The *Animated Physics of* MRI:

The Notes

DB Plewes, PhD
Department of Medical Biophysics and Medical Imaging
University of Toronto

Introduction and using the CD.

The enclosed is an identical copy of the lecture given at the Glasgow, 2001 meeting of the ISMRM. It is run from a MacroMedia Director File which allows convenient migration through the lecture. It needs no additional drivers to run the package. To launch the CD, simply place it in the CD-Rom driver and wait. An auto-start code has been added which starts the CD automatically. It takes about 45 seconds to auto-launch. However, if you want to copy this to your hard drive, you should preserve the folder structures (Director, Imaging and Basics) on the CD. The executable file is called "MRIPhysics.exe" and is located in the "Director" folder. Just click on this to start the lecture.

Once the Director file starts, a Start button is shown at the bottom of the image. Click on this to proceed to the next figure. For each figure, a Next or Previous button can be used to move throughout the lecture. In addition, there are Play, Pause and Rewind buttons (similar to a VCR) to move through the animation. In some cases, the letter "P" is shown beside the Play button which indicates that the graphic is started in a`Paused" mode. To get it going, simply click on the Play button. I have inserted these pauses throughout the lecture to ensure that adequate time occurs to describe each graphic before the animation occurs. In other cases, I have just placed some delays. I have found that the set of delays and pauses works well for the talk with the text below. In addition, a slider has been added which indicates where the current frame resides in the animation. It is useful to time the text with the actions of the animation. Also, it is possible to click and drag the slider tab to move through the animation more quickly or in reverse. Sometimes, I have noticed that if you proceed too quickly through the illustrations, it may require two clicks to move from an animation to the next slide. In this case, you may find that pausing animations before proceeding to the next figure is more reliable.

In this document, I have provided the text which I used to describe each "figure". The figure numbers listed below corresponds to the figure number located on the upper left hand corner of each figure on the CD. You can use text this verbatim or modify your narrative to suit your lecture style or duration. My narration is placed in normal print while any notes about how to run the animation are typed in *italics.* I have found that this lecture requires about 90 minites to give at a reasonable pace. Any comments or suggestions about your experience with this CD would be much appreciated and can be sent to don.plewes@home.com.

I hope you find it useful.

Don Plewes, PhD - Feb, 2003

The Animated Physics of MR1

Lecture narrative for the description of each

figure: Figure 1

Title the presentation. You can skip as may not be

relevant to your presentation.

This is an overview of the lecture as a whole and is composed of two parts. The first will include the fundamentals of proton N1VIR physics, while the second will undertake to provide an intuitive and visual framework to understand the mechanisms of MR image formation.

Figure 3

The first portion of the lecture covers basic NMR physics in six separate sections. These will include:

- A review of the mechanisms for the formation of magnetization,
- The notion of resonance,
- How to detect and excite NMR,
- A description of the rotating frame to understand spin gymnastics -

A discussion of spin dephasing which leads to a,

- A review of the important relaxation mechanisms, T 1 and T2.

Figure 4.

The most important site of the NMR resonance relevant to MRI is the nucleus of the hydrogen atom in water. While other protons occur within biological molecules, water represents the main site for MRI imaging due to the concentration of protons in water and the dynamical properties of water.

The proton is a nuclear particle which exhibits charge, mass and spin. While the first two concepts is familiar, the notion of spin is not as well appreciated. As the name suggests, it can be thought of as a rotation of the nucleus about its axis and in conjunction with the charge of the nucleus, gives the proton a magnetic property similar to a small bar magnet. If this were placed in applied magnetic field, it would, in the absence of any other factors, return slowly to align with the applied magnetic field. However, in addition to the magnetic property of the nucleus, the spin of the proton together with its mass, gives it a property referred to as angular momentum. The notion of angular momentum is familiar and is similar to the effect of a spinning top. In the presence of applied gravitational field, the spinning top will wobble or "precess" at a specific frequency as it moves in the gravitation field. The frequency of this precession is dictated by the details of the spinning top and the strength of the applied gravitational field. In a similar manner,

this phenomenon is also seen NMR arising from the combined effects of proton spin, charge and mass. Specifically, when a proton is placed in an applied magnetic field, it will attempt to realign to the magnetic field, but its angular momentum forces it to precess in a manner similar to that of the spinning gyroscope. The precessional frequency of the proton is known as the "Larmor frequency" and is proportional to the strength of magnetic field and a constant referred to as the "gyromagnetic" ratio. Each nucleus has a different gyromagnetic ratio and reflects the details of the nuclear structure. For example, at a field strength of 1 Tesla (approximately 30,000 times stronger than the earth's magnetic field), the Larmor frequency for protons is 42.57 Mhz. Doubling the magnetic field strength to 2 Tesla would increase the Larmor frequency to 85.14 Mhz.

Figure 5

The gyromagnetic ratio for a number of biologically relevant elements, along with their relative NMR signal is tabulated in this figure. It is noteworthy, that not all nuclei can generate an NMR signal. Only isotopes with an odd number of protons or neutrons have a non zero spin which permits the formation of an NMR signal. This shows that the nucleus of hydrogen, gives the biggest signal largely due to its gyromagnetic ratio and the fact that the most abundant isotope of hydrogen exhibits a spin. In comparison, the isotopes of carbon, sodium or phosphorous with non-zero spin are less abundant and therefore generate a much weaker NMR signal.

Figure 6

When a large number of protons are placed the magnetic field, their collective magnetic effects add to form a single magnetization referred to as the "bulk magnetization". This is a vector which precesses about the magnetic field Bo at the Larmor frequency. As long as the individual nuclei precess in perfect synchrony, the magnetization vector will maintain its magnitude.

Figure 7

In order to detect this magnetization, we use a coil of wire which is connected to a sensitive amplifier which is in turn tuned to the Larmor frequency. The changing flux of the magnetic field through the coil which occurs because of its rotation will induce a tiny oscillating NMR signal in the coil. The frequency of the oscillation is at the Larmor frequency. Only the time varying part of the magnetization is capable of inducing a signal in the coil and as such only the rotating component of the magnetization in a plane orthogonal to the applied magnetic field is detectable by the coil. This plane is given special name and is referred to as the "transverse" plane and as such, this component of the magnetization is referred to as the "transverse" component as opposed to the "longitudinal" component parallel to the Bo field. This also means that the orientation of the receiver coil must be such that its axis lies in the transverse plane in order to detect the changing magnetic field of the transverse component.

Figure 8

In the resting or the "equilibrium" state, the bulk magnetization is aligned parallel to the Bo field (along the z axis) and as such cannot precess. In order to precess and generate an NMR signal, we must tip the magnetization away from this equilibrium alignment so that a component of the magnetization lies in the transverse plane. To achieve this, the spins are exposed to an alternating "B 1" magnetic field which is tuned to the Larmor frequency. As the Larmor frequencies are typically in the Mhz range, these pulses are referred to as radio frequency or "RV pulses. Unlike the Bo field, the direction of the B 1 field is in the transverse plane and therefore is a right-angles to the Bo field. By virtue of this alternating applied magnetic field, the spins can progressively absorb energy and be tipped out of alignment from the longitudinal axis to create a precessing component in the transverse plane. The longer the duration of the B 1 field or the greater its field strength, the greater the tip angle which can be achieved. By careful choice of these two factors, the magnetization can be tipped to any angle relative to the Z axis. Once the RF pulse is turned off, the vector continues to precess about the Z axis and only changes as a result of relaxation processes to be discussed later in this lecture.

Figure 9

We can combine effects of excitation and detection to create very unusual pulse sequences which can exhibit unique features of the spin system. In this illustration, we show how the RF amplifier and coil can be used to excite the magnetization and then detect the resulting NMR signal. The upper box labeled "RF pulse" shows the waveform of pulses used to excite the spin system. Immediately after each pulse, the signal from the RF amplifier is connected to the box entitled "NMR signal" at which point we can observe the detected magnetization.

In this example, we will be using two RF pulses. During the first RF pulse, the spin is tipped from its vertical alignment into the transverse plane. After this, the wire from the RF amplifier is connected to the NMR signal graph where we observe a slowly decaying, oscillating NMR signal_ The signal decays until a second RF pulse is applied. As a result of this second RF pulse, we see that the magnetization begins to reappear to reach a peak value after which it decays again. We have not discussed why this occurs, but only use this example to illustrate how excitation and detection can work in concert to perform the NMR experiment. In fact, what we have illustrated here is an" spin-echo" pulse sequence which is normally used to measure spin-spin relaxation times.

Figure 10 (to start the video, press the play button)

(part one) - The motions of the magnetization vector are complex as they rotate out of alignment with the Z axis and precess in the transverse plane during excitation and relaxation. As seen from our normal point of view, ie the "laboratory frame", we see that the magnetization moves in a spiral trajectory toward the transverse plane. Once the RF pulse is turned-off, the magnetization continues to precess in the transverse plane and would continue to do so, in the absence of any relaxation effects.

(part 2)- In order to simplify our picture of these complex motions, consider that we place

a special frame composed of a turntable which itself rotates about the longitudinal axis at the Larmor frequency. Furthermore, imagine that this frame is carrying a small camera through which we will ultimately observe the rotating magnetization vector. As viewed from the laboratory frame, we see the magnetization, turntable and the camera rotating about the Z axis in synchrony.

(part 3)- Now let us look at the excitation pulse again from the point of the view of the rotating frame camera. In this case, we would see that the magnetization vector slowly tips down into the transverse plane without undergoing any precession. The reason for this should be clear, as our point of view is precisely following the rotation of the vector. As a result, the precessional motion of the spins can be dramatically simplified. A further consequence of this "rotating frame" point-of-view, is that if the spin precesses at a frequency which is slightly slower or faster than that of the rotating frame of reference, the magnetization will appear precess in this frame at a frequency which corresponds to the difference between the spin precession frequency and that of the rotating frame. Thus, the spin can appear to precess in either direction with its frequency dictated by this differential frequency.

Figure 11

Based on these ideas, is now easier for us to understand the fundamental mechanisms which influence contrast in MRI images. There are a number of factors which can influence the size and lifetime of the NMR signal. In the most trivial case, each tissue exhibits a slightly different proton density which will modulate all components of MRI images. However, tissues also exhibit a number of other unique NMR properties which are of greater biological interest which can be selectively imaged with appropriate pulse sequences. These include;

-Spin- Lattice Relaxation (T1), -Spin-Spin Relaxation (T2), -Susceptibility Effects (T2*).

A range of other factors can influence contrast in MR images of which motion is an example. In this presentation, we will focus our attention on understanding the basic mechanisms of relaxation effects T1, T2 and T2*.

Figure 12

In order to appreciate the contrast mechanisms in MRI, it is necessary to understand how the magnetization of the multitude of spins in the object combine to produce the NMR signal. Recall, that vectors which are parallel or "in-phase", add in proportion to their individual lengths. However, if the vectors are not aligned, or are slightly "out-of-phase", their sum will influence both their magnitudes and orientations as shown. Thus as the phase of the spins drift out of alignment, the overall magnetization which can be detected becomes smaller.

Figure 13

To illustrate this concept, let us consider four spins which are processing about the Bo field at exactly the same Larmor frequency. In this case, the resulting magnetization vector would have a vertical component (which is constant) and a precessing transverse component. As indicated earlier, the signal induced in the coil will oscillate at the Larmor frequency.

Figure 14

Now, let us consider these four spins in the laboratory frame and we see that the transverse component is rotating at the Larmor frequency. However when viewed in the rotating frame we see that the magnetization appears to be constant.

Figure 15

If we consider the signals associated with this situation, we see that the signal in the laboratory frame oscillates at the frequency dictated by the Larmor frequency of the spins while the signal as seen from the rotating frame appears to be constant. While the signal induced in the coil is that seen from the laboratory frame, the MR system "demodulates" the NMR signal to create the effect of the rotating frame and would generate this constant signal. The key point is that the demodulation frequency is that which is defined by the frequency of the rotating frame. Choosing a different rotating frame or demodulation frequency would cause the signal in the rotating frame to oscillate.

Figure 16

Now let us consider a different situation in which one of the nuclei is placed in a field which is slightly higher than its three neighbors. In this case, the magnetization from these three identical spins will add to form the large red vector seen on the right of this illustration. If we choose the rotating frame frequency to be equal to the Larmor frequency of these red spins, this will result in a vector in the rotating frame which appears to be motionless. In contrast, as the frequency of the blue spin is slightly higher than its neighboring red spins, it will appear to rotate slowly in the rotating frame. The frequency rotation of the blue spin will appear to be at the difference of frequencies of the blue and read spins as seen from the laboratory frame.

Figure 17 *(to start the video, press the play button)*

To continue on this line, let us consider that all the spins are experiencing slightly different magnetic fields. As a result, each will have its own unique Larmor frequency. If we choose the frequency of the rotating frame to be somewhere intermediate in frequencies of these for spins, we will see that none of the spins appear to be motionless in the rotating frame. As a result, we see the phase or orientation of each of the spins slowly migrates around the transverse plane. With increasing time, the spins continue to move until they become evenly distributed throughout the transverse plane.

Now, if we consider the signal associated from this set of spins, we can see that the vector sum (as represented by the gold vector in the lower right hand corner graph) starts off as initially aligned along the x-axis and then gradually decreases in magnitude as the components of this magnetization vector dephase. If we plot a graph of the magnitude of the gold vector as a function of time (lower left graph), we see that the magnitude of the signal decays exponentially towards zero. The time required reduce the signal to 37% of its original value represents the T2* decay time.

Figure 19

It is worthwhile to summarize what we have said. We have been discussing the effect of constant variations in the Larmor frequency for each spin in the object. In this case, we see that the signal decays with a time constant characterized by T2*. However, it would be natural to ask what causes these variations in the applied magnetic field and what they mean in NM images. First, we should categorize these field variations into two groups; those which are fixed in time as distinct from those which change with time. Fixed inhomogeneities of the magnetic field could be a result of the design of the magnet used to form the Bo field and as such are uninteresting from a biological point of view. However, even with a perfect magnet, the tissues being imaged can distort the fields as a result of their magnetic properties. Changes in tissue magnetic "susceptibility" can generate small field gradients on the order of a few parts per million than can vary throughout the tissue. As these inhomogeneities are fixed in time, the phase of individual spins will accumulate at a constant rate and result in signal loss. In this case, the dephasing due to T2* reports information about the magnet susceptibility of the tissue.

Figure 20

Up to this point, we have been considering that the only influence of the applied magnetic field is that arising from either the tissues or the magnet. However, we must remember that the spins themselves are slightly magnetic which can generate time varying field inhomogeneities. As such, motions of one spin relative to another can cause the magnetic field in the vicinity of each nuclei to undergo subtle time-varying changes in the magnetic field they experience. As spins move under diffusion or thermal motions, the inter-spin spacing and orientations can change and introduce a mechanism to subtly change the magnet field in a time-varying manner. This interaction of spin magnetic fields is referred to as "spin-spin coupling" and can cause the accumulation of a random spin phase which will cause signal decay. In addition, motion of spins through magnetic field gradients in tissue arising from susceptibility variations can be another mechanism for time varying accumulation of phase. Taken together, these effects generate signal loss characterized as "spin-spin relaxation" or T2 relaxation.

In any NMR experiment, spins will experience spin dephasing from both fixed and time varying field changes and both of these factors contribute to the decay time constant T2*. However, using a special combination of excitation pulses, known as a "spin-echo" of more properly a "Hahn echo" sequence, it is possible to measure the signal decay time constant arising only from the time varying changes in the magnetic field to quantify T2.

Figure 21

This figure is a schematic of the spin echo experiment which is worth reviewing before seeing an animation. The spin-echo pulse sequence uses two RF pulses as shown in the panel of the lower left hand side of this figure. The first pulse (a) tips the spins by 90° to force the total magnetization onto the transverse plane. Immediately after this pulse, the spins are in-phase and the NMR signal is maximum. Shortly thereafter, the spins undergo dephasing and the signal decays (at time b). At some arbitrary time (TE/2) after the first RF pulse, a second RF pulse is delivered which rotates all the spins by an additional 180° as shown in panel c. After this pulse, the spins move into an arrangement which mirrors their positions just prior to the 180° pulse. At a time TE seconds after the 90° RF pulse, we find that the signal reappears to form an 'echo'. The reason for this "spin-echo" can be understood by considering the phase of a single spin. During the first TE/2 seconds, a spin will accumulate a phase angle of 0 degrees (relative to the positive Y axis). After the second RF pulse, the phase of this same spin is now 180-0 degrees.
This indicates that the phase of the spin is exactly 0 degrees from the negative Y axis. Thus, assuming that the spin continues to accumulate phase at exactly the same rate, the phase will align with the negative Y-axis in an additional TE/2 seconds after the 180° pulse, or TE seconds after the original 90^0 RF pulse. This argument holds true for all the spins in the system, so that all the spins will re-align along the negative Y-axis to form a spin-echo at a time TE.

Now, if the rate of dephasing for all spins were constant during the TE interval, the magnitude of the transverse magnetization at TE would be identical to the magnetization immediately after the first 90° RF pulse. However, the time varying nature of spin dephasing outlined above, will create small variations in the phase angle accumulation for each spin before and after the 180 degree RF pulse. The net result is that the spins will not perfectly align at time TE resulting in a signal decay. The longer the TE interval, the greater the opportunity to accumulate spin dephasing and results in a greater signal loss.

Figure 22

(Part one)- To illustrate the details of the spin echo sequence with a bit more clarity, this animation shows the behavior of spin dephasing and RF pulses during the sequence. After the 90° pulse the spin are tipped into the transverse plane where they dephase and cause a reduction in the NMR signal. However at some arbitrary later time, we apply a second pulse which tips the spins through additional 180° degrees. We see that this causes the arrangement of the spins to invert into the mirror orientation discussed above and permits the spins to realign at a later time. As shown in the graph of the NMR signal, the echo increases after the 180 pulse to maximum and then decreases again.
(Part 2) If we repeat the same sequence and monitor the signal strength as a function of the echo time, we see that for a short echo time the spins will realign partially to produce an echo signal as seen by the gold NMR signal line. If we now repeat the sequence with a "longer echo time", we see that the time the spins spend in the transverse plane to accumulate random dephasing increases, so that the spin echo alignment is less complete than in the "short echo time" case. As a result the NMR signal during this longer echo time will be smaller. If we repeated this experiment for many different echo times TE, the envelope of the peak echo amplitude (gold

curve) would trace the T2 decay curve. The time taken for this curve to decay to 37% of its peak aptitude characterizes the T2 time constant.
(Part 3) In this view we repeat this last animation from a slightly different point of view so that we can focus on the transverse magnetization components alone. This part is optional and may not add to the clarity of what has already been seen. Depending on the time available, this can be omitted. Again, we see the effect of echo time to allow varying degrees of T2 or spin-spin interactions to influence the echo signal. It should be noted however, that the effects of field variations which are constant in time are no longer present as mechanism of a signal decay. As such, only T2 effects would influence this signal amplitude. By this means, it is feasible to differentiate between spin-spin interactions and signal decay which arises from susceptibility effects.

Figure 23 - T2 Summary

This slide summarizes the fact of the size of the NMR signal vs. the time TE and shows that the decay of the signal amplitude as a function of TE time characterizes the T2 time constant.

Figure 24 - Typical T2 Values in the Head

T2 relaxation times can vary substantially among tissues and is not greatly influenced by the applied magnetic field strength. Here we see some representative T2 times for different tissues in the head and shows that the relaxation times can vary widely.

Figure 25 - T2 Modulation of Image Contrast

In this figure, we show how T2 can modulate image contrast. Two tissues are depicted (blue and gold) which have different T2 times. After the echo formation, the signal is modulated by the effect of the differing T2 times for each tissue and results in a overall signal change for each tissue as shown. In this case, the long T2 components of CSF are clearly rendered as a very bright signal.

Figure 26 - Spin-Lattice (TI) Relaxation

In the preceding section, we reviewed how spin dephasing can cause a loss of detected signal. However, it is important to note that spin dephasing is a loss of the arrangement of spin orientation after the initial excitation pulse. In this case, the energy which has been deposited by the RF excitation pulse has not been dissipated but only lost to detection. However, if we wait longer, this energy will slowly leave the spin system and be distributed throughout the sample or the "lattice". This is analogous to the spin system "cooling off' as the excess energy is dissipated into the surrounding tissues. The time taken for this energy to be dissipated into the environment is referred to as the T1 or "spin lattice" relaxation.

If we have two tissues of differing T 1 times we see that the regrowth of longitudinal magnetization will occur at different rates. If we consider the bulk magnetization available at the time indicated (dotted line), we see that the yellow tissue will generate a smaller NMIZ signal compared to that of the red tissue.

Figure 27- Spin-Lattice (T1) Relaxation -Animation

(Part one) T 1 relaxation is illustrated in this animation and shows the magnetization as it is tipped from the longitudinal axis into the transverse plane after which the transverse component decays. As such, we see the longitudinal component as it gradually grow towards its equilibrium alignment and plotted on the right side of this figure. The time needed to reach 63% of the equilibrium longitudinal magnetization is referred to as the T 1 time constant.

(Part 2) Again, if we consider two tissues of differing T 1 times, we can see that the gold tissue exhibits a shorter T1 time and therefore returns towards its equilibrium magnetization faster than corresponding green tissue.

Figure 28- Spin-Lattice Relaxation Values for Various Tissues

T 1 values for biological tissues depend upon the applied magnetic field, whereas T2 times are relatively constant. Typical values for differing tissues in the head for a field strength of L 5 Tesla range between 500 and 1900 milliseconds. In many cases at 1.5 Tesla, we see that the T1 relaxation times are approximately five to tenfold longer than T2 times the same tissues.

Figure 29 - TI Modulation on Image Contrast

As we will learn, MR images require multiple RF excitations. In the case of the spin echo imaging method, this involves multiple repetitions of 90°-180° RF pulses at intervals of TR seconds. During this TR interval, the longitudinal magnetization will regrow to a value dictated by the tissue T 1 value. The shorter the T 1 value, the greater the regrowth of the magnetization in the TR interval. The net result is that the tissue with a shorter T1 time will be brighter in a spin echo experiment compared to the tissue with a longer T 1 time. This is illustrated in images the lower right where we see the signal from CSF (which exhibits a longer T1 time) is darker than either the surrounding white or grey matter which has a shorter T 1 time.

Figure 30 - Summary of TI and T2 Relaxation

We have seen that there are two main relaxation mechanisms of interest in proton MRI; namely, spin-lattice (T1) and spin-spin(T2) relaxation. Spin-spin relaxation is an example of a relaxation mechanism which is associated with a loss of spin order or phase and is seen as a transverse decay of signal after the initial spin excitation. The spin-spin time constant T2 is measured with a spin-echo experiment. A related time constant, T2*, is also seen as a signal decay from spin dephasing and its time constant is observed by simply observing the signal decay directly. The T2* time constant is always shorter than T2 as it includes dephasing mechanisms from both constant and time varying magnetic field inhomogeneities throughout the tissue.

In contrast, T 1 is the time needed for the spin system to dissipate the energy which was deposited in the tissue by the initial RF excitation pulse. In liquids, this energy dissipation is slow by comparison to the time to cause the spin system to dephase, thus T1 times are generally longer than either T2 or T2*. From this discussion, it is clear that T2 can never be larger than T1.

Image contrast arising from either T2 or T 1 relaxation mechanism is complex and is modulated by the timing of the MRI pulse sequence and size of the flip angles of the excitation pulses. Changing the TE times of spin-echo experiments will alter the T2 weighting while altering the TR interval or flip angle will control the T 1 weighting in MR images.

In the preceding sections, we discussed the very basic physics of how NMR signals are generated and the time constants dictating the nature of the evolution of the transverse and longitudinal magnetization. In this following sections, we will describe how, the unique physics of NMR can be used to create beautiful images of the anatomy.

Figure 31. - Overview of Lecture on the Physics of Image Formation

Now will focus our attention to considering how it is feasible to generate images of the variations of magnetization for tissues in-vivo. From the preceding discussion, it is clear that the relaxation times and the details of the pulse sequence dictates the contrast in images. In this section, we will not consider this interplay further but focus on the mechanisms for image formation. Specifically we will cover six points;

- The details of the structure of MR images,
- Fourier Transforms and their meaning, - The notion of magnetic field gradients, - K-space and its meaning, - How K-space is used in MRI and
- How these concepts work together to develop at MR imaging pulse sequence.

Figure 32 - Structure of MR Images

MRI is unique as a medical imaging method in terms of the relation between the detected signals and the final image. As in any digital imaging method, the challenge of MRI is to define the intensity of the NM signal for an array of pixels corresponding to differing points throughout the anatomy.

Figure 33 - The Question of Localization

Unlike all other imaging methods in current use in medical imaging, the signal detecting device (receiver coils) cannot be collimated to restrict the signal to a specific location as is done in x-ray imaging, ultrasound or radionuclide imaging. Rather, the MR coils detect signals which come from the entire object rather than a point from within it and the signal from each region needs to be individually resolved. In the following sections, we illustrate the mechanisms used to form MR images.

Figure 34 - The Spatial Location Task

We will do this by recognizing that our goal is to find the brightness of pixels located in a three dimensional co-ordinate system (X,Y and Z).

Figure 35 - Techniques for Spatial Localization

To achieve this, we will use three related techniques referred to as:

- Selective excitation,
- Frequency encoding and - Phase encoding.

In every case, these techniques will rely upon the use of magnetic field gradients which deliberately distort the magnetic field in the magnet. However, let us first consider the simplest of these concepts.

Figure 36 -Selective Excitation: The Ingredients

The task of defining the 3D distribution of image brightness generally starts with "selective excitation". As the name implies, this process creates a slab of tissue which is excited so that transverse magnetization is restricted to a specific plane of prescribed location and thickness. The technique involves the combination of NMR resonance, magnetic field gradients and a band limited RF excitation pulse.

Figure 37: Selective Excitation: An Analogy

A simple analogy to understand this concept is to consider an array of tuning forks as illustrated on the right hand side of this figure. Consider an "excitation" tuning fork set to ring at 440 Hz (the note A) as shown on the left side. When the excitation tuning fork is stuck, it moves air at a frequency of 440 Hz and this oscillating pressure wave propagates until all the tuning forks are bathed with oscillating air molecules. However, only the tuning fork set to ring at 440 Hz can absorbing energy in synchrony with the moving air so that it will ring or resonate. However, as the other tuning forks are not of the correct frequency, they cannot resonate with the moving air and remain silent. If we were to dampen the excitation tuning fork, the A note in the tuning fork array would continue to ring and emit its own sound. We know however, that the tuning forks were arranged in a linear array from F to C and since we used A to excite the array, we know that the middle tuning fork must have undergoing excitation. By this means, we can excite a specific tuning fork by the choice of the frequency of the excitation tuning fork. It follows that using an excitation tuning fork which is of higher or lower frequency will move the excited tuning fork to the right or left in the array.

Figure 38: Selective Excitation and NMR Resonance

This simple analogy is perfectly adaptable to NMR selective excitation. In this case, the spins can absorb energy only if the RF frequency is matched to the Larmor frequency and allow the magnetization to be tipped into the transverse plane. In contrast, if the excitation frequency is not matched to the Larmor frequency, no absorption of RF energy occurs so that the orientation of the magnetization within the object remains unaltered.. To create an arrangement similar to the line of tuning forks, we must invoke a means to change the Larmor frequency of the spins in space. To do this, we use a magnetic field gradient.

Figure 39- A Uniform Magnetic Field

Great care is used to build the magnets for MRI so as to achieve a highly homogeneous magnetic field within the magnet bore. This is illustrated in this figure which shows that regardless of the location of measurement of the magnetic field (as represented by the small vectors) it is found to be constant.

Figure 40 - A Magnetic Field Gradient (Gz)

In order to create MR images, the field must be distorted in a precise and controlled manner though the applications of magnetic field gradients. To illustrate this more fully, consider the object in the presence of a Gz gradient which refers to a gradient in the Z direction. A Gz gradient means that the field changes only in the Z direction and is constant for any point in a X-Y plane. In addition, to the Gz gradient having an orientation, it also has a magnitude. In this case, the magnitude of the gradient is rate with which the field changes per unit distance. Typical gradients can have values of 10 mT/m, meaning that the field varies by 10 mT for every meter of distance moved in object. In comparison to the size of the applied magnetic field (-1 Tesla), we see that these gradients represent small perturbations to the overall field.

Figure 41 - Selective Excitation and a Gx Gradient

In order to achieve a selective excitation, we apply a gradient to the object. In this case, we use a Gz gradient so that spins at the front end of the cylinder experience a smaller field and lower Larmor frequency than that near the opposite end.

Figure 42 - The Effect of RF Pulses in Selective Excitation

Now, applying a RF excitation pulse with a frequency to match that near of the center slab will result in a rotation of the magnetization to the transverse plane where it will continue to precess after the RF excitation. By controlling the range of frequencies used in the excitation pulse, we can control the width of the slice while controlling the center frequency of the pulse, allows control of the location of the slice. By this means, we have now generated transverse magnetization in a slab of a specific location and thickness.

Figure 43 - In-Plane Localization

At this point, we have now generated transverse magnetization in a thin slab of tissue somewhere within the body. The remaining task is to identify the brightness of spins within this slab which is achieved with two different techniques referred to as frequency and phase encoding. In general, frequency encoding is a process wherein we measure the NMR signal in the presence of the gradient chosen in a specific direction (for example the X direction). Phase encoding refers to the application of a gradient prior to the frequent encoding gradient and is used to encode information along the orthogonal direction (such as the Y direction.). In reality, the orientation of the gradients use for selective excitation, frequency and phase encoding are arbitrary with the only constraint being that the three orientations are mutually orthogonal. By this means, it is possible to create images in arbitrary orientations.

Figure 44 - The Relation Between the MR System and Image Formation

In order to explain the basics of how the in-plane localization task is performed, we will proceed in a two step manner as illustrated in this figure. First, we will show that MR images can be constructed from so-called "K-space" data by some simple illustrations. Once an intuitive understanding of the nature of 2D K-space has been established, we will then indicate how the MR imaging system generates the image signals in the form of the required K-space data.

Figure 45 - Image Space vs K-Space

To start our understanding of MR images formation, we need to understand the relation between MR images and their "K-space" representation. As shown in this figure, we see that the image has coordinates X and Y while the K-space data has coordinates Kx and Ky. The units of X and Y are in units of distance (ie centimeters) while the units of Kx and Ky are in units of 1/distance (ie centimeters^{-1}). Thus we see that the K-space dimensions are somewhat unfamiliar as they are expressed in reciprocal distances. The gray scale of the K-space data reflects the value of the data at positions Kx and Ky. What is the meaning of the data in K-space?

Figure 46 - A One Dimensional Problem

In order to understand this question, we must make a small diversion to consider the mathematical properties of an operation referred to as a Fourier Transform. In order to build an intuitive understanding of a Fourier transform, let us consider a simple problem of attempting to construct a mathematical formulae for the 1-D object or "target" function shown in this figure. The function is unusual, as it switches discontinuously from 0 to 1 over the region of interest (ROI) shown.

Figure 47 - A Crude Fourier Approximation

In order to appreciate how this can be done, we first consider the average value of this function which spends 50% of its range with a value of 0 and the remaining 50% with a value of 1, to give an average value of 0.5. So our first approximation of the target function is a constant

value of 0.5 over the region. Next, we will add a sinusoidal function of 1 cycle over the ROI and the amplitude of 0.5. This sine functions has a "spatial frequency" of a certain number of cycles per ROI. We can see that in this case, the frequency of our sine function is exactly 1 cycle per ROI. Similarly, the constant value undergoes no cycles per ROI and therefore has a frequency of zero. When we add the constant to this sine curves we see that our approximation oscillates over the right range but does not have the sharp discontinuous edges we need to match our target function (shown dotted).

Figure 48 - A Better Fourier Approximation

In other to improve the accuracy of our approximation, we add another sine function with a frequency of three cycles per ROI and amplitude of -0.15. When this is added to the previous lines, our approximation has been improved.

Figure 49 -The Definition of K-Space

Instead of plotting all the sinusoidal curves over the ROI, let us simplify our representation by plotting a graph of the amplitude of the sine functions we use versus their spatial frequency. This is a short hand graphical notation for the family of sine functions which are needed to be added together to approximate the target function. The horizontal axis has units of spatial frequency (cycles/distance or in this case cycles/ROI) while the vertical axis has units of amplitude. This is the K-space representation of our object! !

At this point, it is composed of three sine functions of frequency 0, 1 and 3 cycles/ROI. We can also represent this 1-D K-space data as a gray scale plot shown on the bottom left hand graph of this figure. In this case, the brightness of the gray scale represents the amplitudes of each of the sinusoidal waveforms while their frequency is expressed in terms of the positions of the gray scales along the horizontal axis.

Figure 50 and 51 - Successively Better Approximations

By continuing to add more sine functions of correct amplitude and frequency, we get arbitrarily close to the desired target function. The K-space representation and image domains are related through a mathematical operation called a Fourier transform. While the details of how this transform operates is beyond the scope of this lecture, the essential point is that it calculates the amplitudes and frequencies needed such than when the various sine functions with these properties are added, we get the desired target function. We see from this simple example, that it is possible to represent our target function which is discontinuous with smooth sine functions of carefully chosen frequency and amplitude. The plot of frequency versus amplitude is the K-space representation of the object.

Figure 52 - Two Dimensional K-Space and Image Space

In order to represent a two dimensional function (ie an image), one needs to use sine functions which oscillate in both directions (ie left-right as well as up-down). This is shown in this figure, where we show the K-space representation of the head image. Again, the relation

between the K-space data and the image data are through a Fourier transform. At any point in the K-space domain, the grey scale tells us the amplitude of the sine function, while the location of the point tells us its frequency and orientation. For example, the yellow point, corresponds to a sine function which oscillates at a certain spatial frequency and has a vertical structure.

Figure 53 - The Meaning of Various Points on K-Space

Other points in the K-space domain show us different frequencies. The center of K-space (interception of the dotted lines) corresponds to a spatial frequency of zero. Considering various other points, we can see that sine patterns of varying frequency and orientation are represented. To simplify our language, lets us refer to these patterns of variously oriented sine functions as "stripes". These stripes can have an orientation, brightness and spatial frequency. For example, points on the Kx axis refer to stripes which are vertical. Points on the Ky axis reflect stripes which are horizontal. For K-space points with both a Kx and a Ky value, the resulting stripes are oblique. The angle of the stripe pattern is such that the strip density in x and y corresponds to the spatial frequency of the Kx and Ky component of the point in K-space. Remarkably, by combining all the points in K-space with their corresponding stripe amplitude and frequencies we generate the head image which we have seen.

In summary, we see that the K-space representation is simply a short hand graphical notation which tells us the family of stripe patterns of varying frequencies and amplitudes needed, such that when added together they provide the desired image. The relation between the K-space representation and the image is through a 2 dimensional Fourier Transform. If we need to create an image with 256x256 pixels in the image domain, the number of points in the K-space domain must also be 256x256 pixels. Similarly, a 3D object can be characterized by a 3D K-space data set.

Figure 54 - The Question of How Stripes are Made in MRI

In the preceding section, we discussed the relation between the K-space domain and the image domain. In this final section we outline how the MR imaging system generates data directly in the K-space domain. We now recognize that the K-space domain, represents the image data as "stripe" functions of varying orientation, spatial frequency and amplitude. The question we address in this section, is how the MR imaging system generates these stripe functions and how it determine their correct amplitude, so that when added together, they form the final image.

Figure 55 - Return to The Relation of the MR System and Image Formation

We now return to our earlier slide which shows the relation of the MR Imaging system, the K-space data and the resultant MR image of the head. We now focus on the role the MR imaging system plays in generating the K-space data.

Figures 56 and 57 - Gradients in X and Y

Once again, we need to remind ourselves of the nature of gradients which will be used to achieve the generation of the stripe patterns mentioned above. In this slide we see the graphical

illustration of a Gx gradient. In this case, the field increases with the +ve X position and decreases with the -ve X location and is constant for all positions within a given Z-Y plane. Finally, the Y gradient causes the field to change only in the Y direction and is constant within a Z-X plane. As gradients have both magnitude and direction they can be represented as vectors and can add to generate gradients in any direction by the simultaneous application of varying amounts of Gx, Gy and Gz gradients.

Figure 58 - An Alternative Representation for Magnetization

Before we proceed further in our discussion of frequency and phase encoding, it is helpful to change our representation of the magnetization from the preceding sections of a rotating vector to something more simple. Specifically, rather than drawing a rotating vector which induces a signal in the coil, let us represent the magnetization by a sphere or ball. The ball will rotate on its axis with one side of the ball colored black, while the other side is white. As the magnetization rotates, the ball revolves about its axis showing a progressive change from the white side to the black side. By looking at the progression of the shading of the ball , we can see the progression of the phase of the spin as it evolves in varying magnetic environments.

Figure 59 - The Effect of a Gradient on an Array of Magnetization Balls

For our first example, let us consider a linear array of balls. After the excitation is created, the magnetization of each spin is in phase, and as such, all the balls are in the same condition, that is, with the white face exposed, (as shown at time a). Next, we will consider the application of a gradient along the direction of the balls. A plot of the time evolution of the gradient waveform (top graph of this illustration) and the resulting magnetic field deviation in space are shown at the bottom of this figure. Initially, the balls have their white sides exposed. However, when the gradient is turned on, the balls experience slightly different magnetic fields. On the extreme left side, the magnetic field deviation is negative and causes the balls to rotate in a clockwise manner in the rotation frame. As we consider balls which are closer to the center of the array, the rotation rate of the balls decrease until we reach the center ball where the magnetic field deviation is zero. Continuing further to the right, we see that the field increases gradually, resulting in a counter clockwise rotation rate of gradually increasing frequency until we reach the extreme right hand side of the array. If we consider a later moment (time b) in the gradient evolution, we see that the balls show a shading from left to right. Now let us observe this evolution carefully as the gradient continues to play in the next animation.

Figure 60 -The Effect of a Gradient on an Array of Magnetization Balls (Animation)

In this illustration, a graph of the gradient waveform is shown (right side) with an arrow indicating the point in time. We see that when the gradient is initiated, the balls rotate with a speed and direction corresponding to the local magnetic field as discussed in the previous slide. However, we note that as the gradient continues, the balls rotate to form various patterns of black and white. To understand the significance of this, we next consider the same effect for twodimensional array of balls.

Figure 61- Creating Vertical "Stripes"

In this case, we have replaced the linear array of balls with a 2D array of balls and a plot of the gradient waveform in X direction is illustrated. As the gradient is played, we see that the array of balls converts from a uniformity white distribution to form a series of vertically oriented stripes of gradually increasing density as the gradient is applied. We see that the density of the stripe pattern (ie the spatial frequency), increases with the duration of the gradient application or the area under the gradient waveform as shown by the red region under the gradient waveform. This allows us to create vertically oriented stripe patterns of gradually increasing spatial frequency. The next question would be how to generate a horizontally oriented stripe pattern.

Figure 62 - Creating Horizontal "Stripes"

As one might anticipate, the only requirement to create horizontally oriented stripes is to orient the gradient in a vertical fashion as shown in this figure. Again, as the gradient waveform is played out, a horizontally oriented pattern of stripes emerges with the density of the stripe pattern (spatial frequency in the Y direction, ie Ky) increasing in proportion to the area swept out by the gradient as shown by the green region. The last remaining question is how to create stripe patterns of arbitrary orientation.

Figure 63 - Creating Oblique Stripes and K-Space

As one might expect, this requires the use of both the Gx and Gy gradients in sequence. In this figure, we show a plot of the gradient in Gy and Gx as a function of time. As shown, the amplitude of the Gy gradient waveform is incremented (with fixed duration) after which a fixed Gx gradient waveform is applied. We also show a graph were the position of a point in the graph corresponds to the area of the gradient waveforms as they evolve (shown in the lower right hand plot). In this plot, the horizontal axis is the area under the Gx gradient while the vertical axis is the area under the Gy gradient. Recall that the spatial frequency of the stripe pattern increases with increasing exposure to a gradient and is represented by the area of the under the gradient waveform. Thus the area of the Gx and Gy gradients corresponds to the spatial frequency Kx and Ky respectively. Thus, as we increment the Gy gradient amplitude, a point on this plot moves progressively along the Ky axis. After each Gy gradient application, the application of the Gx gradient causes that point to progress along the Kx direction. Thus by combined application of the Gx and Gy gradients, we can move throughout all points in the K-space plot.

To start the video, press the play button

Now, if we observe the array balls during the motion of the point in K-space, we see the orientation of the stripe patterns corresponds to the K-space location. Initially if there's no Gy gradient we generate a vertically oriented stripe pattern. After the application of incremented Gy we generate a horizontal stripe pattern which shifts to become oblique with the subsequent application of the Gx gradient. By observing the interplay between the Gx and Gy gradient it

becomes clear that arbitrary stripe patterns are feasible both in terms of their orientation and spatial frequency by the combined application of the Gx and Gy gradients.

Figure 64 - Oblique Stripes: A Summary

To summarize the application of the Gx and Gy gradients we show again the effect of applying Gy gradient (phase encoding gradient) initially to create horizontal stripes which then followed by the Gx gradient (readout gradient) to generate stripe patterns of increasing obliquity and spatial frequency.

Figure 65 - How does the MRI System Measure the K-space Signals?

At this point in our discussion, we have shown how the application of gradients can create the stripe pattens of varying orientation and spatial frequency needed to encode our object. The remaining issue to illustrate, is how the MR imaging system determines the correct amplitude for each point in K-space needed to correctly encode the object. This is done by measuring the signal as viewed from the rotating frame detected by the coil during the application of the Gx gradient.

To start the video, press the play button

In this case, we represent the object as a transverse head image made up of our tiny magnetization balls. During the application of the gradient, the spheres generate the stripe patterns that we have discussed and generate a changing signal which is induced in the RF coil. This signal is periodically sampled by a computer during the application of the Gx gradient. The amplitude of this signal corresponds to the desired K-space amplitude for each point in the Kspace plot. By watching the RF signal and the data which is sampled for various trajectories in Kspace, we can see how three lines of K-space are filled. If we had repeated a larger number of Gy increments, multiple lines in K-space would have been created until the entire K-space region is sampled to form the image. Accordingly, it should be clear that in order to sample 256 lines in K-space, each with its own Gy gradient, it will require 256 separate trajectories through K-space. As a result of the requirement for repeated applications of the Gy gradient, the overall time required to fill K-space from MRI acquisitions can be long.

Figure 66 -A Simple (but incomplete) Pulse Sequence

We can summarize what we have said about the K-space plot and now build a partial MRI pulse sequence as shown in this figure During the RF pulse, a slice is selected in the presence of a Gz gradient. Then an incremented Gy gradient is used to precede the Gx gradient waveform during which the echo is formed and the NMR signal is sampled. The sampled data is then applied to the points in K-space corresponding to the area under the gradients which the spins had experienced at the point of sampling_

Figure 67 - **The Four Quadrants of K-Space**

It will be noticed by looking at the K-space data, that the most intense region occurs near the region were Kx and Ky are both zero. Furthermore, it is clear that using the pulse sequence which we have discussed, only the positive quadrant of K-space can be sampled. Ideally, we would like to sample all four quadrants of K-space which can be achieved by using a more realistic pulse sequence shown in the next figure.

Figure 68 - A More Complete Pulse Sequence

By simply allowing the phase encoding gradient to have negative increments and by preceding readout gradient with a negative lobe, it is possible to sample all four quadrants of Kspace. If we need to form an image with NxN pixels, then we need to take NxN samples in Kspace. Thus N separate periods of data acquisition are needed in order to collect the N incremented Gy gradient applications to form a NxN pixel image. The separation in time between the Gy gradient applications is TR seconds and is an important parameter which is used to control the T 1 weighting of the image.

Figure 69 - Fourier Reconstruction of K-Space: Part A

Finally, it is worthwhile to understand the reconstruction of the K-space data to the final image through a Fourier transform. In the next two figures, we will illustrate that as we the migrate through K-space, the stripe patterns are added to create the final image. First, let us consider a single point (red spot) as it moves through K-space and look at the corresponding "stripe image" which would result. We will multiply the stripe pattern by the amplitude obtained from the K-space date at the location of the red spot. We will then add the resulting stripe patterns into an image called the "sum-of-stripes" image. As we do so, we can observe a partial integration of the K-space data. At this Point click the play button to start the video

First, we noticed that as the spot most from left to right, the stripe patterns are vertical and are very dense. As we progress toward the center of K-space, the stripe pattern becomes more course. Furthermore, all the stripe patterns for this particular trajectory are vertically oriented as the trajectory was for Ky=O. Now by observing the sum of the stripe patterns weighted by the corresponding K-space amplitudes, we can see the sum-of-stripes image shows only horizontal structure.

Next, the spot moves to a different point in K-space and sweeps out a second line with a negative value of Ky. In this case, we see that the stripe space image and the sum-of-stripes image shows a different structure. Specifically, the stripe space shows oblique strips of varying orientation and density as the point moves throughout K-space. Also, the sum-of-stripes image now shows both horizontal and vertical structure. As we continue for a third line in K-space, we can see the same pattern emerging but now the sum-of-stripe image exhibiting finer structure in the vertical direction.

Figure 70 - **Fourier Reconstruction of K-Space: Part B**

In the previous illustration, we showed that by adding the stripe patterns along a line in K-space, a complex patterns emerges in the reconstructed sum-of-stripes image. Furthermore, we can see the orientations of the stripes and how they contribute to the vertical structure within the emerging image. However, in order to see the final reconstructed image, we now consider that the stripe space image represents the summation of all the K-space data along the red line shown overlying the K-space data. By sweeping the red line over K-space and integrating the stripe space data, we will be able to reconstruct all the K-space data as needed to form the final head image.(At this point click on the play button to the video play) You will note that for the red line near the bottom of K-space, only very fine detail of the head image emerges. However, as the red line moves towards the center of K-space and crosses it, the overall contrast patterns of the head image suddenly emerges. Finally, as the red line moves to include the upper portions of Kspace, the fine detail of head image is added. As such, it should be clear from this illustration that the very center K-space provides the majority of the overall contrast in the image, while the contributions from the top and bottom regions of K-space contributes the high-resolution details.

Figure 71 - **Conclusions I**

In conclusion, we have seen that the task of achieving spatial location in MR images is achieved by the applications of three orthogonal magnetic field gradients together with appropriate RF pulses. Selective excitation is used to define a slice with the location and width of the slice defined by the center frequency and frequency range used in the selective excitation pulse. Once a plane of transverse magnetization has been generated, we then precede to use phase encoding and frequency encoding to determine the brightness of points within the plane.

Figure 72 - **Conclusions II**

Frequency encoding is the procedure whereby we measure information in one direction and involves measuring the MRI signal in the presence of the readout gradient. This MRI data is used to define the amplitude of the K-space data for particular line in K-space which corresponds to the area under the readout gradient.

Figure 73 - **Conclusions III**

The phase encoding gradient is used to define information along the last remaining direction which is orthogonal to the readout direction. By incrementing the phase encoding gradient it is feasible to generate varying locations in the Ky direction in proportion to the area under the phase encoding gradient waveform. By the combined application of phase and frequency encoding, it is feasible to traverse the entire K-space region. As a result of the need for incrementing the phase encoding gradient, MRI data acquisition requires repeated RF and gradient pulses on the interval of TR seconds in order to completely sample all of K-space.

Plewes Consulting

Toronto, Ontario
Canada

don.plewes@rogers.com
22