3
1D Liquid State NMR (No J-Couplings)

Lecture notes by Assaf Tal

The Goal of Spectroscopy

In its basic form, the goal of NMR spectroscopy is to quantify the constants appearing in the Hamiltonian: the chemical shifts, J-couplings, quadrupolar constants and so forth. Many of these are important to chemists: chemical shifts yield information about the electronic cloud and the particular chemical groups in our molecule. J-coupling constants tell us conformational information via the Karplus equation, and their existence itself tells us various spins are within 1-3 or so chemical bonds from each other.

There are other constants having to do with relaxation rates and reaction rates that do not enter the Hamiltonian explicitly, and NMR is also used to determine their values. Relaxation rates often tell us interesting things about the temporal dynamics of the system.

As we will see, achieving both of these goals requires a fair bit of ingenuity, which will be the topic of the next few lectures.

Throughout the remainder of this course we will confine ourselves to spin-1/2 spins.

Liquid State NMR

In the previous lectures we learned that the dynamics of a quantum system is governed by the Liouville equation:

$$\frac{d\rho}{dt} = \frac{1}{i\hbar}\left[\hat{H}, \rho\right].$$

and went over the various interactions that make up our system's Hamiltonian:

$$\hat{H} = \hat{H}_Z + \hat{H}_{rf} + \hat{H}_{grad} + \hat{H}_{cs} + \hat{H}_J$$
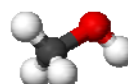$$+ \hat{H}_D + \hat{H}_{para} + \hat{H}_Q$$

In this lecture I will focus on the simplest NMR experiment in the liquid state: 1D NMR. In this experiment most interactions average out to zero due to molecular motion. Furthermore, we won't be using any gradient fields. The Hamiltonian is:

$$\hat{H} = \hat{H}_Z + \hat{H}_{rf} + \hat{H}_{cs} + \hat{H}_J.$$

Methanol

Let's assume our RF field is turned off and look at a concrete example: methanol (CH_3OH):

For methanol we assume the carbon and oxygen nuclei have no spin (^{12}C, ^{16}O). The three methyl hydrogens are chemically equivalent and have the same chemical shift, while the hydroxyl hydrogen has its own chemical shift. Furthermore they are all coupled via J-couplings. The different interactions are:

$$\hat{H}_Z = \overbrace{-\omega_0\hat{I}_{1z} - \omega_0\hat{I}_{2z} - \omega_0\hat{I}_{3z}}^{\text{three methyl protons}} \overbrace{-\omega_0\hat{I}_{4z}}^{\text{hydroxyl proton}}$$

$$\hat{H}_{cs} = \omega_{cs}^{(CH_3)}\hat{I}_{1z} + \omega_{cs}^{(CH_3)}\hat{I}_{2z} + \omega_{cs}^{(CH_3)}\hat{I}_{3z} + \omega_{cs}^{(OH)}\hat{I}_{4z}$$

$$\hat{H}_J = \overbrace{2\pi J_{12}\hat{\mathbf{I}}_1\cdot\hat{\mathbf{I}}_2 + 2\pi J_{13}\hat{\mathbf{I}}_1\cdot\hat{\mathbf{I}}_3 + 2\pi J_{23}\hat{\mathbf{I}}_2\cdot\hat{\mathbf{I}}_3}^{\text{methyl internal couplings}}$$
$$+ \underbrace{2\pi J_{14}\hat{\mathbf{I}}_1\cdot\hat{\mathbf{I}}_4 + 2\pi J_{24}\hat{\mathbf{I}}_2\cdot\hat{\mathbf{I}}_4 + 2\pi J_{34}\hat{\mathbf{I}}_3\cdot\hat{\mathbf{I}}_4}_{\text{methyl-hydroxyl couplings}}$$

By symmetry, the chemical shifts of all methyl protons are the same. We will solve this problem iteratively, and begin this lecture by looking at a single proton, neglecting its couplings.

Single Spin Spectroscopy

The Hamiltonian

Our Hamiltonian has a very simple form:

$$\frac{1}{\hbar}\hat{H} = -\omega\hat{I}_z = \begin{pmatrix} -\frac{\omega}{2} & 0 \\ 0 & \frac{\omega}{2} \end{pmatrix}.$$

Our initial Boltzmann density matrix is

$$\rho_0 = \frac{1}{2}I + \frac{2}{\gamma\hbar}M_0\hat{I}_z = \begin{pmatrix} \frac{1}{2} + \frac{2M_0}{\gamma\hbar} & 0 \\ 0 & \frac{1}{2} - \frac{2M_0}{\gamma\hbar} \end{pmatrix}.$$

The Strategy

Because the Hamiltonian is diagonal we can easily "solve" our problem; that is, calculate the density matrix as a function of time. The propagator is

$$U\left(t\right)=e^{-\frac{i\hat{H}t}{\hbar}}=\begin{pmatrix} e^{-i\omega t/2} & 0 \\ 0 & e^{i\omega t/2} \end{pmatrix}.$$

Therefore, for a general density matrix,

$$\rho=\begin{pmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{pmatrix},$$

the time evolution is

$$\rho\left(t\right)=U\rho U^{\dagger}=\begin{pmatrix} \rho_{11} & \rho_{12}e^{i\omega t} \\ \rho_{21}e^{-i\omega t} & \rho_{22} \end{pmatrix}.$$

This looks quite simple! Remember the meaning of the elements, derive in the first lecture (you can also rederive this by writing out $\frac{1}{2}I+\frac{2}{\gamma\hbar^2}\mathbf{M}\cdot\hat{\mathbf{S}}$ explicitly):

$$\rho=\begin{pmatrix} \frac{1}{2}+\frac{M_z}{\gamma\hbar} & \frac{M_x-iM_y}{\gamma\hbar} \\ \frac{M_x+iM_y}{\gamma\hbar} & \frac{1}{2}-\frac{M_z}{\gamma\hbar} \end{pmatrix}$$

Comparing the two, we see that M_z remains constant while M_{xy} precesses in the xy-plane according to the left hand rule:

$$\rho_{12}\left(t\right)=\frac{M_{xy}\left(t\right)}{\gamma\hbar}=\rho_{12}\left(0\right)e^{i\omega t}$$

$$=\frac{M_{x,0}-iM_{y,0}}{\gamma\hbar}e^{i\omega t}=\frac{\left(M_{xy,0}e^{-i\omega t}\right)^{*}}{\gamma\hbar}$$

This is precisely what the classical (Bloch eqs) picture tells us as well: the z-component of the magnetization remains fixed in time while the x & y components rotate around the B_0 field in the z-direction:
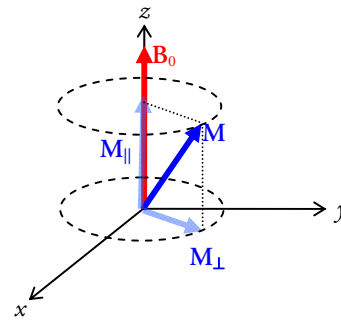
As the upper illustration shows, the magnetic moment $\mathbf{M}$ can be decomposed into components parallel and perpendicular to the main field, $\mathbf{M}=\mathbf{M_\perp}+\mathbf{M_\parallel}$, with the parallel component being static and the perpendicular component rotating in the xy-plane. Their combined motion creates the precession of $\mathbf{M}$ around $\mathbf{B}_0$. This is consistent with the quantum mechanical picture presented above.

Precessing Magnetization Induces Currents in Coils Via The Faraday Effect

This leads us to the following general strategy: if we can "generate" somehow an *x*- or *y*-component for the magnetization, We will then pick up this precession motion via Faraday's law. Indeed, the basic MR experiment can be described as follows:

➤ **Thermal Equilibrium**: At thermal equilibrium, the spins are aligned along B_0 and do not precess.

➤ **Excitation**: The spins are somehow **excited**, that is, tilted to some angle θ with respect to B_0. This usually happens quickly and relaxation can be neglected.

➤ **Precession & Detection**: Once tilted, they precess and give off a time dependent magnetic field. The magnetic field induces a voltage in a nearby RF coil via Faraday's law. We can also further manipulate the spins with magnetic fields during this period to bring out particular contrast types. We usually have a time ~ T_2 before decoherence "eats up" the observable precessing magnetization.

➤ **Thermalization**: Relaxation processes kick in. The transverse magnetization decays with a time constant T_2 while the longitudinal magnetization builds up back up due to T_1 relaxation. If we wait for a time ≈ 5·T_1, the

magnetization will be back at its thermal equilibrium value.

Each such block (excite-acquire-wait) is called a **scan**. It is in fact not mandatory to wait for a time $5 \cdot T_1$ for the spins to return to thermal equilibrium; we'll see later on that waiting a shorter amount of time has both benefits (shorter scan times) and disadvantages (less signal per scan). For now, however, we'll assume that is the case, so M is equal to M_0 and points along the z-axis before the beginning of each scan.

We've already remarked that $B_{RF} \ll B_0$. How can we hope to non-negligibly excite the spins with such a weak RF field? The answer is that we use a **resonant field** that oscillates at the Larmor frequency. Namely, we are going to solve the Bloch equations setting **G**=0, and

$$\mathbf{B}_{RF} = B_1 \cos\left(\omega_{RF}t\right)\hat{\mathbf{x}} - B_1 \sin\left(\omega_{RF}t\right)\hat{\mathbf{y}} .$$

with

$$\omega_{RF} = \omega_0 \quad (\text{"on resonance irradiation"})$$

This means we will need to solve the Bloch equations with a time dependent magnetic field. Although a numerical solution is possible, we will employ a frame transformation trick which will enable us to solve this problem without any approximations. Our strategy will be this: **we will transform to a frame of reference (aka "The Rotating Frame") that rotates with B_{RF} and in which it appears stationary. We will then solve our problem in that frame, and go back to the original frame.**

Most of the concepts will be easier to introduce using the Bloch equations rather than quantum mechanics, so we'll stick with the Bloch equations for the remainder of this lecture. However almost all of the concepts can be re-derived within the framework of QM, and some will be in the upcoming assignment.

Transforming to a Frame Which Rotates At The Same Frequency As The RF Field Makes it Appear Static: The Rotating Frame

In the laboratory frame, this amounts to solving the Bloch equations with a complicated time-dependent magnetic field. The Bloch equations are easier to solve in a frame which rotates around the z-axis with a frequency given by $\omega_{rot} = \omega_{RF}$. We tackle this as follows: consider a static (laboratory) frame with time independent, fixed unit vectors $\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}$, and a rotating frame with unit vectors $\hat{\mathbf{x}}', \hat{\mathbf{y}}', \hat{\mathbf{z}}'$.
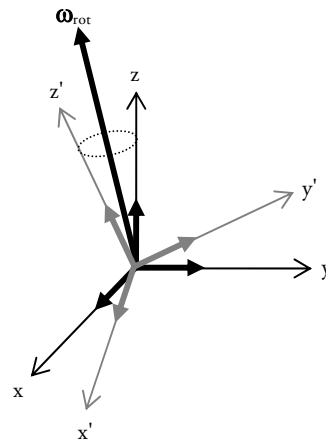
If the rotating frame is rotating with an angular velocity ω_{rot} about an axis given by the unit vector $\hat{\mathbf{n}}$, then each of the axes of the rotating frame precess about the vector $\boldsymbol{\omega}_{rot} = \omega_{rot}\hat{\mathbf{n}}$. This means each obeys a precession equation identical (formally) to the Bloch equation:

$$\frac{d\hat{\mathbf{x}}'}{dt} = \boldsymbol{\omega}_{rot} \times \hat{\mathbf{x}}'$$

$$\frac{d\hat{\mathbf{y}}'}{dt} = \boldsymbol{\omega}_{rot} \times \hat{\mathbf{y}}'$$

$$\frac{d\hat{\mathbf{z}}'}{dt} = \boldsymbol{\omega}_{rot} \times \hat{\mathbf{z}}'$$

The magnetization vector can be expressed in either frame:

$$\begin{aligned} \mathbf{M}(t) &= M_x\hat{\mathbf{x}} + M_y\hat{\mathbf{y}} + M_z\hat{\mathbf{z}} \\ &= M_{x,rot}\hat{\mathbf{x}}' + M_{y,rot}\hat{\mathbf{y}}' + M_{z,rot}\hat{\mathbf{z}}' \end{aligned}$$

For example, if **B**(t) is

$$\mathbf{B}_{RF} = B_1 \cos\left(\omega_{RF}t\right)\hat{\mathbf{x}} - B_1 \sin\left(\omega_{RF}t\right)\hat{\mathbf{y}}$$

as given above, while the rotating frame rotates at the angular frequency $\omega_{rot} = \omega_{RF}$ about the z-axis

($\boldsymbol{\omega}_{\text{rot}} = \omega_{RF}\hat{\mathbf{z}}$) then the **components** of **B** in the two frames are:

$$\begin{pmatrix} B_x \\ B_y \\ B_z \end{pmatrix} = \begin{pmatrix} B_1\cos(\omega_{RF}t) \\ -B_1\sin(\omega_{RF}t) \\ B_0 \end{pmatrix} \qquad \begin{pmatrix} B_{x,rot} \\ B_{y,rot} \\ B_{z,rot} \end{pmatrix} = \begin{pmatrix} B_1 \\ 0 \\ B_0 \end{pmatrix}$$

Note the components change, but the vector is frame-independent since it is a geometrical quantity.

Differentiating **M**(t) with respect to time, we obtain

$$\begin{aligned} \frac{d\mathbf{M}}{dt} &= \frac{dM_{x,rot}}{dt}\hat{\mathbf{x}}' + \frac{dM_{y,rot}}{dt}\hat{\mathbf{y}}' + \frac{dM_{z,rot}}{dt}\hat{\mathbf{z}}' \\ &\quad + M_{x,rot}\frac{d\hat{\mathbf{x}}'}{dt} + M_{y,rot}\frac{d\hat{\mathbf{y}}'}{dt} + M_{z,rot}\frac{d\hat{\mathbf{z}}'}{dt} \\ &= \frac{dM_{x,rot}}{dt}\hat{\mathbf{x}}' + \frac{dM_{y,rot}}{dt}\hat{\mathbf{y}}' + \frac{dM_{z,rot}}{dt}\hat{\mathbf{z}}' \\ &\quad + \boldsymbol{\omega}_{rot}\times\left(M_{x,rot}\hat{\mathbf{x}}' + M_{y,rot}\hat{\mathbf{y}}' + M_{z,rot}\hat{\mathbf{z}}'\right) \\ &= \left(\frac{d\mathbf{M}}{dt}\right)_{rot} + \boldsymbol{\omega}_{rot}\times\mathbf{M} \end{aligned}$$

On the other hand, the Bloch equation says

$$\frac{d\mathbf{M}}{dt} = \mathbf{M}\times\gamma\mathbf{B}\,.$$

Equating, we obtain:

$$\boxed{\left(\frac{d\mathbf{M}}{dt}\right)_{rot} = \mathbf{M}\times\left(\gamma\mathbf{B} - \boldsymbol{\omega}_{rot}\right)}$$

This is precisely the Bloch equation but with an **effective field** $\mathbf{B}_{eff} = \mathbf{B} - \frac{1}{\gamma}\boldsymbol{\omega}_{rot}$.

The above equation is true for any rotating frame. However, in MRI, when we speak of "**the**" rotating frame, we will be referring to a frame which rotates at a constant angular velocity $\omega_{rot} = \omega_{RF}$ about the z-axis according to the left hand rule. For "the" rotating frame, $\boldsymbol{\omega}_{rot} = \omega_{rot}\hat{\mathbf{z}} = \omega_{RF}\hat{\mathbf{z}}$.

When expressed in the rotating frame, the components of the effective field $\mathbf{B}_{eff} = \mathbf{B} - \frac{1}{\gamma}\boldsymbol{\omega}_{rot}$ are:

$$\begin{pmatrix} B_{x,eff} \\ B_{y,eff} \\ B_{z,eff} \end{pmatrix} = \begin{pmatrix} B_1 \\ 0 \\ B_0 - \frac{1}{\gamma}\omega_{rot} \end{pmatrix}$$

(in the rotating frame: $\omega_{rot} = \omega_{RF}$)

If we select $\omega_{RF} = \gamma B_0 = \omega_0$ we are **on resonance**: the RF irradiates the spins at the same frequency as their natural frequency, ω_0. In this case:
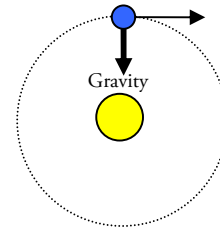
$$\mathbf{B}_{eff} = \begin{pmatrix} B_1 \\ 0 \\ 0 \end{pmatrix}$$

On resonance: $\omega_{RF} = \omega_0$
In the rotating frame: $\omega_{rot} = \omega_{RF}$

An Analogy From Mechanics

Imagine the earth going around the sun in a circle:

This can be understood by an observer in space the following way: the Earth wants to "go forward" but gravity pulls it "inward", curving its path into a circle. In effect, the Earth is continuously "falling" into the sun, but escaping doom thanks to its tangential velocity. All this is all a consequence of Newton's second law, F=ma.

Next, imagine how things would look to an observer standing on the sun and rotating with it. Neglecting for the time being the weather on the surface, the Earth would appear stationary to such an observer:
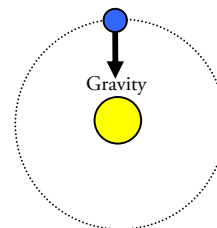
If that observer would try to use Newton's law F=ma to understand his world he would fail: according to F=F$_{gravity}$=ma, earth should be falling towards the sun, but it isn't! The truth is that when you transform to a rotating frame you need to add a **fictitious force**. That is, you need to pre-suppose a force which doesn't arise out of any physical source, called the centripetal force, to explain how it is possible for the earth to remain stationary:
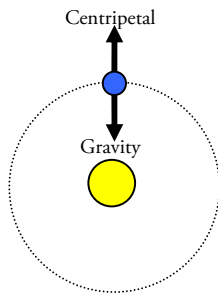
So, in mechanics when you try to understand things in a rotating frame you need to do two things:

1. Understand how things in the "real" frame would look in the rotating frame (e.g., the Earth would remain still).
2. Add fictitious forces (e.g., the centripetal force).

A similar thing happens when you go to a rotating frame in magnetic resonance, rotating with the same angular velocity as the RF field:

1. First, the RF field appears stationary in the rotating field which "matches" its rotation frequency (i.e. because $\omega_{rot} = \omega_{RF}$).
2. Now we need to add the correct fictitious "force" - field, to be precise - given by $\mathbf{B}_{fict} = -\frac{1}{\gamma}\omega_{rot}$. To see, imaging a static spin in the lab frame, with no magnetic field. Now transform to a frame rotating with an angular velocity ω_{rot} about the z-axis. In this frame, the spin would appear to rotate with an angular velocity $-\omega_{rot} = \gamma B_{fict}$, as if there was a fictitious field $B_{fict} = -\frac{\omega_{rot}}{\gamma}$ present along the z-axis.

The Bulk Magnetization Precesses Around The Effective Field In The Rotating Frame

We've seen the magnetization vector obeys the Liouville equations in the rotating frame, only swapping the field for an effective field, $\mathbf{B}_{eff}$. This means $\mathbf{M}$ precesses about $\mathbf{B}_{eff}$ in the rotating frame. Starting from thermal equilibrium at time t=0, $\mathbf{M}$ points along B_0 (taken to coincide with the z-axis) in both the laboratory and the rotating frames, which are also assumed to coincide for t=0:
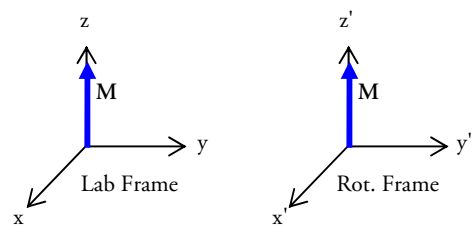

At time t=0 (thermal equilibrium), $\mathbf{M}$ points along the z-axis in both frames.

Now we turn on the resonant RF field in the laboratory frame:

$$\mathbf{B}_{RF} = B_1 \cos\left(\omega_{RF}t\right)\hat{\mathbf{x}} - B_1 \sin\left(\omega_{RF}t\right)\hat{\mathbf{y}} .$$

This field rotates in the xy-plane in the lab frame, and appears stationary in the rotating frame. Furthermore, if we assume our irradiation is on resonance, $\omega_{RF}=\omega_0$, the effective field in the rotating frame has no z-component:
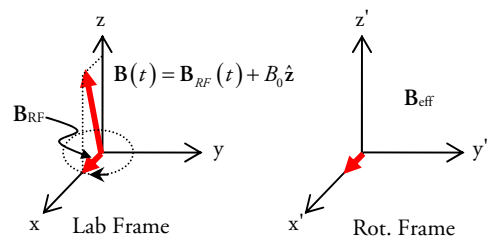
The magnetic field $\mathbf{B}$ in the laboratory frame has a large z-component and a small, rotating xy-component (not shown to scale). In the rotating frame, assuming B$_{RF}$ is on resonance ($\omega_{RF}=\omega_0=\gamma B_0$) the effective field is static.

The magnetization $\mathbf{M}$ precesses about the x axis in the rotating frame. We can thus create any angle

we'd like between it and the z-axis, depending on how long we let it precess and how strong B_1 is. Let's assume we have B_{RF} on for just enough time for the magnetization to tilt to the xy plane - that is, create a 90° angle between B_0 and **M**. Deducing the motion of **M** in the lab frame is now merely a matter of transforming back to the lab frame, which simply rotates at an angular velocity $-\omega_{rot}$ relative to the rotating frame. That is, **M** in the lab frame performs a spiral as it descends and rotates:
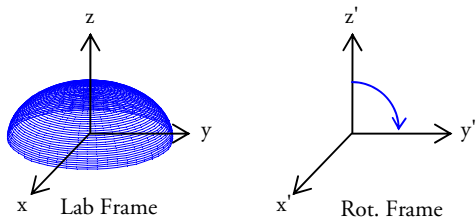
Shown here is the trajectory of the magnetization **M** in the lab (left) and rotating (right) frames. The two frames are connected by a simple rotation.

Setting The Radiofrequency (RF) Pulse's (Area)=(Duration)×(Amplitude) Sets The Flip Angle

We see the spins will perform a rotation about the x-axis in the rotating frame at a frequency $\omega_1 = \gamma B_1$. Note this is **not the same as** ω_{RF}, (one is the amplitude of B_{RF}, the second is its oscillating frequency). After a time τ, **M** will have created an angle $\alpha = \omega_1 \tau = \gamma B_1 \tau$:
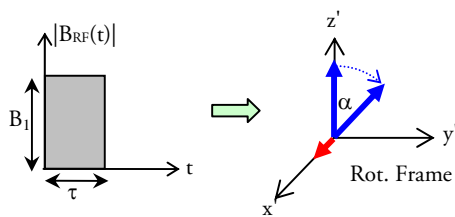
Note that

$$\alpha = \gamma \left(\text{amplitude of RF} \right) \times \left(\text{duration of RF} \right).$$

This relation is true only on resonance, when $\omega_{RF} = \omega_0$, where $\mathbf{B}_{eff}$ has no z-component.

To "tip" the magnetization onto the y axis, we wait a time t_{90} such that:

$$\alpha = \gamma B_1 t_{90} = \tfrac{\pi}{2},$$

or

$$t_{90} = \tfrac{1}{4\gamma B_1}.$$

In the original laboratory (unrotating) frame the spins execute additional motions, but the important thing to realize is that **a spin which is in the xy plane in the rotating frame, must also be in the xy-plane in the laboratory frame** (although *where* in the plane is a different story!).

Example: We've remarked that $B_{1,\max} \sim 10\ \mu T$ for an MRI scanner. For protons, one would need $t_{90} = \tfrac{\pi}{2\gamma B_1} \sim 0.5$ ms to excite the spins onto the xy-plane. For ^{13}C, $t_{90} = \tfrac{\pi}{2\gamma B_1} \sim 2$ ms .

The Phase of the Pulse Determines The Phase of the Excited Magnetization

We have so far modeled B_{RF} in the lab frame as:

$$\mathbf{B}^{(RF,lab)} = \begin{pmatrix} B_1 \cos\left(-\omega_c t\right) \\ B_1 \sin\left(-\omega_c t\right) \\ 0 \end{pmatrix}.$$

Since we have full control over the x and y component we have no problem modulating both $B_1(t)$ and adding a time-dependent phase $\phi(t)$ to the RF field:

$$\mathbf{B}^{(RF,lab)} = \begin{pmatrix} B_1\left(t\right) \cos\left(-\omega_c t + \phi\left(t\right)\right) \\ B_1\left(t\right) \sin\left(-\omega_c t + \phi\left(t\right)\right) \\ 0 \end{pmatrix}$$

In the rotating frame, this will look like this[1]:

---

[1] To prove this, use $\mathbf{B}^{(rot)} = R_z\left(\omega_c t\right)\mathbf{B}^{(lab)}$, where $R_z\left(\omega_c t\right)$ is a RH rotation matrix about the z-axis by an angle $\theta = \omega_c t$ (the rotating frame rotates with a left handed rotation and angular frequency ω_c; in it, it appears the RF field rotates at the same angular frequency but in the opposite direction). There is a bit of algebra and trigonometry involved but the proof is straightforward.

$$\mathbf{B}^{(RF,rot)} = \begin{pmatrix} B_1(t)\cos(\phi(t)) \\ B_1(t)\sin(\phi(t)) \\ 0 \end{pmatrix}.$$

Let's keep $B_1(t)$ and $\phi(t)$ fixed. Then the constant phase $\phi(t)=\phi_0$ is called the **phase of the pulse**, and is equal to the angle the RF field makes with the x-axis. determines where the RF pulse will point in the transverse plane.

The **phase of the magnetization** is defined as the angle made by the transverse component of the magnetization vector (i.e. its projection on the xy plane) with the x-axis.

Because the magnetization gets tipped at right angles to the RF field following the left hand rule, the relation between the pulse's and magnetization's phase is given by:

$$\phi_m = \phi_{RF} + \tfrac{\pi}{2}.$$

The standard notation for a constant RF pulse then assumes the form α_ϕ, where α is its flip angle and ϕ its (constant) phase. The following conventions are also used:

$$\begin{aligned} \phi &= 0^\circ: & x \\ \phi &= 90^\circ: & y \\ \phi &= 180^\circ: & -x \\ \phi &= 270^\circ: & -y \end{aligned}$$

Some examples are shown below (magnetization is assumed to start out from z, and is the blue vector; the RF is the red vector):

$(90°)_x$ pulse

magnetization in yz plane (phase: 90°)

$(45°)_x$ pulse

magnetization in yz plane (phase: 90°)

$(90°)_{-y}$ pulse

magnetization in xz plane (phase: 0°)

$(45°)_y$ pulse

magnetization in xz plane (phase: 180°)
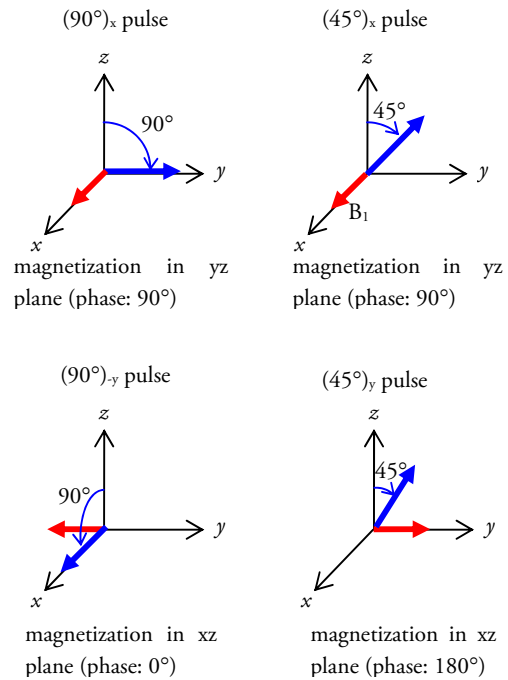
The Signal Is A Decaying Sinusoid

After the magnetization is tipped onto the xy-plane it precesses around the external B_0 field. We can pick it up by measuring the voltage it induces in a nearby coil via Faraday's law. We'll get into the practical aspects of signal detection in the next lecture, but for now I merely remark that this voltage is proportional to the xy-magnetization. That is,

$$s(t) \propto M_{xy}(t) \equiv M_x(t) + iM_y(t).$$

What does a complex signal mean? It just means that we get two signals out of the NMR spectrometer, one proportional to M_x and the other to M_y, and we combine them in the acquisition computer to generate a complex signal which is just easier to deal with.

For a spin having an offset $\omega = \gamma(1+\sigma_{iso})B_0$, we have

$$\mathbf{M}(t) = \begin{pmatrix} M_0\cos(\omega t) \\ -M_0\sin(\omega t) \\ 0 \end{pmatrix}$$

and so

$$s(t) \propto M_{xy}(t)$$
$$= M_0 \cos(\omega t) - iM_0 \sin(\omega t)$$
$$= M_0 e^{-i\omega t}$$

In reality, due to *decoherence*, the signal decays as well with a time constant called T_2:

$$s(t) \propto M_0 e^{-i\omega t} e^{-t/T_2} \,.$$

We discuss this decay next.

Spin Interactions Lead To Relaxation Phenomena

## Spins Are Subjected To Microscopic Fluctuating Magnetic Fields Due To Their Thermal Motion

Each microscopic nuclear magnetic moment **m** "sees" a magnetic field made up of two components: the macroscopic field generated by the coils in the lab, and the microscopic fields given off by its surroundings. For example, the dipolar field generated by one nuclear spin in a molecule will be felt by other nuclear spins in the same molecule.
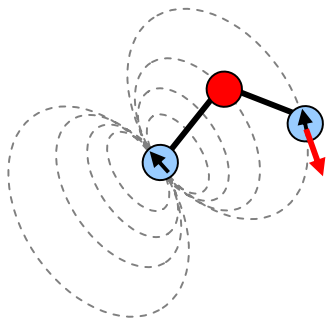
Shown here is the magnetic field (red arrow) felt by one spin due to the dipolar field of the other spin in an H_2O molecule.

It's very important to realize that the orientation of the nuclear magnetic moment has nothing to do with the molecular orientation: if you rotate the molecule by 90°, the nuclear moment will not change, since it's not related to the nuclear charge or mass distribution; it "lives" in its own space and

"talks" to the environment only through the magnetic fields it feels and emits:
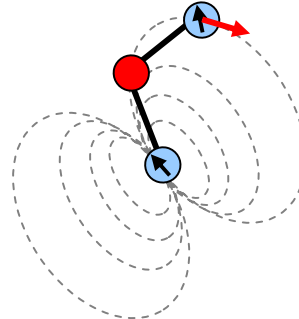
Upon rotation of the molecule, the spins (black arrows) do not change their orientation. Consequently, the spin feels a different magnetic field, in both magnitude and direction.

Since most of the water molecules in the body are in the liquid state in the extra and intracellular matrices[2] –All molecules rotate and tumble around very rapidly. A small water molecule might perform a rotation on picosecond timescales, while larger molecules would rotate more slowly. This molecular rotation leads, by the arguments just laid out, to fluctuating microscopic fields.

Fluctuating Microscopic Fields Lead To Decoherence (T_2) And Return to Thermal Equilibrium (T_1)

The magnetic field felt by a microscopic nuclear magnetic moment can be subdivided into two parts, macroscopic and microscopic:

$$\mathbf{B}(t) = \mathbf{B}_{macro}(t) + \mathbf{B}_{micro}(t),$$

where the macroscopic fields are those generated by the laboratory coils and controlled by the scientist, and the microscopic fields are those fluctuating fields created by other spins in the molecule, electrons, and so forth. Consequently, the Bloch equations which describe the spin's precession become:

---

[2] This is actually not entirely correct, since water molecules often get "stuck" to cell membranes or confined in tight spaces. We will look more into this in later lectures.

$$\frac{d\mathbf{m}}{dt} = \gamma\mathbf{m}\times\mathbf{B} = \gamma\mathbf{m}\times\mathbf{B}_{macro}(t) + \gamma\mathbf{m}\times\mathbf{B}_{micro}(t).$$

Now assume we have N magnetic moments, $\mathbf{m}_1$, $\mathbf{m}_2$, ..., $\mathbf{m}_N$, each experiencing its own unique microscopic field, but all experiencing the same macroscopic one:

$$\frac{d\mathbf{m}_1}{dt} = \gamma\mathbf{m}_1\times\mathbf{B}_{macro}(t) + \gamma\mathbf{m}_1\times\mathbf{B}_{micro}^{(1)}(t)$$

$$\frac{d\mathbf{m}_2}{dt} = \gamma\mathbf{m}_2\times\mathbf{B}_{macro}(t) + \gamma\mathbf{m}_2\times\mathbf{B}_{micro}^{(2)}(t)$$

$$\vdots$$

$$\frac{d\mathbf{m}_N}{dt} = \gamma\mathbf{m}_N\times\mathbf{B}_{macro}(t) + \gamma\mathbf{m}_N\times\mathbf{B}_{micro}^{(N)}(t)$$

We now sum over multiple microscopic spins:

$$\sum_{n=1}^{N}\frac{d\mathbf{m}_n}{dt} = \gamma\sum_{n=1}^{N}\mathbf{m}_n\times\mathbf{B}_{macro}(t) + \gamma\sum_{n=1}^{N}\mathbf{m}_n\times\mathbf{B}_{micro}^{(n)}(t)$$

Since $\mathbf{B}_{macro}$ is common to all summed terms, and since the derivative of the sum equals the sum of the derivatives, we can substitute the microscopic moments by the macroscopic one, $\mathbf{M} = \Sigma_{n=1}^{N}\mathbf{m}_n$ and obtain:

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M}\times\mathbf{B}_{macro}(t) + \gamma\sum_{n=1}^{N}\mathbf{m}_n\times\mathbf{B}_{micro}^{(n)}(t)$$

The last term on the RHS represents the effects of the fluctuating fields and is intractable really. Physically speaking, these fluctuating magnetic

Tissue Type	Nuc.	Mol.	1.5 T		3 T		7 T	
			T_1	T_2	T_1	T_2	T_1	T_2
Gray Matter[b]	^{1}H	H_2O	1188 ± 69	95 ± 8	1820 ± 114	99 ± 7	2132 ± 103	
White Matter[b]	^{1}H	H_2O	656 ± 16	72 ± 4	1084 ± 45	69 ± 3	1220 ± 36	
Cerebrospinal Fluid[b]	^{1}H	H_2O	4070 ± 65				4425 ± 137	
Blood[b]	^{1}H	H_2O	1540 ± 23	290 ± 30	1932 ± 85	275 ± 50	2587 ± 283	
Kidney Cortex[a]	^{1}H	H_2O	966 ± 58	87 ± 4	1142 ± 154	76 ± 4		
Kidney Medula[a]	^{1}H	H_2O	1412 ± 58	85 ± 11	1545 ± 142	81 ± 8		
Liver[a]	^{1}H	H_2O	586 ± 39	46 ± 6	809 ± 71	34 ± 4		
Cartilage, 0° [d]	^{1}H	H_2O	1024 ± 70	30 ± 4	1168 ± 18	27 ± 3		
Cartilage, 55° [d]	^{1}H	H_2O	1038 ± 67	44 ± 5	1156 ± 10	43 ± 2		
Bone marrow (L4 vertebra)[a]	^{1}H	H_2O	549 ± 52	49 ± 8	586 ± 73	49 ± 4		
Prostate[a]	^{1}H	H_2O	1317 ± 85	88 ± 0	1597 ± 42	74 ± 4		
Subcutaneous fat[a]	^{1}H	Fat	343 ± 37	58 ± 4	382 ± 13	68 ± 4		
NAA CH3 (GM)[c]	^{1}H	NAA	1270 ± 50		1470 ± 80	269 ± 7		
NAA CH3 (WM)[c]	^{1}H	NAA	1360 ± 60		1400 ± 150	374 ± 9		

Typical T_1 and T_2 relaxation times from the literature, in milliseconds, in humans. The ± sign indicates standard deviation of the cohort examined. Note that variations may occur within a particular tissue (e.g. cortical vs. deep gray matter), and that numbers provided from different papers might originate from different regions within the same tissue. Also, some skepticism should be practiced when using values obtained for flowing/pulsating media, such as the cerebrospinal fluid.

[a] From: Bazelaire et. al., Radiology 230(3):652-659 (2004)

[b] T_1 values at 1.5 T and 7 T from Rooney et. al., Magn. Reson. Med. 57:308-318 (2007).
T_1, T_2 values at 3 T taken from Rooney et. al., Magn. Reson. Med. 57:308-318 (2007) and Stanisz et. al., Magn Reson Med 54:507-512 (2005).

[c] T_2 values at 3T taken from Kirov et. al., Magn. Reson. Med. 60:790-795 (2008).
T_1 values at 1.5T and 3T from Ethofer et. al., Magn Reson Med 50:1296-1301 (2003)

[d] From: Stanisz et. al., Magn Reson Med 54:507-512 (2005).

fields are the source of (1) **decoherence** (i.e. loss of signal) and (2) **thermalization** (return to thermal equilibrium). Luckily, **phenomenologically** these effects can be respectively embodied by two constants, T_2 and T_1, respectively, which can be integrated into the Bloch equations using simple terms:

$$\begin{cases} \dot{M}_x = \gamma\left(M_y B_{\text{macro},z} - M_z B_{\text{macro},y}\right) - \frac{M_x}{T_2} \\ \dot{M}_y = \gamma\left(M_z B_{\text{macro},x} - M_x B_{\text{macro},z}\right) - \frac{M_y}{T_2} \\ \dot{M}_z = \gamma\left(M_x B_{\text{macro},y} - M_y B_{\text{macro},x}\right) - \frac{M_z - M_0}{T_1} \end{cases}$$

We will omit the subscript *macro* and take the magnetic field appearing in the Bloch equations to signify only the macroscopic (lab-generated) magnetic field.

M_0 is the thermal equilibrium value of the magnetization, as can be seen by turning "off" the macroscopic RF and gradient fields, setting the time derivatives to 0 and solving:

$$\begin{cases} 0 = \gamma M_y B_0 - \frac{M_x}{T_2} \\ 0 = -\gamma M_x B_0 - \frac{M_y}{T_2} \\ 0 = -\frac{M_z - M_0}{T_1} \end{cases} \xrightarrow{\ \text{solve}\ } \begin{cases} M_x = 0 \\ M_y = 0 \\ M_z = M_0 \end{cases}$$

A table of some T_1 and T_2 values has been compiled above. We note that for most tissues, T_1 is on the order of a second, while T_2 is on the order of 100 ms. Furthermore, T_1 tends to increase with increasing field strength, while T_2 tends to decrease. The field-dependence of T_1 and T_2 will await a further chapter which will discuss T_1 and T_2 as sources of contrast.

T_2 Leads To Decoherence

To gain a better understanding of the sort of effect T_2 has on the spins, let us set the macroscopic laboratory field to 0 and examine the time evolution of the magnetization.

$$\begin{cases} \dot{M}_x = -\frac{M_x}{T_2} \\ \dot{M}_y = -\frac{M_y}{T_2} \\ \dot{M}_z = -\frac{M_z - M_0}{T_1} \end{cases}$$

One interesting this is that the transverse (x, y) and longitudinal (z) components of the magnetization become decoupled: M_z does not feature in the equations for M_x and M_y, and M_x and M_y do not appear in the equation for M_z.

The equations for M_x and M_y have simple solutions:

$$M_x(t) = M_x(t = 0)e^{-t/T_2}$$
$$M_y(t) = M_y(t = 0)e^{-t/T_2}$$

This means that, whatever magnetization we start out with, it will decay with a time constant T_2 to zero:
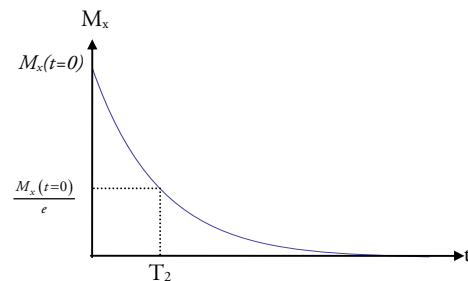
This is called **decoherence**, and represents the physical fact that, unless something specific is done, the spins will point in all possible directions perpendicular to the MRI's static B_0 field, since there is no reason – energetic preference – for them to align in any single particular direction. The time T_2 can be thought of the time it takes M_x (or M_y) to drop to $1/e\approx 37\%$ of its initial value.

T_1 Leads To Thermal Equilibrium

At thermal equilibrium the spins align themselves along the external B_0 field. This is brought about by T_1 relaxation. The solution to the equation involving M_z is:

$$M_z(t) = M_z(t = 0)e^{-t/T_1} + \left(1 - e^{-t/T_1}\right)M_0.$$

We see that, for $t \gg T_1$,

$$M_z(t \gg T_1) \approx M_0.$$

Thus, whatever longitudinal magnetization we start out from at t=0, it will converge back to its thermal equilibrium value M_0:
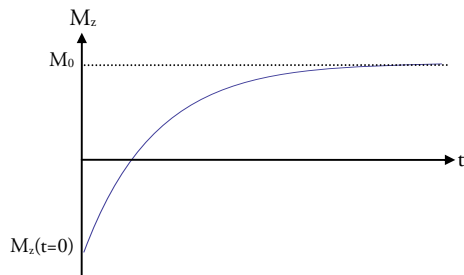
Relaxation Can Be Neglected During Excitation Since Most Pulses Are Shorter Than T_1, T_2

Our calculations in the previous section have shown that excitation mostly happens on the timescale of milliseconds in NMR, which is much shorter than T_1, T_2. Hence, to an excellent approximation, relaxation effects can be neglected during excitation. This might have to be re-examined in solid-state NMR or when dealing with large macromolecules, where T_2s can be prohibitively short (even in the microsecond range!).

The NMR Spectrum

The Fourier Transform

The complex signal from a single nucleus having a given chemical shift is

$$s(t) = s_0 e^{-i\omega_1 t} e^{-t/T_2} .$$

It is fairly easy to deduce ω_1 looking at the signal in the time domain. However, what happens if our signal has, say, three components?

$$s(t) = s_1 \cdot e^{-i\omega_1 t} e^{-t/T_{2,1}} + s_2 \cdot e^{-i\omega_2 t} e^{-t/T_{2,2}}$$
$$+ s_3 \cdot e^{-i\omega_3 t} e^{-t/T_{2,3}} .$$

It is exceedingly difficult for humans to deduce the ω_is by looking at the signal in the time domain. Fortunately, there is a tool that simplifies this, known as the Fourier transform (FT). The FT acts as a "magic box" which reveals the frequency characteristics of a time domain signal s(t) in the form of a spectrum. The spectrum is comprised of

peaks: a peak centered at ω_i tells us s(t) has a frequency component $e^{i\omega_i t}$, and the peak's "size" tells us what its coefficient s_i is.

Given a signal s(t), its Fourier transform is defined as:

$$\hat{s}(\nu) = \int_{-\infty}^{\infty} s(t) e^{2\pi i \nu t} dt .$$

Example: if s(t)=1 for t∈[-T/2, T/2] and 0 elsewhere (a rectangle), then:

$$\begin{aligned}
\hat{s}(\nu) &= \int_{-\infty}^{\infty} s(t) e^{2\pi i \nu t} dt \\
&= \int_{-T/2}^{T/2} e^{2\pi i \nu t} dt \\
&= \left. \frac{e^{2\pi i \nu t}}{2\pi i \nu} \right|_{-T/2}^{T/2} \\
&= \frac{e^{\pi i \nu T} - e^{-\pi i \nu T}}{2\pi i \nu} \\
&= \frac{1}{\pi \nu} \cdot \sin(\pi \nu T) = T \cdot \mathrm{sinc}(\pi \nu T)
\end{aligned}$$

The sinc has a main lobe with width $\Delta \nu = \frac{1}{T}$. This is typical of FTs: the width of the FT is usually inversely proportional to the width of the original function.

The Fourier Transform of a Decaying Exponential

The NMR signal is made up of decaying exponentials. The Fourier transform itself is linear, meaning that if

$$FT[s(t)] \equiv \int_{-\infty}^{\infty} s(t) e^{2\pi i \nu t} dt$$

then

$$FT\left[\sum_n s_n e^{-i\omega_n t - t/T_{2,n}} \right] = \sum_n FT\left[s_n e^{-i\omega_n t - t/T_{2,n}} \right]$$

so we only need to calculate the FT of a single summand. This is easily achieved:

$$\int_{-\infty}^{\infty} s(t) e^{2\pi i v t} dt$$

$$= \int_0^{\infty} s_n e^{\left[2\pi i (v-v_n) - \frac{1}{T_{2,n}}\right] t} dt$$

$$= s_n \left(\frac{-1}{\left[2\pi i (v-v_n) - \frac{1}{T_{2,n}}\right]} \right)$$

$$= s_n \left[\frac{T_{2,n}}{1+\left(2\pi\Delta v T_{2,n}\right)^2} + i \frac{2\pi T_{2,n}^2 \Delta v}{1+\left(2\pi\Delta v T_{2,n}\right)^2} \right]$$

(the last line is obtained by multiplying and dividing by the complex conjugate of the denominator and simplifying.)

The real part is a **Lorentzian** function and is called the **absorptive** part of the spectrum. The imaginary part is called the **dispersive** part of the spectrum. These names are a legacy from optical spectroscopy, from which they were originally borrowed. There, the coefficient of refraction in a material, n, has a real and imaginary part: the real part causes the signal to decay (get absorbed) while the imaginary part causes different frequencies to progress as different speeds through the material, leading to dispersion of the components of the incoming wave packet.

We'll note here in passing that the Lorentzian lineshape is "well behaved": it is fairly well-localized and corresponds to what a peak "should be". The dispersive component is quite the opposite: it transitions sharply from negative to positive, and decays very slowly – i.e., is non-localized. Dispersive components are almost always unwanted but they are a fact of life NMR spectroscopists need to live with. In a well acquired spectrum, the dispersive component will appear in the imaginary part of the spectrum and will not interfere with the absorptive part, but in non-ideal acquisitions the two can get "mixed" and must be dis-entangled via a process known as "phasing". This will be discussed in the next lecture.
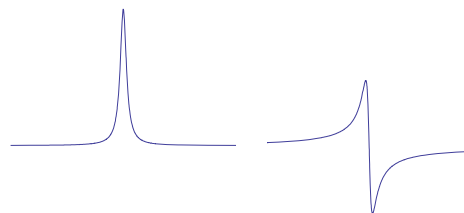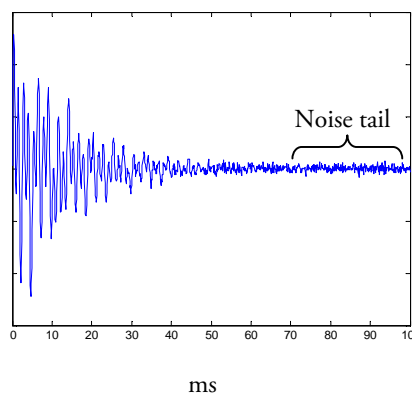

Absorptive (left) and dispersive (right) components of the NMR spectrum.

Practical Aspects of NMR: Acquiring a Spectrum

Setting the Acquisition Time

A typical FID looks like this:

How long should you acquire for? On the one hand, if you acquire for too long you might end up just wasting time and acquiring unnecessary noise. If you don't acquire for long enough you might lose out on important signal. The rule of thumb is to *acquire until your signal decays away*, which happens around $5T_2$ ms. This means you should have an idea of when the signal decays by running a preliminary experiment or knowing something about your sample. In the example above, T_2 was about 15 ms. You'd might have to run a quick reference scan to get a rough idea of how far your signal goes.

Sampling and The Nyquist Criterion

The acquired signal is digitized and we record not the continuous analog signal, but a set of points acquired at equidistant time intervals. This time

interval is called the *dwell time* and usually denoted dt or Δt.

If one acquires for a time T and a dwell time dt, then they will end up with

$$N = \left\lfloor \frac{T}{dt} \right\rfloor$$

at times 0, dt, 2dt, 3dt, ... Consequently, the fourier transform the occurs in the computer is called the *discrete fourier transform*. The MATLAB command that carries it out is called fft.

The effect of the dwell time is to cause **aliasing** in the spectrum: imagine the spectrum not as a linear graph, but as a sheet of paper wrapped around a cylinder of length 1/dt. If the cylinder is too short, the paper will "wrap" onto itself and will make it difficult to read the spectrum. As long as 1/dt is bigger than the length of the sheet of paper you should be ok. This is demonstrated in the next example, in which there are four peaks at -110, -10, 40 and 70 Hz:
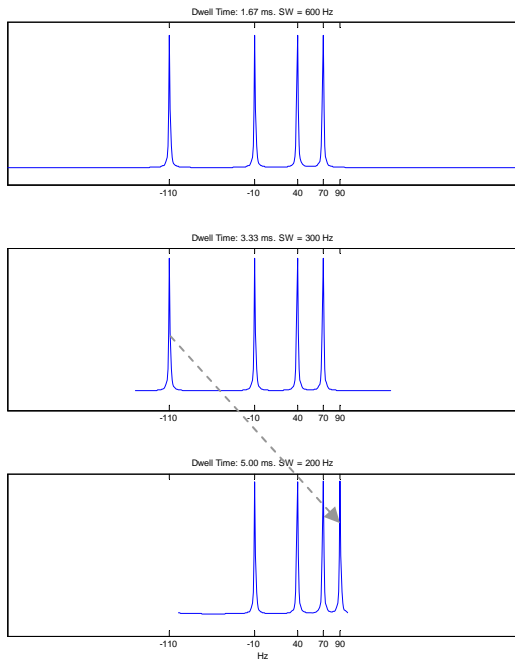
In the third case, dt=5 ms and 1/dt = 200 Hz. The peak at -110 then "folds" back onto 90 Hz (90 = -110+1/dt = -110 + 200).

The reason for aliasing can be understood by looking at the following example, in which the signal $\cos(2\pi\nu t)$ is digitized for different offsets ν

at a dwell time of 4 ms (total acquisition time 40 ms):
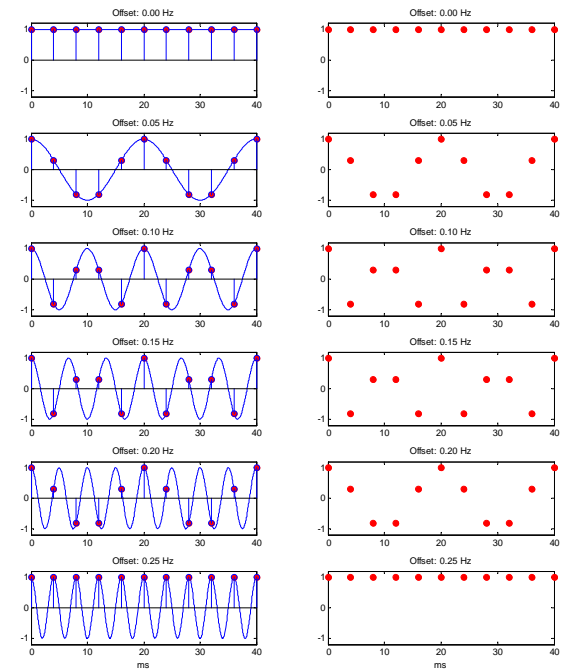
We see that once we progress by 1/dt=0.25 kHz from ν=0, we once again acquire a constant set of points, making it impossible for us to distinguish between ν=0 and ν=0.25 kHz (or ν=0.5 kHz, 0.75 kHz, etc ...). Think of the ADC as a stroboscopic party light: we only observe the scene at equidistant time points (0, dt, 2dt, 3dt, ...), but have no way of knowing what happened between those time points. If we tried to view something that had a periodicity of the stroboscopic light we wouldn't see anything and mistake it for being constant.

The range of non aliased frequencies we observe is called the **spectral width** (denoted SW), and we have just shown that:

$$SW = \frac{1}{dt}.$$

Q: Why not sample really really fast (use tiny Δt) and make the SW really big so we don't have to worry about aliasing?
A: First, all ADCs have a maximal sampling rate, which may or may not allow fast sampling. Second, most ADCs tend to use really small dwell times "behind the scenes", and actually NMR spectrometers tend to *oversample* (use small Δts)

and then digitally *downsample*. This is a slightly complex process which is done because it makes it possible to build simpler analog low pass filters in the ADC. We won't go into the reasons in this course (but you can come and ask me if you're curious).

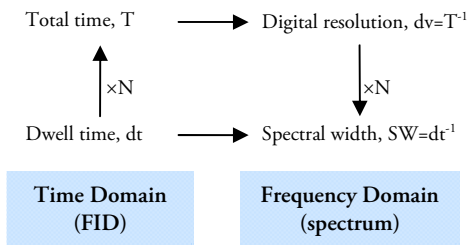## Back to Acquisition Time: Digital Resolution

A complementary parameter to the dwell time is the total acquisition time, T. If we have N points and a dwell time dt, then

$$T = N \cdot dt .$$

The total acquisition time determines the digital resolution: the smallest frequency range one can observe in the Fourier transformed spectrum:

$$\frac{1}{T} = d\nu .$$

Thus we have an "inverse" relationship between the time and frequency domains:

Total time, T → Digital resolution, $d\nu = T^{-1}$

$\times N$ ↑   ↓ $\times N$

Dwell time, dt → Spectral width, $SW = dt^{-1}$

| Time Domain (FID) | Frequency Domain (spectrum) |

Where does the relation $T = 1/d\nu$ come from? We can understand this by examining the FT of a complex exponent, $f(t) = e^{i2\pi\nu t}$. We have seen in the previous lecture that

$$\hat{f}(\nu) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i\nu t} dt = \delta(\nu - \nu_0)$$

and, if we zero out the function outside [-T/2, T/2],

$$\hat{f}_T(\nu) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i\nu t} dt = T \cdot \text{sinc}(\pi\nu T).$$

As $T \rightarrow \infty$ we have $\hat{f}_T(\nu) \rightarrow \delta(\nu - \nu_0)$. However, for a finite T – that is, for a finite acquisition time – we obtain a **broadening** of the signal on the order of $d\nu = T^{-1}$. This means that anything thinner than $d\nu = T^{-1}$ (say, a delta function) will "fatten up" and get a width of $d\nu = T^{-1}$ simply because we acquire for a finite amount of time, T. This can be seen in the following example, in which two chemical shifts at 0 Hz and 200 Hz were simulated with different acquisition times, 512 acquisition points and $T_2 = \infty$. Ideally for an acquisition time of $T = \infty$ we should get a perfect delta function. For $T < \infty$ the delta function is replaced with sinc-like functions. The real part of the spectrum is displayed after 15-fold zero filling of the FID (See below for what zero filling is):

Acq. time: 500 ms

Acq. time: 200 ms

Acq. time: 50 ms

Acq. time: 10 ms
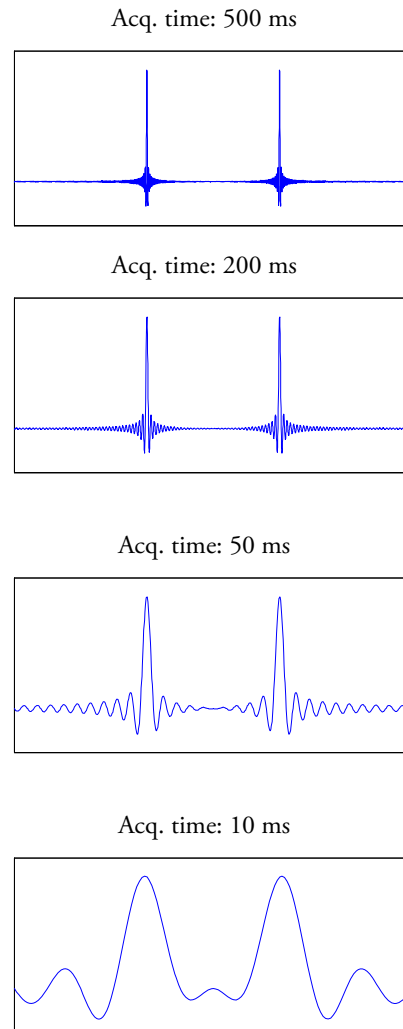
In general, the width of each peak behaves as 1/T (approximately).

The example highlights another interesting feature: **ringing**, which refers to the sinc-like wiggles accompanying each peak. This comes whenever we truncate our signal, which is the case here, since we suddenly stop acquiring after a time T. Ringing is avoided if the FID decays smoothly to zero, which is the case if T>>T_2. Even if T<T_2 the FID can be multiplied by a function that has a smooth decay, such as exp(-t/T_{smooth}). This will make the ringing go away.

The digital resolution is not the only factor affecting resolution. We have already remarked that the width of a Lorentzian peak is determined by T_2 and given by approximately $1/T_2$. This means that even if we acquire for an infinite amount of time (T=∞), our spectral peaks will still be broadened by their natural T_2 decay which also limits our resolution. So the following should be kept in mind: **The fastest decay factor of our signal determines our ultimate resolution**. If T_2<T then our resolution will be $1/T_2$. If T<T_2 then our resolution will be $1/T$. If there is some other factor causing our signal to decay even faster than T, T_2 then that will determine our peaks' widths and, hence, our resolving power.

Lock

NMR samples are prepared in a solvent. Many of these are sold in deuterated forms. For example, D_2O instead of H_2O. This is done for two reasons: to reduce the very large signal from the solvent, which is often at a much larger concentration than the solute and might overpower it; and to provide a signal from the deuterium atoms to "lock" the spectrometer's frequency.

The spectrometer's field is not constant over time but slowly diminishes due to tiny dissipative losses in the superconducting wire. A typical magnet might drift by 10^{-7} Tesla/hour. This might not sound like a lot, but in reality it translates to

$$\gamma \times 10^{-7} T \sim 1 - 10 \; \tfrac{\text{Hz}}{\text{Hour}} .$$

Some NMR experiments are left overnight for many hours to increase the SNR or simply because they are complicated and take a lot of time (see the lecture about 2D NMR). This amounts to drifts much larger than the linewidth and can lead to severe spectral issues.

To overcome this, the signal from the deuterium – which is completely independent from the hydrogen/carbon/nitrogen/phosphorous signals one usually measures in NMR – is acquired in rapid pulses and used to track the field's drift by looking at the frequency of the deuterium nuclei of the solvent. When the spectrometer "sees" this changes, it adjusts the current through a ring which creates a homogeneous main field much like B_0, only not superconducting. This is fine because the changes this ring needs to make are very small, so we don't need it to be superconducting. Samples without any deuterium cannot "lock" the spectrometer's frequency and this should be kept in mind when running long experiments.

Calibrating the Excitation Pulse

As an experimentalist, we can only vary the voltage on the transmitter. How does one give a 90° pulse? That is, how does one know which voltage corresponds to such a pulse? The answer is we need to calibrate it.

The flip angle α is proportional to the B_1 field:

$$\gamma B_1 t_p = \alpha$$

and B_1 is proportional to the applied voltage by the fundamental equations of electrodynamics, known as Maxwell's equations. We now fix t_p at a very short duration and start increasing the voltage in constant steps, looking at a particular peak in the sample. Often this is the solvent which gives off the strongest value when unsuppressed. What we'll get is something that looks like this:
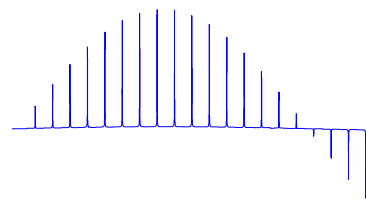
Each peak represents a separate experiment with a different B_1 (which you don't know). The peak amplitude is modeled by

$$A \cdot \sin(\alpha) = A \cdot \sin(\gamma B_1 t_p) .$$

By fitting the maxima of the peaks with this function you can easily find both A (which is meaningless) and B_1, and determine which voltage corresponds to it:
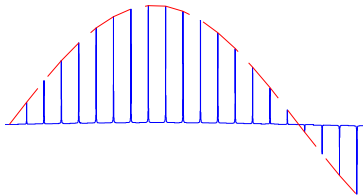

Averaging and SNR

The **signal to noise ratio** (SNR) of an NMR spectrum is one of its most important aspects, particularly because peak strengths are so weak in NMR and often get swallowed up in the noise, becoming unobservable. This is mainly because nuclear paramagnetism is a very weak effect. Because we (usually) can't control the paramagnetic polarization, we have to average over many measurements. The idea is that each measurement has a signal and noise:
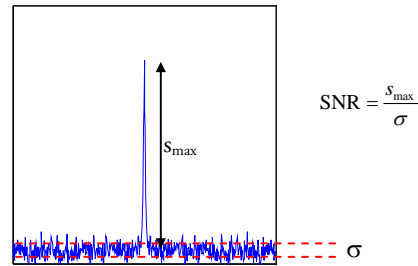
$$s(t) = s_{actual}(t) + n(t)$$

where n(t) is some random noise term:
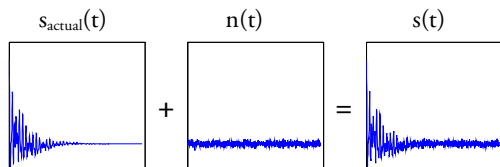
Because the Fourier transform is linear, the same thing happens in the spectral domain (the noise term there will be the FT of the noise term in the time domain, which is ... also noise!):
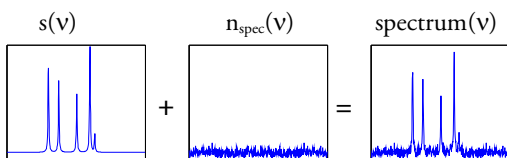
The SNR of a given peak is defined as the ratio of its amplitude to the standard deviation of the noise:

$$SNR = \frac{s_{max}}{\sigma}$$

Two independent measurements will have exactly the same signal $s_{actual}(t)$ but the noise term will be different. What happens when we add them together? The signal doubles in intensity. What happens to the noise? Noise + noise still equals noise, but remember these are random signals: some of the time the signals will cancel out, some of the times they will add constructively, so we won't really get a factor of ×2 in the standard deviation. What we actually get is a factor of $\sqrt{2}$, so the SNR grows by $\sqrt{2}$ as well:

$$SNR = \frac{s_{max}}{\sigma} \xrightarrow{\text{avg. 2 signals}} \frac{2 \cdot s_{max}}{\sqrt{2}\sigma} = \sqrt{2} \cdot SNR \ .$$

In general, for N averages, the SNR will increase by a factor of $\sqrt{N}$. This is the principal of signal averaging. It's not very efficient. For example, if we repeat the same experiment 100 times, we only get ×10 SNR but have to spend ×100 time. Unfortunately, it's often the best we can do.

Shimming

The quality of the spectrum depends greatly on the macroscopic homogeneity of the B_0 field. In inhomogeneous field will lead to a spatial distribution of larmor frequencies (say, for one chemical shift):

$$\omega(\mathbf{r}) = (1 - \sigma)\gamma B_0(\mathbf{r}) \ .$$

Our spectrum will therefore contain an integral over all these peaks:

$$s(t) = \int_{sample} e^{i\omega(\mathbf{r})t - t/T_2} d\mathbf{r} \ .$$

Its Fourier transform will consequently look distorted. When is B_0 inhomogeneity an issue?

When the range of frequencies it creates is wider than a linewidth, which is about 1 Hz in liquid state NMR. Think about what sort of amazing feat it is to achieve this level of homogeneity: we require that

$$\Delta\omega = 1\,\text{Hz}$$
$$\omega = \gamma B_0 = 500\ \text{MHz}$$
$$\frac{\Delta\omega}{\omega} \sim 10^{-9}$$

We require the field to be homogeneous to about 1 part per billion! Think of building a wall 10 cm thick that is so straight it does not deviate (say, by shear forces) by even 10^{-9} of its thickness, which is 1 Å! This is an amazing feat of engineering. While today's modern B_0 main coils can produce a field that's homogeneous to about 10^{-6} over the sample size, further improvements are achieved via **shims**. There are two types of shims:

- **Passive shims** are small ferromagnetic (e.g. iron) elements placed inside the magnet's bore during construction to cancel out spatial inhomogeneities. The shims produce spatial fields which are specifically engineered to cancel out imperfections in the main coil.
- **Active shims** are conducting loops of wire placed around the sample. Current passed through them will generate spatially varying magnetic fields. By adjusting the levels of current we can build spatial patterns that cancel out (some of) the remaining spatial inhomogeneity.

Q: Why do we need active shims? Why aren't passive shims enough?

A: Very simple. Most samples have bulk atomic diamagnetism which will depend on (i) the sample's shape and (ii) composition. This means our magnetic field will be distorted **by the sample itself** and we can't account for it beforehand because we don't know what sort of samples the user will want to test! Active shims let the user fix those sample-specific effects.

Most active shims produce spatial fields which approximate linear combinations of spherical harmonics $Y_{lm}(\phi,\theta)$. The first few are:

| Order (l)	Deg. (m)	Polar function	Cartesian function	Symbol
0	0	1	1	Z^0
1	0	$r\cos\theta$	z	Z
1	1'	$r\sin\theta\cos\phi$	x	X
1	1	$r\sin\theta\sin\phi$	y	Y
2	0	$r^2(3\cos^2\theta-1)$	$2z^2-(x^2+y^2)$	Z^2
2	1	$r^2\sin\theta\cos\theta\sin\phi$	zx	ZX
2	1'	$r^2\sin\theta\cos\theta\cos\phi$	zy	ZY
2	2	$r^2\sin^2\theta\cos2\phi$	x^2-y^2	X^2-Y^2
2	2'	$r^2\sin^2\theta\sin2\phi$	2xy	XY
3	0	$r^3(5\cos^3\theta-3\cos\theta)$	$2z^3-3z(x^2+y^2)$	Z^3
...	...	...	...	...

There are 2l+1 shims of degree l, which are linear combinations of spherical harmonic functions. In theory, if we had an infinite number of shims of orders l=0 until ∞ we could approximate **any** spatial inhomogeneity. In reality:

- We only have a limited number of orders. NMR spectrometers usually have shims up until orders 5 or so, and often not the full set (it might be missing degrees).
- We are limited with the amount of current we can pass through the shim coils.
- Actual shim coils' fields deviate from the perfect spherical harmonics, which complicates things a bit.

Shimming however is a major preliminary part of any NMR experiment, in which one adjusts the active shims to minimize B_0 inhomogeneity. There are many ways to assess the level of B_0 inhomogeneity, but on the most basic level we adjust the shim currents until our peak looks lorentzian and narrow.

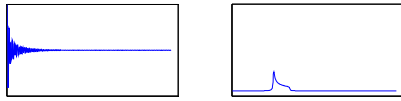The effects of different "types" of field inhomogeneity are shown below.

1. A homogeneous B_0 field and a perfect lorentzian (left: real part of FID; right: real part of spectrum):
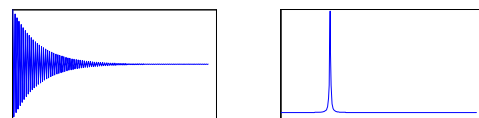
2. A linear inhomogeneity $B_0=\alpha z$: this yields a "beating" in the FID and a loss of signal:

3. A quadratic inhomogeneity, $B_0=\alpha z^2$, will lead to an asymmetric spectral lineshape:

All types of B_0 inhomogeneity lead to **loss of SNR and spectral resolution**, and are extremely detrimental to spectroscopy!

Setting the Filter Bandwidth

Before the ADC there is a low pass filter (LPF) which cuts off high frequencies. Why? Because if it didn't, the high frequencies – which contain only noise – would alias into the spectrum and increase its noise level. The width of this filter is called the filter bandwidth. For the types of electronic/white noise present in the NMR bands, one can say with great precision that
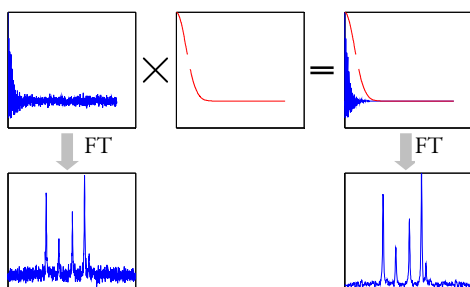
$$\text{FID noise levels} \propto \sqrt{\text{fbw}} \ .$$

The filter bandwidth is directly related to the spectral bandwidth, and in some spectrometers they are implicitly set to be equal.

Practical Aspects of NMR: Processing a Spectrum

Apodization

Apodization means multiplying the FID by a decaying envelope before performing the FT:

Apodization has two effects: it decreases the noise and hence increases the SNR, since it "kills off" the tail of the FID which is usually more dominated by noise than signal; but it also widens the peaks, because it makes the signal decay faster, meaning in increases the effective T_2, reducing spectral resolution.

Zero-filling

Another trick used in NMR post processing is known as zero filling: adding zeros to the end of the FID. This seemingly innocent *extrapolation in the time domain* action is quite useful, and can be shown mathematically to be equivalent to *interpolation in the frequency domain*. Note that zero filling does not change the ADC dwell time and therefore does not change the spectral width.
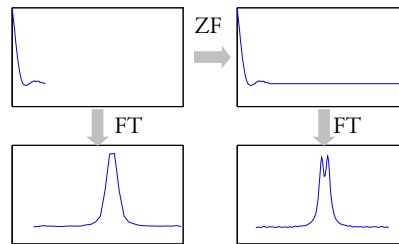
In the above example there were two frequencies present at -0.5 Hz and 2.5 Hz with T_2=125 ms, acquisition time of 512 ms and 32 points. We miss out on the two frequencies not because of T_2 but because of the total acquisition time, i.e. our digital resolution. Zero filling magically made them appear! ZF is no regular linear interpolation in which we "connect the dots" but a special type of interpolation known as "Dirichlet interpolation" which on some mathematical level is ideal for NMR. This magical property increases the resolution by up to several percent to several tens of percent. We won't go into the math of why this happens, but it's almost always a good idea to zero fill a spectrum to twice its size before applying a FT. Zero filling of more than a factor of 2-4 is usually meaningless and should be avoided.
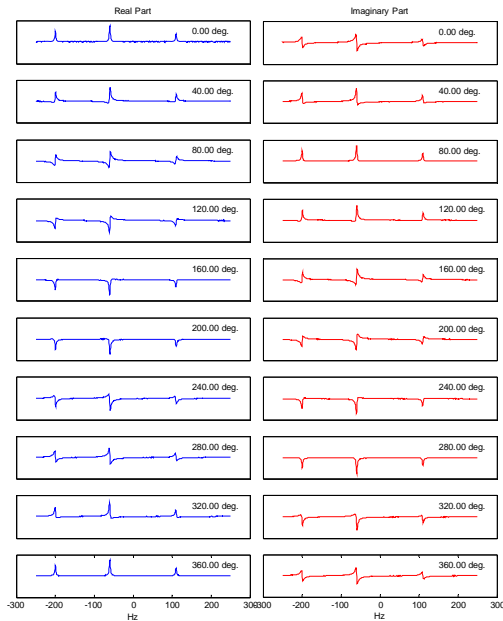
Phasing

Due to hardware constraints, the peaks can have a zero or first order phase. A zero-order phase is a term of the form: $s(t)e^{i\phi}$. Without it, the FID transforms into absorptive (lorentzian) and dispersive parts:

$$s(t) \xrightarrow{\ FT\ } A(\nu) + iD(\nu)$$

With it, the two parts "mix":

$$s(t)e^{i\phi} \xrightarrow{\ FT\ } \left[A(\nu)+iD(\nu)\right]e^{i\phi}$$
$$=\left[A(\nu)\cos(\phi)-D(\nu)\sin(\phi)\right]$$
$$+i\left[A(\nu)\sin(\phi)+D(\nu)\cos(\phi)\right]$$

This looks like this:

Luckily, 0th order phase is easy to correct: just multiply the spectrum by $e^{-i\phi}$. The phase ϕ is not known a-priori so the correction is usually done manually, and terminated when the operator deems his real spectrum "looks absorptive". Alternatively, sophisticated algorithms can do a pretty good job of automating this correction.

First order phases come about due to electronic imperfections as well as finiteness of the RF pulses. Here, a frequency-dependent phase gets added to the peaks. For example, if you have N peaks with frequencies ω_j then your FID will be:

$$s(t) = \sum_{j=1}^{N} A_j e^{i\omega_j t} e^{i\alpha\omega_j} e^{-t/T_2} = \sum_{j=1}^{N} A_j e^{i\omega_j(t+\alpha)} e^{-t/T_2} \ .$$

This seems almost trivial to correct: why not multiply the FID by $e^{-i\alpha\omega}$? But wait: what is ω? We don't know our frequencies a-priori, and even if we did, we can't "access" each summand and fix it independently of the others. When we look at the effects of 1st order phase in the frequency domain, is will look something like this:
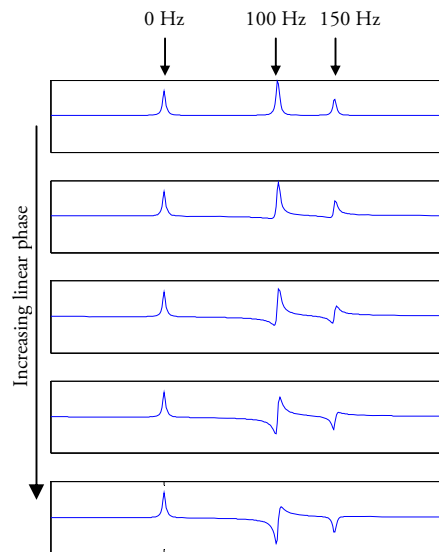
You see the phase at 0 Hz is not affected at all, because its frequency is $\omega=0$ and therefore its linear phase is always $\omega\cdot\alpha=0$ regardless of α. The effect becomes more and more pronounced for peaks farther away from 0 Hz.

Q: Why can't we fix the linear phase by multiplying the spectrum by $e^{i\omega\alpha}$ for some α?

A: Each peak has a constant phase that increases linearly with frequency. This is **not** the same as having a linear phase for the entire spectrum. To illustrate this, just look at what happens to (the real part of) the spectrum when we multiply it by $e^{i\omega\alpha}$:
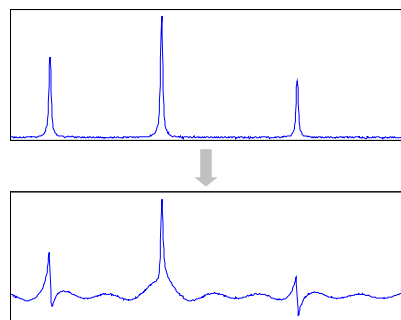
Quite horrible! This happens because $e^{i\alpha\omega}$ affects **the entire lineshape** and not that lineshape's overall phase! You're trying to fix a problem that looks like this:

$$L\left(\omega-\omega_{j}\right)e^{i\alpha\omega_{j}}$$

by doing this:

$$L\left(\omega-\omega_{j}\right)e^{i\alpha\omega_{j}}e^{-i\alpha\omega}$$

where $L(\omega)$ is a Lorentzian peak.

Of course, a real spectrum will have both 0^{th} and 1^{st} order phase issues, and good luck telling them apart! It takes skill, or a good computerized algorithm.