Synthesis Imaging

Claire Chandler, Sanjay Bhatnagar

NRAO/Socorro

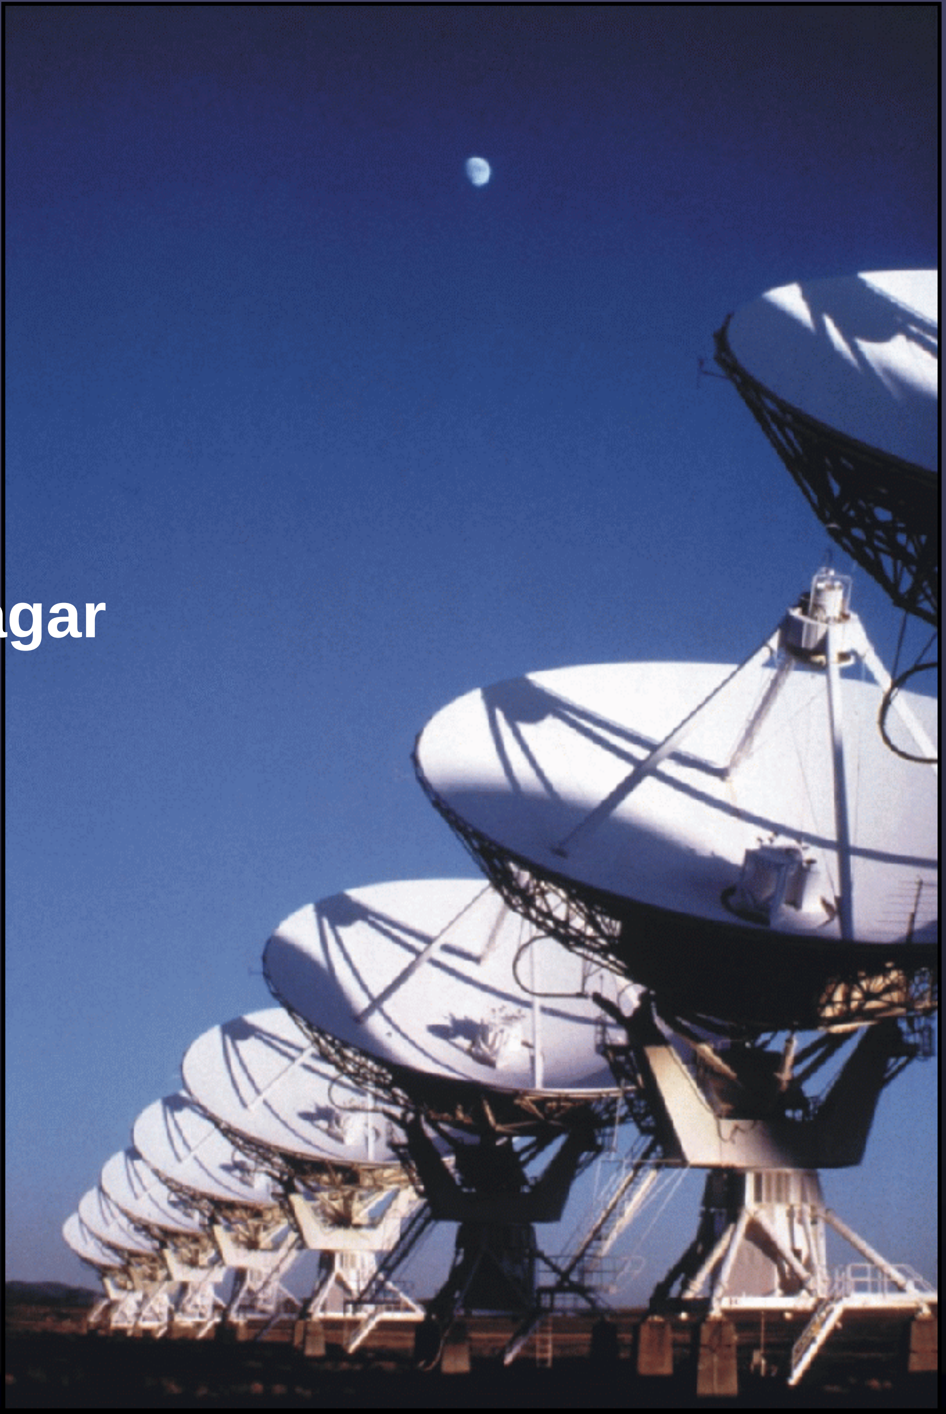

Michelson Summer Workshop

Caltech, July 24-28, 2006

Synthesis Imaging

- Based on the van Cittert-Zernike theorem:
  - The complex visibility of a source, $V(u,v)$, is the Fourier Transform of its intensity distribution on the sky, $I(l,m)$

$$V(u,v) = \iint I(l,m) e^{2\pi i(ul+vm)} dl dm$$

$$I(l,m) = \iint V(u,v) e^{-2\pi i(ul+vm)} du dv$$

  - u,v are spatial frequencies in the E-W and N-S directions, and are the projected baseline lengths measured in units of wavelength, B/λ
  - l,m are direction cosines relative to a reference position in the E-W and N-S directions

Some 2D FT pairs

Image | Visibility amp

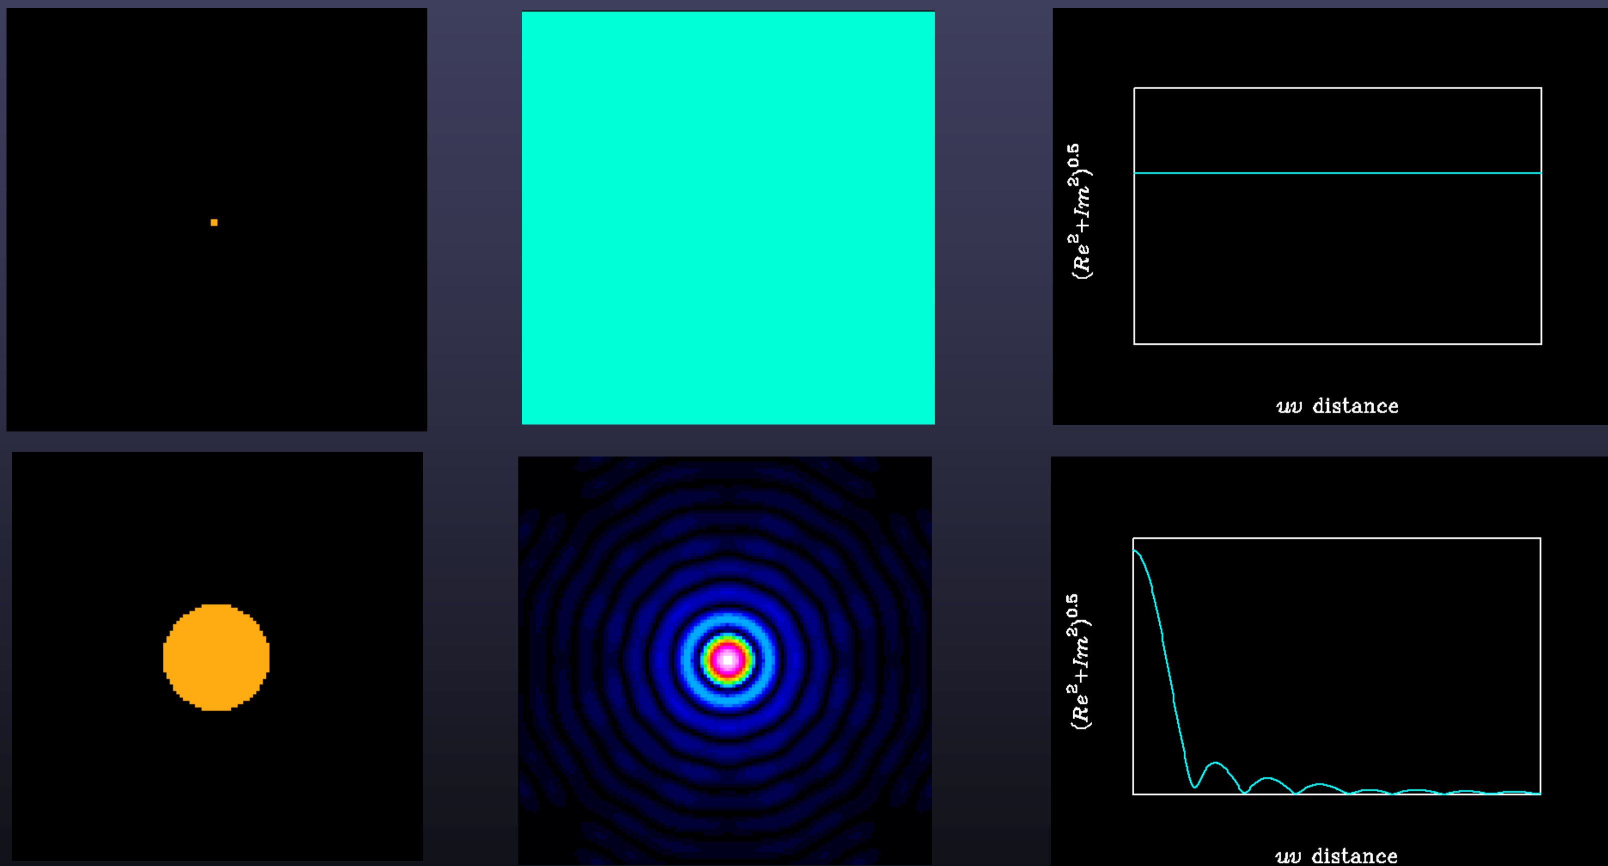

Some 2D FT pairs

Image | Visibility amp

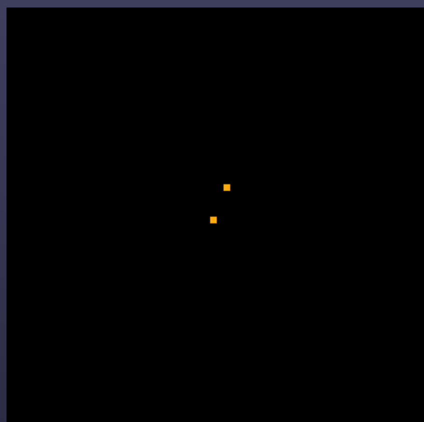

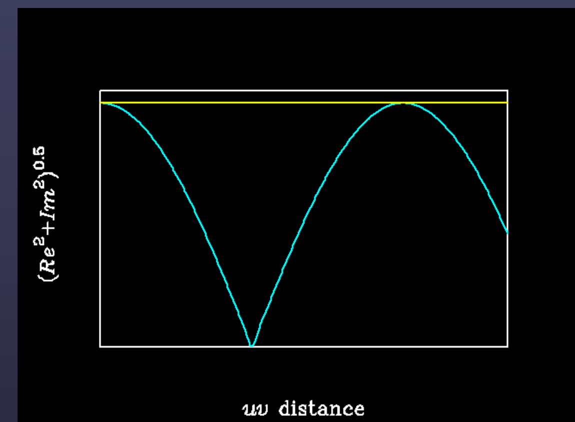

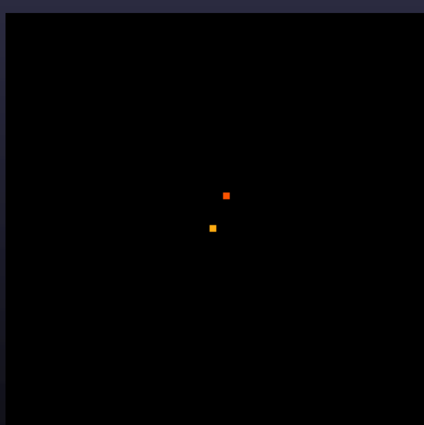

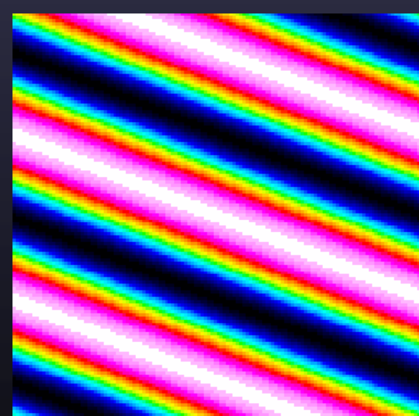

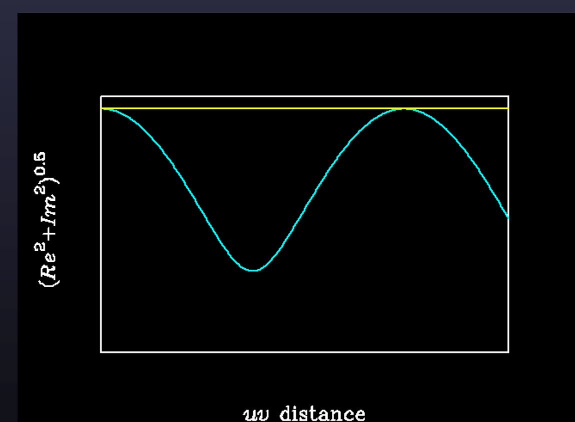

Some 2D FT pairs

Image

Visibility amp

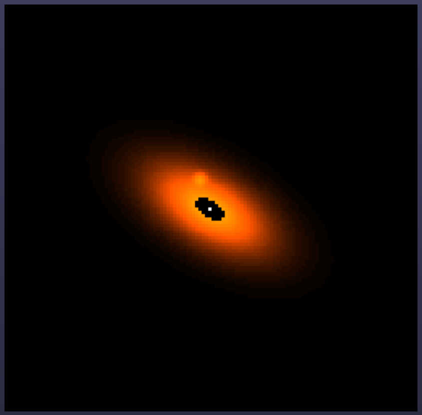

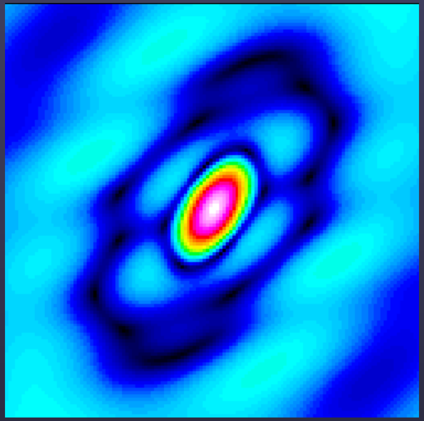

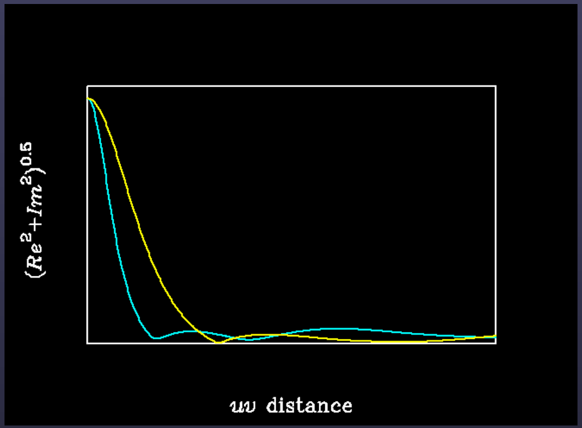

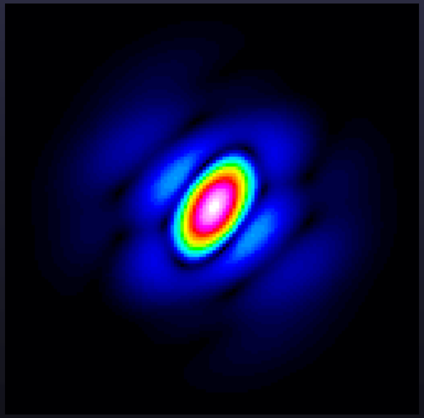

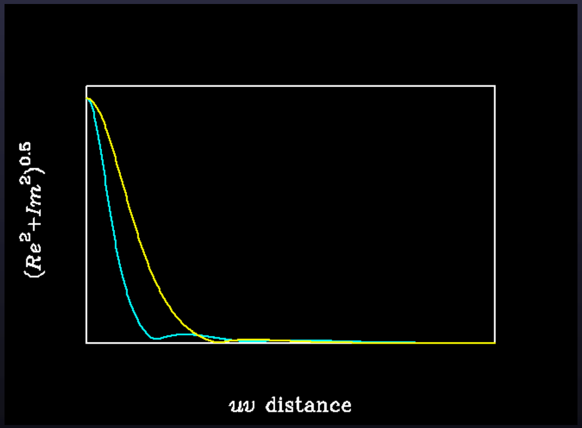

Sampling in the *uv* plane

- Observed sky distribution, $I(l,m)_{obs}$, is the convolution of the true sky distribution, $I(l,m)_{true}$, and the point spread function. This is equivalent to sampling the true visibility function, $V(u,v)$, with some sampling function, $S(u,v)$, in the *uv* plane:

$$I_{obs}(l,m) = I_{true}(l,m) * PSF = \iint S(u,v) V(u,v) e^{-2\pi i(ul+vm)} du dv$$

- For a single telescope, $S(u,v)$ is continuous, and in the absence of seeing is the autocorrelation function of the aperture

Sampling in the *uv* plane

- For an interferometer, $S(u,v)$ is discrete

$$S(u,v)=\sum_k \delta(u_k,v_k)$$

- The measured visibility is the true visibility multiplied by the sampling function

$$V_{meas}(u,v)=S(u,v)V_{true}(u,v)$$

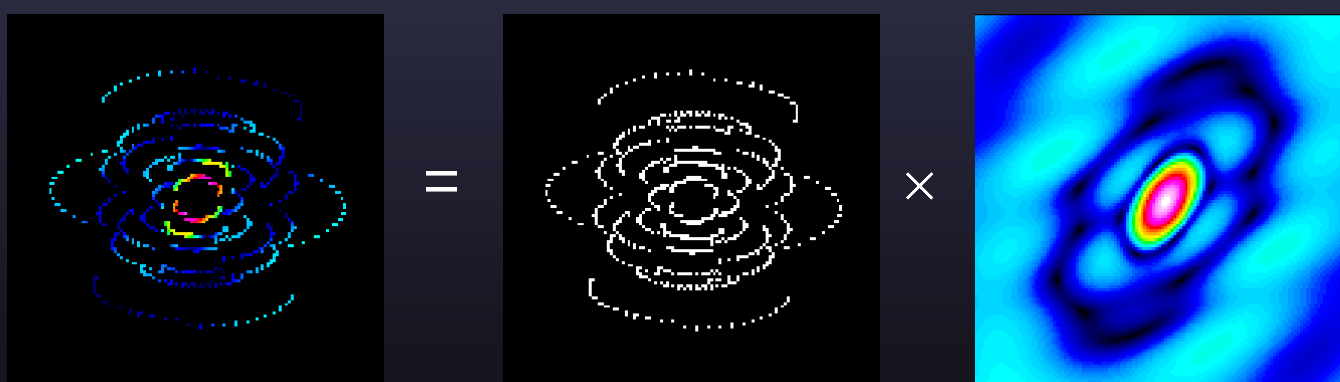

Sampling in the *uv* plane

- Aperture synthesis:
  - Ideally want to make sampling as complete as possible, to synthesize an aperture of diameter B_{max}

Telescope locations

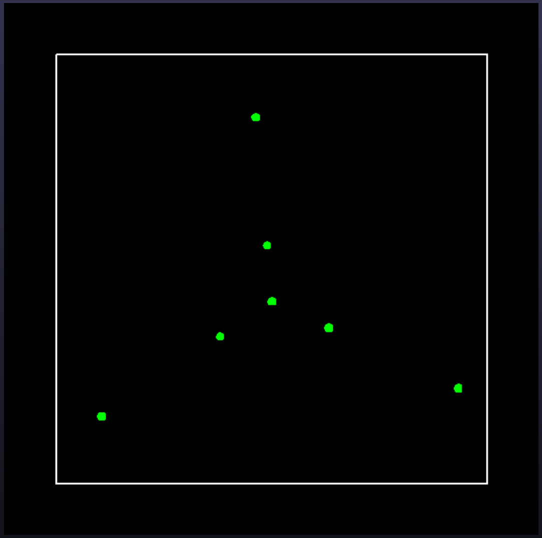

instantaneous *uv* coverage

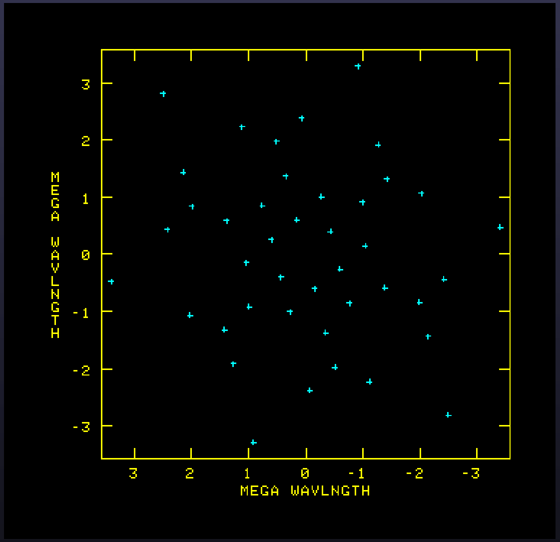

PSF

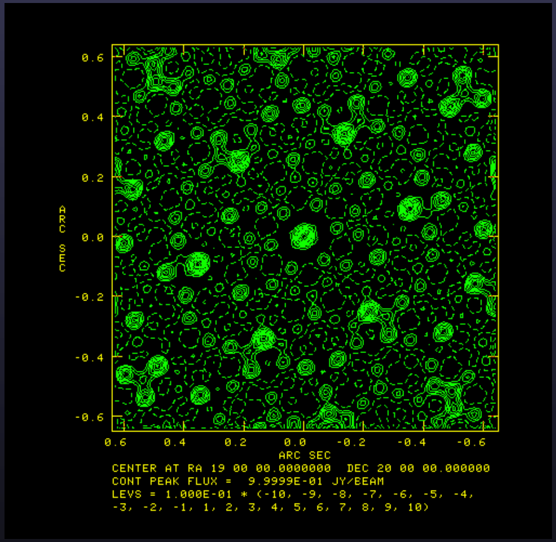

The resulting PSF

- How “nice” the resulting PSF looks depends on how well the *uv* plane is sampled: VLA snapshots

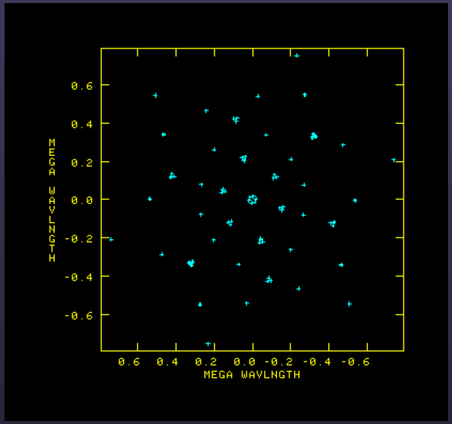

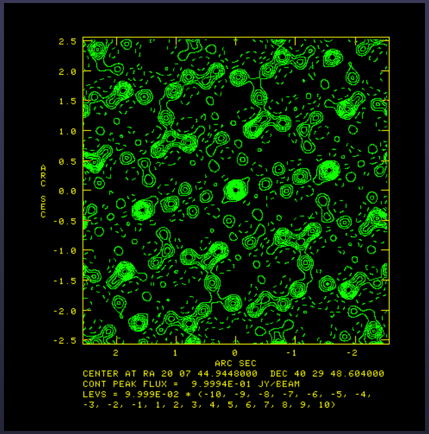

← 3 antennas on each arm

8 antennas on each arm →

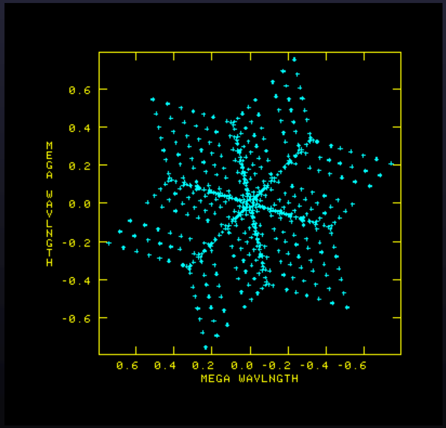

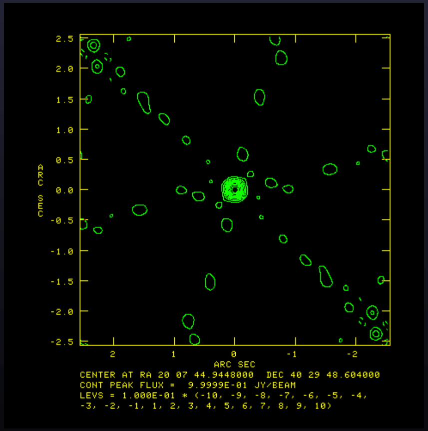

Earth rotation synthesis

- Earth rotation synthesis:
  - We can help to fill in the *uv* plane by making use of the rotation of the Earth. A fixed baseline between telescope 1 and telescope 2, B_{12}, tracking a source from rise to set, will have a changing *projected baseline*, B_{proj}, in the direction of the source, and will trace out *two* arcs in the *uv* plane: one for baseline 1-2, and one for 2-1
  - There are two arcs because the visibility is Hermitian:

$$V(-u,-v) = V^{*}(u,v)$$

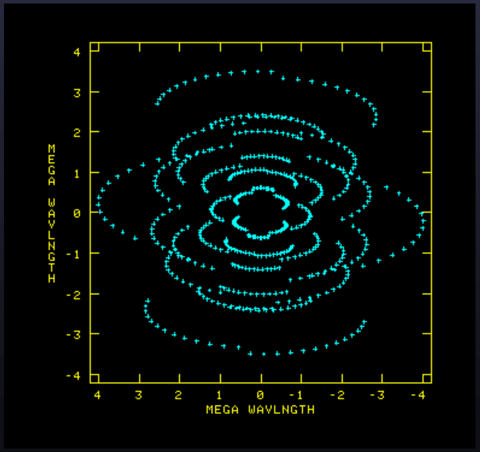

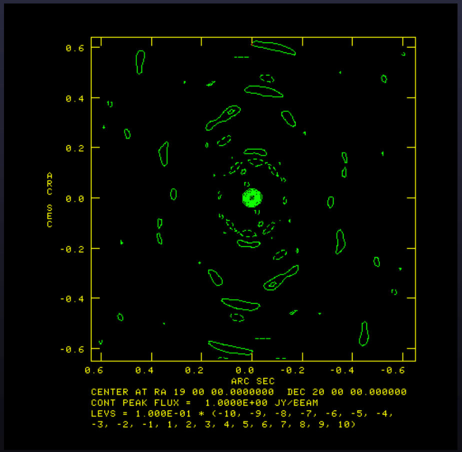

Convolution with the PSF; terminology

- The image obtained from the FT of the sampled visibility is the *dirty image*

$$I_{dirty}(l,m) = FT[S(u,v)V_{true}(u,v)]$$

- The dirty image, I_{dirty}, is the convolution of the true image, I_{true}, and the *dirty beam* (PSF), $B = FT(S)$

$$I_{dirty} = B * I_{true}$$

- [In practice, $I_{dirty} = B*(I_{true}+I_N)$, where $I_N = FT(Vis.\ Noise)$]
- To recover I_{true}, we must deconvolve B from I_{dirty}
- Note: we can do this because S, and therefore B, is well-defined

The dirty image

B

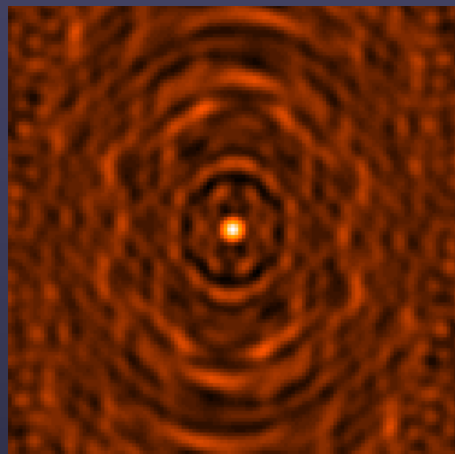

FT

$\Leftrightarrow$

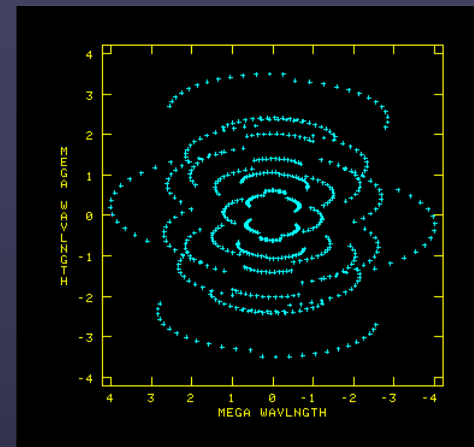

S

$*$

I_{true}

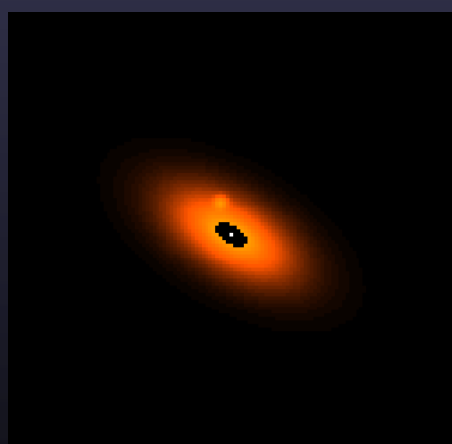

$\rightarrow$

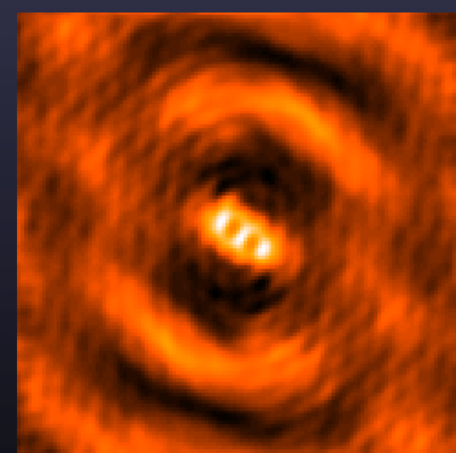

I_{dirty}

The missing information

- Not all parts of the *uv* plane are sampled
- Central hole for $u < u_{min}$ and $v < v_{min}$:
  - Total integrated flux is not measured

$$V(u=0, v=0) = \iint I_{dirty}(l,m)\,dl\,dm = 0$$

  - Upper limit on the largest scale in the image plane
- No measurements for $u > u_{max}$ and $v > v_{max}$:
  - Size of the main lobe of the PSF (the resolution) is finite
- Holes in the *uv* plane:
  - Contribute to the sidelobes of the PSF

The missing information

- Although the total flux is not measured, the flux for scales corresponding to the Fourier components between u_{min} and u_{max} can be measured
- In the presence of extended emission, the observations must be designed keeping in mind:
  - Required resolution ⇒ maximum baseline
  - Largest scale to be reliably reconstructed ⇒ minimum baseline

Recovering the missing information

- To recover information beyond the maximum baseline requires extrapolation (unconstrained)
- Recovering information corresponding to the central hole is possible, but need extra information (e.g., measure total flux using a large single telescope)
- Information corresponding to the *uv* holes requires interpolation

Deconvolution of the PSF in the image plane = interpolation in the visibility plane

- Non-linear methods required

Recovering the missing information

- Note: there is an infinite number of sky distributions consistent with the measurements, need to provide some constraints to the interpolation
- We can assume:
  - The sky brightness is positive (but there are exceptions)
  - The sky is a collection of point sources (weak assertion)
  - The sky could be smooth
  - The sky is mostly blank
- Non-linear deconvolution algorithms search for a model image, I_{model}, such that the residual visibilities $V_{resid} = V_{model} - V_{meas}$ are minimized subject to the constraints given by the assumptions

Practical aspects: overview

- The rest of this lecture addresses some practical aspects of synthesis imaging: choices you will probably be asked to make by any piece of synthesis imaging software
  - FFTs and the need to grid the uv data
  - Forming the dirty beam: weighting
  - An example of a deconvolution algorithm: Clean
  - Finite support: the role of boxes
  - Choosing the image and pixel sizes

Making the dirty image

- The Fast Fourier Transform (FFT) is used for efficient Fourier transformations. However, it requires a regularly-spaced grid of data
- Measured visibilities are irregularly sampled (along tracks in the uv plane)
- Visibilities must be interpolated onto a regular grid using a suitable function

Dirty beam: properties

- The PSF is a weighted sum of cosines corresponding to the measured Fourier components:

$$B(l,m) = \frac{\sum_k w_k \cos(u_k l + v_k m)}{\sum_k w_k}$$

- The visibility weights, w, are also gridded onto a regular grid, FFTed, and used to compute the dirty beam
- The peak of the dirty beam is normalized to unity
- The 'main lobe' has a size of order $dx \sim 1/u_{max}$ by $dy \sim 1/v_{max}$ – this is the resolution of the instrument, or 'clean beam'

NRAO NSF Associated Universities, Inc.

Dirty beam: properties

- Sidelobes extend indefinitely
- Close-in sidelobes are controlled by the envelope of the *uv* coverage: e.g., if the envelope is a circle, the sidelobes near the main lobe must be similar to the FT of a circular disk

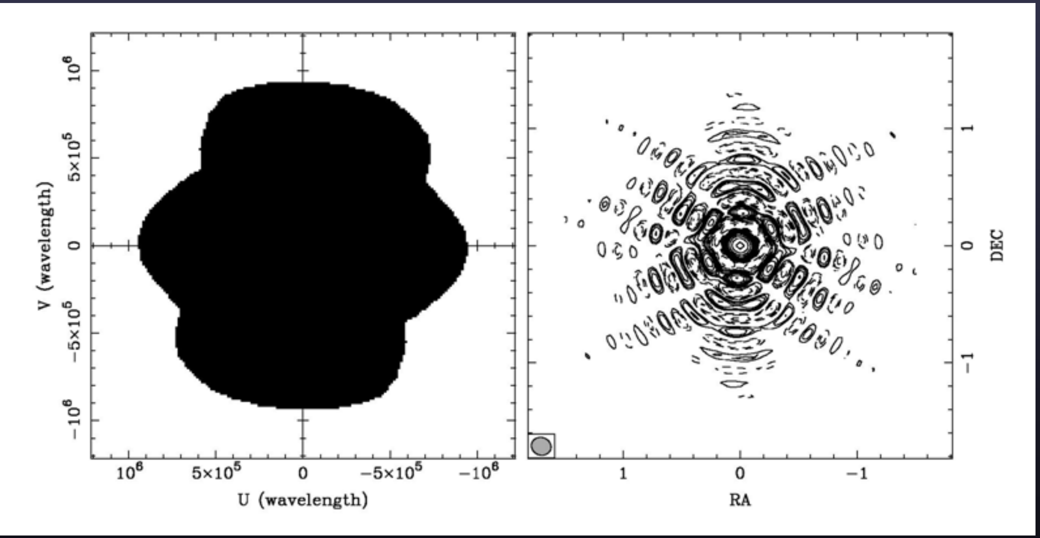

Forming the dirty beam: weighting

- The weighting function , w_k, can be chosen to modify the sidelobe structure of the beam

$$B(l,m) = \frac{\sum_k w_k \cos(u_k l + v_k m)}{\sum_k w_k}$$

- ‘Natural weighting’: $w_k = 1/\sigma_k^2$ where σ_k^2 is the rms noise of the k^{th} gridded visibility
  - Gives the best rms noise across the image
  - Smaller baselines (large spatial scales) have higher weights
  - The effective resolution is worse than the inverse of the longest baseline

Forming the dirty beam: weighting

- Uniform weighting: $w_k = 1/\rho(u_k,v_k)$ where $\rho(u_k,v_k)$ is the density of uv points in the k^{th} cell
  - Short baselines (large scale features in the image) are weighted down
  - Relatively better resolution
  - Increased rms noise

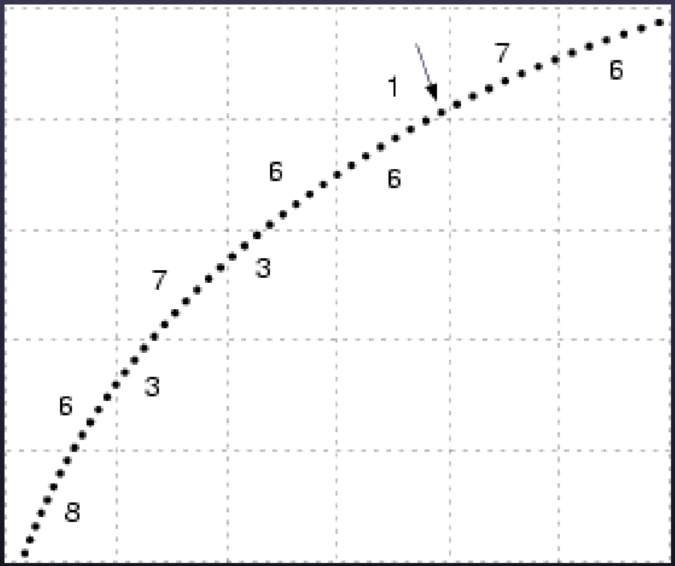

Forming the dirty beam: weighting

- 'Robust' or 'Briggs' weighting: $w_k = 1/[S^2\rho(u_k,v_k) + \sigma_k{}^2]$
  - $S^2 = (5\cdot10^{-R})^2/\rho(u_k,v_k)$ is a parameterized filter that allows continuous variation between optimal resolution (uniform weighting) and optimal noise properties (natural weighting) by varying the robust parameter, R

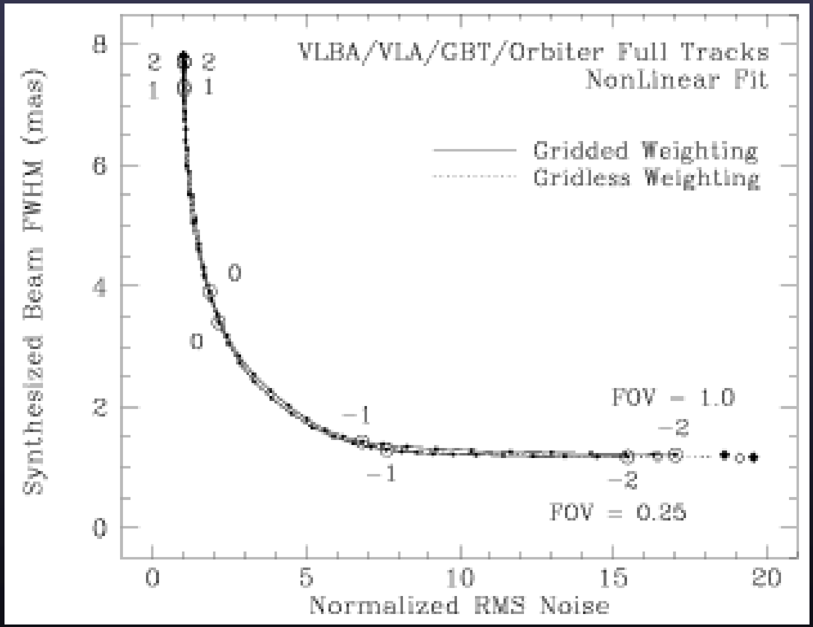

Examples of weighting: VLA

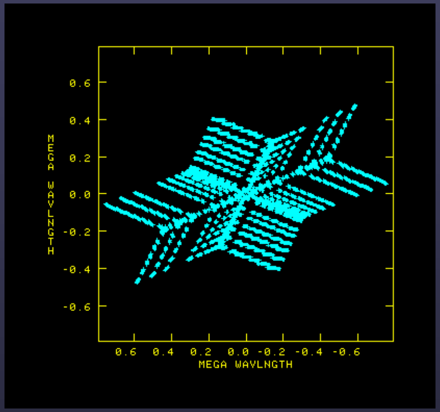

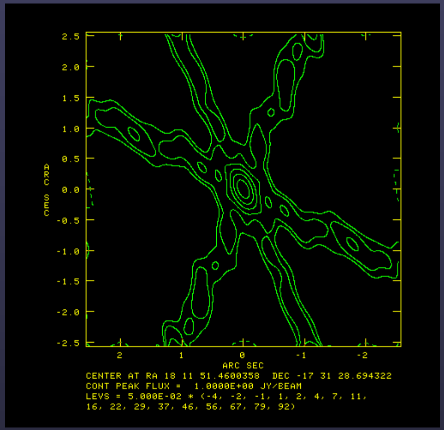

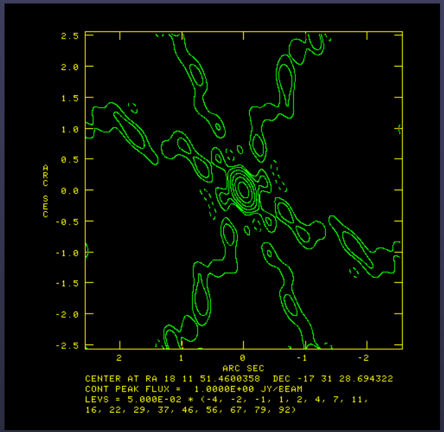

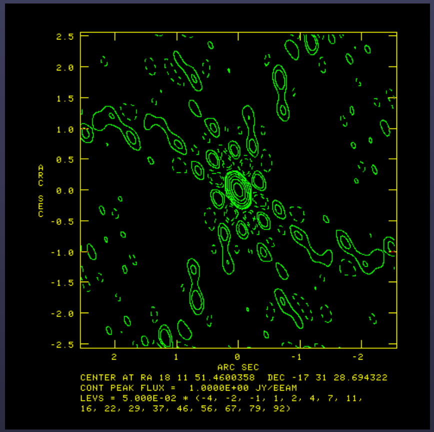

(*uv* coverage)	Natural	Robust	Uniform
Clean beam:	0.56″×0.29″	0.48″×0.26″	0.41″×0.22″
Rms noise:	1.0	1.1	2.0

⇒ tune resolution and sensitivity to suit your science

Examples of weighting: sparse *uv* coverage

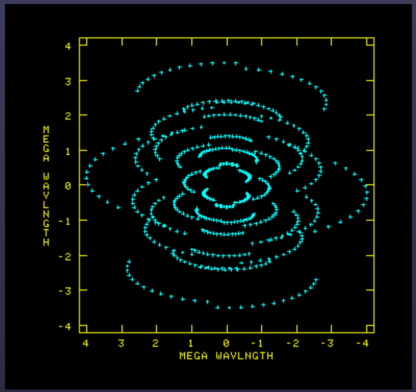

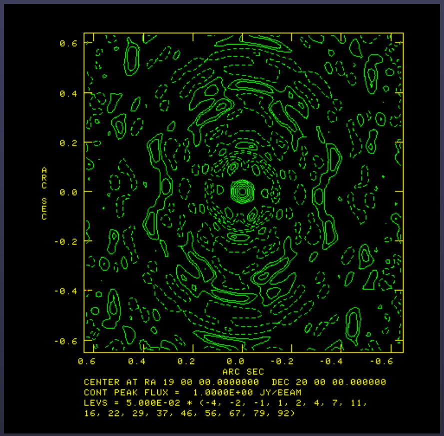

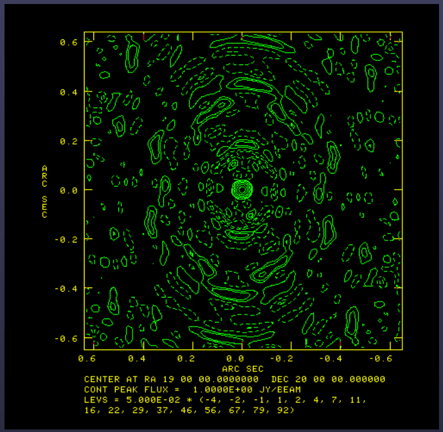

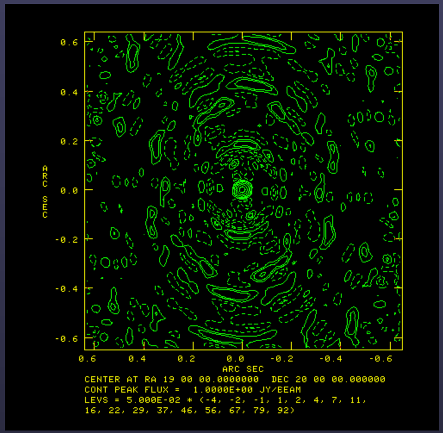

(*uv* coverage)	Natural	Robust	Uniform
Clean beam:	0.049″×0.048″	0.046″×0.045″	0.045″×0.044″
Rms noise:	1.0	1.03	1.08

⇒ natural and uniform weighting similar for sparse *uv* coverage

Forming the dirty beam: tapering

- The PSF can also be further controlled by applying a tapering function to the weights (e.g., such that the weights smoothly go to zero toward longer baselines)

$$w'_k = T(u_k, v_k) w_k(u_k, v_k)$$

- Bottom line on weighting/tapering:
  - They help a bit, but imaging quality is limited by finite sampling of the uv plane

The Clean deconvolution algorithm (Högbom 1974)

- Various deconvolution algorithms are available; Clean is an example of a scale-less algorithm
- Assume the sky is composed of point sources, and is mostly blank; then:
  1. Search for the peak in the dirty image
  2. Subtract a fraction g (the loop gain) of the PSF from the position of the peak (typically $g \sim 0.05–0.1$)
  3. Add g times the peak to a single pixel in the model image
  4. If residuals are not noise-like, go to 1
  5. Smooth the model image by an estimate of the main lobe of the PSF (the clean beam) and add the residuals to make the 'restored image'

The Clean deconvolution algorithm

- Stopping criteria: either specify a maximum number of iterations or the maximum in the residual image (some multiple of the expected noise is typical)
- Search space can be constrained by user-defined windows
- Ignores the coupling between pixels (extended emission)

Clean example: model image

- Model source as the sum of many point sources

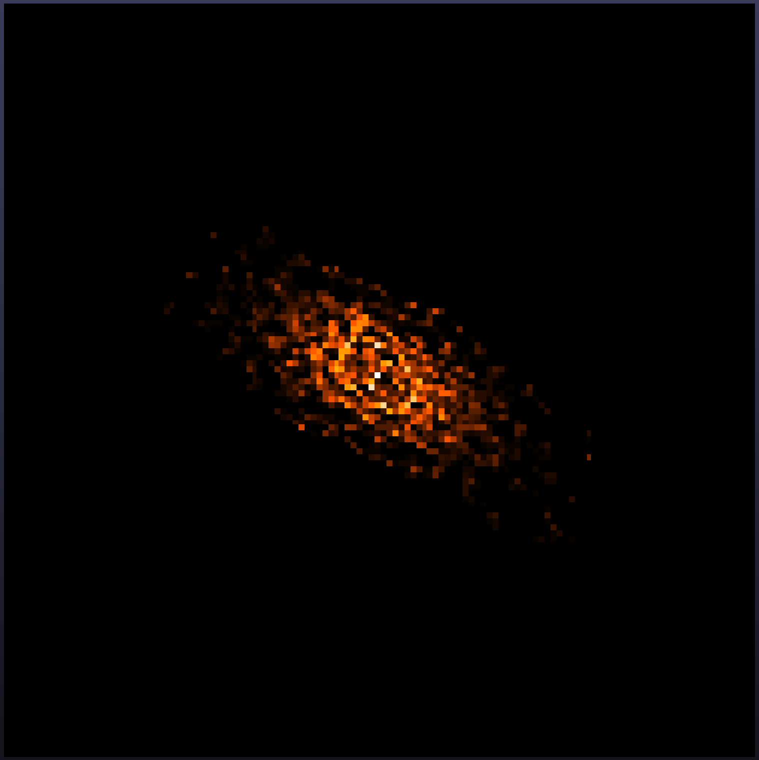

Clean example: residual image

- Subtract (point sources $\times$ PSF) from dirty image to give residual image

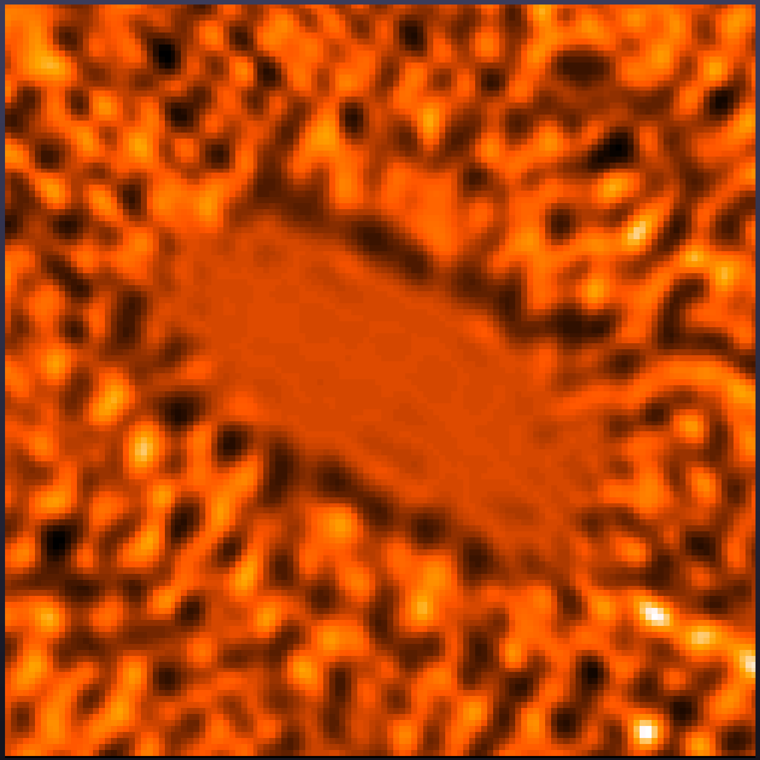

Clean example: restored image

- Smooth the model image by the clean beam and add the residuals to form the restored image

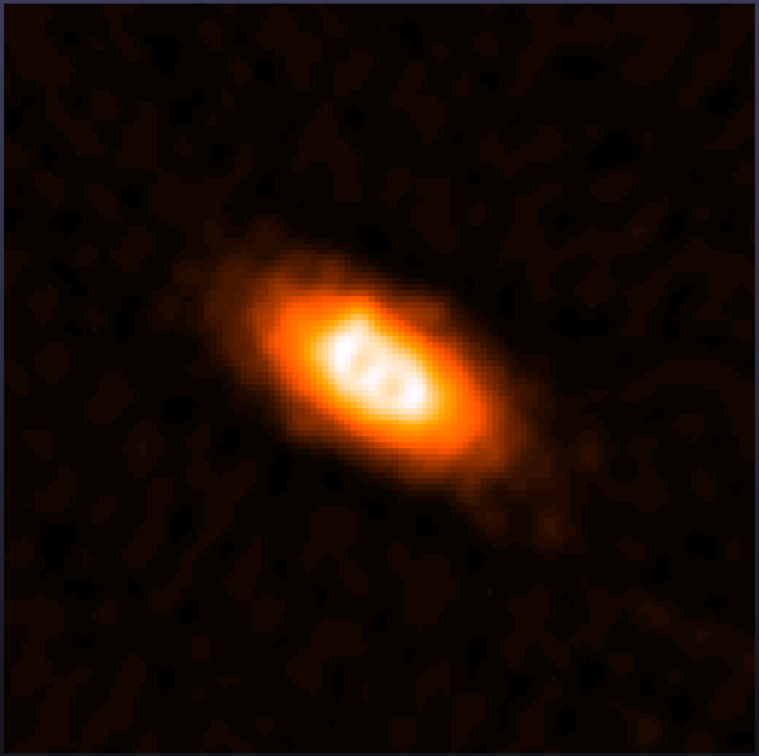

Comparison with I_{true}

- Convolve I_{true} with clean beam for comparison:

$I_{restored}$

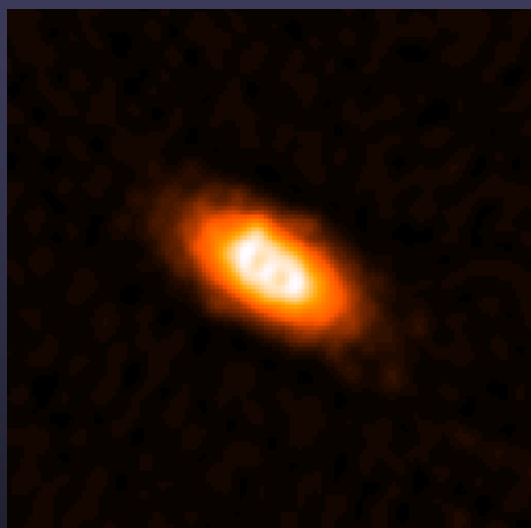

I_{true} * (clean beam)

	$I_{restored}$	I_{true} * (clean beam)
F_{peak} =	7.22	7.65
F_{int} =	120	147

Clean example: visibilities

V_{true}

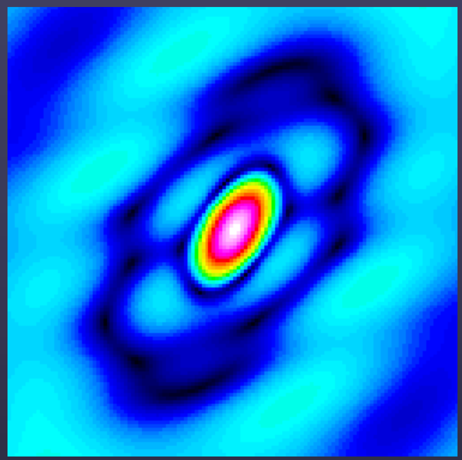

V_{meas}

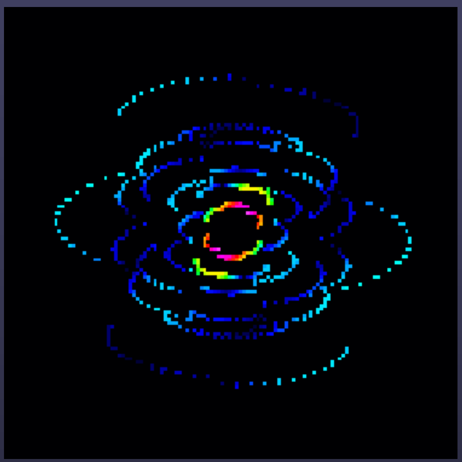

V_{model}

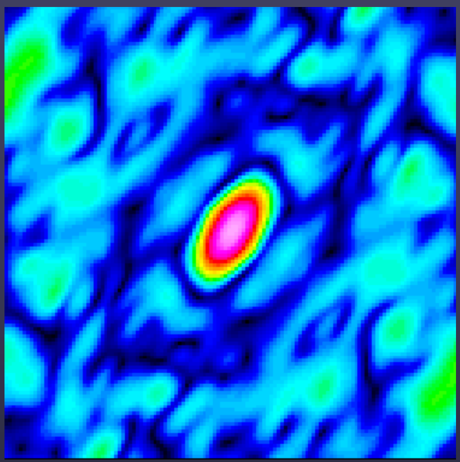

$V_{model} – V_{meas}$

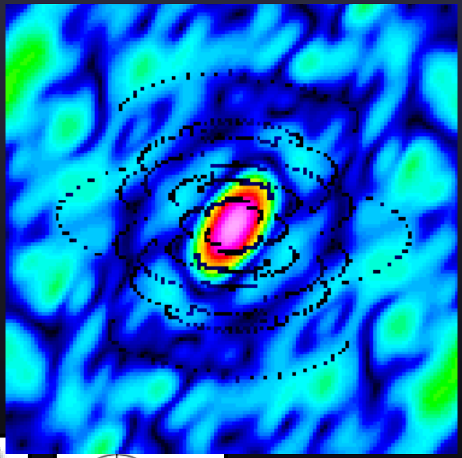

$V_{true} – V_{model}$

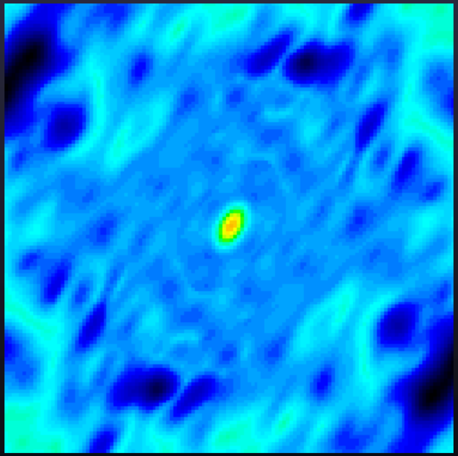

Clean can do a good job of reproducing V_{true} between u_{min} and u_{max}, but generally underestimates the total flux if there is unsampled extended emission

Finite support: the role of boxes

- Limit the search for components to only parts of an image
  - A way to regularize the deconvolution process
- Useful for small numbers of visibilities (VLBI / optical / snapshots using large-N arrays)
- Stop when Cleaning within the boxes has no global effect

With boxes:

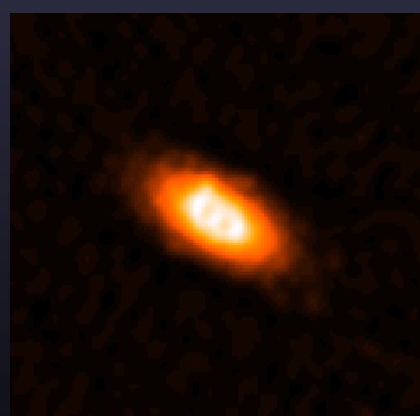

No boxes:

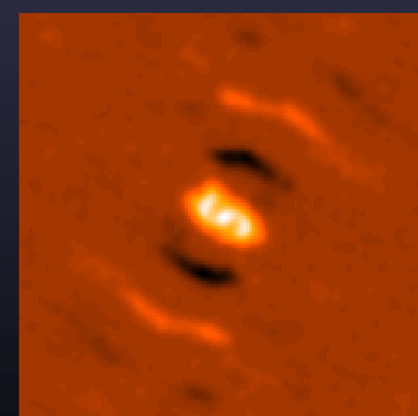

Image and cell sizes

- The size of the cells in the image needs to be chosen so that the main lobe of the dirty beam is at least Nyquist sampled: $\Delta l \leq 1/2u_{max}$, $\Delta m \leq 1/2v_{max}$
- The extent of the dirty image, $l \times m$, is related to the size of the grid cells in the uv plane, through the FT relationship $l = 1/\Delta u$, $m = 1/\Delta v$; if you make the image smaller than $1/\Delta u \times 1/\Delta v$ there may be aliasing
- The size of the image should be big enough to include the largest spatial scale on which there is measured flux (shortest uv spacing)
- If the image is not big enough, sidelobes from sources outside the image may be included and will not be deconvolved properly
- But also: if you have N independent visibilities you can only sensibly image of order $\sim N$ independent beam areas

Final remarks

- Everything I have told you about synthesis imaging assumes that you have visibilities with calibrated amplitudes and phases
- What if you don't?
  - Self-calibration
  - See Chris Haniff's lecture
- There are many other subtleties not covered here; for further reading please see Synthesis Imaging in Radio Astronomy II, ASP Vol. 180 (1998)
- Interferometry and synthesis imaging requires you to think in FT space! This takes practice…