# II
Spins and their Dynamics

Lecture notes by Assaf Tal

## 1. Excitation

1.1 Why Excite?

Let us recant the facts from the previous lecture:

1. We're interested in the bulk (macroscopic) magnetization.
2. When put in a constant magnetic field, this bulk magnetization tends to align itself along the field:


3. The macroscopic magnetization precesses about the external magnetic field.

In MRI, we can only detect a signal from the spins if they precess and therefore induce a current in our MRI receiver coils by Faraday's law. However, this poses a problem: when we put our patient in an MRI machine, the spins in his/her body align along the field. They remain that way indefinitely, so no precession will occur.

A schematic representation of the bulk magnetization vectors in our body in the presence of an external field, at equilibrium.

For the spins to precess, and for us to detect them, we need to somehow force them away from equilibrium – for example, make them perpendicular to the main field:

If we do this and let them be, two things will happen:

1. They will precess about B_0 at a rate[1] γB_0.
2. Eventually they will return to equilibrium because of thermal relaxation. This usually takes ~ 1sec, but can vary.

Until they return to thermal equilibrium their signal is "up for grabs", which is precisely the idea of a basic NMR/MRI experiment:

1. Excite the spins.
2. Measure until they return to equilibrium.

What to exactly **do** with the measured signal and how to recover an image from it is something we'll discuss in subsequent lectures. Meanwhile, let's just ask ourselves *how can one excite a spin*? That is, how can one tilt it from its equilibrium position along the main field and create an angle between them (usually 90 degrees, but not always)?

1.2 The RF Coils

Fortunately (or perhaps unfortunately), our question has been answered a long, long time ago:

> To excite a spin, irradiate it with an external, perpendicular magnetic field at its resonant frequency γB_0.

This is called "on-resonance irradiation". The external resonant field is achieved using an external

---

[1] For correctness we should add the chemical shift here, but for now we'll just focus on water and incorporate the shift later. In fact, everything we say here can be easily generalized by replacing B_0 by $B_0 + \Delta B$. For water, ΔB is a well known quantity.

RF coil built into the MRI machine. These come in various shapes, such as birdcage:

Bird Cage Coil

This coil is ideally capable of generating a homogeneous, time-dependent RF field in the transverse plane (transverse to B_0), often called the B_1-field:

$\mathbf{B}_{RF}$

$\mathbf{B}_0$

The RF field. This picture emphasizes two important points: (i.) B_{RF} is always perpendicular to B_0. (ii.) It is always much smaller than B_0 (in magnitude). Usually, $|\gamma B_{RF}| \sim$ kHz, while $|\gamma B_0| \sim$ MHz. The birdcage coil envelopes the object in question, much like the main field's coils.

The RF field looks like this, analytically:

$$\mathbf{B}_{RF}(t) \equiv B_1\left(\cos\left[\phi_{RF}(t)\right]\hat{\mathbf{x}} + \sin\left[\phi_{RF}(t)\right]\hat{\mathbf{y}}\right)$$

Amplitude   Phase of RF

with

$$\phi_{RF}(t) = \omega_{RF}t + \phi_0$$

Frequency of RF   Constant Phase of RF

Some call the entire $\phi_{RF}(t)$ "the phase of the RF", and others refer to just ϕ_0 as the same. In this course we'll associate ourselves with the first group. When I'd like to make it clear I'm talking about ϕ_0, I'll call it the "constant phase" (it is constant, after all).

The total magnetic field felt by a spin is, therefore:

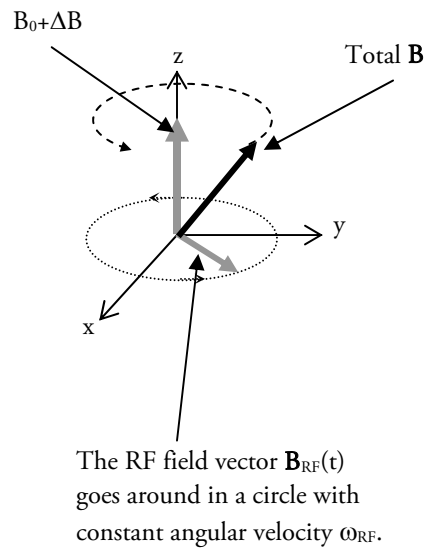$$\mathbf{B}(t) = B_0\hat{\mathbf{z}} + \mathbf{B}_{RF}(t)$$

Main field.   RF field.
Along z-axis.   In xy-plane.
Constant.   Time-dependent.

One often denotes

$$\omega_1 = \gamma B_1$$

The geometric meaning of ω_{RF} can be understood by simply plotting $\mathbf{B}_{RF}(t)$ as a function of time. We then find out that ω_{RF} is the angular frequency of the RF field vector in the xy-plane:

$B_0 + \Delta B$ Total $\mathbf{B}$

z

y

x

The RF field vector $\mathbf{B}_{RF}(t)$ goes around in a circle with constant angular velocity ω_{RF}.

To make this even clearer, let's plot the entire field felt by the spins (each spin feels the same field), which equals the constant B_0 field + the rotating B_{RF} field:

$\mathbf{B}_{RF}$

$\mathbf{B}_0$

1.3 THE ROTATING FRAME

Remember a nuclear spin precesses in a magnetic field B with a frequency given by:

$$\omega = \gamma B$$

In particular, in the main field, B_0, it precesses with the following angular frequency about the z-axis:

$$\omega_L = \gamma B_0 \quad \text{(Larmor Frequency)}$$

There is also the chemical shift, but we'll neglect it for now for simplicity, and come back to it a bit later. What does it mean "to irradiate a spin at its resonant frequency"? It means we need to match the frequency of the RF field to that of the spin:

$$\omega_{RF} = \omega_L$$

Why this works can be best understood by looking at things in a frame that rotates along with the RF field vector, called **the rotating frame**. This is sometimes called "transforming to the rotating frame".

To understand how to perform this transformation, we take an analogy from mechanics. Image 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.

Now 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$_{\text{gravity}}$=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:

$\omega_{rot} = \omega_{RF}$

1. Understand how the magnetic field's vector in the "lab" frame would look in the rotating frame (it would be stationary). The black vector is the total RF field:

In the rot. frame, the RF is constant

2. Add a "fictitious field", $-\gamma\omega_{rot}\hat{\mathbf{z}}$, which – since we've chosen $\omega_{RF} = \omega_L$ - has size $\gamma\omega_{rot} = \gamma\omega_L$, i.e., it's the same size as the main field but opposite in sign and hence cancels it (black vector = total magnetic field):

Adding the fictitious field cancels out the main B_0 field

So *transforming to the rotating frame makes our RF constant and in the xy-plane[2].*

---

[2] Note that it may also have a constant *phase* ϕ_0, meaning it won't point along the x-axis (but still be constant)

Note that the RF can have a constant phase, ϕ_0, as noted above:

$$\phi_{RF}(t) = \omega_{RF}t + \phi_0$$

Frequency   Const. phase

This phase means that, in the rotating frame, the RF itself will have a phase as well:

ϕ_0

Lab frame
(Rotating)

Rot frame
(Constant, no z)

In the lab frame our spins start off along the z-axis. This also holds in the rotating frame:

B_1

(in the rotating frame)

What happens to a magnetic moment in a magnetic field B_1? It precesses around it with an angular frequency $\omega_1 = \gamma B_1$!


$\omega_{precess} = \gamma B_1$

So here is our plan for getting the spins perpendicular to the main field:

1. Irradiate the spins on resonance.
2. Wait for enough time until the spins reach the xy plane (complete a quarter of a circle).
3. Turn off the RF field.

All this happens in the rotating frame. How do things look back in the lab frame? This takes a bit of imagination to see:

Rot. Frame   Lab. Frame

I'll leave that one up to imagination (you need to superimpose the rotation of the frame on the precession).

The angle the spin precesses by is given by

$$\theta = \omega t = \gamma B_1 t$$

Therefore the resonant RF field should be applied for a time

$$\pi/2 = \gamma B_1 t_{\pi/2}$$

that is

$$t_{\pi/2} = \pi/2\gamma B_1$$

Note we still haven't said how to choose the magnitude of B_1. The stronger B_1, the shorter $t_{\pi/2}$. We'll usually desire the shortest pulses possible. One reason is that the longer we wait, the more time we waste and the more the thermal relaxation effects become a nuisance.

Some numbers are in order. For typical MRI scanners,

$B_1 \sim \mu T$

$\gamma = 2\pi \times 42$ MHz/T

yielding

$$t_{\pi/2} = \frac{\pi}{2\gamma B_1} = \frac{1}{4 \times 42 \frac{MHZ}{T} \times 10^{-6} T} \sim 0.5 \text{ ms}$$

which is the "right order of magnitude". Exact values will vary depending on the hardware.

The rotating frame is of utmost importance in MRI because of the following fact:

> **MRI takes place in the rotating frame.**

This is a result of the way the MRI hardware works. All signals we measure are given **as if we've measured them in the rotating frame**. You'll have to take that as an axiom. Therefore, unless otherwise specified, we shall continue "living" in this rotating frame and forget about the lab frame, except at special points which I'll mention.

1.4 A Short Summary

We summarize:

1. MR is easier to understand in a rotating frame.
2. MR "happens" in the rotating frame.
3. The rotating frame's frequency ω_{rot} is always chosen to coincide with the RF's frequency, ω_{RF}.
4. The rotation introduces a "fictitious field" $-\omega_{rot}/\gamma \mathbf{z} = -\omega_{RF}/\gamma \mathbf{z}$.

It is important to make a distinction between all the ω's we've collected so far.

Spin

$\omega_0 = \gamma B_0$	Precession due to main field
$\omega_{cs} = \gamma B_{cs}$	Chemical shift
$\omega_L = \omega_0 + \omega_{cs}$	The Larmor frequency

RF

$\omega_1 = \gamma B_1$	Strength of RF field (in Hz)
ω_{RF}	Frequency of RF field

Rot. Frame

|	
$\omega_{rot} = \omega_{RF}$	Rotating frame angular velocity
$-\omega_{rot}/\gamma$	The "fictitious field" introduced

1.5 THE OFFSET

So far we've assumed

$$\omega_{RF} = \omega_L$$

Pictorially:


However, this isn't possible always when you have more than one chemical species in your sample. For example, fat and water have different ω_L. You can't irradiate on resonance when you have more than one resonant frequency!

Example: on resonance for fat "off-resonance" for water

You can **set** the irradiation frequency on the MRI machine. A "reasonable" choice might be in the middle:

This is called **off-resonance irradiation**, and we need to understand how to think about it. Before we do, some terminology: the difference

$$\Delta\omega = \omega_L - \omega_{RF} = \omega_L - \omega_{rot}$$

is called the **offset** of the spins. Each species has a different offset. For example:

$$\Delta\omega^{(water)} = \omega_L^{(water)} - \omega_{RF}$$
$$\Delta\omega^{(fat)} = \omega_L^{(fat)} - \omega_{RF}$$

Offsets, much like the chemical shifts, are "small": they usually range in the Hz-kHz ranges.

1.6 OFF-RESONANCE IRRADIATION

How do "off-resonance" spins behave in the rotating frame? Let's pick a **particular** spin with a **particular** cs and Larmor frequency. Remember the rotating frame rotates with the same frequency as the RF field, so the RF is constant, and remember it adds a term $-\omega_{RF}/\gamma$ to the z-field.

In lab frame:

$$\mathbf{B}(t) = (B_0 + B_{cs})\hat{\mathbf{z}}$$
$$+ \underbrace{B_1(t)\left[\cos(\omega_{RF}t)\hat{\mathbf{x}} + \sin(\omega_{RF}t)\hat{\mathbf{y}}\right]}_{\mathbf{B}_{RF}(t)}$$

How does **B** look like in the rotating frame?

	Field in Lab.	Field in Rot.
z Axis	$(B_0 + B_{cs})\hat{\mathbf{z}}$	$\left(B_0 + B_{cs} - \dfrac{\omega_{rot}}{\gamma}\right)\hat{\mathbf{z}}$
xy-plane	$\mathbf{B}_{RF}(t)$ rot.	$B_1\hat{\mathbf{x}}$

But!

$$B_0 + B_{cs} - \frac{\omega_{rot}}{\gamma}$$

$$= B_0 + B_{cs} - \frac{\omega_{RF}}{\gamma}$$

$$= \frac{\gamma\left(B_0 + B_{cs}\right)}{\gamma} - \frac{\omega_{RF}}{\gamma}$$

$$= \frac{\omega_L - \omega_{RF}}{\gamma} = \frac{\Delta\omega}{\gamma}$$

So:

So on-resonance ($\Delta\omega=0$), the RF field is along the x-axis (or another axis if it has a phase). Off-resonance spins get an additional small field $\Delta\omega/\gamma$ added. This field is called the effective field:

$$\boxed{\mathbf{B}_{eff} = \frac{\Delta\omega}{\gamma}\hat{\mathbf{z}} + B_1\hat{\mathbf{x}}}$$
The effective field

The "basic rule" of spin dynamics still applies: a spin in a magnetic field B precesses around it with an angular frequency γB. All spins start out along the z-axis and precess about the effective field. Note that since the effective field's magnitude is different for different spins, the precession's angular velocity $\omega=\gamma|\mathbf{B}_{eff}|$ also depends on the spin's chemical shift/offset.

Some graphical examples:

On resonance   Small offset   Large offset
                $\Delta\omega<<\gamma B_1$ $\Delta\omega>>\gamma B_1$

1.7 THE "HARD" $\pi/2$ PULSE

The question: how can we excite all the spins, both on and off resonance?

The answer: irradiate for a short amount of time as possible with a strong-as-possible RF (strong = γB_1>>all offsets $\Delta\omega$).

The rationale: as the graphical examples in the previous section have shown, when the offset is "small" (compared to the RF), the spin almost reaches the xy-plane. *This is good enough for us.* Thus, **we need to make sure the RF is strong enough for the offsets to be small enough to be negligible!**

This short, strong RF irradiation is called a hard $\pi/2$ pulse:

Hard = affects all offsets similarly
$\pi/2$ = brings them to the xy plane (almost!)
Pulse = MRI jargon for RF irradiation

Some notes:

1. It is possible to apply hard RF pulses with different constant phases. This tips the magnetization vector onto different positions in the xy plane. One usually denotes: x = 0°, y = 90°, −x = 180°, −y = 270°. Also, $\pi/2$ is sometimes called a "90-pulse" (unsurprisingly, since it tilts the magnetization by 90 degrees onto the plane).

2. It is possible to irradiate for longer amounts of time and create hard 180 pulses (a.k.a. π-pulses), 270-pulses, and any angle in between.

3. Notational convention: a hard $\pi/2$ pulse with a phase ϕ is often denoted 90_ϕ. Some examples (starting from thermal equilibrium magnetization, & remember the left hand rule!):

90_x 90_{-y} 180_y

(All examples are presented in the rotating frame)

2. Life in the Transverse Plane

2.1 Precession

Once we excite the spins onto the xy-plane and leave them there, they will precess about the effective field. When we turn off the RF, we're only left with the offset's field:

Thermal eq.   Hard $\pi/2$ pulse   Precess in xy with $\Delta\omega$

When viewed "from above", the spin moves in a circle in the xy (transverse) plane:

The **phase** of the spin.

The more we wait, the larger θ becomes:

$$\theta(t) = \Delta\omega \times t.$$

Different chemical species will precess at different angular velocities because of their different offsets, so after the same amount of time they'll point in different directions:


After excitation
water (light blue)
fat (dark blue)

2.2 Thermal Effects

The thermal effects will eventually cause the spin to return to thermal equilibrium along the z-axis. This usually takes ~100 ms to ~ 1 sec. We will discuss them and their origins next lecture.

2.3 Hard π-Pulses

Once in the plane, there is no rule forbidding us from applying additional pulses! For example, we can apply an RF pulse along the x-axis for long enough to induce a π rotation. We make it strong enough to be "hard" – that is, to render the offset insignificant:

(a.) Thermal eq.   (b.) Excite

(c.) Wait (precess)   (d.) π-pulse

This pulse reflects the spin about the x-axis. Here's what happens to the spin after it is excited by a hard $(\pi/2)_x$ pulse, left alone for a time T and then acted upon by a π_y pulse:

(a.) Spin is excited π_y (b.) After time T

(c.) After hard π_y pulse   (d.) After time T
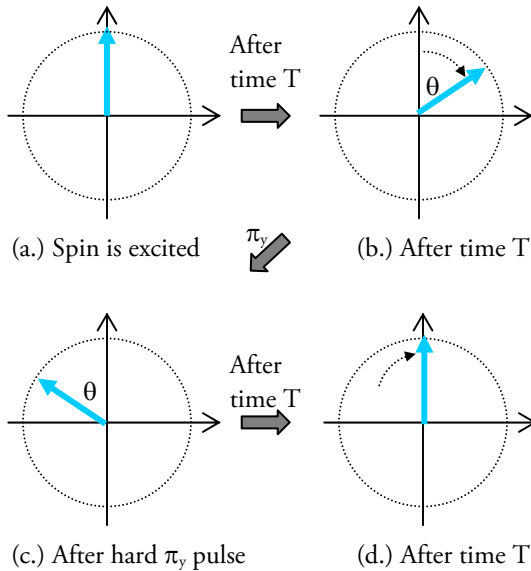
So the spin returns to its original position. This is why π pulses are sometimes said to "reverse time".

2.4 Complex Numbers: A Refresher

It is sometimes useful to switch to complex notation (complex as in "incorporating the imaginary number i"). Recall that i is defined via

$$i^2 = -1$$

Complex numbers are entities of the form a+ib, where a & b are "regular" numbers. Complex numbers are usually denoted by the letter z. For example:

z=4+i5
z=-5+iπ

The trick about complex numbers is to treat them precisely the same way you'd treat regular numbers, by treating i as a variable (just think of it as x) with the special property i^2=-1. For example, addition & multiplication:

z_1=3 + i4, z_2=-2 + i6
z_1+z_2 = 3 + i4 − 2 -6i = 1 +i(4-6) = 1-2i
$z_1 z_2$ = (3 + i4)(-2+i6)
    = 3×(-2)   + 3×i6 + i4×(-2)   + i4×i6
    = -6   +18i   - 8i   +24i^2
    = -6   +18i   -8i   -24

= -30 + 10i

a & b are called the **real** and **imaginary** parts of the complex number, respectively:

$$z = \underset{real}{4} + \underset{imaginary}{8i}$$

Some other things you should know about complex numbers:

1. It doesn't matter whether you write 8i or i8. Remember: think of i as a variable, x. Does it matter if you write 8x or x8? They both mean 8 times x.

2. z =0 can only hold if both its real and imaginary parts equal 0. For example, suppose someone tells you z= x + i(y-4), and that z=0. Then x=0 and y=4.

One very important identity that ties in a lot of different ideas in mathematics is **Euler's identity**:

$$e^{ix} = \cos(x) + i\sin(x)$$

It's very surprising the first time you see it, since it relates two completely different ideas (e and trigonometry). To prove it we need calculus, and we won't be doing it in this course, so you'll have to take my word for it. Because e is an exponent we can use the familiar exponent identities:

$$e^a e^b = e^{a+b}$$
$$\left(e^a\right)^b = e^{ab}$$

These identities hold for any numbers, not just e (for example: $32 = 4 \times 8 = 2^2 \times 2^3 = 2^{2+3} = 2^5$). They also hold for any a & b, even if they are complex numbers! For example:

$$e^{ix} e^{iy} = e^{ix+iy} = e^{i(x+y)}$$

As a fun example[3], you'll be using this in the HW to prove a trigonometric identity.

Complex numbers can be viewed as vectors in the real-imaginary plane, which is just like the xy-

---

[3] This is fun. Trust me.

plane. If z=a+ib, then a is its "x-coordinate" and b is its "y-coordinate":

Once we start to think about complex numbers as vectors, we can start asking vector-related questions:

1. What is the complex number's size (also known as its "magnitude")?
2. What is the angle it makes with the x-axis? This angle is known as the complex number's phase.

We can use basic trigonometry & Pythagoras's theorem to answer these questions:

$$|z| = \sqrt{a^2 + b^2}$$

$$\tan\phi = \frac{b}{a}$$

$$\Rightarrow \quad \phi = \arctan\left(\frac{b}{a}\right)$$

Note that we can completely specify this complex "vector" (number) in two equivalent ways:

1. Specify its real, a, & imaginary, b, parts.
2. Specify its magnitude, |z|, and phase, ϕ.

2.5 COMPLEX MAGNETIZATION

A magnetization vector has three components, M_x, M_y, M_z. Once tilted onto the xy plane (along the x-axis) it will precess. During this precession, M_z will remain constant and M_x and M_y change according to

$$M_x(t) = M_0 \cos(\omega t)$$
$$M_y(t) = -M_0 \sin(\omega t)$$
$$M_z = 0$$

So as long as no RF fields are involved, M_z remains constant and only M_x, M_y change. Instead of dealing with M_x and M_y separately, it's convenient to think of the magnetization in the xy-plane as a complex number M_{xy}.

$$\mathbf{M} = M_0 \cos(\omega t)\hat{\mathbf{x}}$$
$$\quad - M_0 \sin(\omega t)\hat{\mathbf{y}}$$

$$M_{xy} = M_0 \cos(\omega t)$$
$$\quad - iM_0 \sin(\omega t)$$

The good thing about the complex representation is that we can use Euler's identity to go from the cartesian to the polar representation:

$$M_{xy} = M_0 \cos(\omega t) - iM_0 \sin(\omega t)$$
$$= M_0 \left[\cos(\omega t) - i\sin(\omega t) \right]$$
$$= M_0 \left[\cos(-\omega t) + i\sin(-\omega t) \right]$$
$$= M_0 e^{-i\omega t}$$

This representation has no conceptual advantage over the use of vectors, but it does make a lot of calculations much easier. M_{xy} – whether in vector or complex form - is called the transverse magnetization. Whenever presented with a complex transverse magnetization, you can easily recover the x & y components of the magnetization vector by taking the real and imaginary parts of M_{xy}, respectively. For example, given $M_{xy}(t) = M_0 e^{-i\omega t}$, you would use Euler's identity and identity M_x and M_y as follows:

$$M_{xy}(t) = M_0 \left[\cos(-\omega t) + i\sin(-\omega t) \right]$$
$$= M_0 \left[\cos(\omega t) - i\sin(\omega t) \right]$$
$$= \underbrace{M_0 \cos(\omega t)}_{M_x} + i\underbrace{\left[-M_0 \sin(\omega t) \right]}_{M_y}$$

3. Signal Acquisition

3.1 The Physics

We've gathered so far that:
1. The spins align along the main B_0 field at equilibrium.
2. We can excite them onto the xy-plane.
3. When left alone, they precess about B_0 with an angular velocity $\Delta\omega$, called their *offset*, which depends on their chemical shift and the angular frequency at which our rotating frame rotates.

All these manipulations are for naught if we can't **observe** the spins. To this end, the MRI and NMR machines have coils, called ***receiver coils***, wound around the main bore of the magnet in a particular configuration, that can pick up the spins' signals.

The physics involved is all contained in Faraday's law. A quick explanation is in order. Imagine a conductor loop, placed in a magnetic field, and linked to a resistor.


Loop placed in a magnetic field
(the black arrows).

As it is, no current will flow in the system and the resistor will remain cold to touch.

The ***magnetic flux*** ϕ_m through the loop is defined as the magnetic field times the loop's area. In our case, if the loop's area is A and field is B,

$$\phi_m = AB.$$

If we were to **change** the magnetic flux by changing the field B, we would induce a current in the system and heat the resistor up. This would be as if we've replaced the coil with a battery, or voltage source (e.g. an electrical socket):

This is the essence of Faraday's law:

A changing magnetic flux ϕ_m through an area enclosed by a loop of wire generates a voltage in the loop. This voltage is proportional to the flux's rate of change:

$$V \propto \frac{d\phi_m}{dt}$$

The idea of signal reception is this:
1. The nuclear magnetic moments create magnetic fields.
2. Rotating the moment also rotates the field, changing it with time.
3. If we put a coil around the imaged subject, the changing magnetic field will create a changing flux through the coil.
4. The changing flux will induce an observable voltage.

This is precisely what's done in MRI: the body is "wrapped" in a coil that picks up the voltage induced by the spins (often, the same coil is used to create the RF).

Inside a Philips 3T Magnet

Gradient coils (we'll get to those) Body receiver coil Shim iron (not yet discussed)

3.2 An Expression for the Signal

The above qualitative discussion can be given a quantitative formulation. The exact mathematics are a bit too advanced for this course, so I'll try to hand-wave my way through.

Imagine we have a single spin precessing in the xy-plane:

$$m_{xy} = m_0 e^{-i\omega t}$$

For a single spin ½, you'd substitute $m_0 = \frac{1}{2}\gamma\hbar$. A spin is nothing more than a magnetic moment, and magnetic moments generate magnetic fields. The precise form of the field is complicated, but we can say one thing: because the spin rotates periodically, so should the field it creates[4]:

$$B \propto m_0 e^{-i\omega t}$$

The magnetic flux will somehow be proportional to it:

$$\phi_m \propto B \propto m_0 e^{-i\omega t}$$

and the observed voltage will be:

$$V \sim \frac{d\phi_m}{dt} \sim m_0 \omega e^{-i\omega t}$$

This expression for V is complex. In almost all modern MRI machines, two coils are used to pick up data – this is called quadrature detection. One coil picks up a signal proportional to the real part of the magnetization (which is just M_x), and the other coil picks up a signal proportional to its imaginary part (which is just M_y).

For multiple spins, you'd just have to sum the contributions from the different spins:

---

[4] This is actually not a trivial assumption, and has to do with the long wavelengths involved. This wouldn't work so well with short wavelengths.

$$V \sim m_0^{(1)}\omega^{(1)}e^{-i\omega^{(1)}t}$$
$$+m_0^{(2)}\omega^{(2)}e^{-i\omega^{(2)}t}$$
$$+m_0^{(3)}\omega^{(3)}e^{-i\omega^{(3)}t}$$
$$+\dots$$

It doesn't matter where the spin is, just what its Larmor precession is. In general, if we have some bulk magnetization $\mathbf{M}(\mathbf{r})$ (magnetic moment per unit volume at point $\mathbf{r}$), we can divide the object inside the coil – say, the human patient – in small volumes ΔV and sum over all the volumes:

$$s(t) \propto \sum_i M\left(\mathbf{r}_i\right)\omega\left(\mathbf{r}_i\right)e^{i\omega(\mathbf{r}_i)t}$$

When the volume elements become very small, this turns into an integration:

$$s(t) \propto \int M\left(\mathbf{r}\right)\omega\left(\mathbf{r}\right)e^{i\omega(\mathbf{r})t}d\mathbf{r}$$

It's important to note that all this discussion was not carried out in the rotating frame but rather in the lab frame, so, for example, $\omega \sim 136$ MHz on a 3T imaging machine. What can cause $\omega(\mathbf{r})$ to vary from point to point?

1. Magnet imperfections. Usually small variations, on the order of 1-100 Hz.
2. Field gradients that we create – important for imaging, but still small compared to $\omega_0 = \gamma B_0$ (they're on the order of 0-100 kHz).

Variations in $\omega(\mathbf{r})$ are therefore small, so it's approximately constant, $\omega\left(\mathbf{r}\right) \approx \omega_0 = \gamma B_0$, and it can be taken out of the integral:

$$s(t) \propto \omega_L \int M(\mathbf{r}) e^{i\omega(\mathbf{r})t} \, d\mathbf{r}$$

This cannot be said of $e^{i\omega(\mathbf{r})t}$, because even if variations in ω are small, for long enough times they can become significant, so we'll have to leave it as-is.

4. NMR Spectroscopy

While this is a course about NMR *imaging*, I'd like to describe the fundamental experiment at the basis of NMR spectroscopy for three reasons:

1. It'll put to good use most of the concepts we've learned so far.
2. It will introduce the idea of a pulse sequence.
3. We'll revisit it when we talk about spectroscopic imaging (should time allow).

Suppose you have a sample with water, and it has some offset $\Delta\omega$ in the rotating frame. You first excite the water and then let it precess – it will precess with the angular velocity $\Delta\omega$ in the rotating frame. This rotation will induce a voltage in a coil placed around it (in fact, the same RF coil used to excite it can also be used to measure this voltage, but you can use a different coil as well). The periodic circular precession will induce a periodic, sinusoidal voltage in the coil:

Voltage

$\sim\!\text{Sin}(\Delta\omega\times t)$

By observing this signal you can deduce the offset, $\Delta\omega$. Why is this important? Well, think about a sample containing, say, two substances – say, fat and water – having different offsets, $\Delta\omega$, owing to their different chemical structures. Each would give rise to its own signal with its own periodicity:

Water $\sim \sin(\Delta\omega_{water}t)$

time

Fat $\sim \sin(\Delta\omega_{fat}t)$

time
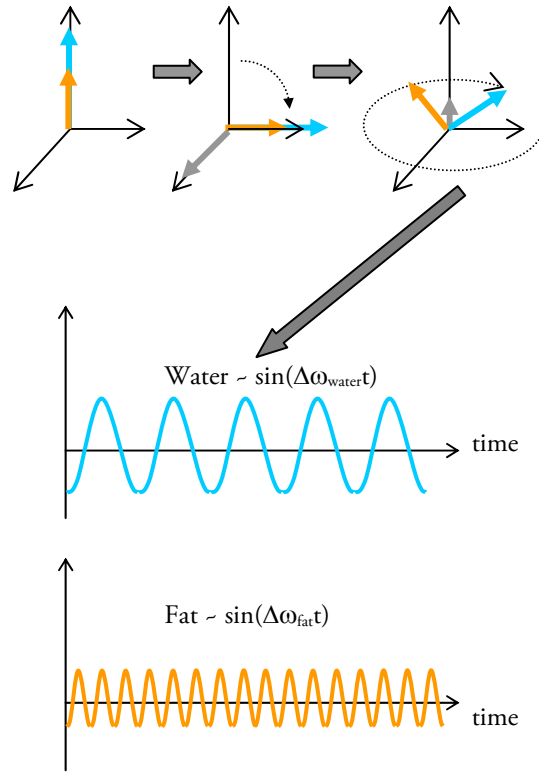
The signal we pick up would be the sum of signals coming from each spin species:

Signal $\sim S_1\sin(\Delta\omega_{fat}t) + S_2\sin(\Delta\omega_{water}t)$

There are post-processing ways for:

1. Uncovering $\Delta\omega_{water}$ and $\Delta\omega_{fat}$ from the total signal.
2. Finding their relative amplitudes (in the above example, the fat (orange) signal was about half the water (blue) signal, indicating in a sense there are less fat molecules in the sample than water molecules).

That's the idea of NMR spectroscopy: you're given a sample with different chemical compounds and you're asked to:

1. Say which compounds there are.
2. Say how much of each compound there is.

The most famous post-processing method is called the Fourier transform. We'll study it in detail in at later date. Now I'd just like to mention that it can decompose a sum of sinuses into "peaks". For the simplified fat+water example I gave above, this Fourier transform – which you can view as a "black box" – would yield a spectrum: