LECTURE 7

T1, T2, T2* CONTRAST

Lecture Notes by Assaf Tal

T_2^* CONTRAST

Sampling k-Space By Reading It Line-By-Line With A Gradient Produces "Gradient Echoes"

We've presented the very basic imaging sequence in the previous two lectures,

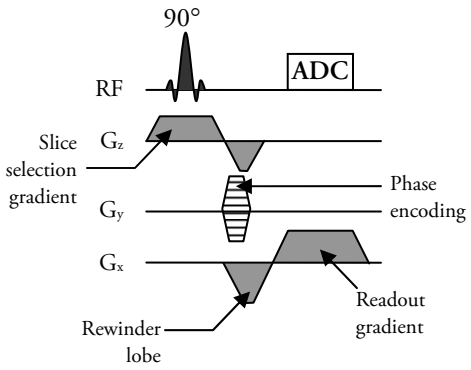

in which a line is acquired in k-space during each excitation:

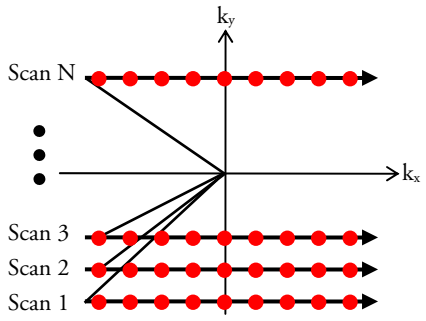

This must be repeated for each excited slice. Let us put aside for the time being the total time required for such a measurement and assume our magnetization is in thermal equilibrium before the start of each k-space line.

Neglecting for a moment the decay of signal due to T_2 and T_2^* effects, the signal as a function of $\mathbf{k}$ is:

$$s(\mathbf{k}(t)) = \int_{\substack{\text{all}\\\text{space}}} M_0(\mathbf{r}) e^{-2\pi i \mathbf{k}(t)\cdot\mathbf{r}} d\mathbf{r} .$$

Recall that measurement doesn't start at t=0 but at some time $t=t_0$ when the ADC opens up. At that point, $\mathbf{k}$ doesn't start from 0 but from some initial value $\mathbf{k}_0=\mathbf{k}(t_0)$, which is the point to which the phase and rewinder gradients take us:

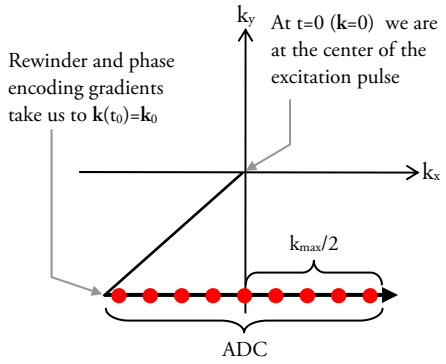

The maximal signal comes from the center of k-space and the lines $k_x=0$ and $k_y=0$. This can be seen by in several ways. First, just by looking at typical k-space data:

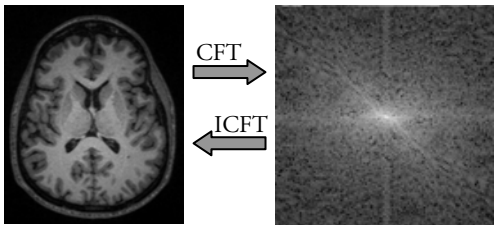

Mathematically, using the fact that $\left|\int f(x)dx\right| \le \int |f(x)|dx$ for any function f(x), we have:

$$\begin{aligned} s(\mathbf{k}) &= \left| \int_{\substack{\text{all}\\\text{space}}} M_0(\mathbf{r}) e^{-2\pi i \mathbf{k}(t)\cdot\mathbf{r}} d\mathbf{r} \right| \\ &\le \int_{\substack{\text{all}\\\text{space}}} \left| M_0(\mathbf{r}) e^{-2\pi i \mathbf{k}(t)\cdot\mathbf{r}} \right| d\mathbf{r} \\ &= \int_{\substack{\text{all}\\\text{space}}} |M_0(\mathbf{r})| d\mathbf{r} = s(\mathbf{0}) \end{aligned}$$

From a intuitive point of view, what we're doing is integrating a function $M_0(\mathbf{r})$ and modulating it by some sinusoid function $e^{2\pi i\mathbf{k}\cdot\mathbf{r}}$. The faster we modulate it, the more it integrates to zero since the positive and negative lobes cancel each other out, assuming $e^{2\pi i\mathbf{k}\cdot\mathbf{r}}$ varies spatially faster than the image:

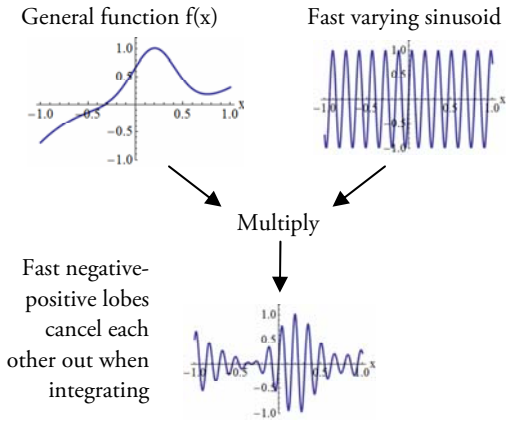

This is why we call this sequence a **gradient echo**: the spins are first *dephased* by taking them to the outskirts of k-space. When we cross the k_x=0 line all of a sudden we get a large amplitude, a so-called gradient echo, which then decays back again as we travel to the other side of k-space. We've seen this mental picture before when discussing the dephasing effect of gradients. If we think of a one dimensional imaging problem and a uniform spatial distribution of spins of size L, the analytical acquired signal is:

$$s(k)=\int_{-\infty}^{\infty}M_0(x)e^{2\pi ikx}dx \sim \mathrm{sinc}(\pi kL)$$

and this is what happens to the spins in the sample after getting excited and "rewound":

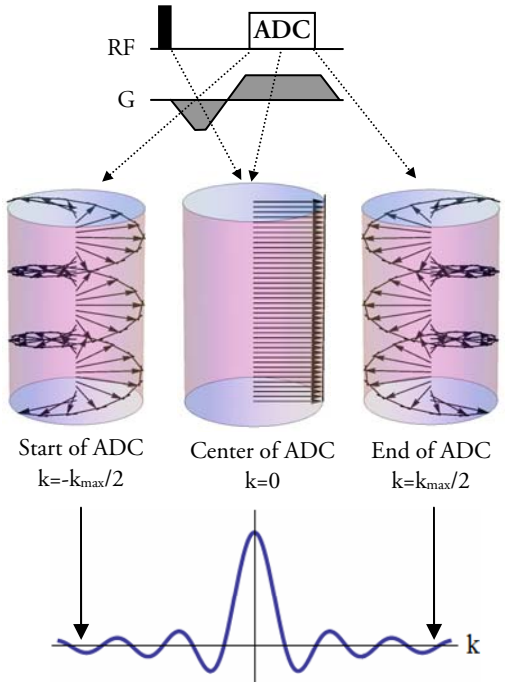

The Gradient-Echo Sequence Produces T_2* Weighting (The Simple Explanation)

We now include relaxation, which occurs via T_2* - the combined result of microscopic T_2 and macroscopic T_2' effects. Relaxation "kicks in" right after excitation. Let us denote by TE the **Time to Echo**; that is, the time between the center of the excitation pulse to the point at which the gradient echo forms at the center of the ADC event (when we cross k_x=0). Then we have:

$$s(\mathbf{k}(t),t)=e^{-t/T_2^*}\int_{\substack{\text{all}\\ \text{space}}}M_0(\mathbf{r})e^{-2\pi i\mathbf{k}(t)\cdot\mathbf{r}}d\mathbf{r}\,.$$

We assumed T_2* does not vary spatially, which is not true in the brain. However, if it does we can just break up our integral into sub-regions in which T_2* is constant.

When viewed in **k**-space, the T_2* decay happens along the readout direction (that would be the x-direction in the pulse sequence shown at the beginning of these lecture notes). Thus we can think of our signal in **k**-space as being the result of multiplying the non-decaying signal, s(**k**), by the additional decay accrued due to the T_2* decay:

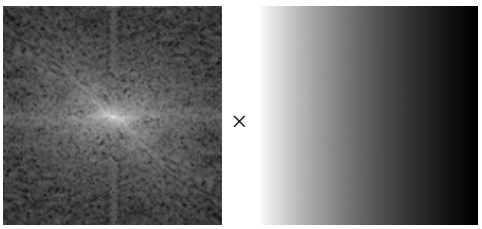

To a first approximation, the value of the decay function $\xi(k_r)$ at the echo time TE when our signal is maximal is the dominant quantity, and we can replace $\xi(k_r)$ merely by its value at the echo time, e^{-TE/T_2^*}, a constant. This implies that the reconstructed signal in each voxel will be remain, as before, the convolution of the spin density $M_0(\mathbf{r})$ with the point spread function PSF($\mathbf{r}$-$\mathbf{r}_{ijk}$), but now multiplied by e^{-TE/T_2^*}.

The Gradient-Echo Sequence Produces T_2^* Weighting (The Complete Explanation)

We now revisit our previous explanation with a more complete mathematical explanation for those readers interested in the details.

Let D be the duration between the center of the excitation pulse and the beginning of acquisition. Some constant signal decay, common to all excitations, is accumulated between excitation and the beginning of the ADC. Since this duration, *D*, is fixed between different scans this is a constant factor of the form $\exp\left[-D/T_2^*\right]$. The acquired k-space signal then decays as e^{-t/T_2^*}, where t=0 corresponds to the beginning of the ADC. Therefore the acquired k-space signal will be the k-space data, multiplied by a decaying envelope:

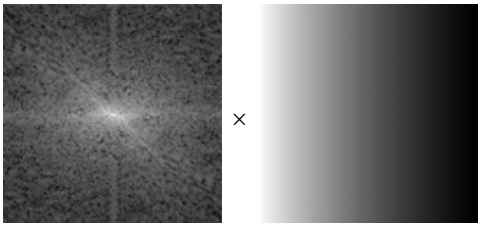

The actual decay only happens along k_r, the readout direction. Hence the "decay function" $\xi(k_r)$ is only a function of k_r. The actual extent of decay will depend on the ratio between T_{ADC}, the total acquisition time, and T_2^* (t=0 corresponds to the **center** of the ADC):

$$\xi(k_r) = \underbrace{e^{-TE/T_2^*}}_{\text{decay until ADC}} \cdot \underbrace{e^{-\frac{t(k_r)}{T_2^*}}}_{\text{decay during ADC}}.$$

Because we set t=0 at the center of the ADC (at the echo), we have

$$k(t) = \frac{t}{T_{ADC}} k_{\max}$$

so k(0)=0, k(-T_{ADC}/2) = -k_{max}/2 and k(T_{ADC}/2)=k_{max}/2. This is a slightly unusual choice of t=0 but it makes subsequent calculations easier. Inverting,

$$t(k) = \frac{T_{ADC}}{k_{\max}} k.$$

Plugging this back into our decay equation, we get:

$$\xi(k_r) = \underbrace{e^{-\frac{TE}{T_2^*}}}_{\text{decay until ADC}} \cdot \underbrace{e^{-\frac{T_{ADC}}{T_2^* k_{\max}} k_r}}_{\text{decay during ADC}}$$

with $k_{max}=\gamma G T_{ADC}=1/\Delta x$ (where Δx is the nominal resolution).

A well known theorem from Fourier analysis states that the Fourier of the multiplication of two functions – such as s(k) and[1] $\xi(k_r)$ – equals the convolution of their Fourier transforms:

$$ICFT\left(a(k)b(k)\right) = \frac{1}{2\pi} ICFT\left(a(k)\right) \otimes ICFT\left(b(k)\right)$$

In our case:

$$ICFT\left(s(k)\xi(k)\right) = ICFT\left(s(k)\right) \otimes ICFT\left(\xi(k)\right) = \frac{1}{2\pi} I(x) \otimes ICFT\left(\xi(k)\right)$$

[1] $\xi(k)$ includes in it already the windowing of the data – namely, that we acquire only from –k_{max}/2 to k_{max}/2.

where the ICFT of s(k), the k-space data, is I(x), the image, and

$$ICFT\left(\xi(k)\right)=\frac{1}{2\pi}\int_{-\infty}^{\infty}\xi(k)e^{2\pi ikx}dk$$

Although it is possible to calculate the ICFT of the decay envelope analytically (try it!), it doesn't yield very good intuition. Instead, let us examine two extremes in the next two sections.

For Short Acquisition Times (T_{ADC}<<T2*), The Effect Of T_2* Is Simply A Constant Factor

In most realistic cases, the acquisition duration T_{ADC} is much shorter than T_2*, so we can neglect it to a first approximation. This means that $\xi(k_r)$ is just a constant function between $-k_{max}/2$ to $k_{max}/2$:

$$\xi(k_r)\approx \underbrace{e^{-\frac{TE}{T_2^*}}}_{\text{decay until ADC}}$$

Then:

$$ICFT\left(\xi(k)\right)=e^{-\frac{TE}{T_2^*}}\operatorname{sinc}\left(\frac{\pi x}{\Delta x}\right)\qquad T_{ADC}<<T_2^*$$

Here we see the "regular" sinc function with a width of Δx, our digital resolution. This is multiplied by e^{-TE/T_2^*}, whatever decay was introduced until the acquisition block, implying the final image is **T_2*-weighted** by a factor e^{-TE/T_2^*}. If T_2* is spatially dependent, different regions in the image will be weighted by their own e^{-TE/T_2^*} factor.

T_2* Signal Decay Places An Ultimate Lower Bound On Our Spatial Resolution Along The Readout Direction

We've discussed the PSF of cartesian sampling and remarked its width, on the order of $\Delta x=1/k_{max}=1/\gamma G T_{ADC}$, broadens our signal and limits our resolution. It would seem that in theory we could get infinite resolution if we had enough signal and acquired for long enough ($T_{ADC}\to\infty$). However, T_2* prevents that from happening, at least along the readout direction. If $T_{ADC}>>T_2$* we can assume the exponent in ξ(k) has decayed to 0 by the time we reach $k_{max}/2$, and its Fourier transform then becomes fairly straightforward since we can extend the upper limit to ∞:

$$\begin{aligned}&ICFT\left(\xi(k)\right)\\&=\frac{1}{2\pi}e^{-\frac{TE}{T_2^*}}\int_{-k_{max}/2}^{k_{max}/2}e^{-\frac{T_{ADC}}{T_2^*k_{max}}k}e^{2\pi ikx}dk\\&\approx\frac{1}{2\pi}e^{-\frac{TE}{T_2^*}}\int_{-k_{max}/2}^{\infty}e^{-\frac{T_{ADC}}{T_2^*k_{max}}k}e^{2\pi ikx}dk\\&=\frac{1}{2\pi}e^{-\frac{\left(TE-\frac{T_{ADC}}{2}\right)}{T_2^*}}\frac{T_2^*k_{max}e^{-\frac{i\pi x}{\Delta x}}}{T_{ADC}}\frac{1+2\pi i\frac{T_2^*}{T_{ADC}}\frac{x}{\Delta x}}{1+\left(2\pi\frac{T_2^*}{T_{ADC}}\frac{x}{\Delta x}\right)^2}\end{aligned}$$

This is a fairly complicated looking function, but it really breaks down into some uninteresting constant factors, times

$$e^{-\frac{\left(TE-\frac{T_{ADC}}{2}\right)}{T_2^*}}$$

which is the T_2* weighting (note that now T_{ADC} appears because it is non-negligible), and

$$\frac{1+2\pi i\frac{T_2^*}{T_{ADC}}\frac{x}{\Delta x}}{1+\left(2\pi\frac{T_2^*}{T_{ADC}}\frac{x}{\Delta x}\right)^2}.$$

i.e., our PSF will now no longer be a sinc but given by the above function. Its real part is known as a **Lorentzian**:

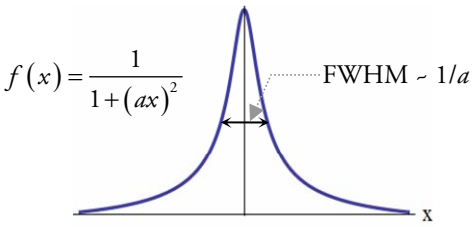

In our case, it will have a width approximately given by $FWHM\sim T_{ADC}\Delta x/T_2^*$. That is, the width of our PSF will become larger and larger the shorter T_2* or the longer T_{ADC} become, effectively blurring our image even if the nominal resolution stays the same.

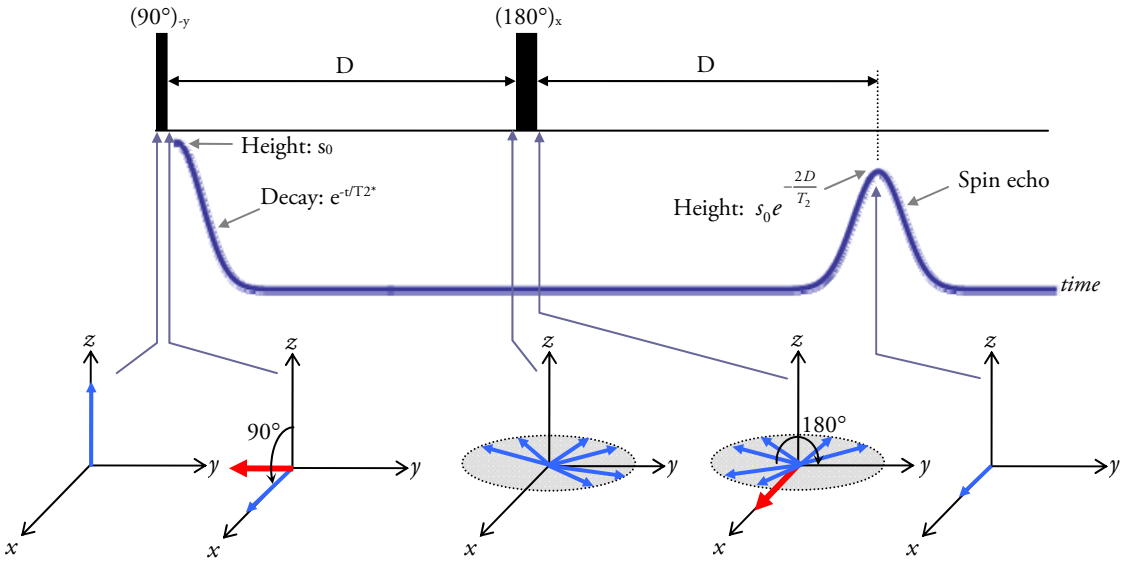

t=0⁻
All the spins start out from thermal equilibrium.

t=0⁺
Following the initial $(90°)_{-y}$ hard pulse, all spins are tilted to the transverse plane.

t=D⁻
The spins dephase due to different spatial offsets (T_2') and microscopic decay (T_2).

t=D⁺
A $(180°)_x$ flips the spins about the y axis and reverses their phases, $\phi \rightarrow -\phi$.

t=2D
All of the spins refocus along the *y*-axis and the signal "revives", forming a **spin echo**. Some signal is lost due to intrinsic, irreversible T_2 decay.

T_2 Contrast

The Spin Echo Refocuses T_2' And Leaves Us Solely With T_2 Related Signal Decay

We now come to the next class of sequences known as **spin echo sequences**, which are T_2-weighted. These sequences rely on a basic element known as a spin echo. Imagine the following simplified pulse sequence, in which the spins are excited, precess around, and then subjected to a 180° pulse – also called a **π-pulse** in MR jargon for obvious reasons. If we run this sequence and plot the signal as a function of time, we will obtain

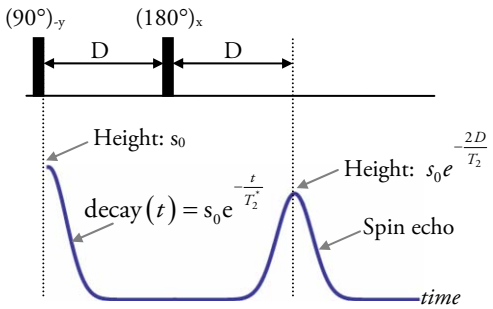

We will now analyze what happens to the signal as a function of time. The π-pulse inverts the phase of the spins in the xy-plane. Focusing on a small, mesoscopic subregion, the offset of those spins due to spatial, constant field inhomogeneities will be $\omega(\mathbf{r},t) = \gamma \Delta B(\mathbf{r})$. Its time evolution will be

$$M_{xy}(\mathbf{r},t) = M_{xy}(\mathbf{r},0)e^{-\int_0^t \omega(\mathbf{r},t')dt'}.$$

At time D, right before the 180° pulse, it will have accumulated a phase

$$\phi(D) = \gamma \Delta B(\mathbf{r}) D.$$

What is the effect of the $(180°)_x$-pulse on it? This can be deduced via a simple calculation: if we write the magnetization vector as

$$\mathbf{M} = \begin{pmatrix} \cos(\phi) \\ \sin(\phi) \\ 0 \end{pmatrix}$$

and use the form of a rotation matrix about x:

$$R_x(\theta)=\begin{pmatrix}1 & 0 & 0\\ 0 & \cos(\theta) & \sin(\theta)\\ 0 & -\sin(\theta) & \cos(\theta)\end{pmatrix}$$

we get

$$R_x(\pi)\mathbf{M}=\begin{pmatrix}1 & 0 & 0\\ 0 & -1 & 0\\ 0 & 0 & -1\end{pmatrix}\begin{pmatrix}\cos(\phi)\\ \sin(\phi)\\ 0\end{pmatrix}=\begin{pmatrix}\cos(-\phi)\\ \sin(-\phi)\\ 0\end{pmatrix}$$

namely, the phase of the spin will be flipped $\phi\to-\phi$. Since in the subsequent delay D it will acquire again the same phase $\phi(D)=\gamma\Delta B(\mathbf{r})D$, these two will cancel out and the spin will return to the x-axis, **regardless of its constant offset**. This experiment is one of the most important experiments in the history of NMR, first carried out by Erwin Hahn in 1950 (in a slightly different variation). Even though the π-pulse negates T_2' decay due to constant inhomogeneities, it does not negate the microscopic T_2 decay due to the microscopic fluctuations of the fields. We thus say that the π-pulse has refocused T_2' but not T_2.

π-Pulses Allow Us To Replace T_2^* Contrast With T_2 Contrast

180° pulses can be used to turn the basic T_2^* based gradient echo (GRE) sequence we've encountered previously into a T_2 based spin echo (SE) sequence as follows:

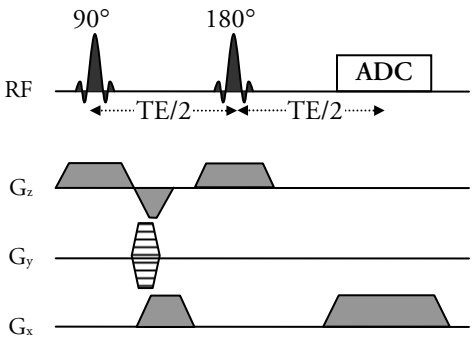

Note several changes compared to the GRE sequence. First, the rewinding gradient lobe along x is now positive. This is because the effect of the 180° is to invert the phase of the spins. This means that spins with a particular **k**-space position,

$$M_{xy}^{(rot)}(\mathbf{r},t)=M_0(\mathbf{r})e^{-2\pi i\mathbf{k}(t)\cdot\mathbf{r}}$$

will get their phase flipped by the 180° pulse:

$$M_{xy}^{(rot)}(\mathbf{r},t)=M_0(\mathbf{r})e^{-2\pi i\mathbf{k}(t)\cdot\mathbf{r}}\xrightarrow{180^\circ}M_0(\mathbf{r})e^{-2\pi i\mathbf{k}(t)\cdot\mathbf{r}}$$

which is equivalent to replacing **k** by **-k**. Thus, the corresponding trajectory in **k**-space for a given scan looks like this:

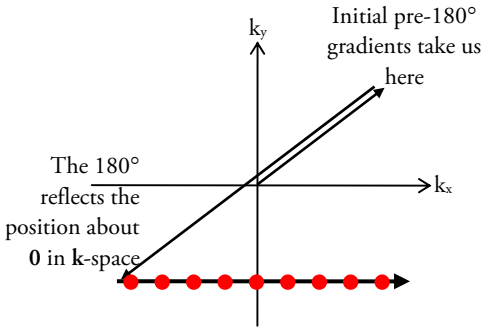

Our simplistic analysis can now be repeated: our signal decays and to a first approximation we need only take its value at the echo time, TE, into account. The decay is now governed by T_2 and not T_2^* thanks to the 180° pulse; This means **our signal in each voxel will be multiplied by** e^{-TE/T_2}.

T_1 Contrast

Waiting For The Magnetization To Return To Thermal Equilibrium Takes Too Much Time

Let's do a simple calculation. Suppose we want to acquire an image with 256×256×192 points, not uncommon in MRI. The T_1 of water in GM and WM in the brain is about 1 second, so after each excitation we would need to water about $5\cdot T_1$ for the magnetization to return to thermal equilibrium along the z-axis, meaning

$$\text{TR} \equiv \text{Time to repeat}$$
$$\approx 5\cdot T_1 \approx 5 \text{ sec}$$

One of the directions would use a readout gradient, so we don't have to "pay" for it with a scan, meaning we acquire all of the points in one shot after the excitation by applying a readout gradient. This still leaves us with 256×192 scans (256 phase encoding steps and 192 slices), leading to a total scan time of

$$256\times192\times TR = 68 \text{ Hours}.$$

Obviously not a very practical protocol. This can be shortened considerably by either reducing the resolution or reducing TR. We will explore what happens during the latter. As we'll see, this will not only shorten the measurement but will also introduce T_1-weighting which is in many times desirable.

Pulsing At A Rate TR~T_1 And Below Introduces T_1 Weighting

Imagine a static bucket with water. The water is said to be in *static equilibrium*, because nothing's happening to it. Next, imaging (i.) poking a hole in the bottom of the bucket, so water start running out, and (ii.) opening a tap just above the bucket, letting water flow in at a constant rate. What will happen? The water may rise or fall, but will eventually reach a new state of equilibrium (remember, the more water there is, the faster it drips out due to pressure; and the less water there is, the slower, so eventually the water flowing in will equilibrate with the water flowing out, even if at first the rates are different). This new equilibrium is termed **dynamic equilibrium**. Something is continuously happening to the system (in and out flow of water), so it's not static anymore, and yet its state doesn't change.

A similar thing happens in MRI when we use rapid pulsing. Consider a train of pulses of flip angle α (that is, $\gamma B_{RF} t = \alpha$) around, say, the $-y$ axis. Let's call the time between pulses TR (for "Time per Repetition"):

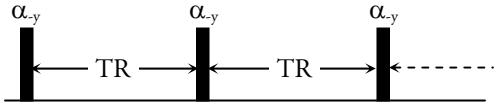

The spins get acted upon by two "forces": the pulses, which repetitively try to take them out of equilibrium, and relaxation, which tries to get them back to equilibrium. There's also precession going on. It can be shown (I won't do it here, but it's not that difficult really) that the spins eventually settle into dynamic equilibrium, regardless of their initial state; that is, after enough pulses have been given, the state of the spins after each pulse is identical. In other words, the magnetization vector at points A, B, C, ... below is the same:

This state will depend on the variables of the system: TR, T_1, T_2 and α, and also the offset of the spins, $\Delta\omega$.

We will assume that TR>>T_2 for the time being. This serves to ensure the transverse magnetization decays to 0 before the next pulse, so we can assume $M_{xy}=0$ just before any of the pulses. In MRI jargon we say that the transverse magnetization is **spoiled** before each excitation.

Denote by $\mathbf{M}^{(A)},\mathbf{M}^{(B)},\mathbf{M}^{(C)}$ the magnetization vectors at A, B and C, respectively. We're interested in computing $\mathbf{M}^{(B)}$ and, subsequently, the magnetization's evolution between the pulses. Note that:

1. By assumption of dynamic equilibrium, $\mathbf{M}^{(A)} = \mathbf{M}^{(C)}$.
2. Since we're assuming the magnetization is spoiled, $M_{xy}^{(A)} = M_{xy}^{(C)} = 0$.

So:

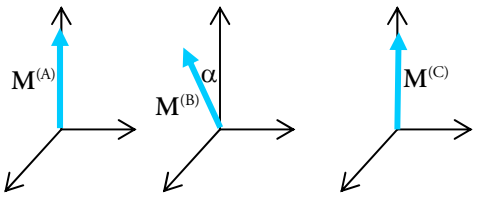

Since $\mathbf{M}^{(B)}$ is a tipped version of $\mathbf{M}^{(A)}$ by an angle α, we have:

$$M_z^{(B)} = M_z^{(A)}\cos(\alpha)$$
$$M_x^{(B)} = M_z^{(A)}\sin(\alpha)$$
$$M_y^{(B)} = 0$$

So $M_{xy}^{(B)} = M_z^{(A)} \sin(\alpha)$. However, we're not interested in the transverse magnetization (yet). We know that the longitudinal component of **M** relaxes back to equilibrium with a time-constant T_1. You've shown in exercise 3 that, following the pulse and after a time TR, M_z will equal

$$M_z^{(C)} = M_z^{(B)} e^{-TR/T_1} + \left(1 - e^{-TR/T_1}\right) M_0$$

where M_0 is the thermal equilibrium value of the magnetization (it's in general **not** equal to $M_z^{(A)}$!). Since $M_z^{(C)} = M_z^{(A)}$ and $M_z^{(B)} = M_z^{(A)} \cos(\alpha)$, we can plug these into the above equation,

$$M_z^{(A)} = M_z^{(A)} \cos(\alpha) e^{-TR/T_1} + \left(1 - e^{-TR/T_1}\right) M_0$$

and solve for $M_z^{(A)}$:

$$M_z^{(A)} = \frac{1 - e^{-TR/T_1}}{1 - \cos(\alpha) e^{-TR/T_1}} M_0 .$$

From this we can compute $M_{xy}^{(B)} = M_z^{(A)} \sin(\alpha)$, and, in fact, deduce $M_{xy}(t)$ for any time between the pulses (t=0 corresponds to the time right after a pulse):

$$\begin{aligned} M_{xy}(t) &= M_{xy}^{(B)} e^{-i\mathbf{k}(t)\cdot\mathbf{r}} e^{-t/T_2} \\ &= \frac{\left[1 - e^{-TR/T_1}\right] \sin(\alpha)}{1 - \cos(\alpha) e^{-TR/T_1}} M_0 e^{-i\mathbf{k}(t)\cdot\mathbf{r}} e^{-t/T_2} \end{aligned}$$

Pulsing At A Rate TR<<T_1 Saturates The Signal

Let us plot the intensity of the **dynamic equilibrium factor**,

$$DEF\left(\frac{TR}{T_1}, \alpha\right) = \frac{\left[1 - e^{-TR/T_1}\right] \sin(\alpha)}{1 - \cos(\alpha) e^{-TR/T_1}},$$

as a function of TR for T_1=1 sec and for several values of α:

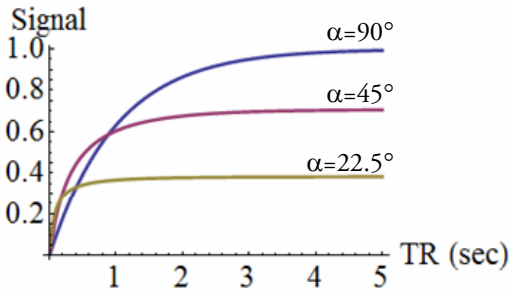

As we take TR to be shorter and shorter three things happen:

1. Our acquisition becomes shorter (good!).
2. We get more T_1 weighting, meaning the intensity becomes more sensitive to T_1, albeit only in a certain range. This is usually good.
3. The signal diminishes (bad).

The first is obvious. We will take a look at the second in a moment. The diminishing intensity is called **saturation**. When we excite the magnetization from thermal equilibrium we diminish the M_z by an amount sin(α). After we do so, M_z will start building up towards thermal equilibrium during TR with a time constant T_1. This means there are two "forces" acting on M_z: T_1 relaxation, which builds it up, and our pulsing, that reduces it. If TR<<T_1 the pulsing wins and M_z reduces to 0 (that is, we saturate the signal). If TR>>T_1, M_z has sufficient time to build towards thermal equilibrium and we start up from the full intensity before the next pulse.

It is not uncommon to see sequences with TR=5 ms. We can repeat our calculation of the sequence's duration and get:

$$256 \times 192 \times TR \approx 4 \text{ minutes} .$$

A significant reduction compared to our previous 68 hours, and also a reasonable scan time.

Getting T_1 Contrast With Rapid Pulsing

The other beneficial aspect of rapid pulsing, other than reducing TR, is introducing T_1 contrast. As we've previously remarked, the dynamic equilibrium factor depends on the ratio between TR and T_1. Let's plot it for a fixed flip angle - say, α=90° - as a function of T_1, for a fixed TR=10 ms:

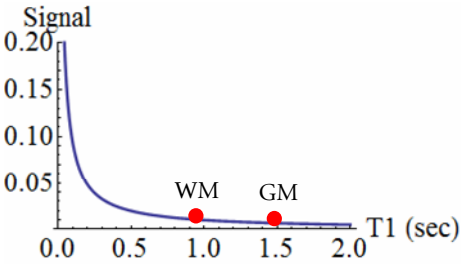

We see that, say, gray matter (T_1~1.5 s), white matter (T_1~1 s) and CSF (T_1~4 s) yield different signal intensities:

$$DEF_{WM} = DEF\left(\tfrac{0.01}{1.0}, \tfrac{\pi}{2}\right) \approx 0.01$$
$$DEF_{GM} = DEF\left(\tfrac{0.01}{1.5}, \tfrac{\pi}{2}\right) \approx 0.007$$
$$DEF_{CSF} = DEF\left(\tfrac{0.01}{4.0}, \tfrac{\pi}{2}\right) \approx 0.002$$

The signal intensity of WM would be higher since it is less saturated. It would thus appear "white", while GM would appear "grey", on a T_1 weighted image – this is a complete coincidence and should not be taken as a rule. On T_2-weighted images, for example, WM has a lower intensity than GM and their visual roles are switched!

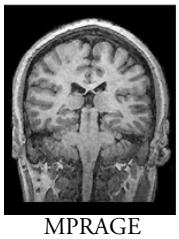
MPRAGE
(T_1-weighted)

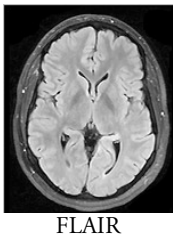
FLAIR
(T_2-weighted)

What happens if we increase TR? Let's set TR=1.0:

$$DEF_{WM} = DEF\left(\tfrac{1.0}{1.0}, \tfrac{\pi}{2}\right) \approx 0.63$$
$$DEF_{GM} = DEF\left(\tfrac{1.0}{1.5}, \tfrac{\pi}{2}\right) \approx 0.47$$
$$DEF_{CSF} = DEF\left(\tfrac{1.0}{4.0}, \tfrac{\pi}{2}\right) \approx 0.22$$

Lowering The Flip Angle Reduces Maximal SNR

The dynamic equilibrium factor really depends only on two parameters: the flip angle and the ratio TR/T_1:

$$DEF\left(\tfrac{TR}{T_1}, \alpha\right) = \frac{\left[1 - e^{-TR/T_1}\right]\sin(\alpha)}{1 - \cos(\alpha)e^{-TR/T_1}}.$$

Let us plot it as a function of the ratio, for several different flip angles:

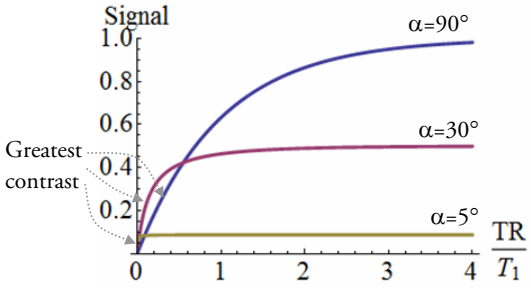

The **maximal signal** is obtained when the signal becomes independent of T_1. Naively, it looks like this might happen when $TR>>T_1$, which is true, since then $e^{-TR/T_1} \approx 0$ and

$$DEF_{\max} = \sin(\alpha) \qquad \begin{pmatrix}\text{for fixed}\\ \text{flip angle}\end{pmatrix}.$$

However, the above plots show this happens even before that for small α! How so? Let's take the degenerate case α≈0, at which cos(α)≈1 and:

$$DEF\left(\tfrac{TR}{T_1}, \alpha\right) \approx \frac{\left[1 - e^{-TR/T_1}\right]\sin(\alpha)}{1 - e^{-TR/T_1}} = \sin(\alpha).$$

So, the smaller α, the more similar the nominator ($1-e^{-TR/T1}$) and denominator ($1-\cos(\alpha)e^{-TR/T1}$) become and consequently the less dependent on TR/T_1. The true criterion is:

$$\frac{TR}{T_1} >> -\log\left(\cos(\alpha)\right) \equiv \left(\tfrac{TR}{T_1}\right)_0$$

Higher Relative CNR Is Obtained At Lower Flip Angles. Higher Absolute CNR Is Obtained At Higher Flip Angles.

The following useful table evalutes some parameters relations and their respective signal intensities:

α	SNR_{max}	$\left(\frac{TR}{T_1}\right)_0$	TR* (ms)	Example: $CNR_{WM,GM}$
1°	0.017	0.0001	0.1	0.001
5°	0.09	0.0038	3.8	0.008
10°	0.17	0.0153	15.3	0.017
30°	0.5	0.1438	144	0.053
90°	1.0	∞	1000	0.145
90°				0.0012 (TR=3.8 ms)

TR* is the value of TR given T_1=1 sec (typical for in-vivo brain GM, WM tissue at 3T). For α=1°, TR*=0.1 ms is too fast for almost all current imaging hardware.

The example is given by setting TR=TR* and calculating the CNR between WM and GM, assuming a standard deviation of noise of unity:

$$CNR_{WM,GM} = \frac{DEF\left(\frac{TR^*}{T_{1,GM}},\alpha\right) - DEF\left(\frac{TR^*}{T_{1,WM}},\alpha\right)}{\sigma_{SNR}}.$$

The above table illustrates two simple points:

1. Lower flip angles result in less SNR.
2. Lower flip angles "force" shorter TRs to get any sort of contrast (i.e. to be in the part of the signal curve which is sensitive to TR/T_1).
3. The resulting "forced" short TRs also lead to quicker sequences.

However, the last line in the table shows that, **for the same TR**, GM/WM contrast is actually enhanced at lower flip angles (compare TR=3.8 ms for α=90° and α=5°). Thus:

> For a **given** TR, the **relative** CNR is (usually) maximized by taking lower flip angles, at the cost of less SNR. The **absolute** CNR is maximized by taking larger flip angles and longer TRs, at the cost of longer sequence duration.

This is also shown in the next page for a "model" of the brain with WM (T_1=1 s), GM (T_1=1.5 s) and CSF (T_1=4 s) for different TRs (50, 400, 1000 ms) and flip angles (10°, 30°, 90°):

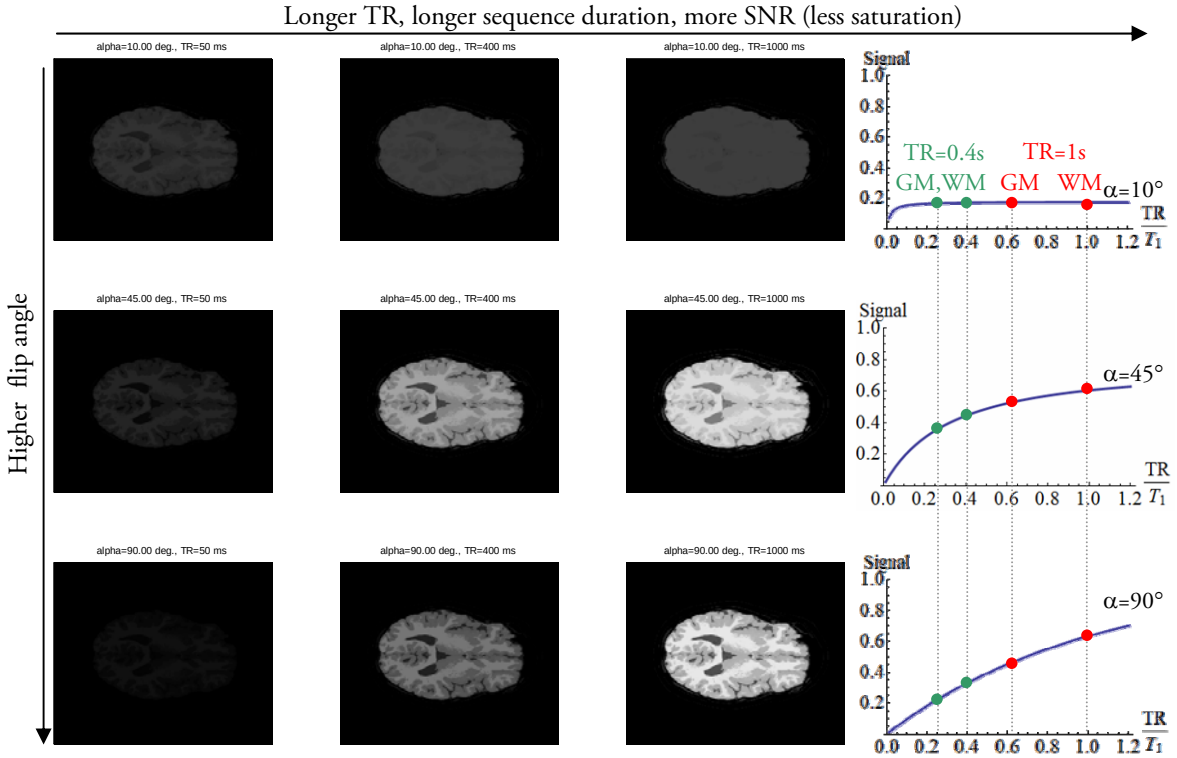

Given TR and T_1, There Is A Flip Angle That Yields Maximal SNR: The Ernst Angle

We're remarked that reducing α reduces the *maximal* SNR achievable, but this does not mean that reducing α always reduces the SNR. This is clearly seen in the above pictures, where, for TR=50 ms, it is the **smallest** α that has the highest SNR. Furthermore, for TR=400 ms, α=45° has more SNR than either α=10° or α=90°. We understand this as follows: While reducing α reduces the amount of magnetization excited, it also leaves more longitudinal magnetization intact before the next excitation pulse. There are thus two competing "effects" at play, implying that there should be some optimal flip angle (for a given TR/T_1 ratio) at which SNR is maximal. This flip angle is called the **Ernst Angle**. To find it, simply differentiate the dynamic equilibrium factor with respect to α and equate to 0 to find its maximum:

$$\frac{d\left\{DEF\left(\frac{TR}{T_1},\alpha\right)\right\}}{d\alpha}=0\,.$$

After some algebra, one finds

$$\alpha_E=\arccos\left(e^{-TR/T_1}\right).$$

For example, for TR=50 ms, T_1=1 sec (WM), we have

$$\alpha_E=18°\,.$$

This indeed shows that we'd expect larger signals for the α=10° case than 45° or 90°. For TR=400 ms, T_1=1 sec again,

$$\alpha_E=48°\,,$$

again, agreeing with our simulation results showing that the α=45° has the largest intensity out of all cases.

Two words of caution, though: (1) the Ernst angle is T_1 specific and therefore tissue specific; (2)

the Ernsr angle maximizes the SNR **but not necessarily the CNR**! Indeed, for the TR=400 ms case, it seems that α=90° yields greater CNR between WM and GM than α=45°. We can verify this numerically, using a noise with a standard deviation of 1:

Flip Angle	Tissue	SNR
α=45°	WM	0.443
	GM	0.361
	$CNR_{WM,GM}$ = 0.082	
α=90°	WM	0.330
	GM	0.234
	$CNR_{WM,GM}$ = 0.096	

Qualitative Analysis of Non-Spoiled Sequences

A quantitative, or even qualitative analysis of non-spoiled sequences, is beyond the scope of these introductory lecture notes. We do however mention a few points of interest regarding these sequences below.

Spoiling Fails When TR~T_2 Or Shorter

All of our discussion up to this point assumed that our sequence is spoiled; namely, that M_{xy}=0 prior to each excitation pulse. This occurs naturally when TR>>T_2 so the transverse magnetization decays to 0 due to T_2 relaxation before the next pulse in the train is applied.

Any unspoiled transverse magnetization prior to the next α excitation pulse will be partially **stored** - that is, converted to longitudinal magnetization - and partially remain in the xy plane. Any longitudinal magnetization will be partially excited and partially remain along the z-axis. The effect of each pulse can be described as:

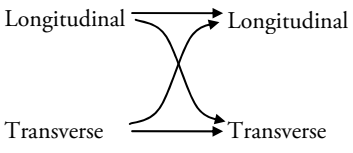

This complicates the analysis and understanding of the sequence considerably; even more so when the flip angle is not 90°. Just to provide a visual picture, here is what happens to an ensemble of spins when the inter-pulse spacing is 50 ms, with α=60°, T_1=500 ms, T_2=100 ms:

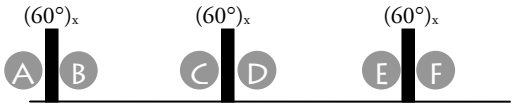

Below I've simulated what happens to an ensemble of spins with a distribution of offsets (different hues of blue correspond to different offsets):

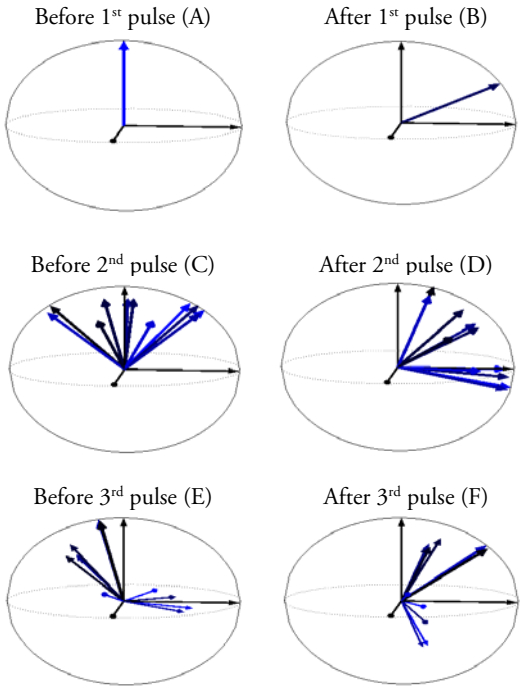

At point (C), just before the second pulse, we still have transverse magnetization – this is the meaning of a **non-spoiled sequence**.

It is quite clear that any attempt to visualize the spins' individual trajectories is impossible after 2-3 pulses. Some spins spend time in the xy-plane and are affected by T_2, while others spend time along z and are affected by T_1, and each spin rotates by a different amount.

What happens if we keep on giving pulses? Below I've simulated the state of the spins just before the 12th (left) and 13th (right) pulses:

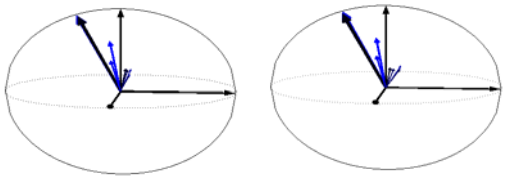

It is clear that the two states are very similar. Even though our sequences are non-spoiled, they still converge to a state of **dynamic equilibrium**.

Spoiling Can Still Be Achieved Even At TR~T_2 Via Either Spoiler Gradients Or RF Pulse Phases

Before taking a look at non-spoiled sequences we note that we can create "effective spoiling" via two mechanisms at our disposal: by introducing **spoiler gradients** at the end of each TR, or by incrementing the phases of the RF pulses between successive excitations. The first approach is called **gradient spoiling** and the second approach is called **RF spoiling**.

The basic idea you should grasp is the following:

> If the transverse components of spins inside a voxel are randomly and/or evenly distributed in the xy plane then their transverse vector sum will add up to zero and we will have effective spoiling. In MRI jargon we say that we're **dephasing** the **intra-voxel magnetization**.

The question now becomes: how can we make sure the spins inside a voxel point in different directions in the xy-plane?

One answer to that would be: apply a strong gradient! A strong (constant, for simplicity) gradient applied for a time t will create a phase of the form:

$$\phi(\mathbf{r}) = \gamma(\mathbf{G}\cdot\mathbf{r})t\ .$$

The signal from the voxel (assuming an ideal boxcar PSF) will be:

$$s = \int_{\text{voxel}} M_{xy}(\mathbf{r}) e^{-i\gamma(\mathbf{G}\cdot\mathbf{r})t} d\mathbf{r}\ .$$

Let's suppose we have just a z-gradient, so

$$\phi(z) = \gamma G z t\ .$$

We've already plotted the shape of the spins as a function of z:

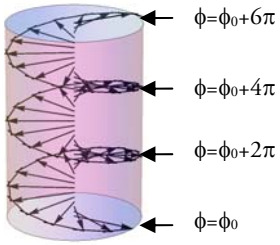

Each "winding" in this helix corresponds to a phase difference of 2π, and means the spins belonging to that helix are evenly distributed in the xy-plane. Adding up the spins will give zero signal if the spins are all equal magnitude. This may or may not be a good approximation, which is why we want as many windings as we can possibly get to null the signal:

$$\Delta\phi_{voxel} >> 2\pi$$

or

$$\gamma G \Delta z_{voxel} t >> 2\pi\ .$$

This would ensure that

$$s = \int_{\text{voxel}} M_{xy}(\mathbf{r}) e^{-i\gamma(\mathbf{G}\cdot\mathbf{r})t} d\mathbf{r} \to 0\ .$$

There is another way of viewing this: by applying a constant gradient we're getting farther and farther away from the center of k-space, and we've already remarked that the farther out we go in k-space the smaller our signal (in the following diagram we assume the strong gradients are applied along the x- and y-axes, not the z-axis):

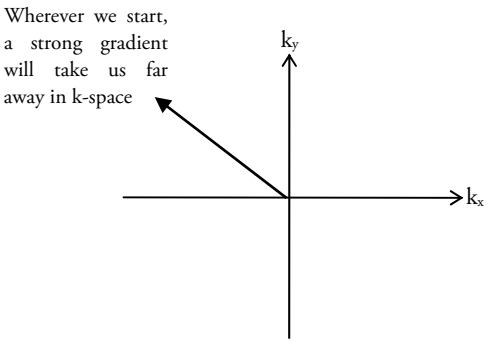

Spoiler gradients are an effective and simple way of ensuring effective spoiling, but they can be tricky to design since they might refocus unwanted signals (see homework assignment).

RF spoiling can also be used to dephase the intra-voxel transverse magnetization. The ideas here are significantly more complex than gradient spoiling, so we will not go into the details[2], but the philosophy is very straightforward: can we cause the newly excited magnetization after each pulse to cancel out the magnetization from all of the previous excitation? The answer is yes. One way to do this is increment the phase of the RF pulses by 117° between excitations:

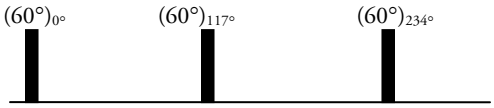

[2] Readers interested in the details are referred to Zur et. al., Magn. Reson. Med. 21(2):251-263 (1991)