

# NMR Quantum Computing

In NMR quantum computing, the computer is a single molecule of a chemical compound. The q-bits are the nuclear spins of the atoms in the molecule (or at least those which are spin-1/2). Quite a few isotopes have spin-1/2 nuclei, notably hydrogen and carbon-13.

A powerful magnetic field is introduced in the  $Z$  direction, resulting in a separation of energy between spin-up and spin-down. This separation is important for practical purposes.

Unlike other models we have considered, however, NMRQC is not done with just one system; rather, one has *many* (e.g.,  $10^{24}$ ) molecules, all performing the same calculation in parallel. And these molecules are not cooled; they are in the liquid state. In NMR, computation is done with a bulk sample at room temperature.

Perhaps surprisingly, despite being immersed in a bulk at room temperature, the nuclear spins can be treated as isolated to a very good approximation. This is because (a) the nuclear magnetic moments are very weak, which means they evolve on a very slow timescale, (b) the nuclei are surrounded by rapidly moving clouds of electrons which screen out other nearby charges, and (c) the molecules tumble rapidly, so that surrounding potentials are averaged out on a timescale short compared to that of the spin evolution time. Because of these effects, it is possible to treat the nuclear spins using the Schrödinger equation; a fact which chemists have been exploiting for years.

To analyze them we need to know their Hamiltonian to a good approximation.

The presence of the magnetic field produces a Hamiltonian of the form

$$\hat{H}_B = \sum_j E_j \hat{Z}_j,$$

where  $\hat{Z}_j$  is the  $Z$  component of spin for the  $j$ th nucleus.  $E_j$  will be roughly the same for all nuclei of the same chemical species, though the configuration of the molecule can produce small differences (*chemical shifts*).

There will also be a (much weaker) Hamiltonian due to the dipole-dipole interactions of the spins. This takes the form

$$\hat{H}_{SS} = \sum_{i,j} \alpha_{ij} \hat{\vec{\sigma}}_i \cdot \hat{\vec{\sigma}}_j,$$

where the interaction strength  $\alpha_{ij}$  between nuclei  $i$  and  $j$  will depend on their chemical species and distance from each other in the molecule. Since these are basically fixed,  $\alpha_{ij}$  can be treated as a constant matrix to be experimentally determined for each chemical.

The full Hamiltonian is

$$\hat{H} = \hat{H}_B + \hat{H}_{SS}.$$

Throughout this lecture, I will be using the shorthand notation

$$\hat{Z}_j = \hat{I}^{\otimes j-1} \otimes \hat{Z} \otimes \hat{I}^{\otimes n-j},$$

and similarly for  $\hat{X}_j$  and  $\hat{Y}_j$ .

(In fact, the spin-spin coupling is much more complicated than I have indicated here. It includes screening effects due to the presence of the electron cloud, indirect couplings via pairs of spins both interacting with the same bound electrons, and averaging effects due to the rapid tumbling of the molecules. However, the effective Hamiltonian can be experimentally measured to give interactions similar to the ones described above. We will see shortly how this can be further simplified.)

## RF pulses

Just as in the ion trap, single-spin transitions can be produced by *resonant driving*. In this case, they are transitions between, e.g.,  $Z$ -up and  $Z$ -down states for spin  $j$ , and have resonant energy  $2E_j$ . As long as  $E_j$  is sufficiently different for different spins, one can tune the resonance so that it affects only a single spin. (It is also possible to tune pulses that affect all spins of a single chemical species.) These pulses can (in principle) produce any single q-bit rotation:  $R_{x,y,z}(\theta)$ . This makes any one-bit gate possible.

Unlike the ion trap, however, the energy differences  $2E_j$  are very small in NMR; so the resonant pulses are not at optical frequencies, but at *radio* frequencies.

These RF pulses are brief compared to the timescale of the Hamiltonian. We will approximate them as instantaneous unitaries.

## Thermal States

Because the molecules are at room temperature, the state of the spins will not be pure; they will be in a particular mixed state called a *thermal state*.

$$\rho_T = \frac{1}{Z(T)} \exp(-\hat{H}/k_B T),$$

where  $\hat{H}$  is the Hamiltonian,  $T$  is the temperature, and  $k_B$  is Boltzmann's constant.  $Z(T)$  is a normalization factor

$$Z(T) = \text{Tr}\{\exp(-\hat{H}/k_B T)\},$$

called the *partition function*. If the eigenvalues of  $\hat{H}$  are  $\lambda_m$ ,

$$Z(T) = \sum_m e^{-\lambda_m/k_B T}.$$

Typically, the energies  $E_j$  are much larger than  $\alpha_{ij}$ , so to a first approximation the energy eigenstates are  $Z$  eigenstates: computational basis states. On the other hand,  $E_j \ll k_B T$ , so  $\rho_B$  is very close to the identity.

$$\rho_T \approx \sum_{x=0}^{2^n-1} p_x |x\rangle\langle x|,$$

$$p_x = \frac{1}{2^n} \left( 1 - \sum_{j=0}^{n-1} (E_j/k_B T)(2x_j - 1) \right)$$

where  $x = x_{n-1} \dots x_1 x_0$  is the binary representation of the  $Z$  basis state  $|x\rangle$ .

## The interaction picture

Because  $E_j \gg \alpha_{ij}$ , the evolution is dominated by the magnetic field, which produces spin precessions on a timescale fast compared to the spin-spin interactions. It is possible to allow for these in a rigorous way, by transforming to a *rotating* picture: the *interaction picture*.

Let  $|\psi(t)\rangle$  be the solution to the Schrödinger equation

$$\frac{d|\psi(t)\rangle}{dt} = -(i/\hbar)(\hat{H}_B + \hat{H}_{SS})|\psi(t)\rangle.$$

Now define the *rotating state*  $|\tilde{\psi}(t)\rangle$

$$|\tilde{\psi}(t)\rangle = \exp(i\hat{H}_B t/\hbar)|\psi(t)\rangle,$$

which obeys the new Schrödinger equation

$$\frac{d|\tilde{\psi}(t)\rangle}{dt} = -(i/\hbar)\hat{H}'_{SS}(t)|\tilde{\psi}(t)\rangle.$$

This new rotating Hamiltonian is

$$\hat{H}'_{SS}(t) = \exp(i\hat{H}_B t/\hbar) \hat{H}_{SS} \exp(-i\hat{H}_B t/\hbar).$$

$\hat{H}_B$  consists of terms like  $E_j \hat{Z}_j$ , and  $\hat{H}_{SS}$  of terms like  $\alpha_{ij}(\hat{X}_i \hat{X}_j + \hat{Y}_i \hat{Y}_j + \hat{Z}_i \hat{Z}_j)$ .  $\hat{H}_B$  commutes with the  $ZZ$  terms, so they will be the same in the rotating frame. The  $XX$  terms will be transformed

$$\begin{aligned} \exp(i\hat{H}_B t/\hbar) \hat{X}_i \hat{X}_j \exp(-i\hat{H}_B t/\hbar) = \\ \hat{X}_i \hat{X}_j \exp(-2i(E_i \hat{Z}_i + E_j \hat{Z}_j)t/\hbar). \end{aligned}$$

As long as the  $E_i$  and  $E_j$  are not too close, this term will be rapidly oscillating and can be approximated as averaging to zero. This leaves us with the approximation

$$\hat{H}'_{SS} = \sum_{ij} \alpha_{ij} \hat{Z}_i \hat{Z}_j,$$

with only  $ZZ$  terms and no explicit time-dependence.

## Refocusing

The remaining Hamiltonian consists of terms that look like  $\hat{Z}_i \hat{Z}_j$ . A term of this type can produce an entangling two-bit gate which (together with one-bit unitaries) is universal for quantum computation:

$$\exp(-i\pi \hat{Z}_i \hat{Z}_j / 4) = (\hat{I} - i \hat{Z}_i \hat{Z}_j) / \sqrt{2}.$$

The problem is that we can't turn these terms off! The spins interact continuously, outside of our control.

There is a technique, however, whereby all but a single term in the Hamiltonian can be made to effectively cancel out. This is called *refocusing*. To illustrate this, suppose we have a Hamiltonian with only two terms:

$$\hat{H} = \alpha_{12} \hat{Z}_1 \hat{Z}_2 + \alpha_{23} \hat{Z}_2 \hat{Z}_3.$$

If we let this system evolve for a time  $2\tau$ , the unitary transformation would be

$$\hat{U} = \exp(-2i\tau(\alpha_{12}\hat{Z}_1\hat{Z}_2 + \alpha_{23}\hat{Z}_2\hat{Z}_3)).$$

Suppose now that we applied an  $\hat{X}$  gate to spin 3 at time  $t = 0$  and again at  $t = \tau$ . The new evolution would be

$$\begin{aligned} \hat{U}' &= \exp(-i\tau(\alpha_{12}\hat{Z}_1\hat{Z}_2 + \alpha_{23}\hat{Z}_2\hat{Z}_3))\hat{X}_3 \\ &\quad \times \exp(-i\tau(\alpha_{12}\hat{Z}_1\hat{Z}_2 + \alpha_{23}\hat{Z}_2\hat{Z}_3))\hat{X}_3 \\ &= \exp(-2i\tau\alpha_{12}\hat{Z}_1\hat{Z}_2) \exp(-i\tau(\alpha_{23}\hat{Z}_2\hat{Z}_3)) \\ &\quad \times \hat{X}_3 \exp(-i\tau\alpha_{23}\hat{Z}_2\hat{Z}_3)\hat{X}_3. \end{aligned}$$

From the algebra of Pauli operators, we know that  $\hat{X}_3 \exp(-i\tau\alpha_{23}\hat{Z}_2\hat{Z}_3)\hat{X}_3 = \exp(+i\tau\alpha_{23}\hat{Z}_2\hat{Z}_3)$ , so the factors involving  $\hat{Z}_2\hat{Z}_3$  cancel out, leaving

$$\hat{U}' = \exp(-2i\tau\alpha_{12}\hat{Z}_1\hat{Z}_2).$$

We have effectively “turned off” the term we didn’t want. This can be done with multiple terms, as well, in the same way.

## Effective Pure States

Generally speaking, it is impossible to do quantum computation in a thermal state, unless the temperature is very low. For liquid-state NMR this is not the case.

Nevertheless, there is a trick that allows interesting computations to be done. This is the preparation of an *effective pure state*

$$\rho = (1 - \epsilon)\hat{I}/2^n + \epsilon|\psi\rangle\langle\psi|,$$

where  $\epsilon \ll 1$ . A state of this type undergoes unitary transformations

$$\rho \rightarrow \hat{U}\rho\hat{U}^\dagger = (1 - \epsilon)\hat{I}/2^n + \epsilon\hat{U}|\psi\rangle\langle\psi|\hat{U}^\dagger.$$

The identity is unchanged, while the second term transforms like a pure state.

Unfortunately, what we start with is not an effective pure state, but a thermal state  $\rho_T$ . While this is close to the identity as well, the remaining terms are not a single pure state. The transformation

$$\rho_T \rightarrow (1 - \epsilon)\hat{I}/2^n + \epsilon|\psi\rangle\langle\psi|$$

is not unitary.

This is where we take advantage of the fact that we are using a bulk sample. We apply *different* unitary transformations to different molecules in the sample, and average the signals from all of them. The average of these will *not* be unitary, and if we choose them carefully they can produce an effective pure state. (In practice, this might be done by imposing a magnetic field that is different at different points in the sample.)

Generally we choose  $|\psi\rangle\langle\psi|$  to be the computational basis state  $|0 \cdots 0\rangle\langle 0 \cdots 0|$ .

## Readout

There is another reason to use a bulk sample: we have no way of measuring the state of the nuclear spins for a single molecule. They are too well isolated from the outside world, and the RF signals a single molecule produces are too weak to be detected.

Instead, the quantum computation is performed on the bulk sample, and then the sample is allowed to relax back to the thermal state. As the spins relax, they emit RF signals at the frequencies  $2E_j/\hbar$ ; the strength of these signals indicates the fraction of molecules that had spin  $j$  in the high-energy rather than the low-energy state.

Since the state of the sample is highly mixed, this would not normally be useful. However, we can take advantage of an important property of effective pure states: *the contribution to the signal from the identity cancels out*. In other words, only the fraction of systems in state  $|\psi\rangle$  produce a net signal.

By running the computation a number of times and applying different unitary transformations at the end, it is possible to completely reconstruct the final state  $|\psi_f\rangle$ . The results of the computation can then be read off from that state. (This procedure is called *quantum state tomography*.)

The strength of the signal received will be proportional to the fraction  $\epsilon$  of the systems in the state  $|\psi\rangle$ .

## Decoherence

There are two major decohering effects in NMR. The first is *thermal relaxation*: the spins can slowly gain or lose energy from the surrounding electrons or other molecules, returning to a thermal state. The second is *dephasing*, in which  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$  changes to


$$|\psi\rangle\langle\psi| \rightarrow |\alpha|^2|0\rangle\langle 0| + |\beta|^2|1\rangle\langle 1|.$$

These two processes have characteristic timescales  $T_1$  and  $T_2$ , respectively.

The decoherence rate for NMR is remarkably low.  $T_1$  can be as long as hours or days (though seconds is more typical), and  $T_2$  can be as long as tenths of a second. This gives enough time to perform hundreds or thousands of two-bit gates before the state is destroyed by decoherence.

## An example

Here is a molecule of trichloroethylene:



The three q-bits correspond to the hydrogen nucleus and the two carbon-13 nuclei. The effective Hamiltonian can be approximated

$$\hat{H} \approx \alpha \hat{Z}_1 \hat{Z}_2 + \beta \hat{Z}_2 \hat{Z}_3.$$

## Experimental set-up


Because liquid-state NMR was developed by chemists decades ago, one can do NMRQC experiments using off-the-shelf equipment.



Because of this well-advanced equipment and set of techniques, NMR quantum computing made spectacularly rapid early progress compared to other techniques. It is still the case that many quantum protocols and algorithms have only been demonstrated in NMRQC, including Shor's factoring algorithm.

As we saw in the homework, 15 is the smallest number for which the factoring problem doesn't reduce to a simplified case. In 2001 the first experimental demonstration of factoring was made at IBM, using a seven-bit NMR quantum computer.

$i$	$\omega_i/2\pi$	$T_{1,j}$	$T_{2,j}$	$J_{7i}$	$J_{6i}$	$J_{5i}$	$J_{4i}$	$J_{3i}$	$J_{2i}$
1	-22052.0	5.0	1.3	-221.0	37.7	6.6	-114.3	14.5	25.16
2	489.5	13.7	1.8	18.6	-3.9	2.5	79.9	3.9	
3	25088.3	3.0	2.5	1.0	-13.5	41.6	12.9		
4	-4918.7	10.0	1.7	54.1	-5.7	2.1			
5	15186.6	2.8	1.8	19.4	59.5				
6	-4519.1	45.4	2.0	68.9					
7	4244.3	31.6	2.0						



The experiment used a simplified version of Shor's algorithm that needed only a relatively small number of RF pulses to carry out. This is still the only practical implementation of Shor's algorithm. (The factors: 3 and 5.)

## Scaleability

In spite of this rapid experimental progress, however, as an approach to building large-scale QIP NMR is essentially a dead end. This is true for a number of reasons.

1. In preparing an effective pure state, the largest that  $\epsilon$  can be is roughly the fraction of systems that are in the correct state  $|0 \cdots 0\rangle$  just by chance at the start of the computation. For an  $n$ -spin molecule, this fraction is roughly  $\epsilon \sim 2^{-n}$ . Since the strength of the signal is proportional to  $\epsilon$ , as  $n$  becomes large the signal goes to zero.

2. Because one cannot do true measurements, error correction is impossible. (Each separate molecule might have a different error syndrome.)

3. Again, because one cannot do true measurements, one extracts the answer by doing quantum state tomography. Describing a state of  $n$  q-bits requires  $2^n$  complex numbers, which will require approximately  $2^n/n$  runs of the computer to determine.

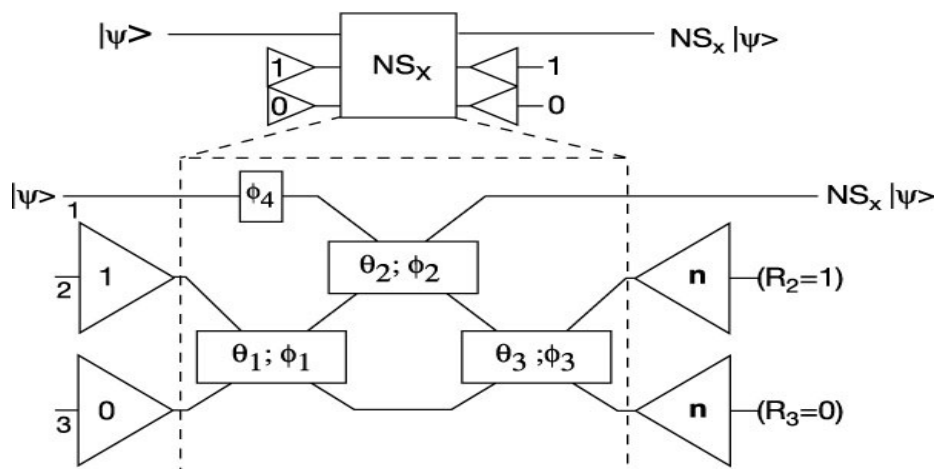
While the most impressive experiments in QIP to date have been done using NMR, in the long run general purpose quantum computers will have to use different systems. The failure of liquid-state NMR to satisfy the DiVincenzo criteria is a flaw that prevents true scalability.

## Solid state q-bits and other systems

In solid-state approaches to QIP the goal is instead of *finding* naturally-occurring q-bits, to *build* q-bits as microscopic solid-state devices. There is a large amount of technological experience in solid-state, and strong financial motives for building ever-smaller devices.

- *Solid-state NMR.* There are various schemes, using the spins of nuclei in crystals; or using the nuclei of phosphorus atoms embedded in silicon, coupled via electron spins.
- *Quantum dots.* These are tiny devices that serve as “wells” to hold electrons. The spin of an electron can serve as the q-bit, or the location of the charge in one of two dots.
- *Superconducting q-bits.* Tiny superconducting loops behave quantum mechanically; the direction of the current (or of the magnetic flux) can serve as a q-bit.

- *Linear optical QC.* It is possible for photons to serve as q-bits (as we saw in looking at quantum cryptography), and they have very low rates of decoherence, but it is hard to make them interact; this makes it difficult to build two-bit gates. It is possible to build *probabilistic* two-bit gates, however, by passing photons through interferometers and measuring some of the outputs.



Many variations of all of these ideas are being actively researched. None of them are as far along at present as the ion trap or liquid-state NMR; but it is too early yet to know which technology may eventually win out.