$$\mathbf{H}^{-\frac{1}{2}}\mathbf{Z}\mathbf{H}^{\frac{1}{2}} = \frac{1}{2}(\mathbf{X}+\mathbf{Z}+\sqrt{2}\mathbf{Y}),$$

and finally

$$\mathbf{H}^{-\frac{1}{2}}\mathbf{Z}^{-\frac{1}{4}}\mathbf{X}^{\frac{1}{4}}\mathbf{H}^{\frac{1}{2}} = \cos^2\left(\frac{\pi}{8}\right)\mathbf{1} - i\sin\left(\frac{\pi}{8}\right)\vec{m}\cdot\vec{\sigma}$$

with

$$\vec{m} = \left(-\frac{1}{\sqrt{2}}\sin\left(\frac{\pi}{8}\right), \sqrt{2}\cos\left(\frac{\pi}{8}\right), \frac{1}{\sqrt{2}}\sin\left(\frac{\pi}{8}\right)\right)$$

from which we see that $\vec{m}^2 = \vec{n}^2$ and $\vec{m}\cdot\vec{n} = 0$. This is again a rotation by the same angle α as before, about an axis orthogonal to the previous axis $\hat{n}$.
The construction in [3] uses the rotations $\mathbf{X}^{\frac{1}{4}}\mathbf{Z}^{\frac{1}{4}}$ and $\mathbf{H}\mathbf{X}^{\frac{1}{4}}\mathbf{Z}^{\frac{1}{4}}\mathbf{H} = \mathbf{Z}^{\frac{1}{4}}\mathbf{X}^{\frac{1}{4}}$, which are quite similar to those used above, but unfortunately the axes of rotation are not orthogonal to each other but only at an angle of 32.65°. In that case the simple Z-Y-Z decomposition of an arbitrary rotation into three factors is not possible, but a decomposition into more than three factors still is.

VI. IMPLEMENTATIONS: HOW TO BUILD QUANTUM COMPUTERS

A. General requirements

Any implementation has to define a quantum mechanical system that provides the quantum register containing N qubits. For a "useful" quantum computer, N should be at least 400, better 1000; limitations on the number N of identifiable qubits will therefore be an important consideration.

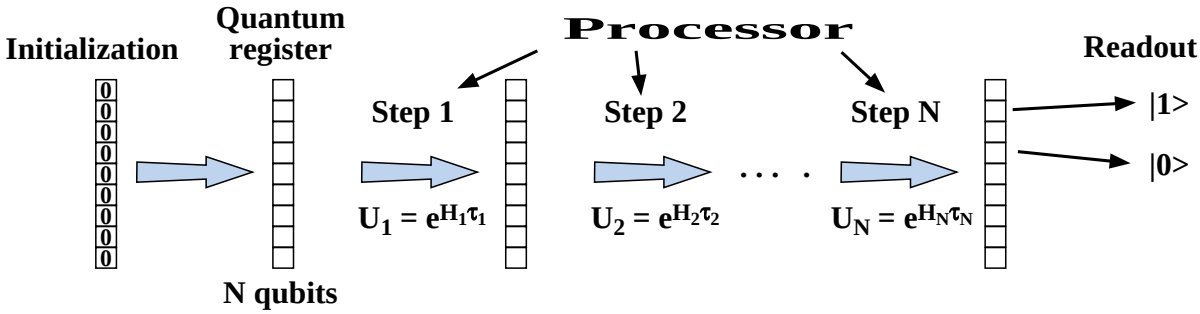

These qubits must be initialized into a well defined state, typically into a ground state $|0>$. This is necessarily a dissipative process. The implementation must then provide a mechanism for applying computational steps to the quantum register. Each of these steps will be implemented by a unitary operation defined by a Hamiltonian H_i that is applied for a time τ_i. After the last processing step, the resulting state of the quantum register must be determined, i.e. the result of the computation must be read out. This would typically correspond to an ideal quantum mechanical measurement, i.e. the projection onto the eigenstate of the corresponding observable. Readout has to be done on each qubit separately.

Today, a single implementation of a quantum computer exists, which uses nuclear spin states of molecules in solution, i.e. liquid-state NMR. Details of this implementation will be discussed in one of the following subsections. In addition, there is a long list of proposed implementations, which includes, as qubits, nuclear and electronic spins, photons, trapped ions, as well as various states of quantum confined structures, mostly in semiconductors, and superconducting devices such as Josephson junctions.

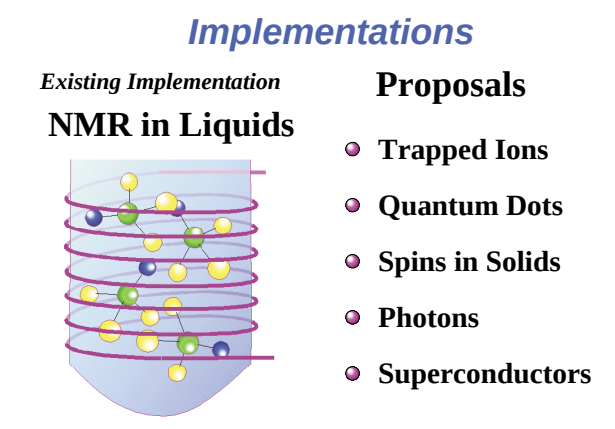

B. DiVincenzo's five criteria

DiVincenzo [30] gives five requirements that a quantum computer must fulfill:
1) A scalable physical system with well characterized qubits. An implementation or embodiment of qubits corresponds to a physical system that has at least two energy levels that can be identified with the two log-

ical states $|0>$ and $|1>$. The system will always be characterized by a number of parameters. They determine the internal state of the qubit, i.e. the energies of the states $|0>$ and $|1>$, but also of any additional states that may exist in the system being considered.

One generally has a large degree of freedom when choosing the identification of states with logical qubits. There are a few caveats, however: In particular, it must be possible to create arbitrary superpositions of these states. This is usually possible unless there is a selection rule that prevents it. As an example, we consider two neighbouring quantum dots, where an electron can tunnel from one dot to the other. It is then possible to identify state $|0>$ with the electron being in dots 1 and state $|1>$ with the electron being in dot 2. However, it is not possible to identify a qubit with each quantum dot, e.g. with the assignment that the presence of an electron corresponds to $|1>$, while its absence would correspond to $|0>$. The superposition of these two states would then correspond to a superposition between states with different particle numbers, which is usually impossible to achieve.

Besides the internal Hamiltonian, the interaction of the system with external fields is also important. External fields are generally required to apply logical operations to the qubits. Finally, the couplings between different qubits must be described, as they are needed for logical operations.

DiVincenzo's second requirement is

2) Initialization into a well defined state. Typically, this state is chosen equal to the logical state $|0>$ for all qubits. In principle thermal relaxation may achieve this, provided that the thermal energy k_BT is small compared to the energy level splitting between states $|0>$ and $|1>$. This may be a slow process in many systems, in particular in the spin systems, where the relaxation times are long. This is not critical for the computation process itself; however, future quantum computers will require error correction schemes. All error correction schemes known to date require an input in the form of freshly initialized qubits. These error correction qubits must be intialized at a rate that is large compared to the dephasing rate. This requirement cannot be fulfilled by thermal relaxtion, where the dephasing processes are always faster than the spin-lattice relaxation. The requirement can be met, however, in many optical systems, such as ion traps, where the initialization goes through optical excitation, which may proceed over a time of the order of nanoseconds.

3) Long decoherence times. The information in the quantum register is subject to decay, due to the interaction with external degrees of freedom. The computation must therefore be completed before this decay has significantly degraded the information. The relevant figure of merit is the number of gate operations that can be completed before a decoherence time. The effect of decoherence can partly be eliminated by quantum error correction. However, error correction also increases the duration of the computation and introduces additional errors. It has been shown that computation can proceed for an arbitrary duration if quantum error correction is used and error-free computation without error correction is possible for a critical minimum duration that is of the order of some tens of thousands of gate operations.

4) A universal set of quantum gates. The unitary operations that act as gates on the qubits must be implemented by Hamiltonians that act on the system for a specified time. Generating the single qubits Hamiltonians is in general relatively straightforward: typically they correspond to external fields acting on the qubits for a spcified duration. Much more complicated are the 2-qubit operations, which cannot be implemented by external fields alone. They involve interactions between the qubits, and in many cases these interactions cannot be switched on and off. Often one has to use static interactions and eliminate the unwanted ones by a procedure called refocusing. This will be discussed in more detail in the section on NMR. Every experimentally realisable gate will include imperfections, i.e. deviations from the ideal behavior. This has the effect of degrading the information in the quantum register and is therefore similar to an additional source of decoherence. Consequently, these errors can also be eliminated by error correction schemes, provided they are small enough.

5) A qubit-selective readout. Such a readout represents a measurement in the quantum mechanical sense. An ideal quantum mechanical measurement collapses the state ψ into an eigenstate ϕ_i of the observable and returns the eigenvalue λ_i of the corresponding state with probability $|c_i|^2$, where c_i is the expansion coefficient of the state $\psi = \sum c_i\phi_i$. Real measurements deviate from this. In many realistic systems, measurements attempts will return no result instead, e.g. when one tries to measure the state of a qubit by scattering a photon from it. If the photon is not scattered, this is not important, one just repeats the attempt. If the photon is scattered but not detected, this is more critical: In this case, an interaction of the qubit with an external system (the photon) has changed the state of the qubit, and a repetition of the measurements may produce a different result.

Several strategies are possible to circumvent this problem: one can try to use a QND (=quantum nondemolition measurement). Such a measurement arranges for the unavoidable influence that the measurement must have on the qubit to be such that it does not affect later measurements of the same variable. Not all variables can be measured this way, but in most cases it should be possible to arrange the system in such a way that QND measurements can be used at least in principle.

Another possibility is to read out not the qubit itself, but a copy of it. If the measurement is not successful, or to check the validity of the measurement result, one can then make an additional copy and read that out. Such a procedure could be repeated many times to achieve very reliable readout even with very unreliable single measurements. The critical part here

is the copy operation, which must be reliable. As we have stressed before, it is not possible to clone a quantum mechanical state, i.e. to make a perfect copy. However, copying just the information of a quantum mechanical state that is relevant for the readout of a specific variable is perfectly possible (in principle!) and can be repeated arbitrarily often.

C. NMR in Liquids

The first implementation of a quantum computer that has been realized - and still the only one in existence - is nuclear magnetic resonance in liquids.

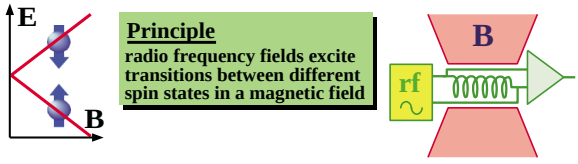

Magnetic resonance is based on the Zeeman effect, which lifts the degeneracy of spin states. For a spin $I = 1/2$, the splitting of the two energy levels is proportional to the magnetic field strength. An alternating magnetic field perpendicular to the static magnetic field can induce transitions between these spin states. The relevant frequency is in the radio frequency range for nuclear spins (10-1000 MHz in fields of 1-20 T), and in the microwave range (10-300 GHz) in fields of 0.1-20 T).

1. System and Hamiltonian

The strongest interaction of nuclear spins is usually the Zeeman interaction with the static magnetic field. It can be described by the Hamiltonian $H_z = -\gamma \mathbf{I} \cdot \mathbf{B}$, where γ is the gyromagnetic ratio of spin I. The usual convention is to orient the z-axis along the static magnetic field. The Hamiltonian then becomes

$$H_z = -\gamma I_z B_0 = -\omega_0 I_z,$$

where B_0 is the strength of the magnetic field and ω_0 the Larmor frequency.
In heteronuclear systems, one often uses I for one species (e.g. ^1H) and S for the other (e.g. ^{13}C). The magnetic field is always treated classically.
Given the commutation relations for angular momentum, we can write the equation of motion (Schrödinger) as

$$dI_x/dt = -\omega_0 I_y \qquad dI_y/dt = \omega_0 I_x \qquad dI_z/dt = 0.$$

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

$$< I_x > (t) = I_{xy}(0) cos(\omega_0 t - \phi)$$

$$< I_y > (t) = I_{xy}(0) sin(\omega_0 t - \phi)$$

$$< I_z > (t) = I_z(0)$$

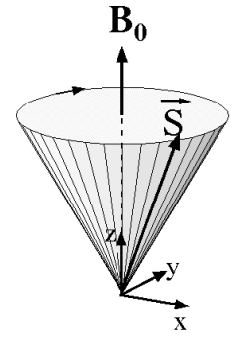

The equation is called the Bloch equation, after one of the discoverers of NMR, who also wrote the theory for it. It can also be derived classically and has applications to many two level systems besides NMR (R.P. Feynman, F.L. Vernon, and R.W. Hellwarth, 'Geometrical representation of the Schrödinger equation for solving maser problems', J. Appl. Phys. 28, 49-52 (1957).).

2. Radiofrequency Irraditation and Rotating Frame

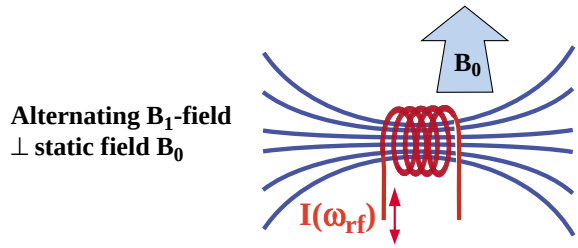

A radio frequency field is applied to the system for inducing transitions between the different spin states. It generates an alternating magnetic field perpendicular to the static field. This alternating magnetic field is best described a s superposition of two fields rotating in opposite directions.
It turns out that to an excellent approximation it is sufficient to consider the effect of that component which rotates in the same direction as the freely precessing spin. It is therefore a convenient fiction to assume that it generates a circularly polarized rf field. The resulting dynamics are then best analyzed in a coordinate system that rotates around the static magnetic field at the radio frequency.
The two coordinate systems are related by

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

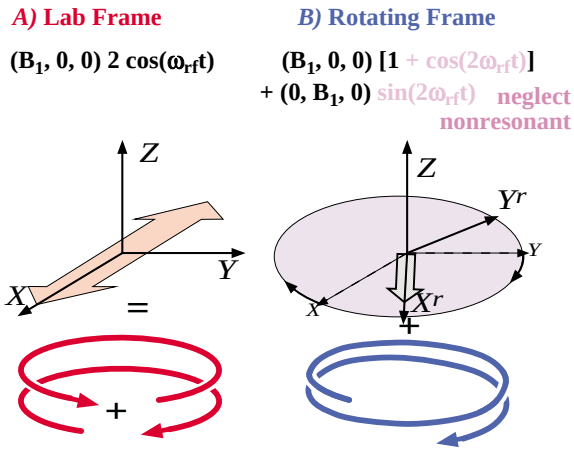

The quantum mechanical states can be transformed from the laboratory frame to the rotaing frame with the operator

$$U(t) = e^{i\omega_L t S_z}.$$

The same operator also transforms the Hamiltonian:

$$H^r = U^{-1} H^{lab} U + i U^{-1} \dot{U}$$

The effective Hamiltonian for this coordinate system therefore contains a reduced magnetic field; the Larmor frequency is reduced to $\Delta\omega_0 = \omega_0 - \omega$.
The rf field appears static, e.g. along the x-axis, so that we obtain a total field

$$\omega_{eff} = (\omega_1, 0, \Delta\omega_0).$$

The resulting evolution of the system is a precession around the direction of the effective field.

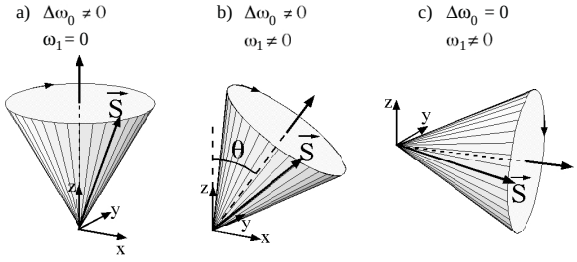

We may consider three specific examples of motion under the Bloch equations. In the absence of rf irradiation ($\omega_1 = 0$), the effective field is aligned along the z-axis and the precession is the same as in the laboratory frame, except that the precession fre-quency is lower.
In the case of resonant 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, the effective field lies along a direction in the xz-plane. Other directions (e.g. along the y-axis) can be chosen by adjusting the phase of the rf field. The rotation axis can therefore be oriented in any arbitrary direction. The angle of rotation $\alpha = \omega_{eff} t$, which is called the flip angle, is given by the product of the effective field strength and the pulse duration t. RF Pulses are a convenient possibility to implement arbitrary rotations about arbitrary axes!

3. NMR Signals

NMR signals are obtained in the time domain, as the response of the system to a rf pulse. We consider first the simplest case, where the system consists of a single spin I = 1/2. The Hamiltonian is (for $\hbar = 1$) $H = -\omega_0 I_z$.
Since NMR systems are always mixed states, we must use a density operator analysis to calculate the signal. The thermal equilibrium density operator is

$$\rho_{eq} \propto exp(-H/kT) \approx 1 - \frac{H}{kT}$$

where the approximate form, derived for the high temperature limit

$$\Delta E = \omega_0 I_z << kT$$

is almost always valid: under typical experimental conditions, it is of the order of 10^{-5}. We have therefore

$$\rho_{eq} = 1 + \frac{\hbar\omega_0}{kT} I_z.$$

In the simplest case, one applies a RF pulse that rotates the spin through an angle of $\pi/2$ into the xy-plane.

$$\rho(0+) = 1 + \frac{\hbar\omega_0}{kT} I_x.$$

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

$$\rho(t) = e^{-iHt}\rho(0)e^{iHt} = 1 + \frac{\hbar\omega_0}{kT}(I_x cos\omega_0 t + I_y sin\omega_0 t).$$

Detection of the signal should not be treated as a quantum mechanical measurement process. There is no reduction of a wavepacket, and the system is virtually unaffected by the measurement. Rather than projecting onto an eigenstate, one measures the expectation value of a specific observable. This is of course closely related to the fact that we are working with an ensemble of spins rather than with individual systems.
Observation of the precessing magnetization is achieved through the Faraday effect. The precessing magnetization changes the flux through the rf coil, thus inducing a voltage proportional to

$$s(t) \propto \frac{d}{dt}\Phi(t) = \omega_0 < F_x >$$

$$= \omega_o \sum_i < I_{ix} >= \frac{\hbar\omega_0^2}{2kT} cos\omega_0 t$$

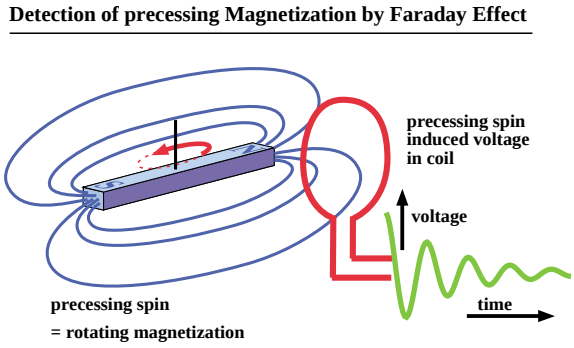

Damping effects, which are not discussed here, cause a decay of the signal,

$$s(t) \propto \frac{\hbar\omega_0^2}{2kT} cos\omega_0 t e^{-t/T_2}$$

For an analysis of the signal one usually considers the Fourier transform. For an FID decaying exponentially with time constant T_2, the spectrum becomes

$$s(\omega) = \sqrt{\frac{1}{2\pi}} \frac{\hbar\omega_0^2}{2kT} \frac{T_2}{1+(\omega-\omega_0)^2 T_2^2},$$

i.e. a Lorentzian with a HWHH $\frac{1}{T_2}$ centered at the Larmor frequency ω_0. The main advantage of the Fourier transform is that it allows one to distinguish different transitions: two distinct transitions usually have different Larmor frequencies ω_{ij}

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

The corresponding resonance lines are therefore separated in frequency space, while the time domain signals overlap. 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, density operator and observable are identical, $\rho(0) = A = I_x$, and amplitudes A_{ij} become

$$A_{ij} \propto |(I_x)_{ij}|^2.$$

4. Coding in Quantum Registers

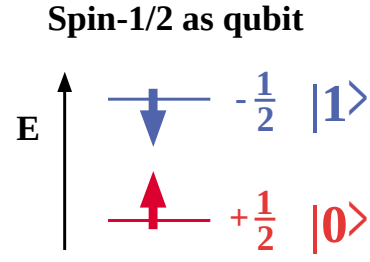

In NMR quantum computing, a single qubit is represented by a spin 1/2. The general terminology has become to use the +1/2 state to represent a logical 0, while the -1/2 state represents a logical 1. To construct a quantum register, one needs several distinguishable qubits, which can be addressed individually.

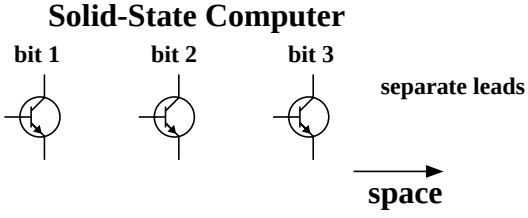

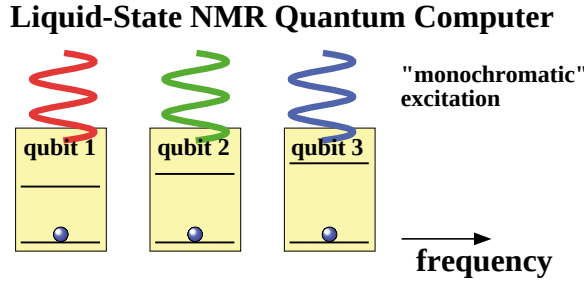

In conventional computers, addressing of qubits (i.e. choosing a qubit for processing) is achieved by fixed wires.

For molecules in solution, this approach is clearly not feasible. Nevertheless, it is possible to address qubits selectively: Since the qubit gates are applied with resonant radio frequency fields, they are only effective when the rf frequency is close to the Larmor frequency of the spin.

Spins whose Larmor frequency differs from the frequency of the radio frequency pulse are not affected by the pulse to a first approximation. The width of the frequency range is of the order of the Rabi frequency, i.e. inversely proportional to the length of the rf pulse.

The difference in Larmor frequencies for different qubits is associated with their gyromagnetic ratio (for heteronuclear spin systems) or with the chemical shift (for homonuclear spin systems). In contrast to conventional computers, where etching localizes different

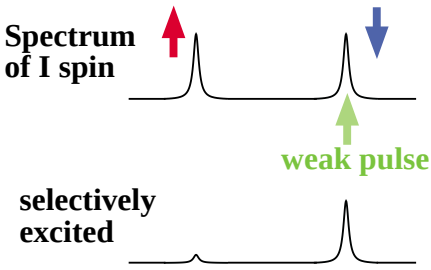

qubits, this may be considered a bottom-up approach, where the molecular design determines the location of the qubit in frequency space.
Nuclei of a given isotopic species have the same gyromagnetic ratio and therefore, to first order, the same resonance frequency. Nevertheless, the spins of a given species show significant differences in resonance frequency. These variations can be traced to local variations of the magnetic field strength that are called chemical shifts. Chemical shifts are determined by the structure of the electrons near the nucleus, but these mechanisms are not of interest here. We only discuss here how chemical shifts can be used to address spins as qubits.
Any NMR pulse sequence must therefore be adjusted to the available range of chemical shifts. The available chemical shift range depends on the isotope under examination. In the case of protons (^1H), this range is of the order of 10 ppm. For ^{13}C, it is about 200 ppm, and similar for ^{15}N. For a typical ^1H NMR frequency, the available frequency range is therefore of the order of 6 kHz, for ^{13}C in the same field 30 kHz.

5. *Coupled Spin Systems*

Implementation of quantum algorithms requires two-qubit gates, which can be implemented by using couplings between qubits / spins.
Such couplings are naturally present in nuclear spin systems and exploited also in NMR spectroscopy. There are two main types of couplings, which are called scalar / indirect / J-couplings, and direct or dipolar couplings. The latter arise from the magnetic dipolar field generated by one spin and felt by the other, while the first is mediated through the electrons and does not depend on orientation of the molecule. The difference in orientation dependence is responsible for the fact that in isotropic liquids, the direct dipole-dipole coupling is averaged to zero, so only the scalar J-coupling is observed in the spectrum. 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

$$H_{AX} = dA_z X_z.$$

The total Hamiltonian is then

$$H = H_z + H_{AX} = -\omega_A A_z - \omega_X X_z + dA_z X_z =$$

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

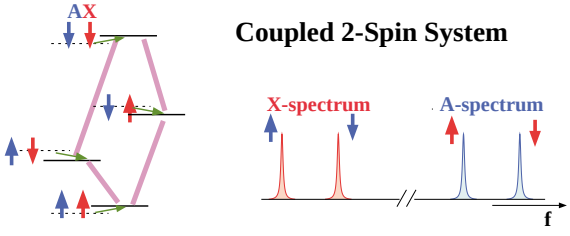

For states with parallel spin orientations, the coupling shifts the energy up, while for states with opposite spin orientation, the energy is shifted down. Transitions are possible between the states $\alpha\alpha \leftrightarrow \alpha\beta$, $\alpha\alpha \leftrightarrow \beta\alpha$, $\alpha\beta \leftrightarrow \beta\beta$, $\beta\alpha \leftrightarrow \beta\beta$. The transition frequencies are

$$\omega_{12} = \omega_{\alpha\alpha\leftrightarrow\alpha\beta} = \omega_X - d/2;$$

$$\omega_{13} = \omega_{\alpha\alpha\leftrightarrow\beta\alpha} = \omega_A - d/2;$$

$$\omega_{24} = \omega_{\alpha\beta\leftrightarrow\beta\beta} = \omega_A + d/2;$$

$$\omega_{34} = \omega_{\beta\alpha\leftrightarrow\beta\beta} = \omega_X + d/2;$$

The spectrum consists of four lines, each of which is associated with a transition of one spin and labeled by the state of the second spin.
To actually calculate the spectrum, we use a density operator analysis. Like in the case of a single spin, we

assume excitation from the equilibrium density operator by an RF pulse. If the two spins are from the same spin species, one can write the density operator after the pulse as

$$\rho(0) \propto A_x + X_x = F_x = \frac{1}{2}\begin{pmatrix} & 1 & 1 & \\ 1 & & & 1 \\ 1 & & & 1 \\ & 1 & 1 & \end{pmatrix}$$

Assuming that this is also the observable, the signal becomes

$$s(t) = Tr\left\{\rho(t)F_x\right\} = Tr\left\{F_x(t)F_x(0)\right\},$$

i.e. the autocorrelation function of the operator F_x. Since we are in the eigenbase of the Hamiltonian, the time dependence is

$$\rho_{lm}(t) = \rho_{lm}(0)e^{i(E_m - E_l)t}.$$

or

$$\rho(t) = \begin{pmatrix} & e^{-\omega_{12}t} & e^{-\omega_{13}t} & \\ e^{\omega_{12}t} & & & e^{-\omega_{24}t} \\ e^{\omega_{13}t} & & & e^{-\omega_{34}t} \\ & e^{\omega_{24}t} & e^{\omega_{34}t} & \end{pmatrix}$$

Including the effect of relaxation, the FID becomes

$$s(t) = Tr\left\{\rho(t)Fx\right\}$$

$$= \left(cos\omega_{12}t + cos\omega_{13}t + cos\omega_{24}t + cos\omega_{34}t\right)e^{-t/T_2}.$$

Using Fourier transformation to obtain the spectrum, we find four Lorentzian-shaped lines at the position of the four transition frequencies. Each of those lines can therefore be associated with a specific off-diagonal density operator element.

6. *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. For liquid state NMR, this is not possible, since the thermal energy is always much higher than the energy separation between the two spin states. Spin states are therefore not pure states, but mixed states, which can not be converted into the pure states required to initiate the quantum register. The success of NMR quantum computing can largely be attributed to the finding that many of these algorithms can also be applied to states that are the sum of a pure state and the unity operator, $\rho_{pp} = 1 + \alpha\rho_p$, where ρ_{pp} is a pseudo-pure state, while ρ_p is a pure state. If the unity operator does not contribute to the signal, the behavior of such a system is exactly equal to that of a pure state.
The coefficient α is largely determined by the polarization of the spin system. Obviously, a single spin is always in a pseudo-pure state. In coupled spin systems, however, the thermal equilibrium states are not even pseudo-pure. There are a number of procedures for preparing pseudo-pure states, which are referred to as "spatial labeling", "temporal labeling" and "logical labeling". Temporal labeling is perhaps easiest to explain, using the example of two coupled spins. In equilibrium, the populations of these four states are

$$\alpha\alpha : 1/4 + \epsilon \qquad \alpha\beta, \beta\alpha : 1/4 \qquad \beta\beta : 1/4 - \epsilon.$$

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

$$\frac{1}{4}(1111) + \epsilon(1 - \frac{1}{3} - \frac{1}{3} - \frac{1}{3})$$

$$= \frac{1}{4}(1111) - \frac{\epsilon}{3}(1111) + \frac{4\epsilon}{3}(1000)$$

and therefore correspond to the sum of the unity operator (=the totally mixed state) and a pseudo pure state.
The well known disadvantages of this process is that one loses signal by destroying polarization. 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. In the case of logical labeling, additional (ancilla) spins are used to create pure states for specific ancilla spin configurations.

7. *One qubit gates*

Single qubit gates are implemented by rf pulses. In the rotating frame, a rf pulse can be represented by its propagator e^{-iHt}, where H is the Hamiltonian during the pulse and t the duration of the pulse. Depending on the phase of the rf, the propagator for a resonant pulse is $e^{-i\phi I_x}$ or $e^{-i\phi I_y}$. The flip angle is

$$\phi_\alpha = \omega_\alpha \tau, \qquad \alpha = x, y,$$

where τ is the duration of the pulse.
Looking at the matrix representation

$$e^{-i\phi I_x} = \begin{pmatrix} cos(\phi/2) & -isin(\phi/2) \\ -isin(\phi/2) & cos(\phi/2) \end{pmatrix}$$

$$e^{-i\phi I_y} = \begin{pmatrix} cos(\phi/2) & -sin(\phi/2) \\ sin(\phi/2) & cos(\phi/2) \end{pmatrix}$$

one notices immediately that the most important single qubit quantum logic gates have straightforward implementations. Combining these two generators, it is possible to implement any SU(2) operation.
An important example is the set of rotations around the z axis, which can be generated e.g. as

$$e^{-i\phi I_z} = \begin{pmatrix} e^{-i\phi/2} & \\ & e^{i\phi/2} \end{pmatrix}$$

$$= e^{-i\pi/2I_x} e^{-i\phi I_y} e^{i\pi/2I_x} = e^{-i\pi/2I_y} e^{i\phi I_x} e^{i\pi/2I_y}$$

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

$$NOT: \quad e^{-i\pi I_x} = \begin{pmatrix} & -i \\ -i & \end{pmatrix} = e^{-i\pi/2} \begin{pmatrix} & 1 \\ 1 & \end{pmatrix}.$$

This implementation of NOT thus differs from the usual representation by an overall phase of $-\pi/2$. 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 $|0>$ and $|1>$ should be an implementation of NOT. However, looking at the propagator for a π_y pulse,

$$e^{-i\pi I_y} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

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

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

can also be implemented by an rf pulse

$$\frac{i}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = e^{-i(\frac{\pi}{\sqrt{2}})(I_x + I_z)}$$

Physically this transformation can be achieved in a number of different ways: either by applying an off-resonant rf pulse, i.e. one where the z-component in the rotating frame does not vanish, or by a sequence of three rf pulses along the y, x and -y axes:

$$H = e^{-i\frac{\pi}{4}I_y} e^{-i\pi I_x} e^{i\frac{\pi}{4}I_y} = e^{-i\pi I_x} e^{i\frac{\pi}{2}I_y}.$$

this is an example that shows how rotations can be rotated. The third version is experimentally the simplest: a $\pi/2_y$-pulse is followed by a πx pulse. Such schemes of combining rotations are very useful and flexible. One often uses the possibility to modify propagators, rather than states by "sandwiching" them between two pulses. For many purposes the Hadamard gate can be replaced by the pseudo-Hadamard gate

$$h = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix} = e^{i\pi/2I_y}$$

and its inverse

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

i.e. by $\pi/2$ rf pulses around the y axis.

8. 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 somewhat different ways of implementing such gates.
One may be referred to as "soft pulses", the other as "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. One applies rf to that component of the I spin spectrum, which is associated with the S-spin in a particular state. This results e.g. in the inversion of spin I conditional to the spin S being in the $|1>$ 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. It has the disadvantage, however, that it requires long pulses, thus causing excess decoherence.

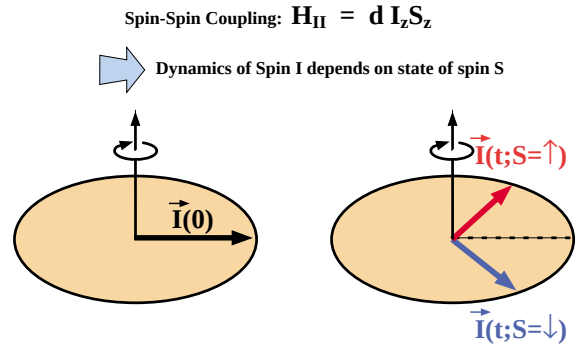

The second approach can also be understood in terms of a vector diagram. We consider a spin A coupled to a control spin X by the interaction dA_zX_z. As shown above, the resulting spectrum has two resonance lines in the A-spectrum, which can be labeled by the states $|\uparrow>$ and $|\downarrow>$ of the X spin. We will assume that pulses can be applied to the A and X spin separately. Starting from the state $|00>$, a $e^{-i\pi/2A_y}$ rf pulse creates a superposition state

$$\Psi(0) = \frac{1}{\sqrt{2}}(|0> + |1>) \otimes |0> .$$

Free precession converts it into a state

$$\Psi(t) = \frac{1}{\sqrt{2}}(|0> e^{-idt/4} + |1> e^{idt/4}) \otimes |0>,$$

where we have assumed that the rf frequency is centered on the spectrum of the A spin, i.e. between the two satellite transitions. After a time $t = \frac{1}{2d}$, the spin has reached a state

$$\Psi(\frac{1}{2d}) = \frac{1}{2}((1-i)|0> + (1+i)|1>) \otimes |0>,$$

If a $e^{-i\pi/2A_x}$ pulse is applied at this time, it returns the spin into its original state $|00>$. 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. It then precesses by 90 degrees to the y-axis, and the subsequent x-puls flips it back to the z-axis.
If we apply the same sequence of pulses to the state $|01>$, the free precession occurs with opposite sign

$$\rho(t) = (|0> e^{idt/4} + |1> e^{-idt/4}) \otimes |1>,$$

and the second pulse rotates the spin to the negative, rather than the positive z-axis. As can be easily checked, the sequence of two pulses with free precession is therefore equivalent to a controlled NOT operation

$$e^{i\frac{\pi}{2}X_y} e^{-i(\frac{\pi}{2}X_z + \frac{\pi}{2}A_z - \pi A_z X_z)} e^{-i\frac{\pi}{2}X_y}$$

$$= (1+i) \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & & 1 \\ & & 1 & \end{pmatrix}.$$

The additional terms of X_z and A_z are for normalization of the relative phases. They can be implemented as composite pulses or as phase shifts.

9. *Three and more Qubits*

Three qubit gates like the Toffoli gate can be constructed in the same way as two-qubit gates. The double-controlled phase shift can be represented as shown. However, since there are no three-spin interactions in nature, these must be created artificially. This is still possible, using e.g. transformations like

$$e^{-i\beta B_y C_z} e^{-i\alpha A_z B_x} e^{i\beta B_y C_z} = e^{-i\gamma A_z B_z C_z}.$$

Alternatively, three- or N-spin gates may be generated using selective pulses. NMR spectroscopists have developed tools to implement arbitrary propagators. See, e.g [31, 32].

10. *Readout*

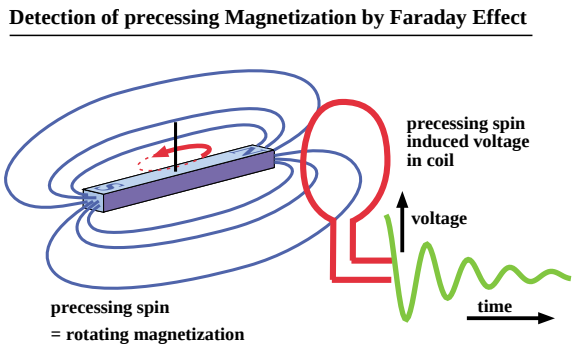

Detection in magnetic resonance is best described in a classical picture: the transverse components of the spin generate a macroscopic magnetisation that precesses around the static magnetic field. The magnetic flux through a coil oriented perpendicular to the field changes therefore sinusoidally. According to Faraday's law, such a temporal variation in the magnetic flux induces a voltage in the coil, which is recorded as the signal.
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. $< \psi_k|\psi_k >$, but rather the evolution of a coherence, i.e. $|\psi_j >< \psi_k|$.
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.
When a quantum algorithm requires the measurement of populations, it can be trivially modified to allow implementation on an NMR quantum computer: It is straightforward to convert populations into coherences that are directly proportional to the populations and are directly observable. Before we discuss these possibilities, we first check which quantities are directly observable. The state

$$\frac{1}{\sqrt{2}}(|0> + |1>) \otimes |0>$$

contains observable coherence (i.e. transverse magnetization), but not the state $|\uparrow>$ or the state $|\downarrow>$. While the instantaneous observable is $\sum_i I_x^i$ or $\sum_i S_y^i$, we need to take into account that a NMR measurement is not instantaneous; rather, one measures an FID signal over a total time of about a second. During this period, the Zeeman part of the Hamiltonian turns I_x and I_y into each other:

$$H_z : I_x \rightarrow I_y \rightarrow -I_x \rightarrow -I_y \rightarrow I_x.$$

Similarly, the coupling Hamiltonian causes evolutions of the type

$$H_{IS} : I_x \rightarrow e^{-i\phi I_z S_z} I_x e^{i\phi I_z S_z} I_x(t) = I_x cos\phi + I_y S_z sin\phi$$

This process is a direct mapping of the different transitions: In a weakly coupled two-spin system, every transition can be labeled with its spin and the state of the other spin. The coherence of a single transition can thus be written as

$$\rho_{12} + \rho_{21} = (1 + I_z) \otimes S_x.$$

Similarly the other S-spin transition corresponds to coherence

$$\rho_{34} + \rho_{43} = (1 - I_z) \otimes S_x.$$

The sum of the two transitions is therefore

$$\rho_{12} + \rho_{21} + \rho_{34} + \rho_{43} = 1 \otimes S_x.$$

Similarly we can take the difference, which corresponds to one of the two lines in emission:

$$\rho_{12} + \rho_{21} - \rho_{34} - \rho_{43} = I_z \otimes S_x,$$

which is the quantity that results from the coupled evolution.

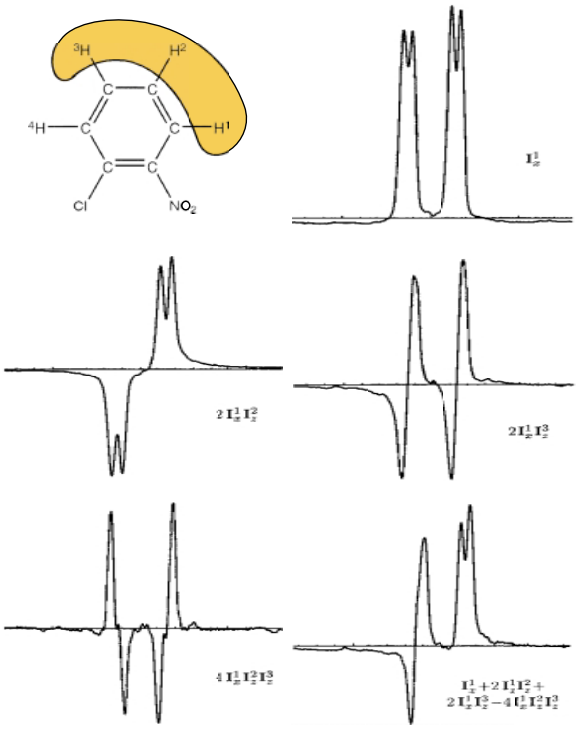

This scheme can easily be extended to more spins. Examples of such terms are given, e.g., in D.G. Cory, M.D. Price, and T.F. Havel, 'Nuclear magnetic resonance spectroscopy: An experimentally accessible paradigm for quantum computing', Physica D 120, 82-101 (1998). It should be noted, though, that the system shown here is not very suitable for quantum computation, since the lines are not well resolved. These resonance lines represent direct measures of the allowed transition, which are determined by the conditions that

$$\Delta m_i = \pm 1, \Delta m_j = 0, i \neq j$$

one spin changes its magnetic quantum number, all other spins remain. If other components of the density operator are to be measured, this is possible by converting them into observable magnetization.
As an example, we consider the AX system. As we have seen before, a $e^{-i\pi/2A_x}$ or $e^{-i\pi/2A_y}$ rf pulse turns the state $|00>$ into observable coherence of the A spin. Similarly a $e^{-i\pi/2X_x}$ or $e^{-i\pi/2X_y}$ rf pulse turns it into observable coherence of the X spin.
If the initial state is $|10>$, the same signals are observed on the A spin, but with opposite phase. On the X spin, this initial state generates a positive signal, but in the opposite transition. Similar observations are possible for the $|01>$ and $|11>$ states. It is possible to measure the complete density operator (I.L. Chuang, N. Gershenfeld, and M. Kubinec, 'Experimental Implementation of Fast Quantum Searching', Phys. Rev. Lett. 80, 3408-3411 (1998).) by "Tomography".

11. DiVincenzo's criteria

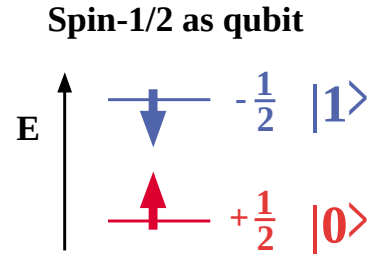

1) Qubits: The usual implementations use nuclear spins I=1/2 and identify $|0> = |+1/2>$ and $|1> = |-1/2>$. The qubits are well characterized in the sense that their energies are well known and the coupling to external fields occurs only through the Zeeman interaction. In the liquid state NMR experiments, 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. There are a few suggestions how this could be avoided, but present state of the art does not allow readout of individual spins with sufficient efficiency to make such quantum computers workable.
The qubits also must be addressed, both for performing operations and for readout. 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 crosstalk 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 fast operations of logical gates.
NMR systems fulfill the "qubit-identification" requirement quite well, but the existing ones fail the first part of DiVincenzo's criterion: they are not scalable. This

point will be discussed again when we address the initialization aspect.

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. This must be combined with the preparation of a pseudo-pure state, by averaging of the populations of the energetically higher states. While this procedures can be used for small spin systems, they are clearly not scalable for larger systems.

3) Long decoherence times. This works quite well for liquid state NMR, where decoherence times are of the order of 1 s. However, typical gate times 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.

D. Solid State NMR/EPR

While NMR has some strong points, it has a large weak point, which is the lack of scalability. To become scalable, pure states are necessary. Pure states correspond to very low spin temperature. While the spin temperature can be quite different from the lattice temperature, it will be very hard to create pure states at room temperature; most proposals for pure state spin systems are therefore geared towards helium temperature. The will therefore most likely be solid state systems.

One particularly interesting suggestions for solid state spin-based quantum computers is due to Kane [B.E. Kane, 'A silicon-based nuclear spin quantum computer', Nature 393, 133-137 (1998).]. He uses ^{31}P impurities in Si, the only $I = 1/2$ shallow (group V) donor in Si. The $Si:^{31}P$ system was exhaustively studied 40 years ago in the first electron-nuclear double-resonance experiments. At sufficiently low ^{31}P concentrations at temperature T = 1.5 K, the electron spin relaxation time is thousands of seconds and the ^{31}P nuclear spin relaxation time exceeds 10 hours. This system would therefore allow a large number of gate operations within a decoherence time.

Complete polarization of the electron spins in such a system can be achieved at a temperature of 100 mK in a magnetic field of 2 T.

Operation of these qubits would be identical to that of a liquid-state NMR system, i.e. in terms of radio frequency pulses. However, since all qubits see the same chemical environment, their resonance frequencies are identical. As a way of addressing them, it may be possible to use small electrodes, which shift the donor electrons closer or farther away from the nuclei.

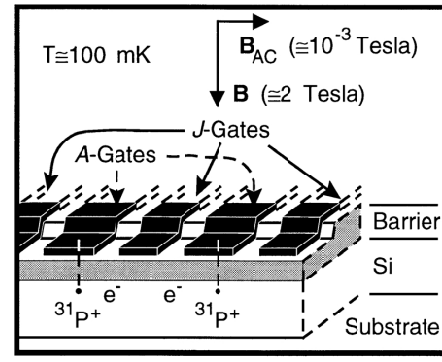

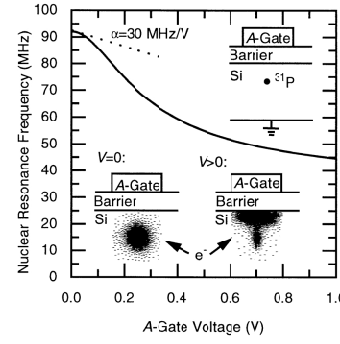

The hyperfine coupling between electrons and nuclei would depend on the position of the electrons. These so-called A-gates could therefore be used for addressing the individual qubits. Similarly J-gates would move electron density between the donor sites, thus inducing an indirect coupling between qubits and allowing the addressing of pairs of qubits.

Very recently, significant progress was reported for the control of interactions between electronic and nuclear spins in a quantum-Hall type system ([33], see also the "news and views" item [34] in the same (17 Jan 2002) issue of *Nature*).