Me myself and MRI:
adventures in not understanding nuclear physics.

Thomas E. Gladwin

August 28, 2007

Contents

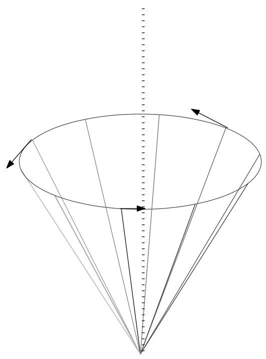

Figure 1: Precession

1 Introduction

The following text contains the current state of what I've pieced together about MRI physics. It's intended for people who've read a few sources and are trying to get a little closer to the bottom of things. The following sources were used heavily in this text.

http:////www.e-radiography.net//mrict//Basic_MR.pdf
University Physics by Young and Freedman.
All You Need to Know About fMRI by Moriel NessAiver, from http://www.simplyphysics.com/TEXTBOOK.HTM.
http://www.cs.unm.edu/~brayer/vision/fourier.html and
http://www.ebyte.it/library/educards/mri/K-SpaceMRI.html

2 Nuclei

2.1 Precession

The nucleus contains an atom's protons and neutrons. In the case of a hydrogen atom, the atom that is targetted by MRI, the nucleus consists of a single proton. Protons have a property called spin which results in a tiny magnetic field. The north - south axis of the field lies along the axis of spin. The proton can be described as a bar magnet, and its field can be characterized by a magnetization vector pointing in the direction of the north - south axis. In an external magnetic field B0, the magnetization vector precesses around an axis aligned with the magnetic fieldlines (see figure 1). *All you Need to Know about MRI Physics* describes precession using a dreidel analogy: the axis along the length of the dreidel doesn't stand up straight, but its top end swings along a circle lying flat in the air, with the centre of the circle lying on a vertical line rising from the tip of the dreidel touching the ground. Alternatively, draw a circle in the air using your finger without moving your hand.

Every type of atom in a field of a given strength is associated with a unique precession frequency for its protons. The frequency is related to the strength of the external field B_0 through the Larmor equation and the gyromagnetic ratio of the nucleus γ [MHz / Tesla]:

$\omega = \gamma B_0$.

For hydrogen, γ is around 50 MHz. So, in a two Tesla field, the magnetization vectors of hydrogen nuclei will precess at around 100 MHz.

2.2 Spin-up and spin-down nuclei

The spin of a proton in an external, static magnetic field field has two energy states: ‘spin-up’, in which case the magnetization vector is parallel, and ‘spin down’ in which case the vector is anti-parallel to the external field. In a given set of very many protons, the high and low energy state will be distributed roughly equally, but with a slight excess of low energy state protons. The stronger the magnetic field, the greater this spin excess (fewer protons "have the energy to go against the field"). The spin excess results in the net magnetization vector (NMV), the sum of all the protons’ individual magnetization vectors. Because more vectors are parallel than anti-parallel, the longitudinal component of the NMV is non-zero. Because the precession phases of the magnetization vectors are not synchronized, the transverse components of the vectors average out to zero.

3 The radio-frequency pulse

At the descriptive level, what happens when a radio-frequency (RF) pulse is emitted by the MRI machine is that the NMV is flipped towards the transverse plane and rotates around the axis of the static magnetic field. However, to understand the processes of relaxation, 180 degree flips for spin echoes, T1- and T2-weighting discussed next, at least a useful lie about the relationship between individual protons and the NMV must be established. The following sections contain the best story I could put together at this point from various sources and considerations. I hope at least the spirit of the explanation is correct.

The ‘pulse’ in ‘RF pulse’ refers to an electromagnetic wave generated by RF coils in the MRI machine. The term radio-frequency refers to the wavelength of the wave, which together with gradient fields allows the selection of slices of nuclei. These concepts are the subject of the following sections.

3.1 Electromagnetic waves

Electromagnetic waves consist of electric and magnetic fields linked together and propagating through space. Propagation means that the region of space in which the fields exist grows. In vacuum, the wave propagates at the speed of light. There is a second kind of movement involved in electromagnetic waves

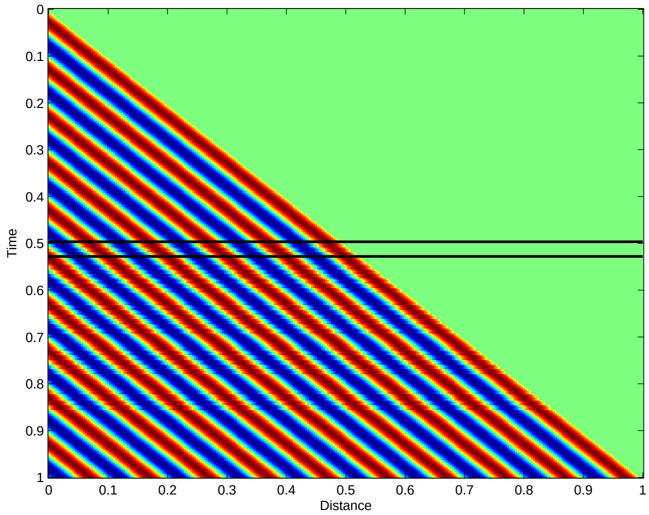

Figure 2: A travelling wave: a function of distance and time

such as radio-frequency pulses: the field strengths at points within the wave can oscillate over time.

Electromagnetic waves arise from the interaction of changing electric and magnetic fields, following from Maxwell's equations for electromagnetism. These four equations turn out to have a solution in the form of the travelling-wave differential equation:

$\frac{\partial^2 y(x,t)}{\partial x^2} = \frac{1}{v^2}\frac{\partial^2 y(x,t)}{\partial t^2}$

To understand the equation, imagine a plane with time t as the vertical and distance x as the horizontal axis, increasing from top to bottom and left to right respectively (see figure 2: to see how the wave changes as time passes, move the 'window' downwards). We start off with all values y(x, t) at zero. Now start a sinusoidal oscillation in the top-left corner. That is, we give the left-most point a speed of 1 and make it oscillate over time. The acceleration of such a change is $\frac{d^2 sin(t)}{dt^2} = \frac{dcos(t)}{dt} = -sin(t)$. As soon as the value at y(0, 0) rises, the acceleration becomes negative. The acceleration of y in the x direction is now also negative by the wave equation, so the derivative of y in the x direction will immediately start to decrease. We now have a non-zero value at $y(0, \delta t)$, that has smoothly generated a negative derivative as it rose. That means that it must have lifted up neighbouring points (otherwise the function would be discontinuous and the second derivative along the x-axis wouldn't exist, let alone fulfill the wave equation). Since these points are now also accelerating in time, they also perturb their neighbourhood. Hence, the wave propagates over space, with a speed dependent on the v parameter that relates spatial to temporal acceleration.

After some time, a region of space will be covered by the wave. All points in this region will be oscillating with the frequency of the source oscillation (the points of zero acceleration propagate with the same speed, preserving the interval in time between them). However, the further points are from the source, the more their acceleration will lag behind. At some distance, the lag will be so great that the oscillation will return to the phase of the original again. Thus spatial waves arise with a constant wavelength, dependent on both the frequency of the source oscillation and the propagation speed.

Maxwell's laws demand that the electric and magnetic fields satisfy the wave equation, playing the role of y and hence pulsing or rippling outwards from an oscillating source. The oscillations at points in the wave result in a wave of fields, travelling along the direction of propagation. Since the propagation speed of electromagnetic waves (in vacuum) is constant, the frequency of oscillation can be described by the wavelength: the distance between peaks in the travelling wave.

3.2 Electromagnetic waves and nuclei

An electromagnetic wave brings energy to the regions of space at which it arrives. When the wave encounters a nucleus, whether anything happens depends on the wavelength. If and only if the wavelength is such that the oscillations in the field resonate with the precession frequency of the nucleus, two things happen.

The first effect is excitation, the transmission of energy to nuclei in the lower energy state. The excitated nuclei go from the 'spin-up' to the 'spin-down' state, reducing the spin excess and hence the longitudinal component of the net magnetic vector, NMV.

The second effect of the pulse is to synchronize the precession phases of the nuclei. The NMV itself swings around the axis of the magnetic field. In other words, the NMV gains a transverse component, circling the axis of the static field.

If exactly half the spin excess is excited, the NMV is zero in the longitudinal axis and all the magnetization is transversal. This is called a ninety degree flip.

3.2.1 Problems

I have some problems with the above depiction. First, I can't see any reason for the transversal magnetization to be equal to the initial longitudinal NMV, although this is generally how the 'flip' of the NMV is portrayed. As discussed later, there are certainly effects of repeating RF pulses before the population of excitated nuclei have reverted to the low-energy state. Yet there seems to be no explanation for a relation between the loss of longitudinal and gain of transverse NMV.

A further point of difficulty to me is that 180 degree flips are also possible (see below), but this seems hard to relate just to exciting nuclei in the above fashion.

An attractive alternative explanation would be that the angle of precession increases in a number of discrete steps from, say, 35 to 145 degrees, as energy

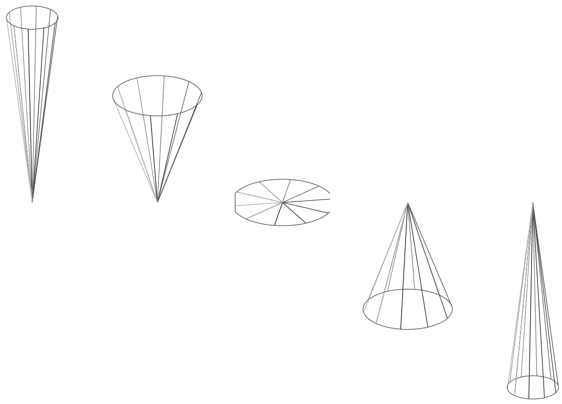

Figure 3: Speculative energy states of the precession angle

is transferred to the nucleus, and that spin-up nuclei become spin-down (and hence change precession direction) if they absorb energy at their maximal angle. This is not the standard explanation and could very well be completely off base in terms of physics. However, so far, it's the best story I can come up with to explain the behaviour of the NMV and relate it to some kind of nucleic energy state.

4 After the pulse

4.1 T1 and T2 relaxation

After the pulse is switched off, two processes occur. First, the energy transferred to the protons is released again, and they return to their low-energy state. As more and more protons release their energy, the spin excess and hence the NMV aligned with B0 is recovered. This recovery is called T1 relaxation, and can be described as an exponential recovery process with a time constant T1 in the order of 1 s.

The second process is the dephasing of the protons' precessions. Under ideal conditions, this is due only to random spin - spin interactions. If one proton moves towards another, its magnetic field may be aligned such that it opposes or adds to the static magnetic field experienced by the other. Hence, the proton will speed up or slow down due to the Larmor equation, and after the protons move apart their phases will have changed. The mass effect of this is the reduction of the transverse NMV. Dephasing is called T2 relaxation, and can be described as an exponetial decay function with a time constant T2, in the order of 10 ms.

Note that the two recovery processes are independent; in particular, they do not preserve the length of the total NMV.

4.2 T2* relaxation and the spin echo

Dephasing also occurs due to non-random inhomogeneities in B0, causing faster dephasing than would be expected from spin - spin interactions alone. This is called T2* relaxation. The effect of non-random dephasing can be cancelled by using (e.g.) spin echoes, which bring back the signal to the T2 relaxation level. Spin echoes are created by flipping the NMV by 180 degrees using a second pulse, given after a period TE/2 following the initial pulse. After a second period of TE/2, the regions with faster precession, which will have rotated further, will have travelled back to the same phase as slower regions, travelling back at their slower speed from a smaller rotation. Therefore the signal at TE, called the spin echo, provides a true measure of the T2 relaxation at that time.

5 What gets measured?

The coils that emit the RF pulse serve as sensors after the emission. The sensors measure the size of the transverse NMV. Due to the rotation of the transverse NMV, two sensors placed at ninety degress to each other in the transverse plane will experience oscillating magnetic fields (causing an alternating current in sensor coils) with a ninety degree phase lag. The importance of this phase lag is discussed in the text on K-space.

Whether the signal in the sensor coils is best described as an induced current due to the changing magnetic field of the NMV, or as the RF energy released and sent back by the excited nuclei, is not clear to me at this point.

6 Contrasts: T1 and TR, T2 and TE

Protons in different types of biological tissue have different values of T1 and T2. By varying the echo time TE and repetition time TR (the interval between RF pulses), the image from the MRI can be made to reflect contrasts between regions with a high versus low density of a chosen tissue.

First, compare tissue A with a short T1, $T1_A$ and tissue B with a long T1, $T1_B$. Let TR be equal to the short $T1_A$. Now, after an initial pulse, 63% of the B0-aligned NMV in tissue A is recovered by time TR. Now let $T1_B$ be long enough that, at TR, only 10% of its B0-aligned NMV has recovered. Both tissues now receive a second RF pulse. Because the pulse 'flips' the longitudinal NMV into the transverse plane, tissue A will now have a greater transverse NMV than tissue B. The 'flipping NMV' model describes what happens, but the causes must have to do with the energy states due to the preceding pulse or pulses. One explanation is that roughly the same protons tend to be excited, due to the temperature of their immediate environment for instance. If the first pulse led to a ninety degree flip from an initial state of 35 degrees precession angle, an immediate second pulse would tend to "click" the same protons through to their 145 degree energy state (in my speculative model), and result in a smaller transverse NMV than the first.

In any case, by choosing a TR such that tissue A can recover its longitudinal NMV while tissue B has only recovered a little, tissue A and B will generate high and low signals respectively. Such an image is called T1-weighted: two regions with different T1's will have different signal strengths.

Second, compare tissue C with a short T2 and tissue D with a long T2. Now, if we lengthen TE, the protons in tissue C will dephase faster and hence tissue D will retain the stronger signal. Thus by choosing a long T2, at which signal from D is still strong but signal from C has decayed, a T2-weighted image can be made. Such an image will show high signal intensity for tissue D and low for tissue C, in other words, the contrast between C and D will be high if they have different T2's.

Say we have tissue X, with long T1 and long T2, and tissue Y with short T1 and short T2. To give X as much of an advantage over Y as possible, we'd use a long TR (so that X can recover its longitudinal NMV) and a long TE (so that the T2 decay in Y has weakened its signal).

7 Slice-selection

Recall that the precession frequency ω of nuclei depends on the strength of the static magnetic field, B0. The relation is given by the Larmor equation:

$\omega = \gamma B_0$,

where γ is the gyromagnetic ratio of the nucleus. Each atom has a unique gyromagnetic ratio. For protons (hydrogen atoms), γ is 42.56 MHz / Tesla. When a gradient is imposed on B0, that is, when the strength of B0 is made to decline linearly from e.g. head to feet, the Larmor equation shows that the precessional frequencies will also run from high to low along the direction of the gradient. Gradients are imposed using additional coils to modulate B0.

Due to the gradient, the resonance frequency of nuclei is different for every slice of the body along the gradient. By fine-tuning the frequency of RF pulses, they will only excite nuclei from a specific region of space. All subsequent processes, such as frequency encoding within the slice (as will be discussed in the context of K-space), will therefore only result in measureable effects from the initially excited slice. The above is called slice selection.

8 K-space

K-space is the 2D Fourier representation of the signal intensities distributed over a slice in the MRI scanner. What this means and how it helps get information on the signal from a specific point in a slice out of the summed NMV is the subject of the next sections.

8.1 Fourier representations in 1 and 2D

So what's a Fourier representation? In 1D, you have a signal sampled at N points, and the Fourier transform describes this signal as a set of N cosines of

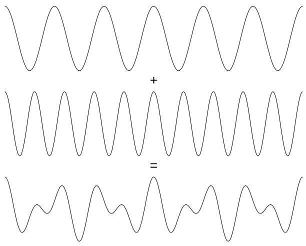

Figure 4: Adding 1D Fourier components

increasing frequencies, each with an amplitude and a phase parameter. When added together, the cosines result in the original signal. The *spectrum* just plots the amplitude against the frequency of each cosine. The cosines have frequencies running from 0, $\frac{1}{T}$ to $\frac{N}{T}$, where T is the time of the signal. Each frequency is $\frac{1}{T}$ higher than the preceding one. What the amplitude spectrum doesn't show is the phase of each cosine, which might well be of secondary interest anyway. As an example, if my spectrum is all zeroes except for a 1 for the 3 Hz cosine and a 2 for the 5 Hz cosine, then the signal is the weighted sum $cos(2\pi 3t) + 2cos(2\pi 5t)$. So, for instance, say you add a random blip to your spectrum. Then the signal it encodes will suddenly include some additional oscillation. So going from a Fourier representation to your signal is very simple: you just add cosines, weighted by the amplitudes, and shifted by the phases. The maths of going from the signal to the Fourier representation involves nothing more obscure than integration and really just measure how much your signal looks like a cosine of frequency ω and how much like a sine of frequency ω. If the signal looks just like a cosine and not at all like a sine, the phase must be 0; if it looks exactly opposite to a cosine and just as much like a sine only in-phase, then the phase must be 135 degrees. And so forth. For an intuition of this, draw some vectors on a unit circle and see how large their projections onto the horizontal (called *real* in Fourier analysis, or *in-phase*) and vertical (*imaginary* or *out-of-phase*) axes are.

Finally, note that the familiar graphical representation of an oscillation over time of a cosine would in reality be measured by a sensor sitting at one point somewhere, sampling a signal, that could also sit at a single point in space, increasing and decreasing in amplitude as time passes. Think of a clock on a tower, where our signal over time is the changing length of the shadow of its hour hand on the ground.

So what is Fourier transformation in two dimensions? Instead of building a signal from cosines, we'll now build it from *sheets* of cosines. These sheets are filled in with cosines, not with $2\pi\omega t$ as input as in a 1D signal, but with $\omega_1 t_1 + \omega_2 t_2$. The cosine now looks like this: $\cos(2\pi(\omega_1 t_1 + \omega_2 t_2))$, which is a two-dimensional function because of the two variables t_1 and t_2 that are required as input. So for every *combination* of frequencies ω_1 and ω_2, we have for each value of t_2 a cosine over t_1 of frequency ω_1 with a phase shift dependent on ω_2 and

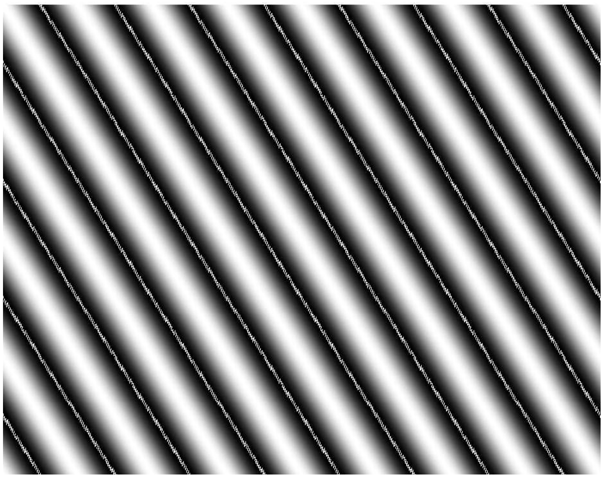

Figure 5: A single 2D-Fourier component

the current value of t_2. The 2D Fourier representation contains two numbers for each combination of frequencies, just as it did in 1D Fourier analysis for every frequency. These numbers still reflect amplitude A and phase ϕ of the cosine: the function with parameters becomes $A\cos(2\pi(\omega_1 t_1 + \omega_2 t_2) + \phi)$. The same information is given by two amplitudes A_1 and A_2, one for a cosine and one for a sine of the specified frequencies (as above, the combination of similarities with the cosine and sine components provides the phase).

Plotting the two input variables t_1 and t_2 as spatial x and y axes, such a 'building block' sheet consists of a stripe pattern (see figure 5). The value of the sheet for frequency-combination (ω_1, ω_2) with amplitudes A_1 and A_2 at time / position (x, y) is $A_1\cos(2\pi(\omega_1 x + \omega_2 y)) + A_2\sin(2\pi(\omega_1 x + \omega_2 y))$. This 2D-to-value mapping, or sheet of values, is the 2D analogy of the cosine 'basis functions' in 1D Fourier transformations. The original matrix is built up from the sum of all such stripe patterns.

Determining the parameters A_1, A_2 from the original data uses the same similarity measurements. For every combination of frequencies (ω_1, ω_2), a pattern can be compared to the $\cos(2\pi(\omega_1 x + \omega_2 y)$ and to the $\sin(2\pi(\omega_1 x + \omega_2 y)$ sheets. If the signal has the same direction and spatial frequency as the cosine sheet, it will be highly similar to that, and not at all similar to the sine sheet. Imagine the signal as vertical blocks rising various heights over a chessboard. The reference sheets consist of diagonal lines of blocks rising above and below the plane. Multiplying the signal and reference heights at each block results in a a new landscape. The new blocks will be strongly positive if both the signal and reference were either both positive or both negative, negative if their sign was different, and close to zero if either was close to zero. Adding up all the heights provides a number that reflects the similarity.

The spectrum is now a matrix showing just the overall amplitude per combination of frequencies. So a single entry in the matrix represents not just a time signal, like a single blip in a 1D spectrum, but a sheet of diagonal stripes (again, there's also some phase shift of the pattern which cannot be known from just the amplitude). A website I just came across mentioned that the stripe pattern of a given point in the spectrum matrix oscillates in the same direction as the direction to the point in the matrix from the centre of the matrix (where frequencies (0, 0), or the overall mean of the data, is stored). That seems to make sense: if the vertical frequency is twice as high as the horizontal, the stripes will be "stretched" vertically. Stripy patterns are an artefact that can occur in MRI images, since the whole pattern only needs a single error-blip in K-space (see below).

Again, we've been looking at a spatial representation of oscillations, which would be a suitable representation when e.g. the data are already present in Matlab. In measurement situations, perhaps a sensor array would be used, if one dimension is time and the other a spatial dimension. In the case of MRI, the rotation of the NMV is exploited and the signal consists of the components of the transverse NMV in the direction of two perpendicular sensor coils. This is explained further in the next section.

So K-space is just the 2D Fourier representation of a two-dimensional pattern. The MRI machine, due to the way gradients work, directly receives the values of elements of K-space. How this works is the subject of the next section, but that's about *acquiring* the values of K-space. What K-space *is* is simply a collections of stripy patterns of different frequency-combinations, of which the sum provides the spatial distribution of signal intensity over voxels.

8.2 Filling in the values of K-space

So now we know what K-space is. The clever thing about MRI is that the signals recieved by the machine *are* values in K-space. The machine is physically performing the 2D Fourier transformation. The way it does this is by imposing gradients in the magnetic field in combinations of directions. As I'll illustrate, this leads to sheets with waves in a given direction, with a spatial frequency that increases at every time step.

8.2.1 Spatial frequency at time steps, temporal frequency at points in space

To illustrate this phenomenon, I'll first use a 1D example (see figure 6). Take a line of length L on which lie points that have a value that oscillates over time. The points oscillate at frequencies running from $-\frac{1}{2}\frac{1}{\Delta t}$ Hz to $+\frac{1}{2}\frac{1}{\Delta t}$ Hz from left to right, for some period Δt. At $t = 0$ the vector is flat. Now take a step forwards to $t = \Delta t$. Now the end points will have completed half an oscillation: on the left, the point has gone down and back up to zero, on the right the point has gone up and then back down to zero. The middle point doesn't move; since its frequency is zero. The points at $\frac{1}{4}$ length from the end points are

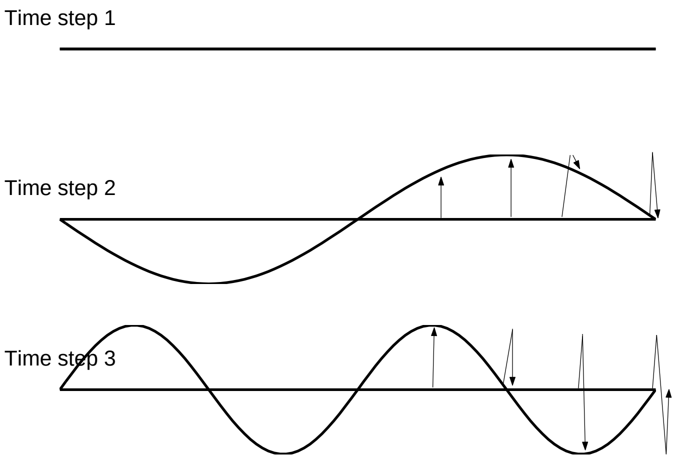

Figure 6: Frequency distribution to spatial pattern in 1D

at 90 degrees of their period, one at -1 the other at +1. The other points lie between these extremes, leading to a spatial wave of period L lying over the line. Now take another step forward, and see what happens. The endpoints and the middle point still end up or remain at zero. But now the quarter points are also at zero, and the one-eights points are at their extremes of + and -1. The frequency of the curve over the line is now at $\frac{2}{L}$. In general, after N steps in time, the points oscillating in time produce a spatial oscillation of frequency $\frac{N}{L}$.

So at every time step, a comparison of the spatial pattern over the line with another 1D distribution of values will extract similarities of the new, "input" curve with spatial oscillations of an increasing frequency in the "reference" curve. This kind of comparison involves two parts. First, take the product of the two vectors (i.e. the list of values over the line) and then take the area under the curve of their product (the areas of regions of values below and above zero get a negative and positive sign respectively). In mathematical terms, taking the area is *integration*: for ever smaller segments of line, get the area of the rectangle above or below that segment, using the mean value of the signal over the segment as height. The more similar the vectors, the higher the area (consider how positive values would then map to positive values, and negative to negative, leading to only highly positive values in the product and therefore a large area; similarly, perfectly opposite vectors would lead to a negative but still large area). Second, shift one of the vectors so it's 90 degrees out of phase with its original and repeat the similarity measurement. This gives us the two numbers we need for the Fourier representation for the frequency of the input curve. So if we just start at $t = 0$, let the points of the reference signal oscillate, and compare every $\frac{1}{\Delta t}$ seconds, we will, point by point, fill in the Fourier representation of the input signal.

In 2D, exactly the same can be done, except now there is a matrix of points

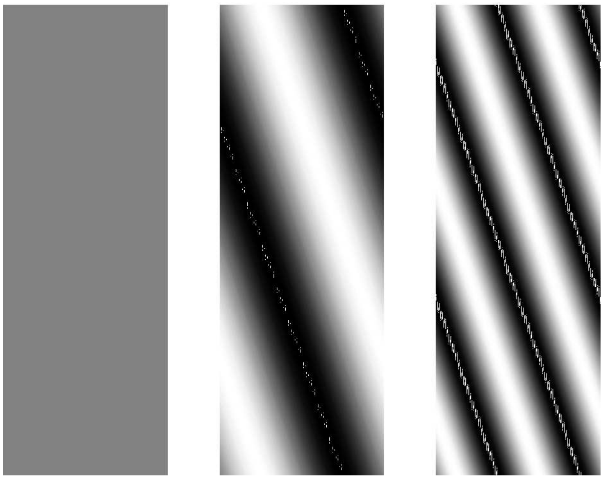

Figure 7: Frequency distribution to spatial pattern in 2D

and frequency increases from left to right and from top to bottom. This leads to the stripy sheets that are the building blocks of 2D Fourier representations. For every combination of frequencies for the left and right direction, the stripes run in the direction determined by the ratio of the frequencies. The spatial frequency in that direction increases, just as in the 1D case, over time steps (see figure 7).

8.2.2 Setting up spatial frequencies using gradients and filling in K-space

Now, how do these 2D Fourier calculations get performed in the MRI machine? The first step is slice selection: a gradient is imposed that linearly modulates the magnetic field strength in the Z direction (head-to-feet for instance, though not necessarily). Due to the gradient, the resonance frequency will be dependent on Z-position, and so an RF pulse of a given frequency band will only flip the magnetic moments of a specific slice of nuclei into the transverse plain. Within this slice, the Y-axis is subjected to a second gradient. The Y-axis is then said to be frequency encoded. The final, X axis, will effectively also get a frequency gradient, but using a different method called phase encoding. It goes like this. The strength of the Y and X gradients, and hence the direction of the stripy reference pattern that will be imposed on the slice, are chosen at $t = 0$. The MRI machine can only detect the summed signal from all the voxels in the slice, but it does so using perpendicular sensors in the transverse plane of the magnetic field. The 'stripy pattern' in this case specifies the precession phase at a given location, in other words, whether a sensor in a given direction will detect that locations signal. So, in terms of the chessboard landscape analogy of the previous section, the pattern provides the block-by-block multiplication *relative to a sensor that detects signals with phase 0 degrees as positive and phase 180 degrees as negative.* The summation is done simply by virtue of the sensory picking up the summed magnetization from all locations in the slice. The second, perpendicular sensor detects the similarity of the pattern with the signals at plus and minus 90 degree phase, i.e. the NMV the first sensor would measure after a ninety degree rotation. This is equivalent to saying that the second sensor receives the comparison of the same signal as the first, but with a phase-shifted reference pattern. Thus the two sensors provide the cosine and sine similarities of the pattern of signal intensity with the imposed patterns of the specified direction and spatial frequency.

Of course, at $t = 0$ the gradients haven't yet been turned on, so the reference pattern is flat, and the similarity measures simply provide the sum, and hence, after normalizing, the mean intensity value. This value gets stored at position (0, 0): both frequencies are zero. Now we take a time step. The frequency encoded gradient exerts its effects during this time, since its constantly effecting the frequencies along its axis. The phase-encoded direction in contrast only bumps the phase of nuclei along its axis between time point. Effectively, this results in a higher frequency over time points - a voxel undergoing double the phase bump every time step undergoes a doubling of the effect on its frequency - exactly the same as the frequency encoded axis. Now, at $t = \Delta t$, the spatial waves arise, and at all subsequent steps their frequency increases. At each time step, the similarity with the currently imposed pattern is directly measured by the sensors and then stored, filling in K-space point by point.

So for a given combination of X and Y gradient strengths, the reference pattern lies in a specific direction, and the spatial frequency of that pattern increases per time step. So the collected values of K-space get placed on a line radiating out from zero, at ever more peripheral positions as time goes on. The frequency combinations will be $(0, 0), (X\frac{1}{L_X}, Y\frac{1}{L_Y}), (2X\frac{1}{L_X}, 2Y\frac{1}{L_Y}), (3X\frac{1}{L_X}, 3Y\frac{1}{L_Y})$ and so forth, until the gradient strengths are changed, so the direction of the spatial pattern changes, and the line being filled veers off in a different direction in K-space. (Some literature suggests that it is rows of K-space that are filled, not lines radiating from zero in all directions. Filling in a row would require constant modulation of one the gradients however.)

Eventually, enough points of K-space get filled in for a sufficient reconstruction of the spatial image to be possible. At heart, these are matrix operations, turning one matrix into another matrix. The dimensions of the spatial matrix need to be stored separately; by itself K-space doesn't know how big the slice it's representing is. By spreading out the gradients over a larger region, more of the body can be imaged, at the cost of the steepness of the slope.