Imaging

What we have:

What we want:

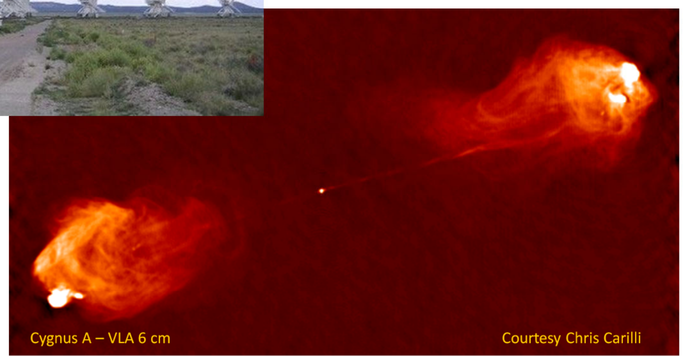

But what is an image? Just a list of location/brightness pairs.

Some approximations I'm going to make:

- I'm going to stick to 1 dimension where possible
- I'm going to ignore imaginary numbers where possible
- I'm going to ignore polarisation
- I'm going to assume perfect calibration
- I'll show very little real data
- No software packages
- As non-mathematical as possible
- (probably more)…

Some impediments to understanding interferometric imaging:

- *Jargon:*
  - Visibilities
  - Baseline
  - Correlation
  - UV plane
  - Dirty beam
  - Dirty image
  - Clean components
  - etc…
- *The Fourier transform.*

I M Stewart, 2012 GLOW School, Bielefeld 3

We'll encounter most of the jargon words during the course of this talk. I won't give a list of definitions. However, before beginning to talk about interferometric imaging, I want to talk in some detail about the Fourier transform.

The Fourier transform.

$$\mathfrak{F}\{g\} = G(\omega) = \int_{-\infty}^{\infty} dt\, g(t) \exp(-i\omega t)$$

- G in general is complex-valued.
- ω is an angular frequency (units: radians per unit t).
- the transform is almost self-inverse:

$$\mathfrak{F}^{-1}\{G\} = g(t) = \int_{-\infty}^{\infty} d\omega\, G(\omega) \exp(i\omega t)$$

- the transform is linear.

I M Stewart, 2012 GLOW School, Bielefeld 4

I think this is the only equation in my talk. I have included it only as a talking point – I won't be doing anything with it. In fact the good news here is that it is entirely possible to build a respectable career as a radio interferometrist without ever needing to evaluate this integral in person.

In the following slides, I want to concentrate on some characteristic features of the Fourier transform (FT). It is useful to become familiar with these features because one finds they appear again and again in interferometric imaging.

There are two points on the present slide I want to emphasize: 'almost self-inverse' means that for purposes of becoming familiar in a qualitative sense with features of the FT, it doesn't matter whether we are talking about the FT of the reverse transform.

The second point is the 'linearity' which just means that the FT of several sources is just the sum of the FTs of the sources taken separately.

The Fourier transform.

It looks ugly - complicated function goes to different complicated function. Complex numbers – ugh.

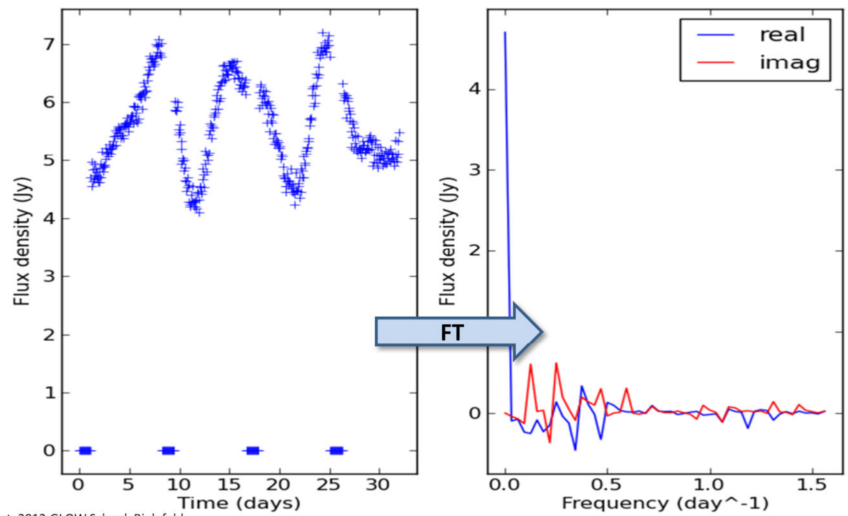

I M Stewart, 2012 GLOW School, Bielefeld 5

This is a typical (non-interferometric) application of the FT. It does not aid us in understanding or enjoying the transform.

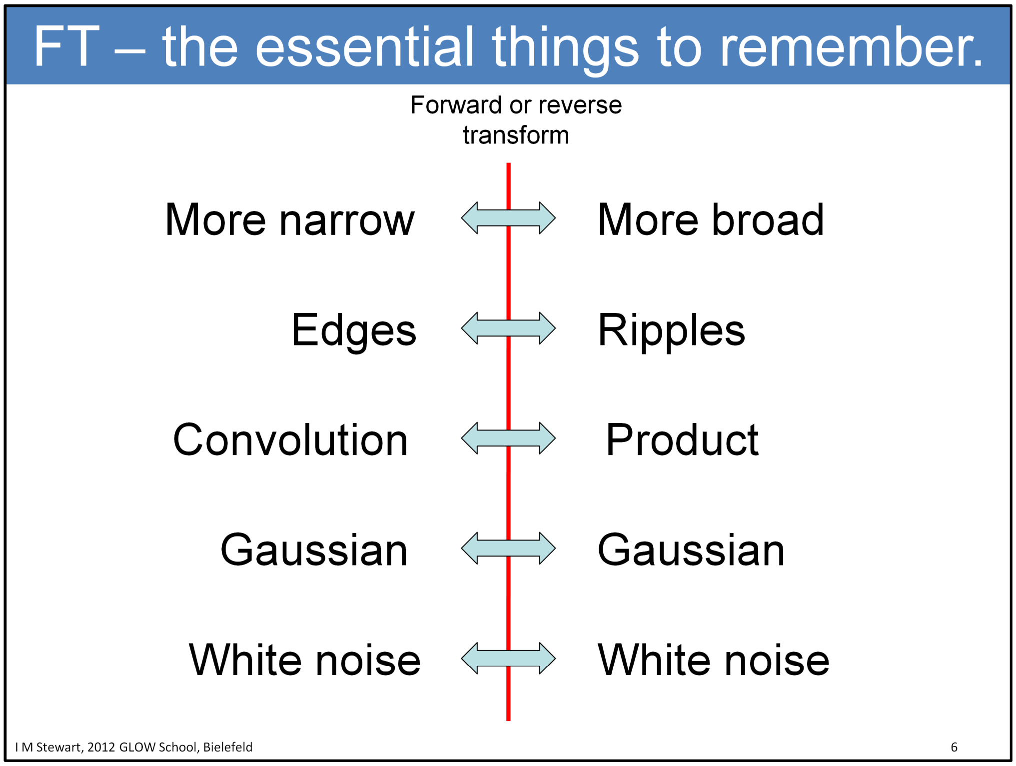

However, there are only a few essential points to remember – we can mostly forget the ugly details. I have listed a few of these essentials here and I will illustrate them with examples in the following slides.

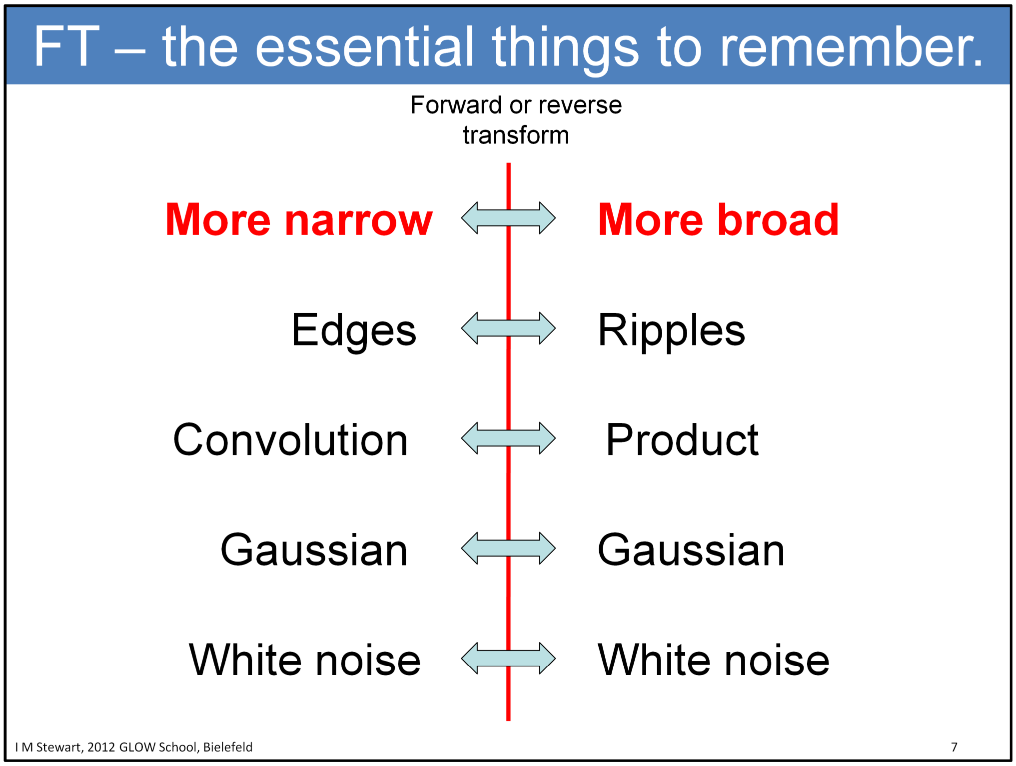

A signal which has a narrow characteristic scale on one side of the FT is expected to have a broad scale on the other (and vice versa).

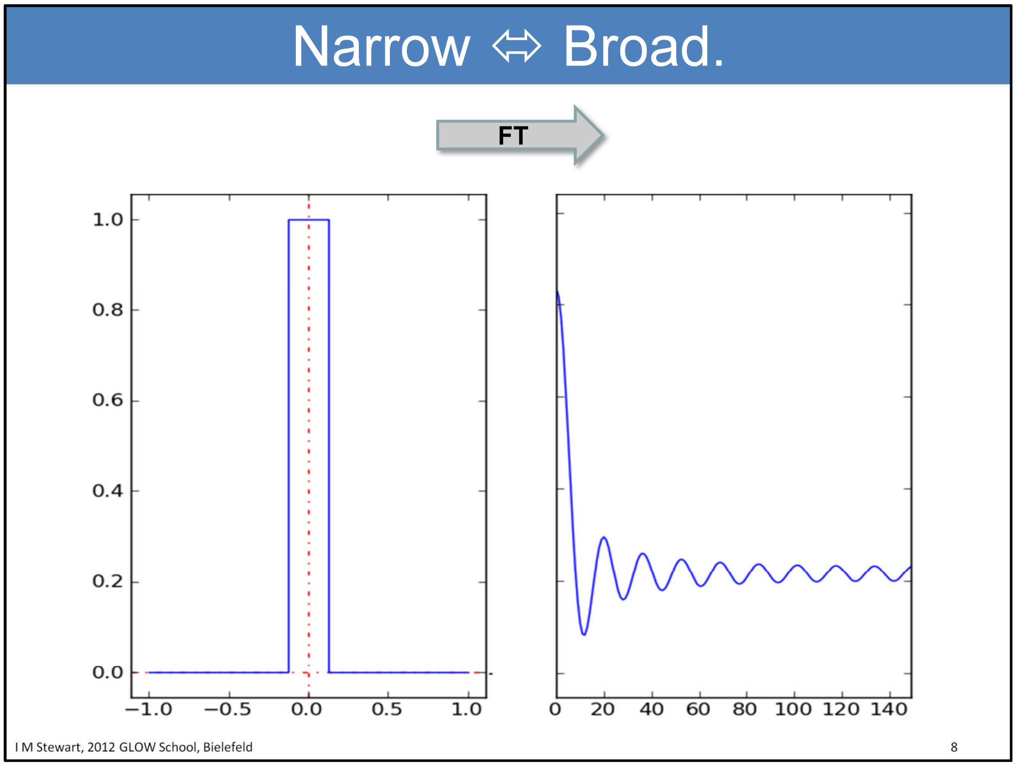

In this animation, I show a ‘top hat’ function on the left hand side (LHS) and its transform on the right hand side (RHS). The function on the RHS is called a sinc function and I probably should have included this top-hat vs sinc pair as one of the frequently-encountered pairs of functions related by the FT. As you can see, as the top hat becomes narrow, the breadth of the waves on the RH function become broad – and vice versa.

This pair also illustrate the next feature of the FT: namely edges ⇔ ripples.

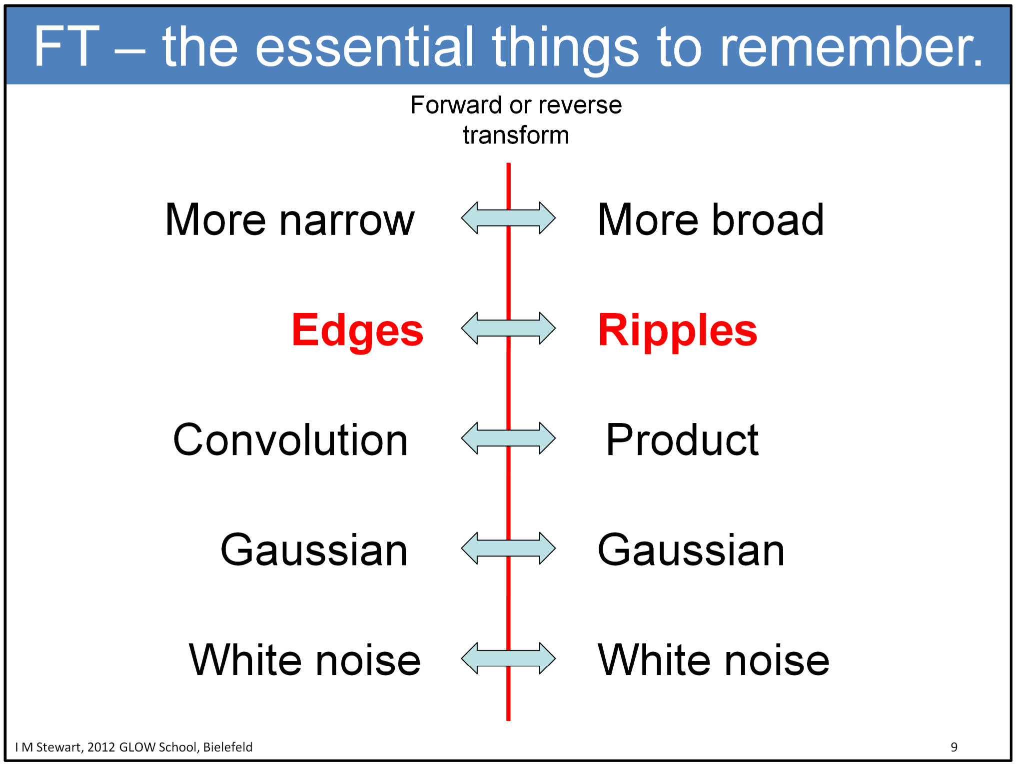

We expect a function which has some sharp edges (discontinuities) to have a FT which has a lot of ripples.

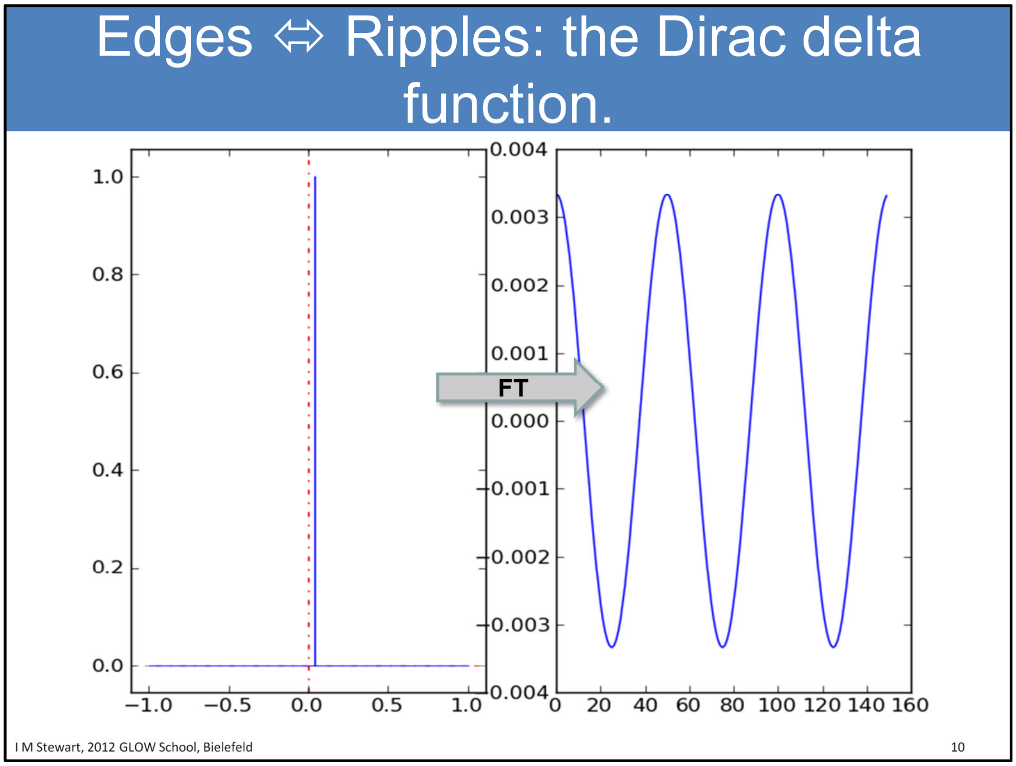

Something which has a non-zero value only at an exact particular location can be represented mathematically by the Dirac delta function. This has about as many edges as one can reasonably pack into a function (mathematicians are capable of being unreasonable and can add many more serious discontinuities!) and so we might expect its FT to be very ripply. In fact the FT of a delta function is a sinusoid which extends to infinity.

Note that I have placed the delta function quite close to the origin. In the principle of narrow ⇔ broad, this gives us moderately widely spaced ripples in the transform. But if we move the delta function further from the origin:

Edges ⇔ Ripples: the Dirac delta function.

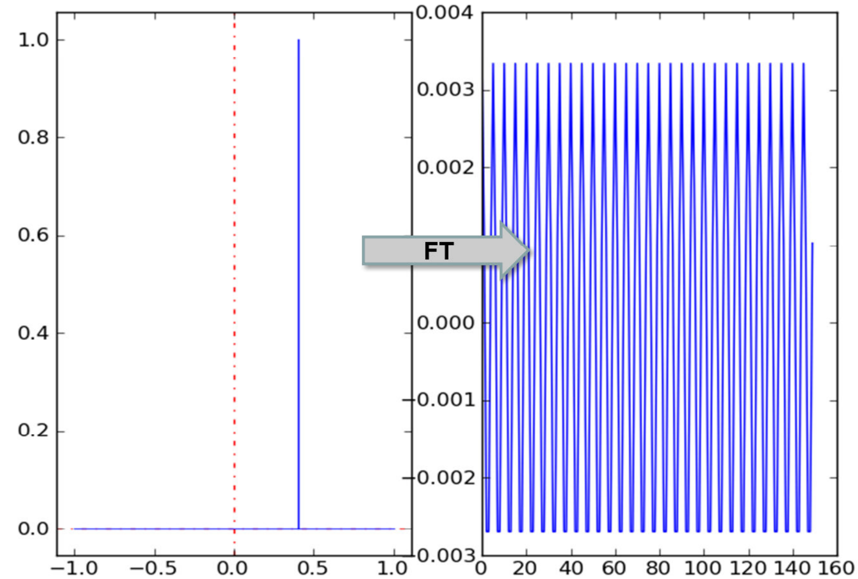

I M Stewart, 2012 GLOW School, Bielefeld 11

…broader ⇔ narrower, thus the ripples become much more closely spaced.

For a delta function exactly at the origin, the breadth of the ripples in its transform becomes infinite and the transform just becomes a flat line.

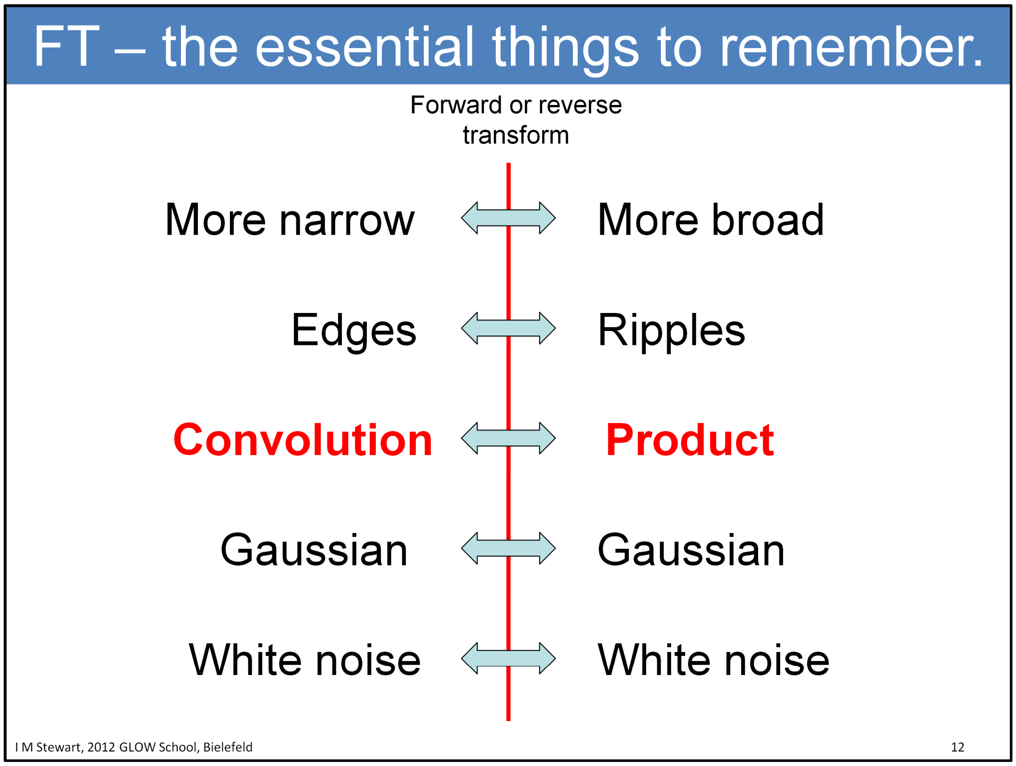

But just what is convolution?

(Note that convolution and correlation are closely related mathematical operations, and for input functions of suitable symmetry, they will yield the same result.)

What exactly does convolution mean?

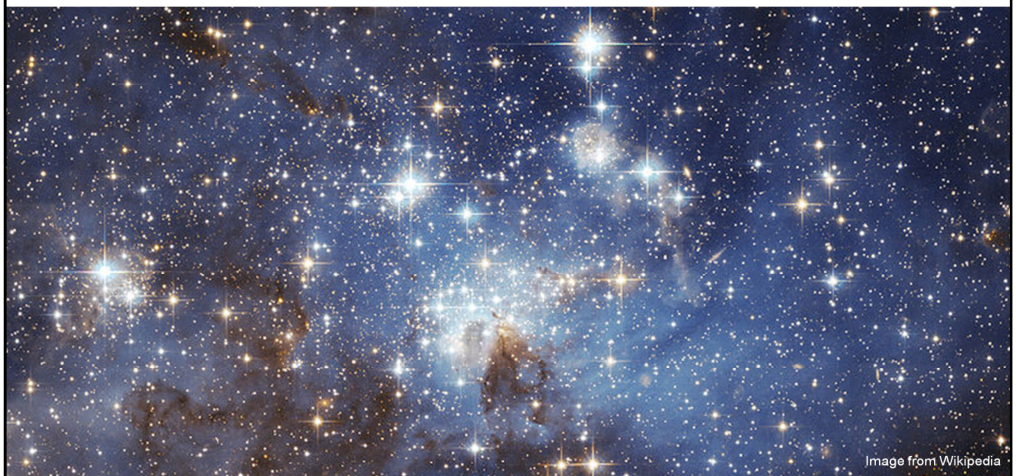

I M Stewart, 2012 GLOW School, Bielefeld 13

This is an example which may be more familiar to optical astronomers. Every star image here is surrounded by spikes which we know (since the stars themselves don't have spikes!) is created by the optics of the telescope, specifically here probably by the 4 legs of the support spider for the secondary mirror. Even the faint stars are surrounded by this pattern, although it becomes too faint to be obvious. The pattern, known to optical astronomers as the point spread function (PSF), is the same around every star. Mathematically we can represent this image as a convolution between the positions of the stars and the PSF.

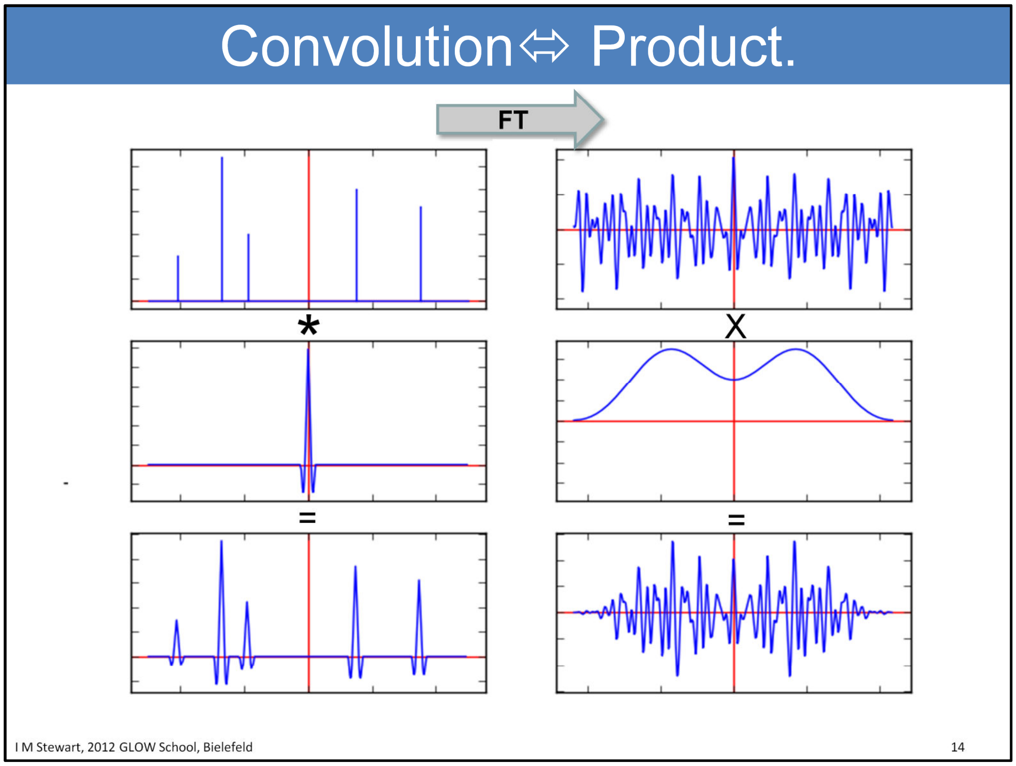

Here we have a 1-dimensional example. On the LHS I have shown some functions and on the RHS their respective FTs. The top left image is a set of 5 delta functions, which can be thought of as an image of 1-dimensional stars in some flatland version of the universe. Its FT at upper right now looks moderately complicated since it is a sum of 5 sinusoids of different size and pitch.

On the 2nd row, on the LHS I have generated a function which is somewhat plausible as an instrumental PSF. It is a narrow function, so we are not surprised that its FT is rather broad.

On the last row, on the LHS I show the convolution between the 'source image' at upper left and the 'PSF' at middle left. The PSF pattern is repeated with appropriate position and height for each of the sources. (The * symbol is usually used to represent the operation of convolution.) The bottom right picture shows the FT of this convolution. It should be apparent that it is the product of the FT at top right and middle right.

FT – the essential things to remember.

Forward or reverse transform

More narrow	⇔	More broad
Edges	⇔	Ripples
Convolution	⇔	Product
Gaussian	⇔	**Gaussian**
White noise	⇔	White noise

Gaussian ⇔ Gaussian.

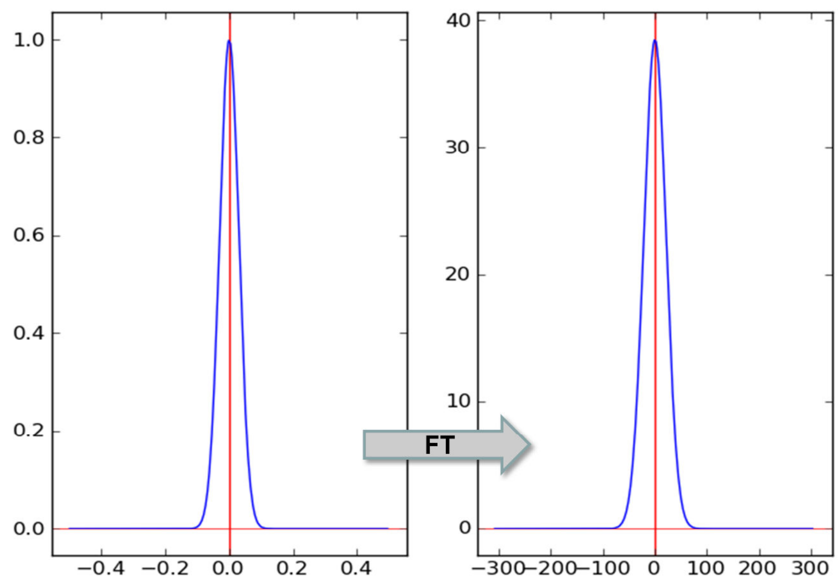

I M Stewart, 2012 GLOW School, Bielefeld 16

This really is a function/FT pair. I did not cheat and use the same function for both.

FT – the essential things to remember.

Forward or reverse transform

More narrow	⇔	More broad
Edges	⇔	Ripples
Convolution	⇔	Product
Gaussian	⇔	Gaussian
White noise	⇔	**White noise**

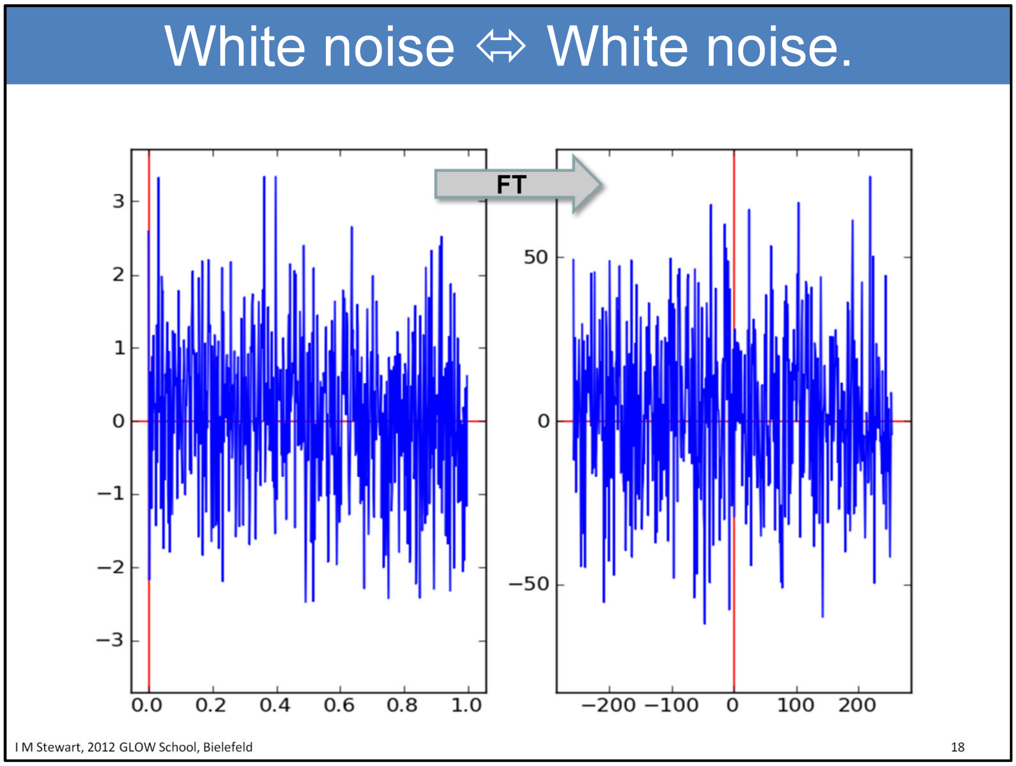

I won't go into mathematical details (to talk really sensibly about noise one has to calculate what is called a power spectrum); the point I want to get across is that the left and right hand sides, which are a FT pair, 'look' about the same.

'White' refers to the fact the FT has about the same amplitude across the breadth of the plot. Since the FT measures the amount of different frequencies (in the mathematical sense of numbers of sinusoid oscillations which fit into the range of the LH plot), the word 'white' is an analogy with white light, which is light which contains about the same amount of power across a range of electromagnetic frequencies.

It is useful to become familiar with the 'look' of white noise. As a help to do this I include some examples in this talk of non-white noise – the first of these on the following slide.

A diversion: red, brown or 1/*f* noise

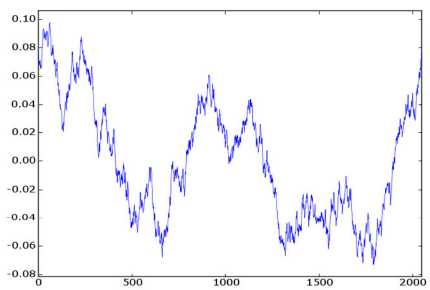

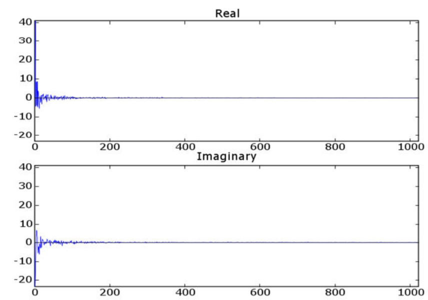

It's fractal – looks the same at all length scales.

If we generate some noise in the Fourier domain (plot on the RHS), but make the amplitude decrease inversely with the mathematical frequency f, then back-transform, the noise (plot on LHS) looks different. These values are now correlated – that means that they are no longer statistically independent.

$1/f$ noise has all kinds of interesting properties. Many natural processes generate this kind of noise.

This is not really relevant to interferometry but it is just kind of interesting.

Nature…?

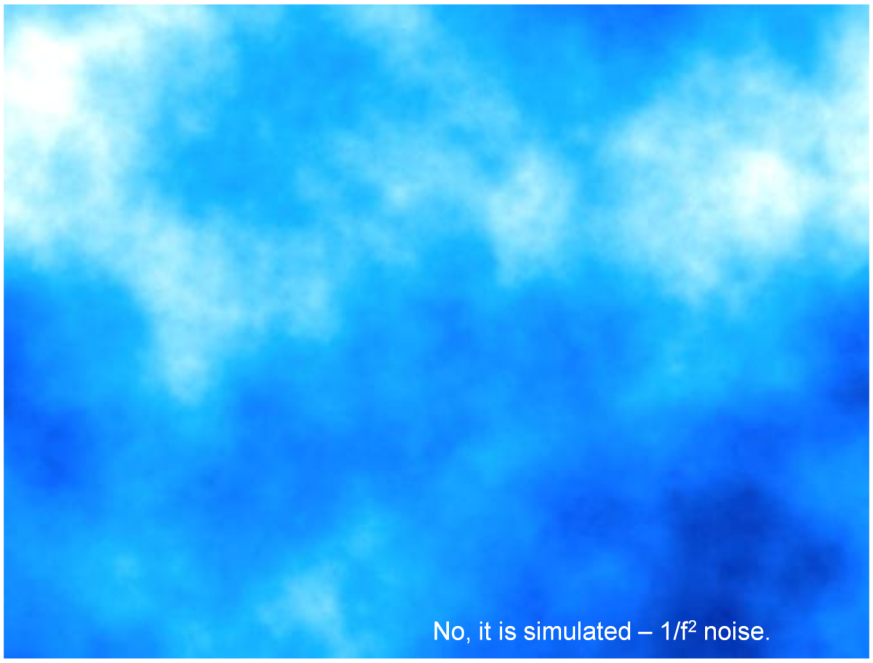

Waves from a source at zenith: phase difference $\varphi=0$

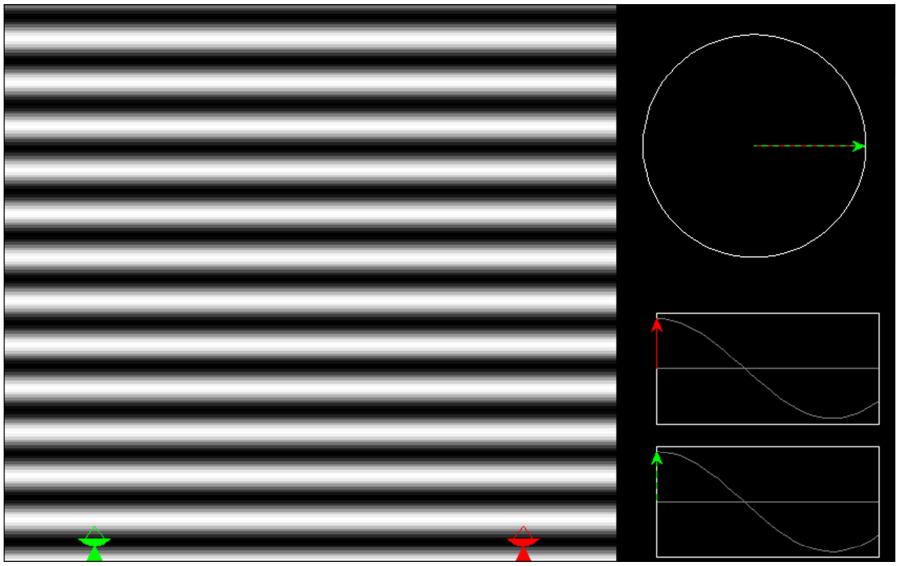

I M Stewart, 2012 GLOW School, Bielefeld 21

Back to work now. I have finished talking about the Fourier transform. Now I want to start talking about interferometry. I won't get to imaging yet though; first I want to quickly recap some of the basics.

In this animation I show in schematic form the radiation impinging on a basic 2-element interferometer from a source at the zenith. The bright and dark bands represent the wavefronts (which are plane waves to excellent approximation for sources at astrophysical distances), at a single frequency, coming from the pointlike source.

A very usual form on antenna is a dipole, i.e. just 2 bits of metal lined up in a plane roughly normal to the direction of the radiation. As the electric waves move past the dipole, the changing field makes the free electrons in the metal slosh back and forth. This in turn generates a voltage between the two dipole elements. The lower-right plots can be though of as 2 chart recorders which display this voltage as it changes with time.

The upper right plot, the circle, is a bit less intuitive. It is in fact an Argand diagram which shows a complex-valued vector – actually two vectors, red and green for the respective antennas. These two vectors are aligned in the present plot. The real parts of these vectors (their horizontal Cartesian dimensions) are matched to the respective voltage at lower right. The imaginary parts have no physical counterpart but we use this complex-valued representation become matters look a lot simpler this way.

The important point to note in this animation is that, because the source lies in a direction normal to the line connecting the antennas, its radiation reaches both at the same time. This is reflected in the fact that the voltage traces line up, and also that the red and green Argand vectors are aligned – one on top of the other.

Waves from an offset source:
phase difference $\varphi>0$

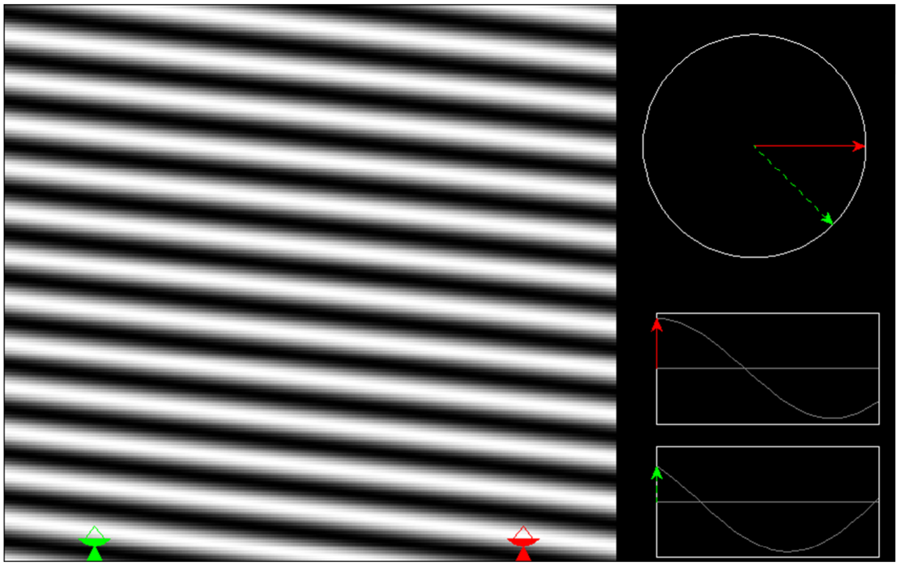

I M Stewart, 2012 GLOW School, Bielefeld 22

Here everything is the same, except I have moved the source slightly away from the zenith. Its radiation now reaches the red antenna before the green. The red voltage plot is ahead and the red vector leads the green one.

In fact the red vector leads by not just about 40 degrees but by 40+360 degrees, i.e. greater than 1 full turn. Notice how much easier it is to see and measure the difference in the Argand diagram than in the voltage plots.

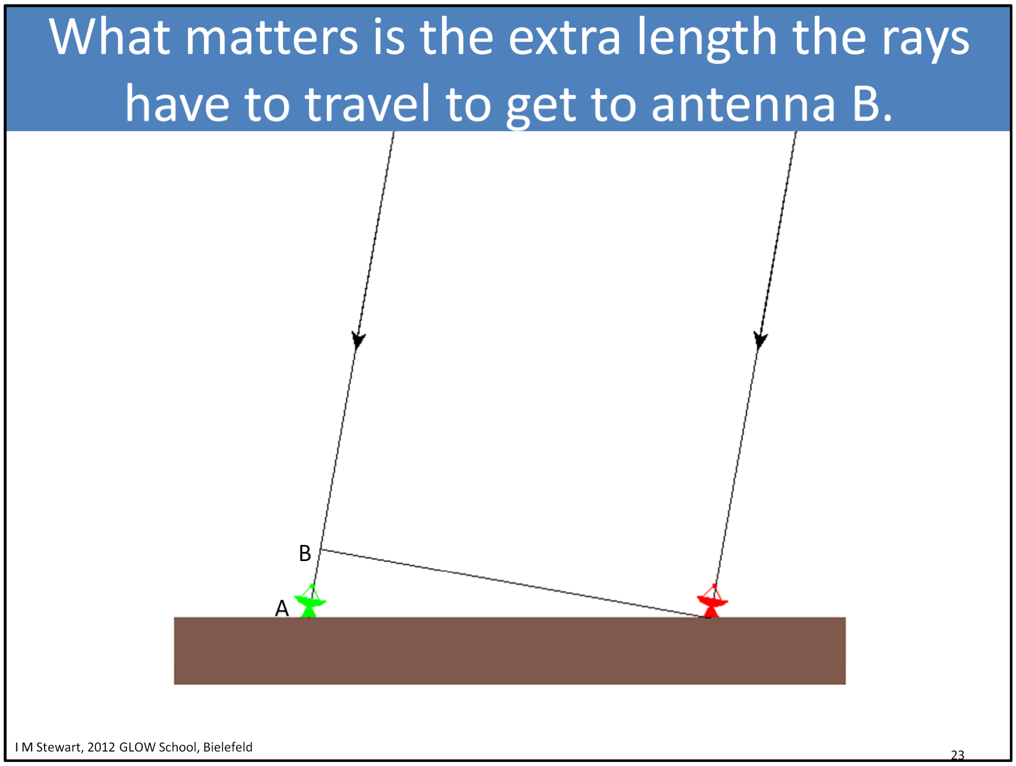

I have stripped away all the busyness now to leave the bare bones. The crucial quantity in interferometry is the length of the short line A-B. There are various ways people talk about this quantity, but they are all equivalent, and, because the geometry is so simple, readily convertible one into another, given a knowledge of the observing wavelength, the speed of light, and the separation between the antennas. Some of these ways to describe A-B are:

-Path difference (in length units)
-Delay (in time units)
-Phase difference (in degrees or radians)

Note that the path/delay/phase difference will increase with (i) increase in the displacement of the source from zenith; (ii) increase in the separation of the antennas; (iii) decrease in the observing wavelength.

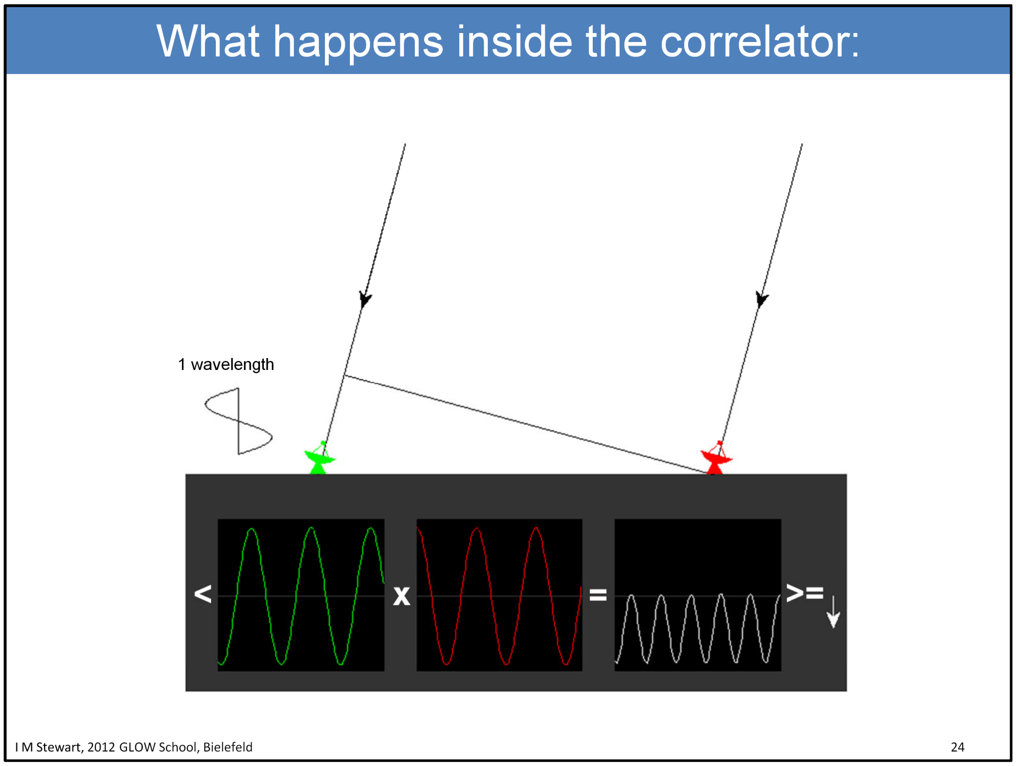

I have added some details back in now to illustrate how the correlator works. Basically it is simple: there are only 2 operations. The red and green plots represent plots of the change of voltage against time for some interval. The correlator first multiplies these together, generating the white plot. The 2nd correlator operation is to calculate the time-average (indicated by the enclosing <> brackets) of the white product plot. This final scalar value is shown as the white arrow on the right. The correlator performs this at intervals on the order of a second or so, thus produces a stream of numbers at that time spacing.

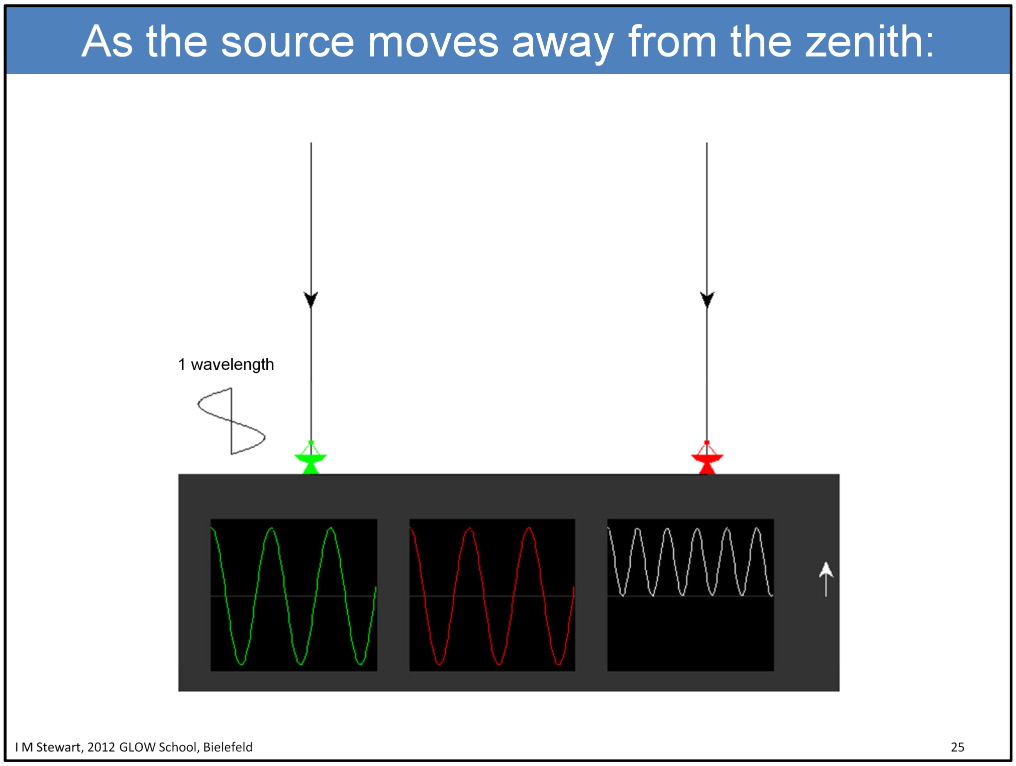

As the source moves away from the zenith, the delay or lag in the reception by the green antenna increases. As the path difference approaches a half-integer number of wavelengths, the red and green signals become 180 deg out of phase and the correlation becomes maximum negative. As the path differences approaches an integer number of wavelengths, the red and green signal become in phase and the correlation becomes maximum positive.

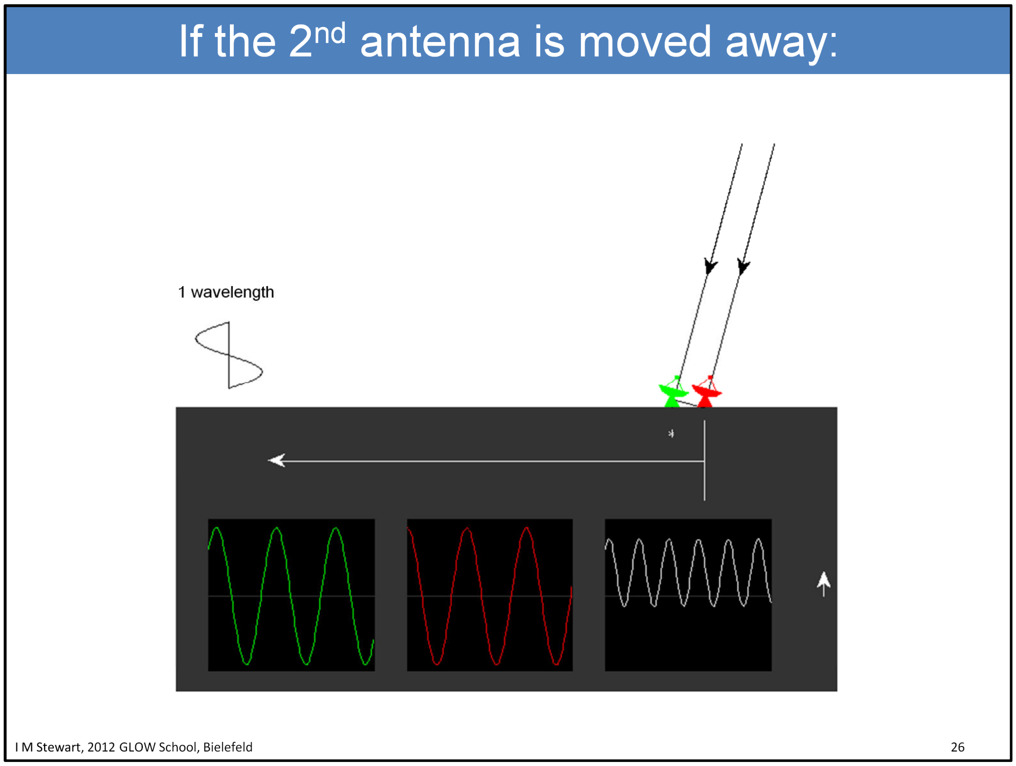

Now I keep the source fixed and imagine moving the green antenna to change the separation between them. I have included now too an additional 'chart recorder' (the white line) showing how the value of the correlation changes as a function of that separation.

This is quite important because it seems to show that the correlation is a sinusoidal function 'laid out on the ground'. If you think about it, hopefully you will realize that, the further the source is from the zenith, the shorter the pitch of the sinusoid will become. Now, think again about slides 10 and 11. In fact, we find that this white plot we have just made, of the way the correlation changes with displacement, is the Fourier transform of the source location – in fact what we have laid out on the ground is nothing less than the FT of the sky brightness distribution, which (recall slide 1) is what we want to get out of this exercise. This FT is known as the visibility function $V(x)$. Each individual measurement of correlation is a point sample of this visibility function.

In fact it is actually a little more abstract than I have said above, because V is not laid out on the ground as such: if we shift both antennas the same distance, the delay will not change, thus the correlation will not change. V is a function of the separation between the antennas. These separations (known as baselines) are usually given in terms of numbers of wavelengths, and the letter used for these quantities is not x but u. In the real world we have 3-dimensional distributions of antennas and the full range of letters for the 3 axes is u, v and w. Often we can, or want to, neglect the vertical displacement coordinate w, hence the frequent use of the term 'uv plane'.

Imaging strategy:

- Measure correlation function (it is called the visibility function) at all values of antenna separation (called baselines).
- Perform the Fourier inversion, using your favour library routine (eg numpy.fft.ifft() in python).

→ **Image!**

- Basically this *is* what we do, but there are a number of inaccuracies which need to be corrected.

Loose ends and inaccuracies

- Sampling function
- Non-coplanarity
- Effects of earth-rotation and multi-frequency aperture synthesis
- Bandwidth/time smearing
- Slowness of direct Fourier inversion: requiring (re)gridding, so can use FFT. Gridding introduces:
  - Aliasing
  - Narrow field
  - Complicates correcting for the uneven sampling function.

I'll discuss all of these in the following slides, but not necessarily in the neat order given on this slide.

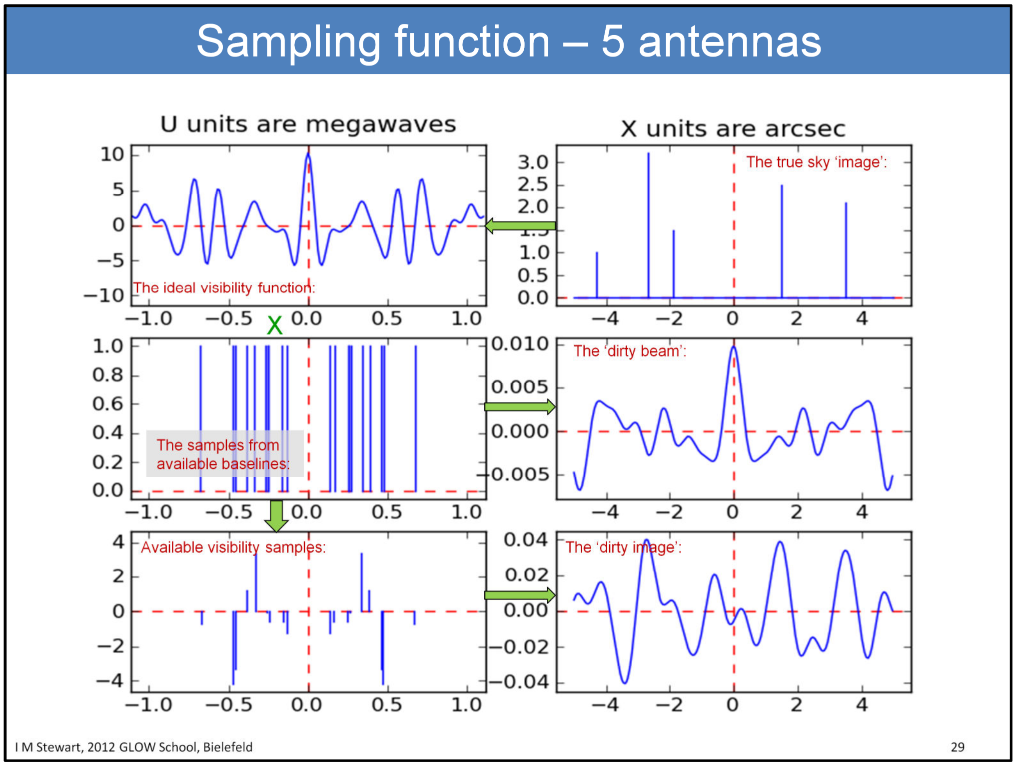

If we could measure the visibility function at all possible values of *u* and *v*, life would be simple. But we only have measurements of *V* at (*u*,*v*) pairs corresponding to the actual separations between antennas (= baselines) of our array. If we have *N* antennas, there are only *N**(*N*-1)/2 unique separations. That is often not many, and has the practical result that the ratio between the information we have to the information we would like to have is small. I have set up an example (in a 1-dimensional simulation) using just 5 antennas, randomly placed along a straight line, with the constraint that 2 antennas may not be closer than a certain minimum (which is enforced on us in practice by the physical size of the antennas).

So. To describe this slide: the starting plot is the one at upper right. I used the same plot as in slide 14, but that is just my laziness. I use it again to represent the true map of brightness distribution of this 1-dimensional sky. As we have shown, our 2-antenna correlation are samples of the Fourier transform of this, which I have placed at upper left. This is the 1-dimensional visibility function *V*(*u*). (Note I have the FTs on the left on this slide, the inverse of slide 14.)

The plot at middle left represents the available *u* values. We can sample *V* at each *u* representing a separation between a particular pair of antennas. So I have calculated all these *u*s and placed a 'spike' (delta function) at each location. For 5 antennas we expect 5*4/2 = 10 baselines. (I can only count 9 spikes, but probably some 2 are almost superimposed.) I have reflected the 9 on the other side of the origin because the FTs are much simpler if we keep everything symmetrical. On the RHS is the inverse FT of this set of sampling spikes. In interferometry jargon, this FT^-1 of the sampling function is called the 'dirty beam'. It plays exactly the same role as the PSF of an optical mirror.

At the bottom left is the product of the true, continuous, infinite visibility function with the samples. The result is ALL the information we have available. It looks pretty sad. When we back-transform, this looks like a pretty sad representation of the sky too. This final result is known as the 'dirty image'. Product ⇔ convolution, so on the RHS the lower right image is the convolution of the true sky (upper right) with the dirty beam (middle right).

The green arrows represent the natural flow of information during the observation.

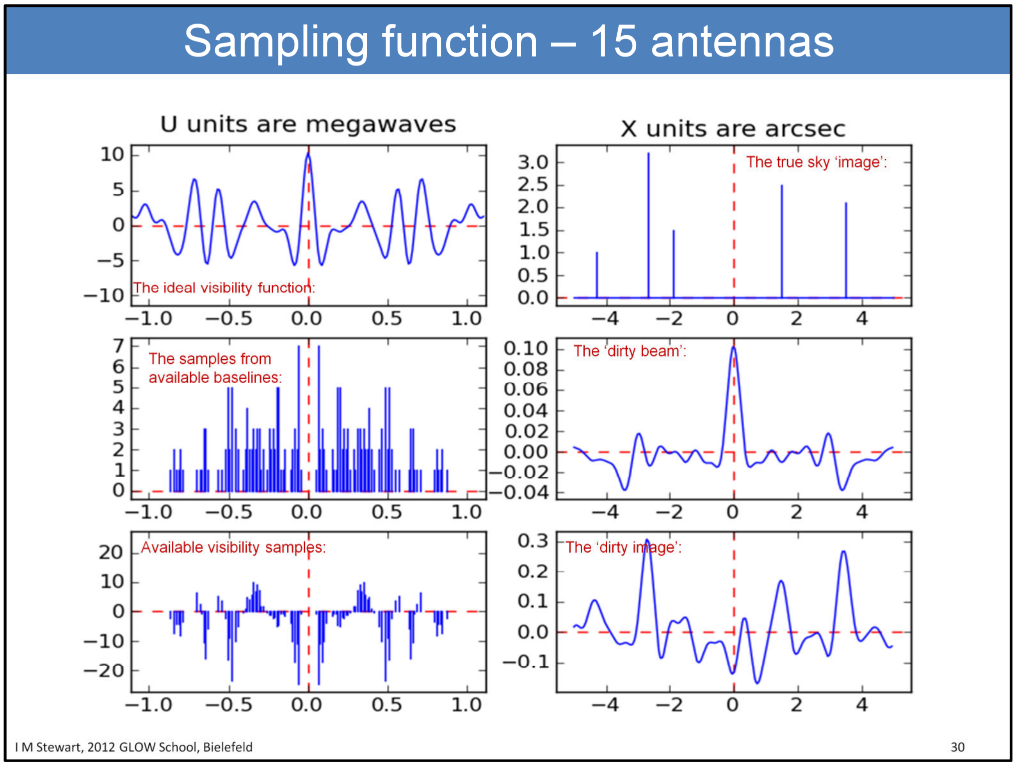

Now I am doing the same exercise with 15 antennas, which gives us 15*14/2 = 105 baselines. Of course this gives us much better sampling of the visibility function and our corresponding dirty beam is better (less power in sidelobes), and thus so is the dirty image.

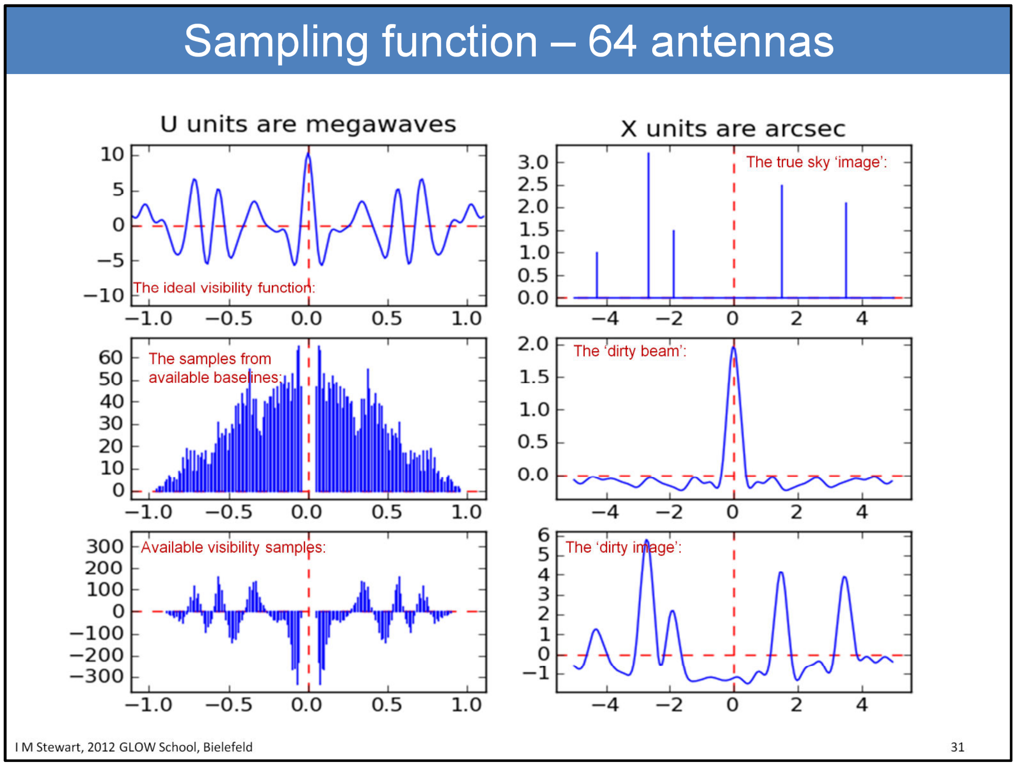

For this final example I have used 64 antennas, the same number as the planned MeerKAT array in South Africa. (The 1-dimensional locations I have chosen here bear no relation to the planned locations for the MeerKAT antennas.) This gives 3000 or so baselines. The dirty image is about as good as we can get without weighting the visibilities (which I will discuss in a few slides time). We can see all the sources now.

The lesson from these 3 examples: We need as many samples of *V* as possible! However, as we have seen, that number seems to be set by the number of antennas. Are there are clever ways to get around this limit?

Wide-band imaging.

How can we increase UV coverage?
…we could get more baselines if we moved the antennas!

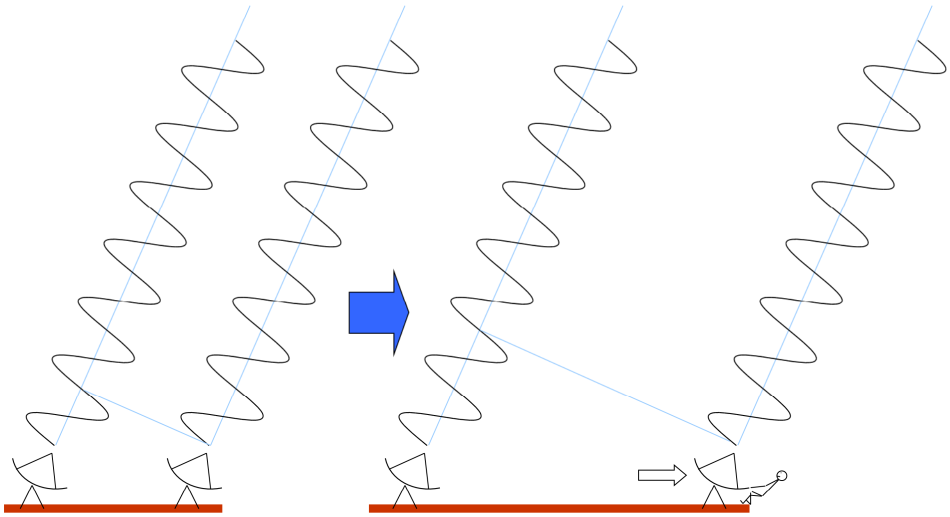

. . . but it is simpler to change the observing wavelength.

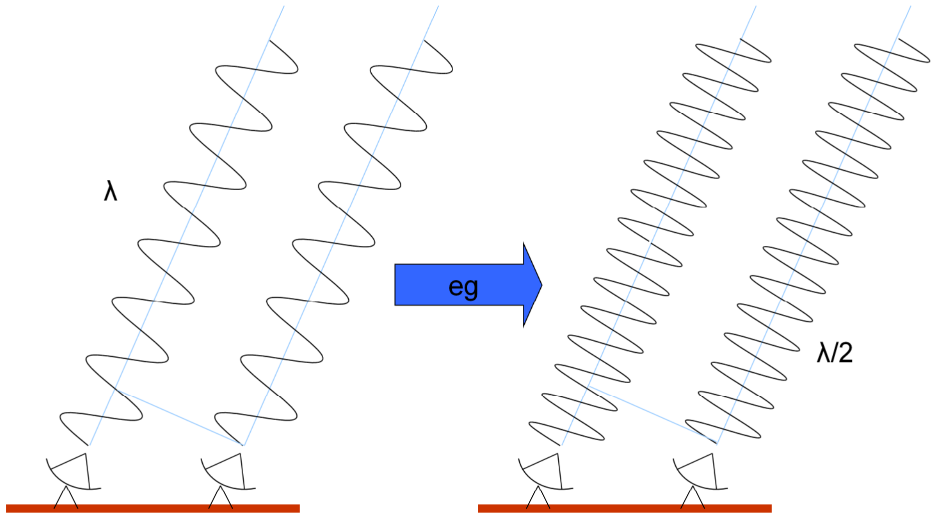

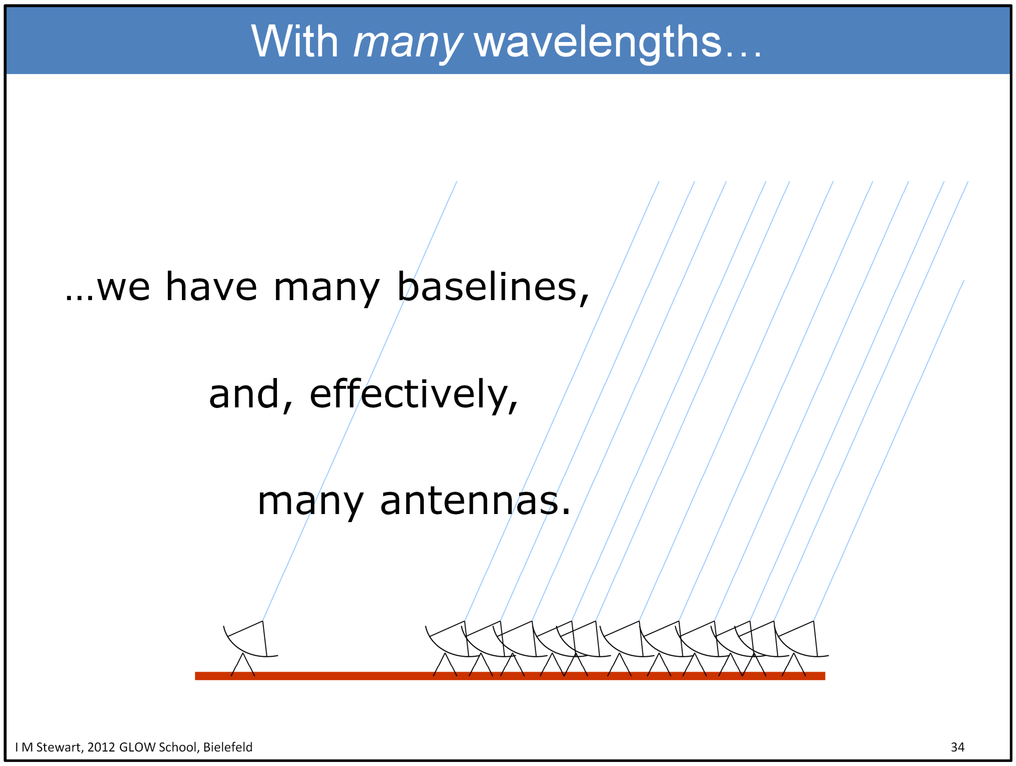

This is called ‘multi-frequency synthesis’. ‘Synthesis’ because we are constructing mathematically, in effect, a big mirror. The other main way to increase the number of samples of *V* is called ‘earth-rotation synthesis’ and I will describe it in a little while. But first I want to describe the problem of non-coplanar arrays.

Problem: non-coplanar arrays.

There is now no common zenith – so there is no place in the sky from which signals arrive at the correlator in phase.

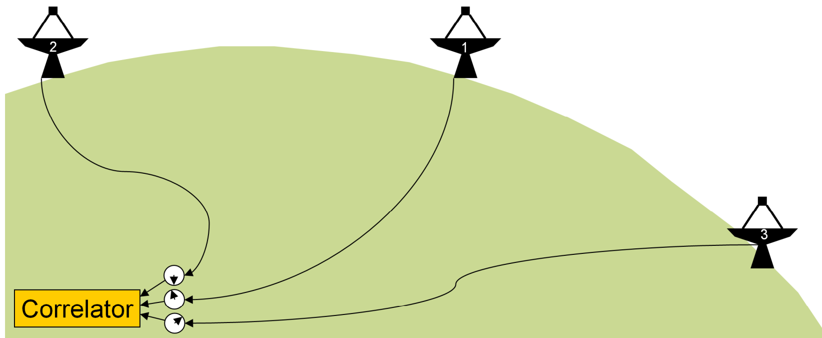

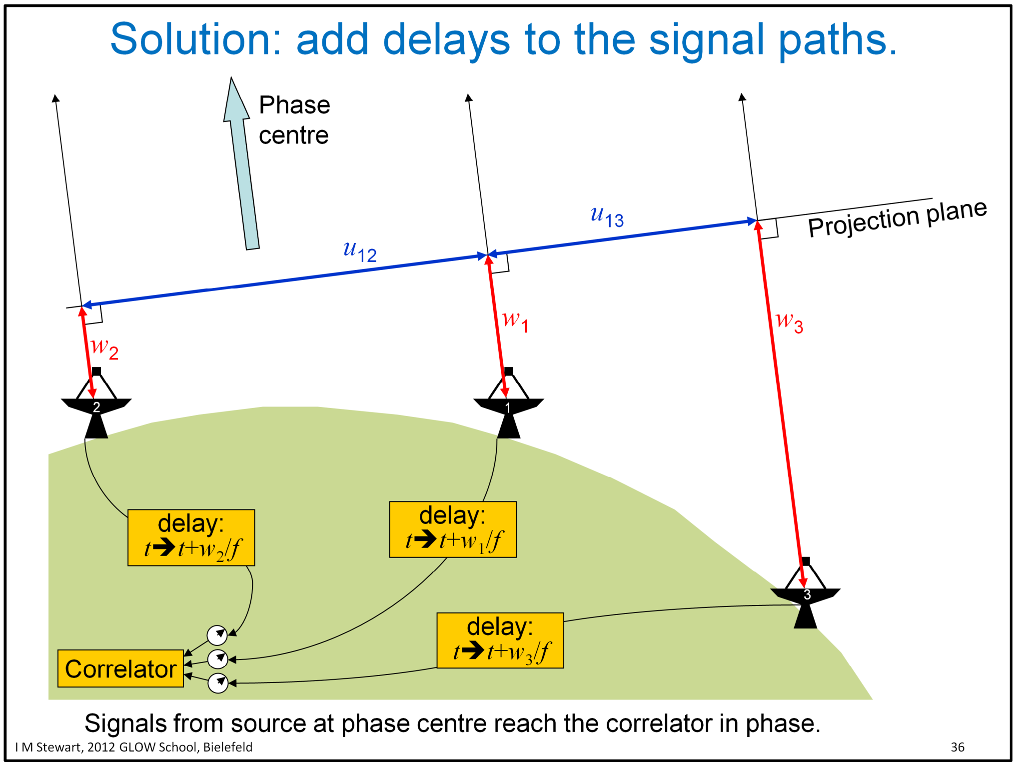

If we do this, we once again have a Fourier relationship between the sky brightness distribution, centred on the phase centre, and the visibility function, mapped out now in terms of <u>displacements projected on the plane shown</u>. The problem is that, if you do the maths, you will find that the Fourier relation which formerly, i.e. with a coplanar array, was exactly true for all points on the visible celestial hemisphere, is now, for the present non-coplanar array, only exactly true for sources lying exactly in the direction of the phase centre. For all other locations it is only approximately true. The approximation gets worse the further one gets from the phase centre. Thus if we back-transform the sampled visibilities, the resulting dirty image is only an acceptable approximation of the sky within a fairly narrow (typically on the order of a degree or so) radius from the phase centre. There are algorithms to correct the resulting aberrations but of course it would be better if the aberration was not there in the first place.

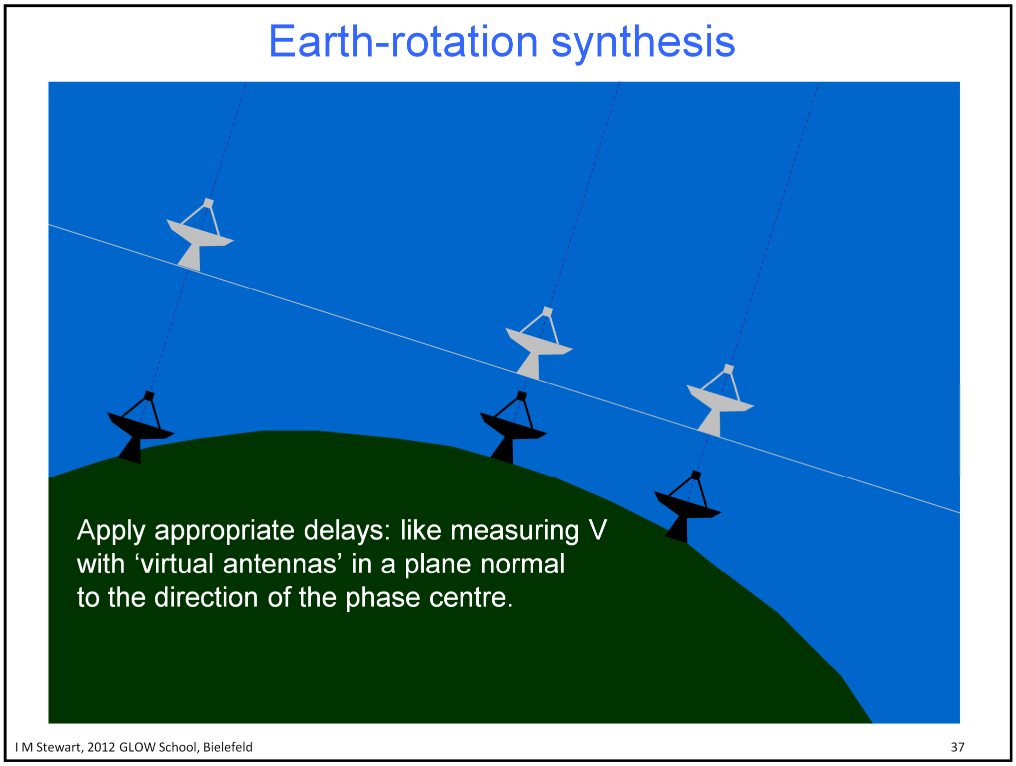

If we observe for some significant fraction of the day, the rotation of the earth changes the projected locations of the antennas. The delays necessary to keep the array pointed to the phase centre also change constantly. In modern instruments these delays are constantly recalculated and updated by the correlator.

Earth-rotation synthesis

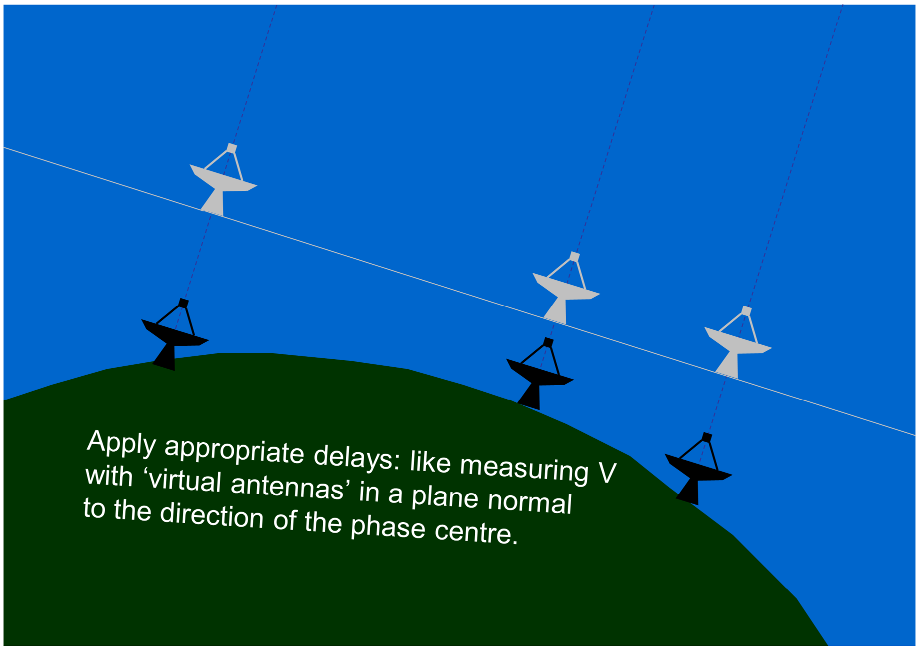

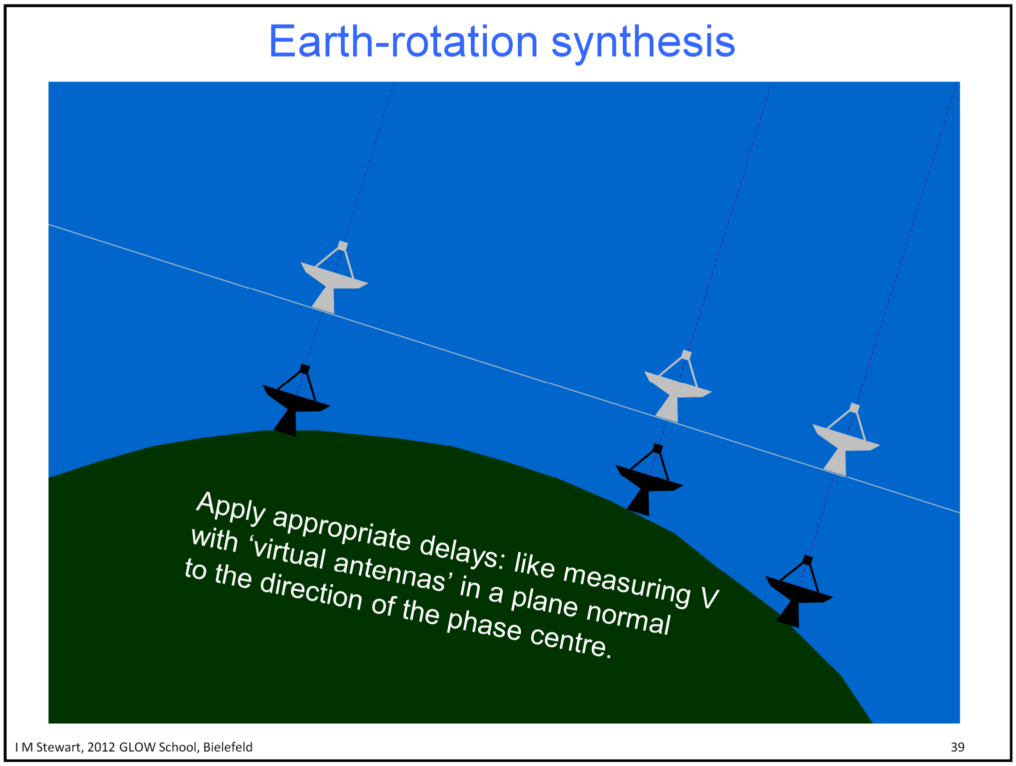

The change in length and orientation of the projected baselines is actually an advantage because it opens a way to make more samples (although no longer simultaneously) of the visibility function. This technique is called earth-rotation aperture synthesis. The following animation illustrates this.

View from the phase centre

I M Stewart, 2012 GLOW School, Bielefeld 40

This is a quasar's eye view of the change in projected baselines as the earth rotates. Note that the sky stays motionless while the earth, carrying the antennas, rotates beneath it.

A simulated example.

The full visibility function $V(u,v)$ (real part only shown).

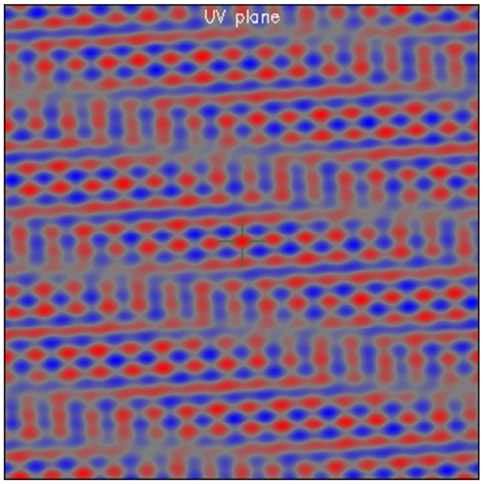

Red positive; blue negative.

A familiar pattern of 'sources'

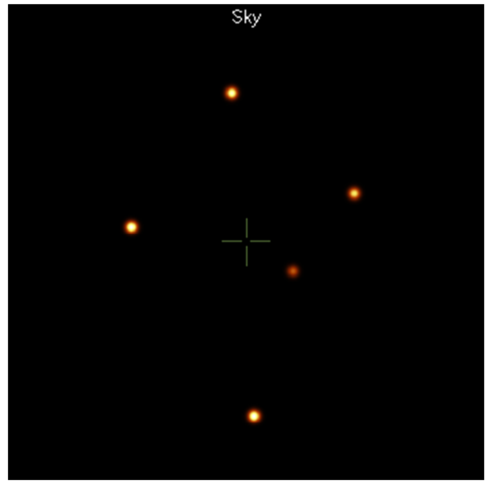

(I've taken some liberties here – obviously the stars of the Southern Cross are not strong radio sources – I've also rescaled their angular separations.)

Now I want to show the effect of different sampling in 2 dimensions. I start here with a simulated true sky brightness distribution on the RHS, with its FT, the full visibility function (or that part of it which fits within the square), on the LHS.

‘Snapshot’ sampling of V is poor.

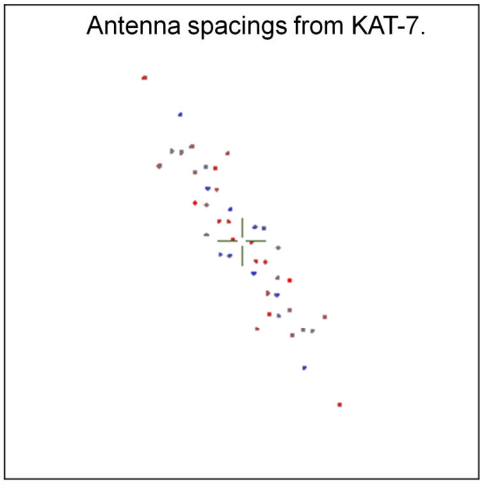

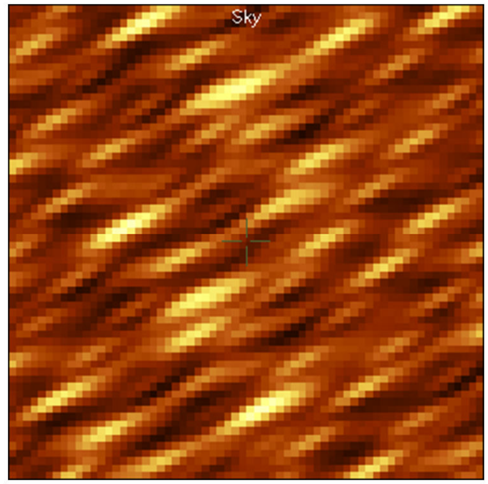

I M Stewart, 2012 GLOW School, Bielefeld 42

Now we turn on the interferometer. We have 7 dishes, thus 7*6/2 = 21 baselines. We are doing neither frequency nor earth rotation synthesis – we just have the 42 (21 + their reflected partners) samples of V. Predictably, the resulting dirty image (RHS) is very poor.

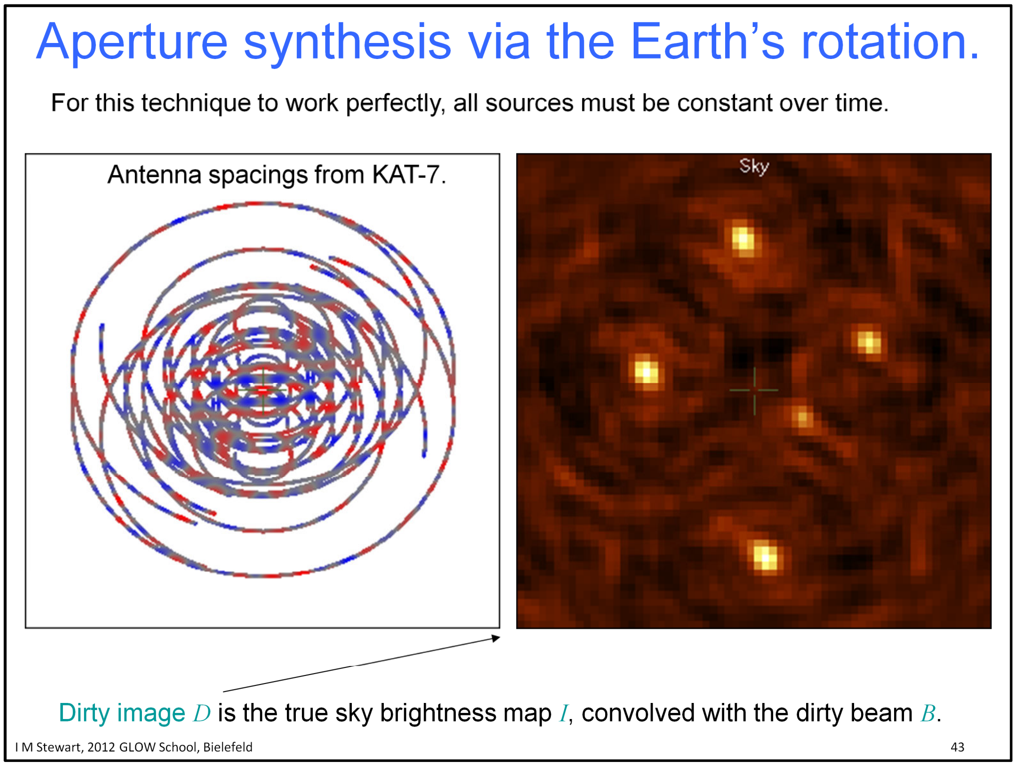

Now we leave the interferometer running during the course of the day. We're doing earth-rotation synthesis and this gives us many more samples. The dirty image is much improved.

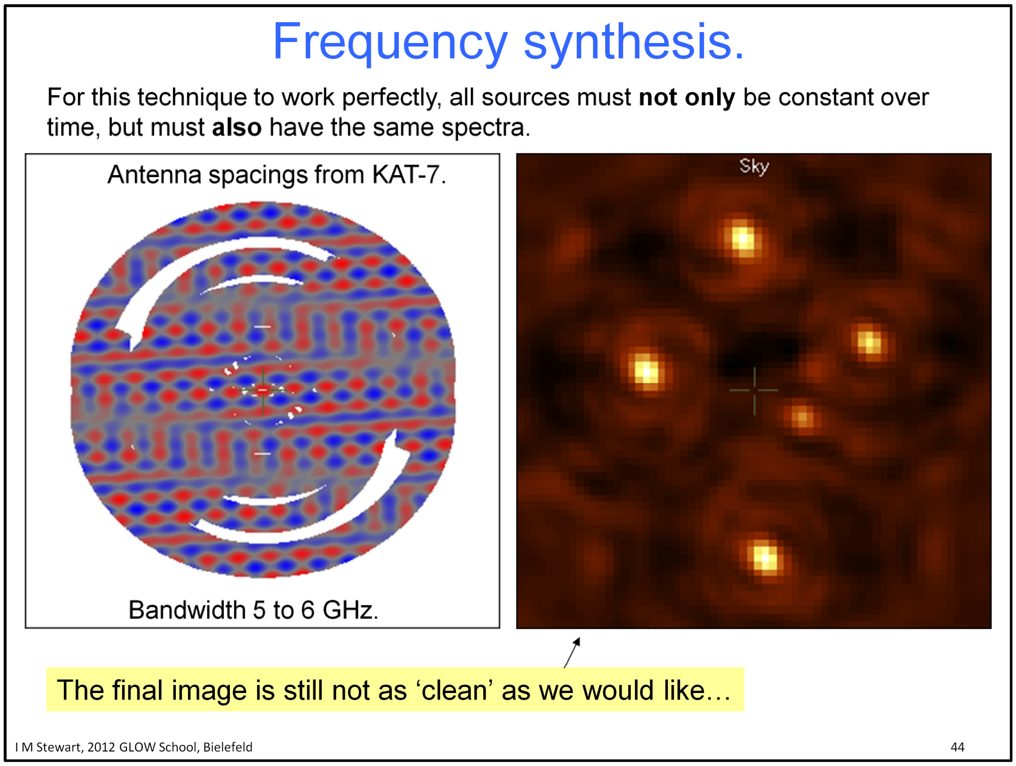

Finally if we do frequency synthesis over a significant bandwidth, we widen the tracks on the UV plane and get even better coverage, with the best dirty image yet. To go further we really have to deconvolve (see deconvolution lecture).

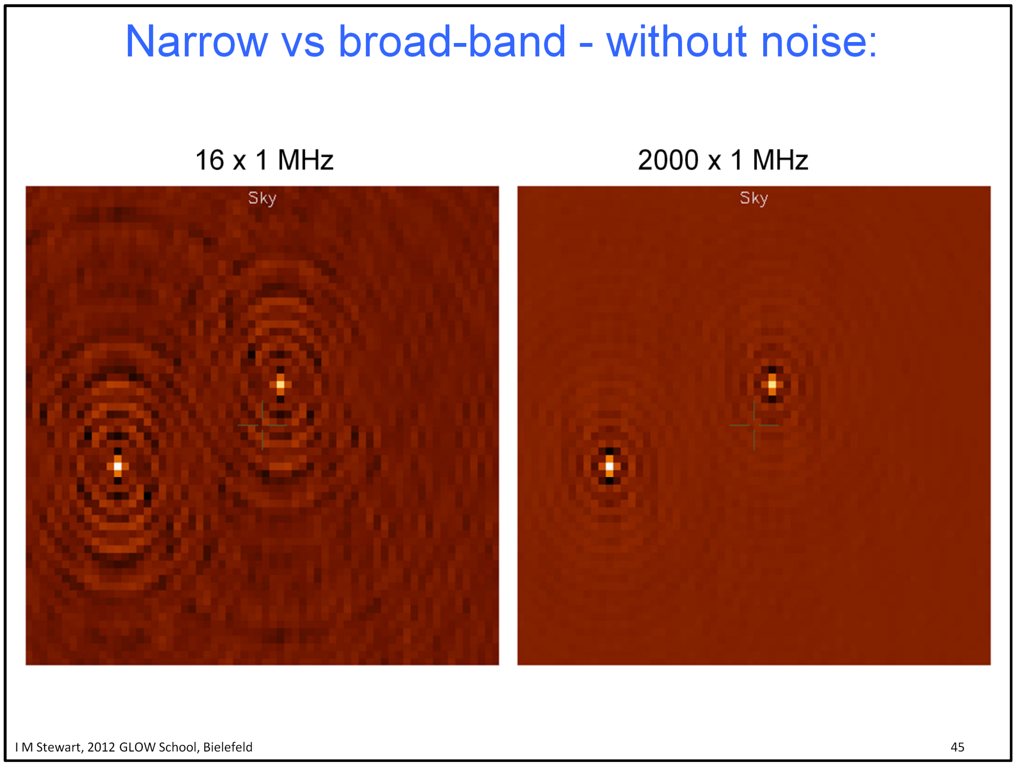

I have changed the source distribution and the array now – I just want to illustrate the good effect of using a wide band on the noise characteristics of the dirty images. So far in my simulations I have included no noise at all. In the following slide I add some.

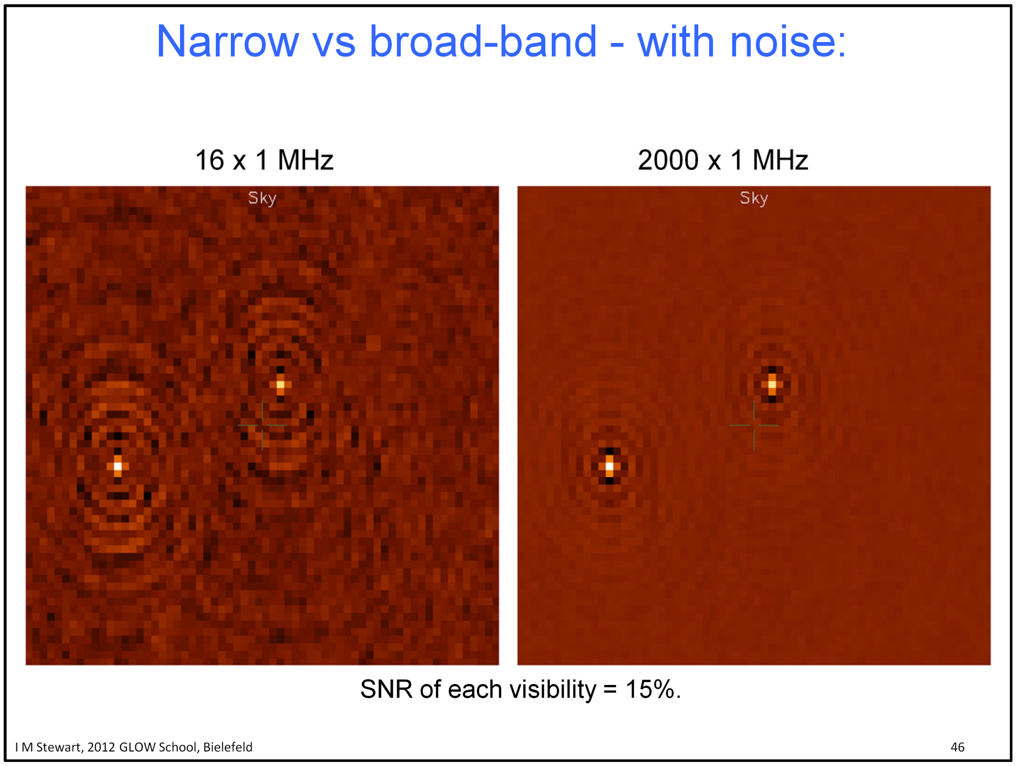

Both the narrow and wide-band plots I show here have the same amount of noise added per visibility sample, but because we have many more samples in the wide-band case (the RH image), the noise is relatively lower.

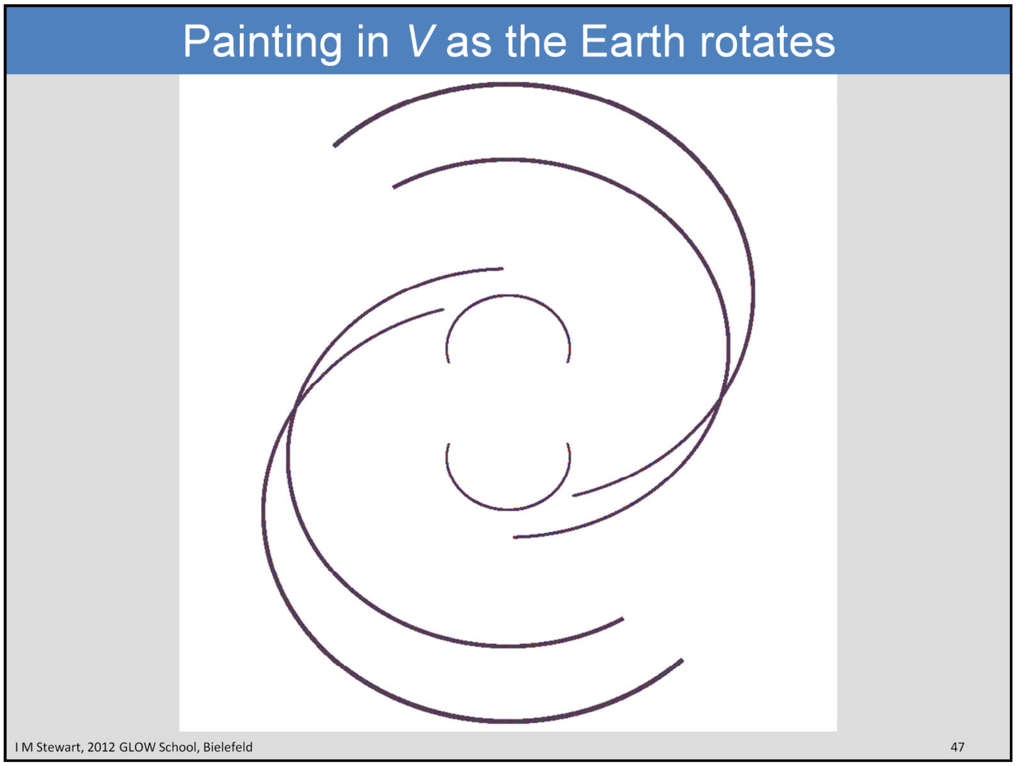

In the next slide I show a little movie to illustrate the earth-rotation synthesis. In this movie I first zoom in to one of the curved tracks in the present slide, then watch as the track slowly lengthens as the earth rotates. We will be able to see how the bandwidth gradually paints in more of the red and blue visibility function as the earth rotates.

V is a continuous function. However, if we want to do some processing of our data on a computer, we cannot deal with a smooth function: we need a list of scalar values. Therefore we have to bin up *V* in some way, average *V* in each bin, and spit out each average as a single scalar value.

Painting in *V* as the Earth rotates

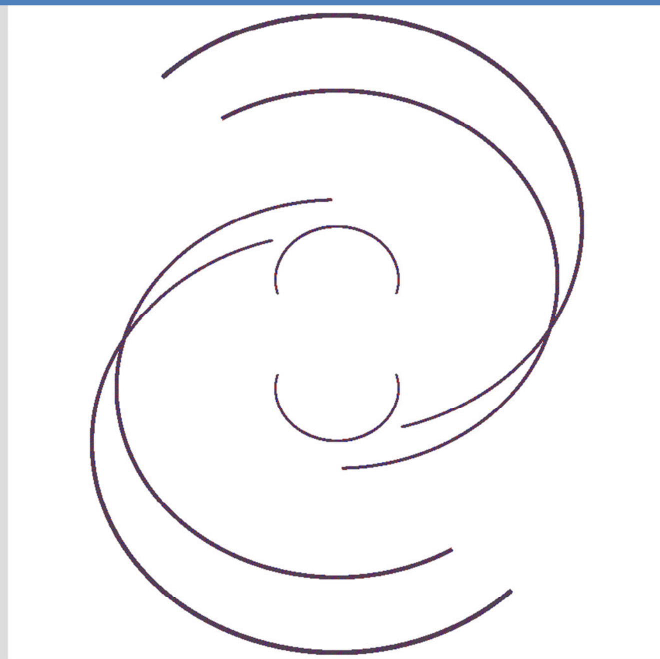

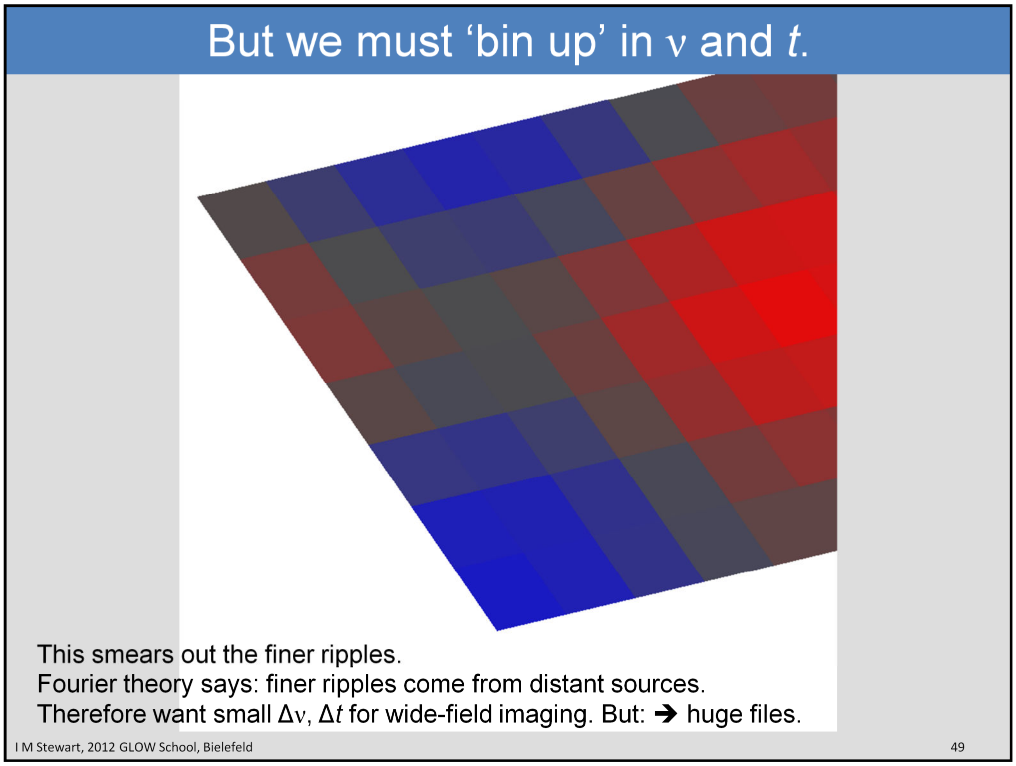

The correlator does this by firstly integrating the signals over some time period (we have already touched on this), and it also divides the total bandwidth into frequency channels and integrating *V* within each channel. The result is to smear out *V* in rhomboidal chunks as shown.

Note that, because the rhomboids are not of the same shape and size over the whole UV plane, this process can NOT be described by a convolution.

A drawback to this necessity is that it places another limit on the breadth of the patch of sky we can image in one go. (The first limit we considered came from non-coplanarity.)

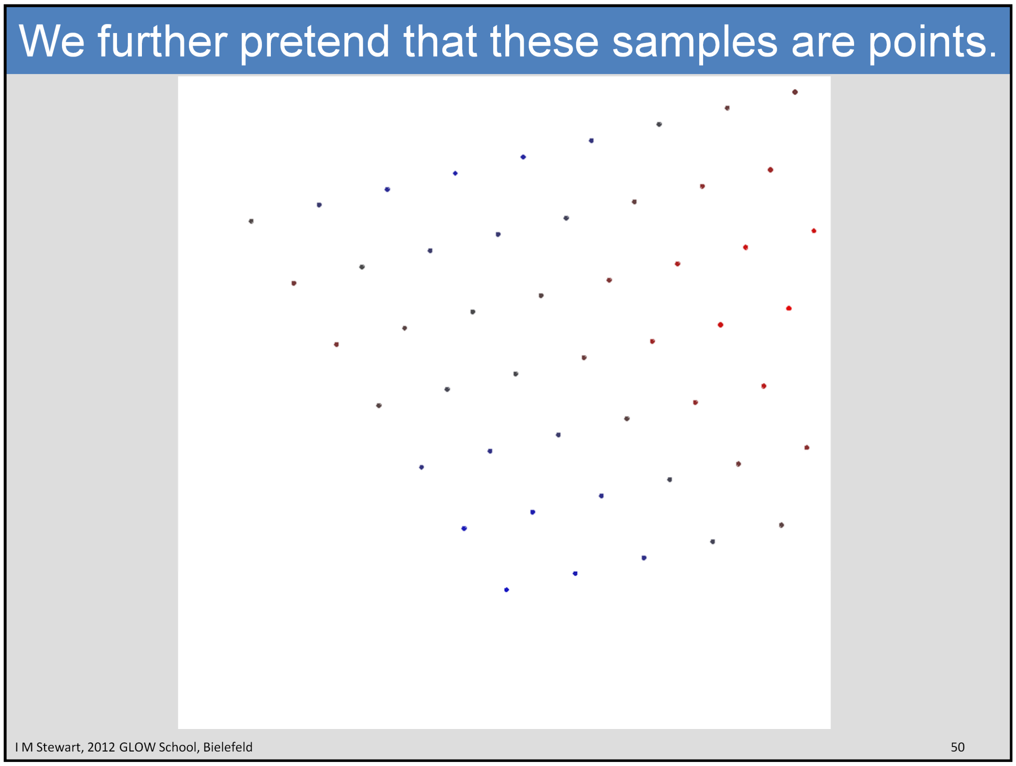

From this point, we want to perform an inverse Fourier transform to obtain our dirty image. It is perfectly possible to do this directly: for each visibility sample, we generate the appropriate sinusoid (remember from slides 10 and 11 that the FT of a delta function, a point, is a sinusoid.) and add that to the sky image we are constructing. The problem is that this is SLOW. If we want to do a fast Fourier transform we need to regrid the samples onto a regularly-spaced orthogonal grid.

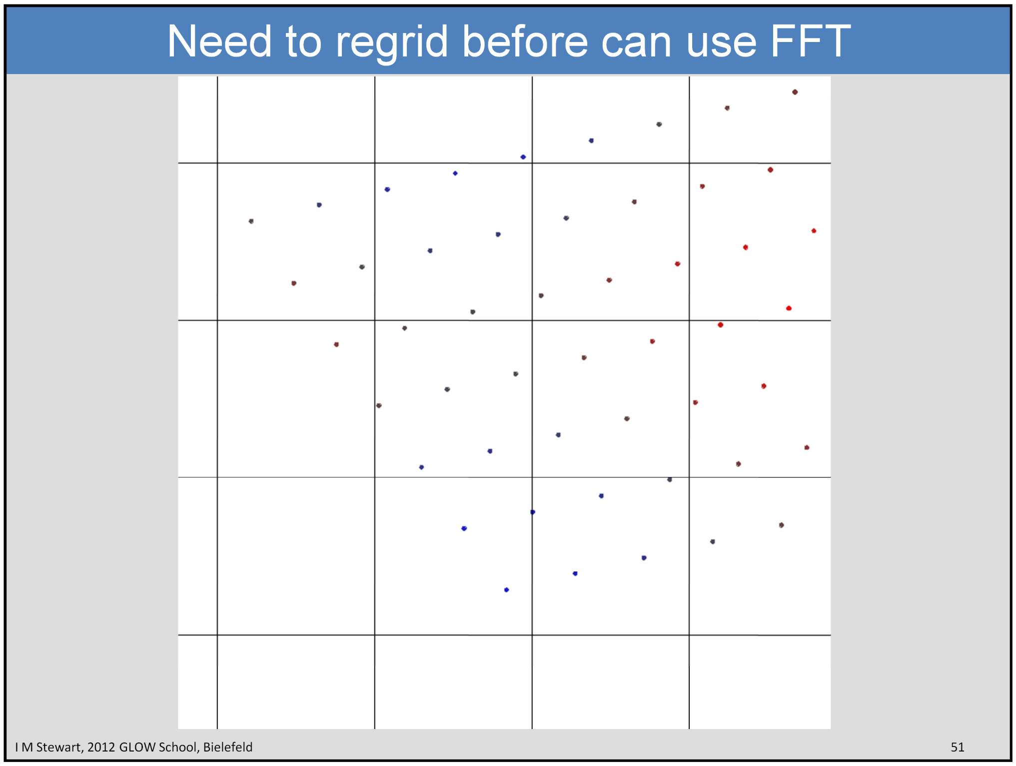

The simplest gridding scheme is to add values in pixels. There are many more elaborate schemes but they all can be expressed as a convolution by a gridding kernel followed by sampling at the pixel centres.

Gridding involves convolution with a 'gridding kernel', then sampling the resulting smooth function on a regular grid

Add-visibilities-in-pixels algorithm is equivalent to convolving with a 'square top-hat' gridding kernel exactly the same size as the pixels.

- Convolution of the visibilities means multiplication in the dirty image (by the FT of the gridding kernel) – DI becomes vignetted.
  - Vignetter has to be divided out afterwards.
- Small kernel gives large vignetter.
- Sampling leads to aliasing (partially offset by the vignetter).

What's the noise in these measurements?

- Theory of noise in a cross-correlation is a little involved... but if we assume the source flux density S is weak compared to sky+system noise, then

$$S_{\text{rms}} \approx \frac{k\sqrt{2}}{\sqrt{A_{\text{e}1} A_{\text{e}2}}} \times \frac{\sqrt{T_{\text{total1}} T_{\text{total2}}}}{\sqrt{t \Delta \nu}}$$

- If antennas the same,

$$S_{\text{rms}} \approx \frac{k\sqrt{2}}{A_{\text{e}}} \times \frac{T_{\text{total}}}{\sqrt{t \Delta \nu}}$$

- Noise is root 2 larger than from a single dish of combined area.
  - Because autocorrelations not done ➔ information lost.

I M Stewart, 2012 GLOW School, Bielefeld 53

I haven't finished talking about the effects of gridding but I need to talk about noise a bit first.

Resulting noise in the image:

I M Stewart, 2012 GLOW School, Bielefeld 54

The noise values are correlated – i.e. not white. You ought to learn to recognize correlated noise.

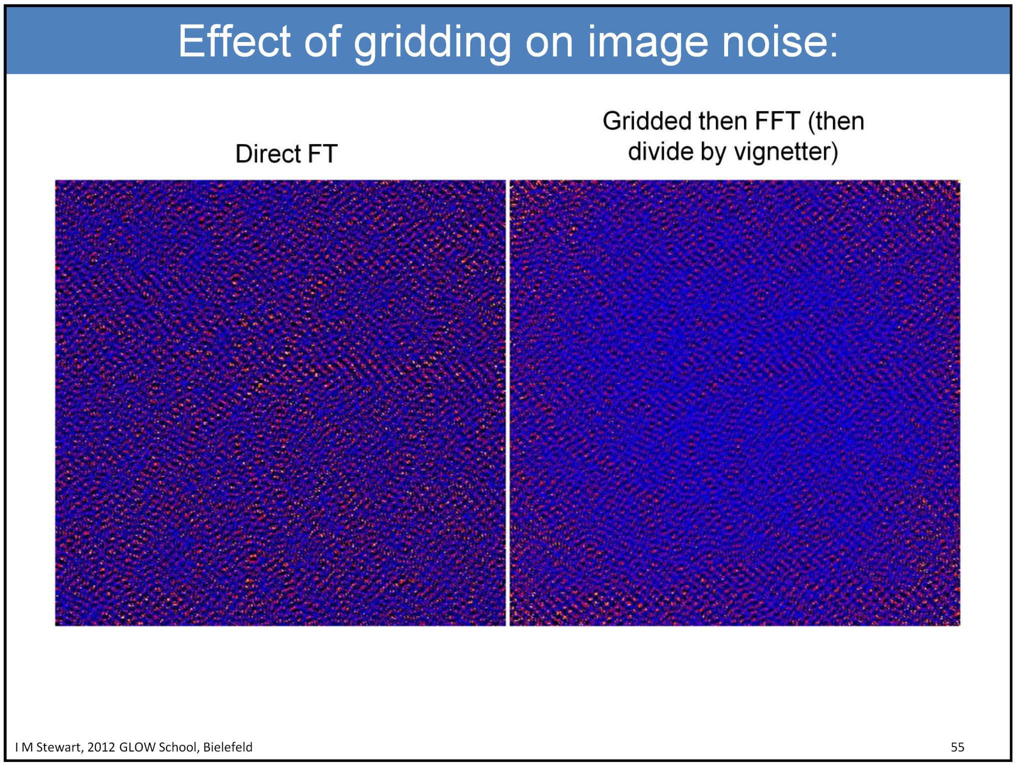

Refer back to slide 52: in the direct (but slow) FT, the noise is uniform across the field (the amount of dirty image we can display is also only limited by the slide of powerpoint slides). For the discrete transform (this is a FT of gridded data), after we divide out the vignetting, we are left with noise which rises near the edges. Usually (e.g. inside AIPS task IMAGR), the outer ¾ of the image area are discarded.

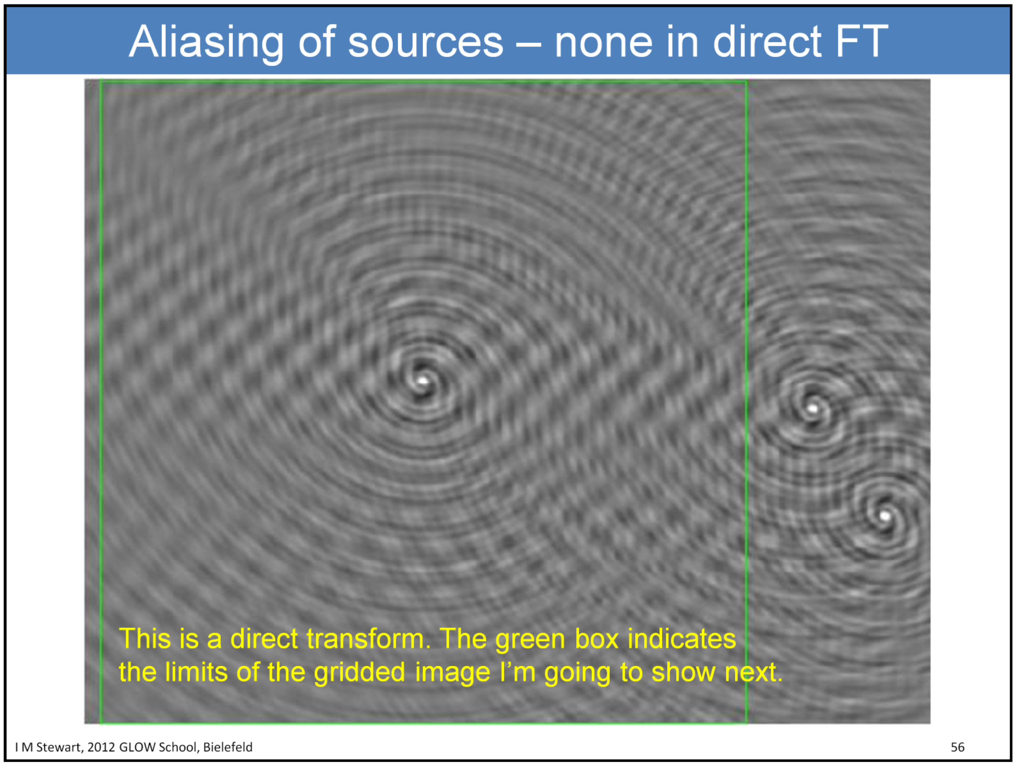

I have again set the noise level to zero. Here I want to illustrate the effects of gridding on sources.

This is a dirty image formed by direct transform of the un-regridded visibilities as they come direct from the correlator. Note that we still have limited sampling of the UV plane, leading to each source image being convolved with the FT^-1 of the sampling function (the dirty beam). You ought to be able to recognize that a convolution is occurring in this image. Just as in slide 13, each source position replicates a copy of the dirty beam pattern (or PSF).

The three sources I start with were all given the same flux density, which is reflected by their equal brightness in the dirty image.

Note again that the size of the patch of sky I can display here (even though non-coplanarity and time-bandwidth-smearing effects will limit the size of the accurately mapped part of the sky) is only limited to how much I can fit on the slide.

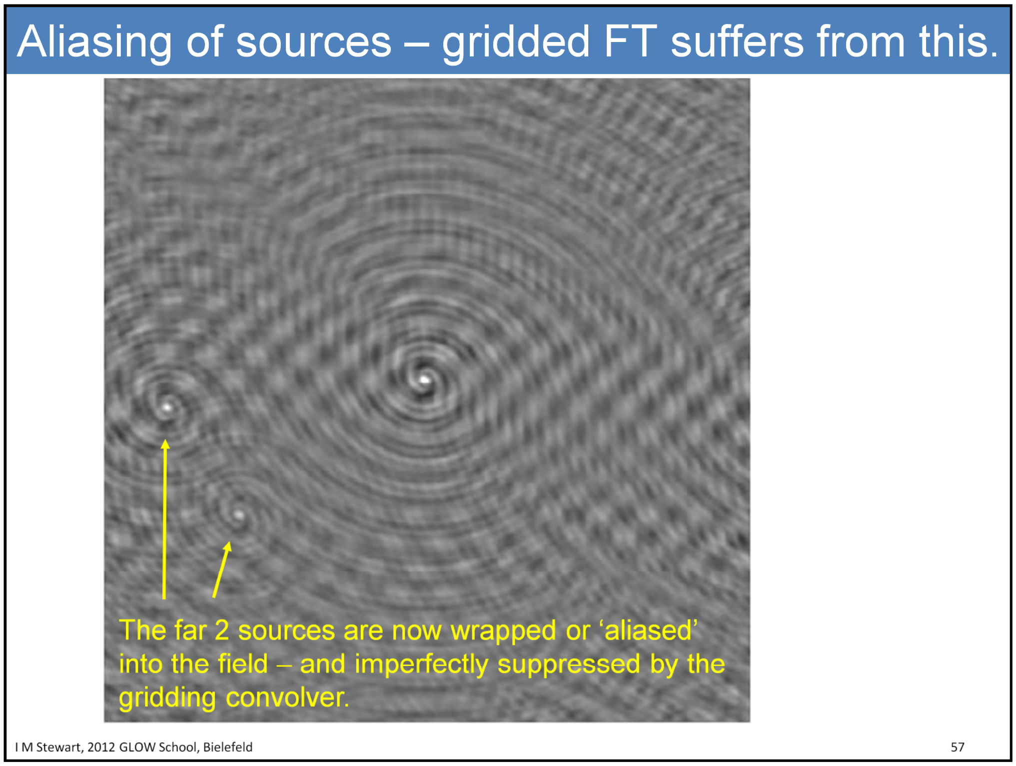

Here the visibilities have been gridded then the fast FT was employed. Effects of this:

(i) The size of the image is limited. One must choose a pixel size and a number of pixels in order to do the gridding.

(ii) The FFT thinks everything is periodic – the practical effect of this is to 'wrap' ANY and ALL sources which in actuality fall outside the bounds of the gridded image back in the image again. Thus the 2 sources which lay to the right outside the green box in the previous slide now wrap back into the image from the LHS.

(iii) This is actually a feature not a bug: before this wrapping occurs, the sky image is vignetted by the Fourier inverse of the gridding kernel. This is the reason that the 2 'wrapped' sources are now fainter than the in-box source.

Weighting

- There are usually far more short than long baselines.

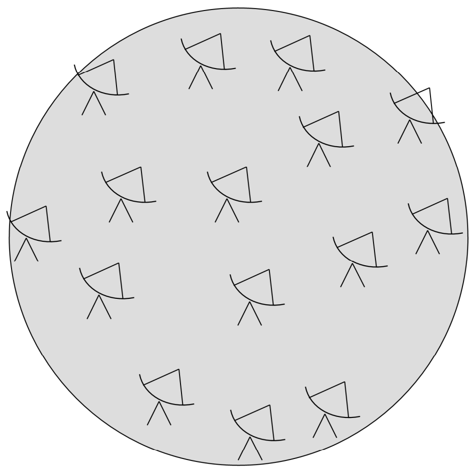

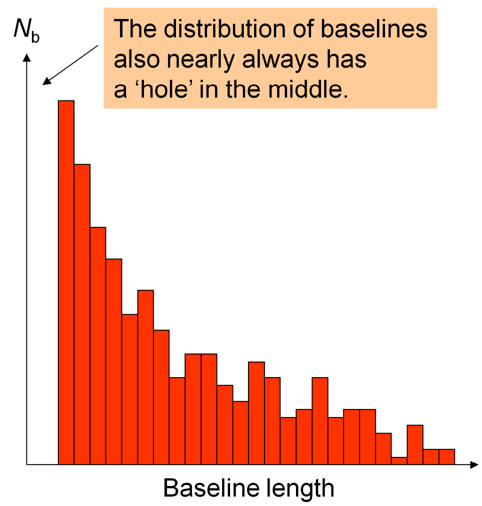

Weighting

- A crude example:

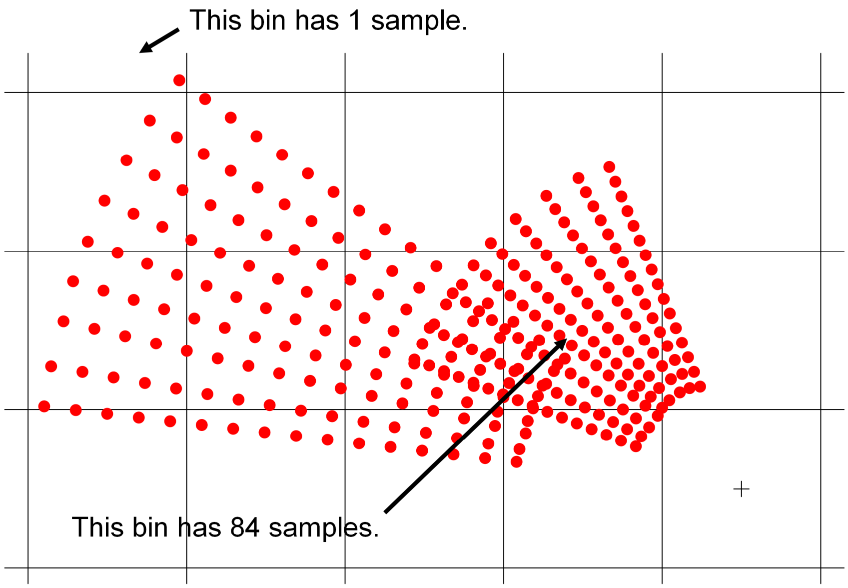

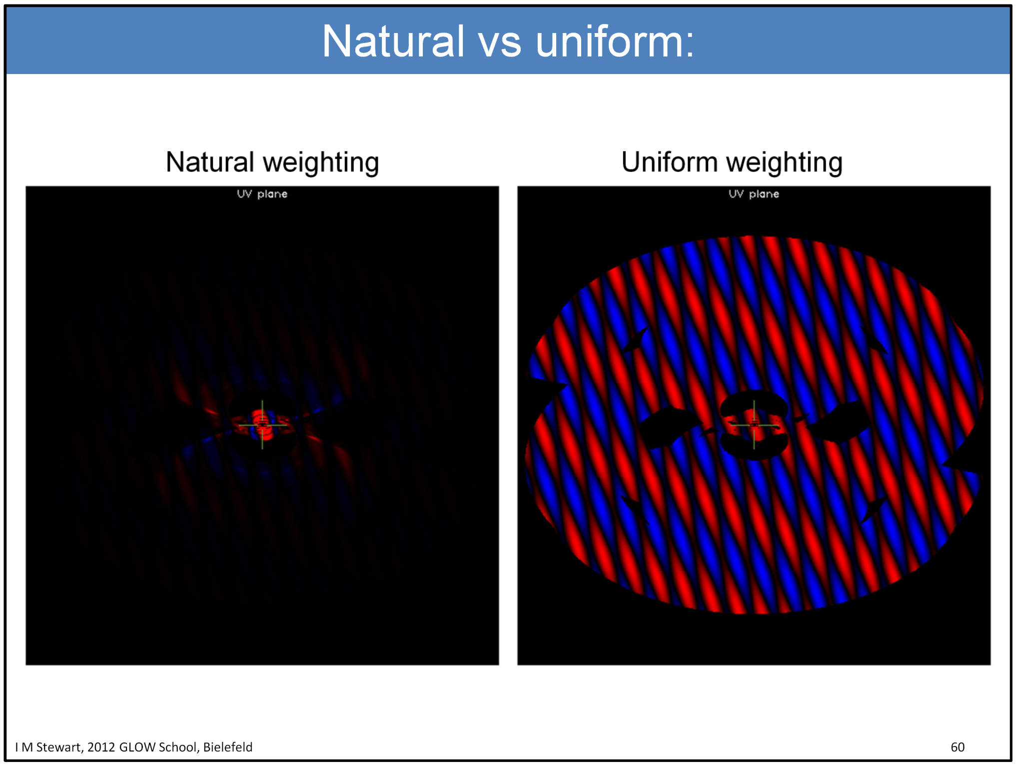

Before gridding (or direct transform) we can multiply each visibility by an appropriate weight. There are many possible choices of weighting scheme. Weighting carries a tradeoff: a weight scheme can improve the dirty beam ('improve' means to make it more compact and more importantly, to reduce its sidelobes) but usually at a cost of increasing the noise in the image. Our choice of weight scheme will be partly dictated by the brightness and structure of the target source.

Here I illustrate the two 'basic' choices of scheme. Natural weighting is not to weight them at all. Uniform weighting is to divide each visibility by effectively the number of visibilities per grid cell. This negates the usual emphasis on the densely-packed baselines towards the centre of the UV plane.

The resulting dirty images:

Natural weighting

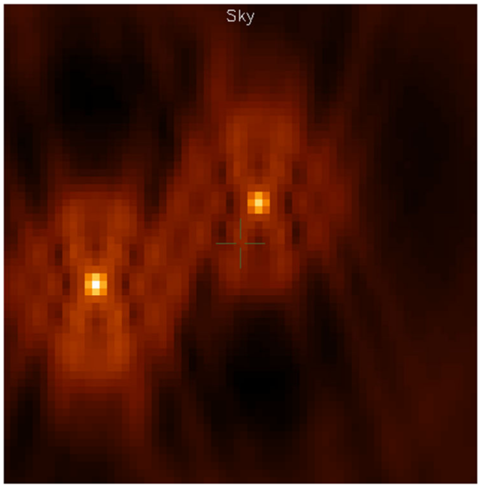

Uniform weighting

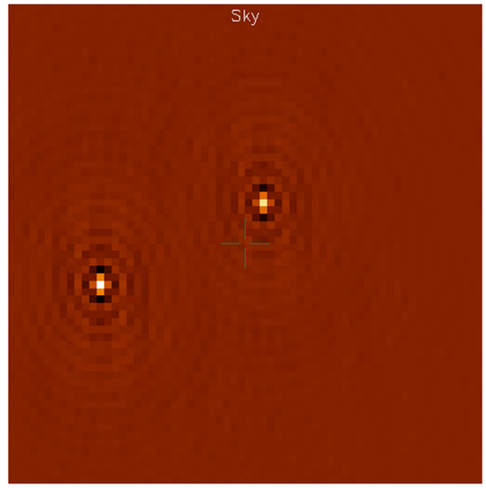

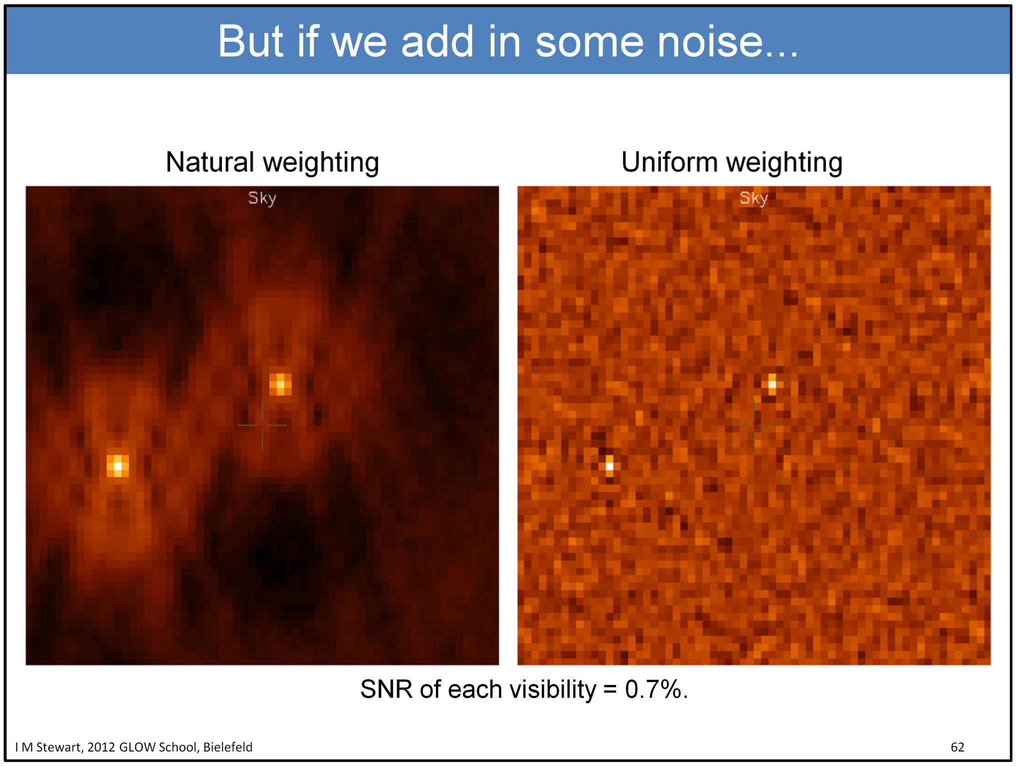

Natural weighting gives THE best signal-to-noise ratio. Uniform weighting can be significantly worse.

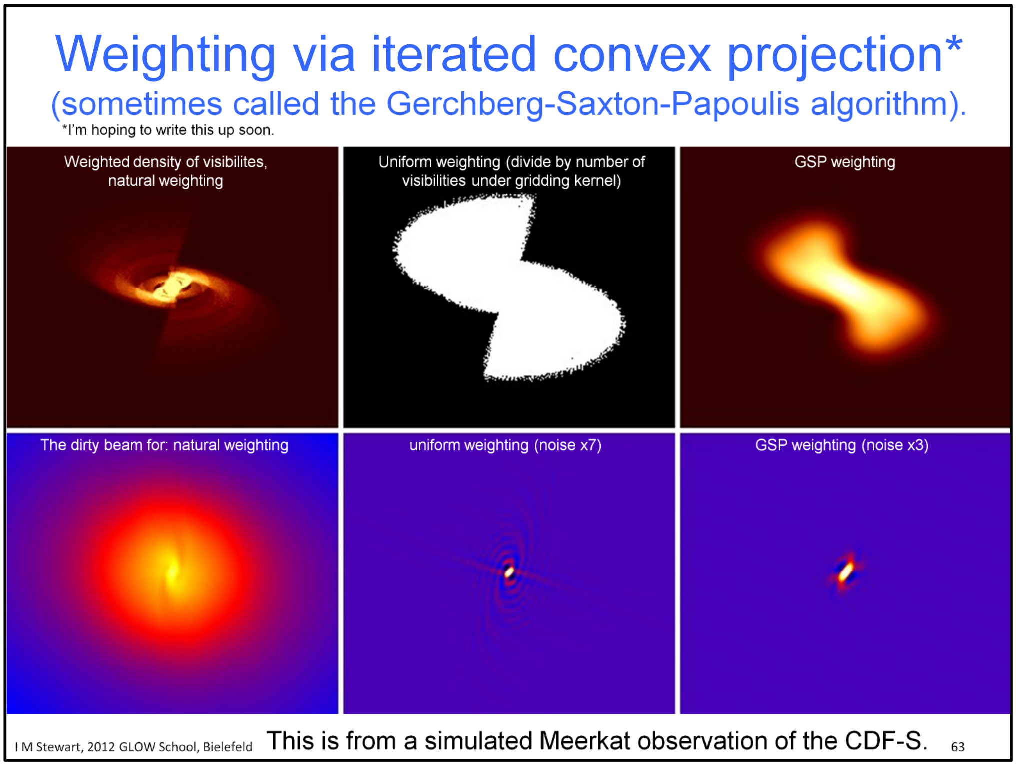

This is a scheme I am currently working on…

…and that is all the slides I have.