10 Liquid-state NMR

The first implementations of a quantum computers used nuclear spins as qubits. The corresponding experimental technique is known as nuclear magnetic resonance (NMR). It can be considered the prototypic al implementation, in the sense that the qubits are actual 2-level systems that are relatively well isolated from other degrees of freedom. On the other hand, it represents a radical departure from most other concepts: While one usually thinks of quantum registers as individual systems (and many projects try to implement such systems), NMR represents qubits by some 10^{20} identical copies of a nuclear spin in a suitable molecule. One therefore refers to this type of quantum information processing as "ensemble quantum computing".

As discussed before, spins $S = 1/2$ are the only physical systems that implement directly the 2-dimensional Hilbert space of a qubit. Furthermore, single-qubit logical operations are best understood as rotations of a spin $1/2$ around a magnetic field. The dynamics of a system of spins $1/2$ is thus a simple, idealized representation of any quantum information processor. We therefore describe its operation here in some detail.

10.1 Basics of NMR

Nuclear magnetic resonance is mainly a spectroscopic tool that is used for the analysis of almost any type of molecule, condensed matter or gases in various environments. In the form of MRI (magnetic resonance imaging) it also has become an important tool in clinical medicine. We start with a review of the basics of NMR spectroscopy before we discuss how this technique can be used for quantum computing.

10.1.1 System and interactions

Magnetic resonance is based on the spin degrees of freedom of electronic and nuclear spins. The spin of charged (and some neutral composite) particles has a magnetic dipole moment associated with it; if such particles are placed in a magnetic field, the energy of these magnetic dipoles depends on their orientation with respect to the field.

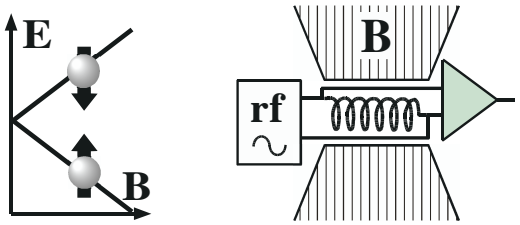

Figure 10.1: Basics of nuclear magnetic resonance (NMR). Left: Zeeman splitting of spin states in a magnetic field. Right: The basic experimental setup consists of a static magnetic field, a radio frequency (RF) generator that creates an alternating magnetic field perpendicular to the static field, and a detector that measures the voltage induced in the coil by the precessing magnetization.

As shown in Figure 10.1, the magnetic field lifts the degeneracy of the spin states. This effect, which is known as the Zeeman effect, is proportional to the strength of the magnetic field. For a spin $S = 1/2$, the splitting of the two energy levels is proportional to the magnetic field strength. Quantum mechanically, it is described by the Hamiltonian

$$\mathcal{H}_z = -\gamma \vec{\mathbf{S}} \cdot \vec{B},$$

where γ is the gyromagnetic ratio of spin S. The usual convention is to orient the z-axis along the static magnetic field. The Hamiltonian then becomes

$$\mathcal{H}_z = -\gamma \mathbf{S}_z B_0 = -\omega_L \mathbf{S}_z, \tag{10.1}$$

where B_0 is the strength of the magnetic field and $\omega_L = \gamma B_0$ the Larmor frequency. For most NMR quantum information processing experiments, we can restrict the discussion to spins $S = 1/2$, for which the Zeeman interaction is the only coupling to external fields.

In magnetic resonance experiments, one uses alternating magnetic fields, which couple to the same magnetic dipole moments, to resonantly excite transitions between these spin states. The resonance condition is that the frequency ω of these alternating fields fulfills the Bohr condition $\hbar\omega = \Delta\mathcal{E}$, where $\Delta\mathcal{E}$ is the separation of the two energy levels ($= \hbar\omega_L$ here). The relevant frequency is in the radio frequency (RF) range for nuclear spins (10–1000 MHz in fields of 1–25 T).

The equation of motion for the spin can be derived classically, by considering the magnetic moment $\vec{\mu} = \gamma\vec{S}$ in a magnetic field $\vec{B}_0$. The interaction between the magnetic field and the magnetic dipole generates a torque

$$\vec{T} = \vec{\mu} \times \vec{B}_0 = \gamma\vec{S} \times \vec{B}_0.$$

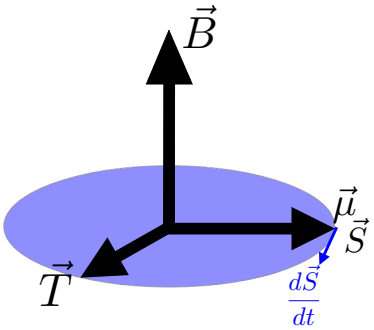

Figure 10.2: Orientation of spin, magnetic field, and the torque generated by the interaction between magnetic moment and magnetic field.

While a torque acting on a classical magnetic moment would rotate it towards the direction of the magnetic field, the spin is also an angular momentum, and the torque is equal to the time derivative of the angular momentum,

$$\frac{d\vec{S}}{dt} = \vec{T}.$$

The resulting equation of motion is thus

$$\dot{\vec{\mu}} = \gamma\vec{\mu} \times \vec{B}_0 = \vec{\mu} \times \vec{\omega}_0.$$

Here, $\vec{\omega}_0 = \gamma\vec{B}_0$ is a representation of the magnetic field in frequency units. Since the time derivative is perpendicular to the direction of the spin and to the magnetic field, the resulting motion is a precession around the magnetic field, rather than a rotation towards the magnetic field.

The same result can also be derived quantum mechanically, from the equation of motion (4.40). Using the commutation relations for angular momentum, the equation of motion becomes

$$\begin{aligned} \frac{d}{dt}\langle \mathbf{S}_x \rangle &= -\omega_L \langle \mathbf{S}_y \rangle \\ \frac{d}{dt}\langle \mathbf{S}_y \rangle &= \omega_L \langle \mathbf{S}_x \rangle \\ \frac{d}{dt}\langle \mathbf{S}_z \rangle &= 0. \end{aligned} \tag{10.2}$$

→ **Problem 1**

The resulting evolution of the spin is a precession around the direction of the magnetic field at the Larmor frequency.

$$\begin{aligned} \langle \mathbf{S}_x \rangle(t) &= S_{xy}(0)\cos(\omega_L t - \phi) \\ \langle \mathbf{S}_y \rangle(t) &= S_{xy}(0)\sin(\omega_L t - \phi) \\ \langle \mathbf{S}_z \rangle(t) &= S_z(0), \end{aligned} \tag{10.3}$$

where $S_{xy}(0)$ is the amplitude of the transverse magnetization and ϕ its phase, i.e., the angle from the x-axis at $t = 0$ (see Fig. 4.2).

As shown in Figure 10.3, this evolution corresponds to a precession around the z-axis, i.e., around the magnetic field. Equation (10.2) is called the Bloch equation, after one of the discoverers of NMR, who also wrote the theory for it

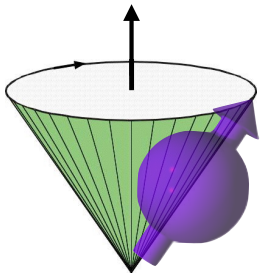

Figure 10.3: Larmor precession of spins in a magnetic field.

[195]. It can also be derived classically and has applications to many two-level systems besides NMR [47]. It also can be used as the equation of motion of a qubit, no matter what the physical basis is.

10.1.2 Radio frequency field

To excite transitions between the different spin states and implement quantum gate operations, one applies a RF magnetic field. It is generated by a current running through a coil that is wound around the sample, as shown in Figure 10.4.

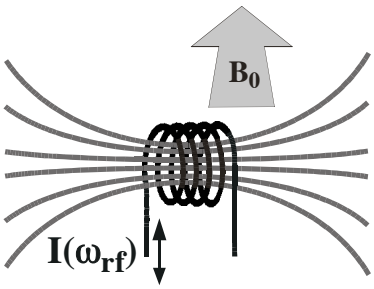

Figure 10.4: An alternating current through a coil generates an RF field perpendicular to the static magnetic field.

The generated RF field is

$$\vec{B}_{rf}(t) \;=\; 2B_1\cos(\omega t)\begin{pmatrix}1\\0\\0\end{pmatrix}\,,$$

where we have chosen the x-axis along the axis of the coil.

This linearly oscillating magnetic field is best described as a superposition of two fields rotating in opposite directions.

$$\vec{B}_{rf}(t) \;=\; B_1\begin{pmatrix}\cos(\omega t)\\ \sin(\omega t)\\ 0\end{pmatrix} + B_1\begin{pmatrix}\cos(\omega t)\\ -\sin(\omega t)\\ 0\end{pmatrix}\,.$$

The first component rotates from x to the y axis (counterclockwise when viewed from the z-axis), the second in the opposite direction.

If we combine this magnetic field with $\vec{B}_0$ into a time-dependent Hamiltonian, we obtain an equation of motion with time-dependent coefficients, which cannot be solved analytically. The same holds true for all qubit systems that are excited by resonantly oscillating control fields. This problem can be solved by moving the time-dependence from the control fields to the coordinate system. For reasons that will become clear, the associated reference frame is known as the rotating frame.

10.1.3 Rotating frame

The resulting dynamics are best analyzed in a coordinate system that rotates around the static magnetic field at the RF frequency. We briefly show here the transformation to this rotating frame since all quantum computing experiments use the rotating frame representation, not the laboratory frame. As shown in Figure 10.5, the two coordinate systems are related by

$$\begin{pmatrix}x\\y\\z\end{pmatrix}^r = \begin{pmatrix}\cos(\omega t) & \sin(\omega t) & 0\\ -\sin(\omega t) & \cos(\omega t) & 0\\ 0 & 0 & 1\end{pmatrix}\begin{pmatrix}x\\y\\z\end{pmatrix}\,.$$

where the vector $\vec{r}^{\,r}$ refers to the rotating coordinate system, the unlabeled one to the laboratory-fixed system.

If we apply this transformation to the RF field, the two circular components become

$$\vec{B}^r_{rf} \;=+B_1\begin{pmatrix}1\\0\\0\end{pmatrix} + B_1\begin{pmatrix}\cos(2\omega t)\\ -\sin(2\omega t)\\ 0\end{pmatrix}\,.$$

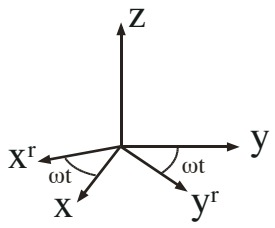

Figure 10.5: Rotating and laboratory-fixed coordinate systems.

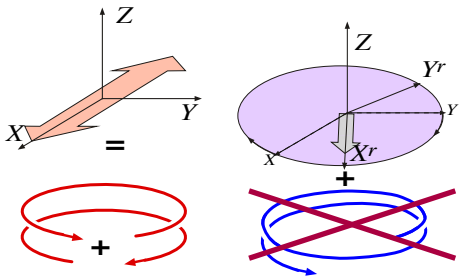

Figure 10.6: Circularly polarized components of the linearly polarized field in laboratory frame (left) and rotating frame (right).

Apparently, one of the two components is now static, while the counter-rotating component rotates at twice the RF frequency. It turns out that, to an excellent approximation, it is sufficient to consider the effect of that component which is static in this coordinate system, while the counter-rotating component can be neglected [72]. It is therefore a convenient fiction to assume that the applied RF generates a circularly polarized RF field, which is static in the rotating frame. The corresponding Hamiltonian is

$$\mathcal{H}^r_{rf} = -\omega_1 \mathbf{S}_x. \tag{10.4}$$

Under most conditions, this approximation yields an excellent description of the actual dynamics. Figure 10.7 compares the exact evolution to the result of the rotating wave approximation. Compared to typical experimental situations, for this figure the parameters have been chosen to exaggerate the deviations by several orders of magnitude. The rapid oscillation occurs at twice the Larmor frequency. In addition, the frequency is shifted slightly, by $\frac{1}{4}\frac{\omega_1^2}{\omega_L}$.

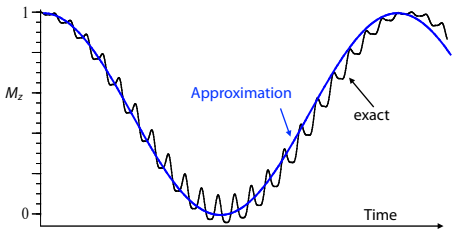

Figure 10.7: Comparison between the exact solution and the rotating wave approximation.

The same reasoning can be used in any type of resonant excitation. In the case of optical spectroscopy (e.g. trapped ion quantum computers), it is known as the rotating wave approximation.

10.1.4 Equation of motion

So far we have transformed the RF field into the rotating frame. We also need to transform the quantum mechanical equation of motion into this reference frame. We start by transforming the state vector, using the unitary operator

$$\mathbf{U}(t) = e^{i\omega t \mathbf{S}_z/\hbar}, \tag{10.5}$$

which defines a rotation around the z-axis. It transforms the laboratory state $|\psi\rangle$ into the rotating frame as

$$|\psi\rangle^r = \mathbf{U}^{-1}|\psi\rangle = e^{-i\omega t \mathbf{S}_z/\hbar}|\psi\rangle. \tag{10.6}$$

To transform an operator A into the same basis, we use

$$A^r = \mathbf{U}^{-1} A \, \mathbf{U}. \tag{10.7}$$

This is valid for all operators, including the density operator or the observables $\mathbf{S}_x$, $\mathbf{S}_y$ and $\mathbf{S}_z$. The only exception that needs special attention is the Hamiltonian. In this case, the transformation has to fulfill the additional requirement that the Hamiltonian remains the generator of the time evolution. Since the transformation $\mathbf{U}$

is time-dependent, the new coordinate system is not an inertial frame of reference. The evolution in this system therefore appears to be subject to additional 'virtual forces' that influence the time evolution and must be accounted for by the transformation. This is in close analogy to centrifugal forces or Coriolis forces that appear if the coordinate system rotates with respect to inertial frames of reference.

Starting with the Schrödinger equation in the laboratory frame

$$\hbar \frac{d}{dt}|\Psi\rangle(t) = -i\mathcal{H}|\Psi\rangle(t),$$

we use eq. (10.6) to substitute

$$|\Psi\rangle(t) = \mathbf{U}|\Psi\rangle^r(t)$$

and obtain an equation of motion for $|\Psi\rangle^r(t)$:

$$\hbar \frac{d}{dt}\left(\mathbf{U}|\Psi\rangle^r(t)\right) = -i\mathcal{H}\mathbf{U}|\Psi\rangle^r(t). \qquad (10.8)$$

The left-hand side can be evaluated with the product-rule:

$$\frac{d}{dt}\left(\mathbf{U}|\Psi\rangle^r(t)\right) = \dot{\mathbf{U}}|\Psi\rangle^r(t) + \mathbf{U}\frac{d}{dt}|\Psi\rangle^r(t)$$

Inserting this into eq. (10.8), rearranging and multiplying with $\mathbf{U}^{-1}$ from the left yields

$$\hbar \frac{d}{dt}|\Psi\rangle^r(t) = -i\mathbf{U}^{-1}\mathcal{H}\mathbf{U}|\Psi\rangle^r(t) - \hbar\mathbf{U}^{-1}\dot{\mathbf{U}}|\Psi\rangle^r(t).$$

The Schrödinger equation in the rotating frame becomes therefore

$$\hbar \frac{d}{dt}|\Psi\rangle^r(t) = -i\mathcal{H}^r|\Psi\rangle^r(t)$$

with the transformed Hamiltonian

$$\mathcal{H}^r = \mathbf{U}^{-1}\mathcal{H}\mathbf{U} - i\hbar\mathbf{U}^{-1}\dot{\mathbf{U}} \qquad (10.9)$$

→ Problem 2

The first term corresponds to the rotation (10.7) of the operator around the z-axis, as for the other operators. The second term takes into account that the rotating coordinate system is not an inertial reference frame, since the rotation is an accelerated motion. Like centrifugal forces, it corrects the equation of motion for the corresponding virtual force. Evaluating this term for the transformation (10.5), we find

$$-i\hbar\mathbf{U}^{-1}\dot{\mathbf{U}} = -i\hbar \cdot \frac{i\omega}{\hbar}\mathbf{S}_z = \omega\mathbf{S}_z.$$

This represents an additional term to the Zeeman operator (10.1). Combining it with the driving Hamiltonian (5.11), we obtain the rotating frame Hamiltonian

$$\begin{aligned}\mathcal{H}^r &= -(\omega_L - \omega)\mathbf{S}_z - \omega_1\mathbf{S}_x = -\Delta\omega_L\mathbf{S}_z - \omega_1\mathbf{S}_x \\ &= -\vec{\omega}_{\text{eff}} \cdot \vec{\mathbf{S}},\end{aligned}$$

wherethe

$$\vec{\omega}_{\text{eff}} = \begin{pmatrix} \omega_1 \\ 0 \\ \Delta\omega_L \end{pmatrix}$$

is the total effective field in the rotating frame, $\omega_1 = \gamma B_1$ is the strength of the RF field in (angular) frequency units and $\Delta\omega_L = \omega_L - \omega$ is the static magnetic field (also in frequency units), reduced by the frequency of the applied field.

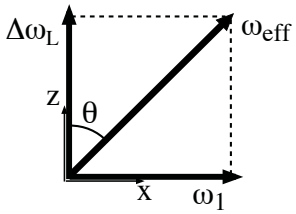

Figure 10.8: Effective magnetic field in the rotating coordinate system.

Figure 10.8 show this vector graphically. The angle θ between $\vec{\omega}_{\text{eff}}$ and the z-axis is given by $\tan\theta = \omega_1/\Delta\omega_L$.

10.1.5 Evolution

The resulting evolution of the spins in the rotating frame is exactly the same as if a (small) static

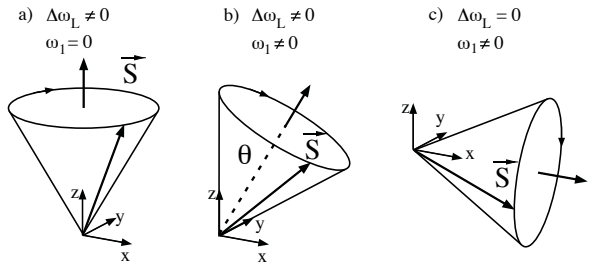

Figure 10.9: Spin precession for the cases of free precession($\omega_1 = 0$, left), resonant irradiation ($\Delta\omega_L = 0$, right), and the general case (center).

field were applied in this direction in the laboratory frame: they undergo a precession around the magnetic field $\vec{\omega}_{eff}$.

Figure 10.9 shows three specific examples for the motion of spins in this effective field. In the absence of RF irradiation ($\omega_1 = 0$), the effective field is aligned with the z-axis and the precession is the same as in the laboratory frame, except that the precession frequency is reduced by ω, the frequency of the applied RF field. In the case of resonant irradiation (shown on the right), the field along the z-axis vanishes and the effective field lies along the x-axis. In the general case b), the effective field lies along a direction in the xz plane.

So far, we have assumed that the direction of the RF field coincides with the x-axis in the rotating frame. This can be changed by shifting the phase of the applied rf signal. As a function of this phase, the coupling Hamiltonian becomes

$$\mathcal{H}^r_{rf} = -\omega_1(\cos\varphi\,\mathbf{S}_x + \sin\varphi\,\mathbf{S}_y)$$

and the effective field

$$\vec{\omega}_{\text{eff}} = \begin{pmatrix} \omega_1\cos\varphi \\ \omega_1\sin\varphi \\ \Delta\omega_L \end{pmatrix}. \tag{10.10}$$

As a simple example, we consider the case that the RF is applied on resonance, with $\varphi = 0$, such that $\vec{\omega}_{\text{eff}} = (\omega_1, 0, 0)$. If the spin is initially aligned with the z-axis, it rotates around the x-axis as

$$\rho(t) = \mathbf{S}_z\cos(\omega_1 t) + \mathbf{S}_y\sin(\omega_1 t). \tag{10.11}$$

The RF field thus rotates it to the y-axis and from there to the negative z-axis. Such a rotation by an angle $\omega_1\tau_p = \pi$, with τ_p the duration of the pulse, corresponds to an inversion of the spins. If the field is left on, the spins continue to precess, returning to the $+z$ axis, again to the negative and so on. This process of successive inversions is called Rabi flopping, in reference to Rabi's molecular beam experiment [73]. The frequency ω_1 at which this process occurs is called the Rabi frequency.

The primary use of RF irradiation in NMR quantum computers is to create logical gate operations. As discussed in Chapter 5, single-qubit gates correspond to rotations of the (pseudo)spins. Pulses of RF radiation are a convenient means for implementing such rotations around arbitrary axes. According to eq. (5.18), the rotation axis $\vec{\omega}_{eff}$ can therefore be oriented in any arbitrary direction by adjusting frequency (and thereby $\Delta\omega_L$) and phase φ of the RF field. The angle of rotation $\alpha = \omega_{\text{eff}}\tau_p$ around the effective field, which is called the flip angle, is given by the product of the effective field strength ω_{eff} and the pulse duration τ_p.

10.1.6 NMR signals

Most NMR signals are obtained in the time domain, as the response of the system to an RF pulse. We assume that the system is initially in thermal equilibrium which is given by Boltzmann statistics:

$$\boldsymbol{\rho}_{\text{eq}} \propto \exp(-\frac{\mathcal{H}}{k_B T}) \approx \mathbf{1} - \frac{\mathcal{H}}{k_B T},$$

where the approximate form, derived for the high-temperature limit

$$\Delta\mathcal{E} = \hbar\omega_L \ll k_B T$$

is always valid in liquid state NMR: under typical experimental conditions, $\frac{\hbar\omega_L}{k_BT}$ is of the order of 10^{-5}. We have therefore

$$\boldsymbol{\rho}_{\text{eq}} = \frac{1}{2}\left(\mathbf{1} + \frac{\hbar\omega_L}{k_BT}\mathbf{S}_z\right).$$

In the simplest NMR measurement, one applies an RF pulse that rotates the spins through an angle $\frac{\pi}{2}$ into the xy plane. According to (10.11), this yields

$$\boldsymbol{\rho}(0+) = \frac{1}{2}\left(\mathbf{1} + \frac{\hbar\omega_L}{k_BT}\mathbf{S}_y\right).$$

After the pulse, the system undergoes Larmor precession under the Zeeman Hamiltonian

$$\begin{aligned}\boldsymbol{\rho}(t) &= e^{-i\mathcal{H}t/\hbar}\boldsymbol{\rho}(0+)e^{i\mathcal{H}t/\hbar} \\ &= \frac{1}{2}\left(\mathbf{1} + \frac{\hbar\omega_L}{k_BT}(\mathbf{S}_y\cos\omega_L t - \mathbf{S}_x\sin\omega_L t)\right).\end{aligned}$$

Detection of the signal should not be treated as a quantum mechanical measurement process. There is no reduction of a wavefunction, and the system is virtually unaffected by the measurement. Rather than projecting onto an eigenstate, one measures the expectation value of a specific observable as a function of time, without disturbing the free evolution of the quantum system. This is of course closely related to the fact that the system consists of an ensemble of many spins rather than a single particle. It is thus more appropriate to use a classical picture for the detection of the signal.

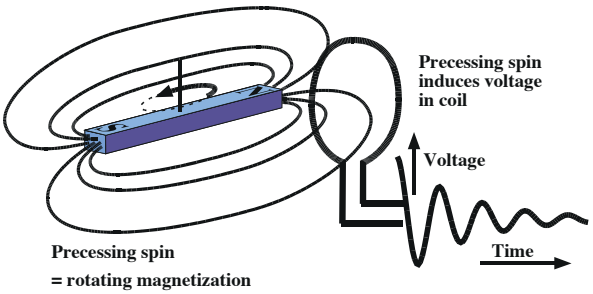

Figure 10.10: Detection of freely precessing spins through the Faraday effect.

Figure 10.10 shows how observation of the precessing spins is achieved through the Faraday effect. The polarized spin ensemble is a macroscopic magnetization; as it precesses, it changes the flux through the RF coil, thus inducing a voltage signal proportional to

$$s(t) \propto \frac{d}{dt}\Phi(t) \propto \frac{d}{dt}\sum_i \langle \mathbf{S}_x^i\rangle \propto \cos(\omega_L t). \quad (10.12)$$

Damping effects, which are not discussed here, cause a decay of the signal,

$$s(t) \propto \cos(\omega_L t)e^{-t/T_2}.$$

This signal, which is generated by freely precessing magnetization that slowly decays is known as free induction decay (FID).

10.1.7 Resonance lines

For an analysis of the signal one usually considers not the time domain signal, but its Fourier transform. For an FID decaying exponentially with time constant T_2, the spectrum becomes

$$s(\omega) = \sqrt{\frac{1}{2\pi}}\frac{T_2}{1+(\omega-\omega_L)^2T_2^2},$$

i.e., a Lorentzian with a half-width at half height $\frac{1}{T_2}$ centered at the Larmor frequency ω_L.

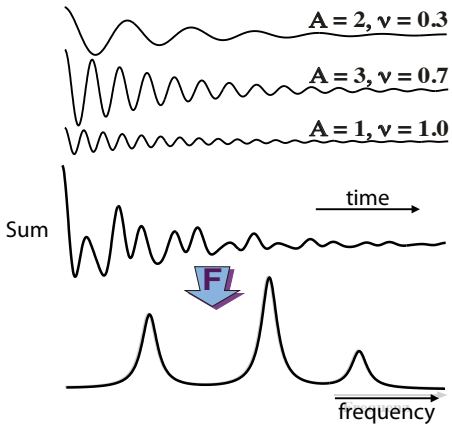

Figure 10.11: Superposition of 3 time-domain signals and its Fourier transform.

While the frequency-domain signal contains the same information as the time-domain FID, it is still very useful to do this transformation. The main advantage of the Fourier transform is that it allows one to distinguish different transitions: two distinct transitions usually have different Larmor frequencies

$$\omega_{ij} = \frac{E_i - E_j}{\hbar}.$$

Here, E_i and E_j are the energies of the eigenstates connected by the transition. The corresponding resonance lines are therefore separated in frequency space, while the time domain signals overlap. As an example, fig. 10.11 shows 3 FID signals

$$s_{ij}(t) = A_{ij} e^{i\omega_{ij} t}$$

with different amplitudes A_{ij} and frequencies $\nu_{ij} = \omega_{ij}/2\pi$ as well as their sum. Clearly, the sum is very hard to interpret. If the total time domain signal is Fourier-transformed, the resulting spectrum, shown in the bottom trace, clearly shows the three different signal contributions as separate resonance lines.

The amplitude of each resonance line is determined by the product of a density operator element with an element of the observable; in the simplest case, where the nontrivial part of the initial density operator and the observable are identical,

$$\boldsymbol{\rho}(0) - \frac{1}{2}\mathbf{1} = \mathbf{A} = \mathbf{S}_x,$$

and the amplitudes A_{ij} of the individual transitions in the spectrum become

$$A_{ij} \propto |(\mathbf{S}_x)_{ij}|^2.$$

10.1.8 Refocusing

In many NMR experiments, and particularly in (NMR-) quantum computation, it is necessary to eliminate unwanted interactions. This includes unwanted environmental perturbations, such as magnetic field inhomogeneities, or unwanted terms in the system Hamiltonian. As an example for the latter, consider a system of multiple qubits coupled by an interaction such as a bilinear coupling $S_z^i S_z^k$ between qubit i and k. Such couplings are essential for 2-qubit gate operations, but they are unwanted when a single-qubit gate is to be generated. The terms that we consider, are either linear or bilinear terms, i.e. they contain a single spin operator (e.g. S_x) or a product of two spin operators acting on two different spins, such as $S_z^1 S_z^2$.

Refocusing can eliminate such terms. This is usually achieved by a sequence of RF pulses that modulates the evolution in such a way that the total effect of the interaction on the system vanishes. The first such experiment is the "Hahn-echo" observed in liquid state NMR by Erwin Hahn [122]. We consider here only the simplest cases, which must fulfill the two conditions

- The interaction $\mathcal{H}_1$ can be inverted by a control operation, $\mathcal{H}_c : \mathcal{H}_1 \to -\mathcal{H}_1$
- The interaction commutes with the static system Hamiltonian, $[\mathcal{H}_0, \mathcal{H}_1] = 0$.

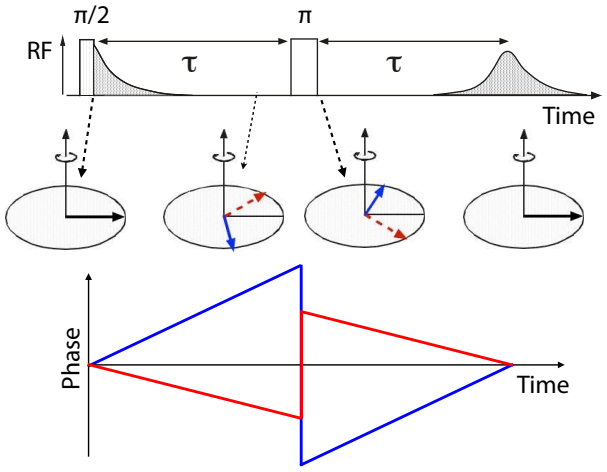

Figure 10.12: Refocusing of magnetic field inhomogeneities in a Hahn echo experiment.

Figure 10.12 shows a typical experiment. The initial $\frac{\pi}{2}$ RF pulse converts longitudinal into transverse magnetization $\propto \mathbf{S}_x$ that subse-

quently precesses in the magnetic field. For a system of uncoupled spins, the density operator after the RF pulse is

$$\begin{aligned}\boldsymbol{\rho}(\tau) - \frac{1}{2}\mathbf{1} &\propto e^{-i\mathcal{H}\tau/\hbar}\mathbf{S}_x e^{i\mathcal{H}\tau/\hbar} \\ &= \mathbf{S}_x\cos(\Delta\omega_L\tau) + \mathbf{S}_y\sin(\Delta\omega_L\tau).\end{aligned}$$

As shown in the lower part of figure 10.12, the phase $\Delta\omega_L\tau$ (which represents the orientation of the magnetization in the xy plane) increases linearly with time. If two spins experience different magnetic fields, their precession frequency differs. In the figure, the blue and red lines indicate the evolution of the phase of two spins that experience different magnetic fields (e.g., due to magnetic field inhomogeneity). In the central part of the figure, the blue and red arrows indicate the orientation of these spins. If a distribution of such Larmor frequencies is present, the overall effect will be destructive interference and a loss of signal, as indicated in the upper part of Figure 10.12 and discussed in more detail in section 7.2.

To refocus this destructive interference process, one can apply a second RF pulse. A π_x pulse, i.e. a rotation around the x-axis by an angle π, leaves the x-component of the density operator invariant but inverts the y-component:

$$\begin{aligned}\boldsymbol{\rho}(\tau+) - \frac{1}{2}\mathbf{1} &\propto \mathbf{S}_x\cos(\Delta\omega_L\tau) - \mathbf{S}_y\sin(\Delta\omega_L\tau) \\ &= \mathbf{S}_x\cos(-\Delta\omega_L\tau) + \mathbf{S}_y\sin(-\Delta\omega_L\tau).\end{aligned}$$

Apparently, the pulse inverts the phase of the xy magnetization vector, as indicated in the lower part of Figure 10.12. After the pulse, the spins continue to precess in the magnetic field. If the Larmor frequency remains constant over time, the total phase acquired during the time τ after the refocusing pulse is equal to the phase that the spin acquired between the two pulses, before its phase was inverted. As a result, the total phase

$$\varphi_t = -\Delta\omega_L\tau + \Delta\omega_L\tau = 0$$

vanishes, independently of the Larmor frequency of the spin. The destructive interference is then eliminated, and a "spin-echo" is observed.

In a similar way, unwanted couplings between spins (qubits) can be eliminated by suitable refocusing sequences. In an AX system (see Section 10.2.3, e.g.,) the coupling term can be eliminated by applying a refocusing pulse to one of the spins. For a Hamiltonian

$$\mathcal{H}_{AX} = \omega_A\mathbf{A}_z + \omega_X\mathbf{X}_z + d\mathbf{A}_z\mathbf{X}_z, \quad (10.13)$$

the initial condition $\boldsymbol{\rho}(0) - \frac{1}{2}\mathbf{1} \propto \mathbf{A}_x + \mathbf{X}_x$, and equal precession periods before and after a π pulse on the X-spin, the system evolves to

$$\begin{aligned}\boldsymbol{\rho}(2\tau) - \frac{1}{2}\mathbf{1} &= \mathbf{U}(\tau)e^{-i\pi\mathbf{X}_x/\hbar}\mathbf{U}(\tau)(\mathbf{A}_x + \mathbf{X}_x)\mathbf{U}^\dagger(\tau) \\ &\quad \cdot e^{i\pi\mathbf{X}_x/\hbar}\mathbf{U}^\dagger(\tau) \\ &= \mathbf{U}(\tau)e^{-i\pi\mathbf{X}_x/\hbar}\mathbf{U}(\tau)e^{i\pi\mathbf{X}_x/\hbar}(\mathbf{A}_x + \mathbf{X}_x) \\ &\quad \cdot e^{-i\pi\mathbf{X}_x/\hbar}\mathbf{U}^\dagger(\tau)e^{i\pi\mathbf{X}_x/\hbar}\mathbf{U}^\dagger(\tau),\end{aligned}$$

where $\mathbf{U}(\tau) := e^{-i\mathcal{H}_{AX}\tau/\hbar}$ is the time evolution operator describing the precession. Using

$$e^{-i\pi\mathbf{X}_x/\hbar}\mathcal{H}_{AX}e^{i\pi\mathbf{X}_x/\hbar} = \omega_A\mathbf{A}_z - \omega_X\mathbf{X}_z - d\mathbf{A}_z\mathbf{X}_z,$$

we find that the the refocusing pulse eliminates the effect of the Zeeman term $\mathbf{X}_z$ as well as the coupling term $\mathbf{A}_z\mathbf{X}_z$, but leaves the Zeeman term of the A spin. Similar refocusing schemes are possible to eliminate different terms in larger spin systems.

10.2 NMR as a molecular quantum computer

10.2.1 Spins as qubits

The two quantum states that represent a qubit correspond naturally to the two states of a spin-1/2 – the only quantum system whose Hilbert space has exactly two states. It is therefore always possible to use the Feynman–Vernon–Hellwarth picture [47] to describe the qubit as a virtual spin-1/2. In this chapter, however, the virtual spin is a real nuclear spin of a molecule in solution: we study NMR systems to show how quantum computers can

be implemented. It should be realized, however, that the quantum computers that can be built this way still have very limited capabilities. They should not be compared to conventional computers, which have been developed over half a century, but to early prototypes, whose development only started twenty years ago.

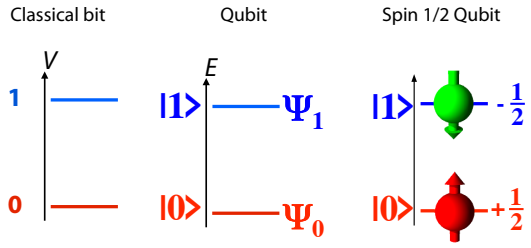

Figure 10.13: Identification of bits with voltage levels (classical computer, left), quantum mechanical states (generic quantum computer, center), and states of a spin-1/2 (right).

Using the spins as qubits requires a mapping of the logical qubit states to the spin states. As shown in Figure 10.13, the spin states take over the role of voltage levels in classical computers. Conventionally, one chooses the $|m_S = +1/2\rangle$ state to represent a logical 0, while the $|m_S = -1/2\rangle$ state represents a logical 1. To construct a quantum register, one needs several distinguishable qubits.

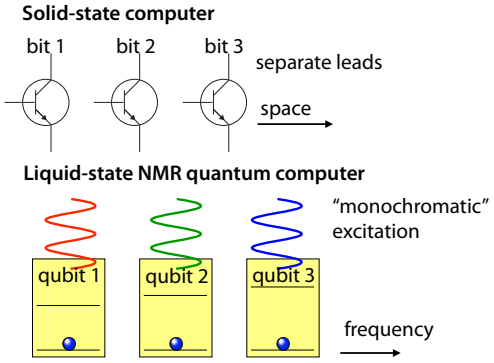

Figure 10.14: Addressing of qubits in NMR quantum computers vs. solid state computers.

Every computer using the network architecture must be able to selectively address individual bits or qubits. As indicated in Figure 10.14, conventional electronic computers (e.g. Si-based) use wires for this purpose. In liquid state NMR quantum computers, the qubits are nuclear spins of freely floating molecules; clearly it is not feasible to use wires for addressing in this case. Nevertheless, it is possible to address qubits selectively. Since the qubit gates are applied with resonant RF fields, they are only effective when the RF frequency is close to the Larmor frequency of the spin.

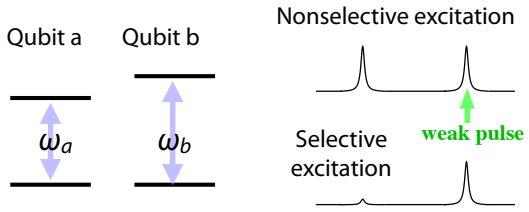

Figure 10.15: Left: qubits with different resonance frequencies. Right: Selective vs. nonselective excitation.

As shown in Fig. 10.15, spins whose Larmor frequency differs from the frequency of the RF pulse are not affected by the pulse to a first approximation, provided the frequency separation $\Delta\omega_0 = |\omega_a - \omega_b|$ is larger than the Rabi frequency ω_1 of the excitation pulse. The width of the affected frequency range is therefore inversely proportional to the duration of the RF pulse.

For nuclear spins of different isotopes, this condition can be readily fulfilled. Typical resonance frequencies for typical nuclear spin qubits are, in a field of $B_0 = 14$ T:

Isotope	^{1}H	^{13}C	^{19}F
$\omega_0/2\pi$ [MHz]	600	151	565

The frequency differences are this in the MHz range while typical Rabi frequencies are of the order of $\omega_1 \approx 10$ kHz, so the condition $\Delta\omega_0 \gg \omega_1$ is always fulfilled if the qubits are associated with different nuclear spin species.

10.2.2 Chemical shift

Since the number of nuclear isotopes is limited and the number of suitable nuclear isotopes is very limited, it is desirable to have other mechanisms for distinguishing the qubits. The easiest possibility is the chemical shift. The nuclear Larmor frequency is $\omega_L = \gamma B_0$. The gyromagnetic ratio γ is the same for every spin of a given isotope. However, the magnetic field B_0 can differ, since it corresponds not to the externally applied field, but to the local field at the site of the nucleus.

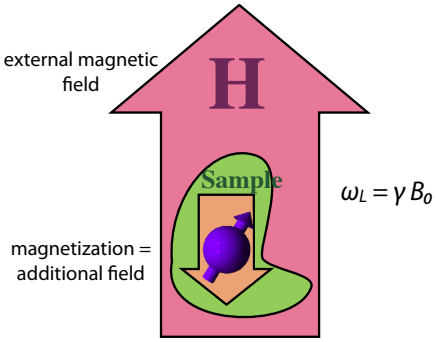

Figure 10.16: Chemical shift: local magnetic fields shift the Larmor frequency.

The magnetic field strength at the site of the nucleus differs in general from the externally applied magnetic field: the electron system, in which the nucleus is embedded, has a non-vanishing magnetic susceptibility. These shifts depend therefore on the electronic structure and are generally known as "chemical shift". The Hamiltonian that describes such a system of qubits can be written as

$$\mathcal{H}_Z = -\sum_i \omega_i \, \mathbf{S}_z^i,$$

where the index i runs over all spins (qubits). These frequency shifts are proportional to the magnetic field strength and can be used to distinguish different qubits.

Figure 10.17 shows a typical example. The molecule ethylbenzene is a standard used for calibrating NMR spectrometers. It contains three

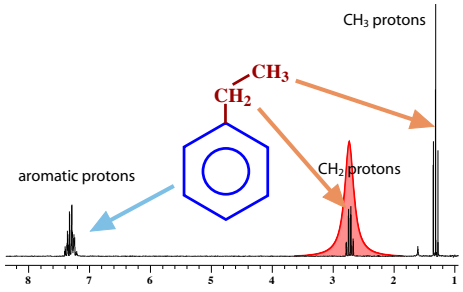

Figure 10.17: Chemical shift of protons in the NMR spectrum of ethylbenzene.

types of protons, which cover a chemical shift range of about 6 ppm. In general, the available chemical shift range depends on the isotope considered. In the case of protons (^{1}H), the range $|\omega_i - \omega_j|/\omega_i$ is of the order of 10 ppm. For ^{13}C, it is about 200 ppm, and similar for ^{15}N. For a typical ^{1}H NMR frequency, the available frequency range is therefore of the order of 6 kHz, for ^{13}C in the same field 30 kHz. In contrast to conventional computers, where lithographic patterns localize different bits, this may be considered a bottom-up approach, where the molecular structure determines the location of the qubit in frequency space.

10.2.3 Coupled spin systems

Implementation of quantum algorithms requires two-qubit gates, which must rely on couplings between qubits / spins. Such couplings are naturally present in nuclear spin systems and are exploited also in NMR spectroscopy.

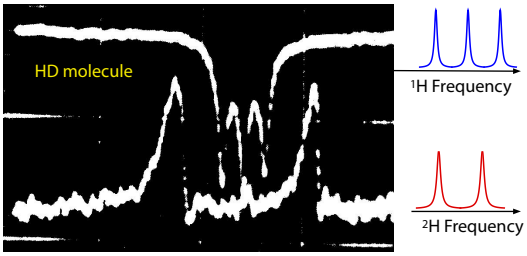

Figure 10.18: Splitting of resonance lines in the HD-molecules by the coupling between ^{1}H and D=^{2}H.

Figure 10.18 shows the effect of the coupling on the spectrum of the HD-molecule, which contains two different nuclear spin species. The ^{1}H nucleus has a spin 1/2, while the D=^{2}H nucleus, which contains one proton as well as one neutron, has spin $I = 1$. As a result of the interaction, the ^{1}H spectrum (upper trace) splits into three resonance lines, while the spectrum of the ^{2}H splits into two. As we show below, the number of resonance lines is given my the multiplicity $2I + 1$ of the coupling partner.

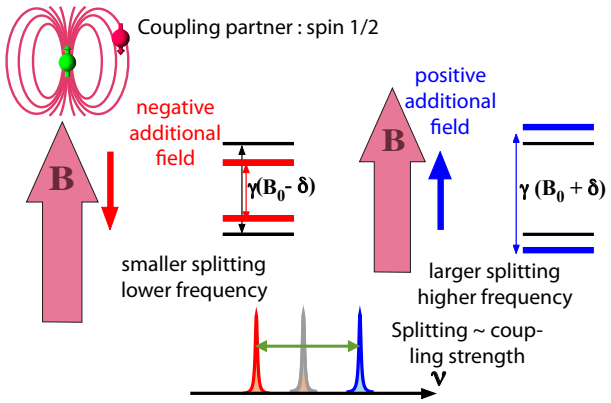

Figure 10.19: Coupling between nuclear spins as an effective field.

The effect of the coupling can qualitatively be understood in a semiclassical picture, where every spin generates a small additional field, which is felt by the coupling partner(s). Its direction depends on the orientation of the spin that generates the field, as shown in Fig. 10.19. Depending on the sign of the additional field, it increases or reduces the Zeeman splitting of the coupling partner and therefore shifts its resonance frequency. This shift is proportional to the magnetic quantum number of the source spin. Accordingly, every line of the multiplet can be labeled by the spin state of the source spin.

There are two main types of couplings; the first is called scalar, indirect, or J-coupling, the second type is the direct or dipolar coupling. In the case of the dipolar coupling, the coupling energy can be calculated from the distance and orientation of the spins:

$$\mathcal{E}_{dd} = \frac{\mu_0}{4\pi r_{12}^3}\mu_1\mu_2(1 - 3\cos^2\theta).$$

Here, r_{12} is the distance between the two nuclei, μ_i are their magnetic moments, and θ is the orientation angle of the internuclear vector with respect to the direction of the external magnetic field. The isotropic average of the Legendre polynomial $P_2(\theta) = 1-3\cos^2\theta$ vanishes. In isotropic liquids, where molecules rotate freely, these interactions are therefore averaged to zero and do not contribute to the evolution of the spins (but to their relaxation).

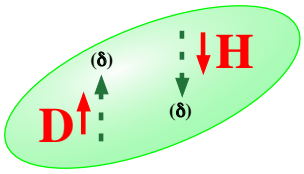

Figure 10.20: J-coupling between the nuclear spins of the HD-molecule.

The isotropic coupling (J-coupling), however, does not depend on the molecular orientation and is therefore not affected by the molecular motion. As a result, only the scalar J-couplings are observed in the spectrum. These couplings are mediated by the electrons in the chemical bonds and are therefore often called indirect couplings. Fig. 10.20 illustrates the underlying mechanism: the hyperfine interaction between the nuclear spins and the electron spins of the chemical bonds lifts the degeneracy of the two electronic spin states and leads to a (very small!) polarisation of the electron spins in that bond, with the energetically favored orientation having a slightly higher density near the nucleus. Accordingly, the opposite spin state has a slightly higher density at the other nucleus. The hyperfine interaction therefore lowers the energy of the opposite nuclear spin state in the second nucleus. The energy of the two nuclear spins therefore includes a term

$$\mathcal{E}_J = -J\vec{\mu}_1 \cdot \vec{\mu}_2.$$

In both cases, the coupling between two spins can be understood as a small additional magnetic field generated by spin A and acting on spin X, as well as in the opposite direction. We consider here only the simplest case (which is most useful for NMR quantum computing), where the interaction can be written as

$$\mathcal{H}_{AX} = d\mathbf{A}_z\mathbf{X}_z,$$

where A and X are the two spins and d the coupling constant. The total Hamiltonian is then

$$\mathcal{H} = \mathcal{H}_z + \mathcal{H}_{AX} = -\omega_A\mathbf{A}_z - \omega_X\mathbf{X}_z + d\mathbf{A}_z\mathbf{X}_z$$

$$\frac{1}{2}\begin{pmatrix} -\omega_A-\omega_X+\frac{d}{2} & & & \\ & -\omega_A+\omega_X-\frac{d}{2} & & \\ & & \omega_A-\omega_X-\frac{d}{2} & \\ & & & \omega_A+\omega_X+\frac{d}{2} \end{pmatrix}$$

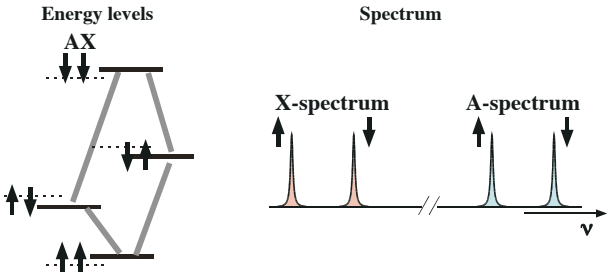

Figure 10.21: Energy levels and spectrum of a system of two qubits A and X. The dashed horizontal lines indicate the energy levels of the Zeeman Hamiltonian (no coupling), the solid lines the energies of the full Hamiltonian.

Figure 10.21 shows the energy levels and the spectrum of such a two-spin system. The dashed horizontal lines in the left-hand part indicate the energy levels of the Zeeman Hamiltonian alone (no coupling), the solid lines represent the energies of the full Hamiltonian. The coupling shifts the states with parallel orientation of the two spins upwards (for a positive sign of the coupling constant d), the states with antiparallel orientation downwards.

As discussed in section 10.1.6, NMR signals are generated by precessing magnetization. Observable transitions must therefore have non-vanishing matrix elements of the transverse spin operators S_x and S_y. These allowed transitions correspond to the flip of a single spin by one quantum. In the present spin system, such transitions are those between the states $\uparrow\uparrow\leftrightarrow\uparrow\downarrow$, $\uparrow\uparrow\leftrightarrow\downarrow\uparrow$, $\uparrow\downarrow\leftrightarrow\downarrow\downarrow$, $\downarrow\uparrow\leftrightarrow\downarrow\downarrow$. The transition frequencies are

$$\begin{aligned} \omega_{12} &= \omega_{\uparrow\uparrow\leftrightarrow\uparrow\downarrow} = \omega_X - d/2; \\ \omega_{13} &= \omega_{\uparrow\uparrow\leftrightarrow\downarrow\uparrow} = \omega_A - d/2; \\ \omega_{24} &= \omega_{\uparrow\downarrow\leftrightarrow\downarrow\downarrow} = \omega_A + d/2; \\ \omega_{34} &= \omega_{\downarrow\uparrow\leftrightarrow\downarrow\downarrow} = \omega_X + d/2; \end{aligned}$$

The spectrum consists of four lines, each of which is associated with a transition of one spin and labeled by the (invariant) state of the second spin.

10.2.4 Pseudo / effective pure states

Before NMR quantum computing was demonstrated, all algorithms for quantum computers assumed that quantum computers use individual quantum systems, which are initially prepared in a specific quantum state. Unfortunately, detecting individual spins is extremely difficult and has only been achieved in a few specific systems [196, 197, 198, 199, 200, 201]. In most cases, signals can be detected only from macroscopic ensembles of spins, containing some 10^{20} spins. These spins are not in identical quantum mechanical states and therefore cannot be described by a pure state. For the description of the mixed states, one has to use a density operator.

Many quantum algorithms require pure quantum states for their implementation and can therefore not be applied to NMR systems in thermal equilibrium. Nevertheless, mixed states can be made to mimic pure states and therefore allow the implementation of these algorithms. For this purpose, the target system has to be prepared in an initial state that can be written as the sum of the unit operator and an operator representing a pure state:

$$\boldsymbol{\rho}_{pp} \propto \beta\mathbf{1} + \alpha\boldsymbol{\rho}_p,$$

where $\boldsymbol{\rho}_{pp}$ is referred to as a "pseudo-pure" state, or "effective pure state", while $\boldsymbol{\rho}_p$ is a pure state.

The density operator component proportional to the unity operator has vanishing magnetization and therefore does not contribute to the signal. In addition, it commutes with the Hamiltonian and therefore does not undergo time evolution. Accordingly, the behavior of such a system is completely determined by the second component and therefore is exactly equal to that of a pure state. The coefficient α is largely determined by the polarization of the spin system.

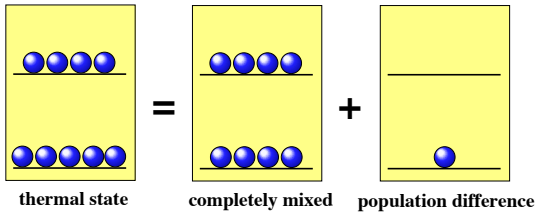

Figure 10.22: A thermal state of a single qubit is a pseudo-pure state.

As shown in Fig. 10.22, a single spin is always in a pseudo-pure state (compare (4.31)).

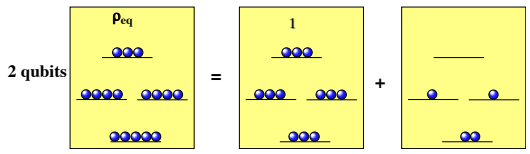

Figure 10.23: A thermal state of a 2-qubit system.

In multi-qubit spin systems, however, the thermal equilibrium states are not even pseudo-pure. Figure 10.23 illustrates this for a 2-qubit system. Unitary operations cannot be used to bring such a system into a pseudo-pure state. Instead one has to average over a number of different mixed states to make the pseudo-pure state.

There are a number of procedures for implementing such an averaging scheme, which are referred to as "spatial labeling" [10], "temporal labeling"[202] and "logical labeling" [203]. Temporal labeling is perhaps easiest to explain, using the example of two coupled spins. If the two spins are from the same type of nuclear isotope, the populations of the four states are

$$\uparrow\uparrow\colon 1/4+\epsilon \qquad \uparrow\downarrow,\downarrow\uparrow\colon 1/4 \qquad \downarrow\downarrow\colon 1/4-\epsilon,$$

as illustrated in Figure 10.23. Here $\epsilon \approx \hbar\omega_L/k_BT \approx 10^{-5}$ is determined by the ration of the Zeeman energy to the thermal energy of the system.

To obtain a pseudo-pure state, one can equalize the populations of three levels (e.g., $\uparrow\downarrow,\downarrow\uparrow,\downarrow\downarrow$) by cyclically permuting them and adding the results of three experiments with different initial conditions. The time-averaged populations would then be

$$\frac{1}{4}\begin{pmatrix}1\\1\\1\\1\end{pmatrix}+\epsilon\begin{pmatrix}1\\-\frac{1}{3}\\-\frac{1}{3}\\-\frac{1}{3}\end{pmatrix}$$
$$=\ (\frac{1}{4}-\frac{\epsilon}{3})\begin{pmatrix}1\\1\\1\\1\end{pmatrix}+\frac{4\epsilon}{3}\begin{pmatrix}1\\0\\0\\0\end{pmatrix}.$$

The corresponding averaged density operator corresponds to the sum of the unit operator (= the totally mixed state) and a pure state.

The main disadvantage of this procedure is that the averaging process reduces the polarization and therefore the signal amplitude. In the case of spatial labeling, one turns the population differences of states 2, 3, 4 into transverse magnetization, which is destroyed by pulsed field gradients. It was soon realized [204] that this loss of polarization, which increases exponentially with the number of spins in the quantum register, severely restricts the usefulness of liquid-state NMR quantum computing. Similarly, the number of operations required increases exponentially with the number of qubits. This can be reduced to polynomial overhead by logical labeling [203], which uses additional (ancilla) spins to create pure states for specific ancilla spin configurations. For the related techniques POPS [205] or SALLT [206], the overhead is independent of the number of qubits.

While this loss of signal is a severe problem for scalability, it is not always necessary to prepare a pseudo-pure state. Many quantum algorithms can also be applied to mixed states, and some of these mixed state algorithms actually run faster than the corresponding pure state algorithms [207, 208]. These algorithms, which can provide an exponential speedup even over optimal quantum algorithms, use a combination of quantum parallelism with classical parallelism: the ensemble of nuclear spins corresponds then to a large number of quantum computers running in parallel.

10.2.5 Single-qubit gates

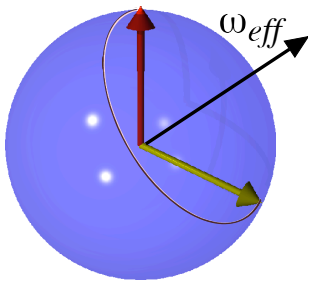

Figure 10.24: Single-qubit gate implemented as a rotation around the effective field ω_{eff}.

Single-qubit gates are implemented by RF pulses. In the rotating frame, an RF pulse can be represented by its propagator

$$U = e^{-i\mathcal{H}\tau/\hbar} = e^{i\vec{\omega}_{eff}\cdot\vec{S}\tau/\hbar},$$

where $\mathcal{H}$ is the Hamiltonian during the pulse and τ the duration of the pulse. Depending on the phase φ of the RF field, the propagator for a resonant pulse is

$$e^{-\frac{i}{\hbar}\beta(\mathbf{S}_x\cos\varphi+\mathbf{S}_y\sin\varphi)}$$

The flip angle is $\beta = \omega_1\tau$, where ω_1 is the amplitude of the RF field (the Rabi frequency) and τ the duration of the pulse.

As discussed in section 2.2.5, the shortest possible duration of a single-qubit gate is given by the energy level splitting. In the case of spins, the energy level splitting is given by the strength ω_1 of the interaction with the RF field and the gate duration τ is inversely proportional to the interaction strength, $\tau = \beta/\omega_1$, where β is the rotation angle of the gate. In section 10.2.1, we showed that addressing of individual qubits requires that the strength of the rf field is weaker than the difference between their resonance frequencies. Accordingly, the duration of single-qubit gates will always be longer than the inverse of the frequency difference between the qubits.

Combining two such generators (rotations) with different axes, it is possible to implement any SU(2) operation. An important example is the set of rotations around the z-axis, which cannot be generated by RF pulses directly. They can, however, be realized by combining three rotations around axes in the xy plane:

$$\begin{aligned} e^{-i\phi\mathbf{S}_z/\hbar} &= \begin{pmatrix} e^{-i\phi/2} & \\ & e^{i\phi/2} \end{pmatrix} \\ &= e^{-i\frac{\pi}{2}\mathbf{S}_x/\hbar}e^{-i\phi\mathbf{S}_y/\hbar}e^{i\frac{\pi}{2}\mathbf{S}_x/\hbar} \\ &= e^{-i\frac{\pi}{2}\mathbf{S}_y/\hbar}e^{i\phi\mathbf{S}_x/\hbar}e^{i\frac{\pi}{2}\mathbf{S}_y/\hbar} \end{aligned} \quad (10.14)$$

We now consider the most important single-qubit gates. Using the conventional choice of relative phases between states, the NOT gate may be implemented, up to an irrelevant overall phase, by

$$\begin{aligned} \text{NOT}: \quad e^{-i\pi\mathbf{S}_x/\hbar} &= \begin{pmatrix} & -i \\ -i & \end{pmatrix} \\ &= e^{-i\frac{\pi}{2}}\begin{pmatrix} & 1 \\ 1 & \end{pmatrix}. \end{aligned}$$

This implementation of NOT thus differs from the usual representation by an overall phase of $-\frac{\pi}{2}$. Since such overall phases do not correspond to observable quantities, we will not consider them here and regard all implementations that differ by such a phase factor as equivalent.

One might first think that any 180 degree pulse, which inverts the two states $|0\rangle$ and $|1\rangle$ should

be an implementation of NOT. However, looking at the propagator for a π_y pulse,

$$e^{-i\pi \mathbf{S}_y/\hbar} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

one sees that this differs from the NOT in terms of the *relative* phase that it applies to the two states.

The Hadamard gate

$$\mathbf{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

can also be implemented by an RF pulse

$$\frac{i}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = e^{-i(\frac{\pi}{\sqrt{2}})(\mathbf{S}_x + \mathbf{S}_z)/\hbar}$$

Physically this transformation can be achieved in a number of different ways: either by applying an off-resonant RF pulse with $\Delta\omega_L = \omega_1$, or by a sequence of RF pulses along the y, x and $-y$ axes:

$$\begin{aligned} \mathbf{H} &= e^{i\frac{\pi}{4}\mathbf{S}_y/\hbar} e^{-i\pi \mathbf{S}_x/\hbar} e^{-i\frac{\pi}{4}\mathbf{S}_y/\hbar} \\ &= e^{-i\frac{\pi}{4}\mathbf{S}_y/\hbar} e^{-i\pi \mathbf{S}_z/\hbar} e^{i\frac{\pi}{4}\mathbf{S}_y/\hbar} \\ &= e^{-i\pi \mathbf{S}_z/\hbar} e^{i\frac{\pi}{2}\mathbf{S}_y/\hbar}. \end{aligned} \tag{10.15}$$

The three-pulse version is also interesting: as in the case of the composite z-rotation (10.14), it can be understood as a "rotated rotation". The central pulse executes the desired π rotation around an axis in the xy plane. The first and last pulses then rotate the axis from the xy plane into the xz plane. This scheme is experimentally easier to implement since it only requires resonant pulses.

The last version in eq. (10.15) is the shortest: a $\left(\frac{\pi}{2}\right)_y$ pulse is followed by a π_z rotation, which can be implemented, e.g., by a phase shift. If the z-rotation is omitted, this gate is known as the pseudo-Hadamard gate

$$\mathbf{h} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix} = e^{i\frac{\pi}{2}\mathbf{S}_y/\hbar},$$

which can replace the Hadamard gate in many cases. It is not it's own inverse, but

$$\mathbf{h}^{-1} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} = e^{-i\frac{\pi}{2}\mathbf{S}_y/\hbar}.$$

They correspond to $\pm\frac{\pi}{2}$ rotations around the y axis.

→ **Problem 3**

10.2.6 Two-qubit gates

Two-qubit gates require couplings between the spins to apply transformations to one spin conditional on the state of the other spin. There are two different ways of implementing such gates. One may be referred to as "soft pulses", the other as "hard pulses plus free precession". The first uses the fact that weak RF fields affect only transitions whose resonance frequency is close to the RF frequency. As we discussed in Section 10.2.3, the transitions of a nuclear spin that is coupled to another spin can be labeled by the state of the coupling partner. A weak RF field whose frequency matches the frequency of one resonance of spin A (e.g.) therefore excites spin A on the condition that spin X is in the $|1\rangle$ state – a CNOT gate.

$$\text{CNOT} = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & & 1 \\ & & 1 & \end{pmatrix}.$$

This variation is conceptually simple since it can be described in terms of two-level systems, and it can be extended to more complicated spin systems. The condition that the pulse must be selective requires that the RF field ω_1 be weak compared to the coupling d, $\omega_1 \ll d$, and therefore that the gate duration τ be long compared to the inverse coupling strength, $\tau d \gg 1$. As a result, this type of gate operation is more susceptible to decoherence.

Figure 11.24 shows, as an example, how a CNOT gate can be implemented by a selective π-pulse

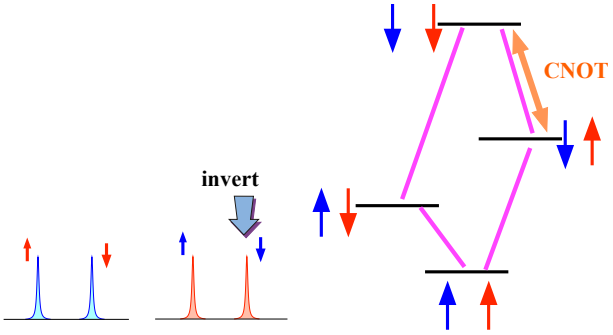

Figure 10.25: CNOT gate implemented by a selective pulse on the $\downarrow\downarrow\leftrightarrow\downarrow\uparrow$ transition.

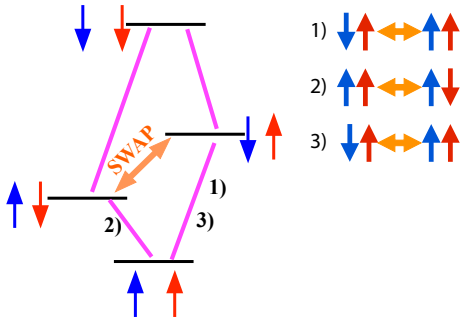

Figure 10.26: SWAP gate implemented by a sequence of three selective π-pulses.

applied to the $\downarrow\downarrow\leftrightarrow\downarrow\uparrow$ transition, i.e. the $11 \leftrightarrow 10$ transition.

Figure 10.26 shows a second example. In this case, the SWAP operation $\uparrow\downarrow\leftrightarrow\downarrow\uparrow$ cannot be implemented directly, since the matrix element of the magnetic dipole operator vanishes for the transition $\uparrow\downarrow\leftrightarrow\downarrow\uparrow$. It can be replaced by a sequence of three π-rotations, e.g. the combination given in fig. 10.26.

10.2.7 Two-qubit gates with nonselective pulses

The second approach consists of a combination of single qubit gates with periods of free precession. We consider a spin A coupled to a control spin X by the interaction $d\mathbf{A}_z\mathbf{X}_z$. The corresponding Hamiltonian (10.13) yields a spectrum with two resonance lines in the A-spectrum, which can be labeled by the states $|\uparrow\rangle$ and $|\downarrow\rangle$ of the X spin. We will assume that pulses can be applied to the A and X spin separately – a condition which must be satisfied for the one-qubit gates. In contrast to the first implementation, however, the pulses used here always act on all transitions of a given spin, independent of the state of its coupling partner(s), as in any single-qubit gate.

Starting from the state $|00\rangle = |X =\uparrow, A =\uparrow\rangle$, an $e^{-i\frac{\pi}{2}\mathbf{A}_y/\hbar}$ RF pulse creates a superposition state

$$|\Psi_\uparrow(0)\rangle = \frac{1}{\sqrt{2}}[|0\rangle \otimes (|0\rangle + |1\rangle)].$$

Free precession converts it to

$$|\Psi_\uparrow(t)\rangle = \frac{1}{\sqrt{2}}[|0\rangle \otimes (|0\rangle e^{-i\hbar dt/4} + |1\rangle e^{i\hbar dt/4})],$$

where we use a rotating frame that is resonant with the Zeeman frequency for the A and (independently) for the X spin.

After a time $t = \frac{\pi}{d\hbar}$, the phase factor is $e^{\pm i\pi/4} = (1 \pm i)/\sqrt{2}$ and the system has reached the state

$$\begin{aligned} |\Psi_\uparrow(\frac{\pi}{d\hbar})\rangle &= \frac{1}{2}[|0\rangle \otimes ((1-i)|0\rangle + (1+i)|1\rangle)]. \\ &= \frac{1-i}{2}[|0\rangle \otimes (|0\rangle + i|1\rangle)], \end{aligned}$$

which corresponds to an $\pi/2$ rotation of the A-spin around the z-axis. An $e^{-i\frac{\pi}{2}\mathbf{A}_x/\hbar}$ pulse applied at this time returns the system to its original state $|00\rangle$ (apart from an overall phase factor).

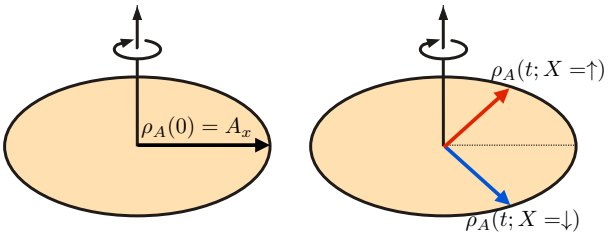

Figure 10.27: Evolution of nuclear spin coherence under a coupling to another spin-1/2.

As shown in fig. 10.27, this can be readily followed in terms of a vector model. The initial

y-pulse turns the spin from the z-axis to the x-axis. The red vector, which corresponds to the control qubit being in the $\uparrow\rangle$ state, then precesses by 90 degrees to the y-axis, and the subsequent x-pulse flips it back to the z-axis.

If we apply the same sequence of pulses to the state $|10\rangle = |X =\downarrow, A =\uparrow\rangle$, the free precession occurs with opposite sign

$$|\Psi_\downarrow(t)\rangle = \frac{1}{\sqrt{2}}[|1\rangle \otimes (|0\rangle e^{i\hbar dt/4} + |1\rangle e^{-i\hbar dt/4})],$$

as shown by the blue vector in fig. 10.27. Accordingly, the second pulse rotates the spin to the negative, rather than the positive z-axis. Clearly, this corresponds to an inversion conditioned on the control qubit being in the $|1\rangle$ state.

The free precession period under the Hamiltonian

$$\mathcal{H}_{AX} = \frac{d}{\hbar}\mathbf{A}_z\mathbf{X}_z$$

implements the transformation $e^{i(\pi\mathbf{A}_z\mathbf{X}_z)/\hbar^2}$. Together with the two pulses, this implements a CNOT gate:

$$\begin{aligned} & e^{i\frac{\pi}{2}\mathbf{A}_y/\hbar} e^{-i(\frac{\pi}{2}\mathbf{X}_z + \frac{\pi}{2}\mathbf{A}_z - \pi\mathbf{A}_z\mathbf{X}_z/\hbar)/\hbar} e^{-i\frac{\pi}{2}\mathbf{A}_y/\hbar} \\ = \; & e^{-i(\frac{\pi}{2}\mathbf{X}_z + \frac{\pi}{2}\mathbf{A}_x - \pi\mathbf{A}_x\mathbf{X}_z/\hbar)/\hbar} \\ = \; & (1+i)\begin{pmatrix} 1 & & & \\ & 1 & & \\ & & & 1 \\ & & 1 & \end{pmatrix}. \end{aligned}$$

The additional terms of $\mathbf{X}_z$ and $\mathbf{A}_z$ are for normalization of the relative phases. They can be implemented as composite z-pulses [209] or by an appropriate choice of the reference frequency.

Three qubit gates like the Toffoli gate can be constructed in the same way as two-qubit gates - either by selective pulses or by a combination of single-qubit gates and free precession periods. Formally, a three-qubit operation involves three-particle interactions, corresponding to Hamiltonian terms $\mathcal{H}_3 = ABC$, where A, B and C are single-qubit operators of the three involved qubits. Such interactions do not exist on the fundamental level, but they can be created artificially, by using transformations like

$$e^{-i\beta\mathbf{B}_y\mathbf{C}_z} e^{-i\alpha\mathbf{A}_z\mathbf{B}_x} e^{i\beta\mathbf{B}_y\mathbf{C}_z} = e^{-i\gamma\mathbf{A}_z\mathbf{B}_z\mathbf{C}_z}.$$

Here, $\mathbf{A}_\alpha$, $\mathbf{B}_\beta$, and $\mathbf{C}_\gamma$ refer to the three qubits and the three factors can be generated in the same way as the CNOT operation discussed above, by combining free precession periods under the effect of a two-qubit coupling Hamiltonian with single-qubit gates. Each factor corresponds to a 2-qubit gate operation.

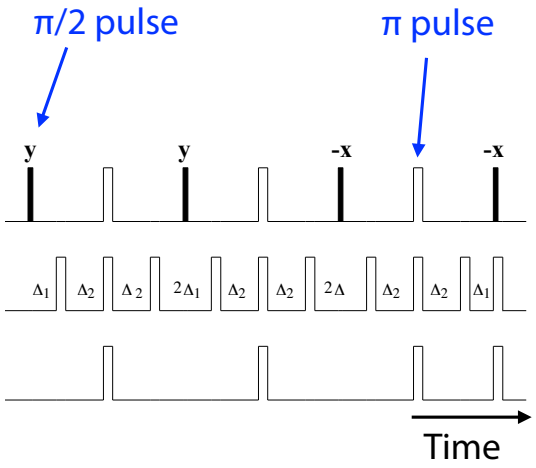

Figure 10.28: Pulse sequence for implementing a Toffoli gate.

Figure 10.28 shows a possible pulse sequence for implementing the Toffoli gate with hard pulses. Alternatively, three- or N-spin gates may be generated using selective pulses [210, 211, 212]. The Toffoli gate, e.g., can be implemented by a selective inversion pulse applied to the transition $|110\rangle \leftrightarrow |111\rangle$.

10.2.8 Qubit readout

As discussed in Section 10.1.6, detection in magnetic resonance is best described in a classical picture: the transverse components of the spin generate a macroscopic magnetization that precesses around the static magnetic field. Obviously such a detection scheme is not compatible with the usual description of a quantum mechanical measurement, which involves the collapse of

a wavefunction. Instead, one observes the system continuously, without significantly affecting its behavior. This difference is closely related to the fact that the system is an ensemble, rather than the usually assumed single-particle system. In addition, the observed quantity is not the population of some state, i.e., $|\psi_k\rangle\langle\psi_k|$, but rather the evolution of a coherence, i.e., $|\psi_j\rangle\langle\psi_k|(t)$, where $|\psi_{j,k}\rangle$ are eigenstates of the Zeeman Hamiltonian.

According to equation (10.12), the signal contribution of a specific coherence is proportional to the corresponding matrix element of the total spin operator $\sum_i \mathbf{S}_y^i$. This matrix element vanishes unless exactly one of the spins changes its magnetic quantum number, i.e., unless the transition occurs between two states

$$|i\rangle = |m^0, m^1, ...m^{N-1}\rangle$$

and

$$|f\rangle = |m^{0\prime}, m^{1\prime}, ...m^{N-1\prime}\rangle$$

with $m^{j\prime} = m^j$ for all but one j.

The observed signal is the sum over the contributions of the individual spins. A measurement at a single instant in time therefore does not determine the values of the individual qubits, but only the sum over all qubits. However, the possibility for continuous measurements makes it possible to distinguish the contributions from the individual qubits, since they evolve at different frequencies. As we discussed in Section 10.2.1, all spins in an NMR qubit register must have different Larmor frequencies to allow addressability for logical operations. This condition also implies that their precession frequencies during detection are different. As discussed in section 10.1.7, Fourier transformation of the FID from such a system therefore separates the contributions from different qubits in frequency space.

Measuring the FID is a straightforward way to measure the expectation value of transverse spin components S_x and S_y. When a quantum algorithm requires the measurement of populations, it can be trivially modified to allow for implementation on an NMR quantum computer. One adds an RF pulse that converts the populations into transverse coherence and again measures the FID of the system.

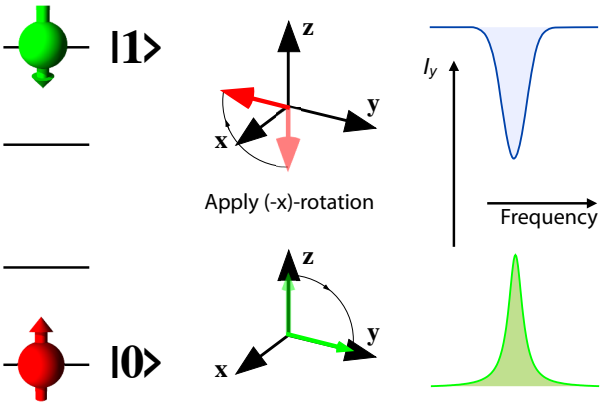

Figure 10.29: Readout of populations with the help of an RF pulse for the two basis states. The vector diagram shows how the spin is rotated by the RF pulse and the (single line) spectra show how the resulting amplitudes identify the qubit state.

Figure 10.29 shows, as an example, the signal that one observes from a single qubit if it is in one of the two eigenstates before the RF pulse is applied. If it is in the ground state, which corresponds to the spin pointing along the direction of the magnetic field, the RF pulse rotates it to the positive y-axis. Since $\mathbf{S}_y$ is the observable, we expect a positive signal at the Larmor frequency of this qubit. If the spin is in the logical $|1\rangle$ state instead, it always points in the opposite direction and the signal becomes negative.

There are cases in quantum computation, where the readout process hinges on the collapse of a wavefunction. For those cases, which include Shor's algorithm, the algorithm must be modified when it is applied to an NMR system. The non-existence of a collapse is handled by appending an additional step, which is polynomial in the number of bits and allows one to obtain the result from ensemble measurements [9, 213].

10.2.9 Readout in multi-qubit systems

As the number of qubits increases, the number of resonance lines in the associated NMR spectra also increases. While the addressability criterion mandates an increase in the number of lines that is proportional to the number of qubits, the couplings between the spins (which are needed for two-qubit gates) increase the number of lines much more rapidly. Every coupling partner doubles the number of resonance lines. If all N qubits are coupled to all other qubits (which is usually not the case), each qubit gives rise to 2^{N-1} resonance lines, corresponding to the 2^{N-1} states of its coupling partners and the total number of lines is $n_L = N2^{N-1}$.

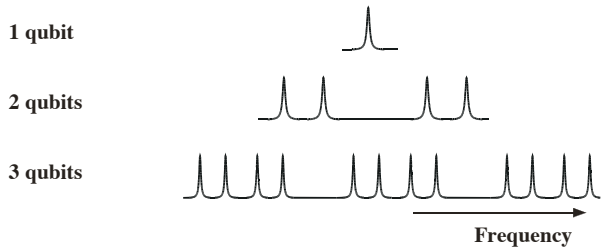

Figure 10.30: Increase in the number of resonance lines in N spin systems.

Figure 10.30 shows the number of resonance lines for $N = 1$, 2, and 3 qubits. In most real systems, not all couplings are large enough to be measurable, resulting in a smaller number of lines.

This exponential increase in the number of lines in a finite frequency bandwidth limits the number of useful qubits. However, it does have the advantage that the spectrum contains much more information about the state of the quantum mechanical system than the simple readout of individual qubits. Every group of lines associated with transitions of qubit $|j\rangle$ provides information about the state of the corresponding qubit but, in addition, it can also yield information about the states of the other qubits. To illustrate this, we consider the two-qubit system of Section 10.2.3 and assume that we are interested in the readout of the states

$$|00\rangle, |01\rangle, |10\rangle, |11\rangle.$$

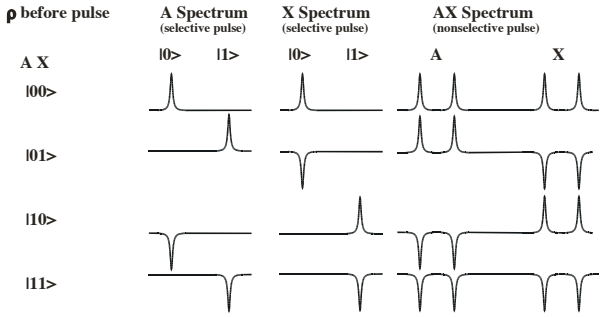

Figure 10.31: Signals in NMR readout for different spin states.

Figure 10.31 shows how these states can be distinguished by applying an RF pulse, measuring the FID and calculating its Fourier transform. If we apply the pulse only to the A or X spin, we measure only a partial spectrum. Each partial spectrum consists of two resonance lines that can be labeled with the quantum state of the coupling partner. If the coupling partner X is in state $|0\rangle$, e.g., the spectrum of the A spin only shows the single resonance line associated with this state. Starting from the state $|00\rangle$, e.g., and applying a $\pi/2$ rotation to the A-spin, we obtain the state

$$\Psi_{00+} = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |0\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |10\rangle),$$

which evolves with the single frequency $(\mathcal{E}_{10} - \mathcal{E}_{00})/\hbar$ and therefore generates a single resonance line in the spectrum, as shown in the upper left of fig. 10.31. Even the partial spectrum of either spin provides therefore a clear distinction between all four possible cases.

It is also possible to apply an RF pulse that excites both spins simultaneously. The resulting nonselective spectrum, shown in the last column, again allows for a clear distinction between the four cases.

This scheme can easily be extended to more spins; examples are given, e.g., in [214]. In general, a spectrum of a weakly coupled N-spin system contains $N2^{N-1}$ resonance lines. Taking into account that the usual NMR experiments measure not only $\sum_i \mathbf{S}_y^i$, but also $\sum_i \mathbf{S}_x^i$, this

number doubles to $N2^N$. The number of resonance lines is thus even larger than 2^N, the total number of coefficients that describe a pure state of N qubits. This shows that the resonance line amplitudes are not independent of each other.

10.2.10 DiVincenzo's criteria

DiVincenzo[8] listed five criteria that implementations of quantum computers should fulfill to be considered "useful". We summarize here to what degree liquid state NMR fulfills these criteria:

1. **Well-defined qubits.**

   The usual implementations use nuclear spins $S = 1/2$ and identify $|0\rangle = |\uparrow\rangle$ and $|1\rangle = |\downarrow\rangle$. The qubits are well characterized in the sense that their energies are well known and the coupling to external fields occurs only through the Zeeman interaction. In the liquid state NMR experiments, logical qubits are not represented by individual spins, but by collections of spins of the order of Avogadro's number. This is in contrast to the usual assumption of quantum computation theory, and some consequences of this need to be addressed in the context of readout and initialization.

   In liquid state NMR, the individual qubits are distinguishable by their resonance frequency. The resonance frequencies of the different spins may be shifted by chemical shift effects or the qubits may be represented by different isotopes. The latter is clearly preferable, since it avoids cross-talk between qubits. However, since the number of useful isotopes is limited, assigning different isotopes to different qubits is clearly not a scalable procedure. When one uses chemical shift differences, the separation should be as large as possible to allow for fast operations of logical gates.

   In summary, NMR systems fulfill the "qubit-identification" requirement quite well, but liquid-state NMR appears to fail the scalability criterion.

2. **Initialization into a well defined state.**

   In liquid state NMR, initialization is achieved by relaxation, which provides for an excess of spins in the ground state. For algorithms designed to work with pure states, this must be combined with the preparation of a pseudo-pure state. While these procedures can be used for small spin systems, they are clearly not scalable for larger systems. Furthermore, thermal initialization is not sufficient for repeated quantum error correction, which will represent an essential part of scalable quantum computing.

3. **Long decoherence times.**

   The long decoherence time (of the order of a second) of liquid state NMR is one of its biggest advantages. However, typical durations of two-qubit gates are at least several milliseconds, so the number of gates that can be applied is limited to approximately 100.

4. **A universal set of quantum gates.**

   At this point, liquid state NMR scores very well: the implementation of unitary transformations is well established and rather straightforward.

5. **A qubit-selective readout.**

   Another strong point, as discussed above. The differentiation of qubits requires chemical shift separation, but is much easier to achieve than the addressing during gating. It is even possible to read out the full density operator, rather than only the populations, as in standard quantum computing algorithms.

10.2.11 Scaling behavior

Liquid state NMR was the first experimental technique that allowed the implementation of quantum algorithms. Nevertheless, there are serious obstacles to advancing this system much farther. One difficulty is associated with the

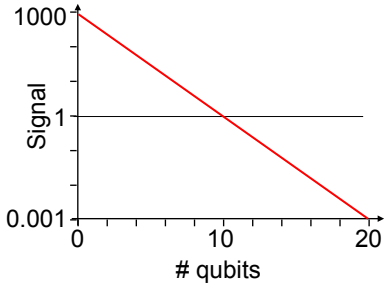

Figure 10.32: Loss of signal amplitude due to preparation of pseudo-pure states as a function of the quantum register size.

preparation of pseudo-pure states [204]: The procedure averages all populations but one. As long as the spin system can be described by the high-temperature approximation, the population of an individual spin state is inversely proportional to the number of states. It therefore decreases as 2^{-N} with the number of spins N. The detectable signal size therefore limits the possible number of spins to be used in such a quantum information processor. Figure 10.32 shows the resulting decrease of the signal amplitude.

The reduction of sensitivity associated with the preparation of pseudo-pure states can be avoided by using algorithms that do not require pure states to work with. For this purpose, variations of algorithms have been developed that can be applied directly to mixed states [215, 207, 216, 208]. For the purpose of database search, such modified algorithms can even be exponentially faster [207, 208] than the original algorithm developed by Grover [133].

Another approach to beating the exponential decrease of the signal size due to the pseudo-pure state preparation would be to work with sufficiently high spin polarization that one can create good approximations of pure states. Virtually complete polarization of the electron spins by thermal relaxation can be achieved at a temperature of 100 mK in a magnetic field of 2 T, where $\frac{\hbar\omega}{k_B T} = 27 \gg 1$. High enough nuclear spin polarization, in contrast, cannot be achieved in thermal equilibrium within the currently accessible experimental conditions.

Highly spin polarized hydrogen nuclei can be obtained by several non-equilibrium techniques, e.g., by separating the ortho- and para- components in molecular hydrogen gas [217]: The energy of para-hydrogen is lower than that of ortho-hydrogen. Accordingly, it is possible to bring pairs of protons in hydrogen molecules into the singlet spin state by cooling them to low temperature. When the symmetry between the two nuclei in the molecule is broken, e.g., through a chemical reaction, it is possible to achieve truly entangled nuclear spin states [218, 219]. Other approaches to pure state preparation include optical pumping [220, 221] or polarization exchange with electron spins at very low temperature [222, 223]. All these techniques require that the system be kept at low temperature to avoid competing processes that reduce the polarization. This also implies that the material that contains the spins be a solid rather than a liquid.

Another aspect of liquid state NMR that may make it difficult to scale up to larger numbers of qubits, is the addressing of the individual qubits. Current implementations use the natural chemical shift range of the nuclear spins to distinguish them by their resonance frequency. Since the chemical shift range is limited, this procedure cannot be extended to arbitrarily large numbers of spins. The larger the number of qubits, the smaller is therefore the separation of their resonances and therefore the slower the switching speed. It appears therefore necessary to design an addressing scheme that does not rely on chemical shift differences.

10.3 NMR Implementation of Shor's algorithm

When Peter Shor published his algorithm for factorization in polynomial time (Ref. [19], for details see Section 8.3), it generated enormous boost to the field of quantum information processing. Similarly, its first physical implementa-

tion convinced more people that quantum computing may not remain a field of science fiction but become reality. The first experimental implementation was published by a group at IBM Almaden Research Center near San Jose, California, using an NMR quantum computer [213]. They factorized 15, the smallest integer to which the Shor algorithm can be applied (remember: N must be odd and not the power of a prime).

10.3.1 Qubit implementation

For the implementation of Shor's factoring algorithm, Vandersypen *et al.* used a custom-designed molecule with five ^{19}F and two ^{13}C nuclear spins.

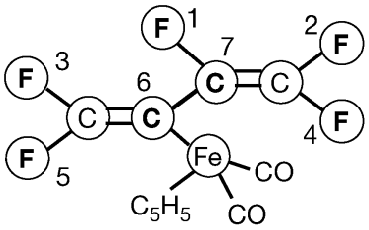

Figure 10.33: Custom designed molecule with seven nuclear spin qubits [213].

The use of carbon and fluorine nuclei spreads the frequencies over a relatively wide range and therefore allows for fast processing. ^{19}F and ^{13}C are both spins-1/2, have generally long decoherence times and a large chemical shift range that allows for fast gating of the qubits. As shown in fig. 10.33, the custom-built molecule contains five fluorine and four carbon nuclei. As actual qubits, five fluorine and two carbon nuclei were used; two additional carbon nuclei were not used in this experiment.

As shown in figure 10.34, the chemical shift separation between the qubits is typically of the order of 1 kHz, thus allowing for single-qubit gate switching times of the order of 1 millisecond. Each qubit is coupled to every other qubit, although some of the coupling constants are relatively small. While the large number of coupling constants allows for direct implementation of all

Frequencies, relaxation times

	i	$\omega_i/2\pi$	$T_{1,i}$	$T_{2,i}$
^{19}F	1	−22052.0	5.0	1.3
	2	489.5	13.7	1.8
	3	25088.3	3.0	2.5
	4	−4918.7	10.0	1.7
	5	15186.6	2.8	1.8
^{13}C	6	−4519.1	45.4	2.0
	7	4244.3	31.6	2.0

Coupling constants

i	J_{7i}	J_{6i}	J_{5i}	J_{4i}	J_{3i}	J_{2i}
1	−221.0	37.7	6.6	−114.3	14.5	25.16
2	18.6	−3.9	2.5	79.9	3.9	
3	1.0	−13.5	41.6	12.9		
4	54.1	−5.7	2.1			
5	19.4	59.5				
6	68.9					
7						

Figure 10.34: Resonance frequencies, relaxation times and coupling constants of the molecule.

two-qubit gates, it leads to a rather complicated spectrum: since every spin is coupled to six other spins, we expect $2^6 = 64$ resonance lines for every spin or a total of $7 \cdot 64 = 448$ lines. Most of these transitions can actually be observed, but several resonance lines are so close in frequency that they are difficult to distinguish.

Figure 10.35 shows the multiplet structure for the first qubit, which has resolved couplings to the other six qubits. Another consequence of the many couplings is that for every gate most of the couplings must be refocused.

Shor's algorithm consists of two main blocks: the period-finding algorithm and the quantum Fourier transform (QFT). It requires a quantum register consisting of n workspace qubits and m qubits to store the number N to be factorized. For $N = 15$, m must be at least 4, since $2^4 = 16 > 15$ and n in the general case 8. However, using specific properties of the $N = 15$ case (for details, see section 10.3.3), n can be reduced to 2. In their implementation, Vandersypen *et al.* chose $n = 3$, to find additional periods.

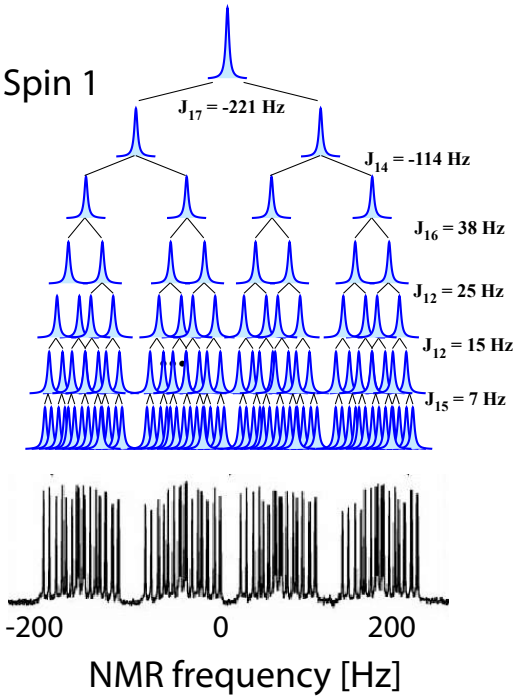

Figure 10.35: Qubit 1 multiplet structure from the couplings to 6 other qubits. The bottom trace shows the experimental spectrum.

10.3.2 Initialization

Shor's algorithm starts with the initial state

$$|\psi_0\rangle = |0000001\rangle, \tag{10.16}$$

i.e., a pure state. The NMR system must therefore be first be brought from the thermal to a pseudo-pure state. In this case, Vandersypen *et al.* used temporal averaging. As we discussed in Section 10.2.4, the temporal averaging process for two spins involves a sum over three different experiments. For the seven-qubit system used for the factorization experiment, the total number of states is 128. Out of these, the population of 127 must be averaged, while that of the last one is kept. Equalizing 127 populations could be done be averaging over 127 cyclic permutations. Using the fact that the populations of many of these states are already equal, it was possible to reduce the number of individual experiments to 36.

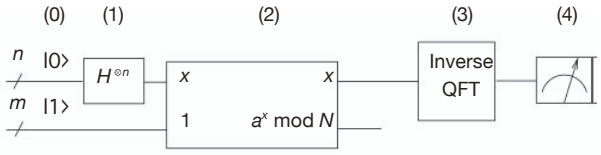

Figure 10.36: Simplified network model of Shor's algorithm [213].

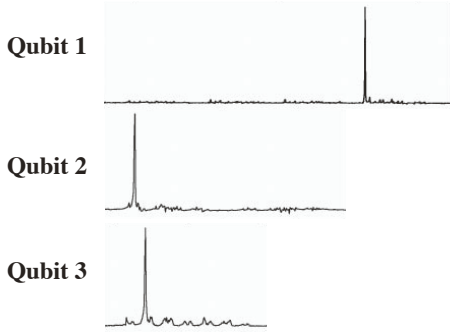

Figure 10.37: Demonstration of pure state preparation in the spectra of qubits 1–3 [213].

The success of the preparation scheme can be checked by applying a selective readout pulse to the system, measuring the resulting FID and converting it into a spectrum. We recall that the different resonance lines in the multiplet of lines originating from a single spin can be labeled by the state of its coupling partners. If these coupling partners are in a pure state that is also an eigenstate of the Hamiltonian (e.g. $|000000\rangle$, this corresponds to a single line in the spectrum of the fist spin. Starting from a pure eigenstate like (10.16), we therefore expect to find only a single line in the spectrum of a selectively excited qubit. As Figure 10.37 shows, this is fulfilled to an excellent approximation in the spectra of the first three qubits. The small additional lines that would not be present in the ideal case can be used to quantitate the degree of purity achieved by the state preparation.

While the source register is initiated in the state $|0\rangle$, the target register is initially in state $|1\rangle$. This is achieved by first initiating it into state $|0\rangle$ and subsequently flipping bit 7. In the im-

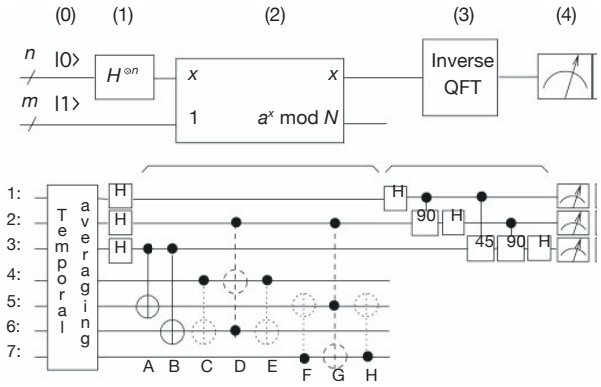

Figure 10.38: Implementation of Shor's algorithm by gates for N=15 and a=7 [213].

plementation details of fig. 10.38, qubits 1-3 hold the workspace bits, with qubit 1 the msb and qubit 3 the lsb, while qubits 4-7 hold the register m, with qubit 7 representing the lsb.

The next step is the generation of the superposition of all spin states of qubits 1–3 (the input qubits) through the Hadamard transformation. The Hadamard gates were implemented by spin-selective $\frac{\pi}{2}$ pulses on the first three qubits.

10.3.3 Computational steps

One of the crucial steps of Shor's algorithm (as well as of corresponding classical algorithms) is the modular exponentiation $f(q) = a^q \mod N$ for 2^n values in parallel. As discussed in Section 8.3.3, this is done qubit by qubit with the help of the identity

$$a^q = a^{2^{n-1}q_{n-1}} \ldots a^{2q_1} a^{q_0}, \tag{10.17}$$

where q_n are the bits of the binary representation of q. While the period of $f(q)$ can be as large as N, only the values 2 and 4 appear for N=15. Since a must be coprime with N, the possible choices of a for N=15 are 2, 4, 7, 8, 11, 13 and 14. For the choices $a = 2$, 7, 8, and 13, one finds $a^4 \mod 15 = 1$, while $a^2 \mod 15 = 1$ for $a =$ 4, 11 and 14. Since we only require the values 0, 1, 2 and 3 for q, it can be encoded in a 2-qubit quantum register. Vandersypen *et al.* chose to use three qubits for encoding q; the additional qubit may be used for test purposes. Together with the $m = 4$ ($\geq \log_2 15$) qubits needed to encode $f(q)$, a total of seven qubits were used. To implement the exponentiation efficiently, the powers of a were precomputed on a classical computer. The eight values of q are stored as a superposition in the qubits labeled 1, 2, 3 in Figure 10.38. The exponentiation is then computed in the target register through CNOT operations.

The first step is a multiplication mod 15 with a^{q_0}. q_0 is encoded in qubit 3, so multiplication by a^{q_0} corresponds to multiplication with a if qubit 3 is 1 and to multiplication with $a^0 = 1$,i.e. to NOOP, if qubit 3 is 0. Since the target register is initialized into state $|1\rangle$, multiplication by a can be done by adding $(a-1)$, again controlled by qubit 3. This addition can be implemented by two CNOT operations: for $a = 7$, qubits 5 and 6 must be changed from 0 to 1. The controlled addition is therefore achieved by the operation CNOT $(3,5)$ CNOT $(3,6)$, as shown in Figure 10.38, where the operations are labeled A and B. For $a = 11$, qubits 4 and 6 must be incremented, which is done as CNOT $(3,4)$ CNOT $(3,6)$.

The second step is multiplication with $a^{2q_1} \mod 15$. For $a = 7$, this corresponds to multiplication by $7^2 - 15 \cdot 3 = 4$, controlled by q_1 or qubit 2 in Figure 10.38. Multiplication by 4 corresponds to a shift of the bits in the register by 2 positions. In the case of modular multiplication, the shift is replaced by a rotation. In the case of the four-bit register m, multiplication by 4 can thus be implemented by a rotation by 2 positions, which corresponds to swapping bits 0 with 2 and 1 with 3. In Figure 10.38, this corresponds to SWAP operations of 4 with 6 and 5 with 7, both controlled by qubit 2. Each SWAP operation can be decomposed into 3 CNOT operations, of which the second is turned into a CCNOT for the controlled SWAP. These CNOT and CCNOT operations are labeled CDE and FGH in Figure 10.38. Vandersypen *et al.* used a number of simplifications (="compiler optimizations") to simplify or eliminate specific gates, taking advantage of the

special situation. These simplifications are indicated in the figure as dotted gates (can be eliminated) or dashed gates (can be simplified). Gate C can be eliminated because the control qubit is zero, thus reducing the gate to the unity operation. The doubly controlled gates D and G act on target bits that are in basis states (not superposition states), which allows for additional simplifications. Gate F can be simplified to a NOT operation, since the control qubit is always 1. Finally, gates E and H can be omitted, since they act on qubits that are no longer accessed afterwards and therefore do not affect the result.

After the multiplication step, Shor's algorithm requires an (inverse) QFT of the register n. It contains Hadamard gates and phase gates (i.e., z-rotations) of 45 and 90 degrees. In practice, the phase gates are usually turned into rotations of the coordinate axes: rather than apply actual z-pulses (which can be implemented by composite rotations), one simply shifts the phases of all earlier pulses by the corresponding amount. This reduces the power deposition of the system and the overall duration of the algorithm, and the resulting fidelity is higher, since phase shits are essentially ideal rotations.

10.3.4 Readout

At the end of the standard algorithm, the information is stored in the populations of the spin state. As discussed in Section 10.2.8, one obtains the populations by applying an RF pulse, measuring and Fourier transforming the FID.

The three spectra shown in Figure 10.39 display the resulting state of the three qubits for an input of $a = 11$. They contain only positive lines for qubits 1 and 2, indicating that they are in state $|0\rangle$ at the end of the computation. Qubit 3 has one positive and one negative line, indicating that it is in a superposition state $|0\rangle + |1\rangle$.

After the inverse QFT, qubit 3 is the most significant qubit. The exponentiation therefore generates the states $|100\rangle = |4\rangle$ and $|000\rangle = |0\rangle$. This indicates that the period of the probability (8.44)

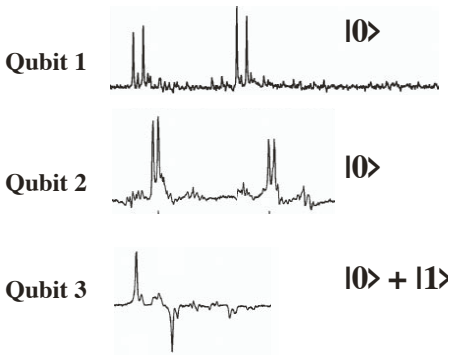

Figure 10.39: Spectra of the three result-qubits for the input $a = 11$ [213].

is $p = 4$. Since $n = 3$ qubits were used, the desired number r is given by (see Section 8.3.3) $r = 2^n/p = 2$. A classical calculation yields the greatest common divisor of $11^{2/2} \pm 1$ and 15 as 3 and 5, and thus directly the prime factors of N.

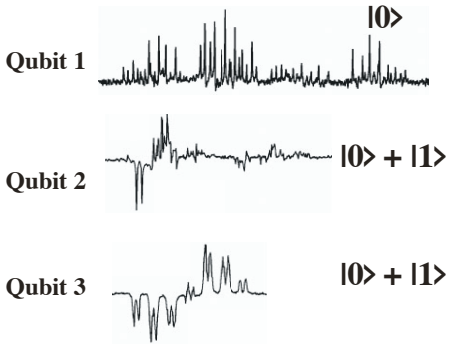

Figure 10.40: Spectra of the three result-qubits for the input $a = 7$ [213].

If the input $a = 7$ is used instead, the observed spectra shown in Figure 10.40 show that both qubits 2 and 3 are in superposition states, while qubit 1 is again in state $|0\rangle$. The possible results are therefore the states $|000\rangle = |0\rangle$, $|010\rangle = |2\rangle$, $|100\rangle = |4\rangle$, and $|110\rangle = |6\rangle$, indicating a period of 2. We conclude that $r = 8/2 = 4$ and $\gcd(7^{4/2} \pm 1, 15) = 3, 5$ as before. Obviously both trial values for a produce the expected result.

10.3.5 Decoherence

The experimental implementation of Shor's algorithm represented a milestone for quantum information processing, not because of the result itself, but because it provided the possibility of studying limitations to quantum information processing in a working example.

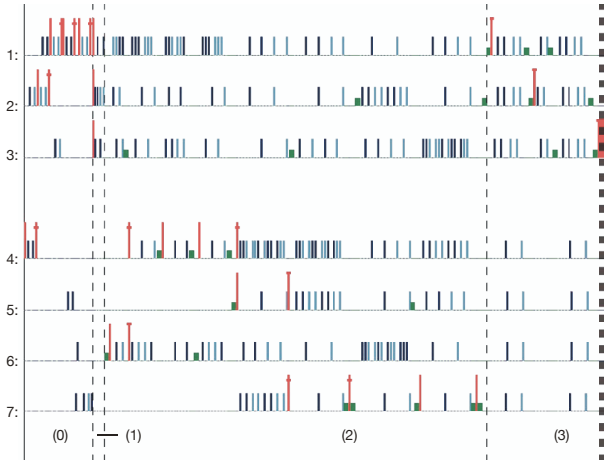

Figure 10.41: Pulse sequence used for the implementation [213].

The IBM group used some 300 RF pulses to implement the algorithm. Most of the pulses were used not for the processing itself, but to compensate for unwanted effects, such as spin-spin couplings and magnetic field inhomogeneity.

The overall sequence lasted almost 1 second, which is longer than some of the relevant relaxation times (=decoherence times). This caused a significant loss of information and therefore deviations of the experimental measurements from the idealized behavior. Vandersypen *et al.* analyzed these deviations with a model for the relevant decoherence processes and found that they could explain most of the differences with their model.

10.4 Spin chains

As discussed in a separate lecture by J. Stolze, linear chains of coupled spins represent interesting model systems that allow one to study various aspects of quantum mechanical dynamics. Such systems are also useful test systems for various quantum algorithms and have therefore often been implemented by NMR quantum computers. A few examples of such implementations will be discussed here.

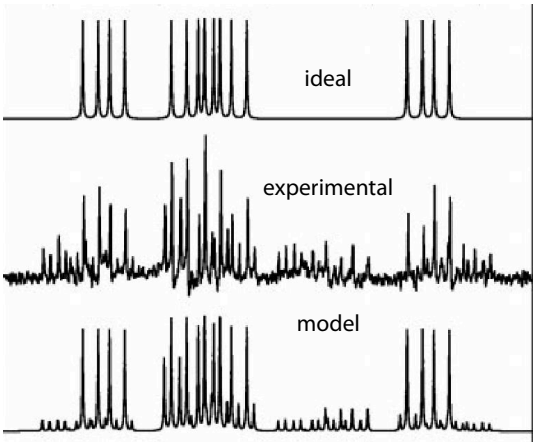

Figure 10.42: Comparison of the ideal spectrum, the experimental and the result of a model calculation including the effect of decoherence processes [213].

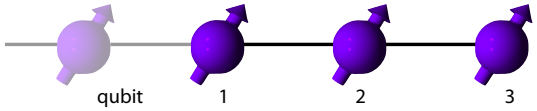

Figure 10.43: Model of a spin chain.

10.4.1 Quantum state transfer

Quantum-state transfer (QST), i.e., the transfer of an arbitrary quantum state $\alpha|0\rangle + \beta|1\rangle$ from one qubit to another, is an important element in quantum computation and quantum communication. The most direct method to implement QST is based on SWAP operations. For a pair of coupled qubits, this is an elementary operation. For a chain of qubits with only nearest-neighbor interactions, a series of SWAP operations between neighboring qubits can be used until the quantum state arrives at the target qubit.

For specific systems, it is possible to transfer quantum information without applying gate operations, but instead relying on a static coupling network. If a quantum state is prepared at one end of such a chain, it will travel through the chain without requiring control operations. The main difficulty with this approach is the required precision with which the couplings have to be realized in order to generate a transfer with high fidelity. This requirement can be relaxed significantly, without compromising the fidelity of the transfer, by applying gate operations to the receiving end of the spin chain that effects the transfer. The capability for applying such gate operations is not an additional requirement, since such operations are required anyway if the spin chain is to be used for communication between quantum registers. This gate accumulates any amplitude of the initial state that is transferred along the chain. The protocol allows one, in principle, to obtain unit fidelity for the transfer, even if the couplings along the chain have arbitrary fluctuations, as long as a finite amplitude reaches the end of the chain. Obtaining a large transfer amplitude requires multiple iterations, each of which includes the evolution of the spin chain and the two-qubit gate operation. This procedure is known as iterative quantum state transfer [224].

For the spin chain without the end-qubit, we assume that the Hamiltonian is

$$\mathcal{H} = 2\pi \sum_{i,k} J_{ik} \left(S_x^{(i)} S_x^{(k)} + S_y^{(i)} S_y^{(k)} \right) .$$

In addition, we assume that the coupling between the end-qubit and the one next to it is switchable, i.e. it can be turned on and off. With these requirements, the protocol can be implemented.

Figure 10.44 shows how the amplitude that can be transferred to the final spin increases with each iteration. The different curves correspond to different coupling strengths. While the differences lead to a different transfer speed, they all increase towards full transfer with increasing number of iterations. Apart from the increased transfer, decoherence leads to an overall loss of signal amplitude and therefore to an optimal number of iterations.

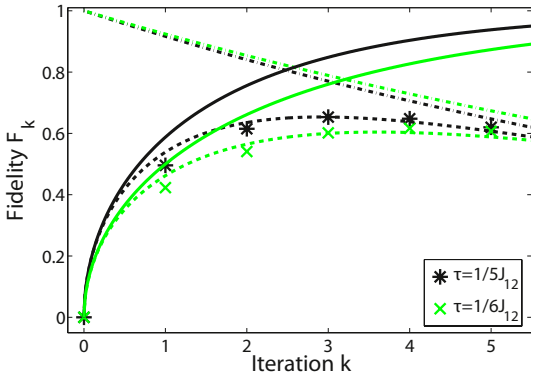

Figure 10.44: Achieved transfer of the quantum state with increasing number of iterations.

10.4.2 Two- and three-body interactions

In most cases, spins (and also most other qubits) are subject to pairwise interactions. This means that the Hamiltonian contains only terms depending on single qubits and pairs, such as

$$\mathcal{H} = a\, S_z^1 + b\, S_x^2 S_y^3 + \dots .$$

Three-body interactions, such as

$$\mathcal{H}_3 = c\, S_z^1 S_x^2 S_y^3$$

are only rarely found in nature, but they can arise as effective interactions or they can be generated in quantum simulators. They can significantly change the static as well as the dynamical properties. They can, e.g., speed up the transfer of quantum states along a spin chain. As a specific example, we consider a minimal model a spin chain consisting of three members coupled by the Hamiltonian

$$\begin{aligned} \mathcal{H} &= S_x^1 S_x^2 + S_y^1 S_y^2 + S_x^2 S_x^3 + S_y^2 S_y^3 \\ &+ \frac{\lambda}{2} \left(S_x^1 S_z^2 S_y^3 + S_y^1 S_z^2 S_x^3 \right) . \end{aligned} \tag{10.18}$$

Here, the parameter λ represents the strength of the three-body interaction, relative to the two-body interaction.

With this Hamiltonian, we can consider different types of state transfer: we can initialize the left-hand spin and transfer its state to the right-hand side, or we can do the opposite, or we can let the state be transferred from left to right and back. All three types of transfer can take place, with or without the three-body interaction. In the absence of the three-body interaction ($\lambda = 0$) the system is symmetric, i.e. the transfers left→right and right→left run at the same speed, and the round-trip time is twice the time of the individual transfer.

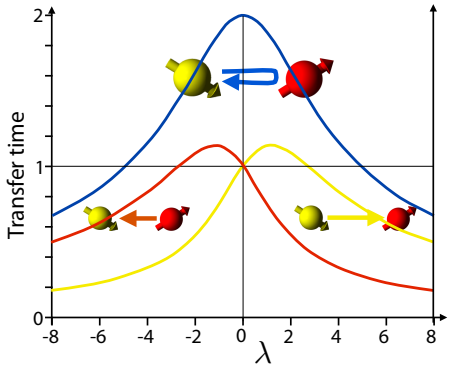

Figure 10.45: Duration of QST as a function of the coupling strength of the 3-qubit interaction.

This is no longer the case in the presence of the three-body coupling, as shown in fig. 10.45.

The experimental implementation can be simplified by decomposing the Hamiltonian into two commuting parts. Each of these two parts generates a harmonic evolution and the superposition is still quite simple [212].

Figure 10.46 shows the pulse sequence that was used to implement half of the Hamiltonian [212].

Figure 10.47 shows the observed transfer of the state S_x^1 along the chain. When it reaches the final qubit, the total system is in the state $S_z^1 S_z^2 S_x^3$. This final state is reached for all values of λ, but fastest when $\lambda = 4$ and slowest for $\lambda = 0$.

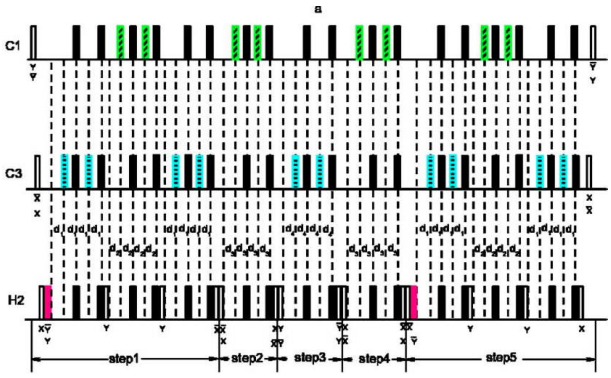

Figure 10.46: Pulse sequence used to implement part of the Hamiltonian of eq. (10.18).

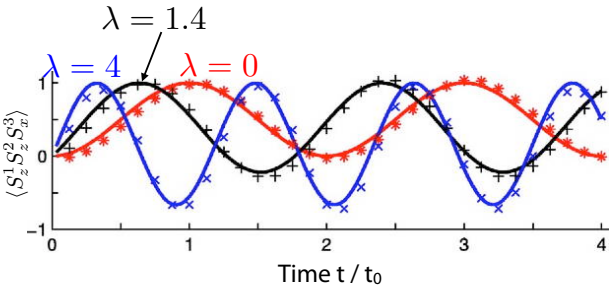

Figure 10.47: Predicted (curves) and measured (data points) transfer of the state along the spin chain for three different strengths of the 3-body interaction.

Problems and Exercises

1. Derive the equations of motion (10.2) from the Schrödinger equation using the Hamiltonian (10.1).
2. Show that the effective Hamiltonian (10.9) is the correct Hamiltonian $\mathcal{H}^r$ in the rotating frame by transforming the solution of the Schrödinger equation from the laboratory frame to the rotating frame and demanding that the transformed operator $\mathcal{H}^r$ generates the correct time evolution.
3. Verify equations (10.14) and (10.15) using the properties of spin matrices.