Orbital Angular Momentum (OAM)
An Antenna Engineer's Perspective
By Neill Tucker

After some very interesting e-mail exchanges with a student researching Orbital Angular Momentum (OAM) I decided to have a closer look myself to see what all the excitement is about.

Anyone who has worked with antennas may already be familiar with the momentum properties of EM radiation. This can be in the form of linear polarised EM waves which carry Linear Momentum (LM), or circularly polarised EM waves which carry Spin Angular Momentum (SAM). The polarisation properties of EM radiation are described in most antenna theory textbooks e.g. Balanis [1]. By contrast OAM, which refers to the rotating nature of the propagating wave-front itself rather than its EM components, is conspicuous by its absence.

The polarisation property of EM radiation is used to good effect in most of our modern wireless communication systems. Either, to mitigate the effects of multipath in non line-of-sight communication systems, such as mobile phones. Or, to allow the direct reuse of frequency allocations in line of sight applications, such as satellite tv. Whether it is the polarisation properties of the EM radiation itself or the modulation schemes applied to it, the goal is invariably the same, to increase the channel capacity for a given bandwidth.

Any property of EM radiation that can be applied in an orthogonal (one doesn't interfere with the other) manner is potentially useful in terms of making maximum use of available bandwidth. Antennas can transmit and receive two completely different streams of information on exactly the same frequency by using polarisation to discriminate between the two streams. This can be vertical / horizontal polarisation or left-hand / right-hand circular polarisation. It should be noted that there is only limited discrimination (3dB) between linear and circular polarisation. So, although there are 4 polarisation possibilities they can only be used as orthogonal pairs (VP/HP or LHCP/RHCP) therefore only doubling the channel capacity.

To date, it is really only the advanced digital modulation schemes that have enabled significant increases in capacity. Notice how many more tv channels there are available, now it's all gone digital. Whether the content of all the extra channels justifies the technical effort is another question. However, assuming our precious bandwidth would be filled absorbing and worthwhile material. Imagine multiplying the capacity gains afforded by modulation techniques by not two polarisation options, but by twenty, thirty, or more OAM options. Unlike the standard polarisation properties, OAM can theoretically exist in an almost infinite number of orthogonal modes. The huge potential increase in channel capacity is at the root of what all the excitement is about.

What is OAM and how do you transmit and receive it?

It has long been known that EM radiation can carry momentum. Photons of light exert a small but measurable force when they strike an object. EM radiation of much lower radio frequencies can also exist as photons, and have the same effect, except that it is much smaller. The energy of a photon is given by $E = hf$ where E is the photon energy in joules, h is Plank's constant and f is frequency in Hertz. Photons with OAM also exert a measurable force except this force is rotational about the photon beam axis. An entry in Wikipedia [2] gives a very good explanation of the angular momentum properties of light and therefore EM radiation in general.

However, for communications applications we are more interested in how the EM radiation interacts with charged particles (electrons) that are present within the conductor of the antenna. It is these that are swept back and forth by the oscillating EM field and either enter a Low Noise Amplifier (LNA) in the receive case, or lose energy to free space in the transmit case.

This interaction between EM waves and electrons underpins all of our existing radio architectures and therefore influences how we 'look' at potential improvements. A possible problem with using OAM modes as some form of discrimination between channels using the same transmit frequency, is how the OAM manifests itself in the EM radiation. A few examples should help to explain things.

Using some simple phased array software called ArrayCalc [3]; a circular array of rectangular patches has been modelled (figure1). The array is operates at 2.4GHz and all patches are arranged with their E-field components vertically orientated in line with the z-axis. The patches are arranged such that their centres lie on a circle in the y-z plane of radius 0.55*lambda (68.8mm at 2.4GHz).

The model was run for 2 different excitation modes :

Mode#0 = All patches have equal phase and amplitude.
Mode#1 = All patches have equal amplitude but are phased in 45deg increments around the circle

Results for Mode#0 Excitation

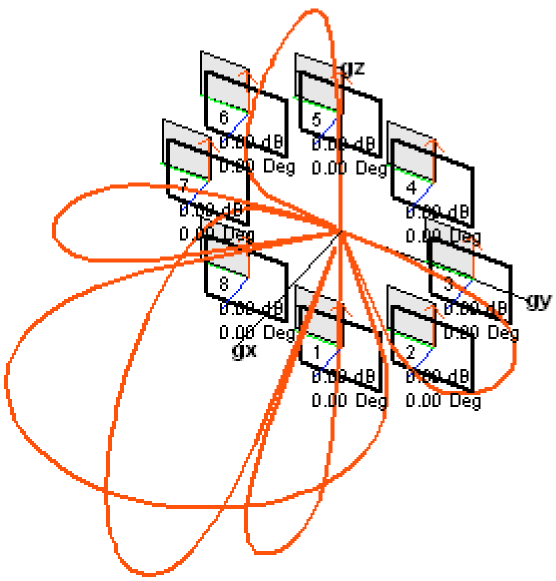

Figure 1 Circular array of rectangular patches

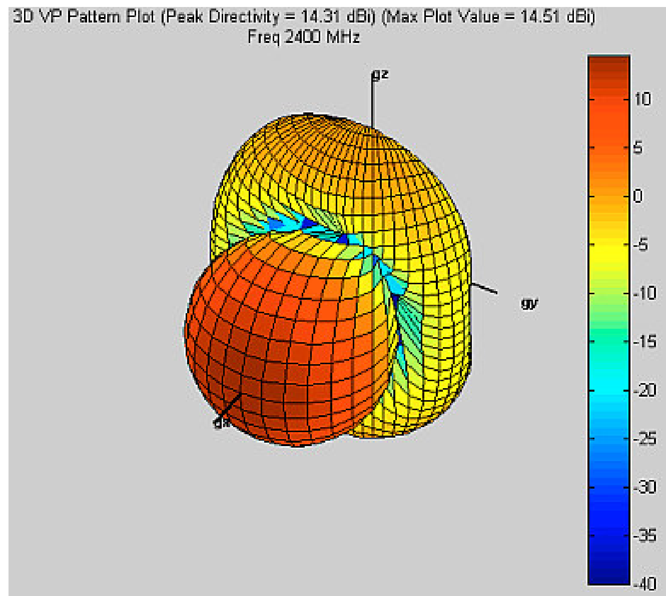

Figure 2 3D Pattern (Vertically polarised component)

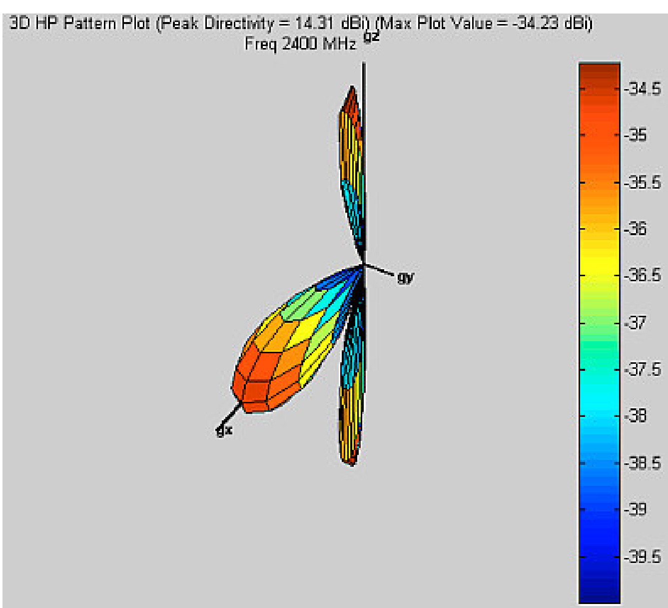

Figure 3   3D Pattern (Horizontally polarised component)

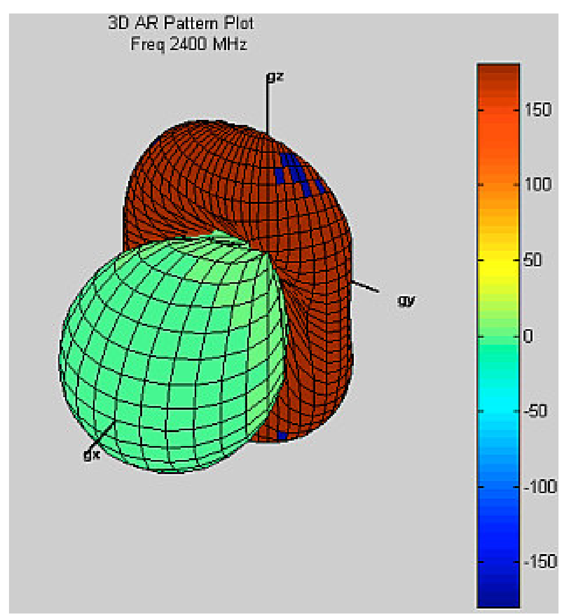

Figure 4 3D Plot of phase (overlaid on total field pattern)

The mode#0 excitation produces a very typical antenna pattern in that most of the power is concentrated in a single direction, the x-axis in this case (figure2). The radiated power is also limited for the most part to a single polarisation, vertical in this case. There is very little horizontally polarised or 'cross-polar' radiation (figure3). In this model there would have been zero cross-polar component except I deliberately set the patches to be 0.2 deg off vertical. In reality patches will always radiated some level of cross-polar, the deliberate error was to introduce some horizontal polarisation component, which is what you would see in practice.

In ArrayCalc v2.5 it is possible to plot the phase of the radiated field overlaid on a 3D pattern plot of the total field intensity (figure 4). Notice that the phase front is nominally flat over the entire main lobe of the pattern. Also that there is a 180deg phase reversal in the side lobe.

Results for Mode#1 Excitation

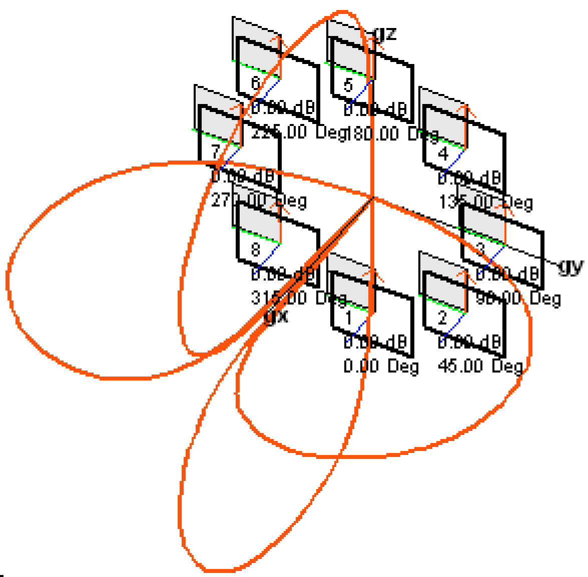

Figure 5 Circular array of rectangular patches

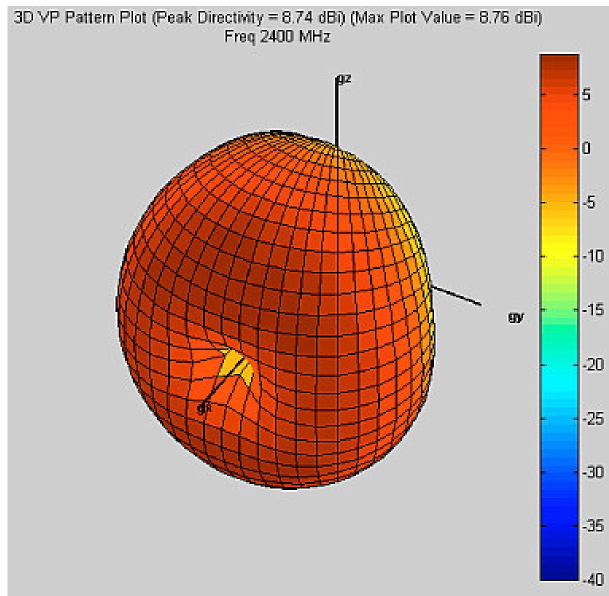

Figure 6 3D Pattern (Vertically polarised component)

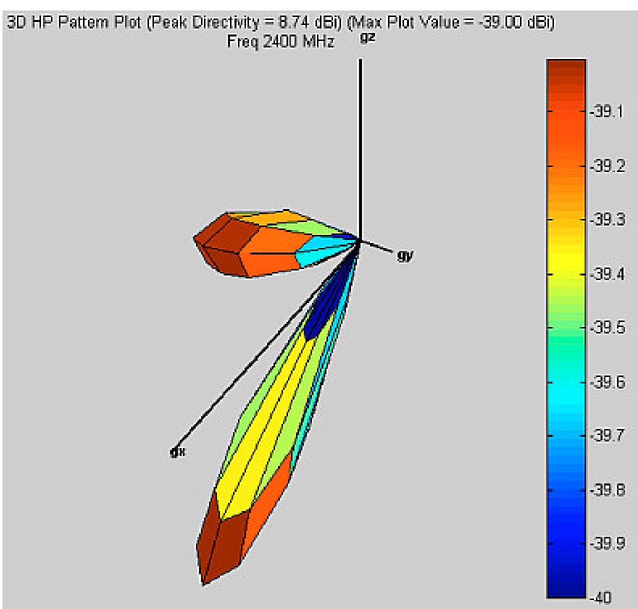

Figure 7 3D Pattern (Horizontally polarised component)

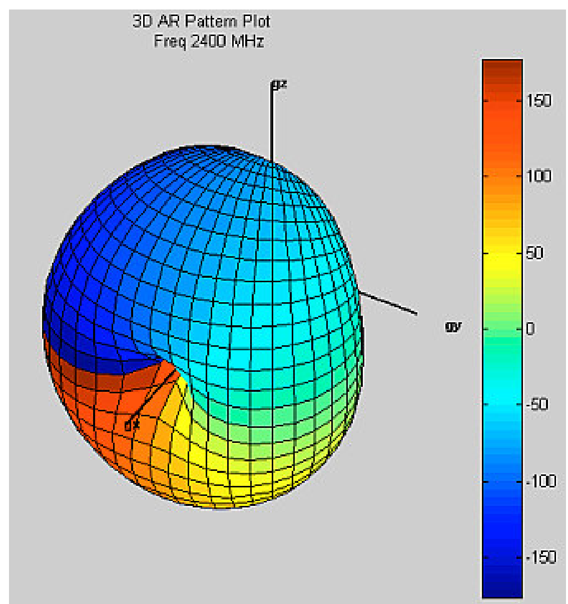

Figure 8 3D Plot of phase (overlaid on total field pattern)

Exciting the array in mode#1 (figure5) produces a very different field pattern (figure6). Most significantly, there is now a null on the bore-sight of the antenna, where there was previously a maximum. The radiated field is still vertically polarised (compare figs 6 & 7) but is now distributed in annular beam around the x-axis. Most importantly however, the phase front of the radiated field is no longer flat, but is rotating about the axis of propagation (figure 8) and is one of the principal characteristics of OAM. By using multiple rings of elements with higher order mode excitations, multiple modes can be combined into the same beam, see later.

Practical application (long range)

Assuming we limit discussion to the main lobe of antenna patterns (nulls and sidelobes complicate things due to rapid phase reversals). We can see how these modes would compare in a practical line-of-sight application such as satellite tv.

<u>Standard polarisation discrimination</u>

In the UK the majority of satellite tv is provided by the Astra2 satellite which has downlink frequencies between 10 and 12 Ghz. By using both vertical and horizontal polarisation (VP and HP), the channel capacity is effectively doubled for the available bandwidth. Antennas that radiate a particular type of polarisation, do so in the entire main lobe. In other words, it doesn't matter which section of the main lobe pattern you intersect with your receiving antenna; it is all 'encoded' with the same polarisation attributes such as VP or HP. This is relevant, because even in line-of-sight applications such as satellite tv, the receive antenna only intercepts a small portion of the transmit antenna coverage. The Astra2 satellite covers the whole of the UK but a receive dish is only around 0.3m^2. Despite this, the VP/HP field probes in the Low Noise Block Down Converter (LNB feed head) are still able to discriminate between the two polarisations.

<u>OAM mode discrimination</u>

One possible problem with using OAM, is that it is a characteristic of the main lobe as a whole. If you were only able to intersect a small portion of the main lobe pattern, the only identifying attribute would be a small phase gradient across the aperture of the receiving antenna. I would say it is debatable whether you could identify the EM radiation as having a particular OAM mode attributed to it. The schematics in figures 9 and 10 illustrate how I see the difference between standard polarisation and OAM. In the first case you can read 'Vertical' in whatever direction you look towards the transmitting antenna. In the second case you only can only see the 'd' part of the description 'OAM mode 1'. Not only that, the proportion of the word that is visible depends on the distance from the source, so you may not see even a complete letter if you were further away.

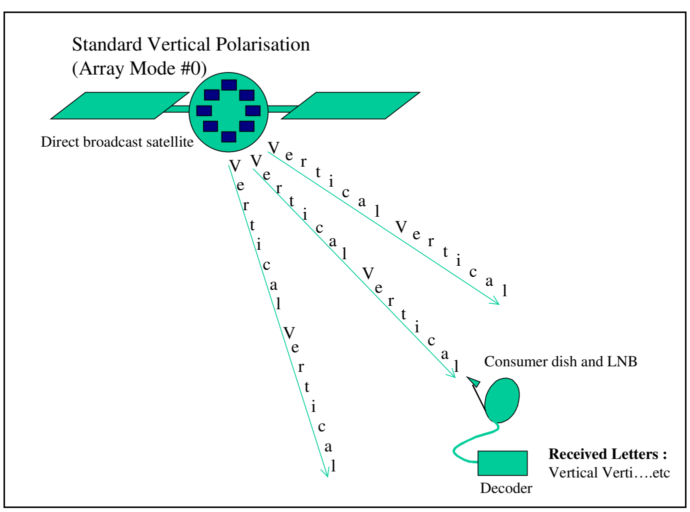

Figure 9 Standard VP Propagation (Mode#0)

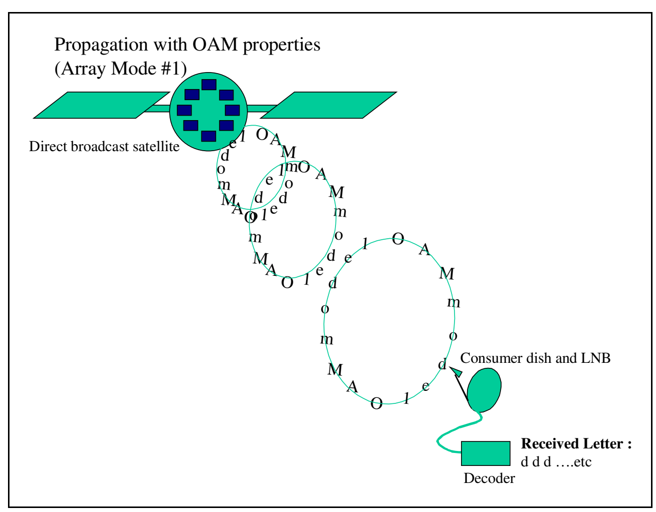

Figure 10 OAM VP Propagation (Mode#1)

To view this in terms of antenna patterns, we can go back and take a closer look at figure 8. Here we see that the phase information has been plotted on a grid. Each of these grid squares represents 5deg Azimuth by 5deg Elevation, measured from the centre of the antenna (axis origin). If we re-calculate at 1deg resolution and zoom in on a 5deg by 5deg section, we can get an idea of what a receive antenna aperture (subtending a one grid-square solid angle) might ‘see’. The square in figure 11 is taken from an area of the plot (figure 8) where the radiated field is at a maximum. Over this 5 by 5deg square there are 3 changes in colour which corresponds to around 17deg of phase change across the aperture. For more precise analysis of this see the exoam2.m example in ArrayCalc v2.5.

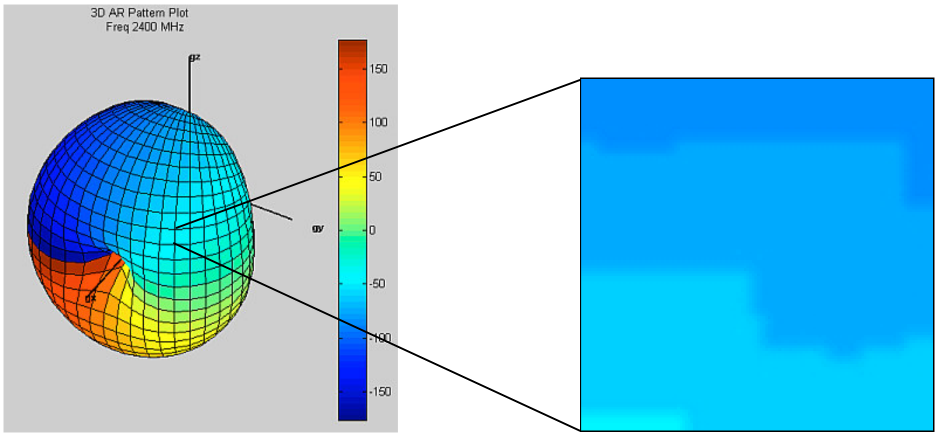

Figure 11   Zoomed 5 by 5 deg grid square from fig 8 (1deg AZ/EL plot resolution)

Bear in mind though that this 5deg square corresponds to (5*5)/(180*180)*100=0.077% of the whole half hemisphere of radiation, which is a lot. If the main lobe of a Direct Broadcast Satellite (DBS) antenna covered the UK with a 1000km diameter spot beam, a 60cm receiving dish would intercept only (pi*0.3^2)/(pi*500000^2)*100 = 3.6e-11% of it and the phase shift across it would not be measurable!

From the previous section we have seen that OAM manifests itself in 2 ways. Firstly, radiated power is concentrated in an annular beam around a central null. Second, the phase front of the propagating wave rotates about the axis of propagation. In order to decode the OAM mode, a significant proportion of the radiated field would have to be captured by the receive antenna. To do this, both transmit and receive antennas would have to be electrically large, to generate narrow transmit beams and have large collecting apertures respectively. At extremely high, near-light frequencies, electrically large antennas (apertures of many wavelengths) can still be physically small and practical. However, at such high frequencies, there is significantly greater bandwidth available so reducing the need for OAM in the first place.

While the satellite tv example is somewhat extreme in terms of distance, it does illustrate another important point. Existing line-of-sight communication systems that try to incorporate OAM into their architecture will face not only the usual signal strength (link budget) issues, but the ability to decode the OAM modes would also challenged by increasing link distance.

Practical application (short range)

The greatest need for increased channel capacity is in the lower frequency bands used for mobile communications and WiFi coverage, which are much shorter range compared to satellite tv.

To see how OAM might fit in with short range communications systems we need to take a look at the propagation of the OAM wave-fronts themselves. The diagram in figure 12 shows how waves from a standard antenna, such as the circular array operating in mode#0, would look. Although the wave leaves the antenna as a spherical wave, viewing a small section of it at a distance, it approximates to a flat wavefront known as a plane wave. The important thing to note is that the plane of the wave (or phase front) is normal to the direction of propagation, as shown in figure 12.

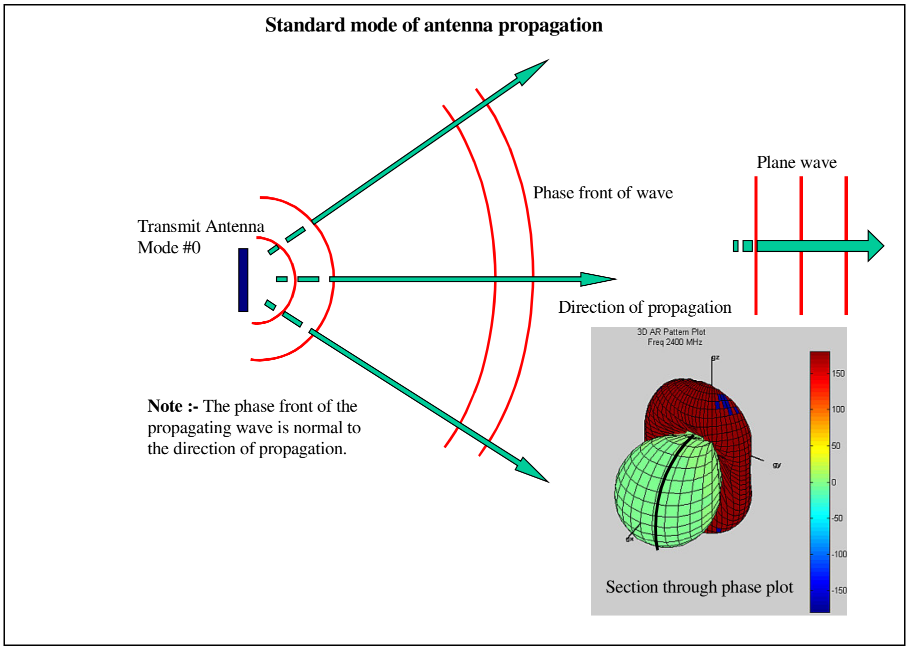

Figure 12 Phase fronts for standard mode #0 propagation

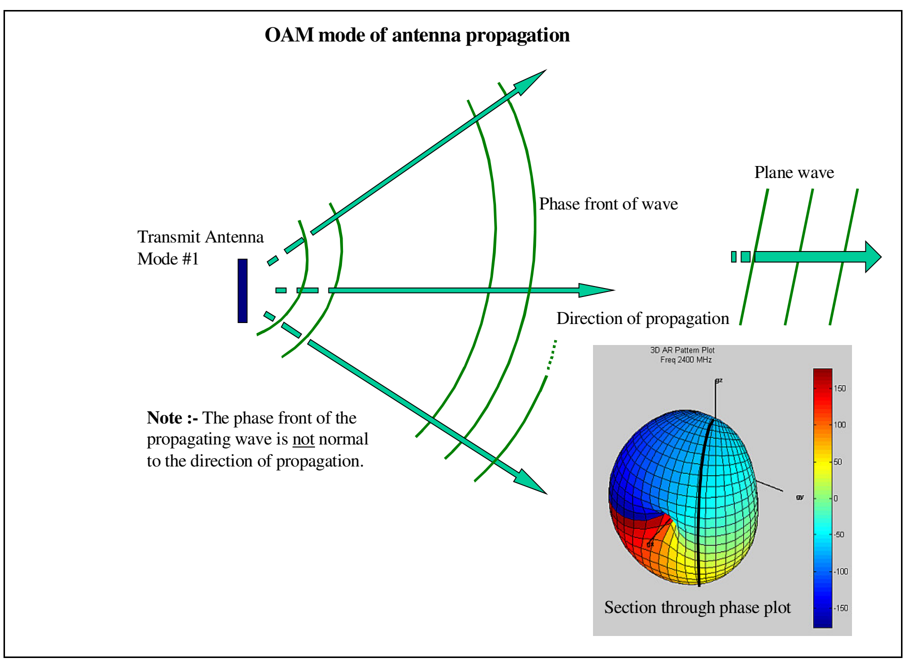

Figure 13 Phase fronts for OAM mode #1 propagation

By contrast, when the circular array operates in mode#1, producing OAM, something different happens. The plane waves have a phase gradient across them, which means that the phase front of the wave is no longer normal to the direction of propagation, as shown figure 13. However, it should be noted the phase gradient per unit distance would reduce with increasing distance from the transmitter, so as previously mentioned, range will be an issue.

How could the phase gradient be useful?

If we look at how a phased array steers a transmit beam, we can get some insight into how signals with OAM properties might be used to advantage. To transmit in a direction normal to the plane of the array (the mechanical boresight), all elements are equi-phased. To steer the beam off mechanical boresight a phase gradient is applied to the array. Since antennas are reciprocal, similar logic can be applied to the receive beam, with interesting results.

If an array is equi-phased and orientated towards the source of the plane waves it will receive standard mode#0 waves according to the radiation pattern shown in red in figure 14. However, mode#1 (OAM) waves would be received according the green pattern. In fact, if the phase gradient and aperture size were just right, there might be a null in the direction of the transmitter, and almost no signal would be received.

If the array is now phased to steer the beam off mechanical boresight, the converse will happen. The mode#0 waves will now be received according to the green pattern and the mode#1 waves like the red pattern. In other words, by suitable phasing of the receive array it may be possible to discriminate between modes

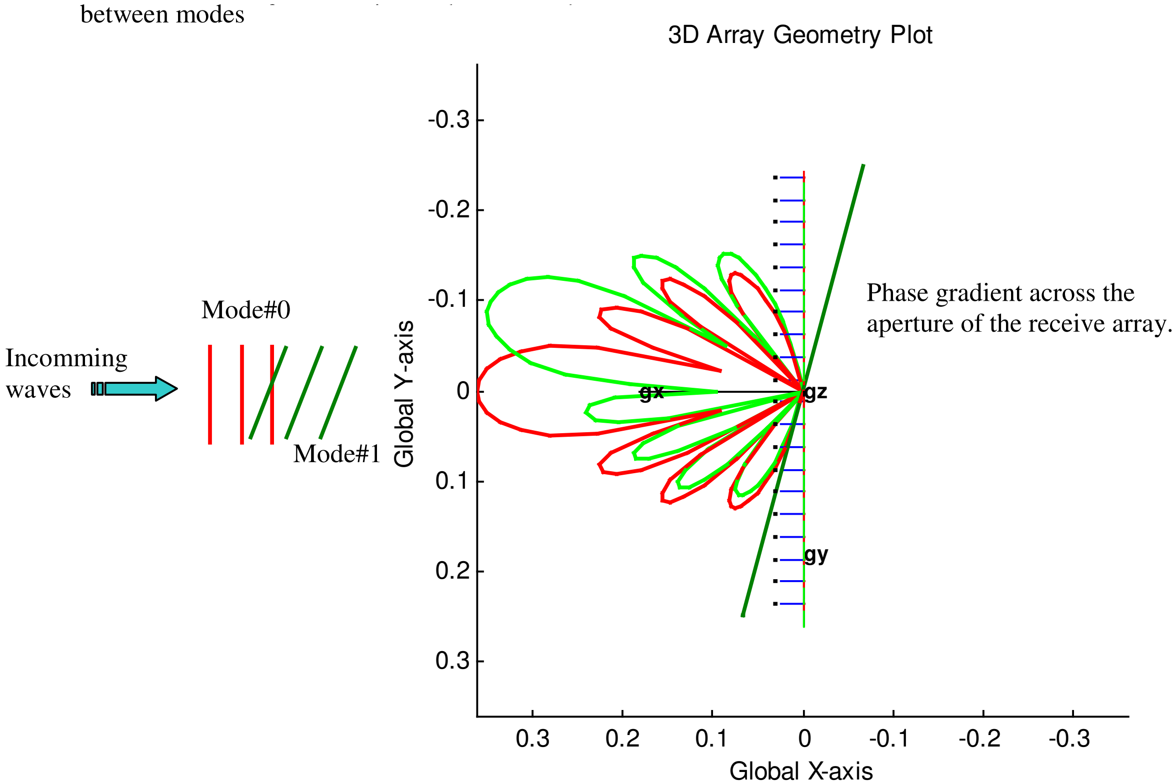

Figure 14 Receive antenna patterns

In non-line-of sight situations that experience multi-path fading, Multiple Input Multiple Output (MIMO) is often used to increase channel capacity. The MIMO system uses multiple transmit and receive antennas in different physical locations to take advantage of the multiple signal paths, increasing the channel capacity. The phase gradient in OAM signals gives the impression to a receiving antenna that the signal is coming from another direction than the actual physical location of the source. This quality could be used to increase the diversity of signal paths, there by increasing channel capacity in a MIMO system.

Conclusion

From my initial look into OAM, I would say that it certainly offers some interesting possibilities for increasing wireless channel capacity, and warrants serious investigation. Whether existing radio architectures can make full use of the OAM characteristics is another question? At present I would say application would be limited to short range MIMO systems and possibly high gain, line-of-sight links. The critical difference from the standard HP/VP and LHCP/RHCP polarisation schemes is that the OAM characteristics appear to be range limited when applied to existing architectures. To make full use of OAM properties some lateral thinking will probably be required, but then that's what makes it interesting.

References :

[1] 'Antenna Theory Analysis and Design' 2nd edition by Constantine A. Balanis. Published Wiley ISBN 0-471-59268-4

[2] 'Angular momentum of light' http://en.wikipedia.org/wiki/Angular_momentum_of_light

[3] ArrayCalc v2.5 MathWorks File Exchange

[4] 'Antenna diversity' http://en.wikipedia.org/wiki/Antenna_diversity

[5] 'MIMO' http://en.wikipedia.org/wiki/MIMO