2.27 Measurements: Consider the state

$$|\psi\rangle = \frac{1}{\sqrt{3}}|00\rangle + \frac{1}{3}|01\rangle + \frac{1}{\sqrt{3}}|10\rangle + \frac{\sqrt{2}}{3}|11\rangle . \tag{2.175}$$

The first qubit is measured to be in the state $|0\rangle$. What's the state of the system after this measurement? What's the probability that a subsequent measurement of qubit 2 will be measured in the state $|1\rangle$?

3 Nuclear Magnetic Resonance

Many physical degrees of freedom can encode a qubit. The first one we'll study is spin. There exist different approaches to quantum computing with spin states. These include liquid-state nuclear magnetic resonance (NMR), solid-state NMR, electron spin resonance, semiconductor quantum dots, and nitrogen-vacancy centers. This chapter looks at the earliest of these approaches, liquid-state NMR. In this approach, the qubits are encoded using the nuclear spins of atoms in a molecule. NMR quantum computers played an important role in the early days of quantum computing.

Felix Bloch and Edward Purcell first observed NMR in 1946, long before the advent of quantum computing. The two would win the Nobel Prize for this discovery in 1952. Since its discovery, the applications of NMR have become broad and diverse. These applications range from familiar medical ones, such as medical imaging, to lesser-known industrial ones, such as natural resource exploration. By 1997, the engineering and techniques of NMR had become well honed, and many NMR spectrometers were commercially available. That year, two groups proposed a protocol for quantum computing with NMR, one consisting of David Cory, Mark Price, and Timothy Havel, and the other of Neil Gershenfeld and Isaac Chuang. Also, in that year, an NMR spectrometer implemented a quantum algorithm. NMR spectrometers would be the first to perform many other algorithms and protocols.

This chapter explores how a liquid-state NMR system can satisfy the DiVincenzo criteria. Section 3.1 introduces the basic idea of spin and NMR spectroscopy. In Section 3.2 we explain how a qubit can be encoded in the spin states of nuclei in molecules. Next, we show how NMR control techniques are harnessed to perform quantum algorithms with the molecules. This is done in two parts. First, we see in Section 3.3 that single-qubit gates are implemented by applying external magnetic fields. We take this as an opportunity to introduce recurring topics in the text: the rotating-wave approximation, the Rabi cycle, and pulse shaping. Second, we see in Section 3.4 that two-qubit gates are implemented by regulating the natural coupling between neighbouring nuclei in the molecule. Together, these techniques give us universal quantum control. Section 3.5 introduces how the bulk nuclear magnetization of the NMR sample is used to determine the state of the qubits. This gives us all the elements needed to describe how the qubits are initialized, explained in Section 3.6. Two types of noise are pertinent to NMR quantum computing and are introduced in Section 3.7. This section also discusses the different interpretations of what "decoherence" means and introduces the T_1 and T_2 times. We conclude in Section 3.8 with a summary of the chapter, an overview of the strengths and weaknesses of liquid-state NMR for quantum computing, and a discussion on the role of entanglement in quantum computing.

3.1 NMR Background

Classical particles can have angular momentum describing the rotational element of dynamics. It is defined as $\vec{L} = \vec{r} \times \vec{p}$, where $\vec{r}$ is the distance of a particle from the origin and $\vec{p}$ its momentum. This definition makes sense for extended objects or ones at a distance from the origin. In quantum mechanics, a new form of angular momentum appears called **spin**. It is an intrinsic angular momentum and a fundamental property of matter, like mass or charge. Particles with spin can possess an intrinsic magnetic dipole moment, called the **magnetic moment**, which determines how they interact with a magnetic field.

The measured values of spin are quantized. They come in positive half-integer multiples of Planck's reduced constant $\hbar$. The **Stern–Gerlach experiment** demonstrates the quantization of spin. It was conceived by Otto Stern and implemented by Walther Gerlach. Otto Stern was awarded the Nobel Prize for this work in 1943. This chapter focuses on sets of systems with spin $\frac{\hbar}{2}$ (and thereon, we assume the $\hbar$ is implicit). Electrons, protons, and neutrons are the prototypical examples of spin-$\frac{1}{2}$ particles. The spin, s, of a spin-$\frac{1}{2}$ system is quantized such that $s = \pm\frac{1}{2}$, also known as **spin-up** $(+\frac{1}{2})$ and **spin-down**$(-\frac{1}{2})$.

Spin is represented in quantum mechanics by a Hermitian operator:

$$\hat{\vec{S}} := \hat{S}_x\vec{x} + \hat{S}_y\vec{y} + \hat{S}_z\vec{z}. \tag{3.1}$$

It is a vector and thus can point in different directions in three dimensions. Spin operators have components on the x, y- and z-axes. The spin operators' components are proportional to the Pauli matrices:

$$\hat{S}_x = \frac{\hbar}{2}\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}, \quad \hat{S}_y = \frac{\hbar}{2}\begin{bmatrix}0 & -i\\ i & 0\end{bmatrix}, \quad \text{and} \quad \hat{S}_z = \frac{\hbar}{2}\begin{bmatrix}1 & 0\\ 0 & -1\end{bmatrix}. \tag{3.2}$$

An NMR experiment requires an **NMR sample**. Preparing an NMR sample consists of diluting a sample of molecules in a solvent. This solution is placed in a small tube about the size of a pencil and positioned inside the magnet.

The other requirement for an NMR experiment is a **spectrometer**. One of the main components of a spectrometer is a magnet, typically a superconducting one, depicted in Fig. 3.1. The primary role of this magnet is to provide a strong and constant magnetic field to the sample. The magnet is a few feet in diameter and height, with the exact size depending on the magnet's strength. NMR quantum computing uses a magnetic field of about 100,000 gauss or 10 tesla. For reference, the magnetic field of the Earth is about 0.5 gauss. A circular coil of superconducting wire made of niobium–titanium alloy produces the spectrometer's magnetic field. This wire carries a current of about 100 amps, roughly the current of an entire household if all electrical devices are on at once. For the wires to become superconducting, they are immersed in a vessel filled with liquid helium. The helium is at a temperature of a few degrees kelvin. This helium vessel is surrounded by another container filled with liquid nitrogen. This second container reduces the heat exchange between the helium and the room. The magnetic field produced by the cooled wires is stable and has a highly homogeneous part. This homogeneous part is about 1 cm wide, a few centimetres tall, and located at the magnet's center. This is where the sample is placed.

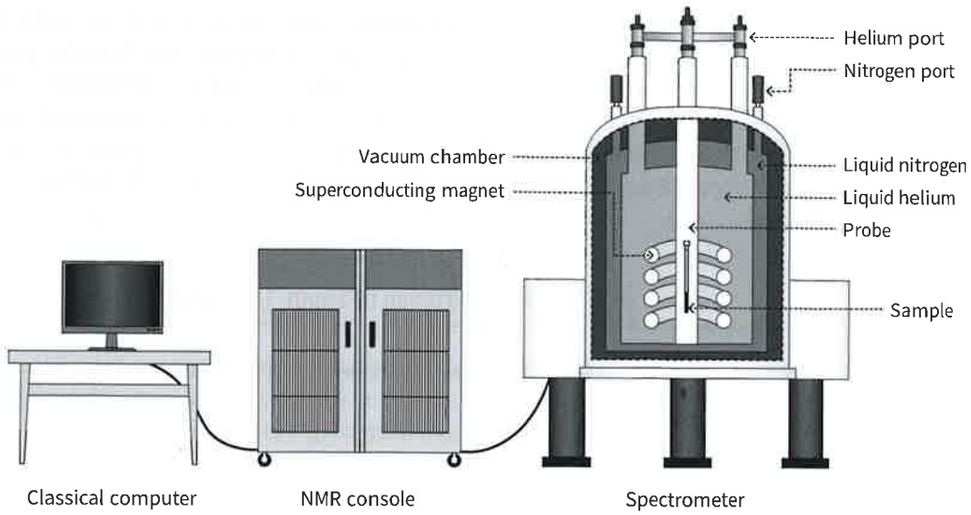

Fig. 3.1 **NMR spectrometer.** The NMR spectrometer has four distinct components: the superconducting magnet, the NMR console and corresponding electronics, the probe which sends or receives radiation to or from the sample, and a classical computer. The superconducting magnet is immersed in a vessel filled with liquid helium. The helium is held at a few degrees kelvin to keep the magnet at superconducting temperatures. This helium vessel is surrounded by another container filled with liquid nitrogen. This second container reduces the heat exchanged between the helium and the room. The sample is placed at the center of the magnet and inside the probe.

Around where the sample is placed are two sets of coils. One set is known as the **shimming coils**. These coils are used to improve the field's homogeneity, utilizing a process called **shimming**. In the old days, shimming was done by adding and removing small magnets called shims, hence the name. Today shimming is done by adjusting the current in the shimming coils. The second set of coils are the radio-frequency (RF) coils, which have a very different function. They are used to control and observe the nuclei.

An NMR spectrometer also requires a classical computer and an **NMR console**. These requirements are not unique to NMR quantum computing. Every quantum computer requires a classical computer for control and post-processing. The classical computer is the interface between the experimentalist and the quantum computer. The NMR console contains all the control and read-out electronics. Every quantum computer requires these kinds of electronics. Since these components exist for all quantum computers in some fashion, we won't mention them again in the following chapters.

3.2 Qubit

This section explains how we can encode a qubit in the distinct nuclear spins within a molecule. First, we present the Hamiltonian of a spin-$\frac{1}{2}$ nucleus in a magnetic field

(Section 3.2.1). We use the eigenstates of this Hamiltonian to encode the qubit. Next, we find that the states of the qubits are not stationary. The magnetic field from the spectrometer induces a torque on the spins that causes the nuclei to precess in the magnetic field. This precession is what allows us to distinguish different nuclei and thus encode different qubits (Section 3.2.2). Finally, the nuclei are not floating in a solution by themselves; they are part of a molecule. We will mention what makes a good molecule for NMR quantum computing (Section 3.2.3).

3.2.1 Hamiltonian of a Spin-1/2 Nucleus

We first consider the Hamiltonian of a nuclear spin in a magnetic field. Since the spin provides a magnetic dipole moment, a nucleus with spin will interact with a magnetic field. The Hamiltonian of a nucleus with spin interacting with a magnetic field $\vec{B}$ is

$$\hat{H} = -\vec{\hat{\mu}} \cdot \vec{B}, \tag{3.3}$$

where $\vec{\hat{\mu}} = \gamma \vec{\hat{S}}$ and γ is the **gyromagnetic ratio**. It is characteristic of each nucleus: For hydrogen nuclei (protons) $\gamma_H = 2\pi \cdot 42.58$ MHz/T and for carbon-13 nuclei $\gamma_{^{13}C} = 2\pi \cdot 10.71$ MHz/T.

First, let's look at the Hamiltonian when a spin-$\frac{1}{2}$ nucleus interacts with a strong constant magnetic field along the z-axis, $B_0\vec{z}$, originating from the magnet of the spectrometer. The nuclei we consider are not isolated but part of a molecule. They will be surrounded by an electronic cloud that will partially shield the nuclei from the magnetic field by a factor of $(1 - \delta_{CS})$. δ_{CS} is known as the **chemical shift**. Thus, the magnetic field felt by the nucleus is $\vec{B} = (1 - \delta_{CS})B_0\vec{z}$ and the Hamiltonian of the spin-$\frac{1}{2}$ nucleus used to encode the qubit is

$$\hat{H}_z = -\frac{\hbar}{2}(1 - \delta_{CS})\gamma B_0 \hat{\sigma}_z = -\frac{\hbar\omega_L}{2}\hat{\sigma}_z, \tag{3.4}$$

where $\omega_L := (1 - \delta_{CS})\gamma B_0$ is an angular frequency known as the **Larmor precession frequency** of the nucleus. This Hamiltonian induces a rotation of the magnetic moment around the z-axis. The energy difference between the two eigenstates of the Hamiltonian, i.e. the two spin states, is $\hbar\omega_L$.

We encode the qubit in the two energy eigenstates of Eq. (3.4). These are the spin-up state $|0\rangle$ with energy $-\hbar\omega_L/2$ and the spin-down state $|1\rangle$ with energy $\hbar\omega_L/2$. The magnetic field has thus separated the energy of the spin-up and spin-down state of the spin-$\frac{1}{2}$ sample. The energy levels $\pm\hbar\omega_L/2$ are known as the **Zeeman levels**, $\hbar\omega_L$ is known as **Zeeman splitting**, and $\hat{H}_z$ is known as the **Zeeman Hamiltonian**. Many different spin-$\frac{1}{2}$ nuclei have been used to encode qubits in this way, such as ^{1}H, ^{13}C, ^{15}N, ^{19}F, and ^{31}P. Note that we use ^{13}C instead of the more common ^{12}C isotope since ^{12}C is not a spin-$\frac{1}{2}$ nucleus.

As we will show in Section 3.3, nuclei with different Larmor frequencies will allow us to have a multi-qubit system if the difference in frequencies is sufficiently large. Equation (3.4) shows that increasing the field B_0 can increase that difference.

3.2.2 Larmor Precession

In addition to distinguishing the energy of the two qubit states, the magnetic field applies a torque to the nucleus, causing it to precess around the direction of the field, $\vec{B}$. We can see that the evolution is a precession by looking at the expectation of the spin operators. Consider a spin-$\frac{1}{2}$ nucleus initially pointing along some arbitrary axis. Its state can be written as (Eq. (2.61))

$$|\psi(0)\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle\,. \tag{3.5}$$

Since the system's Hamiltonian is time-independent, the system will evolve according to the unitary evolution $\hat{U}(t) = \exp(i\omega_L\hat{\sigma}_z t/2)$. After some time t the state $|\psi(0)\rangle$ will evolve to

$$\begin{aligned}|\psi(t)\rangle &= \hat{U}(t)\,|\psi(0)\rangle \\ &= e^{i\omega_L\hat{\sigma}_z t/2}\cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\omega_L\hat{\sigma}_z t/2}e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle \\ &= e^{i\omega_L t/2}\cos\left(\frac{\theta}{2}\right)|0\rangle + e^{-i\omega_L t/2}e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle\,,\end{aligned} \tag{3.6}$$

where in the second line we used Eq. (2.45). The expectation value of the spin component pointing in the x-direction is

$$\begin{aligned}\left\langle\hat{S}_x\right\rangle &= \langle\psi(t)|\hat{S}_x|\psi(t)\rangle \\ &= \frac{\hbar}{2}\Big(\langle 0|\,e^{-i\omega_L t/2}\cos\left(\tfrac{\theta}{2}\right) + \langle 1|\,e^{i\omega_L t/2}e^{-i\phi}\sin\left(\tfrac{\theta}{2}\right)\Big)\big(|0\rangle\langle 1| + |1\rangle\langle 0|\big) \\ &\quad\times\Big(e^{i\omega_L t/2}\cos\left(\tfrac{\theta}{2}\right)|0\rangle + e^{-i\omega_L t/2}e^{i\phi}\sin\left(\tfrac{\theta}{2}\right)|1\rangle\Big) \\ &= \frac{\hbar}{2}\Big(e^{i(-\omega_L t+\phi)}\cos\left(\tfrac{\theta}{2}\right)\sin\left(\tfrac{\theta}{2}\right)\langle 0|0\rangle + e^{-i(-\omega_L t+\phi)}\cos\left(\tfrac{\theta}{2}\right)\sin\left(\tfrac{\theta}{2}\right)\langle 1|1\rangle\Big) \\ &= \frac{\hbar}{2}\sin(\theta)\cos(-\omega_L t + \phi).\end{aligned} \tag{3.7}$$

Thus the spin-$\frac{1}{2}$ component along the x-axis oscillates. By performing a similar calculation, we find the expectation value of the spin pointing in the y- and z-directions to be

$$\left\langle\hat{S}_y\right\rangle = \frac{\hbar}{2}\sin(\theta)\sin(-\omega_L t + \phi) \quad \text{and} \quad \left\langle\hat{S}_z\right\rangle = \frac{\hbar}{2}\cos(\theta). \tag{3.8}$$

This means that the y-component experiences the same oscillation as the x-component but with a $\pi/2$ phase difference, and the z-component does not change with time. Thus, we can see that the spin-$\frac{1}{2}$ system is precessing around the z-axis. The axis of precession was fixed by the constant background magnetic field's direction provided by the large magnet of the spectrometer.

A spin-$\frac{1}{2}$ with a positive magnetic moment will precess around the negative direction of an applied magnetic field, as can be observed from Eqs. (3.7) and (3.8). This is a result we will use later.

3.2.3 Molecules Used

The nuclei are not isolated systems, but exist within the structure of a molecule. A molecule that contains n spin-$\frac{1}{2}$ nuclei can be used as an n-qubit quantum computer, so long as the spin-$\frac{1}{2}$ have different precession frequencies. As we will see, if two spin-$\frac{1}{2}$ nuclei have the same precession frequencies, controlling or measuring them independently is impossible. Consider again the expression for the precession frequency $\omega_L := (1 - \delta_{CS})\gamma B_0$. Nuclei can have different precession frequencies for two reasons. One is that the nuclei are of different species, such as ^{1}H and ^{13}C. In that case, they have different γ, and their frequencies in a field of many teslas will be separated by many megahertz, such as chloroform in Fig. 3.2a. However, nuclei of the same species can also have different precession frequencies if they have different chemical shifts, δ_{CS}. Consider, for example, the ^{13}C-labelled trichloroethylene in Fig. 3.2b. One carbon is bound to two Cl while the other is bound to one Cl and a ^{1}H. The magnetic field felt by each ^{13}C is slightly influenced by the electrons of its neighbouring atoms. This different chemical environment cause the two ^{13}C to have different chemical shifts and thus precess at different frequencies. As a result, trichloroethylene is a three-qubit sample.

One of the challenges in NMR quantum information processing is finding suitable molecules. We need to find nuclei with sufficiently large chemical shift differences. To control one nucleus' state without affecting the others on a timescale of ~1 ms, we need ~1 kHz range between the qubit frequencies. Fortunately, many catalogues tabulate the chemical shifts for many molecules.

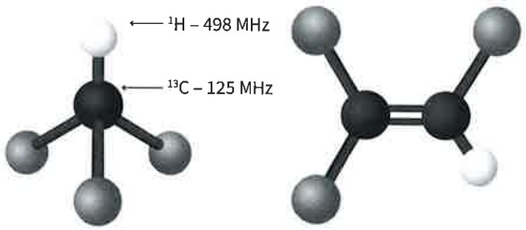

(a) Chloroform (b) Trichloroethylene

Fig. 3.2 **Molecular structure of chloroform and trichloroethylene.** (a) Chloroform and (b) trichloroethylene. Chloroform is a two-qubit sample, where the two-qubits are ^{13}C (black) and the ^{1}H (white). The chlorine nuclei (grey) interact very weakly with the hydrogen and carbon nuclei, so they're neglected. The magnetic moments of hydrogen and carbon nuclei have different precession frequencies ω_L since they have different gyromagnetic ratios γ. Trichloroethylene is a three-qubit sample. The three qubits are the two ^{13}C and the one ^{1}H. The two ^{13}C have different precession frequencies due to their chemical shifts. The approximate precession frequencies of the ^{13}C and ^{1}H nuclei given above are for a spectrometer with an 11.7 T magnet.

3.3 Single-Qubit Gates

Now that we have qubits, we need to learn how to control them. We can implement single-qubit gates with radio-frequency (RF) pulses (Section 3.3.1). Describing these pulses provides a natural opportunity to introduce two important concepts drawn on many times throughout this text: the rotating-wave approximation (Section 3.3.2) and the Rabi cycle. With these concepts, we see in Section 3.3.3 how a generic single-qubit gate is achieved with RF pulses. The quality of these RF pulses can be improved with techniques that we introduce in Section 3.3.4 and refer back to in future chapters.

3.3.1 Radio-Frequency Pulses

In Section 3.2.2 we saw the effect that a constant magnetic field has on a spin-$\frac{1}{2}$ nucleus. It applies a torque to the nuclear spin that causes it to precess at the Larmor frequency around the direction of the applied field. We will take advantage of the nuclei having different Larmor frequencies to address nuclei individually. By an **RF pulse**, here we mean an electromagnetic field that is turned on for a fixed period of time and then turned off. Applying an RF pulse generated by a coil pointing in, for example, the x-axis in the laboratory frame of reference at the Larmor frequency of a particular nucleus will result in changing its spin state. This rotation is depicted in Fig. 3.3.

Adding the magnetic field of an RF pulse complicates our system's Hamiltonian and, thus, its time evolution. The magnetic field now has two components. One is the constant magnetic field (in the $\vec{z}$-direction) that we previously studied, and the other is the applied RF pulse. We assume that the RF pulse is being applied along the x-axis with frequency and strength ω_{rf} and B_{rf}, respectively. Then the overall magnetic field is

$$\vec{B}(t) = (1 - \delta_{cs})B_0\vec{z} + B_{rf}\cos(\omega_{rf}t)\vec{x}. \tag{3.9}$$

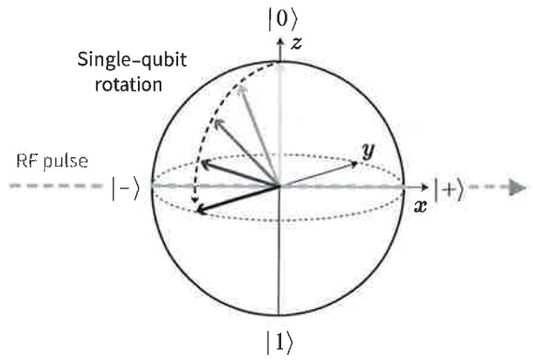

Fig. 3.3 **Single-qubit rotation in the rotating frame of reference.** An incident RF pulse is applied along the desired axis of rotation in the rotating frame of reference. In this figure, that direction is the x-axis. The RF pulse applies torque to the nucleus' spin, rotating the spin around the negative of the applied axis. This changes the state of the qubit. The degree of this rotation can be controlled by adjusting the power and/or duration of the RF pulse.

Due to the added field, the nucleus is no longer simply precessing about the z-axis. Determining its evolution is more complicated since we now have a time-dependent Hamiltonian,

$$\hat{H}(t) = -\vec{\hat{\mu}} \cdot \vec{B} = -\frac{\hbar\omega_L}{2}\hat{\sigma}_z - \hbar\omega_R \cos(\omega_{rf}t)\hat{\sigma}_x, \tag{3.10}$$

where $\omega_R := \gamma B_{rf}/2$ is the Rabi frequency, whose significance we discuss later. To make matters worse, the Hamiltonian doesn't commute with itself at different times, $[\hat{H}(t), \hat{H}(t')] \neq 0$. This is the complicated third type of solution to the Schrödinger equation that we encountered in Chapter 2 and said we would solve using approximate methods.

3.3.2 Rotating-Wave Approximation and Rabi Oscillations

We can approximate the effect of the time-dependent Hamiltonian in Eq. (3.10) using the **rotating-wave approximation** (RWA). We introduce it here and use it again in Chapters 5 and 6. To understand it, we will rewrite the time-dependent part of the Hamiltonian in terms of two components:

$$\vec{B}_{\pm}(t) = \frac{B_{rf}}{2}\left(\cos(\omega_{rf}t)\,\vec{x} \pm \sin(\omega_{rf}t)\,\vec{y}\right), \tag{3.11}$$

where $\vec{B}_{\pm}(t)$ can be thought of as positive and negative rotating fields around the z-axis. Intuitively, we would expect the negative rotation term to have a larger effect since it's in phase with the rotation resulting from the constant field. To see this, we will move to a rotating frame of reference rotating at angular frequency ω_{rf}. The operator that transforms the laboratory frame into this rotating frame is $\hat{R}_z(\omega_{rf}t) = \exp(-i\omega_{rf}t\hat{\sigma}_z/2)$. The state of the nucleus in the rotating frame is then

$$|\psi'(t)\rangle = \hat{R}_z(\omega_{rf}t)\,|\psi(t)\rangle\,. \tag{3.12}$$

We can now calculate the Hamiltonian in this reference frame. We can find the Schrödinger equation in the rotating frame by taking the time derivative of $|\psi'(t)\rangle$ and using the Schrödinger equation for $|\psi(t)\rangle$,

$$\begin{aligned} i\hbar\frac{\mathrm{d}}{\mathrm{d}t}|\psi'(t)\rangle &= i\hbar\frac{\mathrm{d}}{\mathrm{d}t}\left(\hat{R}_z(\omega_{rf}t)\,|\psi(t)\rangle\right) \\ &= i\hbar\left(\frac{\mathrm{d}}{\mathrm{d}t}\hat{R}_z(\omega_{rf}t)\right)|\psi(t)\rangle + i\hbar\hat{R}_z(\omega_{rf}t)\frac{\mathrm{d}}{\mathrm{d}t}|\psi(t)\rangle \\ &= \left(i\hbar\frac{\mathrm{d}}{\mathrm{d}t}\hat{R}_z(\omega_{rf}t) + \hat{R}_z(\omega_{rf}t)\hat{H}(t)\right)\hat{R}_z(\omega_{rf}t)^\dagger\,|\psi'(t)\rangle \qquad (3.13) \\ &= \hat{H}'(t)\,|\psi'(t)\rangle\,. \qquad (3.14) \end{aligned}$$

$\hat{H}'(t)$ is the effective Hamiltonian in the rotating frame of reference. The first term of $\hat{H}'(t)$ evaluates to

$$i\hbar\frac{\mathrm{d}\hat{R}_z(\omega_{rf}t)}{\mathrm{d}t}\hat{R}_z(\omega_{rf}t)^\dagger = i\hbar\left(-i\frac{\omega_{rf}}{2}\hat{\sigma}_z\right)\exp\left(-i\frac{\omega_{rf}}{2}t\hat{\sigma}_z\right)\exp\left(i\frac{\omega_{rf}}{2}t\hat{\sigma}_z\right) = \frac{\hbar\omega_{rf}\hat{\sigma}_z}{2}. \tag{3.15}$$

The second term of $\hat{H}'(t)$ evaluates to

$$\begin{aligned} \hat{R}_z(\omega_{rf}t)\hat{H}(t)\hat{R}_z(\omega_{rf}t)^\dagger &= e^{-i\omega_{rf}t\hat{\sigma}_z/2}\left[-\frac{\hbar\omega_L}{2}\hat{\sigma}_z - \hbar\omega_R\cos(\omega_{rf}t)\hat{\sigma}_x\right]e^{i\omega_{rf}t\hat{\sigma}_z/2} \\ &= -\frac{\hbar\omega_L}{2}\hat{\sigma}_z - \hbar\omega_R\cos(\omega_{rf}t)\left[\hat{\sigma}_x\cos(\omega_{rf}t) + \hat{\sigma}_y\sin(\omega_{rf}t)\right] \\ &= -\frac{\hbar\omega_L}{2}\hat{\sigma}_z - \frac{\hbar\omega_R}{2}\hat{\sigma}_x - \frac{\hbar\omega_R}{2}[\hat{\sigma}_x\cos(2\omega_{rf}t) + \hat{\sigma}_y\sin(2\omega_{rf}t)], \end{aligned} \tag{3.16}$$

where, as a reminder, $\omega_R := \gamma B_{rf}/2$. The Hamiltonian $\hat{H}'(t)$ is then:

$$\hat{H}'(t) = -\frac{\hbar(\omega_L - \omega_{rf})}{2}\hat{\sigma}_z - \frac{\hbar\omega_R}{2}\hat{\sigma}_x - \frac{\hbar\omega_R}{2}[\hat{\sigma}_x\cos(2\omega_{rf}t) + \hat{\sigma}_y\sin(2\omega_{rf}t)]. \tag{3.17}$$

The first term of the Hamiltonian corresponds to the remnant field of the magnet in the rotating frame. The second comes from the term $\vec{B}_-(t)$, rotating with the Larmor precession, and the last two terms come from counter-rotating $\vec{B}_+(t)$.

The RWA is valid when ω_{rf} is near ω_L and for $\omega_{rf} \gg \omega_R$. For now, we'll set $\omega_{rf} = \omega_L$. Later, we'll see what happens when ω_{rf} has a small deviation from ω_L. Thus, the Hamiltonian is

$$\hat{H}'(t) = -\frac{\hbar\omega_R}{2}\hat{\sigma}_x - \frac{\hbar\omega_R}{2}[\hat{\sigma}_x\cos(2\omega_{rf}t) + \hat{\sigma}_y\sin(2\omega_{rf}t)]. \tag{3.18}$$

The evolution consists of one time-independent term (the $\vec{B}_-(t)$ term) and two other terms that are rapidly oscillating (the $\vec{B}_+(t)$ term). The RWA consists of dropping the fast-oscillating terms. While an entirely rigorous proof to justify this approximation is beyond the scope of this book, we can give an intuitive motivation by considering the form of the time-evolution operator, which is the exponential of the integral $\int \hat{H}(t)\mathrm{d}t$. On the timescales of $t \gg 1/\omega_{rf}$, we see that the contribution of the fast-oscillating terms are strictly bounded, e.g. $|\int \cos(2\omega_{rf}t)\mathrm{d}t| \leq 1$, while the contribution of the static term continues to grow linearly with time, eventually overwhelming the oscillating terms. The RWA is depicted conceptually in Fig. 3.4. We also explore it numerically in the exercises.

When $\omega_L = \omega_{rf}$, then the Hamiltonian in the rotating frame reduces to

$$\hat{H}'(t) \approx -\frac{\hbar\omega_R}{2}\hat{\sigma}_x. \tag{3.19}$$

This Hamiltonian is what we set out to implement, a rotation of the qubit's state around the x-axis, but note that it is the x-axis in the rotating frame. We will therefore need to track this rotating frame as we quantum compute.

The continual application of this RF field will result in oscillations between the qubit states $|0\rangle$ and $|1\rangle$, known as **Rabi oscillations**. This cycling behaviour is called the **Rabi cycle** and the inverse of its duration is the **Rabi frequency**, which depends on the amplitude of the oscillating field. In the above example this was $\omega_R = \gamma B_{rf}/2$.

The Hamiltonian in Eq. (3.19) corresponds to a rotation around the x-axis for the qubit which has $\omega_{rf} = \omega_L$; this is called **on-resonance**. However, the pulse from Eq. (3.11) irradiates all the nuclei of the quantum computer. So what happens to the nuclei that are **off-resonance**, i.e. when the **detuning** $\Delta = \omega_{rf} - \omega_L \neq 0$? In that case the Hamiltonian

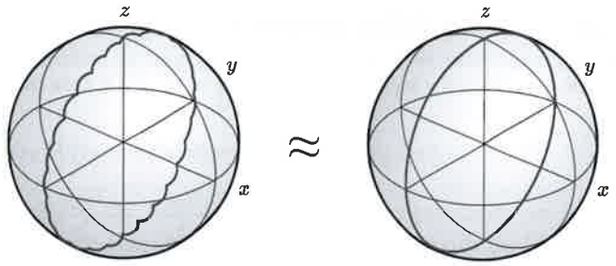

Fig. 3.4 **The rotating-wave approximation (RWA).** The trajectory on the sphere's surface depicts the evolution of the state. The trajectory on the left shows that the system's Hamiltonian has two terms. One term causes the system to orbit the sphere, and the other causes sinusoidal oscillations that are orthogonal to this orbit. As ω_L gets larger than ω_R, the second term oscillates faster and faster, and it can, on average, be ignored. This is the RWA depicted by the trajectory on the right.

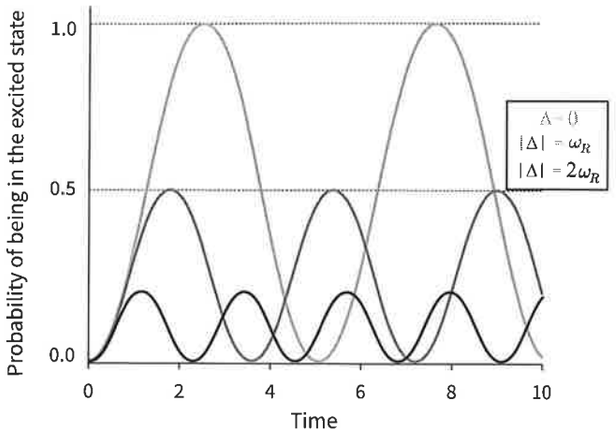

Fig. 3.5 **Rabi oscillations.** When the field's frequency ω_{rf} is on-resonance with the two-level systems ω_L, the state coherently oscillates between its ground and excited state. The larger the detuning between the field and systems frequency, the smaller the amplitude of the oscillations.

would be $\hat{H}'(t) = -\frac{\hbar\Delta}{2}\hat{\sigma}_z - \frac{\hbar\omega_R}{2}\hat{\sigma}_x$. This is no longer a rotation around the x-axis alone. It would be a rotation around the axis $\vec{n} \propto \omega_R\vec{x} + \Delta\vec{z}$. Thus, if the nucleus begins in the state $|0\rangle$, it will oscillate around some axis between the x- and z-axes and, in each cycle, never reach the state $|1\rangle$ (see Fig. 3.5). This is why the curves in Fig. 3.5 with $\Delta \neq 0$ never reach a probability of 1 of being in the excited state. As the detuning increases, the axis of rotation is further tilted toward the z-axis, and the maximum probability of reaching the excited states decreases even further. For $|\Delta| \gg |\omega_R|$, i.e. a large detuning, the rotation is approximately around the z-axis and has no effect on the nuclei. Thus, to target a specific qubit, we need to be on-resonance.

3.3.3 Universal Single-Qubit Control

We have just seen how to implement gates of the form $\exp\left(-i\frac{\hbar\omega_R}{2}t\hat{\sigma}_x\right)$, i.e. arbitrary rotations around the x-axis. By applying the RF field for a time $t = \frac{\pi}{\hbar\omega_R}$, we implement a single-qubit $\hat{\sigma}_x$ gate up to some global phase. Pulses of this duration are known as **π-pulses**. For a rotation by $\pi/2$ around the x-axis in the Bloch sphere, we would apply the RF field for half that time.

The next question is how to obtain a universal set of gates. One way to do so is by being able to rotate the qubit's state around at least one other axis. This can be achieved by modifying the phase of the RF pulse. The RF pulse considered in the previous subsection assumed that the magnitude of the RF field was at its maximum at $t = 0$. Repeating the above derivation for a more general RF pulse that has some phase ϕ, i.e.

$$\vec{B}(t) = B_0\vec{z} + B_{rf}\cos(\omega_{rf}t - \phi)\vec{x}, \tag{3.20}$$

we find that, after the RWA approximation and for $\omega_L = \omega_{rf}$, the final Hamiltonian in the rotating frame is

$$\hat{H}' = -\frac{\hbar\omega_R}{2}(\cos(\phi)\hat{\sigma}_x + \sin(\phi)\hat{\sigma}_y). \tag{3.21}$$

Thus, by changing the phase of the RF pulse to $\phi = \frac{\pi}{2}$, we obtain a $\hat{\sigma}_y$-component. Using this Hamiltonian, we can generate a universal set of gates for a single qubit.

As an example, we can implement rotations around the z-axis. By using three pulses with phases ϕ, 0, $\pi/2$ and 0, we get:

$$\begin{aligned}\hat{R}_x(\tfrac{\pi}{2})\hat{R}_y(\theta)\hat{R}_x(\tfrac{-\pi}{2}) &= \frac{1}{\sqrt{2}}\left[\hat{\mathbb{1}} - i\hat{\sigma}_x\right]\left[\cos(\tfrac{\theta}{2})\hat{\mathbb{1}} - i\sin(\tfrac{\theta}{2})\hat{\sigma}_y\right]\frac{1}{\sqrt{2}}\left[\hat{\mathbb{1}} + i\hat{\sigma}_x\right] \\ &= \cos(\tfrac{\theta}{2})\hat{\mathbb{1}} - i\sin(\tfrac{\theta}{2})\hat{\sigma}_z \\ &= \hat{R}_z(\theta).\end{aligned} \tag{3.22}$$

This series of rotations is depicted in Fig. 3.6. However, there is a better way to implement z-rotations.

Note that a z-rotation is equivalent to a change of reference frame. Remember that the gates are in the rotating frame of the qubit, so a z-rotation, let us say of $\pi/2$, changes the x-axis to the y-axis, the y-axis to the negative x-axis, and so on. Thus, instead of doing the rotation physically, we can redefine the rotating frame. This changes the labels of future rotations, i.e. the phase of the pulse. z-rotations are then just a question of **book-keeping** which rotating frame the qubit is in. An example of how to do this is to note that a z-rotation after a pulse in the xy-plane is equivalent to changing its phase and preceding it by the z-rotation:

$$\hat{R}_z(\theta)\hat{R}_x(\tfrac{\pi}{2}) = \hat{R}_z(\theta)\hat{R}_x(\tfrac{\pi}{2})\hat{R}_z(-\theta)\hat{R}_z(\theta) = \hat{R}_{\vec{n}}(\tfrac{\pi}{2})\hat{R}_z(\theta), \tag{3.23}$$

where $\vec{n} = \cos(\theta)\hat{x} + \sin(\theta)\hat{y}$. We can then move all the z-rotations to the beginning of the computation (presuming the two-qubit gates also commute with the z-rotations). Assuming that we start in the initial state $|0\rangle$, the z-rotations gate doesn't affect the initial state, so we

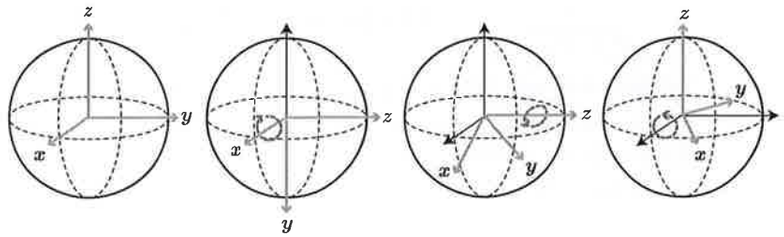

(a) Initial state (b) π-rotation around x-axis (c) θ-rotation around y-axis (d) π-rotation around x-axis

Fig. 3.6 **Rotation around the z-axis done with a combination of x- and y-axis rotations.** (a) The reference frame of a state which we would like to rotate by θ around the z-axis. This is achieved by a series of x- and y-axis rotations. First (b) a $-\pi/2$-rotation is done around the x-axis. Then (c) a θ-rotation is done around the y-axis. Finally, (d) a $\pi/2$-rotation is done around the x-axis. The final state is the initial state rotated by θ around the z-axis.

can ignore those gates. In this approach, the algorithm being implemented is reorganized before it is implemented to avoid having to do any z-rotations.

3.3.4 Techniques to Reduce Pulse Errors

Finally, we conclude our discussion on single-qubit control by considering some techniques used to reduce the error in pulses. These techniques are now used in many quantum computing implementations. We will introduce them here and refer back to them in the subsequent chapters.

So far we have assumed that we can do perfect pulses. In practice, imperfections such as miscalibrations and field inhomogeneities introduce errors. **Composite pulses** improve the accuracy of pulses by sending multiple pulses that vary in their power, phase, and direction. We'll illustrate the effect of composite pulses with the following example. Assume you have calibrated your spectrometer so that a pulse of a fixed duration and power implements a $\frac{\pi}{2}$-pulse. However, the calibration might be off by some ϵ. If you attempt to implement a π-pulse along the y-axis starting from the $|0\rangle$ state using two $\frac{\pi}{2}$-pulses, you'll implement a rotation of $\hat{R}_y(\pi + 2\epsilon)$. To mitigate this error, you can insert a $\hat{R}_x(\pi + 2\epsilon)$ pulse between the two $\frac{\pi}{2}$-pulses. If ϵ is negative, "under-shooting" the correct pulse duration, a $\hat{R}_x(\pi + 2\epsilon)$ would cause you to have now "over-shot." If ϵ is positive, "over-shooting" the correct pulse duration, an $\hat{R}_x(\pi + 2\epsilon)$ would cause you to have now "under-shot." Thus, a $\hat{R}_x(\pi + 2\epsilon)$ reverses the sign of the error. You can then apply the next $\hat{R}_y(\frac{\pi}{2} + \epsilon)$, the two errors will partially cancel out, and you will now have completed a pulse that is closer to being a $\hat{R}_y(\pi)$. There will still be an error left due to the error in $\hat{R}_x(\pi + 2\epsilon)$, but it will be of order ϵ^2. We will explore this further in the exercises. The particular composite pulse we have described works only if we start in the state $|0\rangle$. However, a universal pulse that works on all states can be achieved using five pulses.

Up to now we have assumed that gates are implemented by turning RF pulses on and off instantaneously. However, the physical pulse generator can only approximate this behaviour.

These idealized instantaneous pulses are sometimes called **square pulses**. We can improve the precision with which we apply pulses with **pulse shaping**. We use computational techniques to find more complex pulses to emulate the desired rotation with power and phase that changes smoothly – called **shaped pulses**. These pulses make it possible to target a single qubit even when the difference between the Larmor precession frequencies of that qubit and others is small, e.g. when there is a small difference in chemical shifts. Thus, shaped pulses can be engineered to mitigate **crosstalk**, an effect in which a signal transmitted to one qubit has an undesired effect on other qubits. Moreover, shaped pulses can also make rotations more robust against decoherence or miscalibration. The pulse-shaping technique has been used in NMR for some time and is now used in many other quantum computing implementations. A common technique for pulse shaping is gradient ascending pulse engineering (GRAPE). Square pulses between two different species of nuclei (heteronuclear) take a few microseconds, but shaped pulses will take a few milliseconds between the same species.

To summarize this section, single-qubit rotations in NMR are implemented by sending a pulse at the Larmor frequency of the targeted qubit. Four parameters describe a pulse. The first parameter is the pulse frequency, ω_{rf}, which tells us which qubit is affected. The second parameter is the pulse's phase, ϕ, which is used to specify the rotation axis in the xy-plane. The last two parameters are the pulse's power (the strength of B_{rf} or ω_R) and time. The product of these, $\omega_R t$, determines the degree of the rotation around the defined axis. As part of programming an NMR quantum computer, these parameters will be specified for each qubit pulse. Here we have implicitly assumed we used square pulses. If we used shaped pulses instead, we would also have had to give the shape itself. During the quantum computation, each qubit is followed in its rotating frame corresponding to its Larmor frequency, where the Hamiltonian is time-independent during the pulse.

3.4 Two-Qubit Gates

We now complete a universal gate set by adding a two-qubit gate. There exists a natural interaction between the spins of two nuclei that produces a two-qubit interaction. We study this interaction in Section 3.4.1. In Section 3.4.2, we'll see how to use this interaction to implement a CNOT gate. Unlike other quantum computing platforms, the two-qubit gate is always "turned on" in NMR. We'll present a technique to "turn off" this interaction in Section 3.4.3.

3.4.1 J-Coupling

To implement a two-qubit gate, we'll consider the natural interactions between the nuclei. We'll first investigate the common interaction between two magnets, the dipole–dipole interaction. It is given by

$$\hat{H}_D = -\frac{\mu_0 \gamma_1 \gamma_2 \hbar^2}{4\pi |\vec{r}|^3} \left(3(\vec{\sigma}^{(1)} \cdot \vec{r})(\vec{\sigma}^{(2)} \cdot \vec{r}) - \vec{\sigma}^{(1)} \cdot \vec{\sigma}^{(2)}\right), \tag{3.24}$$

where μ_0 is the permeability of free space, $\vec{r}$ is the vector connecting the two nuclei, and $\vec{\sigma}^{(i)}$ is the vector of Pauli matrices for nucleus i. For this discussion, we have included the superscripts on operators to clarify what qubit the operator is acting on. In liquid-state NMR, molecules are rapidly tumbling on a timescale of nanoseconds. Thus, this interaction averages to zero and cannot be used for a two-qubit gate.

However, there exists a weaker indirect interaction between the spins of two nuclei which is known as the **scalar J-coupling**. This interaction is mediated by surrounding electrons and is independent of the molecules' orientation. The Hamiltonian for the J-coupling between the ith and jth nuclei is

$$\hat{H}_J = \frac{\hbar\pi J_{ij}}{2}\vec{\sigma}^{(i)} \cdot \vec{\sigma}^{(j)}. \tag{3.25}$$

If the difference between the Larmor frequencies of nuclei 1 and 2 is much larger than the value of J_{12}, then $\vec{\sigma}^{(1)} \cdot \vec{\sigma}^{(2)}$ reduces to $\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}$. The proof of this is an exercise. Since we mostly work in this limit in liquid-state NMR experiments, we will make this assumption for the rest of the chapter. The strength of the J-coupling is typically much larger between neighbouring nuclei, and weaker for more distant nuclei. Thus, in NMR, we typically can only implement two-qubit gates between the nearest neighbours.

To get intuition on the effect of the J-coupling, we'll study a simple example. First, consider a two-qubit NMR system, such as ^{13}C–chloroform from Fig. 3.2. The Hamiltonian is given by,

$$\hat{H}_{2q} = -\frac{\hbar\omega_{L1}}{2}\hat{\sigma}_z^{(1)} - \frac{\hbar\omega_{L2}}{2}\hat{\sigma}_z^{(2)} + \frac{\hbar\pi J_{12}}{2}\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}. \tag{3.26}$$

To understand the effect of the interaction, we can ask: What will happen to the first qubit if the second one is in the state $|0\rangle$? Then it is an eigenstate of $\hat{\sigma}_z^{(2)}$ with eigenvalue 1. The effective Hamiltonian for the first qubit will simplify to

$$\langle 0|_2\hat{H}_{2q}|0\rangle_2 = -\frac{\hbar\omega_{L1}}{2}\hat{\sigma}_z^{(1)} - \frac{\hbar\omega_{L2}}{2}\hat{\mathbb{1}} + \frac{\hbar\pi J_{12}}{2}\hat{\sigma}_z^{(1)} \tag{3.27}$$

$$= -\frac{\hbar\omega_{L2}}{2}\hat{\mathbb{1}} + \left(-\frac{\hbar\omega_{L1}}{2} + \frac{\hbar\pi J_{12}}{2}\right)\hat{\sigma}_z^{(1)}. \tag{3.28}$$

Up to a constant $(-\frac{\hbar\omega_{L2}}{2})$, this Hamiltonian is of the same form as the Zeeman Hamiltonian for the first qubit, but with a lower frequency (if J_{12} is positive). If the second qubit had been in the state $|1\rangle$ (an eigenstate of $\hat{\sigma}_z^{(2)}$ with eigenvalue -1) then the Hamiltonian would have reduced to

$$\langle 1|_2\hat{H}_{2q}|1\rangle_2 = \frac{\hbar\omega_{L2}}{2}\hat{\mathbb{1}} + \left(-\frac{\hbar\omega_{L1}}{2} - \frac{\hbar\pi J_{12}}{2}\right)\hat{\sigma}_z^{(1)}, \tag{3.29}$$

which again is of the same form as the original Zeeman Hamiltonian, Eq. (3.4), except that the first qubit is now precessing faster. Thus, we can see that the effect of the J-coupling is that of a controlled rotation; if the second qubit is in the state $|0\rangle$ the first qubit rotates slower, and if the second qubit is in the state $|1\rangle$ the first qubit rotates faster.

As the phase in the z-basis of the qubit tracks rotation around the z-axis, the J-coupling thus implements a control-phase gate. To see this, we can rewrite the two-qubit unitary evolution for a time $t = \frac{1}{2J_{12}}$ (where $\hat{U}_J(t) := e^{-i\hat{H}_J t/\hbar}$):

$$\hat{U}_J\left(\frac{1}{2J_{12}}\right) = e^{-i\frac{\pi}{4}\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}} = e^{-i\frac{\pi}{4}\hat{\sigma}_z^{(1)}} \otimes |0\rangle\langle 0|_2 + e^{i\frac{\pi}{4}\hat{\sigma}_z^{(1)}} \otimes |1\rangle\langle 1|_2\,, \tag{3.30}$$

which shows the control phase explicitly. If the second qubit is in the state $|0\rangle$ then the first qubit is rotated by $\frac{\pi}{2}$, and if the second qubit is in the state $|1\rangle$ the first qubit is rotated by $\frac{-\pi}{2}$. This follows from the rotation operator, where rotation by angle θ is given by $e^{-i\frac{\theta}{2}\hat{\sigma}_z}$. As the Hamiltonian is symmetric through the exchange of qubits 1 and 2, the same argument works for qubit 1 being the control.

One great simplification in implementing the two-qubit gate is that we don't need to worry about the effect of the J-coupling when we're only implementing one-qubit gates. This is because they're implemented on a much faster timescale than the two-qubit gates.

3.4.2 Implementing a CNOT

Along with single-qubit gates, the two-qubit interaction from the J-coupling (a CZ gate) is sufficient for universal quantum computing. We know from Chapter 2 that we can implement a CNOT gate with a CZ gate and single-qubit pulses. Here we illustrate how this is done in NMR.

Let us consider the first qubit being the control and the second qubit being the target, as seen in Fig. 3.7. Then, the CNOT gate can be written in the form

$$\hat{U}_{\text{CNOT}} = |0\rangle\langle 0|_1 \otimes \hat{\mathbb{1}} + |1\rangle\langle 1|_1 \otimes \hat{\sigma}_x, \tag{3.31}$$

which has some similarities with Eq. (3.30). We will show how to transform Eq. (3.30) into Eq. (3.31). First, the control and the target of the CNOT are reversed, which is easy to fix by relabelling the qubits. The second step is that we would like to have nothing happening if the control qubit is in $|0\rangle$. This can be fixed by implementing a $\hat{R}_z^{(2)}(-\frac{\pi}{2})$ after $\hat{U}_J(\frac{1}{2J_{12}})$. The next step is to transform $e^{i\pi\hat{\sigma}_z/2} = i\hat{\sigma}_z$ into a $\hat{\sigma}_x$ gate. This can be done by rotating the gate. This is accomplished with the pulses $\hat{R}_y^{(2)}(\frac{\pi}{2})$ and $\hat{R}_y^{(2)}(-\frac{\pi}{2})$ before and after the J-coupling. We are nearly done, except that we are implementing a control $-i\hat{\sigma}_x$ instead of $\hat{\sigma}_x$. This can be fixed by changing the phase of the first qubit by applying a $\hat{R}_z^{(1)}(\frac{\pi}{2})$ gate. Note that $\hat{R}_y^{(2)}(-\frac{\pi}{2})\hat{R}_z^{(2)}(-\frac{\pi}{2}) = \hat{R}_x^{(2)}(\frac{\pi}{2})\hat{R}_y^{(2)}(-\frac{\pi}{2})$. Putting all of this together, we end up with:

$$\hat{R}_z^{(1)}\left(\tfrac{\pi}{2}\right)\hat{R}_x^{(2)}\left(\tfrac{\pi}{2}\right)\hat{R}_y^{(2)}\left(-\tfrac{\pi}{2}\right)\hat{U}_J\left(\tfrac{1}{2J_{12}}\right)\hat{R}_y^{(2)}\left(\tfrac{\pi}{2}\right) = e^{\frac{i\pi}{4}}\hat{U}_{\text{CNOT}}. \tag{3.32}$$

This evolution gives you the CNOT up to a global phase.

To implement this quantum circuit, we need to translate it into the spectrometer's language. Here is an example of what it looks like in a spectrometer from the company Bruker:

```
(C2_90:sp9 ph13):f1
3 u
3 u ipp13
0.71365 ms
3 u
```

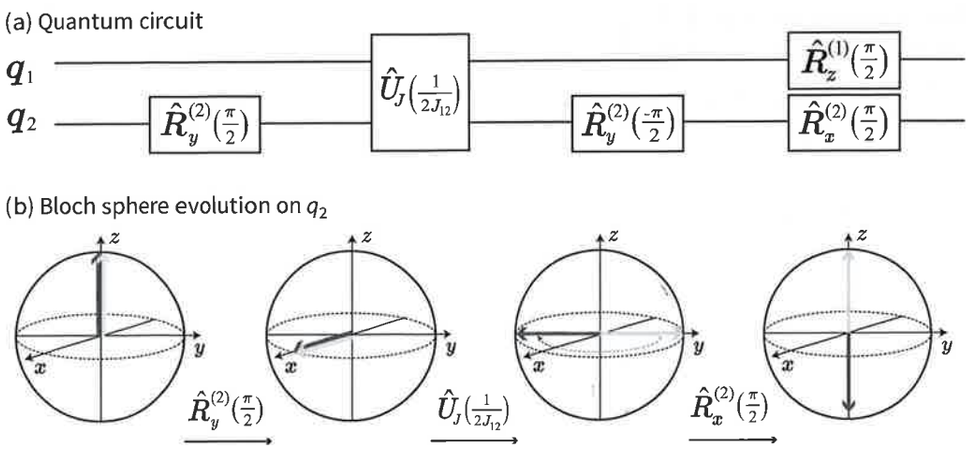

Fig. 3.7 **CNOT gate in NMR.** q_1 is the control and q_2 is the target qubit. (a) Decomposition of the CNOT in the natural gates of NMR; $\pi/2$ rotations in the xy-plane and generic z-rotations, with two-qubit $\hat{\sigma}_z{}^{(1)}\hat{\sigma}_z{}^{(2)}$ rotations. (b) The corresponding evolution of the target qubit depicted on the Bloch sphere. In this depiction, q_2 begins in the initial state $|0\rangle$. The lighter grey arrow depicts the evolution of q_2 when q_1 is in the state $|0\rangle$, and the darker grey arrow depicts its evolution when q_1 is in the state $|1\rangle$.

```
(C2_90:sp9 ph19):f1
6 u ipp13
8 u
(C2_90:sp9 ph20):f1
6 u ipp20
```

where we have used two kinds of commands. The first command, (C2_90:sp9 ph13):f1, is an example of a pulse. The inputs to this command are, in the order that they appear in the algorithm: the qubit we are acting on and the angle of rotation around the Bloch sphere (C2_90), the pulse shape (sp9), the phase of the pulse (ph19), and the channel which applies the pulse (f1). A three-microsecond delay (3 μs) is needed for the electronics to settle down from the pulse and update the phase (ipp13). Now it is the turn of the J-coupling, where the spectrometer waits for 0.71 365 milliseconds. Then we have the y and x pulses interspersed with delays of 3 μs. You have now seen how a quantum algorithm is implemented.

3.4.3 Refocusing

One of the challenges with the NMR two-qubit gates is that they're always present since it's part of the system's natural Hamiltonian. This raises the question of how to selectively "turn off" the coupling when it's not part of an algorithm. **Refocusing** is a technique that achieves this by inserting pulses at judicious places to cancel the unwanted evolution.

If we have two qubits, this can be achieved by waiting for the J-coupling to evolve the system for a time t, then applying pulses that reverse the sign of a second J-coupling

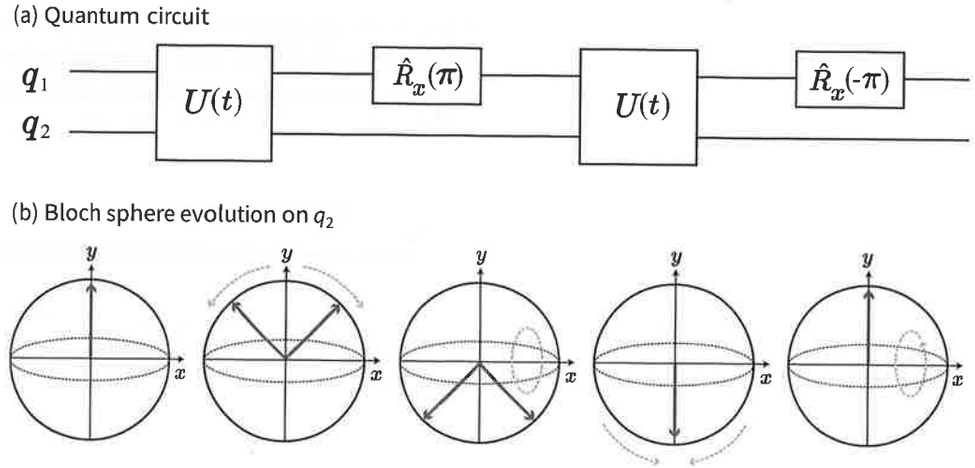

Fig. 3.8 **Refocusing in NMR.** (a) The two-qubits refocusing circuit that is used to "turn off" the J-coupling (Section 3.4.3). (b) The evolution of q_1 on the Bloch sphere, where we show only the projection on the xy-plane. The one arrow is replaced by two arrows to depict the superposition of spin-ups and spin-downs accumulating different phases from the second spin. In (b), the initial state of the first qubit is $\frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$.

evolution of time t. After an evolution of time $2t$, the qubit evolves as if $J = 0$. The pulse sequence for refocusing a two-qubit system is

$$\hat{U}_{\text{refocus}} = \hat{R}_x^{(1)}(-\pi)\hat{U}_J(t)\hat{R}_x^{(1)}(\pi)\hat{U}_J(t), \tag{3.33}$$

and is depicted in Fig. 3.8. We can work out the explicit corresponding evolution. In the Larmor frequency rotating frame of each qubit, the Hamiltonian consists of only the J-coupling terms $\hat{H} = \frac{\hbar\pi}{2}J_{12}\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}$. The free evolution is then $U(t) = e^{-i\hat{H}t/\hbar} = e^{-\frac{i\pi J t}{2}\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}}$. The later component of the refocusing evolution is

$$\begin{aligned}\hat{R}_x^{(1)}(-\pi)\hat{U}_J(t)\hat{R}_x^{(1)}(\pi) &= i\hat{\sigma}_x^{(1)}\left(e^{-\frac{i\pi J_{12}t}{2}\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}}\right)(-i\hat{\sigma}_x^{(1)}) \\ &= \hat{\sigma}_x^{(1)}\left(\cos\left(\frac{\pi}{2}J_{12}t\right)\hat{\mathbb{1}} - i\sin\left(\frac{\pi}{2}J_{12}t\right)\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}\right)\hat{\sigma}_x^{(1)} \\ &= \cos\left(\frac{\pi}{2}J_{12}t\right)\hat{\mathbb{1}} + i\sin\left(\frac{\pi}{2}J_{12}t\right)\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)} \\ &= e^{i\frac{\pi}{2}J_{12}t\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}}.\end{aligned} \tag{3.34}$$

This is the J-coupling evolution with a negative sign, i.e. the J-coupling in reverse. Thus, the entire refocusing procedure is

$$\hat{U}_{\text{refocus}} = \hat{R}_x^{(1)}(-\pi)\hat{U}_J(t)\hat{R}_x^{(1)}(\pi)\hat{U}_J(t) = e^{i\frac{\pi}{2}J_{12}t\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}}e^{-i\frac{\pi}{2}J_{12}t\hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}} = \hat{\mathbb{1}}. \tag{3.35}$$

We find that this is equivalent to no evolution occurring, as we hoped. This simple refocusing procedure can be generalized to the case with three or more qubits to remove undesired J_{ij}-coupling evolutions.

Figure 3.8 depicts this evolution and shows the potentially divergent spin directions coming back together at the end of the sequence. The effect when the nuclear spins realign is called a **spin echo**. The name comes from the signal reappearing after some time, as an echo does. We will revisit spin echo in Section 3.7.2.

3.5 Measurement

Once we have finished an algorithm, implemented by a series of pulses on a set of qubits, the qubits will typically have some component of their magnetization in the xy-plane rotating at their Larmor frequencies. To measure this magnetization, we use the same coil used for pulsing. Without a pulse, the nuclei will keep rotating in the laboratory frame and induce a changing flux in the direction of the coil. Let's assume it's in the x-direction (see Fig. 3.9). Faraday's law tells us that a time-dependent magnetic flux inside a coil induces a voltage and produces a current in this coil. By measuring this current, we measure the state of the qubits.

We first review the basics of NMR measurements (ensemble measurement) in Section 3.5.1. These measurements give some information about the final state of the system. To recover the exact state we need to introduce tomography (Section 3.5.2).

3.5.1 Ensemble Measurements

Up until this point, we have been thinking of the quantum computer as a single molecule. Unfortunately, a single molecule induces a current that is much too weak to be detected with current technology. We need an ensemble of about 10^{15} spins to be able to measure

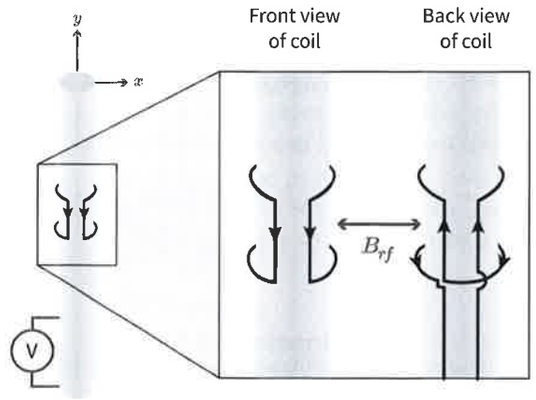

Fig. 3.9 **RF coils used to control and measure an NMR sample's signal.** When the NMR sample is placed in the spectrometer, it will be positioned so that the coils surround it as shown. The small arrow shows the direction of the current at a given time. A detailed analysis of this coil shows that an oscillating current produces an oscillating magnetic field in the x-direction, as shown by the double arrow labelled B_{rf}. To implement a gate, an RF pulse is sent through a coil. To measure the sample's state, no pulse is sent and the current in the wire is measured.

the current. This ensemble is made of many "quantum computing" molecules that are dissolved in a solvent. These different quantum computing molecules have the same evolution and don't interact with one another. What we measure is the current induced by the **bulk magnetization**, coming from the magnetization in the xy-plane of all 10^{15} quantum computing molecules. In NMR, what we're measuring are expectation values and not single-shot measurements. This feature is unlike the other quantum computers we study in this text.

We begin by calculating the expectation value of the nuclei's magnetization. We denote by $\hat{\rho}(t)$ the state of the system at time t in the laboratory frame. We denote by $\hat{M}_x$ and $\hat{M}_y$ the operators for the magnetization along the x-direction and along the y-direction, respectively. The expectation values of these operators for an n-qubit molecule are

$$\langle \hat{M}_x(t) \rangle = m\mathrm{tr}\left[\hat{\rho}(t)\sum_{i=1}^{n}\hat{\sigma}_x^{(i)}\right] \quad \text{and} \quad \langle \hat{M}_y(t) \rangle = m\mathrm{tr}\left[\hat{\rho}(t)\sum_{i=1}^{n}\hat{\sigma}_y^{(i)}\right], \tag{3.36}$$

where m depends on the size of the ensemble and the magnetic moments of the nuclei. We combine the magnetization measured along each axis to create a complex value that represents the magnetization in the xy-plane, $\hat{M} = \hat{M}_x + i\hat{M}_y$. We define $\hat{\sigma}_- := (\hat{\sigma}_x + i\hat{\sigma}_y)/2 = |0\rangle\langle 1|$ and $\hat{\sigma}_+ := (\hat{\sigma}_x - i\hat{\sigma}_y)/2 = |1\rangle\langle 0|$ in the computational basis. The symbols $\hat{\sigma}_\mp$ are chosen in this way to be consistent with the notation in other chapters. Thus, the expectation value of the magnetization in the xy-plane is

$$\langle \hat{M}(t) \rangle = 2m\mathrm{tr}\left[\hat{\rho}(t)\sum_{i=1}^{n}\hat{\sigma}_-^{(i)}\right]. \tag{3.37}$$

Measurements in NMR are **weak ensemble measurements**. Usually, in quantum mechanics, measurements affect the state of the system. Here, the measurements are very weak. This means that the interaction of the nuclei with the macroscopic coils is very small. The current induced from one nucleus gives little information about the state of that nucleus. Thus we treated the field as a classical precessing magnetization.

First, we will evaluate how the magnetization's expectation value, for two qubits, changes in time using the two-qubit Hamiltonian from Eq. (3.26). We will ignore the J-coupling term for the moment and reintroduce it later. Doing so, the magnetization evaluates to

$$\begin{aligned}\left\langle \hat{M}(t) \right\rangle &= 2m\mathrm{tr}\left[\hat{\rho}(t)(\hat{\sigma}_-^{(1)} + \hat{\sigma}_-^{(2)})\right] \\ &= 2m\mathrm{tr}\left[\hat{\rho}(0)e^{i\hat{H}t/\hbar}(\hat{\sigma}_-^{(1)} + \hat{\sigma}_-^{(2)})e^{-i\hat{H}t/\hbar}\right] \qquad (3.38) \\ &= 2m\left(e^{-i\omega_{L1}t}\mathrm{tr}\left[\hat{\rho}(0)(\hat{\sigma}_-^{(1)})\right] + e^{-i\omega_{L2}t}\mathrm{tr}\left[\hat{\rho}(0)(\hat{\sigma}_-^{(2)})\right]\right). \qquad (3.39)\end{aligned}$$

Thus we have a magnetization consisting of two time-dependent signals with different frequencies.

The current produced by this magnetic moment is

$$I = \eta\frac{\mathrm{d}\langle \hat{M}(t) \rangle}{\mathrm{d}t} = 2\eta m\left(-i\omega_{L1}e^{-i\omega_{L1}t}\mathrm{tr}\left[\hat{\rho}(0)(\hat{\sigma}_-^{(1)})\right] - i\omega_{L2}e^{-i\omega_{L2}t}\mathrm{tr}\left[\hat{\rho}(0)(\hat{\sigma}_-^{(2)})\right]\right), \tag{3.40}$$

where η is a factor that depends on the detection efficiency of the coil. The terms in the trace contain some information about the state $\hat{\rho}(0)$. The time-varying signal from the

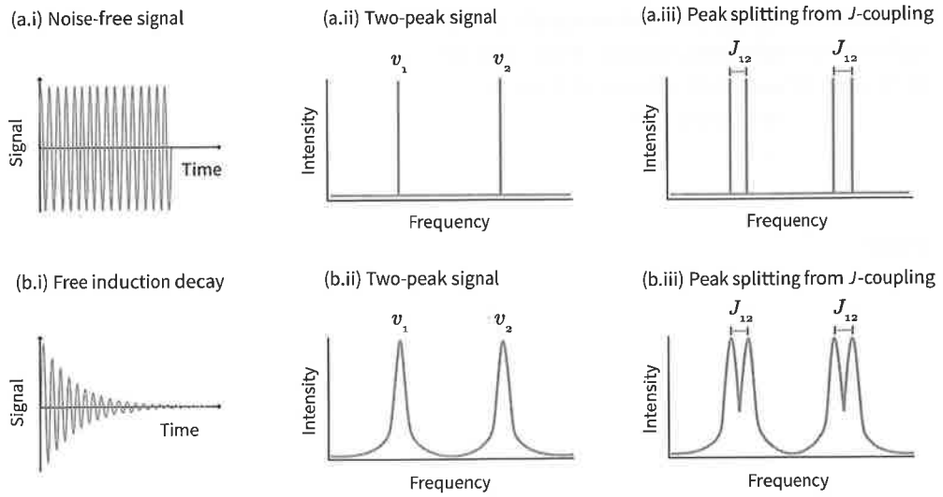

Fig. 3.10 **NMR signal and the free-induction decay.** In an NMR experiment, the signal is displayed in terms of frequency, not angular frequency. We denote the frequency of the ith nucleus as $\nu_i = \frac{\omega_i}{2\pi}$. (a) The ideal signal is measured without noise. (b) A realistic signal that includes noise. (a.ii) The Fourier transform of the signal is taken to see the frequencies present in the signal. This spectrum shows two noninteracting spin-$\frac{1}{2}$. The first/second peak shows the components of the first/second qubit that's in the state $\hat{\rho} = \hat{\sigma}_x^{(1/2)}$. (a.iii) The spectrum of an NMR signal with two interacting spin-$\frac{1}{2}$. The peaks are split by the J-coupling. The two peaks on either side of ω_{L1} and ω_{L2} correspond to whether the other qubit is in the state $|0\rangle$ or $|1\rangle$. (b.i) The signal measured from the nuclei is in the shape of a decaying sinusoid, called the free-induction decay (FID). (b.ii) and (b.iii) are analogous plots as (a.ii) and (a.iii) but in the presence of noise. The Fourier transform of the signal is a Lorentzian in this case instead of a Dirac delta function.

magnetization will be the combination of one sinusoid with frequency ω_{L1} and another with ω_{L2}.

To extract information from the current, we take its Fourier transform:

$$F(I) = C_1\delta(\omega - \omega_{L1}) + C_2\delta(\omega - \omega_{L2}), \tag{3.41}$$

where $C_j \coloneqq -i\sqrt{2\pi}\,\eta m\omega_{Lj}\,\mathrm{tr}[\hat{\rho}(0)(\hat{\sigma}_-^{(j)})]$ and $\delta(x)$ is the Dirac delta function. The Fourier transform gives us a signal with two peaks, located at the frequencies ω_{L1} and ω_{L2}. These peaks are depicted in Fig. 3.10. The location of the peaks tells us what qubit they correspond to. A sample of n noninteracting qubits with different frequencies will have n distinct peaks.

The real part of the Fourier transform gives us the overlap of the density matrix $\hat{\rho}(0)$ with $\hat{\sigma}_y$ and the imaginary part, its overlap with $\hat{\sigma}_x$. By measuring these strengths of the delta function (C_j) we can determine four components of the density matrix, $\hat{\sigma}_{x,y}^{(1,2)}$.

We'll now repeat the calculation we did above, but with the J-coupling included. We first split the magnetization into the component from the first and second qubit:

$$I = \eta m\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{tr}\left[\hat{\rho}(t)(\hat{\sigma}_-^{(1)} + \hat{\sigma}_-^{(2)})\right] = \eta\frac{\mathrm{d}\left\langle\hat{M}_1(t)\right\rangle}{\mathrm{d}t} + \eta\frac{\mathrm{d}\left\langle\hat{M}_2(t)\right\rangle}{\mathrm{d}t}. \tag{3.42}$$

We denote by $P_0 \coloneqq \left[\begin{smallmatrix}1&0\\0&0\end{smallmatrix}\right]$ and $P_1 \coloneqq \left[\begin{smallmatrix}0&0\\0&1\end{smallmatrix}\right]$ the projectors onto $|0\rangle$ and $|1\rangle$. Next, we look at the magnetization of the first qubit,

$$\left\langle\hat{M}_1(t)\right\rangle = 2m\mathrm{tr}\left[\hat{\rho}(t)\hat{\sigma}_-^{(1)}\otimes\hat{\mathbb{1}}^{(2)}\right] = 2m\mathrm{tr}\left[\hat{\rho}(t)\left(\hat{\sigma}_-^{(1)}\otimes(P_0^{(2)} + P_1^{(2)})\right)\right]. \tag{3.43}$$

As we have seen, the effect of the J-coupling is to increase or decrease the Larmor precession frequency by $\frac{\pi J_{12}}{2}$ depending on the state of the second qubit. Thus, it's not surprising that evaluating this expression again with the J-coupling we find

$$\left\langle\hat{M}_1(t)\right\rangle = 2me^{-i(\omega_{L1}-\pi J_{12}/2)t}\mathrm{tr}\left[\hat{\rho}(0)(\hat{\sigma}_-^{(1)}\otimes P_0^{(2)})\right] \tag{3.44}$$

$$+ 2me^{-i(\omega_{L1}+\pi J_{12}/2)t}\mathrm{tr}\left[\hat{\rho}(0)(\hat{\sigma}_-^{(1)}\otimes P_1^{(2)})\right]. \tag{3.45}$$

The J-coupling has split the peak at frequency ω_{L1} into two peaks of frequencies $\omega_{L1} \pm \frac{\pi J_{12}}{2}$. Thus, we now have four peaks. Two corresponding to each qubit. The splits in the peaks can be seen in Fig. 3.10(a.iii). If we instead looked at the expectation value of $\hat{\sigma}_-^{(1)} \otimes \hat{\sigma}_z^{(2)}$ in Eq. (3.43), we would have obtained the second term in Eq. (3.45) with a minus sign.

Before the peak splitting, we could measure four components of the density matrix. The real component of the signal from the first qubit gives us $\hat{\sigma}_x^{(1)} \otimes \hat{\mathbb{1}}^{(2)}$. The imaginary component of the signal from the first qubit gives us $\hat{\sigma}_y^{(1)} \otimes \hat{\mathbb{1}}^{(2)}$ and a similar relation holds for the second qubit. When the J-coupling is present, the peaks split, and we can measure eight components of the density matrix. The real component of the signal from the first qubit gives us $\hat{\sigma}_x^{(1)} \otimes \hat{\mathbb{1}}^{(2)}$ and $\hat{\sigma}_x^{(1)} \otimes \hat{\sigma}_z^{(2)}$. The imaginary component of the signal from the first qubit gives us $\hat{\sigma}_y^{(1)} \otimes \hat{\mathbb{1}}^{(2)}$ and $\hat{\sigma}_y^{(1)} \otimes \hat{\sigma}_z^{(2)}$. Similar relations hold for the second qubit.

Consider the spectra for the first qubit. The spectra for each one of the four components have a unique shape, given in Fig. 3.11. These four shapes form a basis for any possible spectra which we could measure in this scenario. For a generic density matrix, the NMR spectrometer returns some signal which is a linear combination of these different peak shapes. The spectra are then run through software that fits the shape of the spectra to a linear

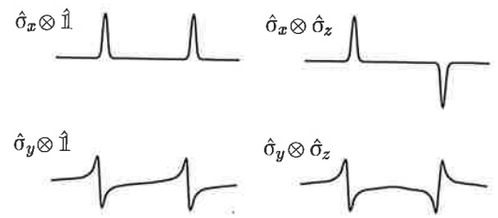

Fig. 3.11 **NMR spectra for a two-qubit sample.** The NMR spectra for the four components of the density matrix that are measured for a two-qubit sample: $\hat{\sigma}_x \otimes \hat{\mathbb{1}}$, $\hat{\sigma}_x \otimes \hat{\sigma}_z$, $\hat{\sigma}_y \otimes \hat{\mathbb{1}}$ and $\hat{\sigma}_y \otimes \hat{\sigma}_z$. These four shapes provide a basis for all two-qubit NMR spectra.

combination of the four shapes in Fig. 3.11, thus giving us four of the components of the density matrix.

Suppose we have an n-qubit sample, and the kth qubit is being observed. In that case, any component where the kth Pauli matrix is $\hat{\sigma}_x$ or $\hat{\sigma}_y$ and all others are $\hat{\mathbb{1}}$ or $\hat{\sigma}_z$ will be observed, so long as there are J-couplings between the kth qubits and the qubits who have $\hat{\sigma}_z$.

In Eq. (3.38) we assumed that the evolution was noiseless. In reality, this is not the case. The noise introduces a decaying function of the form e^{-t/T_2} in Eq. (3.38). This is due to decoherence, and we will discuss this effect in Section 3.7. The resulting curve of decaying current is called the **free-induction decay** (FID). Through the Fourier transform, the exponential decay will turn the delta functions into a Lorenztian curve in Eq. (3.41). The spectra with noise are shown in Fig. 3.10b.

3.5.2 Tomography

In the last subsection we found that by calculating the Fourier transform of the FID we can determine some of the components of the density matrix for each qubit. In the two-qubit example, these components were $\hat{\sigma}_x^{(1)} \otimes \hat{\mathbb{1}}^{(2)}$, $\hat{\sigma}_x^{(1)} \otimes \hat{\sigma}_z^{(2)}$, $\hat{\sigma}_y^{(1)} \otimes \hat{\mathbb{1}}^{(2)}$, and $\hat{\sigma}_y^{(1)} \otimes \hat{\sigma}_z^{(2)}$ for the first qubit. Similar components were found when measuring the second qubit. However, this isn't sufficient to determine the full two-qubit density matrix. To do so, more work needs to be done, and this process is called **state tomography**.

We'll illustrate how tomography is done using a one-qubit example, but this procedure generalizes to multiple qubits. We know from Eq. (2.63) that the density matrix for a single qubit can be written in terms of the Pauli matrices as

$$\hat{\rho} = \frac{1}{2}\left(\hat{\mathbb{1}} + \vec{n}\cdot\vec{\sigma}\right), \tag{3.46}$$

where $\vec{n} = [n_1, n_2, n_3]$ is a vector of length ≤ 1. From the result in Section 3.5.1, we find that a single NMR experiment will measure n_1 and n_2, the $\hat{\sigma}_x$ and $\hat{\sigma}_y$ components of $\hat{\rho}$. To know the state of the system we also need to measure n_3. This can be achieved by running the experiment again, but right before we measure we would apply a $\hat{R}_y(\frac{\pi}{2})$ rotation,

$$\hat{\rho}' = \hat{R}_y\left(\frac{\pi}{2}\right)\hat{\rho}\hat{R}_y\left(\frac{\pi}{2}\right)^\dagger = \frac{1}{2}\left(\hat{\mathbb{1}} + n_3\hat{\sigma}_x + n_2\hat{\sigma}_y - n_1\hat{\sigma}_z\right). \tag{3.47}$$

This will rotate the magnetization along the z-axis to point in the x-axis, and thus be measurable. Measuring $\hat{\sigma}_x$ for the final state now gives us the z-component (n_3) of the original density matrix. In this way, each component of the density matrix can be measured.

For two or more qubits, NMR measurements will determine more components of the density matrix. The density matrix for two qubits can be written in terms of tensor products of each of the two-qubit Pauli matrices (where we include the unit matrix here), giving rise to 16 components. Performing the full tomography for a two-qubit experiment would require rotating different components of the magnetization into the measurable observable until all $16 - 1$ components of the density matrix are measured. The missing one comes from the fact that the identity operator is not observable in NMR. For multi-qubit systems,

the rotations required to transform different components of the magnetization into the measurable observable will need both one- and two-qubit rotations.

3.6 Initialization

We have now explored what the NMR qubits are and how they can be controlled and measured. We now have all the elements necessary to describe how the state is initialized. The goal is to prepare a fiducial state such as $|0\rangle^{\otimes n}$ or an equivalent one. To do this, we begin with a thermal state (Section 3.6.1). We then perform a nonunitary action on this state to prepare a state that is equivalent to the fiducial state, known as the pseudo-pure state (Section 3.6.2).

3.6.1 Thermal State

If we leave an NMR sample idling for a few seconds, it will end up in a state known as the **thermal state**. It is given by

$$\hat{\rho}_{\text{th}} = \frac{e^{-\hat{H}/k_BT}}{\text{tr}\left[e^{-\hat{H}/k_BT}\right]}, \tag{3.48}$$

where $k_B = 1.38 \times 10^{-23}\text{m}^2\text{kgs}^{-2}\text{K}^{-1}$ is the Boltzmann constant. At room temperature, the matrix elements of $\hat{H}/k_BT$ are small, $\mathcal{O}\left(\frac{\hbar\omega_{Li}}{k_BT}\right) \approx 10^{-5}$. By taking the series expansion of $\hat{\rho}_{\text{th}}$, we find that

$$\hat{\rho}_{\text{th}} = \frac{1}{2^n}\left[\hat{\mathbb{1}} - \frac{\hat{H}}{k_BT} + \mathcal{O}\left(\frac{\hat{H}^2}{k_B^2T^2}\right)\right]. \tag{3.49}$$

This thermal state is a small deviation away from being the maximally mixed state, $\hat{\mathbb{1}}$.

We can also neglect the effect of the J-couplings since their values are around 100 Hz and the Larmor frequencies are many hundreds of MHz. Thus we approximate the thermal state to be

$$\hat{\rho}_{\text{th}} \approx \frac{1}{2^n}\left[\hat{\mathbb{1}} + \sum_i \frac{\hbar\omega_{L_i}}{2k_BT}\hat{\sigma}_z^{(i)}\right]. \tag{3.50}$$

For the following discussion, we will write this expression as an equality.

To satisfy the DiVincenzo criteria of being able to initialize the system, we need a procedure to transform the thermal state in NMR into a fiducial state to start a quantum computation. This is typically the $|0\rangle^{\otimes n}$ state. From Eq. (3.50), we see that to achieve this we would require that $\hbar\gamma B_0/k_BT \gg 1$. This, in turn, would require either a temperature that is on the order of 10 mK or below or a background magnetic field which is 10^5 times stronger than what we've considered. At that low temperature, the sample would not be liquid anymore. Also, a magnetic field of that strength is far beyond what today's technology can achieve. Instead, we prepare a different fiducial state to start a quantum computation.

3.6.2 Pseudo-Pure State

Clearly, the thermal state is very far from a pure state. However, we can still use it to perform interesting computations. Using NMR tools, we can transform it into a state of the form

$$\hat{\rho}_{\text{PPS}} = \frac{1-\epsilon_n}{2^n}\hat{\mathbb{1}} + \epsilon_n |0\rangle\langle 0|^{\otimes n}. \tag{3.51}$$

The states $\epsilon_n |0\rangle\langle 0|^{\otimes n}$ are called **pseudo-pure states** (PPS) and for all practical purposes behave like the usual initial state required for quantum computation. This is based on three key observations.

The first observation is that the observable in NMR, the magnetization $\hat{M}$, is a traceless operator. Thus, it does not depend on the part of the density matrix proportional to the identity operator $\hat{\mathbb{1}}$. If we define a **deviation matrix** $\hat{\delta}$ such that $\hat{\delta} - \hat{\rho} = \lambda\hat{\mathbb{1}}$ for some λ, then we find

$$\text{tr}[\hat{\delta}\hat{M}] = \text{tr}[(\hat{\rho} + \lambda\hat{\mathbb{1}})\hat{M}] = \text{tr}[\hat{\rho}\hat{M}] + \lambda\text{tr}[\hat{M}] = \text{tr}[\hat{\rho}\hat{M}]. \tag{3.52}$$

The state of a multi-qubit system can be divided into two parts, the identity operator for the whole system, $\hat{\mathbb{1}}$, and the rest represented by a sum of Pauli matrices. The second observation is that unitary evolutions transform the identity operator into itself,

$$\hat{U}\hat{\mathbb{1}}\hat{U}^\dagger = \hat{\mathbb{1}}, \tag{3.53}$$

and Pauli matrices into themselves,

$$\hat{U}\left(\sum_i \alpha_i \hat{P}_i\right)\hat{U}^\dagger = \sum_j \beta_j \hat{P}_j, \tag{3.54}$$

where $\hat{P}_i$ are Pauli matrices and α_i and β_j are constants.

The third observation is that scales are relative, i.e. that the probability that a final computation is in a specific state, let's say $|1\rangle^{\otimes n}$, can be written as a ratio of two magnetizations. The absolute value of the magnetization is not important, it only needs to be large enough to be observable over the noise. Let's see how the third observation works for one qubit. The thermal state has a deviation,

$$\hat{\delta} := \hat{\rho} - \frac{1-\epsilon}{2}\hat{\mathbb{1}} = \epsilon |0\rangle\langle 0|, \tag{3.55}$$

where $\epsilon = \frac{\hbar\omega_L}{2k_BT}$. We can first measure $a = \text{tr}[\hat{\delta}\hat{\sigma}_z]$. Say we want to find the probability of an initial state $|0\rangle$ to be in the state $|1\rangle$ after an evolution $\hat{U}$. To do so, we would evaluate

$$\begin{aligned} p_1 &= \langle 1|\hat{U}|0\rangle\langle 0|\hat{U}^\dagger|1\rangle \\ &= \text{tr}[\hat{U}|0\rangle\langle 0|\hat{U}^\dagger|1\rangle\langle 1|] \\ &= \tfrac{1}{2}\text{tr}[\hat{U}|0\rangle\langle 0|\hat{U}^\dagger(\hat{\mathbb{1}} - \hat{\sigma}_z)] \\ &= \tfrac{1}{2}(1 - \text{tr}[\hat{U}|0\rangle\langle 0|\hat{U}^\dagger\hat{\sigma}_z]). \end{aligned} \tag{3.56}$$

If we call $a' = \text{tr}[\hat{\delta}'\hat{\sigma}_z] = \epsilon\text{tr}[\hat{U}|0\rangle\langle 0|\hat{U}^\dagger\hat{\sigma}_z)]$, the probability to find, let's say, the system to be in the state $|1\rangle$ is $p_1 = (1 - a'/a)/2$. You will notice that p_1 is independent of ϵ. It is also independent of m and η that we encountered in Section 3.5.1.

To generalize this argument to n qubits, we need to be able to prepare the states in Eq. (3.51). For $n \geq 2$ systems, the preparation of a PPS state requires the use of a combination of unitary and nonunitary transformations.

One technique for preparing a PPS is **temporal averaging**. This approach consists of running multiple experiments and then later averaging the results of these experiments. The final outcomes of all the experiments are equivalent to doing one experiment with a PPS. We illustrate this procedure below for two qubits. The goal is to find the magnetization for a pseudo-pure state $\epsilon_2 |00\rangle\langle 00|$ that has evolved under $\hat{U}_E$. We will find it by adding the results of three experiments starting with slightly different initial states.

We begin with the thermal state $\hat{\rho}_{\text{th}}$, where we have labelled its diagonal elements in the following way:

$$\hat{\rho}_{\text{th}} = \begin{bmatrix} a & 0 & 0 & 0 \\ 0 & b & 0 & 0 \\ 0 & 0 & c & 0 \\ 0 & 0 & 0 & d \end{bmatrix} = \hat{\rho}_1. \tag{3.57}$$

$\hat{\rho}_1$ is the initial state for the first run of the experiment. The final state after the evolution is then $\hat{U}_E\hat{\rho}_1\hat{U}_E^\dagger$.

For the second experiment, we want to permute the diagonal elements b, c, d to d, b, c. This can be achieved by first applying a CNOT_{21} gate and then a CNOT_{12} gate to obtain:

$$\hat{\rho}_2 = (\text{CNOT}_{12}\text{CNOT}_{21})\,\hat{\rho}_{\text{th}}\,(\text{CNOT}_{12}\text{CNOT}_{21})^\dagger = \begin{bmatrix} a & 0 & 0 & 0 \\ 0 & d & 0 & 0 \\ 0 & 0 & b & 0 \\ 0 & 0 & 0 & c \end{bmatrix}. \tag{3.58}$$

We then conduct the experiment and get the final state $\hat{U}_E\hat{\rho}_2\hat{U}_E^\dagger$.

For the initial state in the third run of the experiment, we want to permute b, c, d to c, d, b. This is achieved by applying a CNOT_{12} gate and then a CNOT_{21} gate to obtain:

$$\hat{\rho}_3 = (\text{CNOT}_{21}\text{CNOT}_{12})\,\hat{\rho}_{\text{th}}\,(\text{CNOT}_{21}\text{CNOT}_{12})^\dagger = \begin{bmatrix} a & 0 & 0 & 0 \\ 0 & c & 0 & 0 \\ 0 & 0 & d & 0 \\ 0 & 0 & 0 & b \end{bmatrix}. \tag{3.59}$$

We then conduct the experiment and get the final state $\hat{U}_E\hat{\rho}_3\hat{U}_E^\dagger$. The average magnetization over the three results is

$$\langle\hat{M}\rangle = \frac{1}{3}\sum_{i=1}^{3}\text{tr}\left[\hat{M}(\hat{U}_E\hat{\rho}_i\hat{U}_E^\dagger)\right] = \text{tr}\left[\hat{M}\hat{U}_E(\frac{1}{3}\sum_{i=1}^{3}\hat{\rho}_i)\hat{U}_E^\dagger\right], \tag{3.60}$$

where

$$\frac{1}{3}\sum_{i=1}^{3}\hat{\rho}_i = \frac{1}{3}\begin{bmatrix} 3a & 0 & 0 & 0 \\ 0 & b+c+d & 0 & 0 \\ 0 & 0 & b+c+d & 0 \\ 0 & 0 & 0 & b+c+d \end{bmatrix}$$
$$= \frac{(1-a)}{3}\hat{\mathbb{1}} + \frac{(4a-1)}{3}|00\rangle\langle 00|$$
$$= \frac{(1-\epsilon_2)}{4}\hat{\mathbb{1}} + \epsilon_2 |00\rangle\langle 00|, \tag{3.61}$$

where we used the property that $\hat{\rho}_{\text{th}}$ has trace 1 to substitute $b+c+d = 1-a$ and set $\epsilon_2 = \frac{(4a-1)}{3}$. Then,

$$\left\langle \hat{M} \right\rangle = \epsilon_2 \text{tr}[\hat{M}\hat{U}_E |00\rangle\langle 00| \hat{U}_E^\dagger]. \tag{3.62}$$

Thus, the average of the three signals is equivalent to having run one experiment with the pseudo-pure state $\epsilon_2 |00\rangle\langle 00|$. This is temporal averaging.

Another common technique is that of spatial averaging, where the PPS is prepared in a single experiment. This involves the use of magnetic field gradients which we do not cover in this text. Temporal and spatial averaging methods for preparing a pure state are useful for up to a dozen qubits, but they are not scalable. One scalable method called algorithmic cooling will be explored in the exercises.

3.7 Noise

In the previous section we studied the signal we measured from the NMR sample. We said that this signal gradually decays (Fig. 3.10a). The cause of this decay is noise, which we'll explore in this section. First, we introduce the concepts of the T_1 and T_2 times (Section 3.7.1). These times give a way of thinking about decoherence that originated in NMR and is now present in most other implementations. We will also use the T_1 and T_2 times to describe decoherence in Chapters 4–6. Section 3.7.2 identifies the noise sources in NMR that contribute to the T_1 and T_2 times.

3.7.1 Introduction to T1 and T2

We can partially characterize the quality of a quantum computer by its T_1 and T_2 times. T_1 and T_2 are decay constants that measure different things. T_1 is related to the time a state takes to return to its thermal equilibrium. T_2 is related to the time taken for the phase of a superposition state to be randomized. The T_1 time measures the loss of energy, and the T_2 time measures the loss of coherence from the system. T_1 and T_2 are decay constants that tell us the probability of certain types of errors occurring over time.

The T_1 time goes by different names, including; relaxation time, energy relaxation time, longitudinal relaxation time, spontaneous emission time, or amplitude damping time.

The T_2 time has just as many different names, such as the coherence time, dephasing time, spin–spin relaxation time, or transversal relaxation time. In this text, we will primarily use the terms **relaxation time** and **dephasing time** in conjunction with T_1 time and T_2 time.

3.7.2 Noise in NMR

Since we measure the magnetization to determine the state of the qubits, we need to consider the effect of noise on the sample's magnetization. For pedagogical clarity, we'll limit ourselves here to a single-qubit sample. The procedures for determining the relaxation rates are more involved when we have more qubits.

The **Bloch equations** are equations used to calculate the magnetization as a function of time in the presence of noise. For a single qubit, these equations are

$$\frac{\mathrm{d}\left\langle \hat{M}_z \right\rangle}{\mathrm{d}t} = \gamma\left(\left\langle \hat{M} \right\rangle \times \vec{B}\right)_z - \frac{\left\langle \hat{M}_z \right\rangle - \left\langle \hat{M}_{\text{th}} \right\rangle}{T_1}, \tag{3.63}$$

$$\frac{\mathrm{d}\left\langle \hat{M}_{x,y} \right\rangle}{\mathrm{d}t} = \gamma\left(\left\langle \hat{M} \right\rangle \times \vec{B}\right)_{x,y} - \frac{\left\langle \hat{M}_{x,y} \right\rangle}{T_2}, \tag{3.64}$$

where $\hat{M}_{\text{th}}$ is the thermal state's polarization. In both equations, the first term on the right-hand side of the equation recovers the one-qubit rotation we studied in Section 3.3. The second term describes the impact of noise on nuclear spins.

The T_1 relaxation causes the bulk magnetization to realign with the z-axis. This results in losing the detectable signal since measurements are done in the xy-plane. In NMR, the value of T_1 corresponds to the rate of energy exchange between the nuclei and the environment as the nuclei relax back to their thermal state. A single nucleus can emit one quantum of energy into the environment or absorb one quantum from the environment. The emission rate is determined by thermal noise.

The **inversion-recovery method** is commonly used to measure the T_1 time. It starts with a thermal state and the application of a $\hat{R}_x(\pi)$ pulse to invert the polarization. After applying this pulse, the magnetization is pointing entirely in the $-z$-direction. This is shown in Fig. 3.12. We then wait for some time t and then measure the magnetization. We repeat this experiment with different values of t. By doing so, we can find the rate at which the system returns to the thermal state. The solution of the Bloch equation after the $\hat{R}_x(\pi)$ pulse is

$$\left\langle \hat{M}_z \right\rangle = \left\langle \hat{M}_{\text{th}} \right\rangle \left(1 - 2e^{-\frac{t}{T_1}}\right). \tag{3.65}$$

By fitting the measured data to this equation, we can extract T_1.

The T_2 relaxation results from nuclei from the different molecules precessing at slightly different frequencies and no longer being in sync. It follows that the magnetizations of the nuclei are not adding constructively, decreasing the bulk magnetization. This occurs due to inhomogeneities of the magnetic field (in space and time). In theory, the overall magnetic field is uniform with high accuracy, but will, in practice, also include some deviations. These deviations can result from imperfections in the superconducting magnet or the presence of other magnetic fields. These other magnetic fields include those created by the sample itself.

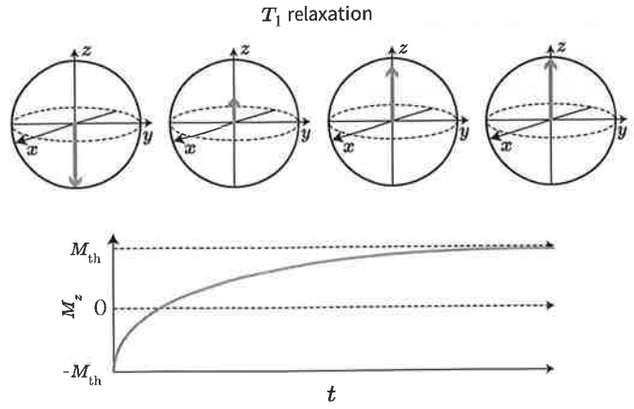

Fig. 3.12 **Inversion-recovery method for determining T_1.** After a $\hat{R}_x(\pi)$ pulse, the magnetization of the sample will realign with the external magnetic field on a timescale of T_1. As this realignment occurs, the magnetization along the z-axis gradually returns to its equilibrium value.

Since the sample's nuclei have a magnetic moment and are constantly in motion, they will disturb the magnetic environment. Thus the magnetic field will vary slightly at different locations and at different moments in time. As a result, the nuclei's Larmor frequencies will vary. Over time these differences become notable as the spins will fall out of sync with one another, and the magnetization in the xy-plane will be lost (Fig. 3.13). The direction of the magnetization in the xy-plane of the nuclei (for the same qubit) in different molecules will differ, leading to different phases in the qubit's state. This is why this form of relaxation is also called dephasing.

There are two relevant constants for dephasing. The constant related to the dephasing we observe is known as T_2^* (pronounced "T_2 star"). This is sometimes what people refer to when they say "dephasing time." However, part of the signal lost to dephasing can be regained using the refocusing technique we introduced in Section 3.4.3. This is depicted in Fig. 3.8b. The other part of the signal loss comes from random fluctuations and thus can't be regained. This more fundamental and limiting dephasing relates to the T_2 time. If we have a perfectly uniform and constant magnetic field, T_2 would equal T_2^*. The T_2 and T_2^* are related by

$$\frac{1}{T_2} = \frac{1}{T_2^*} - |\gamma \Delta B|, \tag{3.66}$$

where ΔB is the field inhomogeneity. As a word of caution, many references use the term T_2 time in place of T_2^* time.

A common way to measure the T_2^* is with a **Ramsey experiment**. We begin with the system in the thermal state and apply a $\frac{\pi}{2}$-pulse along any axis in the xy-plane. This will rotate the sample's magnetization into the xy-plane. We then continuously measure the magnetization. By doing so, we will see an exponential decay of the FID. In this case, the solution of the Bloch equation is

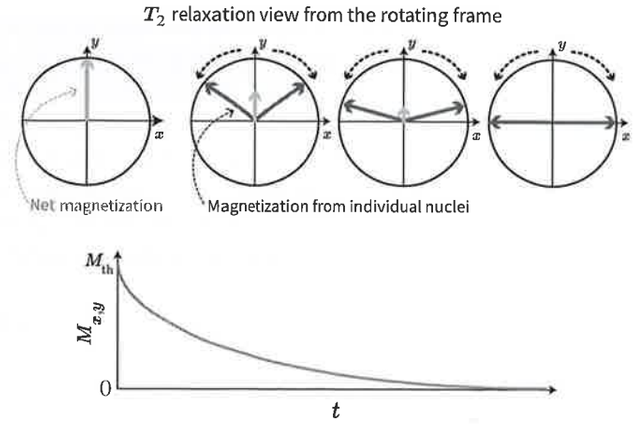

Fig. 3.13 **NMR signal with dephasing.** Four snapshots at increasing time are shown. The lighter grey arrow is the average magnetization of many individual spins' magnetization, which are represented by the darker grey arrows. Inhomogeneities in the magnetic field will cause the spins of the individual nuclei to precess at different rates. All nuclei's magnetizations start in the same direction, but as the individual nuclei gradually fall out of phase with one another the net magnetic moment will decrease and the magnetization in the xy-plane is lost with time.

$$\left\langle \hat{M}_{x,y} \right\rangle = \left\langle \hat{M}_{\text{th}} \right\rangle e^{-\frac{t}{T_2^*}}. \tag{3.67}$$

By fitting the measured data to this equation, we can extract T_2^*. Here we have measured the observed dephasing.

The values of T_1 and T_2 are not entirely independent. T_2 is always smaller than $2T_1$ because relaxation induces some dephasing,

$$T_2 \leq 2T_1. \tag{3.68}$$

However, typically, $T_2 \ll T_1$.

We have described the noise using two parameters, T_1 and T_2. Their effect on the density matrix can be evaluated using the Kraus operators introduced in Chapter 2. However, a detailed analysis of the number of parameters to describe the evolution (e.g. a gate) of a quantum state of one qubit would tell us that 12 parameters are needed. Out of these 12 parameters, 3 correspond to the unitary part of the evolution and 9 to the nonunitary, or the noisy, part. An astute reader might well ask why we have described only two of them here. The answer is that in NMR and many other systems we have a symmetry around the z-axis of the Bloch sphere (or the axis defined by the eigenstates of the Hamiltonian). This symmetry reduces the number of noisy parameters to only two. For a multi-qubit system, every qubit has its own T_1 and T_2, but can also have some correlations between these. As the system grows, the number of parameters to describe the noise grows exponentially.

In NMR experiments, single-qubit gates are implemented on the order of microseconds, and two-qubit gates are implemented on the order of milliseconds. The main limit to the length of the experiment is the number of two-qubit gates. Experimentally T_1 times are on the order of a few seconds and T_2 times are on the order of a second. Thus, quantum

computations can be run before significant coherence is lost. As a result of their lengths, the T_2 time creates the upper bound on the length of the experiment and the T_1 time provides the lower bound on the time needed to wait between experiments in order to reset the state.

The Larmor frequencies are of the order of hundreds of megahertz, so it takes about 1 ns (nanosecond) for a nuclear spin to go around the *xy*-plane. The spins will rotate about one billion times during a typical experiment. For NMR quantum information experiments where the phase information is crucial, it is pretty impressive that the clock of the spectrometer tracks these phases to a precision much better than one part in one billion.

It will also be useful to distinguish between technical and intrinsic sources of decoherence. The first kind is decoherence caused by noise in classical control parameters. The variation of the background magnetic field in the spectrometer or the power of the RF fields are examples of this kind of noise. Intrinsic sources of decoherence can come from interactions with elements of the environment that are necessarily present. In NMR, the predominant sources of noise are technical.

3.8 Conclusion

We conclude this chapter with a summary of how we can use the nuclear spins of a solution of molecules for quantum computing (Section 3.8.1), a discussion on the strengths and weaknesses of computing with NMR (Section 3.8.2), and a brief discussion on the role of entanglement in quantum computing (Section 3.8.3). For further reading on NMR quantum computing, see Laflamme et al. (2002) or Jones (2001), and on the physics of NMR, see Levitt (2013).

3.8.1 Summary

In liquid-state NMR quantum computing, the qubits are encoded in the eigenstates of spin-$\frac{1}{2}$ nuclei in a magnetic field. The nuclei are part of a molecule. The number of nuclei in the molecule with distinct Larmor frequencies determines the maximum number of qubits in the computer. Single-qubit control is achieved through the application of RF pulses. The pulse applies a torque to the nuclei's spin, rotating the spin, and thus changing the qubit's state. Within the molecule, nuclei will interact via the J-coupling. This coupling can be regulated to implement two-qubit gates. Together these operations give us universal control. Before an experiment, an NMR sample will be in a thermal state. We perform a nonunitary action to this state to prepare it in a pseudo-pure state. To determine the qubits' final state, we measure the magnetization of the nuclei in the *xy*-plane. Since the signal from a single molecule is too weak to detect, we perform ensemble measurements of the many molecules in the sample. Finally, the main sources of decoherences in NMR are thermal relaxation (causing the signal to realign with the *z*-axis) and inhomogeneities of the magnetic field (causing the nuclei from different molecules to precess at different frequencies).

An overview of the contents of this chapter is provided in Table 3.1.

Table 3.1 A summary of how an NMR spectrometer satisfies each of the DiVincenzo criteria for quantum computation

Criteria	Realization
1. A scalable physical system with well characterized qubits.	Qubit: Nuclear spins within a molecule with different Larmor frequencies. Scalability: Limited by molecules and the pseudo-pure state.
2. The ability to initialize the	Pseudo-pure state. state of the qubits to a simple fiducial state.
3. A universal set of quantum gates.	One-qubit gate: Rabi oscillations with RF pulses. Two-qubit gate: using the J-coupling between nuclear spins.
4. A qubit-specific measurement capability.	Ensemble measurements.
5. Long relevant decoherence times, much longer than the gate operation time.	Decoherence time: seconds. One-qubit gate time: microseconds. Two-qubit gate time: milliseconds.

3.8.2 Relative Strengths and Weaknesses

NMR's technological maturity has endowed it with notable strengths for quantum computing research. The first is that commercial spectrometers are commonly available and have been used long before the advent of NMR quantum computing. Second, the many years of refinements have allowed for high-fidelity control and measurements. Many techniques for spin control have since transferred to other quantum technologies. For example, the use of composite and shaped pulses for improving gate fidelity originated with NMR. Today, trapped-ion and superconducting qubits also benefit from these techniques. NMR quantum computing's availability and high fidelity caused them to play a pivotal role in developing quantum computing as a field. When skepticism for the field of quantum computing was at a high, NMR quantum computers demonstrated many small-scale quantum algorithms. Such algorithms include three-qubit quantum error correction, the Deutsch–Jozsa algorithm, Grover's algorithm, and Shor's algorithm (for numbers with a few digits) and some quantum simulations. Aside from its maturity, NMR also benefits from encoding the qubit in spin-$\frac{1}{2}$ nuclei. By doing so, one obtains a well characterized two-level system. In NMR, leakage to higher states is nonexistent, unlike in trapped-ions or superconducting qubits.

NMR's key flaw is its scalability. This flaw has many causes. The first is a result of how qubits are addressed, which is by their Larmor precession frequencies. An n-qubit sample will require n precession frequencies that are sufficiently separated to allow for individual control. The limited bandwidth thus creates a limit to the size of the system. The second

C-13 labeled 12-qubit system

	C1	C2	C3	C4	C5	C6	C7	H1	H2	H3	H4	H5
C1	30020											
C2	57.58	8779										
C3	-2.00	32.70	6245									
C4	0	0.30	0	10333								
C5	-1.25	2.62	1.11	33.16	15745							
C6	5.54	-1.66	0	-3.53	33.16	34381						
C7	1.25	37.48	0.94	29.02	21.75	34.57	11928					
H1	0	0	2.36	166.6	4.06	5.39	8.61	3310				
H2	4.41	1.86	146.6	2.37	0	0	0	0	2468			
H3	1.81	3.71	146.6	2.37	0	0	0	0.18	-12.41	2158		
H4	-13.19	133.6	-6.97	6.23	0	5.39	3.78	-0.68	1.28	6.00	2692	
H5	7.87	-8.35	3.35	8.13	2.36	8.52	148.5	8.46	-1.06	-0.36	-1.30	3649
T_1	7.99	3.61	1.83	3.72	9.89	7.80	3.64	3.83	2.13	2.28	2.65	3.47
T_2^*	0.40	0.31	0.44	0.25	0.25	0.40	0.38	0.29	0.39	0.34	0.15	0.30

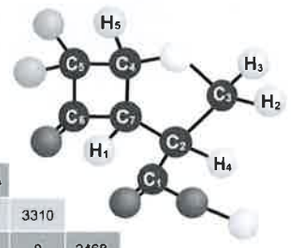

Fig. 3.14 **12-qubit sample for NMR quantum computing.** The molecular structure of a type of dichloro-cyclobutanone, specifically per-^{13}C labelled (1S,4S,5S)-7,7-dichloro-6-oxo-2-thiabicyclo[3.2.0]heptane-4-carboxylic acid. The seven carbon and five hydrogen nuclei can be used as qubits providing a 12-qubit quantum processor. The diagonal elements are the chemical shifts (Hz), the off-diagonal elements are the J-couplings between two spins (Hz), and the bottom rows contains the relaxation times T_1 and T_2^* (seconds).

limit to scalability is from the pseudo-pure state. The system's polarization characterizes the strength of the pseudo-pure state. It scales as 2^{-n}, where n is the number of qubits in the system. This means that for a large system, the pseudo-pure state is undetectable. The polarization could increase from cooling the system. However, the needed cooling would result in the sample solidifying. This would bring us to solid-state NMR, which brings a different set of challenges. The polarization could also increase using a more powerful magnet. But this quickly becomes technologically difficult. These are the major factors that limit scalability. An NMR sample with 12 qubits, such as the one included in Fig. 3.14, is one of the largest samples that has been used to date for NMR quantum computations.

NMR quantum computing has a few other challenges which are less pressing. Consider the last DiVincenzo criterion. In NMR, the decoherence time is on the order of 1–10 seconds. Yet, the time for two-qubit gates is proportional $\sim\frac{1}{2J}$, which for typical values of J results in gates of a few milliseconds. This ratio of decoherence to gate time limits the circuit depth achievable. Single-qubit gates are much shorter, on the order of microseconds. Second, in NMR quantum information processing, measurements are nonprojective and done only at the end of computations. This first feature necessitates algorithms requiring projective measurements to be reformulated. Measurements only occurring at the end of the computation, compounded with how qubits are initialized, means that the selective reset of qubits during computation is impossible. This means that performing quantum error

correction in NMR will require more qubits. A final criticism of NMR that has been brought up is the lack of entanglement when only a few qubits are used.

3.8.3 Role of Entanglement in Quantum Computation

We conclude with a comment on the role of entanglement in quantum computation. Entanglement is lacking in NMR quantum computing experiments with less than a dozen qubits. Sufficiently mixed states are separable and thus don't exhibit entanglement. The pseudo-pure state is a highly mixed state and is separable for $\epsilon < 2/2^{2n}$, where n is the number of qubits.[1] However, for polarizations above a certain threshold, $\epsilon > 2/\sqrt{2}^n$, some entanglement is possible.

In the early day of NMR quantum computing, a controversy arose from this lack of entanglement. At the time, entanglement was taken as synonymous with being quantum. The controversy arose when Artur Ekert and Richard Jozsa proved that a class of algorithms that act on a pure state and don't perform any entangling gates can be simulated efficiently by a classical computer. However, this result only holds for pure states.

Interestingly, an increasing amount of entanglement is not a sufficient condition for making the algorithm hard to simulate classically. The Gottesman–Knill theorem (Chapter 2) states that computations with Clifford gates alone can be efficiently simulated classically. However, Clifford gates can generate entanglement. Thus, entanglement alone is not sufficient for achieving a computational advantage.

Furthermore, an increasing amount of entanglement does not seem necessary for achieving a computational advantage. This is demonstrated by the DQC1 (**Deterministic Quantum Computation with 1 qubit**) model. DQC1 uses a single qubit that's prepared in a pseudo-pure and $n-1$ other qubits which are in completely mixed states. By using a controlled unitary gate, where the pseudo-pure state is controlling the other qubits, DQC1 can perform physics simulations that have no known efficient classical algorithms. A quantum circuit for such a DQC1 experiment is depicted in Fig. 3.15. DQC1 algorithms can

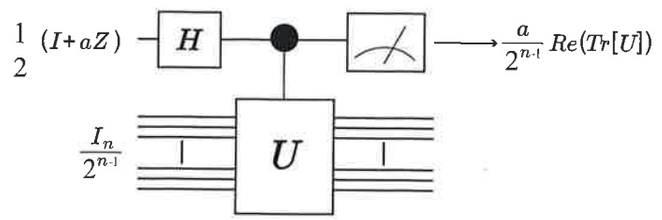

Fig. 3.15 **DQC1 algorithm.** It measures the trace of a unitary operator $\hat{U}$ that can be efficiently implemented in a quantum computer. It uses one qubit in a pseudo-pure, and $n-1$ qubits in the completely mixed state. By measuring the state of the first qubit, the trace of the operator can be found. This computation has no known efficient classical algorithm. Thus, it seems that even without much entanglement, quantum computers can perform efficient tasks that classical computers might not be able to carry out.

[1] For reference, ϵ for the thermal state, using an 11.7 tesla spectrometer, is on the order of 10^{-5} for hydrogen.

measure properties of quantum chaotic systems, knot invariants, partition functions, and other quantities that are believed to be hard classically.

3.9 Exercises

3.1 **Stored energy:** Estimate the energy stored in the magnetic field of an NMR magnet. Use the typical parameters of length = 10 cm, radius = 10 cm, number of windings = 10^4, and current = 100 A. For reference, a stick of dynamite contains about 1 MJ of chemical energy. The energy of an inductor is $E = LI^2/2$, where L is the inductance and I is the current running through the inductor. The inductance of a solenoid is $L = \mu_0 N^2 A_S / l$, where μ_0 is the magnetic constant, N is the number of turns in the coil, A_S is the cross section area of the coil, and l is the length of the coil.

3.2 **Eigenstates of the Hamiltonian:** Given the Hamiltonian

$$\hat{H} = -\frac{\hbar\omega_{L1}}{2}\hat{\sigma}_z^{(1)} - \frac{\hbar\omega_{L2}}{2}\hat{\sigma}_z^{(2)} + \frac{\hbar\pi J_{12}}{2}\left(\hat{\sigma}_x^{(1)}\hat{\sigma}_x^{(2)} + \hat{\sigma}_y^{(1)}\hat{\sigma}_y^{(2)} + \hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)}\right), \quad (3.69)$$

find the eigenstates and energy eigenvalues (also known as the **eigenenergies**) using the computational basis. Show that two eigenstates are separable and two are entangled states.

3.3 **Composite pulses:** Here we will explore how composite pulses can be used to mitigate the effect of small miscalibrations, denoted by ϵ. The following two pulse sequences are equivalent to $\hat{R}_y(\pi)$ if the error was zero, $\epsilon = 0$:

(a) $R_y(\frac{\pi}{2}+\epsilon)R_y(\frac{\pi}{2}+\epsilon)$; and
(b) $R_y(\frac{\pi}{2}+\epsilon)R_x(\pi+2\epsilon)R_y(\frac{\pi}{2}+\epsilon)$.

For an initial state $|0\rangle$, calculate the final state after pulse sequence a (denoted by $|\psi_a\rangle_f$), after pulse sequence b (denoted by $|\psi_b\rangle_f$), and after applying the ideal gate $R_y(\pi)$ (denoted by $|\psi\rangle_f$). One way to compare how close two states are is by taking the square of their inner product. This value is known as the **fidelity** and we study it in Chapter 7:

$$F_p(|\psi\rangle, |\phi\rangle) = |\langle\psi|\phi\rangle|^2. \quad (3.70)$$

Using your software of choice, plot the fidelity between $|\psi_a\rangle_f$ and $|\psi\rangle_f$ and plot the fidelity between $|\psi_b\rangle_f$ and $|\psi\rangle_f$. Overlap these plots to show that the second pulse sequence is more robust to miscalibration, i.e. that changes in ϵ result in smaller changes to the fidelity.

(c) *Challenge:* Show that this fails when the initial state is not $|0\rangle$.
(d) *Challenge:* For a general initial state $|\psi\rangle$, find a pulse sequence that works.

3.4 **On- and off-resonance effects:** For this exercise you can assume that the chemical shift $\delta_{CS} = 0$.

(a) What is the Larmor frequency for protons if $B_0 = 7\,\mathrm{T}$? If the length of a $\pi/2$ pulse on-resonance is $t = 1\,\mathrm{ms}$, compute the magnitude of the RF-field B_x.

(b) Now imagine that the pulse is detuned by 0.5 kHz away from the resonance frequency, but we have the same magnitude B_x and same pulse duration t. What is the effect of the pulse?
(c) Show that an off-resonance pulse, $|\omega_L - \omega_{rf}| \gg \omega_R$, will have no effect.

3.5 **RWA and the Bloch equations:** In this problem we will explore the validity of the rotating-wave approximation by looking at numerical simulations of the semi-classical Bloch equations. The Bloch equations are the equations of motion for the expectation value of the magnetization in an NMR setup. Ignoring decoherence, they are

$$\frac{d\vec{M}}{dt} = \gamma(\vec{M} \times \vec{B}), \quad (3.71)$$

where $\vec{M}$ is the magnetization of the sample and $\vec{B}$ is the applied magnetic field. For the Rabi problem we assume $\gamma B_z = \omega_0$, $\gamma B_x = 2\omega_R \cos(\omega t)$, and $\gamma B_y = 0$.

(a) Using the form of $\vec{B}$ above, write out the Bloch equations as a set of coupled differential equations for the components of $\vec{M}$. Using your favourite software package, write a simple program to numerically integrate these equations. *Note that many packages, such as MATLAB, have built-in equation solvers.*
(b) Taking $\omega_0 = \omega = 10\omega_R$ with the initial condition $M_z = M_0$, $M_x = M_y = 0$, integrate the Bloch equations long enough to see a few periods of oscillations in M_z. How does the numerical solution for M_z compare to the analytical solution derived in Section 3.3 using the RWA? Integrate the equations for a couple of larger values of ω_R. How do the solutions change? Can you explain the small oscillations that develop? (*Hint*: think of counter-rotating terms.) Investigate plots of $M_z(t)$ for at least three values of the ratio of ω_0/ω_R and comment on the behaviour.

3.6 **Conjugating Z-rotations:** Consider the following sequence of single-qubit rotations and free evolutions:

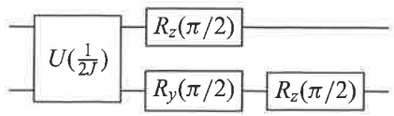

(3.72)

Using the necessary commutation relations, move the Z-rotations to the beginning of the circuit. In what way would this problem become more complex if the rotations were not rotations of integer multiples of $\pi/2$?

3.7 **Universal single-qubit control:** In Section 3.3.2 we found the effect of an RF pulse along the x-axis with a phase of $\phi = 0$. We found the Hamiltonian for this pulse in the rotating frame given by $\hat{R}_z(\omega_{rf}t)$ is

$$\hat{H}' = -\frac{\hbar(\omega_L - \omega_{rf})}{2}\hat{\sigma}_z - \frac{\hbar\omega_R}{2}\hat{\sigma}_x. \quad (3.73)$$

Prove that if the initial pulse was

$$\vec{B}(t) = B_0\vec{z} + B_{rf}\cos(\omega_{rf}t - \phi)\vec{x}, \quad (3.74)$$

then the final Hamiltonian in that given frame of reference would have been

$$\hat{H}' = -\frac{\hbar(\omega_L - \omega_{rf})}{2}\hat{\sigma}_z - \frac{\hbar\omega_R}{2}(\cos(\phi)\hat{\sigma}_x + \sin(\phi)\hat{\sigma}_y). \tag{3.75}$$

3.8 **Refocusing:** In the text we presented a refocusing scheme to simulate a process in which two qubits are not evolving.

(a) Find a refocusing scheme for three qubits that stops just the first qubit from evolving, i.e. one that cancels the effects of J_{12} and J_{13} but not J_{23}.

(b) Find a scheme that cancels all three couplings.

(c) How many pulses did your scheme require? How would your scheme's number of pulses scale with the number of qubits? If it scales exponentially, can you find a scheme that does not?

3.9 **Weak-coupling approximation:** A typical two–nuclear spin Hamiltonian is given by

$$\hat{H}_J = -\frac{\hbar\omega_{L1}}{2}\hat{\sigma}_z^{(1)} - \frac{\hbar\omega_{L2}}{2}\hat{\sigma}_z^{(2)} + \frac{\hbar\pi J_{12}}{2}\vec{\sigma}^{(1)} \cdot \vec{\sigma}^{(2)}. \tag{3.76}$$

Show that it can be approximated by

$$\hat{H}_J = -\frac{\hbar\omega_{L1}}{2}\hat{\sigma}_z^{(1)} - \frac{\hbar\omega_{L2}}{2}\hat{\sigma}_z^{(2)} + \frac{\hbar\pi J_{12}}{2}\sigma_z^{(1)}\sigma_z^{(2)}. \tag{3.77}$$

when in the **weak-coupling regime**, $J_{12} \ll |\omega_{L1} - \omega_{L2}|$. This approximation is known as the **secular approximation**.

3.10 **Generation of the B_{rf} field:** Explain how the wires making the coil in Fig. 3.9 produce a magnetic field in the x-direction at the midpoint of the coil. This geometry allows for the production of the field in the desired direction and makes it easier to insert and remove samples.

3.11 **FID and Fourier transform:** In NMR experiments we measure the free-induction decay (FID) and then perform a Fourier transform to get the NMR spectra.

(a) Denote the x- and y-components of the detected signal by $S_x(t) = S_0\cos(\omega_0 t)$ and $S_y(t) = S_0\sin(\omega_0 t)$ in the absence of decoherence. What is the function for the FID? What is its Fourier transform?

(b) In the presence of dephasing, $S_x(t)$ and $S_y(t)$ both have an extra term, e^{-t/T_2}. What is the function for the FID? What is its Fourier transform? Find the half-width of the NMR spectra (the width of the Lorentzian at half its maximum value).

3.12 **Measurements in NMR:** Consider a three-qubit experiment, where qubits 1 and 2 have coupling J_{12} and qubits 2 and 3 have coupling J_{23}. Both couplings can be resolved.

(a) Find a pulse sequence that uses only single-qubit rotations so that the signal you measure on qubit 1 determines the matrix element $\hat{\sigma}_x \otimes \hat{\sigma}_x \otimes \hat{\sigma}_x$.

(b) Based on what you know about peak splitting for two qubits, what should the three-qubit spectra from part (a) look like?

3.13 **The initial state in NMR:** The thermodynamic initial state in NMR is $\hat{\rho} = \exp\left(-\hat{H}/kT\right)/Z$, where the partition function is $Z = \mathrm{tr}[\exp\left(-\hat{H}/kT\right)]$. For a field in the z-direction, the energy splitting for a nuclear spin is $\hbar\omega_L$, where ω_L is the Larmor frequency.

(a) Calculate $\hbar\omega_L/kT$ for $T = 300$ K and $\omega_L = 500$ MHz.

(b) Justify the approximation $\hat{\rho} \approx 1/2(\hat{1} - \hat{H}/kT)$.

3.14 **Intensity of NMR signal:** In an NMR sample, how does the ratio between spins up and down depend on the external field B_0? (*Hint*: Boltzmann factor.) Considering that the maximum signal observed is proportional to the difference between the number of spins that are up and down, what is the percentage of spins that can be detected with a magnetic field $B_0 = 7$ T? Study the difference if we study protons and ^{13}C nuclei.

3.15 **Algorithmic cooling:** In this exercise we explore **algorithmic cooling**, where we effectively cool one qubit by shifting some of its entropy to other qubits. Here we will cool one qubit using two others. Each qubit is in contact with a bath and is in the state

$$\hat{\rho}_b = \frac{1}{2}\begin{bmatrix} 1+\epsilon & 0 \\ 0 & 1-\epsilon \end{bmatrix}. \tag{3.78}$$

The polarization of each qubit is $P = \mathrm{tr}[\hat{\rho}_b\hat{\sigma}_z] = \epsilon$. The state of the three qubits is $\hat{\rho}_{123} = \hat{\rho}_b^{(1)} \otimes \hat{\rho}_b^{(2)} \otimes \hat{\rho}_b^{(3)}$.

(a) If $\epsilon > 0$, show that the probability of being in the state $|011\rangle$ is smaller than the probability of being in $|100\rangle$.

(b) It is possible to reduce the entropy of the first spin by swapping the state $|011\rangle$ and $|100\rangle$ with the unitary transformation:

$$\hat{1\!\!1} - |100\rangle\langle 100| - |011\rangle\langle 011| + |100\rangle\langle 011| + |011\rangle\langle 100|. \tag{3.79}$$

Show that the new density matrix is diagonal with the following diagonal elements:

$$\hat{\rho}_{123}^{\text{after}} = \frac{1}{8}\begin{bmatrix} (1+\epsilon)^3 \\ (1+\epsilon)^2(1-\epsilon) \\ (1+\epsilon)^2(1-\epsilon) \\ (1+\epsilon)^2(1-\epsilon) \\ (1+\epsilon)(1-\epsilon)^2 \\ (1+\epsilon)(1-\epsilon)^2 \\ (1+\epsilon)(1-\epsilon)^2 \\ (1-\epsilon)^3 \end{bmatrix}. \tag{3.80}$$

(c) Find the reduced density matrix of the first qubit for small value of ϵ:

$$\hat{\rho}_1^{\text{after}} = \mathrm{tr}_{23}[\hat{\rho}_{123}^{\text{after}}]. \tag{3.81}$$

Show that the polarization ($P = \mathrm{tr}[\hat{\rho}\hat{\sigma}_z]$) is bigger after the swapping of states operation (as long as $\epsilon > 0$).

(d) Calculate the entropy of the first qubit before and after the swap operation. Remember $S = -\mathrm{tr}[\hat{\rho}\log\hat{\rho}] = -\sum_i p_i \log p_i$, where the log is base 2 and p_i are the density matrices eigenvalues. Show that the operation has reduced the entropy of the first qubit.

(e) Show that the entropy of qubits 2 and 3 has increased (for some ϵ).