in duration, a net coupled evolution takes place corresponding to the excess forward or reverse evolution. In principle, therefore, we can organize any refocusing scheme such that it incorporates any desired amount of coupled evolution for each pair of qubits.

Third, refocusing sequences can also be used to remove the effect of I_z^i terms in the Hamiltonian. Of course, these terms vanish in principle if we work in the multiply rotating frame [see Eq. (21)]. However, there may be some spread in the Larmor frequencies, for instance, due to magnetic-field inhomogeneities. This effect can then be reversed using refocusing pulses, as is routinely accomplished in spin-echo experiments (Sec. V.A.4).

5. Pulse sequence simplification

There are many possible pulse sequences which in an ideal world result in exactly the same unitary transformation. Good pulse sequence design therefore attempts to find the shortest and most effective pulse sequence that implements the desired transformations. In Sec. IV, we shall see that the use of more complex pulses or

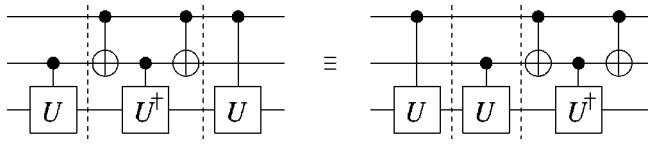

FIG. 14. Commutation of unitary operators can help simplify quantum circuits by moving building blocks around such that cancellation of operations as in Fig. 13 becomes possible. For example, the three segments (separated by dashed lines) in these two equivalent realizations of the TOFFOLI gate (doubly-controlled NOT) commute with each other and can thus be executed in any order.

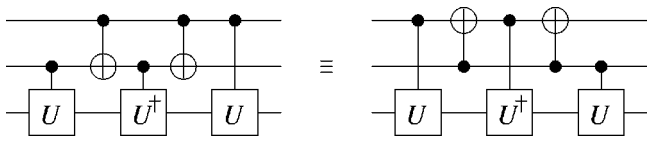

FIG. 15. Choosing one of several equivalent implementations can help simplify quantum circuits, again by enabling cancellation of operations as in Fig. 13. For instance, the two control qubits in the TOFFOLI gate play equivalent roles, so they can be interchanged.

pulse sequences may sometimes increase the degree of quantum control. Here, we look at three levels of pulse sequence simplification.

At the most abstract level of pulse sequence simplification, careful study of a quantum algorithm can give insight into how to reduce the resources needed. For example, a key step in both the modified Deutsch-Jozsa algorithm (Cleve *et al.*, 1998) and the Grover algorithm (Grover, 1997) can be described as the transformation $|x\rangle|y\rangle \rightarrow |x\rangle|x \oplus y\rangle$, where $|y\rangle$ is set to $(|0\rangle - |1\rangle)/\sqrt{2}$, so that the transformation in effect is $|x\rangle(|0\rangle - |1\rangle)/\sqrt{2} \rightarrow (-1)^{f(x)}|x\rangle(|0\rangle - |1\rangle)/\sqrt{2}$. Thus we might as well leave out the last qubit as it is never changed.

At the next level, that of quantum circuits, we can use simplification rules such as those illustrated in Fig. 13. In this process, we can fully take advantage of commutation rules to move building blocks around, as illustrated in Fig. 14. Furthermore, gates that commute with each other can be executed simultaneously. Finally, we can take advantage of the fact that most building blocks have many equivalent implementations, as shown, for instance, in Fig. 15.

Sometimes, a quantum gate may be replaced by another quantum gate, which is easier to implement. For instance, refocusing sequences (Sec. III.A.4) can be kept simple by examining which couplings really need to be refocused. Early on in a pulse sequence, several qubits may still be along $\pm\hat{z}$, in which case their mutual $I_z^i I_z^j$ couplings have no effect and thus need not be refocused. Similarly, if a subset of the qubits can be traced out at some point in the sequence, the mutual interaction between these qubits does not matter anymore, so only their coupling with the remaining qubits must be refocused. Figure 16 gives an example of such a simplified refocusing scheme for five coupled spins.

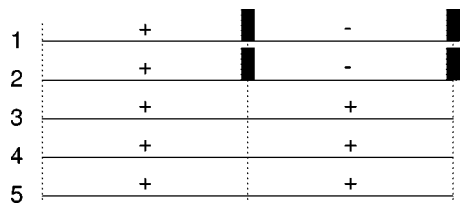

FIG. 16. Simplified refocusing scheme for five spins, designed such that the coupling of qubits 1-2 with qubits 3-5 is switched off, i.e., J_{13}, J_{14}, J_{15}, J_{23}, J_{24}, and J_{25} are inactive whereas J_{12}, J_{34}, J_{35}, and J_{45} are active.

More generally, the relative phases between the entries in the unitary matrix describing a quantum gate are irrelevant when the gate acts on a diagonal density matrix. In this case, we can, for instance, implement a CNOT simply as $X_2U_J(1/2J)Y_2$ rather than the sequence of Eq. (31).

At the lowest level, that of pulses and delay times, further simplification is possible by taking out adjacent pulses which cancel out, such as X and $\bar{X}$ (an instance of the first simplification rule of Fig. 13) and by converting "difficult" operations to "easy" operations.

Cancellation of adjacent pulses can be maximized by properly choosing the pulse sequences for subsequent quantum gates. For this purpose, it is convenient to have a library of equivalent implementations for the most commonly used quantum gates. For example, two equivalent decompositions of a CNOT_{12} gate (with $J_{12}>0$) are

$$Z_1\bar{Z}_2X_2U_J\left(\frac{1}{2J}\right)Y_2, \tag{37}$$

as in Eq. (31), and

$$\bar{Z}_1\bar{Z}_2\bar{X}_2U_J\left(\frac{1}{2J}\right)\bar{Y}_2. \tag{38}$$

Similarly, two equivalent implementations of the HADAMARD gate on qubit 2 are

$$X_2^2Y_2 \tag{39}$$

and

$$\bar{Y}_2X_2^2. \tag{40}$$

Thus, if we need to perform a HADAMARD operation on qubit 2 followed by a CNOT_{12} gate, it is best to choose the decompositions of Eqs. (37) and (40), such that the resulting pulse sequence,

$$Z_1\bar{Z}_2X_2U_J\left(\frac{1}{2J}\right)Y_2 \quad \bar{Y}_2X_2^2, \tag{41}$$

simplifies to

$$Z_1\bar{Z}_2X_2U_J\left(\frac{1}{2J}\right)X_2^2. \tag{42}$$

An example of a set of operations that is easy to perform is the rotations about $\hat{z}$. While the implementation of $\hat{z}$ rotations in the form of three RF pulses [Eq. (28)] takes more work than a rotation about $\hat{x}$ or $\hat{y}$, rotations about $\hat{z}$ need in fact not be executed at all, provided the coupling Hamiltonian is of the form $I_z^iI_z^j$, as in Eq. (21). In this case, $\hat{z}$ rotations commute with free evolution under the system Hamiltonian, so we can interchange the order of $\hat{z}$ rotations and time intervals of free evolution. Using equalities such as

$$Z\bar{Y}=XY\bar{X}\bar{Y}=XZ, \tag{43}$$

we can also move $\hat{z}$ rotations across $\hat{x}$ and $\hat{y}$ rotations, and gather all $\hat{z}$ rotations at the end or the beginning of a pulse sequence. At the end, $\hat{z}$ rotations do not affect the outcome of measurements in the usual $|0\rangle$, $|1\rangle$ "computational" basis. Similarly, $\hat{z}$ rotations at the start of a pulse sequence have no effect on the usually diagonal initial state. In either case, Z rotations do not then require any physical pulses and are in a sense "for free" and perfectly executed. Indeed, Z rotations simply define the reference frame for $\hat{x}$ and $\hat{y}$ and can be implemented by changing the phase of the reference frame throughout the pulse sequence.

It is thus advantageous to convert as many X and Y rotations as possible into Z rotations, using identities similar to Eq. (28), for example,

$$XY=XY\bar{X}X=ZX. \tag{44}$$

A key point in pulse sequence simplification of any kind is that the simplification process must itself be efficient. For example, suppose an algorithm acts on five qubits with initial state $|00000\rangle$ and outputs the final state $(|01000\rangle+|01100\rangle/\sqrt{2}$. The overall result of the algorithm is thus that qubit 2 is flipped and that qubit 3 is placed in an equal superposition of $|0\rangle$ and $|1\rangle$. This net transformation can obviously be obtained immediately by the sequence $X_2^2Y_3$. However, the effort needed to compute the overall input-output transformation generally increases exponentially with the problem size, so such extreme simplifications are not practical.

6. Time-optimal pulse sequences

Next to the widely used but rather naive set of pulse sequence simplification rules of the previous subsection, there exist powerful mathematical techniques for determing the minimum time needed to implement a quantum gate, using a given system and control Hamiltonian, as well as for finding time-optimal pulse sequences (Khaneja *et al.*, 2001). These methods build on earlier optimization procedures for mapping an initial operator onto a final operator via unitary transformations (Sørenson, 1989; Glaser *et al.*, 1998), as in coherence or polarization transfer experiments, common tasks in NMR spectroscopy.

The pulse sequence optimization technique expresses pulse sequence design as a geometric problem in the space of all possible unitary transformations. The goal is to find the shortest path between the identity transformation I and the point in the space corresponding to the desired quantum gate, U, while traveling only in directions allowed by the given system and control Hamiltonian. Let us call K the set of all unitaries k that can be produced using the control Hamiltonian only. Next we assume that the terms in the control Hamiltonian are much stronger than the system Hamiltonian (as we shall see in Sec. III.B.2, this assumption is valid in NMR only when using so-called hard, high-power pulses). Then, starting from I, any point in K can be reached in a negligibly short time, and similarly, U can be reached in no time from any point in the coset KU, defined by $\{kU|k\in K\}$. Evolution under the system Hamiltonian for a fi-

nite amount of time is required to reach the coset KU starting from K. Finding a time-optimal sequence for U thus comes down to finding the shortest path from K to KU allowed by the system Hamiltonian.

Such optimization problems have been extensively studied in mathematics (Brockett, 1981) and have been solved explicitly for elementary quantum gates on two coupled spins (Khaneja *et al.*, 2001) and a three-spin chain with nearest-neighbor couplings (Khaneja *et al.*, 2002). For example, a sequence was found for producing the trilinear propagator $\exp(-i2\pi I_z^1 I_z^2 I_z^3)$ from the system Hamiltonian $\hbar 2\pi J(I_z^1 I_z^2 + I_z^2 I_z^3)$ in a time $\sqrt{3}/2J$, the shortest possible time (Khaneja *et al.*, 2002). This propagator is the starting point for useful quantum gates such as the doubly controlled NOT or TOFFOLI gate. The standard quantum circuit approach, in comparison, would yield a sequence of duration $3/2J$ (it uses only one coupling at a time while refocusing the other coupling), and the common NMR pulse sequence has duration $1/J$.

Clearly, the time needed to *find* a time-optimal pulse sequence increases exponentially with the number of qubits n involved in the transformation, since the unitary matrices involved are of size $2^n \times 2^n$. Therefore the main use of the techniques presented here lies in finding efficient pulse sequences for building blocks acting on only a few qubits at a time, which can then be incorporated in more complex sequences acting on many qubits by adding appropriate refocusing pulses to remove the couplings with the remaining qubits. While the examples given here are for the typical NMR system and control Hamiltonian, the approach is completely general and may be useful for other qubit systems too.

B. Experimental limitations

Many years of experience have taught NMR spectroscopists that while the ideal control techniques described above are theoretically attractive, they neglect important experimental artifacts and undesired Hamiltonian terms which must be addressed in any actual implementation. First, a pulse intended to selectively rotate one spin will to some extent also affect the other spins. Second, the coupling terms $2\pi J_{ij} I_z^i I_z^j$ cannot be switched off in NMR. During time intervals of free evolution under the system Hamiltonian, the effect of these coupling terms can easily be removed using refocusing techniques (Sec. III.A.4), so long as the single-qubit rotations are perfect and instantaneous. However, during RF pulses of finite duration, the coupling terms also distort the single-qubit rotations. In addition to these two limitations arising from the NMR system and control Hamiltonian, a number of instrumental imperfections cause additional deviations from the intended transformations.

1. Cross-talk

Throughout the discussion of single- and two-qubit gates, we have assumed that we can selectively address each qubit. Experimentally, qubit selectivity could be accomplished if the qubits were well separated in space or, as in NMR, in frequency. In practice, there will usually be some cross-talk, which causes an RF pulse applied on resonance with one qubit to slightly rotate another qubit or shift its phase. Cross-talk effects are even more complex when two or more pulses are applied simultaneously.

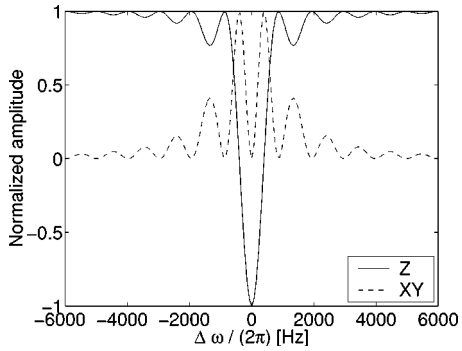

FIG. 17. Simulation of the spin response to a 1-ms constant-amplitude RF pulse as a function of the frequency offset $\Delta\omega$ between ω_0 and ω_{rf}. The spin starts off in $|0\rangle$ (along $+\hat{z}$ in the Bloch sphere) and $\omega_1/2\pi = 500$ Hz is chosen such that the rotation angle amounts to 180° for an on-resonance pulse.

The frequency bandwidth over which qubits are rotated by a pulse of length t_{pw} is roughly speaking of order $1/t_{pw}$. Yet, since the qubit response to an RF field is not linear (it is sinusoidal in $\omega_1 t_{pw}$), the exact frequency response cannot be computed using Fourier theory.

For a constant-amplitude (rectangular) pulse, the unitary transformation as a function of the detuning $\Delta\omega$ is easy to derive analytically from Eqs. (16) and (17). Alternatively, we can exponentiate the Hamiltonian of Eq. (15) to get U directly. An example of a qubit response to a rectangular pulse is shown in Fig. 17.

It is evident from Fig. 17 that short rectangular pulses (known as "hard" pulses) excite spins over a very wide frequency range. The frequency selectivity of a pulse can of course be increased by increasing t_{pw} while lowering B_1 accordingly (thus creating what is known as a "soft" pulse), but decoherence effects become more severe as the pulses get longer. Fortunately, as we shall see in Secs. IV.A and IV.B, the use of shaped and composite pulses can dramatically improve the frequency selectivity of the RF excitation.

Even if a pulse is designed not to produce any net $\hat{x}$ or $\hat{y}$ rotations of spins outside a specified frequency window, the presence of RF irradiation during the pulse still causes a shift $\Delta\omega_{BS}^i$ in the precession frequency of spins i at frequencies well outside the excitation frequency window (Emsley and Bodenhausen, 1990). As a result, each spin accumulates a spurious phase shift during RF pulses applied to spins at nearby frequencies.

This effect is related to the Bloch-Siegert shift mentioned in Sec. II.B.1 and is known as the *transient generalized Bloch-Siegert shift* in the NMR community. It is related to the ac Stark effect in atomic physics. At a deeper level, the acquired phase can be understood as

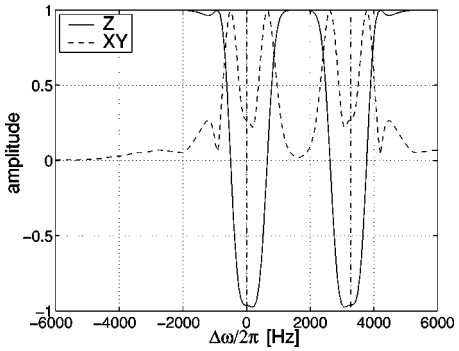

FIG. 18. Simulation of the spin response to two simultaneous pulses with carrier frequencies at 0 Hz and 3273 Hz (vertical dashed lines) away from the spin-resonance frequency, with a calibrated pulse length of 2650 μs (as for an ideal 180°). The amplitude profile of the pulses is Hermite shaped (Sec. IV.A) in order to obtain a smooth spin response. For ideal inversion, the solid line should be −1 at the two frequencies, and the dashed line should be zero. From Steffen *et al.*, 2000.

an instance of Berry's phase (Berry, 1984): the spin describes a closed trajectory on the surface of the Bloch sphere and thus returns to its initial position, but it acquires a phase shift proportional to the area enclosed by its trajectory.

The frequency shift is given by

$$\Delta\omega_{BS} \approx \frac{\omega_1^2}{2(\omega_0 - \omega_{rf})} \tag{45}$$

(provided $\omega_1 \ll |\omega_0 - \omega_{rf}|$), where $\omega_0/2\pi$ is the original Larmor frequency (in the absence of the RF field). In typical NMR experiments, the frequency shifts can easily reach several hundred Hz in magnitude. We see from Eq. (45) that the Larmor frequency shifts up if $\omega_0 > \omega_{rf}$ and shifts down if $\omega_0 < \omega_{rf}$.

Fortunately, the resulting phase shifts can be easily computed in advance for each possible spin-pulse combination, if all the frequency separations, pulse amplitude profiles, and pulse lengths are known. The unintended phase shifts $R_z(\theta)$ can then be compensated for during the execution of a pulse sequence by inserting appropriate $R_z(-\theta)$, which can be executed at no cost, as we saw in Sec. III.A.5.

Cross-talk effects are aggravated during simultaneous pulses, applied to two or more spins with nearby frequencies ω_0^1 and ω_0^2 (say $\omega_0^1 < \omega_0^2$). The pulse at ω_0^1 then temporarily shifts the frequency of spin 2 to $\omega_0^2 + \Delta\omega_{BS}$. As a result, the pulse on spin 2, if applied at ω_0^2, will be off-resonance by an amount $-\Delta\omega_{BS}$. Analogously, the pulse at ω_0^1 is now off the resonance of spin 1 by $\Delta\omega_{BS}$. The resulting rotations of the spins deviate significantly from the intended rotations.

The detrimental effect of the Bloch Siegert shifts during simultaneous pulses is illustrated in Fig. 18, which shows the simulated inversion profile for a spin subject to two simultaneous 180° pulses separated by 3273 Hz. The centers of the inverted regions have shifted away from the intended frequencies and the inversion is incomplete, which can be seen most clearly from the substantial residual $\hat{x}$–$\hat{y}$-magnetization (>30%) over the whole region intended to be inverted. Note also that since the frequencies of the applied pulses are off the spin resonance frequencies, complete inversion cannot be achieved no matter what tip angle is chosen (see Sec. II.B.2).

In practice, simultaneous soft pulses at nearby frequencies have been avoided in NMR (Linden, Kupče, and Freeman, 1999) or the poor quality of the spin rotations was accepted. Pushed by the stringent requirements of quantum computation, several techniques have meanwhile been invented to generate accurate simultaneous rotations of spins at nearby frequencies (see Secs. IV.A.2 and IV.B.2).

2. Coupled evolution

The spin-spin couplings in a molecule are essential for the implementation of two-qubit gates (Sec. III.A.3), but they cannot be turned off and are thus also active during the RF pulses, which are intended to be just single-qubit transformations. Unless ω_1 is much stronger than the coupling strength, the interactions strongly affect the intended nutation. For couplings of the form $JI_z^iI_z^j$, the effect is similar to the off-resonance effects illustrated in Fig. 7: the coupling to another spin shifts the spin frequency to $\omega_0/2\pi \pm J/2$, so a pulse sent at $\omega_0/2\pi$ hits the spin off-resonance by $\mp J/2$.

In practice, J coupling terms can only be neglected for short, high-power pulses used in heteronuclear spin systems: typically $J < 300$ Hz while ω_1 is up to ≈ 50 kHz. For low-power pulses, often used in homonuclear spin systems, ω_1 can be of the same order as J and coupling effects become prominent. The coupling terms also lead to additional complications when two qubits are pulsed simultaneously. In general, the qubits become partially entangled (Kupče and Freeman, 1995).

As was the case for cross-talk, NMR spectroscopists have developed special shaped and composite pulses to compensate for coupling effects during RF pulses while performing spin-selective rotations. In recent years, the use of such pulses has been extended and perfected for quantum computing experiments (Secs. IV.A and IV.B).

3. Instrumental errors

A number of experimental imperfections lead to errors in the quantum gates. In NMR, the most common imperfections are inhomogeneities in the static and RF magnetic field, pulse length calibration errors, frequency offsets, and pulse timing and phase imperfections.

The static field B_0 in modern NMR magnets can be made homogeneous over the sample volume (a cylinder 5 mm in diameter and 1.5 cm long) to better than 1 part in 10^9. This amazing homogeneity is obtained by meticulously adjusting the current through a set of so-called "shim" coils, which compensate for the inhomogeneities produced by the large solenoid. At $\omega_0 = 500 \cdot 2\pi$ MHz, linewidths of 0.5 Hz can thus be obtained, correspond-

ing to a dephasing time constant T_2^* (see Sec. V.A.2) of $1/(2\pi 0.5)=0.32$ s. In the course of long pulse sequences (of order 0.1–1 s), even the tiny remaining inhomogeneity would therefore have a large effect, so its effect must be reversed using refocusing sequences (Sec. III.A.4)

The RF field homogeneity is typically very poor, due to constraints on the geometry of the RF coils: the envelope of Rabi oscillations (Sec. V.A.1) often decays by as much as 5% per 90° rotation, corresponding to a quality factor of only ≈5. In sequences containing only a few pulses, this is not problematic, but in multiple-pulse experiments, the RF field inhomogeneity is often the dominant source of errors and signal loss.

Imperfect pulse length calibration has an effect similar to B_1 inhomogeneity: the qubit rotation angle is different than was intended. Only the correlation time for the error is different. Miscalibrations are constant throughout an experiment, whereas the RF field experienced by any given molecule changes on the time scale of diffusion through the sample volume.

Frequency offsets occur in different contexts. In traditional NMR experiments, the Larmor frequencies are often not known in advance. RF pulses are then expected to rotate the spins over a wide range of frequencies, quite the opposite case to that of quantum computing, where the Larmor frequencies are precisely known and rotations should be spin selective. However, we have seen earlier that $I_z^i I_z^j$ coupling terms act as a frequency offset of one spin, which depends on the state of the other spin. Qubit-selective rotations of qubit i thus require a uniform rotation over a range $\omega_0^i \pm \Sigma_{j\neq i} |J_{ij}|/2$.

Various approaches have been developed to reduce the sensitivity of RF pulses and pulse sequences to these instrumental errors, sometimes in combination with solutions to cross-talk and coupling artifacts. These advanced techniques are the subject of the next section.

IV. ADVANCED PULSE TECHNIQUES

The accuracy of quantum gates that can be achieved using the simple pulse techniques of the previous section is unsatisfactory when applied to multispin systems, where the given NMR system and control Hamiltonian lead to undesired cross-talk and coupling effects. In addition, the available instrumentation can only imperfectly approximate ideal pulse amplitudes, timings, and phases, for realistic sample geometries and coil configurations, and any real molecule includes additional Hamiltonian terms such as couplings to the environment, which are undesired. Nevertheless, extremely precise control can be achieved despite these imperfections, and this is accomplished using the art of *shaped pulses, composite pulses,* and *average Hamiltonian theory*, the subject of this second major section of this review.

These advanced techniques are based on the assumption that errors are, at least on some accessible time scale, systematic, rather than random. This assumption clearly holds for the terms in the ideal NMR Hamiltonian of Eqs. (21) and (22), and applies also to most instrumental errors. Then, by using the special properties of evolution in unitary groups, such as the SU(2^n) which describes the space of operators acting on n qubits, the systematic errors can in principle be canceled out.

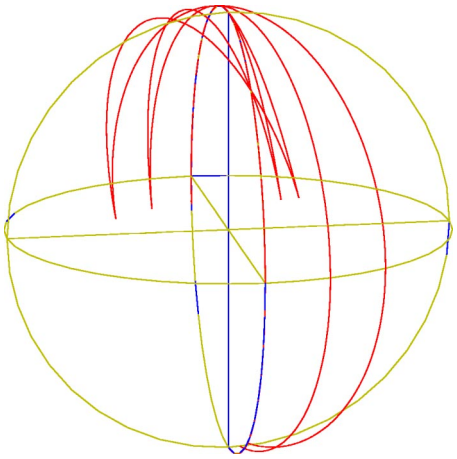

FIG. 19. (Color in online edition) Trajectory on the Bloch sphere of a qubit initially in $|0\rangle$, when a so-called IBURP1 pulse (Geen and Freeman, 1991) is applied, of duration 1 ms and $\omega_1=3342$ Hz, with a frequency offset (analogous to $I_z^i I_z^j$ coupling) of 0, 100, and 200 Hz. This pulse is intended to rotate the qubit from $|0\rangle$ $(+\hat{z})$ to $|1\rangle$ $(-\hat{z})$. We see that the effect of the frequency offset is largely removed by the specially designed pulse shape; all three trajectories terminate near $-\hat{z}$.

A. Shaped pulses

The amplitude and phase profile of RF pulses can be specially tailored in order to ease the cross-talk and coupling effects discussed in Secs. III.B.1 and III.B.2. In practice, the pulse is divided into a few tens to many hundreds of discrete time slices; to achieve an arbitrarily shaped pulse, it suffices to control the amplitude and phase of the slices separately. Furthermore, multiple shaped pulses applied at various frequencies can be combined into a single pulse shape, since a linear vector sum of pulse slices also results in a valid pulse. Here, we consider simple amplitude and phase shaped pulses.

1. Amplitude profiles

The frequency selectivity of RF pulses can be much improved compared to standard rectangular pulses with sharp edges, by using pulse shapes that smoothly modulate the pulse amplitude with time. Such pulses are typically especially designed to excite or invert spins over a limited frequency region, while minimizing $\hat{x}$ and $\hat{y}$ rotations for spins outside this region (Freeman, 1997, 1998).

Furthermore, specialized pulse shapes exist which minimize the effect of couplings during the pulses. Such self-refocusing pulses (Geen and Freeman, 1991) take a spin over a complicated trajectory in the Bloch sphere, in such a way that the net effect of couplings between the selected and nonselected spins is reduced (Fig. 19). It

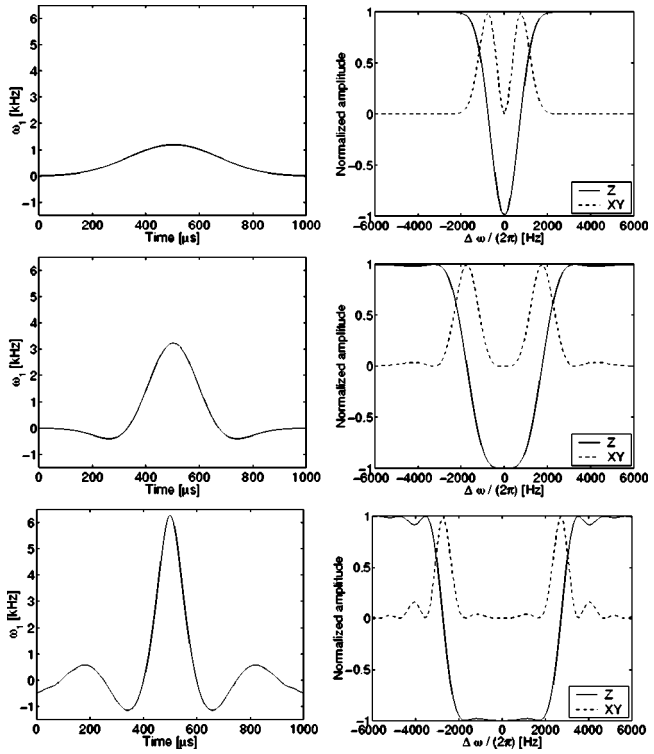

FIG. 20. (Left) Time profile for, from top to bottom, a 1-ms Gaussian, Hermite 180, and REBURP shaped pulse. (Right) Corresponding frequency response of a qubit initially along $+\hat{z}$, displaying the $\hat{z}$ and $\hat{x}$-$\hat{y}$ components of the qubit after the pulse.

is as if those couplings are only in part or even not at all active during the pulse (couplings between pairs of nonselected spins will still be fully active but their effect can be removed using the standard refocusing techniques described in Sec. III.A.4). As a general rule, it is relatively easy to make 180° pulses self-refocusing, but much harder to do so for 90° pulses.

The self-refocusing behavior of certain shaped pulses can be intuitively understood to some degree. Nevertheless, many actual pulse shapes have been the result of numerical optimizations. Often, the pulse shape is expressed in a basis of several functions, for instance, a Fourier series (Geen and Freeman, 1991),

$$\omega_1(t) = \left\{ A_0 + \sum_n \left[A_n \cos\left(n\frac{2\pi}{t_{pw}} t \right) + B_n \sin\left(n\frac{2\pi}{t_{pw}} t \right) \right] \right\}, \tag{46}$$

and the weights of the basis functions, A_n and B_n, are optimized using numerical routines such as simulated annealing.

Comparison of the performance of various pulse shapes is facilitated by computing the corresponding spin responses. This is most easily done by concatenating the unitary operators of each time slice of the shaped pulse, as the Hamiltonian is time independent within each time slice. Figure 20 presents the amplitude profile and pulse response for three standard pulse shapes of equal duration, illustrating that different pulse shapes produce strikingly different spin response profiles.

Properties relevant for choosing a pulse shape include:

- frequency selectivity: product of excitation bandwidth and pulse length (lower is more selective),
- transition range: the width of the transition region between the selected and nonselected frequency region,
- power: the peak power required for a given pulse length and tip angle (lower is less demanding),
- self-refocusing behavior: degree to which the J coupling between the selected spin and other spins is refocused (the signature for self-refocusing behavior is a flat top in the excitation profile),
- robustness to experimental imperfections such as pulse length errors,
- universality: whether the pulse performs the intended rotation for arbitrary input states or only for specific input states.

TABLE II. Properties of relevant pulse shapes. The Gaussian (Bauer *et al.*, 1984) and Hermite (Warren, 1984) shapes are described by analytical functions and were identified early on. The BURP family of pulses (Geen and Freeman, 1991) resulted from numerical optimization routines. Continued work in this area has produced a large number of additional pulse shapes, such as the AV 90 (Abramovich and Vega, 1993).

	Selectivity	Transition range	Power	Self-refocusing	Robustness
Rectangular	poor	very wide	minimal	no	good
Gauss 90	excellent	wide	low	fair	good
Gauss 180	excellent	wide	low	fair	good
Hrm 90	moderate	moderate	average	good	fair
Hrm 180	good	moderate	average	very good	fair
UBURP 90	poor	narrow	high	excellent	poor
REBURP180	poor	narrow	high	excellent	poor
AV 90	fair	moderate	average	good	fair

Table II summarizes these properties for a selection of widely used pulse shapes. Only universal pulses (also known as general-rotation pulses) are included in the table, since quantum computations must work for any input state.

Obviously, no single pulse shape optimizes for all properties simultaneously, so pulse shape design consists of finding the optimal tradeoff for the desired application. For quantum computing experiments, we can select molecules with large chemical shifts, so sharp transition regions are not so important. Furthermore, the probe and spectrometer can deal with relatively high powers. The crucial parameters are the self-refocusing behavior, the selectivity (short, selective pulses minimize decoherence), and to some extent the robustness.

It is also possible to start from a desired frequency response and invert the transformation to find the pulse shape that produces this response. Again, given the nonlinear nature of the response, the inverse transformation is not given by a Fourier transform, but it can nevertheless be computed directly (Pauly *et al.*, 1991).

Even self-refocusing shaped pulses do not generally remove the coupling terms completely. Furthermore, when two spins are pulsed simultaneously with self-refocusing pulses, the refocusing effects are often destroyed (Kupče and Freeman, 1995). In both cases, the remaining coupled evolution that takes place during the pulses must be reversed at an earlier and/or later stage in the pulse sequence.

If we could decompose the evolution during an actual pulse into an idealized, instantaneous X or Y rotation with no coupling present, followed and/or preceded by a time interval of free evolution, we could compensate for the coupling effects simply by adjusting the appropriate time intervals of free evolution in between the pulses (Sec. III.A.4). However, $\mathcal{H}_{rf}$ and $\mathcal{H}_J$ do not commute, so such a decomposition is not possible.

Nevertheless, the coupled evolution can still be unwound to first order (Knill *et al.*, 2000; Vandersypen *et al.*, 2001), when a time interval of reverse evolution both before and after the pulse is used:

$$e^{+i\mathcal{H}_J\tau/\hbar}e^{-i(\mathcal{H}_{rf}+\mathcal{H}_J)t_{pw}/\hbar}e^{+i\mathcal{H}_J\tau/\hbar} \approx e^{-i\mathcal{H}_{rf}t_{pw}/\hbar}, \tag{47}$$

where τ is chosen such that the approximations are as good as possible according to some distance or fidelity measure (see Sec. V.C). The optimal τ is usually close to but not equal to $t_{pw}/2$. In comparison, a negative time interval only before *or* after the pulse,

$$\begin{aligned} e^{+i\mathcal{H}_J\tau/\hbar}e^{-i(\mathcal{H}_{rf}+\mathcal{H}_J)t_{pw}/\hbar} &\approx e^{-i\mathcal{H}_{rf}\,t_{pw}/\hbar} \\ &\approx e^{-i(\mathcal{H}_{rf}+\mathcal{H}_J)t_{pw}/\hbar}e^{+i\mathcal{H}_J\tau/\hbar}, \end{aligned} \tag{48}$$

is much less effective.

2. Phase profiles

An alternative to amplitude shaping that is often useful is frequency or phase shaping. One specific phase-shaping method utilizes fixed, small increments $\Delta\phi$ to the phase of successive slices of a pulse to achieve an excitation profile centered at a frequency that differs from the RF carrier frequency ω_{rf} by $\Delta\phi/\Delta t$, where Δt is the duration of each time slice. This technique for shifting the RF frequency is known as phase ramping (Patt, 1991). We can express the effect of phase ramping mathematically by replacing Eq. (11) by

$$\begin{aligned} \mathcal{H}_{rf}/(-\hbar\omega_1) &= \cos\left[\omega_{rf}t + \left(\phi_0 + \frac{\Delta\phi}{\Delta t}t\right)\right]I_x \\ &\quad - \sin\left[\omega_{rf}t + \left(\phi_0 + \frac{\Delta\phi}{\Delta t}t\right)\right]I_y \\ &= \cos\left[\left(\omega_{rf} + \frac{\Delta\phi}{\Delta t}\right)t + \phi_0\right]I_x \\ &\quad - \sin\left[\left(\omega_{rf} + \frac{\Delta\phi}{\Delta t}\right)t + \phi_0\right]I_y. \end{aligned} \tag{49}$$

The use of phase shifts thus permits us to obtain an RF field at a different frequency than is generated by the signal generator. Furthermore, the displaced frequency can be chosen to be different for every pulse and can even be varied in the course of a pulse.

A useful application of phase ramping lies in compensation for Bloch-Siegert effects during simultaneous pulses, where the RF applied at ω_0^i shifts the resonance frequency of spin j to $\omega_0^j + \Delta\Omega_{BS}$ (Sec. III.B.1). The rotations of both spins can be significantly improved simply by shifting the RF excitation frequencies via phase ramping such that they track the shifts of the corresponding spin frequencies (Steffen *et al.*, 2000). In this way, the pulses are always applied on-resonance with the respective spins. The calculation of the frequency shift throughout a shaped pulse is straightforward and needs to be done only once, at the start of a series of experiments.

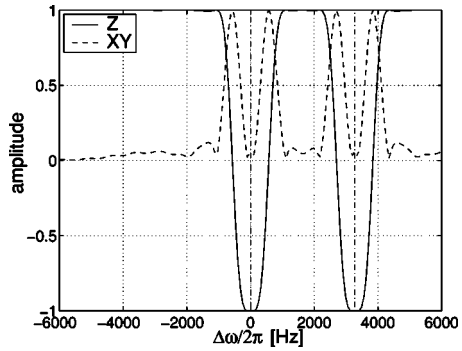

FIG. 21. Similar to Fig. 18 but with frequency shift correction. From Steffen *et al.*, 2000.

Figure 21 shows the simulated inversion profiles for the same conditions as in Fig. 18, but this time using the frequency shift corrected scheme. The inversion profiles are clearly much improved and there is very little left-over $\hat{x}-\hat{y}$ magnetization. Simulations of the inversion profiles for a variety of pulse shapes, pulse widths, and frequency separations confirm that the same technique can be used to correct the frequency offsets caused by three or more simultaneous soft pulses at nearby frequencies. The improvement is particularly pronounced

when the frequency window of the shaped pulse is two to eight times the frequency separation between the pulses (Steffen *et al.*, 2000).

B. Composite pulses

Another practical method for compensating for systematic control errors in NMR experiments is the application of a sequence of pulses instead of a single pulse. This method of *composite pulses* arises from the observation that concatenation of several pulses can produce more accurate rotations than is possible using just a single pulse, due to strategic cancellation of systematic errors and other unwanted systematic effects. Composite pulses work particularly well for compensating errors arising from the RF field inhomogeneity, frequency offsets, imperfect pulse length calibration, and other instrumental artifacts introduced in Sec. III.B.3. They leverage the ability to control one parameter precisely to compensate for the inability to control another parameter well. We describe two approaches to construction of composite pulses: an analytical method and one employing numerical optimization.

1. Analytical approach

The three parameters that characterize a hard pulse are its frequency offset $\Delta\omega$, phase ϕ, and area $\hbar\gamma B_1 t_{pw}$, given by the product of the pulse amplitude B_1 and pulse duration t_{pw} (Sec. II.B.1). In terms of qubit operations, errors in these parameters translate directly into errors in the axis $\hat{n}$ and angle θ of rotation, such that the actual operation applied is not the ideal $R_{\hat{n}}(\theta)$ of Eq. (23), but rather

$$\tilde{R}_{\hat{n}}(\theta) = \exp\left[-i\frac{f(\theta,\hat{n})\cdot\vec{\sigma}}{2}\right], \tag{50}$$

where $f(\theta,\hat{n})$ is a function which characterizes the systematic error. For example, under- and over-rotation errors caused by pulse amplitude miscalibration or RF field inhomogeneity may be described by $f(\theta,\hat{n})=\theta(1+\epsilon)\hat{n}$, while RF phase errors may be described by $f(\theta,\hat{n})=\theta[\hat{n}_x\cos\epsilon+\hat{n}_y\sin\epsilon,\hat{n}_y\cos\epsilon-\hat{n}_x\sin\epsilon,\hat{n}_z]$, where ϵ is a fixed, but unknown parameter. The essence of the composite pulses technique is that a number of erroneous operations are concatenated, varying $\hat{n}$ and θ, to obtain a final operation which is as independent of ϵ as possible. This is done without knowing ϵ.

This technique can be illustrated by considering the specific case of linear amplitude errors, in which

$$\tilde{R}_{\hat{n}}(\theta) = \exp\left[-\frac{i\theta(1+\epsilon)\hat{n}\cdot\vec{\sigma}}{2}\right]. \tag{51}$$

Let the goal be to obtain $R_x(\pi/2)$. Using as a measure of error the average gate fidelity, defined in Eq. (113), we find that $\bar{F}(R_x(\pi/2),\tilde{R}_x(\pi/2))=[2+\cos(\epsilon\pi/2)]/3\approx 1-\pi^2\epsilon^2/24$, so the error is quadratic in ϵ for small ϵ. Consider, in contrast, the sequence

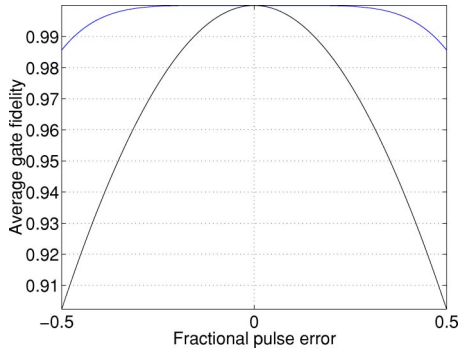

FIG. 22. (Color in online editon) Comparison of gate fidelities for average and composite unitary transforms: Lower curve, plot of the average gate fidelity between the ideal $R_x(\pi/2)$ and actual unitary transforms $\tilde{R}_x(\pi/2)$; upper curve, plot of fidelity between the ideal $R_x(\pi/2)$ and the composite sequence $BB1_{90}$, as a function of the fraction of over-rotation error ϵ. Note how much higher fidelity the $BB1$ sequence has (the best possible fidelity is 1), over a wide range of errors.

$$BB1_\theta = \tilde{R}_\phi(\pi)\tilde{R}_{3\phi}(2\pi)\tilde{R}_\phi(\pi)\tilde{R}_x(\theta), \tag{52}$$

where $\tilde{R}_\phi(\cdot)$ denotes a rotation about the axis $[\cos\phi,\sin\phi,0]$, and the choice $\phi=\cos^{-1}(-\theta/4\pi)$ is made. This sequence gives average gate fidelity $\bar{F}(R_x(\pi/2),BB1_{90})\approx 1-21\pi^6\epsilon^6/16384$, which is much better than for the single pulse, even for relatively large values of ϵ, as shown in Fig. 22. The operation of the $BB1_{90}$ sequence is illustrated graphically in Fig. 23.

A few comments about this result are in order. This result is the best that has been presented in the literature to date (Wimperis, 1994; Cummins and Jones, 2000; Jones, 2003b). Currently, no pulse sequence that cancels out errors to higher order (for all possible initial states) has yet been published. It is also fairly general; $BB1_\theta$ approximates $R_x(\theta)$. Also, while composite pulses have been widely studied and employed in the art of NMR, this sequence is special in that it is *universal* (also termed *fully-compensating* or *class A*): the amount of error cancellation is independent of the starting state of the spin (Tycko, 1983; Tycko *et al.*, 1985). Other examples of such universal composite pulses are the sequence

$$\tilde{R}_{60}(180)\tilde{R}_{300}(180)\tilde{R}_{60}(180), \tag{53}$$

which performs a X^2 rotation with compensation for pulse length errors, and

$$\tilde{R}_y(385)\tilde{R}_y(-320)\tilde{R}_y(25), \tag{54}$$

which performs a Y rotation compensating for off-resonance errors and to some extent for pulse length errors as well.

Earlier, in the original work which introduced the concept of composite pulses into NMR (Levitt and Freeman, 1979; Levitt, 1986), only limited pulse sequences were known, which worked only for particular initial states; for example, there is the common $\tilde{R}_x(\pi/2)\tilde{R}_{-y}(\pi)\tilde{R}_x(\pi/2)$, used to approximate $R_x(\pi)$. Fig-

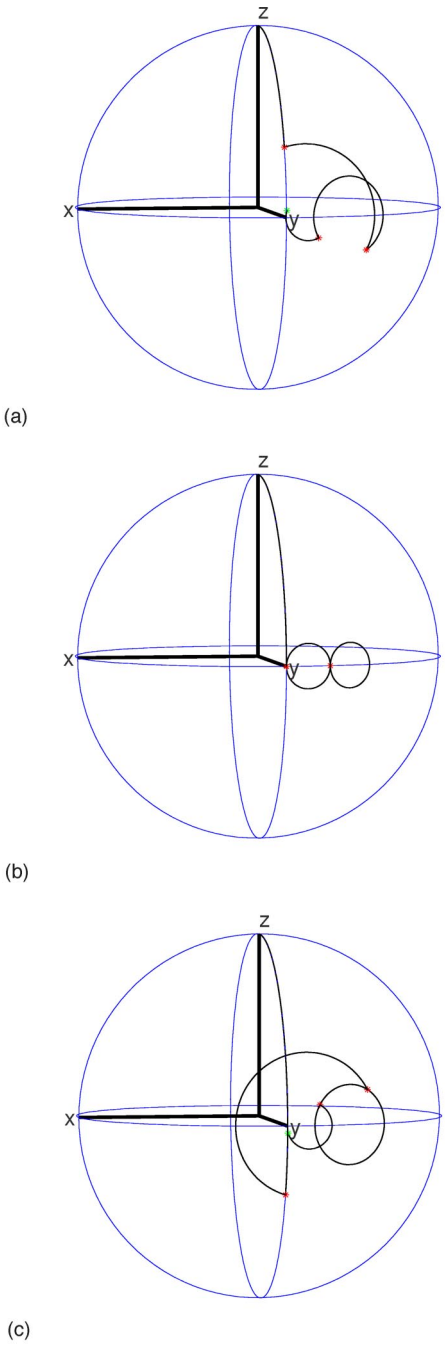

FIG. 23. (Color in online edition) Illustration of the trajectories of a spin as it transforms under the $BB1$ pulse sequence of Eq. (52), starting initially in the $|0\rangle$ state. Three trajectories are shown, in which the error is 50%: (a) under-rotation; (b) zero, and (c) 50% over-rotation. Plotted symbols denote the endpoints of each pulse in the sequences.

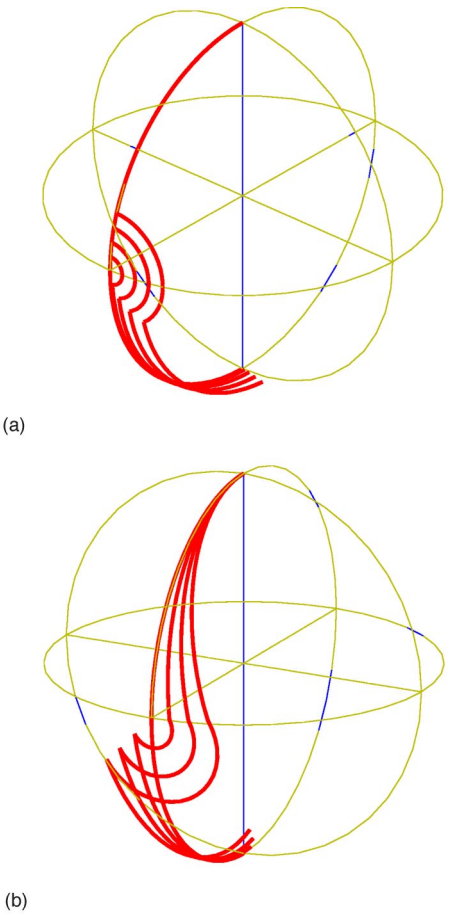

FIG. 24. (Color in online edition) Trajectory in the Bloch sphere described by a qubit initially in $|0\rangle$, when a composite 180° rotation is applied, consisting of three imperfect rotations, $\tilde{R}_x(\pi/2)\tilde{R}_{-y}(\pi)\tilde{R}_x(\pi/2)$: (a) The tip angles are set $0,5,\ldots,20\%$ too short; (b) The pulse is applied off-resonance, with $(\omega_0-\omega_{rf})/\omega_1=0,0.05,\ldots,0.20$. In both cases, the effect of the errors in the individual pulses is largely removed by the composite pulse.

ure 24 illustrates how this simple sequence removes the effect of errors in either the rotation angle or the rotation axis.

Systematic errors in the coupling strengths can also be tackled using composite rotations, in order to obtain accurate two-qubit gates. This was shown explicitly for the case of Ising couplings (Jones, 2003a).

Similar compensation for slowly fluctuating errors can be achieved during a train of pulses, separated by time intervals of free evolution. The simplest instance of such a pulse train uses only 180° pulses. Off-resonance effects in such pulses can be largely reversed by properly choosing the phases of the pulses. For instance, and at first sight surprisingly, the errors from off-resonant pulses $X^2\bar{X}^2$ roughly add up, while they largely compensate each other in X^2X^2. This cancellation can be appreciated via a simple Bloch sphere picture (Fig. 25). The remaining errors are further reduced for a properly chosen train of four pulses, $X^2X^2\bar{X}^2\bar{X}^2$, which performs markedly better than $X^2\bar{X}^2X^2\bar{X}^2$ (Levitt *et al.*, 1982). Further reduction of the effect of off-resonance errors can be obtained by using even longer trains of 180° pulses (Levitt *et al.*, 1982).

Evidently, quantum computing sequences are not as transparent as just a train of 180° pulses. Even throughout a quantum computing sequence, the effect of RF inhomogeneities can be removed to a large extent (Vandersypen, Steffen, Sherwood, Yannoni, Breyta, and

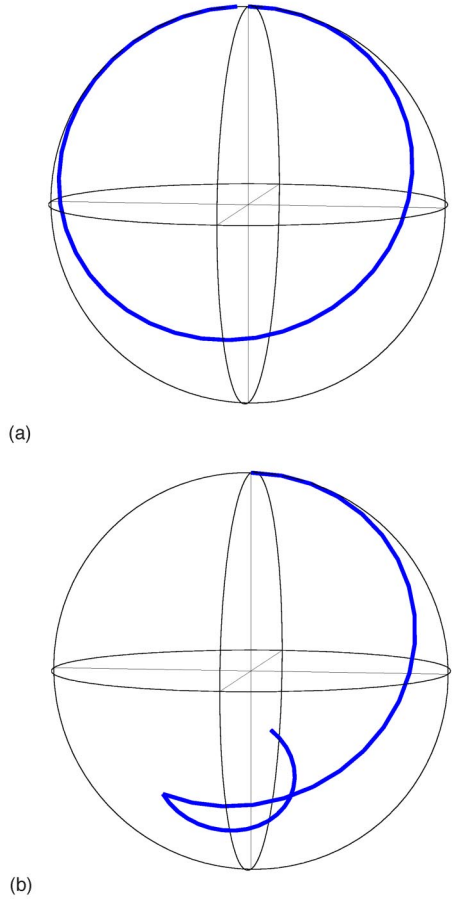

FIG. 25. (Color in online edition) Trajectory in the Bloch sphere of a qubit initially in $|0\rangle$, subject to two consecutive 180° pulses, applied off-resonance with $(\omega_0-\omega_{rf}/\omega_1)=0.5$: (a) If the two pulses are applied with the same phase (X^2X^2), the qubit is taken simply along a circular trajectory through $|0\rangle$, and reaches a point near $|0\rangle$; to be precise, the 50% resonance offset makes the rotation angle $\sqrt{(2^2+1^1)/2^2}=\sqrt{5/4}$ larger than 360°. (b) In contrast, if the two pulses are applied with opposite phases ($X^2\bar{X}^2$), the qubit is left far from $|0\rangle$.

Chuang, 2000), as illustrated in Fig. 26. After completion of a routine involving the equivalent of about 1350 90° pulses, the measured amplitude was about 15% of the full amplitude. Without removal of the effect of RF inhomogeneity, the signal would have been buried in the noise very rapidly.

This level of error cancellation was achieved partly due to a judicious choice of the phases of the refocusing pulses. Nevertheless, a more detailed description and understanding of the error operators is needed in order to fully exploit the potential for error cancellation in arbitrary pulse sequences.

2. Numerical optimization

The composite pulses we discussed in the previous subsection are designed to compensate for certain types of errors (mostly over- or under-rotations and frequency offsets), and work even when the exact Larmor frequencies, spin-spin coupling strengths, and the magnitude of

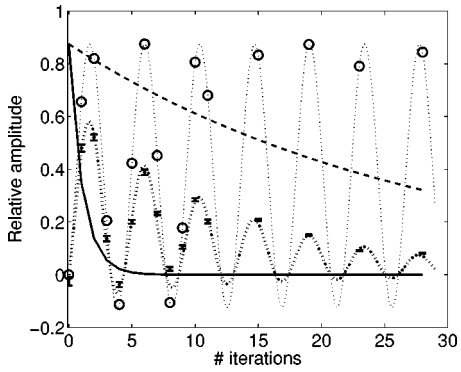

FIG. 26. Experimental (error bars) and ideal (circles) amplitude of $|\downarrow\uparrow\downarrow\rangle$, as a function of the number of iterations of a quantum search algorithm (Nielsen and Chuang, 2000), for three qubits, executed on $^{13}CHFBr_2$. Each iteration contains the equivalent of almost 50 90° pulses. The dotted lines serve to guide the eye. Dashed line: the signal decay for ^{13}C due to decoherence, which represents a lower bound on the decay rate. Solid line: the signal strength retained after applying a continuous RF pulse of the same cumulative duration per search iteration as the pulses in the actual experiment (averaged over three spins, measured up to four iterations and then extrapolated). Similar observations have been reported in Vandersypen *et al.* (1999). From Vandersypen, Steffen, Sherwood, Yannoni, Breyta, and Chuang, 2000.

the errors are unknown. This is the usual case in NMR spectroscopy. However, in quantum computing experiments, detailed knowledge of the system Hamiltonian is usually available and can be used to tailor the composite pulses to the system specifics, taking the degree of quantum control one step further.

Following the notation of Fortunato *et al.* (2002), we consider the concatenation of a number of rectangular pulses, each described by four parameters: the pulse duration τ^m, a constant amplitude ω_1^m, the transmitter frequency ω_{rf}^m, and the initial phase ϕ^m, where m indexes the pulse. These parameters may be strongly modulated from one pulse to the next.[3] Via a numerical optimization procedure, the values of τ^m, ω_1^m, ω_{rf}^m, and ϕ^m are chosen such that the resulting net unitary evolution U_{net} is as close as possible to the ideal unitary transformation U_{ideal}, according to some fidelity measure (Sec. V.C).

In practice, the number of time slices in the composite pulse is increased starting from one, until a satisfactory solution is found. While the fidelity function may have many local maxima, and finding the global maximum may therefore take a long time, suitable algorithms such as the Nelder-Mead Simplex algorithm (Nelder and Mead, 1965) often succeed in finding a reasonably good solution. Furthermore, the optimization routine can incorporate penalties on high powers, large frequencies,

[3]Jumps in the transmitter frequency can be conveniently realized with phase-ramping techniques. As discussed in Sec. IV.A.2, this is done by phase shifting the raw RF excitation in fixed increments per time so that a different RF frequency is obtained.

and negative or very long time periods, in order to prevent the algorithm from returning infeasible solutions.

Computation of U_{net} uses the fact that the Hamiltonian during a fixed-amplitude RF pulse can be made time independent by transforming into a reference frame rotating at the transmitter frequency, as we have seen in Sec. II.B.2. We shall call $\mathcal{H}_{rot}^m$ the effective Hamiltonian in the frame rotating at ω_{rf}^m during segment m. Given that ω_{rf}^m may be different for every segment of the pulse, it is most convenient to transform back to a common reference frame at the end of every time slice. This can be the frame of the raw RF frequency, or the laboratory frame of the n-spin system. In the lab frame, the time evolution during segment m is described by

$$U^m = e^{-i\omega_{rf}^m \Sigma_{k=1}^n I_z^k \tau^m} e^{-i\mathcal{H}_{rot}^m \tau^m}. \tag{55}$$

Since all U^m are expressed in the same reference frame, we can simply multiply them together to get $U_{net}=\Pi_m U^m$ and compare the result directly with U_{ideal}, expressed in the laboratory frame as well.

Two representative examples of composite pulses designed for spin-selective rotations in homonuclear spin systems are given in Fig. 27. The gate fidelity (Sec. V.C.2) obtained with these two pulses is displayed in Fig. 28. Naturally, the fidelity is close to unity only near the resonance frequencies for which the gate was designed to work.

Composite pulses can thus effectively generate accurate single- and multiple-qubit Hamiltonians, using detailed knowledge of the system Hamiltonian and only limited knowledge about the errors. Often, however, full knowledge of the system parameters is not available, and thus methods beyond composite pulses must be employed.

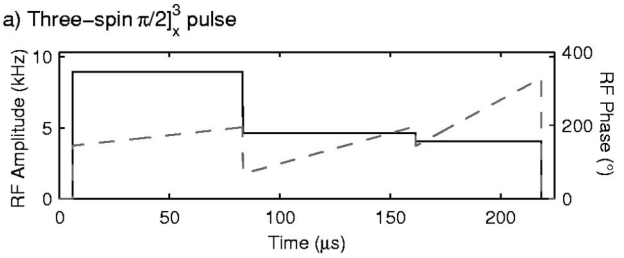

Pulse Parameters:	1	2	3
Time (μs)	77.3	78.2	56.9
Power (kHz)	8.96	4.62	4.06
Initial Phase (°)	145	68	144
Frequency (kHz)	1.81	4.59	9.03

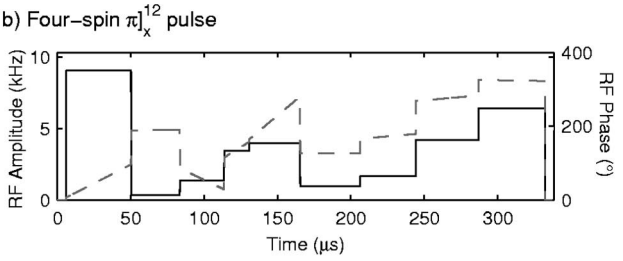

Pulse Parameters:	1	2	3	4	5	6	7	8	9
Time (μs)	44.3	33.3	30.1	16.9	34.7	40.9	37.9	42.7	45.4
Power (kHz)	9.06	0.31	1.36	3.45	3.99	0.94	1.68	4.21	6.42
Initial Phase (°)	6	9	85	115	164	126	166	-92	-34
Frequency (kHz)	5.66	0.07	-5.29	6.98	9.30	-0.04	0.96	0.99	-0.19

FIG. 27. The ideal RF waveform for two examples of strongly modulated pulses: solid lines, amplitude of the waveform; dashed lines, phase. Details of the pulse parameters, as in Eq. (55) are listed below each wave form. The 6-μs time interval with zero RF power before and after the composite pulses is needed due to experimental implementation issues. The composite pulse in (a) performs a 90° rotation on one of the ^{13}C nuclei of ^{13}C-labeled Alanine and the pulse in (b) performs a simultaneous 180° rotation on two ^{13}C nuclei of ^{13}C-labeled Crotonic acid. Courtesy of D. G. Cory. Reproduced from Fortunato *et al.*, 2002.

C. Average-Hamiltonian theory

The average-Hamiltonian formalism offers a versatile framework for understanding how to effectively create or remove arbitrary terms in the Hamiltonian by periodic perturbations, without requiring full knowledge of the system dynamics. The refocusing sequences presented in Sec. III.A.4 and more general multiple-pulse sequences designed to neutralize the effect of dipole-dipole couplings can be explained within this framework. Reduction of full dipole-dipole coupling given by Eq. (5) to the simplified forms of Eqs. (6) and (7) can also be understood with average-Hamiltonian theory.

Following Ernst *et al.* (1987), we first introduce the Magnus expansion and then see how we can modify a time-independent Hamiltonian via a time-dependent perturbation. We use two concrete examples to illustrate the concepts.

1. The Magnus expansion

The essence of average-Hamiltonian techniques is that the evolution $U(t)$ under a time-dependent Hamiltonian $\mathcal{H}(t)$ can be described by an effective evolution under a time-*independent* average Hamiltonian $\bar{\mathcal{H}}$, under two conditions (Haeberlen and Waugh, 1968; Ernst *et al.*, 1987): (i) $\mathcal{H}(t)$ is periodic and (ii) the observation is stroboscopic and synchronized with the period t_c of $\mathcal{H}(t)$.

We can then calculate $\bar{\mathcal{H}}$ exactly from

$$U(t_c) = \exp(-i\bar{\mathcal{H}}t_c), \tag{56}$$

by diagonalizing $U(t_c)$ and taking the logarithm of the resulting eigenvalues (Nielsen and Chuang, 2000).

In practice, it is often more convenient to compute $\bar{\mathcal{H}}$ approximately. Let us assume that $\mathcal{H}(t)$ is piecewise constant [the analysis can be easily generalized to the case of continuously changing Hamiltonians (Ernst *et al.*, 1987)]: $\mathcal{H}(t)=\mathcal{H}_k$ for $\Sigma_0^{k-1}\tau_i < t < \Sigma_0^k \tau_i$, and $t_c=\Sigma_0^n \tau_k$, so

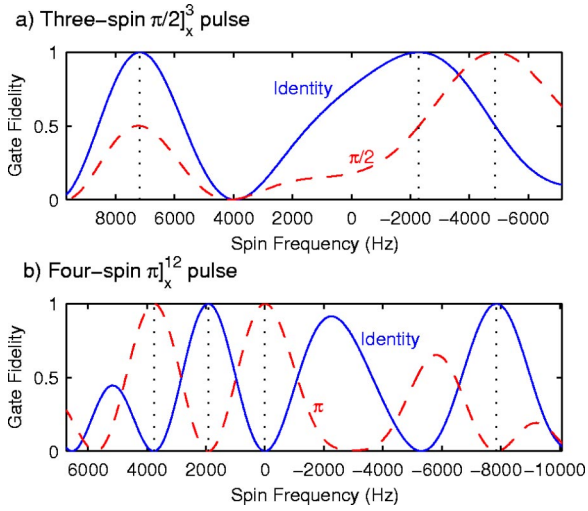

FIG. 28. (Color in online edition) Gate fidelity of the two example pulses of Fig. 27 as the resonance frequency of a test spin is varied. The solid (dashed) line is calculated with identity (desired transformation) as the intended transformation. The vertical dotted lines denote the actual chemical shifts for each spin. Courtesy of D. G. Cory. Reproduced from Fortunato *et al.*, 2002.

$$U(t_c) = \exp(-i\mathcal{H}_n\tau_n)\cdots\exp(-i\mathcal{H}_0\tau_0). \tag{57}$$

Repeated application of the Baker-Campbell-Hausdorff relation,

$$e^B e^A = \exp\left\{A + B + \frac{1}{2}[B,A] + \frac{1}{12}([B,[B,A]] + [[B,A],A]) + \cdots\right\} \tag{58}$$

gives

$$\bar{\mathcal{H}} = \bar{\mathcal{H}}^{(0)} + \bar{\mathcal{H}}^{(1)} + \bar{\mathcal{H}}^{(2)} + \cdots, \tag{59}$$

where

$$\bar{\mathcal{H}}^{(0)} = \frac{1}{t_c}\{\mathcal{H}_0\tau_0 + \cdots + \mathcal{H}_n\tau_n\}, \tag{60}$$

$$\bar{\mathcal{H}}^{(1)} = \frac{-i}{2t_c}\{[\mathcal{H}_1\tau_1,\mathcal{H}_0\tau_0] + [\mathcal{H}_2\tau_2,\mathcal{H}_0\tau_0] + [\mathcal{H}_2\tau_2,\mathcal{H}_1\tau_1] + \cdots\}, \tag{61}$$

and so forth. This expansion, called *the Magnus expansion* (Magnus, 1954), forms the basis of average-Hamiltonian theory.

2. Multiple-pulse decoupling

Let us consider a pulse sequence of n infinitesimally short pulses U_k separated by time intervals τ_k of free evolution under the system Hamiltonian $\mathcal{H}_0$, and such that $U_n\cdots U_2U_1 = I$ (for pulses of finite length, the duration of the pulses must also be included in the average). The pulses correspond to basis transformations, and we can thus describe the system evolution via a sequence of time intervals τ_k of free evolution under the Hamiltonian $\tilde{\mathcal{H}}_0(k)$, with

$$\tilde{\mathcal{H}}_{0(0)} = \mathcal{H}_0, \tag{62}$$

$$\tilde{\mathcal{H}}_{0(1)} = U_1^{-1}\mathcal{H}_0U_1, \tag{63}$$

$$\tilde{\mathcal{H}}_{0(2)} = U_1^{-1}U_2^{-1}\mathcal{H}_0U_2U_1, \tag{64}$$

and so forth. Note that the order in which the U_k are applied to $\mathcal{H}_0$ is reversed and that the U_k themselves are reversed as well. If we let $t_c = \Sigma_0^n\tau_k$, then the overall transformation $U(t_c)$ is given by

$$U(t_c) = \exp(-i\tilde{\mathcal{H}}_{0(n)}\tau_n)\cdots\exp(-i\tilde{\mathcal{H}}_{0(0)}\tau_0). \tag{65}$$

We can now use the Magnus expansion of Eq. (59) and Eqs. (60) and (61), where we replace $\mathcal{H}_k$ by $\tilde{\mathcal{H}}_{0(k)}$, to obtain the average Hamiltonian $\bar{\mathcal{H}}_0$ which describes the net time evolution during t_c. The zeroth-order average Hamiltonian is given by

$$\bar{\mathcal{H}}_{(0)}^0 = \frac{1}{t_c}\sum_{k=0}^{n}\tau_kU_1^{-1}\cdots U_k^{-1}\mathcal{H}_0U_k\cdots U_1. \tag{66}$$

The crux of average-Hamiltonian theory is that, by properly choosing the pulse U_k, we can ensure that $\bar{\mathcal{H}}_0^{(0)}$ contains only the desired terms.

Sophisticated pulse sequences (Mehring, 1983) can also remove undesired contributions from the higher-order terms in the expansion, although this is generally harder since $\bar{\mathcal{H}}_0^{(1)},\bar{\mathcal{H}}_0^{(2),},\ldots$ contain cross-terms between the various $\tilde{\mathcal{H}}_{0(k)}$. The commutators involved in the higher-order terms do become smaller for shorter cycle times, though, so fast cycles result in better averaging.

We also point out that pulse sequences which satisfy

$$\tilde{\mathcal{H}}_{0(k)} = \tilde{\mathcal{H}}_{0(n-k)} \tag{67}$$

or equivalently

$$U_{k+1} = U_{n-k}^{\dagger} \tag{68}$$

contain no contributions of odd orders to $\bar{\mathcal{H}}_0$,

$$\bar{\mathcal{H}}_{(k)}^0 = 0 \quad \text{for} \quad k = 1,3,5,\ldots, \tag{69}$$

and thus perform significantly better than other sequences.

Let us now illustrate the operation of multiple-pulse decoupling via two examples. First, the original multiple-pulse sequence for removal of dipole-dipole interactions is the WAHUHA-4 sequence (Waugh *et al.*, 1968),

$$\tau\bar{X}\tau Y2\tau\bar{Y}\tau X\tau, \tag{70}$$

where the pulses are applied to all qubits involved, τ stands for free evolution under the system Hamiltonian for a duration τ, and the unitaries are ordered from right

to left, as usual. The pulses rotate the Zeeman terms in the Hamiltonian from $-\hat{z}$ to $-\hat{y}$, $-\hat{x}$, $-\hat{y}$, and back to $-\hat{z}$ [see Eqs. (62)–(64)] for a duration τ, τ, 2τ, τ, and τ, respectively. The zeroth-order average Zeeman term is thus oriented along $-(\hat{x}+\hat{y}+\hat{z})$, with strength scaled down by a factor $1/\sqrt{3}$. The dipolar Hamiltonian of Eq. (6) goes through the forms $[3I_z^iI_z^j-I^i\cdot I^j]$, $[3I_y^iI_y^j-I^i\cdot I^j]$, and $[3I_x^iI_x^j-I^i\cdot I^j]$ for equal durations, and is thus zero on average.

By selectively not pulsing specific qubits, it is also possible to reintroduce some of the couplings as desired. In Fig. 12, we already saw explicitly how to do this for $I_z^iI_z^j$ couplings.

A second example is an extension of the conventional spin-echo sequence (Sec. V.A.4) to three-component spin echoes (Augustine and Hahn, 1997). In conventional echo sequences, 180° pulses about $\hat{x}$ or $\hat{y}$ remove the effect of a Hamiltonian of the form $c_z\sigma_z$. Now we ask ourselves what sequence of pulses would freeze the evolution under a Hamiltonian of the form

$$\mathcal{H}=c_x\sigma_x+c_y\sigma_y+c_z\sigma_z, \tag{71}$$

where c_x,c_y,c_z are arbitrary coefficients. We can use Eq. (66) to verify that the sequence

$$X^2\tau\bar{X}^2Y^2\tau\bar{Y}^2Z^2\tau\bar{Z}^2\tau, \tag{72}$$

or equivalently, after simplification,

$$X^2\tau Z^2\tau X^2\tau\bar{Z}^2\tau, \tag{73}$$

gives a zeroth-order average Hamiltonian $\bar{\mathcal{H}}^{(0)}=0$, and thus in effect corresponds to a three-component echo sequence. Another way to show this is to note that

$$X^2\mathcal{H}\bar{X}^2=+c_x\sigma_x-c_y\sigma_y-c_z\sigma_z, \tag{74}$$

$$Y^2\mathcal{H}\bar{Y}^2=-c_x\sigma_x+c_y\sigma_y-c_z\sigma_z, \tag{75}$$

$$Z^2\mathcal{H}\bar{Z}^2=-c_x\sigma_x-c_y\sigma_y+c_z\sigma_z. \tag{76}$$

Clearly, $\mathcal{H}+X^2\mathcal{H}\bar{X}^2+Y^2\mathcal{H}\bar{Y}^2+Z^2\mathcal{H}\bar{Z}^2=0$, and so the sequence of Eq. (72) gives, to zeroth order, no net evolution. Again, if τ is sufficiently short, the higher-order contributions will be negligible.

3. Reversing errors due to decoherence

Can we apply multiple-pulse sequences to reverse the effect of interactions of a qubit with degrees of freedom in the environment? It is not clear *a priori* that this is possible: standard average-Hamiltonian theory assumes that we can manipulate both interacting particles involved, for instance, via RF pulses. However, we have no control of degrees of freedom in the environment.

Remarkably, it is actually possible to remove the effect of unwanted interactions with degrees of freedom in the environment, even when applying operations to the system only (Viola and Lloyd, 1998; Viola *et al.*, 1998; Duan and Guo, 1999; Vitali and Tombesi, 1999), provided the control operations are applied faster than the fluctuations. Knowledge about the nature of the interactions can be applied to simplify the sequence of decoupling operations, and such knowledge can even be experimentally extracted (Byrd and Lidar, 2003), in part using a procedure known as *process tomography*, described in Sec. V.B.2.

If the fluctuations are faster than the accessible control operations, errors can be corrected using quantum error correction (Shor, 1995; Steane, 1996; Nielsen and Chuang, 2000), or they can be prevented by encoding the qubits in a subspace that is not affected by decoherence (Zanardi and Rasetti, 1997; Lidar *et al.*, 1998). This is discussed further in Sec. V.D.

V. EVALUATION OF QUANTUM CONTROL

The pulse control methods presented in the last two sections can have impressive performance, but this is very much contingent on having an accurate model of the system under control. A variety of techniques have been used in NMR to characterize the system dynamics and to evaluate the performance of control sequences. In this section, we review some of these techniques, beginning with a set of standard experiments to determine how quantum a qubit system is, then proceeding to tomographic methods for fully characterizing system dynamics, and concluding with fidelity metrics for control and their implications for scalability to large systems.

A. Standard experiments

In NMR spectroscopy as in atomic physics, a number of standard experiments serve to test the quantum-mechanical behavior of a given system and to determine the extent of its isolation from the environment (see Sec. II.C), in terms of its phase coherence time T_2 and its energy relaxation time T_1, as well as the decay time in the rotating frame $T_{1\rho}$.

1. Coherent oscillations driven by a resonant field

The dynamics of a single spin, driven resonantly by a coherent field, were presented in Secs. II.B.1 and II.B.2. From Eq. (27), we have that in the ideal case the RF field induces transitions from $|0\rangle$ to $|1\rangle$, where a qubit initially in $|0\rangle$ will be found in $|1\rangle$ after an RF pulse of duration t_{pw} with probability

$$\Pr[|1\rangle]=\sin^2(\gamma B_1t_{pw}/2)=\sin^2(\omega_1t_{pw}/2). \tag{77}$$

The probability initially increases over time, until it reaches a maximum $\Pr[|1\rangle]=1$ and then decreases again, by stimulated emission, a cycle which keeps repeating itself.

Such oscillations of a two-level quantum system driven by a resonant field are known as Rabi oscillations (Rabi, 1937), and the Rabi frequency $\omega_1/2\pi$ is proportional to the amplitude of the control field. Observation of Rabi oscillations is usually accepted as a signature of quantum coherent behavior.

In reality, the envelope of the Rabi oscillation signal is always damped, due to decoherence as well as instrumental imperfections; measurement of this decay time is useful, and known as a *nutation* experiment. In NMR, the Rabi decay time is often much shorter than the intrinsic phase randomization time constant T_2, due to the inhomogeneity across the macroscopic sample of the RF field driving the Rabi oscillation. In other systems, the Rabi decay time may be longer than T_2, because (i) a long pulse can be seen as a concatenation of many 180° pulses, which can have a refocusing effect (see Sec. III.A.4), and (ii) the qubit is near $\pm\hat{z}$, where phase randomization has no effect, for roughly half the time during Rabi oscillations.

Coherent oscillations driven by a resonant field have been observed in NMR and in many atomic systems for a long time. Recently, however, observations of such coherent dynamics have been made in other qubit systems, including systems made from Josephson junctions (Nakamura *et al.*, 1999), molecular vibrational states (Tesch and de Vivie-Riedle, 2002; Vala *et al.*, 2002), and excitons in semiconductor quantum dots (Stievater *et al.*, 2001).

2. Coherent oscillations initiated by a kick

A quantum system starting off in a state that is not an eigenstate of the (static) system Hamiltonian will precess about the quantization axis of the system Hamiltonian, a motion known as Larmor precession (see Sec. II.A.1). Such a situation could be realized by abruptly changing the system Hamiltonian, e.g., by suddingly applying a strong static field along $\hat{x}$ instead of along $\hat{z}$. Alternatively, and more realistically in NMR, Larmor precession can be initiated by suddenly kicking the qubit out of the Hamiltonian eigenbasis. For a nuclear spin with Hamiltonian $-\hbar\omega_0 I_z$, as in Eq. (1), this is done by applying a 90° RF pulse, causing a transition for instance from $|0\rangle$ to $(|0\rangle+|1\rangle)/\sqrt{2}$, which initiates the time evolution

$$|\psi(t)\rangle = \frac{e^{i\omega_0 t/2}|0\rangle + e^{-i\omega_0 t/2}|1\rangle}{\sqrt{2}}, \tag{78}$$

as illustrated in Fig. 2. Like Rabi oscillations, the observation of Larmor precession is also a signature of quantum coherent behavior.

The Larmor precession is also damped, but in contrast to the Rabi decay time, the Larmor decay time, termed T_2^*, is never longer than T_2. Usually, $T_2^* < T_2$; in particular, for NMR,

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}, \tag{79}$$

where T_2' is the dephasing time constant due to static magnetic-field inhomogeneities or other instrumental imperfections.

Larmor oscillations initiated by a kick have been observed recently in a variety of systems, including those driven resonantly mentioned earlier, and in addition a system of charges in coupled quantum dots (Hayashi *et al.*, 2003). The oscillations can be observed directly if the measurement basis lies in the $\hat{x}$-$\hat{y}$ plane, as is the case in NMR. If the measurement takes place along $\pm\hat{z}$, we must first change the basis via a second 90° pulse.

3. Ramsey interferometry

The double-pulse experiment

$$X\ \tau\ X, \tag{80}$$

where time goes from right to left (as always, for unitary transformations given in this article) and τ denotes a free evolution period, under the evolution operation $e^{-i\mathcal{H}_{sys}\tau/\hbar}$ is known as a Ramsey interference experiment (Ramsey, 1950). Originally, this "method of separated oscillatory fields" was applied to electronic states of molecular beams traversing through two microwave excitation zones. In NMR, two pulses are involved, separated by a delay time τ. Ramsey interference is most naturally described in the rotating frame of the RF field driving the transition. If the qubit starts off along $\hat{z}$, the first X pulse rotates it to $-\hat{y}$. Then the qubit precesses about $\hat{z}$ for a time τ. Finally the second X pulse rotates the $\pm\hat{y}$ component of the qubit state to $\mp\hat{z}$. Components along $\pm\hat{x}$ at the end of the interval τ remain along $\pm\hat{x}$ after the second X pulse.

If only a single qubit is considered and the RF field is exactly on-resonance with the qubit precession, the qubit stays in place in the rotating frame during the time interval in between the two pulses, and the final state does not vary with τ. However, if the RF and the qubit are detuned in frequency by $\Delta\omega$, both the $\hat{x}$ and the $\hat{z}$ components of the final state display a beating pattern as a function of $\Delta\omega\tau$, the so-called Ramsey fringes. The decay time of the envelope of the Ramsey fringes is T_2^*, the same as that for Larmor precession.

For coupled qubits, the beating pattern contains information on the coupling strengths. This fact forms the basis for two-dimensional correlation spectroscopy (Jeener, 1971; Ernst *et al.*, 1987), a widely used range of two-pulse techniques for molecular structure determination.

4. Measurement of T_2

The intrinsic T_2 time can be extracted in an experiment that is based on Larmor or Ramsey experiments. Certain imperfections which cause the Ramsey or Larmor decay time T_2^* to be smaller than T_2 can be removed by applying refocusing sequences.

The simplest instance of such a refocusing sequence consists of a single 180° pulse applied halfway through the time interval of free evolution initiated by a 90° pulse. The entire sequence is thus

$$\frac{\tau}{2} X^2 \frac{\tau}{2} X. \tag{81}$$

A second X pulse should be added at the end if the measurement takes place in the $\pm\hat{z}$ basis. In multispin

systems, the pulses must be applied selectively to one spin, in order to measure the T_2 of that spin.

The X^2 refocusing pulse not only removes simple scalar spin-spin couplings, as described in Sec. III.A.4, but also undoes the effect of spatial variations of the static magnetic field along $\hat{z}$. Such field inhomogeneities cause spins in different regions of the sample to become progressively out of phase with each other during the first time interval $\tau/2$. As a result, their magnetic moments cancel each other out and the NMR signal vanishes. Provided the magnetic-field variations are constant throughout the experiment, all the spins come back into phase again (now along $+\hat{y}$) by the end of the second time interval $\tau/2$, because of the 180° refocusing pulse. As a result, the signal recovers, producing the well-known spin-echo. A generalization of this technique known as *three-component refocusing* (Sec. IV.C.2) undoes effects from any static spin Hamiltonian terms.

The echo signal decays as a function of τ, and the decay time constant is a measure of T_2. However, terms in the Hamiltonian fluctuating on a time scale shorter than τ are not removed by a single refocusing pulse. Their effect can still be removed if the fluctuations are slow compared to τ/n and a train of n refocusing pulses is applied, each preceded and followed by a time interval $\tau/2n$ of free evolution. This so-called Carr-Purcell sequence (Carr and Purcell, 1954),

$$\frac{\tau}{2n}X^2\frac{\tau}{n}\cdots X^2\frac{\tau}{n}X^2\frac{\tau}{2n}X, \tag{82}$$

produces a first echo along $+\hat{y}$ after τ/n, a second echo along $-\hat{y}$ after $2\tau/n$, a third along $+\hat{y}$ after $3\tau/n$, and so forth. The magnitude of the echo signal decays exponentially throughout this sequence, and the echo signal left at the end of this sequence decreases exponentially as a function of the total time τ. To the extent that slow fluctuations in the Hamiltonian have been refocused, the decay time constant is the intrinsic T_2.

As we have seen in Sec. III of this review, small but fixed errors in the pulse amplitude or duration may accumulate throughout a multiple-pulse sequence such as the Carr-Purcell sequence. However, if the phase of the refocusing pulses is shifted by 90° with respect to the initial 90° pulse, pulse length errors are compensated on even-numbered echoes and are thus not cumulative. In this sequence,

$$\frac{\tau}{2n}Y^2\frac{\tau}{n}\cdots Y^2\frac{\tau}{n}Y^2\frac{\tau}{2n}X, \tag{83}$$

known as the Carr-Purcell-Meiboom-Gill or CPMG sequence (Meiboom and Gill, 1958), the echoes appear all along $-\hat{y}$. Again, the decay time constant of the echo signal gives T_2.

Since T_2 indicates how long a qubit can remain phase coherent, it is usually called the coherence time, although the terms *phase randomization time, dephasing time*, and *transverse relaxation time* are also used. In NMR, T_2 is also known as the *spin-spin relaxation time*. In any case, T_2 is an important number for evaluating the potential of quantum computers, as the ratio of T_2 over the typical duration of a quantum logic gate expresses the number of operations that can be completed coherently.

5. Measurement of T_1

Energy exchange with the environment makes a qubit that is out of equilibrium gradually return to thermal equilibrium. In thermal equilibrium, the qubit is in a statistical mixture of $|0\rangle$ and $|1\rangle$, with probabilities set by the temperature and the energy difference between $|0\rangle$ and $|1\rangle$. The time constant of this equilibration process, T_1, is often called the *energy relaxation time*, the *longitudinal relaxation time*, or simply the relaxation time. In NMR, T_1 is often termed the *spin-lattice relaxation time*.

Two standard experiments for measuring T_1 are inversion recovery and saturation recovery. The sequence for the inversion recovery experiment is

$$X\tau X^2. \tag{84}$$

The 180° pulse inverts the $|0\rangle$ and $|1\rangle$ probabilities, then during time τ, relaxation takes place, and finally a 90° read-out pulse is applied if necessary (i.e., when the measurement basis is in the $\hat{x}$-$\hat{y}$ plane). In saturation recovery, a strong RF field is applied for a long enough time that it saturates the qubit transition and equalizes the $|0\rangle$ and $|1\rangle$ probabilities. As in inversion recovery, the original $|0\rangle$ and $|1\rangle$ populations are altered, and we can monitor the populations' return to their equilibrium value as a function of τ. The time constant of this equilibration process is T_1.

Note that both the inversion recovery and saturation recovery experiments bring the qubit out of equilibrium, but to a state that has no coherence. As a result, phase randomization does not affect the equilibration process—we measure purely the effect of energy exchange with the bath. In contrast, Ramsey and spin-echo experiments for measuring T_2^* and T_2 pick up contributions from phase randomization both without and with energy exchange with the bath. If energy exchange dominates phase randomization, the measured T_2 is $2T_1$.

The relevance of T_1 is twofold. First, it sets an upper bound for T_2 and second, it tell us how much time we have to perform a measurement in the $\{|0\rangle,|1\rangle\}$ basis. Phase randomization does not change the $|0\rangle$ and $|1\rangle$ probabilities, so T_2 is irrelevant during such a measurement. In many cases, $T_1 \gg T_2$, in which case we have more time to measure than to perform coherent operations.

6. Measurement of $T_{1\rho}$

A third decay time useful to characterize the degree of isolation between a qubit and the environment is $T_{1\rho}$. This time constant can be measured via a technique called *spin locking*, in which the spin is first rotated into the $\hat{x}$-$\hat{y}$ plane, say, by a Y pulse, and next continuous irradiation is applied, phase shifted by 90° with respect

to the pulse, so it is aligned with the spin state (along the $\hat{x}$ axis):

$$R_x(\text{continuous})Y. \tag{85}$$

The continuous irradiation along $\hat{x}$ locks the spin to the $\hat{x}$ axis, in the following sense. Whenever the spin starts to diverge from the $\hat{x}$ axis due to interactions with the environment, the RF field rotates it to the opposite side of the $\hat{x}$ axis within a time π/ω_1. Provided that the spin is still moving in the same direction after this time, it will thus return to the $\hat{x}$ axis (note that spin locking thus also inhibits evolution due to J couplings and moderate frequency offsets). Only if the spin moves in the opposite direction after π/ω_1 will it continue to depart from the $\hat{x}$ axis due to the rotation by the RF field. So the amplitude along $\hat{x}$ decays, and the decay time constant is termed $T_{1\rho}$, known in NMR as the spin-lattice relaxation time in the rotating frame.

We see thus that, whereas T_2 is governed by low-frequency fluctuations in the environment and T_1 depends on fluctuations at the Larmor frequency, the decay during spin locking arises from fluctuations at the Rabi frequency used during spin locking. The spin-locking experiment thus gives additional information on the spectral density of the interactions with the environment.

B. Measurement of quantum states and gates

The standard experiments presented in the previous section give only partial information on the system dynamics. Here we show that in fact the full relaxation superoperator can be determined systematically by a procedure known as *process tomography*, which builds upon *state tomography*, as described below.

1. Quantum state tomography

The density matrix ρ completely describes our knowledge of the state of a system. Measurement of the density matrix is therefore extremely helpful when testing or claiming the preparation of specific quantum states.

One-time measurement of each of n qubits, in a given basis of 2^n states $|m\rangle$, gives very little information on ρ. All that can be inferred from a measurement outcome m is that $\Pr[|m\rangle] \neq 0$.

Repeated measurements of n qubits, each time prepared in the same state and measured in the same basis, reveals the probability distribution for the measurement basis states,

$$\Pr[|m\rangle] = \langle m|\rho|m\rangle = \mathrm{Tr}(\rho|m\rangle\langle m|) = \mathrm{Tr}(\rho M), \tag{86}$$

where M is an observable or measurement operator. If we repeatedly measure each qubit in the $\{|0\rangle,|1\rangle\}$ basis, we thus obtain all the diagonal entries of ρ, ρ_{ii}.

Quantum state tomography (Chuang, Vandersypen, *et al.*, 1998; Chuang *et al.*, 1998a, 1998b) is a method that allows *all* the elements of the density matrix ρ to be determined. This method consists of repeating the measurement of the same state in various measurement bases, until all the elements of ρ can be determined, by solving a set of linear equations. In practice, it is often more convenient to first rotate the qubits via a unitary transformation and then perform the measurement in a fixed basis. This is equivalent to measuring in different bases, since

$$\mathrm{Tr}[\rho(UMU^\dagger)] = \mathrm{Tr}[(U^\dagger\rho U)M]. \tag{87}$$

Specifically, we can expand the density matrix of a single qubit ρ as

$$\begin{bmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{bmatrix} = \rho_{00}|0\rangle\langle 0| + \rho_{01}|0\rangle\langle 1| + \rho_{10}|1\rangle\langle 0| + \rho_{11}|1\rangle\langle 1|. \tag{88}$$

Measurements of a single qubit in the $\{|0\rangle,|1\rangle\}$ basis give us ρ_{00} and $\rho_{11}=1-\rho_{00}$. However, after changing basis via a 90° rotation about $\hat{x}$, transforming ρ to $X\rho X^\dagger$, we obtain access to $\mathrm{Im}(\rho_{10})=-\mathrm{Im}(\rho_{01})$. Similarly, measurement after transformation by Y reveals $\mathrm{Re}(\rho_{10})=\mathrm{Re}(\rho_{01})$. Thus, by measuring the qubit state first directly, then measuring the same state again after an X read-out pulse, and then again after a Y read-out pulse, we can reconstruct ρ completely.

Similarly, for n qubits, we can expand ρ as

$$\rho = \sum_{i=0}^{2^n-1}\sum_{j=0}^{2^n-1} \rho_{ij}|i\rangle\langle j| \tag{89}$$

and choose a set of basis changes that gives access to all 4^n-1 degrees of freedom in ρ.

However, it is much easier to find a suitable set of basis changes if we use the Pauli expansion of ρ instead of Eq. (89). The Pauli expansion for a single-qubit state is

$$\rho = c_0\sigma_0 + c_1\sigma_1 + c_2\sigma_2 + c_3\sigma_3, \tag{90}$$

where $c_0=1$ for normalization, and we use $\sigma_0=I/2$, $\sigma_1=\sigma_x/2$, $\sigma_2=\sigma_y/2$, $\sigma_3=\sigma_z/2$. Measurement in the computational basis, described by the observables $\sigma_0\pm\sigma_3$, gives us $\Pr(|0\rangle)=(c_0+c_3)/2$, and $\Pr(|1\rangle)=(c_0-c_3)/2$ so we can extract c_3. Since

$$X\rho X^\dagger = c_0\sigma_0 + c_1\sigma_1 - c_3\sigma_2 + c_2\sigma_3, \tag{91}$$

$$Y\rho Y^\dagger = c_0\sigma_0 + c_3\sigma_1 + c_2\sigma_2 - c_1\sigma_3, \tag{92}$$

we indeed obtain $(c_0\pm c_2)/2$ after applying X and $(c_0\mp c_1)/2$ after using Y.

For n qubits, Eq. (90) generalizes to

$$\rho = \sum_{i=0}^{3}\sum_{j=0}^{3}\cdots\sum_{k=0}^{3} c_{ij\cdots k}\,\sigma_i\otimes\sigma_j\otimes\cdots\otimes\sigma_k, \tag{93}$$

where $c_{00\cdots0}=1$. Measurement in the computational basis is described by observables of the form

$$(\sigma_0\pm\sigma_3)\otimes(\sigma_0\pm\sigma_3)\otimes\cdots\otimes(\sigma_0\pm\sigma_3) \tag{94}$$

and returns the probabilities

$$\sum_{i,j,\ldots,k\in\{0,3\}} \pm \frac{c_{ij\cdots k}}{2^n}. \tag{95}$$

For example, for two qubits, these are

$$\Pr(|00\rangle) = (c_{00} + c_{03} + c_{30} + c_{33})/4, \tag{96}$$

$$\Pr(|01\rangle) = (c_{00} - c_{03} + c_{30} - c_{33})/4, \tag{97}$$

$$\Pr(|10\rangle) = (c_{00} + c_{03} - c_{30} - c_{33})/4, \tag{98}$$

$$\Pr(|11\rangle) = (c_{00} - c_{03} - c_{30} + c_{33})/4. \tag{99}$$

After measurement of the four $\Pr[|m\rangle]$, we can solve for c_{03}, c_{30}, c_{33} from this overdetermined set of linear equations. Again, we can determine the other $c_{ij\cdots k}$ by transformation of the corresponding $\sigma_{ij\cdots k}$ to an observable, for instance,

$$X_1 Y_2(\sigma_0 + \sigma_2) \otimes (\sigma_0 + \sigma_1) X_1^\dagger Y_2^\dagger = (\sigma_0 + \sigma_3) \otimes (\sigma_0 - \sigma_3). \tag{100}$$

We end this discussion of state tomography with three additional comments:

First, in order to obtain all the basis-state probabilities as in Eqs. (96)–(99), we must each time read out all the qubits. If it is only possible to read out any one single qubit in each experiment, we obtain n bit-wise probabilities instead of 2^n basis-state probabilities, giving spin-spin correlations. The measurement operators are then of the form

$$2^{n-1}[\sigma_0 \otimes \sigma_0 \otimes (\sigma_0 \pm \sigma_3) \otimes \cdots \otimes \sigma_0] \tag{101}$$

and we measure probabilities

$$\frac{1}{2}(c_{0\cdots 000\cdots 0} \pm c_{0\cdots 030\cdots 0}). \tag{102}$$

It is now no longer possible to rotate arbitrary components of ρ into observable positions using just single-qubit rotations. Two-qubit gates are necessary to obtain all $c_{ij\cdots k}$.

Second, the measurement basis obviously need not be the computational basis. In NMR experiments, for instance, the single-qubit measurement operator can be written as $-i\sigma_1 - \sigma_2$. For two coupled spins, the measurement operators are

$$2(-i\sigma_1 - \sigma_2) \otimes (\sigma_0 \pm \sigma_3), \tag{103}$$

$$(\sigma_0 \pm \sigma_3) \otimes 2(-i\sigma_1 - \sigma_2), \tag{104}$$

and so forth. Since NMR experiments are normally performed on a large ensemble of molecules, the expectation value of the observables can be read out by acquiring a single spectrum. The four operators in Eq. (104) correspond to the four lines in the spectrum of a two-spin system (two doublets). Phase-sensitive detection permits us to separately record the real and imaginary component of each spectral line and distinguish σ_1 and σ_2 contributions.

Third, errors in the gates used for changing basis during state tomography lead to a measured density matrix which differs from the actual state of the system. If the errors are known and can be modeled accurately, they can be incorporated in the state tomography procedure and the actual state can nevertheless be determined accurately.

Quantum state tomography has been experimentally implemented in many atomic systems, notably the early work mapping out photon states (Smithey *et al.*, 1993) and vibrational cat states of trapped atoms (Meekhof *et al.*, 1996). Recently, it has become a common tool used to evaluate NMR states (Chuang, Vandersypen, *et al.*, 1998; Chuang *et al.*, 1998b), states of optical photon qubits (Thew *et al.*, 2002), and even vibrational states of molecules (Skovsen *et al.*, 2003).

2. Quantum process tomography

Now that we know how to determine experimentally the state of a quantum system, it is only a short step to the characterization of a quantum process, such as a quantum logic gate, communication channel, storage device, and so forth. In general, let us consider a quantum-mechanical black box whose input may be an arbitrary quantum state, and whose output is the result of the internal dynamics of the black box, as well as interactions with the outside world. Then can we ascertain the transfer function of this black box?

The answer is yes (Chuang and Nielsen, 1997; Poyatos *et al.*, 1997; D'Ariano and Lo Presti, 2001; Boulant *et al.*, 2003). The outline of the procedure is to first determine the output state of the black box for a set of input states which form a basis for the system Hilbert space, and then to use the fact that quantum mechanics is linear to compute the entire transfer function from this finite set of input-output pairs.

An arbitrary quantum state transformation is a linear map $\mathcal{E}$,

$$\rho \mapsto \frac{\mathcal{E}(\rho)}{\mathrm{Tr}[\mathcal{E}(\rho)]}, \tag{105}$$

where we can express $\mathcal{E}(\rho)$ in the operator-sum representation or Kraus representation (Kraus, 1983; Nielsen and Chuang, 2000)

$$\mathcal{E}(\rho) = \sum_k A_k \rho A_k^\dagger. \tag{106}$$

This is an alternative to the superoperator formalism widely used in NMR (Ernst *et al.*, 1987). The A_k are operators acting on the system alone, yet $\mathcal{E}$ completely describes the possible state changes of the system, including unitary operations, generalized measurements, and decoherence (for trace-preserving processes, $\Sigma_i A_k^\dagger A_k = 1$). The expansion of Eq. (106) is in general not unique. In fact, we can always describe $\mathcal{E}$ using a fixed set of operators $\tilde{A}_k$ which form a basis for the set of operators on the state space, so that (Chuang and Nielsen, 1997)

$$\mathcal{E}(\rho) = \sum_{p,q} \chi_{pq} \tilde{A}_p \rho \tilde{A}_q^\dagger, \tag{107}$$

where χ_{pq} is a positive Hermitian matrix. Since the $\tilde{A}_k$ are fixed, $\mathcal{E}$ is completely described by χ. In general, χ will contain $16^n - 4^n$ independent real parameters, where n is the number of qubits.

In order to determine χ experimentally, we choose a basis of 4^n linearly independent density matrices ρ_j which span the system Hilbert space, and determine $\mathcal{E}(\rho_j)$ for each j. We can then write down a set of linear equations of the form of Eq. (107), where we plug in the measurement outcomes $\mathcal{E}(\rho_j)$ and solve for the χ_{mn}.

The most convenient choice for the ρ_j depends on the implementation of the qubits and on the observables, as was the case for state tomography. Clearly, the effort needed to perform quantum process tomography increases even more rapidly with the number of qubits than quantum state tomography. This procedure has been used experimentally only for one- and two-qubit NMR (Childs *et al.*, 2001; Boulant *et al.*, 2003) and optical-photon systems (Altepeter *et al.*, 2003; Mitchell *et al.*, 2003).

The operation elements A_k in the operator sum representation of Eq. (106) can describe arbitrary quantum operations, but among these a select subset are useful to identify. For example, when A_k is a unitary matrix, this corresponds to perfect, closed-system Hamiltonian evolution. Phase damping (T_2) is described by

$$A_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{\gamma} \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix}, \tag{108}$$

and amplitude damping (T_1) by

$$A_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{\gamma} \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0 & \sqrt{1-\gamma} \\ 0 & 0 \end{bmatrix}, \tag{109}$$

where $\gamma \sim e^{-t/\tau}$ parametrizes the strength of the damping, for time t, in terms of a time constant τ. These and other relaxation parameters (Nielsen and Chuang, 2000) can be obtained by process tomography.

Such results can, in turn, be useful for approximate numerical simulation of relaxation and decoherence processes in spin systems. Phase damping and energy relaxation can be simulated in alternation with unitary evolution under the system and control Hamiltonian, taking sufficiently short time slices to obtain a good approximation of the true dynamics. This permits an n-spin system to be simulated using $n4^n$ steps, compared to 16^n for fully general quantum operations. Experimental results have shown this method to be predictive of the system dynamics throughout sequences containing hundreds of RF pulses (Vandersypen, 2001; Vandersypen *et al.*, 2001).

C. Fidelity of quantum states and gates

The methods of the previous section give us full knowledge of the system state and dynamics, but sometimes a more succinct measure for comparing theoretical expectations with experimental measurements is useful. This can be provided by quantum state fidelity and gate fidelity.

1. Quantum state fidelity

One elementary goal of quantum control is to create some pure state $|\psi\rangle$. However, suppose the final output is instead the pure state $|\phi\rangle$. Does $|\phi\rangle$ represent $|\psi\rangle$ with high fidelity?

Classically, the fidelity of two probability distributions $\{p_x\}$ and $\{q_x\}$ is given by $F(p_x, q_x) = \Sigma_x \sqrt{p_x q_x}$; when they are equal, the fidelity is one. The analogous quantum measure of fidelity for two pure states $|\psi\rangle$ and $|\phi\rangle$ is

$$F(|\psi\rangle, |\phi\rangle) = |\langle\phi|\psi\rangle|, \tag{110}$$

which is simply the absolute value of the overlap between the two states.

More generally, the output state of a control sequence is often described by a density matrix ρ; this is useful because density matrices can describe classical statistical mixtures of quantum states, arising from decoherence processes, for example. The fidelity between a pure state $|\psi\rangle$ and a mixed state ρ is

$$F(|\psi\rangle, \rho) = \sqrt{\langle\psi|\rho|\psi\rangle}, \tag{111}$$

which reduces to Eq. (110) when $\rho = |\phi\rangle\langle\phi|$.

The most general case is the fidelity between two density matrices, ρ and σ, which is defined as (Nielsen and Chuang, 2000)

$$F(\sigma, \rho) \equiv \mathrm{Tr}\sqrt{\sqrt{\sigma}\rho\sqrt{\sigma}}. \tag{112}$$

Despite the apparent asymmetry in this expression, it is actually symmetric in ρ and σ and, furthermore, reduces properly to Eq. (111) when one density matrix is pure.

Note that in the literature, sometimes the square of Eq. (110) is defined as the fidelity (Bowdrey *et al.*, 2002); this departs from the usual classical definition for fidelity, but is convenient because $F(|\psi\rangle, |\phi\rangle)^2$ can be interpreted as the probability that a system in the state $|\phi\rangle$ is found to be in the state $|\psi\rangle$ when measured in the $\{|\psi\rangle, |\psi_\perp\rangle\}$ basis. Such probabilities are meaningful in the accuracy thresholds discussed in Sec. V.D.

Other metrics for comparing two states have been used to quantify the relative error between theoretical and experimental states, such as the simple two-norm (Vandersypen *et al.*, 2000) and other expressions (Fortunato *et al.*, 2002). These were used because the diagonal elements of the density matrix were suppressed; such metrics are inferior to the fidelity measure, which should be used when possible, due to its direct connection to quantum information measures and fault-tolerance theorems.

It is worthwhile to consider a specific example relating control precision to state fidelity. Suppose we desire $|\psi\rangle = |1\rangle$, but obtain $|\phi\rangle = R_y(\pi + \epsilon)|0\rangle = -\sin\epsilon/2|0\rangle + \cos\epsilon/2|1\rangle \approx -\epsilon/2|0\rangle + (1 - \epsilon^2/8)|1\rangle$. The resulting error

probability is $1-|\langle\phi|\psi\rangle|^2=\epsilon^2/4$. This example makes the point that, for small rotation angle errors ϵ, the gate failure probability goes as ϵ^2.

2. Quantum gate fidelity

A more complex goal of quantum control is to accomplish a desired quantum operation. Perhaps the most common scenario is one in which the desired operation U is a unitary transform on a single qubit, whereas the actual transform accomplished is some quantum operation $\mathcal{E}$ (given in the operator sum representation).

A natural way to evaluate control precision is through the *average gate fidelity*

$$\bar{F}(\mathcal{E},U)=\int F(|\psi\rangle,U^\dagger\mathcal{E}(|\psi\rangle\langle\psi|)U)^2\,d\psi \tag{113}$$

$$=\int\langle\psi|U^\dagger\mathcal{E}(|\psi\rangle\langle\psi|)U|\psi\rangle\,d\psi, \tag{114}$$

where the integral is over the uniform (Haar) measure $d\psi$ on the Hilbert space of the system. For a single qubit, this formula can be reduced to a simple expression (Bowdrey *et al.*, 2002; Nielsen, 2002),

$$\bar{F}(\mathcal{E},U)=\frac{1}{2}+\frac{1}{12}\sum_{k=\{1,2,3\}}\mathrm{Tr}[U\sigma_kU^\dagger\mathcal{E}(\sigma_k)], \tag{115}$$

where σ_k are the three Pauli matrices. Similar simple formulas can be obtained for higher-dimensional systems (Nielsen, 2002). Note that by convention (Nielsen and Chuang, 2000), the average gate fidelity is defined such that it goes as the *square* of the usual state fidelity; thus it can be interpreted as a probability.

A more difficult quantity to calculate is the *minimum* gate fidelity,

$$F(\mathcal{E},U)=\min_{|\psi\rangle}F(U|\psi\rangle,\mathcal{E}(|\psi\rangle\langle\psi|)); \tag{116}$$

the square of this quantity gives the worst-case gate failure probability that is relevant for fault-tolerance threshold theorems.

D. Evaluating scalability

This article has been concerned with the control of complex systems composed of multiple distinct physical pieces. Given some degree of control over a few such pieces, how controllable is a very large quantum system composed of many pieces? Normally, one would expect that a system composed of unreliable pieces would itself be unreliable, and that the overall probability of failure increases rapidly with the number of pieces. Unexpectedly, however, arbitrarily reliable quantum systems can be built from unreliable parts as long as certain criteria are met.

The main criterion for being able to construct a reliable system is that the probability of error per operation p be below the "accuracy threshold" (Aharonov and Ben-Or, 1997; Gottesman, 1997; Kitaev, 1997; Knill *et al.*, 1998a; Preskill, 1998) p_{th}. When $p<p_{th}$ is satisfied, a quantum error correction circuit [for instance, as demonstrated by NMR experiments (Cory, Mass, *et al.*, 1998; Leung *et al.*, 1999; Knill *et al.*, 2001)] can be constructed using the unreliable components; this circuit performs computations on encoded qubits, such that a net decrease in error is achieved even when error correction itself is done with the faulty gates.

The probability of failure per operation p must of course be defined, and this is done in terms of the fidelity metrics discussed in the previous section, which incorporate decoherence (e.g., T_1, T_2, gate times, etc.) and control imperfections. Thus, for example, p is bounded from above by the gate fidelity $p\leqslant 1-F(\mathcal{E},U)^2$.

Remarkably, no reliable resources need be utilized for the fault-tolerant construction. Through k levels of recursive application of error correction, the device error p can be reduced to p^{2^k}, using physical resources (space, time, and energy), which scale as d^k for some constant d. Thus a small increase in resources exponentially reduces the overall error. Many assumptions are made in obtaining p_{th}, such as the availability of local, fast, parallel classical control resources, but the generally accepted theoretical optimal value of p_{th} is about 10^{-4} (Knill *et al.*, 1998a, 1998b), with optimistic estimates ranging as high as 10^{-3} with additional restrictions (Steane, 2002). As we have seen at the end of Sec. V.C.1, this implies that, for instance, rotation angles must be precise to order $\sim\sqrt{10^{-4}}=10^{-2}$. In principle, p_{th} can be experimentally measured by implementing a recursive error correction circuit and testing its probability of failure, but this has not yet been accomplished.

The fault-tolerance threshold p_{th} and its relative value compared with state and gate fidelities give a crisp criterion for system scalability for specific implementations. Modern classical systems are robust mainly because component failures can be controlled; similarly, the future of control over quantum systems hinges on our ability to evaluate p_{th} and to build components that fail with probability $p<p_{th}$.

VI. DISCUSSION AND CONCLUSIONS

In this review, we have presented a diverse set of tools intended to compensate for undesired or uncontrolled terms in the Hamiltonian of coupled qubits, as well as for instrumental limitations. These tools are most powerful and easiest to design when all the terms in the system Hamiltonian commute with each other and the control terms can be much stronger than the system Hamiltonian. The common theme of the control techniques is careful tailoring of the amplitude, phase, and frequency of the time-dependent terms in the Hamiltonian, whether in the form of shaped pulses, composite pulses, or multiple-pulse sequences. We now discuss the effectiveness and applicability of these advanced control techniques, with a look at where they could be used in other quantum systems.

- *Pulse shaping* is particularly attractive because of the modular and scalable design approach. Amplitude

profiles are selected from a library of standard or specially designed shapes in order to minimize cross-talk (frequency-selective pulses) and coupling effects (self-refocusing pulses). Robustness to experimental imperfections can also be considered in the choice of pulse shape. Once suitable amplitude profiles have been chosen, the pulse lengths are set as short as possible while maintaining qubit selectivity. The same amplitude profiles and pulse lengths are then used throughout the pulse sequence.

Remaining cross-talk effects can be further reduced at a small cost (quadratic in the number of qubits). From the amplitude profiles and pulse lengths, unintended phase shifts produced by single RF pulses as well as off-resonance effects during simultaneous pulses can be precomputed, once for every pair of qubits.

The main disadvantage of the standard pulse-shaping techniques is that often the coupled evolution during the pulses (in particular 90° pulses or simultaneous pulses) cannot be completely frozen. The remaining coupled evolution can be unwound to a large extent during the time intervals before and after the pulse, but such reversal is never perfect because the RF terms in the Hamiltonian, I_x^i and I_y^i, do not commute with the coupling terms, $I_z^i I_z^j$. Furthermore, shaped pulses must often be quite long in order to remain spin-selective, which means that decoherence has more effect. This problem clearly gets worse as the Larmor frequencies of the spins approach each other. Nevertheless, the combination of pulse-shaping and phase-ramping techniques has been very successful in practice. It has enabled the implementation of the most complex sequences of operations realized to date, acting on up to seven nuclear spins.

- *Composite pulses* have proven to be a versatile tool in NMR spectroscopy, mostly for compensating systematic errors such as RF field strength variations and frequency offsets. Another useful application of composite pulses is the effective creation of unitary operators which are otherwise not accessible or not easily accessible. A good example is the composite $\hat{z}$ rotation, created from a sequence of $\hat{x}$ and $\hat{y}$ rotations.

  Even so, the use of (hard) composite pulses in NMR quantum computing experiments has been limited so far. Their main drawback is that in multispin homonuclear molecules, single-frequency but high-power and rectangular pulses will rotate spins in a large frequency window, about an axis and over angles which depend on RF field strength and the respective resonance offsets. This severely limits straightforward application of hard composite pulses in homonuclear spin systems.

  Nevertheless, it is in principle possible to take advantage of the differences in resonance offsets in order to rotate one spin while the other spins undergo no net rotation. Such effective frequency selectivity despite the use of hard pulses was demonstrated in a quantum computation on a homonuclear two-spin system, first using single hard pulses (Jones and Mosca, 1999) and later using composite hard pulses (Cummins and Jones, 2000).

  The same idea underlies the operation of composite pulses tailored to achieve any rotation of one or more spins about an independent axis, using detailed knowledge of the system Hamiltonian. Furthermore, short, high-power pulses can be used, so the effect of decoherence is reduced compared to the case of the long, low-power shaped pulses. Even more attractive here is the fact that *all* the coupling terms can be effectively frozen and that other types of cross-talk, such as Bloch-Siegert effects, are automatically taken care of, unlike the case of shaped pulses.

  The main disadvantage of such strongly modulated composite pulses is that the time needed to find near-optimal pulse parameters increases exponentially with the number of qubits n, as it involves computing unitary matrices of size 2^n by 2^n. Nevertheless, for small numbers of qubits, this technique can be very useful.

- *Average-Hamiltonian* techniques underlie the operation of widely used multiple-pulse refocusing sequences. In the context of liquid NMR quantum computing, couplings are of the form $I_z^i I_z^j$ and refocusing sequences consist simply of a train of 180° pulses. Such refocusing sequences are an essential ingredient of all NMR quantum computing experiments involving more than two spins.

  More complex decoupling sequences exist to remove the effect of coupling Hamiltonians of a different form, as is the case of solid-state NMR and many other qubit implementations. Even errors arising from interactions with the environment, i.e., decoherence, can be removed using multiple-pulse sequences.

  In all cases, the refocusing operations (e.g., the 180° pulses) must be fast compared to the fluctuations they are intended to cancel, and they must also be repeated at a rate faster than the fluctuations.

- *Perspective.* In early NMR quantum computing experiments on heteronuclear spin systems, where short, high-power RF pulses were used, errors in the time evolution were usually dominated by various instrumental limitations. Most experiments on homonuclear systems, in contrast, made use of long, low-power pulses, and were limited by cross-talk and coupling effects. As the pulse techniques for coping with limitations of the Hamiltonian and instrumentation became more advanced, the field reached the point where errors due to imperfect quantum control were smaller than errors caused by decoherence.

  Reaching this point in many-qubit systems must be a prime objective for any implementation of quantum computers, along with reduction of decoherence itself. Quantum information and computation theory offers a common language which can facilitate transfer and translation of the techniques for coherent

control of coupled nuclear spins to other fields of physics. Such cross-fertilization has already started, in systems as diverse as trapped ions (Gulde *et al.*, 2003), excitons in quantum dots (Chen *et al.*, 2001), and Cooper pair boxes (Collin *et al.*, 2004), and is likely to accelerate the progress towards the elusive goal of complete control over quantum systems.

REFERENCES

Abragam, A., 1961, *Principles of Nuclear Magnetism* (Clarendon Press, Oxford).

Abramovich, D., and S. Vega, 1993, J. Magn. Reson., Ser. A **105**, 30.

Aharonov, D., and M. Ben-Or, 1997, in *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing* (ACM Press, New York), p. 176.

Altepeter, J. B., D. Branning, E. Jeffrey, T. C. Wei, P. Kwiat, R. Thew, J. L. O. Brien, M. A. Nielsen, and A. G. White, 2003, Phys. Rev. Lett. **90**, 193601.

Augustine, M., and E. Hahn, 1997, J. Chem. Phys. **107**, 3324.

Bauer, C., R. Freeman, T. Frenkiel, J. Keeler, and A. Shaka, 1984, J. Magn. Reson. (1969-1992) **58**, 442.

Bennett, C. H., and D. P. DiVincenzo, 2000, Nature (London) **404**, 247.

Berry, M., 1984, Proc. R. Soc. London, Ser. A **392**, 45.

Bloch, F., and A. Siegert, 1940, Phys. Rev. **57**, 522.

Boulant, N., T. Havel, M. Pravia, and D. Cory, 2003, Phys. Rev. A **67**, 042322.

Bowdery, M. D., D. K. Oi, A. J. Short, K. Banaszek, and J. A. Jones, 2002, Phys. Lett. A **294**, 258.

Braunstein, S., and H.-K. Lo, 2000, Eds., *Scalable Quantum Computers—Paving the Way to Realization* (Wiley, Berlin), reprint of Fortschr. Phys. **48**, 2000. Special issue on the implementation of quantum computers.

Bremner, M., C. Dawson, J. Todd, A. Gilchrist, A. Harrow, D. Mortimer, M. Nielsen, and T. Osborne, 2002, Phys. Rev. Lett. **89**, 247902.

Brockett, R., 1981, in *New Directions in Applied Mathematics*, edited by P. Hilton and G. Young (Springer, New York), pp. 11–27.

Brüschweiler, R., J. Madsen, C. Griesinger, O. Sørensen, and R. Ernst, 1987, J. Magn. Reson. (1969-1992) **73**, 380.

Byrd, M., and D. Lidar, 2003, Phys. Rev. A **67**, 012324.

Carr, H., and E. Purcell, 1954, Phys. Rev. **94**, 630.

Chen, P., C. Piermarocchi, and L. J. Sham, 2001, Phys. Rev. Lett. **87**, 067401.

Childs, A., I. Chuang, and D. Leung, 2001, Phys. Rev. A **64**, 012314.

Chuang, I. L., N. Gershenfeld, and M. Kubinec, 1998a, Phys. Rev. Lett. **80**, 3408.

Chuang, I. L., N. Gershenfeld, M. G. Kubinec, and D. W. Leung, 1998b, Proc. R. Soc. London, Ser. A **454**, 447.

Chuang, I. L., and M. A. Nielsen, 1997, J. Mod. Opt. **44**, 2455.

Chuang, I. L., L. M. K. Vandersypen, X. L. Zhou, D. W. Leung, and S. Lloyd, 1998, Nature (London) **393**, 143.

Clark, R., 2001, Ed., *Proceedings of the 1st International Conference on Experimental Implementation of Quantum Computation*, Sydney, Australia, January 2001 (Rinton, Princeton, NJ).

Cleve, R., A. Ekert, C. Macchiavello, and M. Mosca, 1998, Proc. R. Soc. London, Ser. A **454**, 339.

Collin, E., G. Ithier, A. Aassime, P. Joyez, D. Vion, and D. Esteve, 2004, Phys. Rev. Lett. **93**, 157005.

Cory, D., A. Fahmy, and T. Havel, 1996, in *Proceedings of the Fourth Workshop on Physics and Computation* (New England Complex Systems Institute, Boston, MA), pp. 87–91.

Cory, D., R. Laflamme, E. Knill, L. Viola, T. Havel, N. Boulant, G. Boutis, E. Fortunato, S. Lloyd, R. Martinez, C. Negrevergne, M. Pravia, *et al.*, 2000, Fortschr. Phys. **48**, 875.

Cory, D. G., A. F. Fahmy, and T. F. Havel, 1997, Proc. Natl. Acad. Sci. U.S.A. **94**, 1634.

Cory, D. G., W. Mass, M. Price, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo, 1998, Phys. Rev. Lett. **81**, 2152.

Cory, D. G., M. D. Price, and T. F. Havel, 1998, Physica D **120**, 82.

Cummins, H., and J. Jones, 2000, New J. Phys. **2.6**, 1.

D'Ariano, G., and P. Lo Presti, 2001, Phys. Rev. Lett. **86**, 4195.

Duan, L.-M., and G. Guo, 1999, Phys. Lett. A **261**, 139.

Ekert, A., and R. Jozsa, 1996, Rev. Mod. Phys. **68**, 1.

Emsley, L., and G. Bodenhausen, 1990, Chem. Phys. Lett. **168**, 297.

Ernst, R. R., G. Bodenhausen, and A. Wokaun, 1987, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Oxford University, Oxford).

Fortunato, E., M. Pravia, N. Boulant, G. Teklemariam, T. Havel, and D. Cory, 2002, J. Chem. Phys. **116**, 7599.

Freeman, R., *Spin Choreography* (Spektrum, Oxford).

Freeman, R., 1998, Prog. Nucl. Magn. Reson. Spectrosc. **32**, 59.

Galindo, A., and M. Martin-Delgado, 2002, Rev. Mod. Phys. **74**, 347.

Geen, H., and R. Freeman, 1991, J. Magn. Reson. (1969-1992) **93**, 93.

Gershenfeld, N., and I. Chuang, 1998, Sci. Am. (Int. Ed.) **278** (6), 66.

Gershenfeld, N., and I. L. Chuang, 1997, Science **275**, 350.

Glaser, S. J., T. Schulte-Herbrüggen, M. Sieveking, O. Schedletzky, N. C. Nielsen, O. W. Sørensen, and C. Griesinger, 1998, Science **280**, 421.

Gottesman, D., 1997, *Stabilizer Codes and Quantum Error Correction*, Ph.D thesis (California Institute and Technology).

Grant, D., and R. Harris, 2001, Eds., *The Encyclopedia of NMR* (Wiley, West Sussex, England).

Grover, L. K., 1997, Phys. Rev. Lett. **79**, 325.

Gulde, R., M. Riebe, G. Lancaster, C. Becher, J. Eschner, H. H. and F. Schmidt-Kaler, I. Chuang, and R. Blatt, 2003, Nature (London) **421**, 48.

Haeberlen, U., and J. Waugh, 1968, Phys. Rev. **175**, 453.

Hayashi, T., T. Fujisawa, H. D. Cheong, Y. H. Jeong, and Y. Hirayama, 2003, Phys. Rev. Lett. **91**, 226804.

Jeener, J., 1971, *Lectures Ampère International Summer School* (Basko Polje, Yogoslavia).

Jeener, J., 1982, Adv. Magn. Reson. **10**, 1.

Jones, J., 2000, Fortschr. Phys. **48**, 909.

Jones, J., 2001, Prog. Nucl. Magn. Reson. Spectrosc. **38**, 325.

Jones, J., 2003a, Phys. Rev. A **67**, 012317.

Jones, J., 2003b, Philos. Trans. R. Soc. London, Ser. A **361**, 1429.

Jones, J., and M. Mosca, 1999, Phys. Rev. Lett. **83**, 1050.

Jones, J. A., and E. Knill, 1999, J. Magn. Reson., Ser. A **141**, 322.

Jones, J. A., M. Mosca, and R. H. Hansen, 1998, Nature (London) **393**, 344.

Khaneja, N., S. Glaser, and R. Brockett, 2001, Phys. Rev. A **63**,

032308.
Khaneja, N., S. Glaser, and R. Brockett, 2002, Phys. Rev. A **65**, 032301.
Khitrin, A., V. Ermakov, and B. Fung, 2002, Chem. Phys. Lett. **360**, 161.
Kitaev, A. Y., 1997, Russ. Math. Surveys **52**, 1191.
Knill, E., R. Laflamme, R. Martinez, and C. Negrevergne, 2001, Phys. Rev. Lett. **86**, 5811.
Knill, E., R. Laflamme, R. Martinez, and C.-H. Tseng, 2000, Nature (London) **404**, 368.
Knill, E., R. Laflamme, and W. H. Zurek, 1998a, Science **279**, 342.
Knill, E., R. Laflamme, and W. H. Zurek, 1998b, Proc. R. Soc. London, Ser. A **454**, 365.
Kraus, K., 1983, *States, Effects, and Operations: Fundamental Notions of Quantum Theory*, Lecture Notes in Physics, No. 190 (Springer, Berlin).
Kupče, E., and R. Freeman, 1995, J. Magn. Reson., Ser. A **112**, 261.
Leibfried, D., R. Blatt, C. Monroe, and D. J. Wineland, 2003, Rev. Mod. Phys. **75**, 281.
Leskowitz, G., N. Ghaderi, R. Olsen, and L. Mueller, 2003, J. Chem. Phys. **119**, 1643.
Leung, D., L. Vandersypen, X. Zhou, M. Sherwood, C. Yannoni, and I. Chuang, 1999, Phys. Rev. A **60**, 1924.
Leung, D. W., I. L. Chuang, F. Yamaguchi, and Y. Yamamoto, 2000, Phys. Rev. A **61**, 042310.
Levitt, M., 1986, Prog. Nucl. Magn. Reson. Spectrosc. **18**, 61.
Levitt, M., 2001, *Spin Dynamics* (Wiley, West Sussex, England).
Levitt, M., and R. Freeman, 1979, J. Magn. Reson. (1969-1992) **33**, 473.
Levitt, M., R. Freeman, and T. Frenkiel, 1982, J. Magn. Reson. (1969-1992) **47**, 328.
Lidar, D. A., I. L. Chuang, and K. B. Whaley, 1998, Phys. Rev. Lett. **81**, 2594.
Linden, N., H. Barjat, R. Carbajo, and R. Freeman, 1999, Chem. Phys. Lett. **305**, 28.
Linden, N., E. Kupče, and R. Freeman, 1999, Chem. Phys. Lett. **311**, 321.
Lloyd, S., 1995, Sci. Am. (Int. Ed.) **273** (4), 140.
Magnus, W., 1954, Commun. Pure Appl. Math. **7**, 649.
Maklin, Y., G. Schön, and A. Schirman, 2001, Rev. Mod. Phys. **73**, 357.
Meekhof, D., C. Monroe, B. King, W. Itano, and D. Wineland, 1996, Phys. Rev. Lett. **76**, 1796.
Mehring, M., 1983, *High Resolution NMR Spectroscopy in Solids* (Springer, Berlin).
Meiboom, S., and D. Gill, 1958, Rev. Sci. Instrum. **29**, 688.
Mitchell, M., C. Ellenor, S. Schneider, and A. Steinberg, 2003, Phys. Rev. Lett. **91**, 120402.
Nakamura, Y., Y. A. Pashkin, and J. S. Tsai, 1999, Nature (London) **398**, 786.
Nelder, J., and R. Mead, 1965, Comput. J. (UK) **7**, 308.
Nielsen, M. A., 2002, Phys. Lett. A **303**, 249.
Nielsen, M., and I. Chuang, 2000, *Quantum Computation and Quantum Information* (Cambridge University, Cambridge, England).
Nielsen, M. A., E. Knill, and R. Laflamme, 1998, Nature (London) **396**, 52.
Osborne, I., and R. Coontz, 2002, Science **298**, 1353.
Patt, S., 1991, J. Magn. Reson. (1969-1992) **96**, 94.
Pauly, J., P. L. Roux, and D. Nishimura, 1991, IEEE Trans. Med. Imaging **10** (1), 53.
Poyatos, J. F., J. I. Cirac, and P. Zoller, 1997, Phys. Rev. Lett. **78**, 390.
Preskill, J., 1998, Proc. R. Soc. London, Ser. A **454**, 385.
Rabi, I., 1937, Phys. Rev. **51**, 652.
Ramsey, N., 1950, Phys. Rev. **78**, 695.
Redfield, A., 1957, IBM J. Res. Dev. **1**, 19.
Shor, P. W., 1995, Phys. Rev. A **52**, R2493.
Skovsen, E., H. Stapelfeldt, S. Juhl, and K. Molmer, 2003, Phys. Rev. Lett. **91**, 090406.
Slichter, C. P., 1996, *Principles of Magnetic Resonance* (Springer, Berlin).
Smithey, D., M. Beck. M. Raymer, and A. Faridani, 1993, Phys. Rev. Lett. **70**, 1244.
Somaroo, S., C. H. Tseng, T. F. Havel, R. Laflamme, and D. G. Cory, 1999, Phys. Rev. Lett. **82**, 5381.
Sørenson, O., 1989, Prog. Nucl. Magn. Reson. Spectrosc. **21**, 503.
Steane, A., 1998, Rep. Prog. Phys. **61**, 117.
Steane, A., 2002, quant-ph/0207119.
Steane, A. M., 1996, Phys. Rev. Lett. **77**, 793.
Steffen, M., 2003, *A Prototype Quantum Computer Using Nuclear Spins in Liquid Solution*, Ph.D. thesis (Stanford University).
Steffen, M., L. Vandersypen, and I. Chuang, 2000, J. Magn. Reson., Ser. A **146**, 369.
Steffen, M., L. Vandersypen, and I. Chuang, 2001, IEEE Micro. **21** (2), 24.
Stievater, T., X. Li, D. Steel, D. Katzer, D. Park, C. Piermarocchi, and L. Sham, 2001, Phys. Rev. Lett. **87**, 133603.
Tesch, C. M., and R. de Vivie-Riedle, 2002, Phys. Rev. Lett. **89**, 157901.
Thew, R., K. Nemoto, A. White, and W. Munro, 2002, Phys. Rev. A **66**, 012303.
Tycko, R., 1983, Phys. Rev. Lett. **51**, 775.
Tycko, R., H. Cho, E. Schneider, and A. Pines, 1985, J. Magn. Reson. (1969-1992) **61**, 90.
Vala, J., Z. Amitay, B. Zhang, S. R. Leone, and R. Kosloff, 2002, Phys. Rev. A **66**, 062316.
Vandersypen, L. M. K., 2001, *Experimental Quantum Computation with Nuclear Spins in Liquid Solution*, Ph.D. thesis (Stanford University).
Vandersypen, L. M. K., M. Steffen, G. Breyta, C. S. Yannoni, R. Cleve, and I. L. Chuang, 2000, Phys. Rev. Lett. **85**, 5452.
Vandersypen, L. M. K., M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, 2001, Nature (London) **414**, 883.
Vandersypen, L. M. K., M. Steffen, M. H. Sherwood, C. Yannoni, G. Breyta, and I. L. Chuang, 2000, Appl. Phys. Lett. **76**, 646.
Vandersypen, L. M. K., C. Yannoni, and I. Chuang, 2002, in *Encyclopedia of Nuclear Magnetic Resonance: Supplementary Volume*, edited by D. Grant and R. Harris (Wiley, West Sussex, England), p. 687.
Vandersypen, L. M. K., C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, 1999, Phys. Rev. Lett. **83**, 3085.
Viola, L., E. Knill, and S. Lloyd, 1998, Phys. Rev. Lett. **82**, 2417.
Viola, L., and S. Lloyd, 1998, Phys. Rev. A **58**, 2733.
Vitali, D., and P. Tombesi, 1999, Phys. Rev. A **59**, 4178.
Walmsley, I., and H. Rabitz, 2003, Phys. Today **56** (8), 43.
Warren, W., 1984, J. Chem. Phys. **81**, 5437.
Waugh, J., L. Huber, and U. Haeberlen, 1998, Phys. Rev. Lett.

20, 180.
Wieman, C., D. Pritchard, and D. Wineland, 1999, Rev. Mod. Phys. **71**, S253.
Wimperis, S., 1994, J. Magn. Reson., Ser. B **109**, 221.
Yannoni, C., M. Sherwood, L. Vandersypen, M. Kubinec, D. Miller, and I. Chuang, 1999, Appl. Phys. Lett. **75**, 3563.
Zanardi, P., and M. Rasetti, 1997, Phys. Rev. Lett. **79**, 3306.
Zeilinger, A., 1999, Rev. Mod. Phys. **71**, S288.
Zhang, W., and D. Cory, 1998, Phys. Rev. Lett. **80**, 1324.
Zutic, I., J. Fabian, and S. D. Sarma, 2004, Rev. Mod. Phys. **76**, 323.