


Radio Astronomy

Lecture 8

The Techniques of Radio Interferometry III: Imaging

Lecturer: Michael Wise (wise@astron.nl)

May 12th, 2017




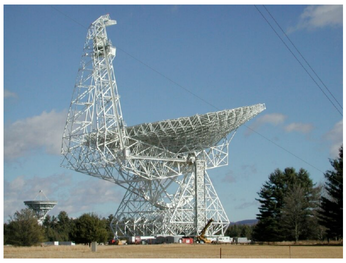

This lecture will be the second of two lectures on some of the analysis techniques used in radio interferometry.
In the previous lecture we focused on the calibration process, here we will discuss the process of making images once the data has been calibrated.
As in the previous lecture, the goal is understand the process but also to develop your intuition about what a "good" image looks like and why.

Outline

- Imaging and Deconvolution
- Image Quality, Noise, Dynamic Range
- Wide-band imaging
- Wide-field imaging
- Mosaicing

Imaging and Deconvolution

Basic Imaging

How do we go from the measurement of the visibility function to images of the sky?

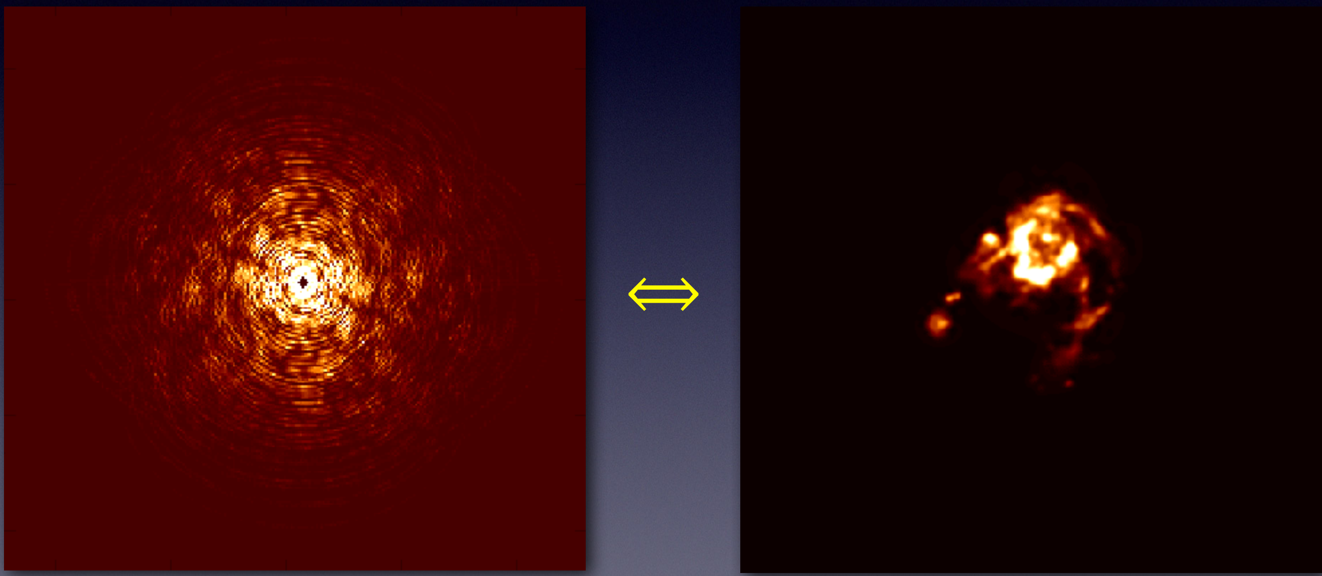

V(u,v) ⇔ *I(x,y)*

In radio astronomy, imaging is process whereby we transform from measured visibilities to images of the sky.
The signatures of the measurement process are encoded in the visibilities we measure.
These signatures are reflected in the resulting surface brightness image.
Much of the imaging process is devoted to correcting for these measurement effects (or trying to do so).

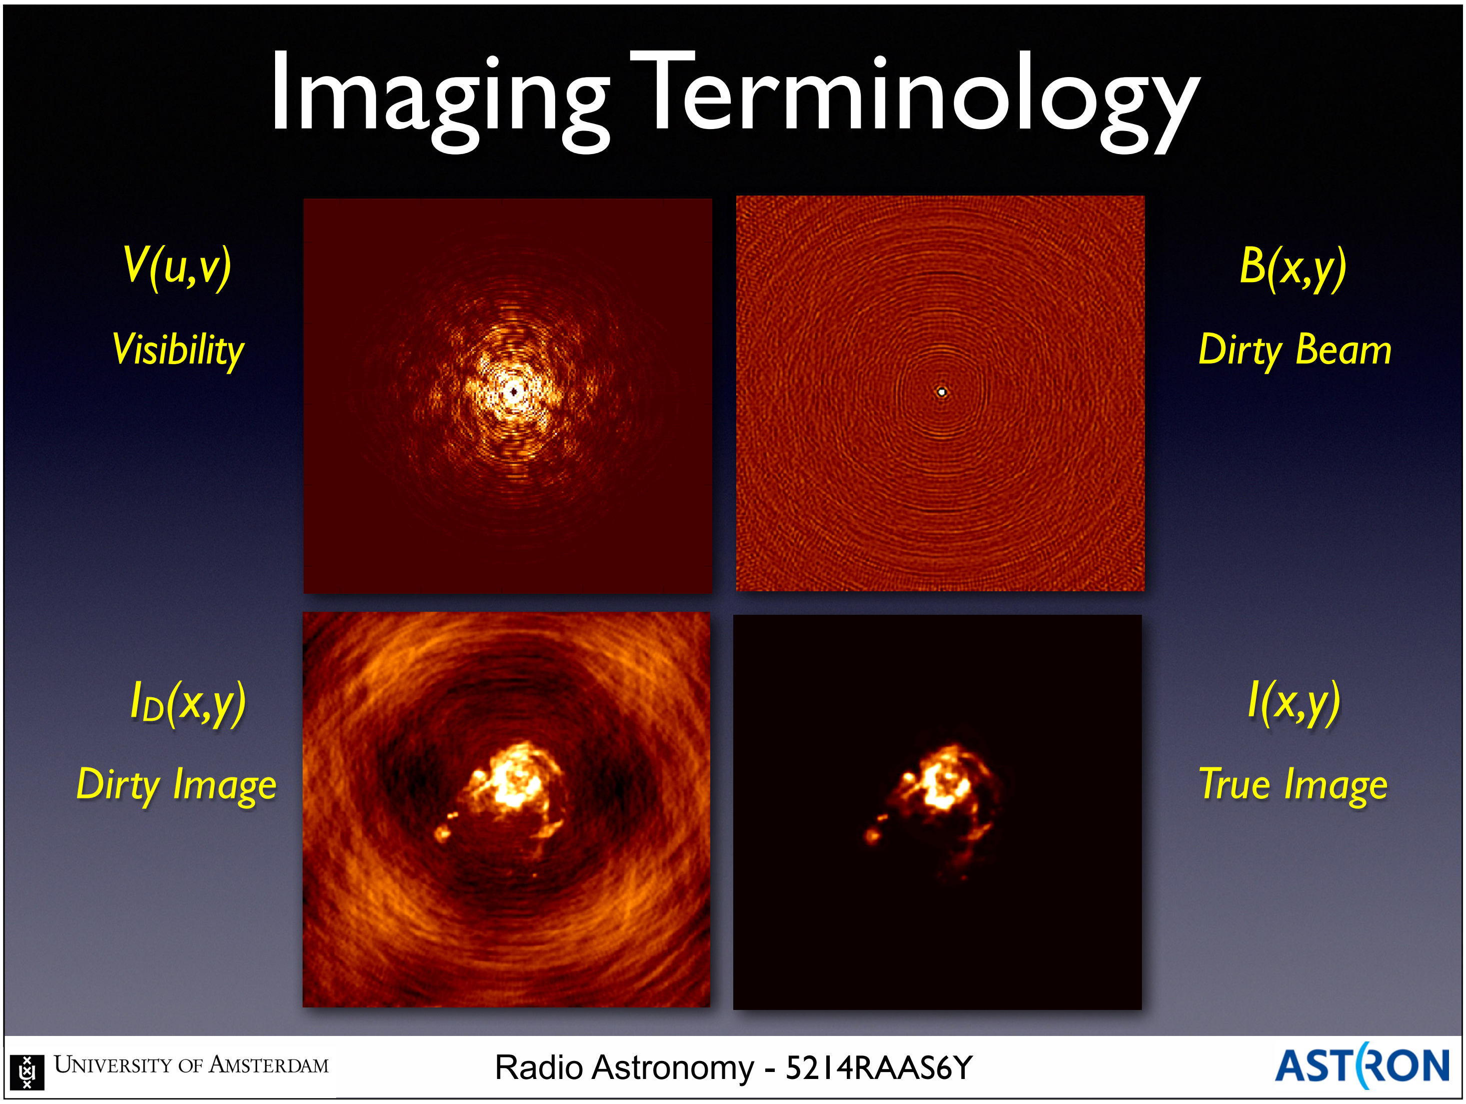

Some basic terms. We measure visibilities directly. How well we sample all the points in the (u,v) plane ultimately determine our image quality.

The dirty beam is the Fourier transform of the (u,v) sampling function. Essentially the point spread function of the telescope.

The dirty image is the convolution of the dirty beam with the true surface brightness of our source.

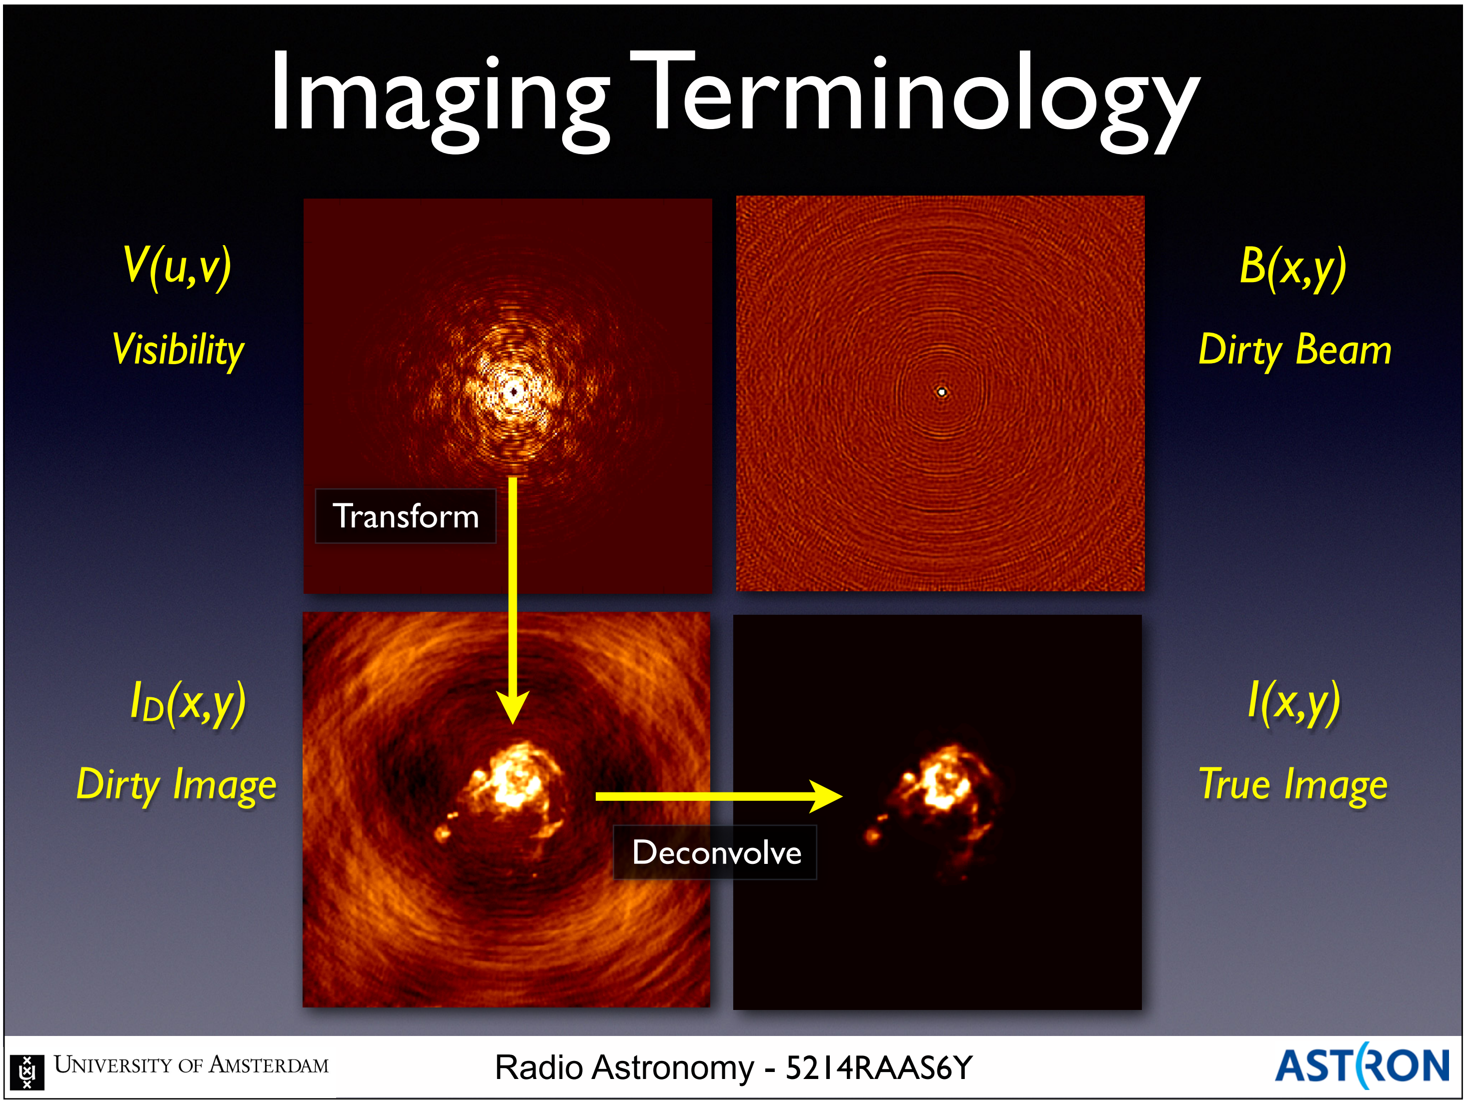

We measure V(u,v) directly. We can calculate B(x,y) because we know what our telescope looks like.
So to derive I(x,y), we must deconvolve the dirty image by the beam model.
Beam models are not always simple! Deconvolution is computationally expensive.

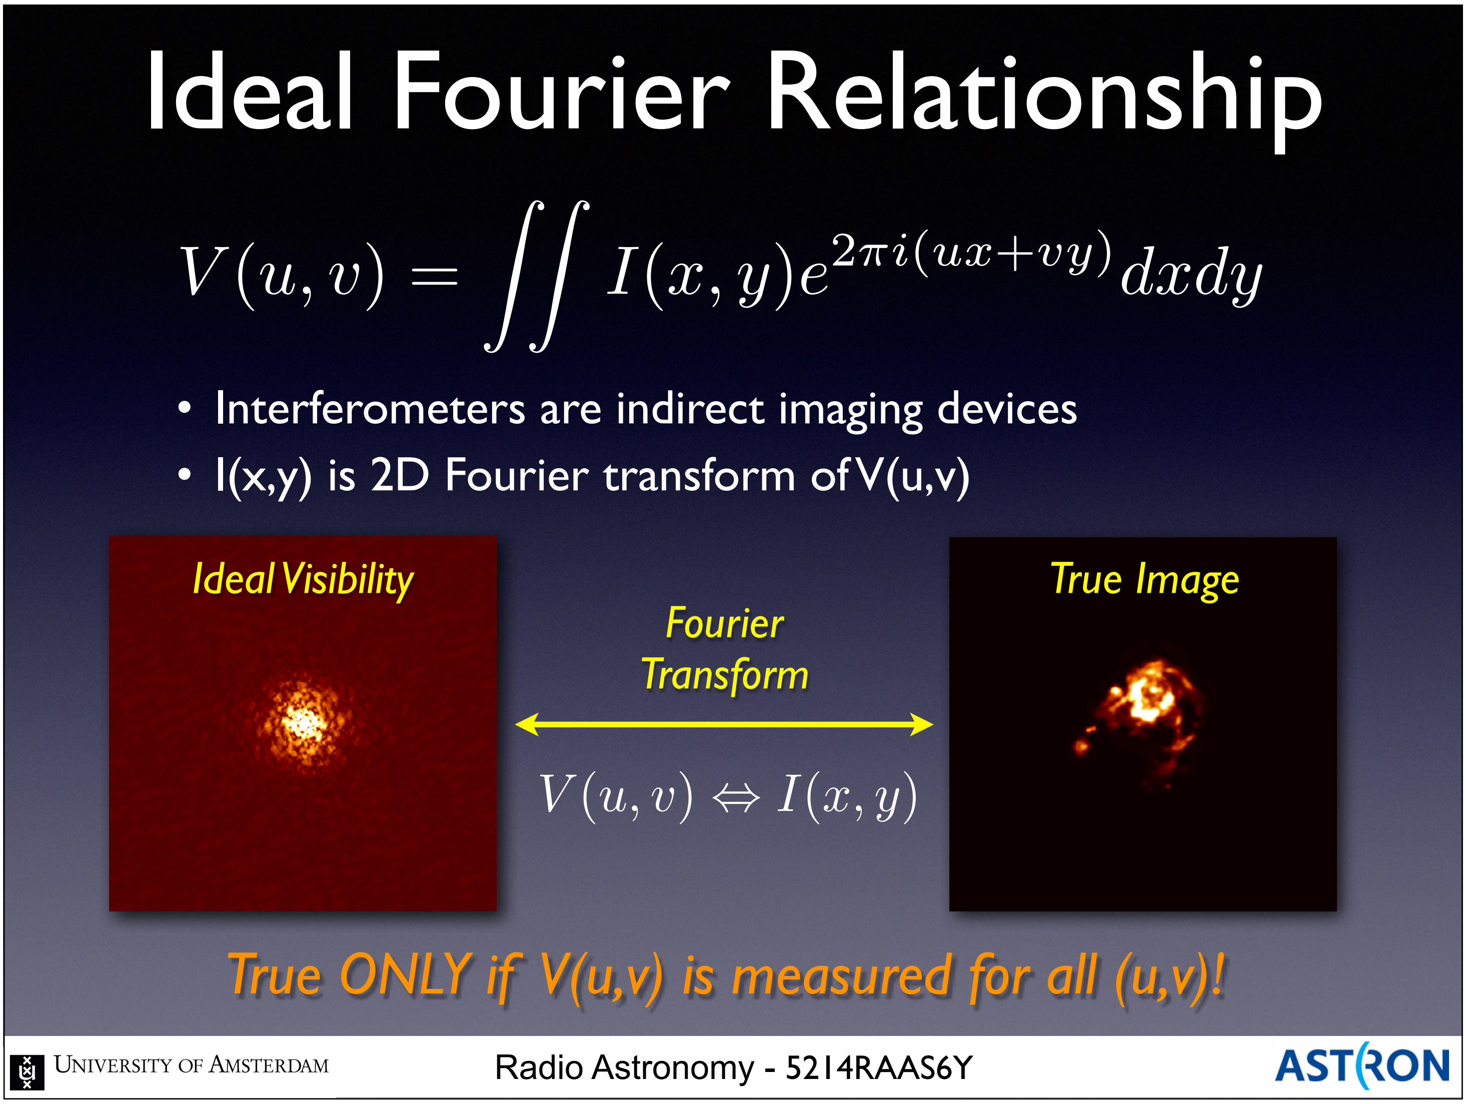

The Fourier relationship only holds exactly if we sample the (u,v) perfectly....and we never do. Every single observation, even with the same telescope, is unique.
Different exposure times, atmospheric effects, antenna problems, etc. make each (u,v) sampling unique.

(u,v) Plane Sampling

- With a limited number of antennas, the uv-plane is sampled at discrete points:

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

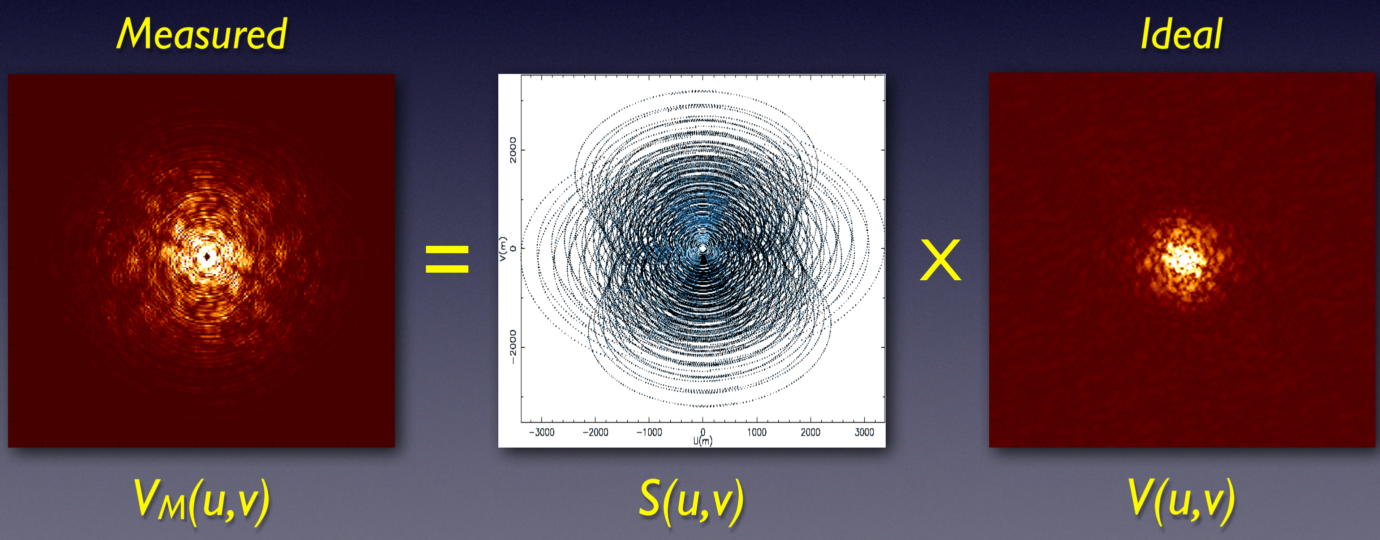

Because we have a finite number of antennas and baselines, we sample the V(u,v) function at discrete points.
The measured visibility function is therefore a limited subset of the true visibility function.
We are missing information about the sky.
So the FT of the measured visibility, or image of the sky we derive, will be missing information about our source.

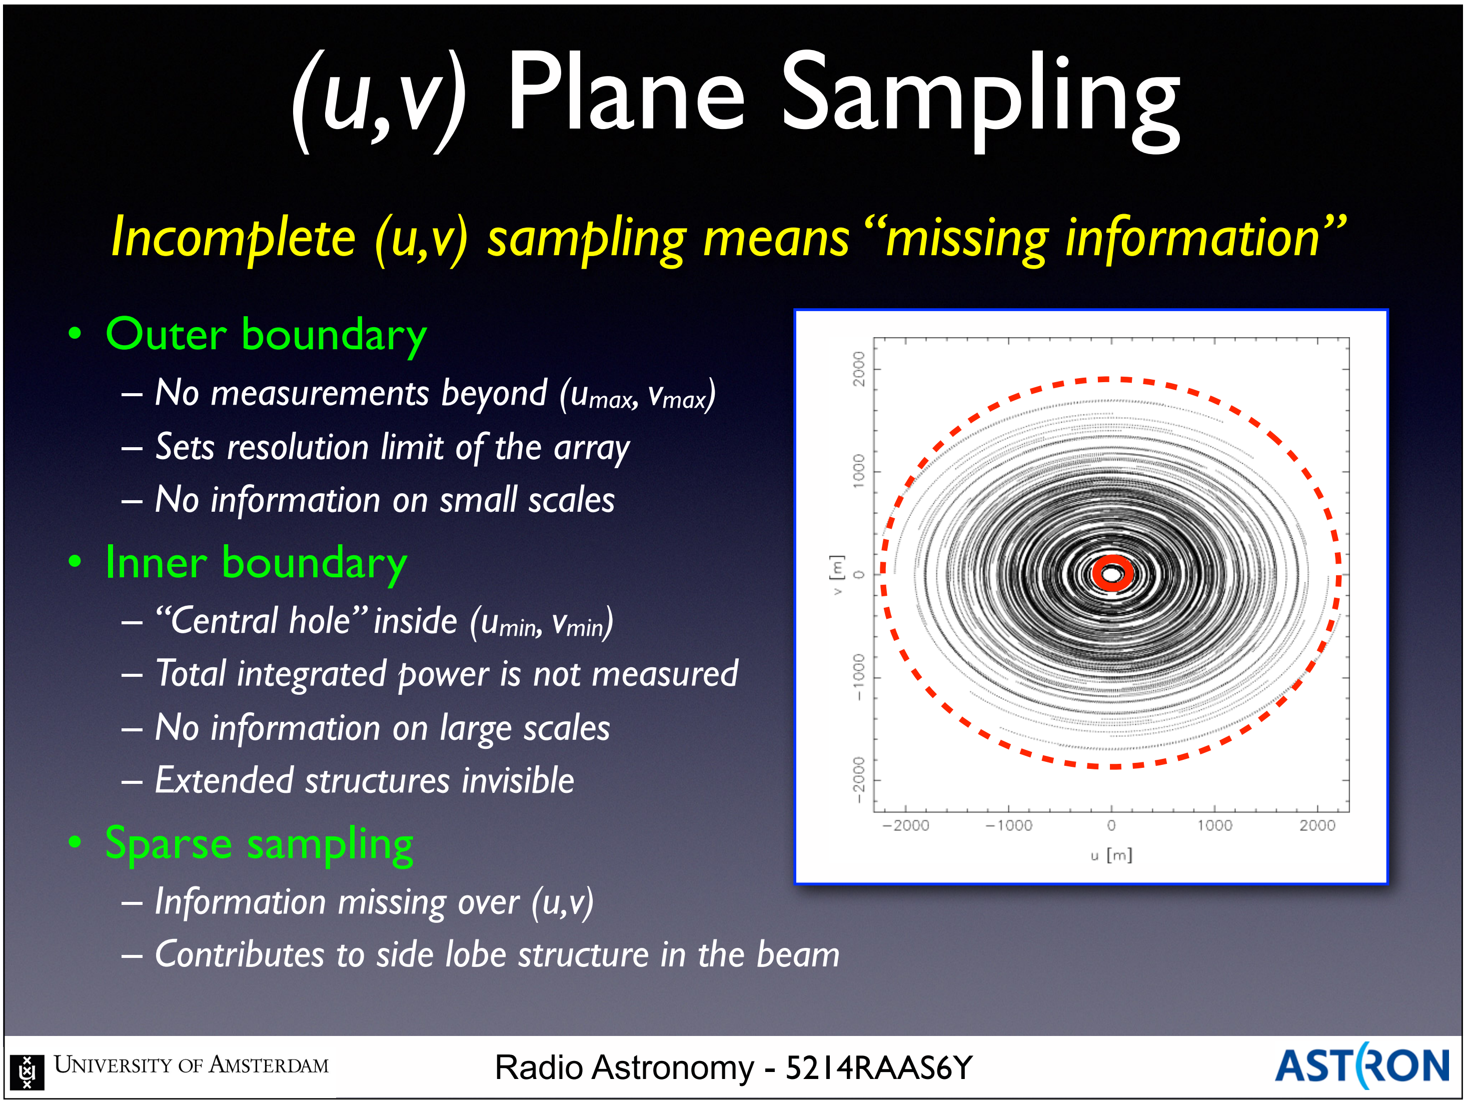

Small (u,v) scales sample large physical scales, and vice versa.
Limited long and short baselines means we only sample the source flux on a limited number of spatial scales.
Downside, it is possible to miss some aspects of a source completely with the wrong choice of baselines.
Upside, can "tune" observations to only see the parts of the source you want to study.

Effect of (*u*,*v*) sampling

- Transforming gives the dirty image $I_D(x,y)$

$$I_D(x,y) = FT^{-1}[V_M(u,v)] = FT^{-1}[S(u,v)V(u,v)]$$

- Using the convolution theorem gives:

$$I_D(x,y) = B(x,y) * I(x,y) \qquad B(x,y) = FT^{-1}[S(u,v)]$$

- Dirty image is convolution of true image with dirty beam $B(x,y)$

To recover $I(x,y)$, we must deconvolve $B(x,y)$ from $I_D(x,y)$

The convolution theorem says that the convolution of two functions is equal to the product of their FT's. Multiplication (division) are computationally simpler than Fourier transforms.

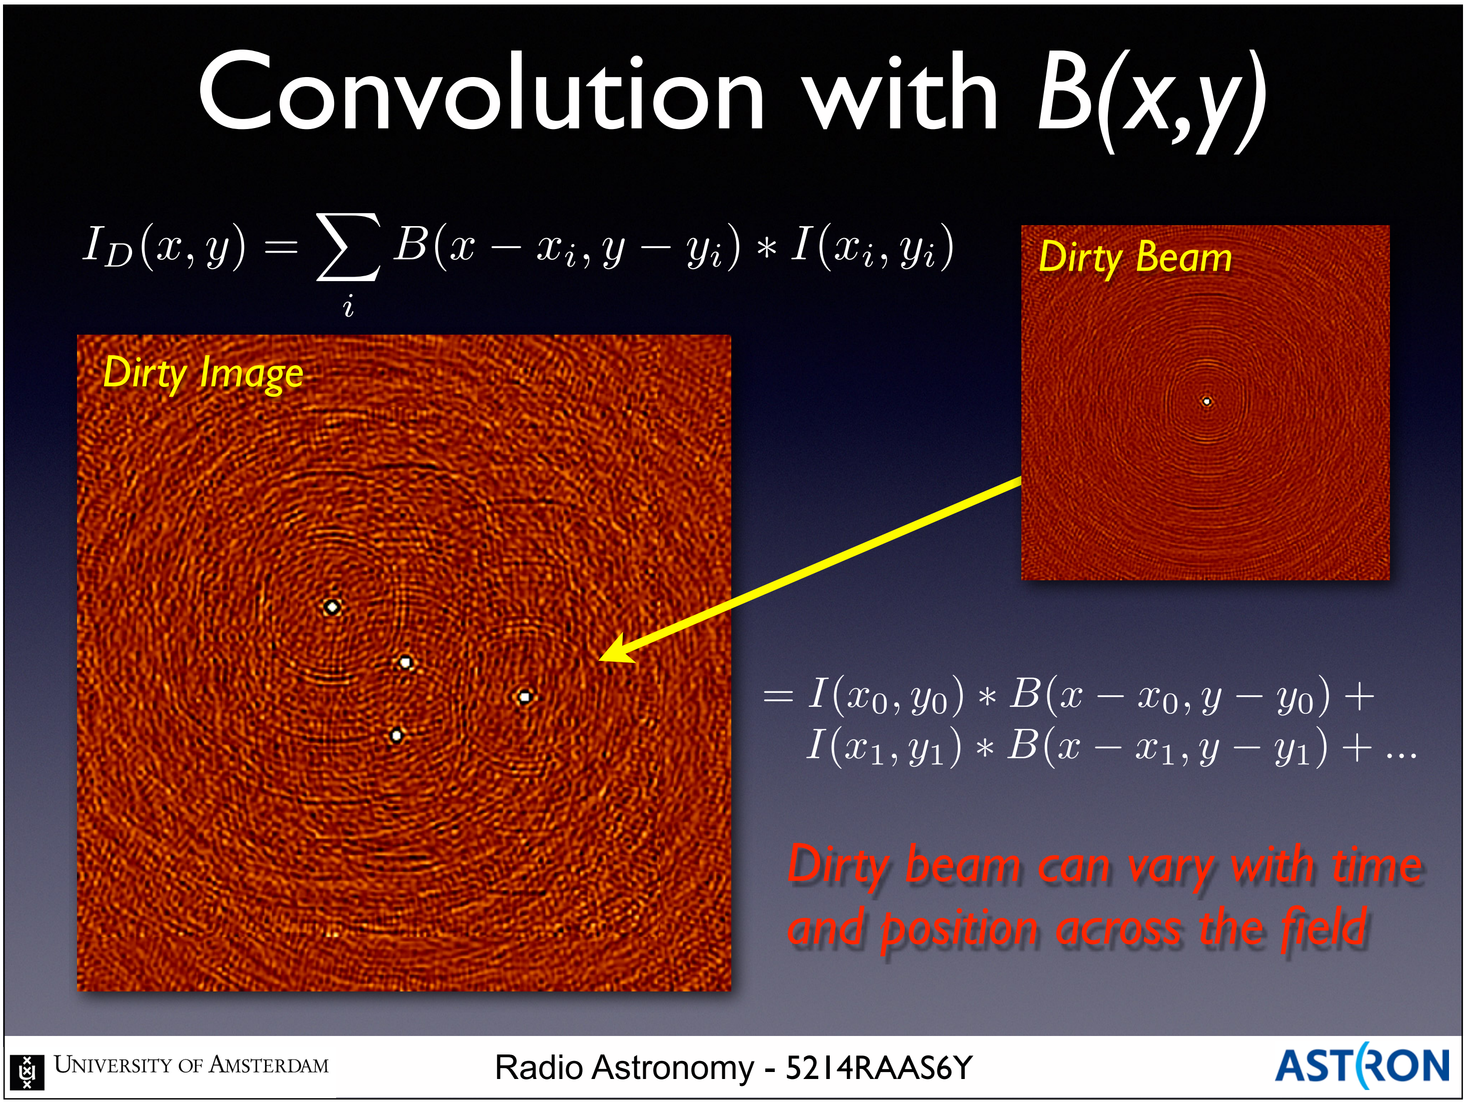

The dirty image we obtain is the sum of all sources on the sky convolved with the dirty beam.
For a set of point sources, this sum is easy to visualize.
Complicated sources can be described as the sum of simpler functions (like delta functions, Gaussians, etc.).
The dirty beam will in general be a function of time, position, and frequency.

Computing the Dirty Image

- "Fourier Transform"
  - *Use Fast Fourier Transform (FFT) algorithm*
  - *Compute scales as ~O(NlogN) for (N x N) image*
  - *FFT requires data on a regularly spaced grid*
  - *Radio arrays sample V(u,v) on irregular grids, so.....*
- "Gridding"
  - *Used to resample V(u,v) for FFT*
  - *Convolutional gridding used to resample $V_M(u,v)$*
  - *Gridding function affects resulting dirty image*
- "Weighting"
  - *Weighting function W_k can be chosen to modify the side lobes*
  - *Different weights $\Longrightarrow$ different B(x,y)*
  - *Can "tune" for resolution or sensitivity*

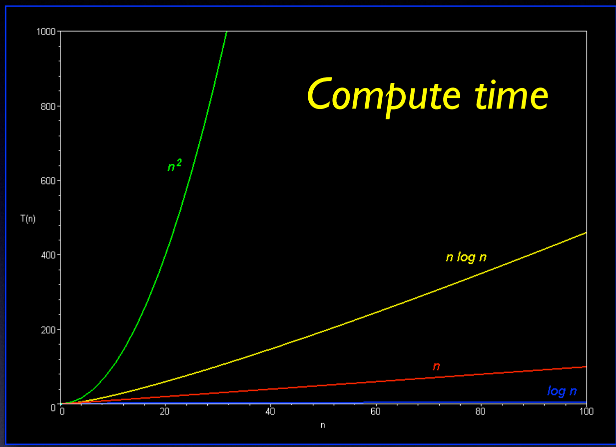

Dirty beam is a weighted sum of the measured Fourier components

$$B(x,y) = \frac{\sum_k W_k cos(u_{kl} + v_{km})}{\sum_k W_k}$$

The Fourier transform is a computationally expensive procedure…goes as N*N for N samples.
In practice we use a faster algorithm called the Fast Fourier Transform (FFT) that scales as N* log N.
To use this algorithm, we must sample the V(u,v) data onto a regularly spaced grid.
The resampling or regridding process can also be computationally expensive.
We can control the properties of the beam by introducing the concept of "weights".
Different weighting schemes allow us to control which baselines contribute and by how much.

Weighting Schemes

Observed image is a weighted-average of the data

$$I_D(x,y) = \frac{\sum_k F^{-1}[W_k(u,v)\ S(u,v)\ V(u,v)]}{\sum_k W_k(u,v)}$$

$$W_k = \frac{1}{\sigma_k^2}$$

- Natural
  - *Maximizes the sensitivity, degrades angular resolution*

$$W_k = \frac{1}{\sigma_k^2\ \rho(u_k, v_k)}$$

- Uniform
  - *Best angular resolution, reduced point source sensitivity*

$$W_k = \frac{(1+s)}{\sigma_k^2\ [1 + \frac{s\rho(u_k,v_k)}{\sigma_k^2}]}$$

- Robust
  - *Smooth, tunable combination of natural and uniform*

$$W_k = \frac{1}{\sigma_k^2} e^{-\frac{(u^2+v^2)}{t^2}}$$

- Tapering
  - *Similar to smoothing, degrades angular resolution*

A number of different weighting schemes are in common use.
Each has their own specific advantages and disadvantages.
You can choose which one you use when making your image, depending on your scientific goals.

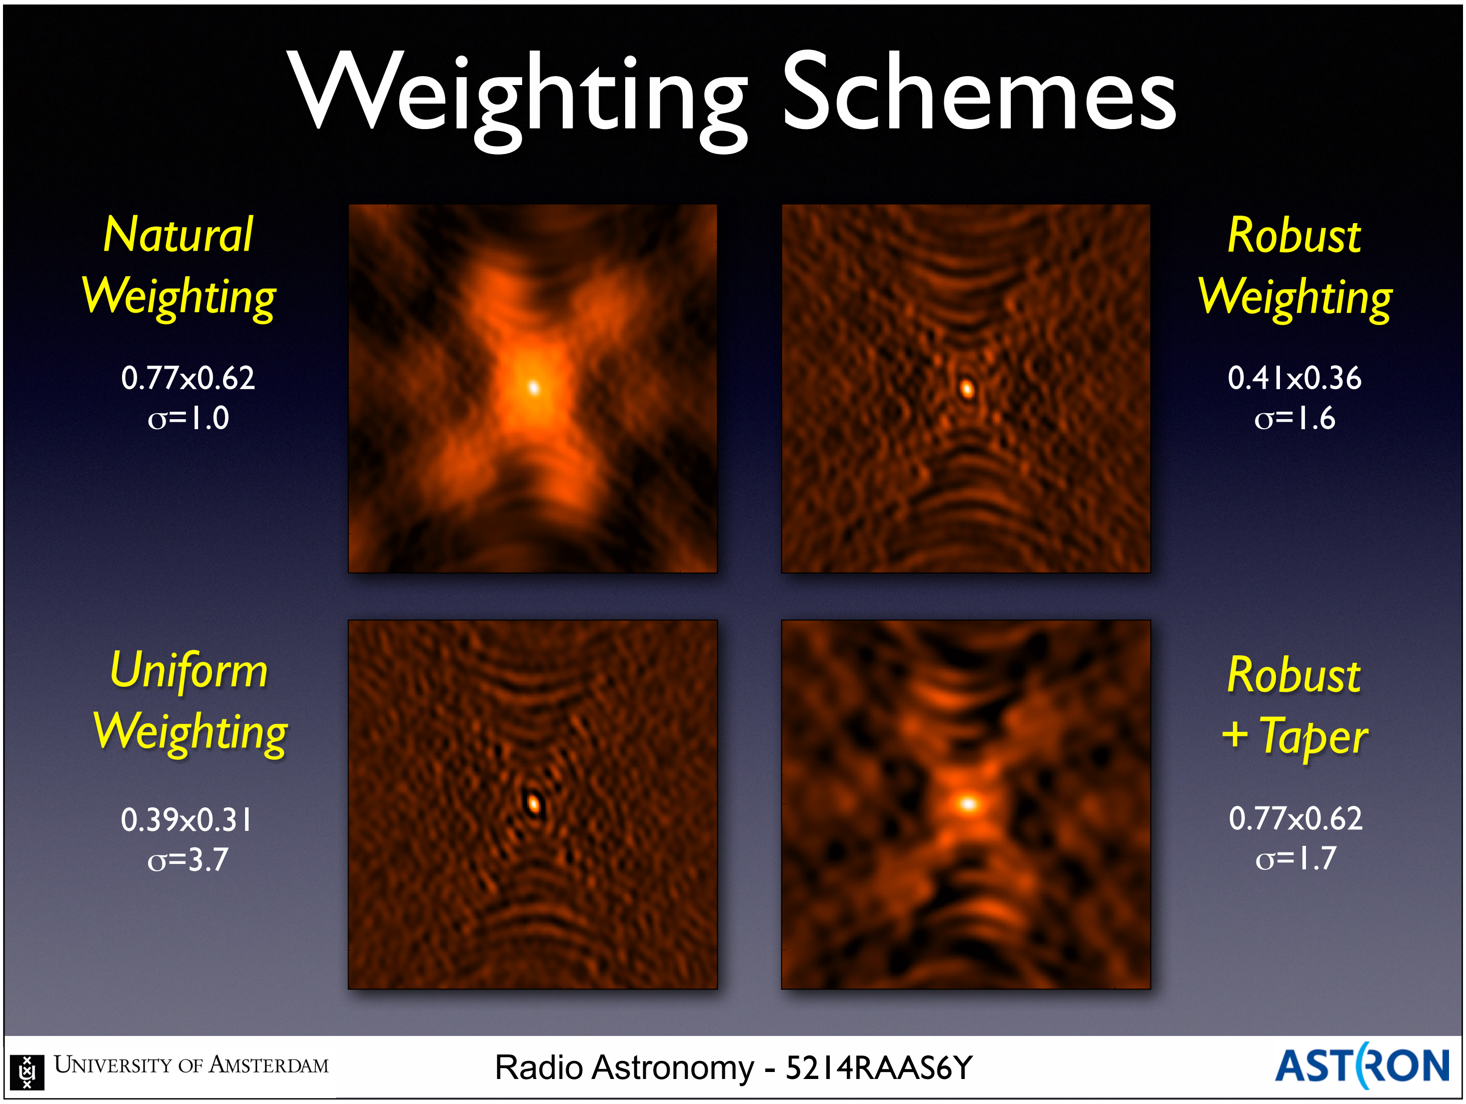

Some examples showing the range of different images you can derive by changing the weighting used.

Essentially you can have maximum resolution or maximum sensitivity, but not both.

Example: WSRT

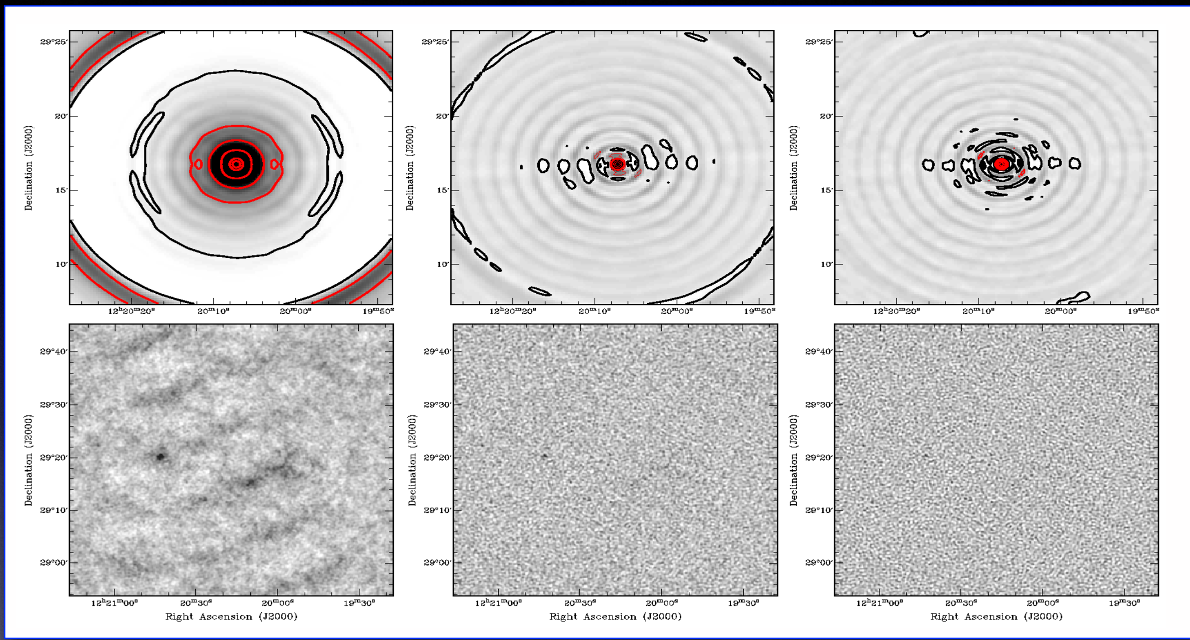

Natural weighting
σ = 0.5

Robust = 0
σ = 0.6

Uniform weighting
σ = 0.7

Difference in noise of 40% (factor 2 in observing time!)

An example of the effect of different weighting schemes on a Westerbork data set.
Can have a big effect on the resulting noise in your image!

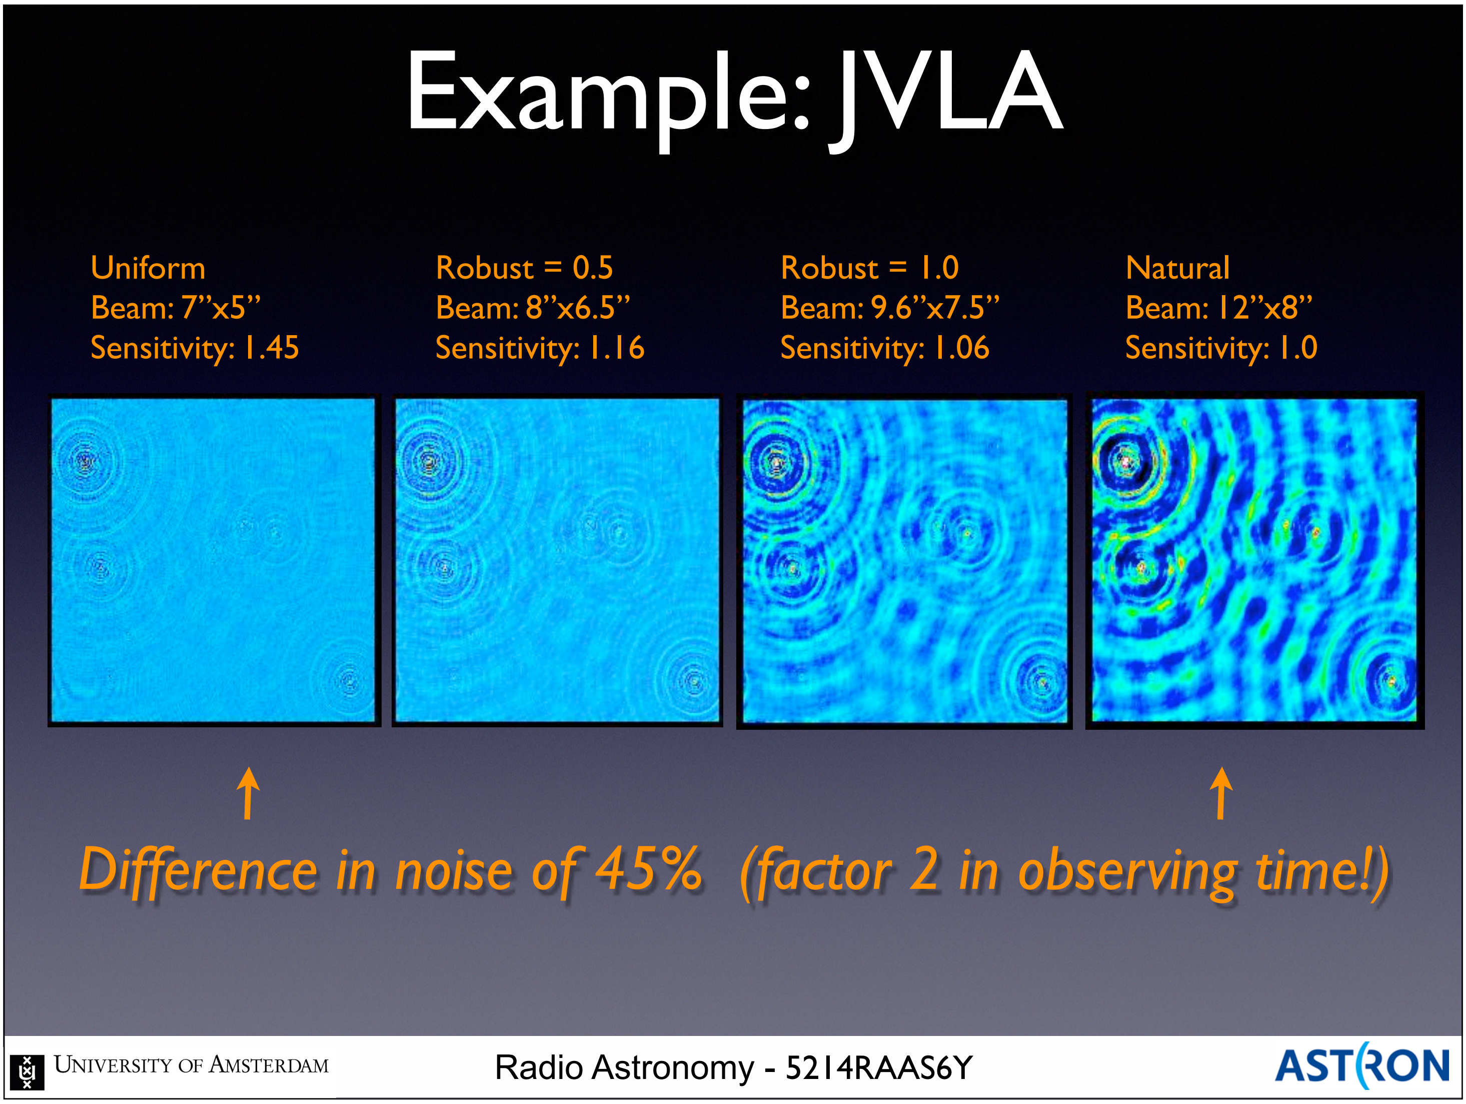

Example based on JVLA data. We go from uniform weighting where we try to make the uv plane uniformly weighted and every point is scaled to the same level. This gives the lowest side-lobe levels because you've got the smoothest structure in the uv plane. It gives you high resolution but at the cost of increased noise. On the other side of the scale you can do natural weighting you use the data as you observe it, give them the weights they were observed with, and don't re-weight it in any way. You get lower resolution because you have more uv data at short baselines and maximum sensitivity. You can see a 45% improvement in sensitivity in going from uniform to natural weighting. Then there's robust weighting which is a kind of optimal combination of the two and it has a robustness parameter that you can adjust which gives you most of your resolution without losing too much sensitivity. If you move it more towards natural weighting you get some more sensitivity but at the cost of increased sidelobe structure in your image.

Weighting Summary

	Uniform/Robust *All spatial- frequencies get equal weight*	Natural/Robust *All data points get equal weight*	Tapering *Lower spatial freqs. get higher weight*
Resolution	Higher	Medium	Lower
Sidelobes	Lower	Higher	Depends
Point Source Sensitivity	Lower	Maximum	Lower
Extended Source Sensitivity	Lower	Medium	Higher

- Imaging parameters provide a lot of freedom
- Appropriate choice depends on science goals

Summary of the different weighting schemes and their advantages and disadvantages.
The imaging software we use has all these options built-in, the choice is up to you.
Know what the defaults you used are, someone will always ask!

Deconvolution

- Calibration and Fourier transform: $V(u,v) \Rightarrow I_D(x,y)$
- Deconvolve $B(x,y)$ from $I_D(x,y)$ to recover $I(x,y)$ for science
- Information is missing, so be careful (there's noise, too)

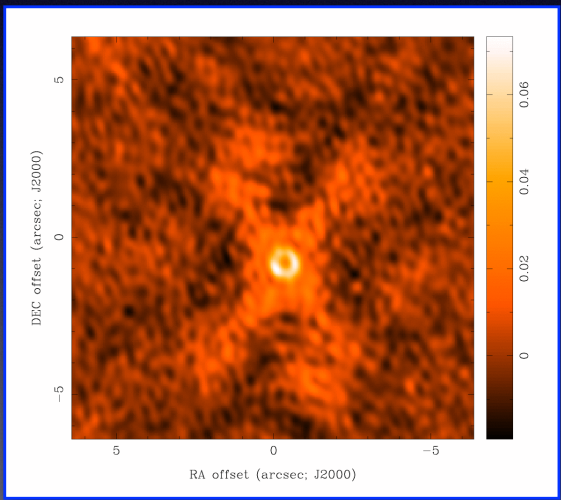

Dirty Image

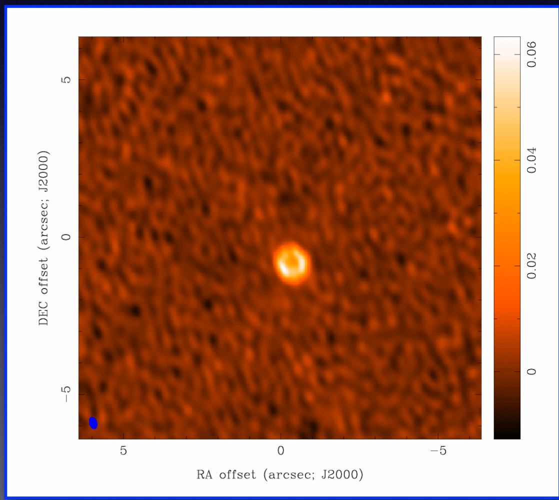

Cleaned Image

Once we have a dirty image, we can try to remove the artifacts caused by the sampling process. These artifacts are represented by the dirty beam. We correct for the beam by deconvolution. This process is usually called “cleaning”. Incomplete sampling and the presence of noise make it an imperfect process.

Deconvolution Issues

Iteratively fit a sky-model to the observed visibilities

- Reconstruction Issues
  - No unique solution. In fact, there are infinite solutions.
  - There will always be un-resolved structure ⇒ Unphysical to believe structure < FWHM of beam
  - Total integrated power is never measured ⇒ Reconstruction of largest spatial scales is always an extrapolation
  - Requires iterative, non-linear fitting process ⇒ Compute intensive
  - No unique prescription for extracting optimal solution

⇒ Constrain the solution using astrophysical plausibility

Since we don’t know the true surface brightness on the sky, we try to model it.
The final image we derive is our “best” model.
Best is defined as the model that reproduces the observed data the most accurately.
We find this best model through an iterative fitting process.
Every radio image ever shown is just someone’s best model!

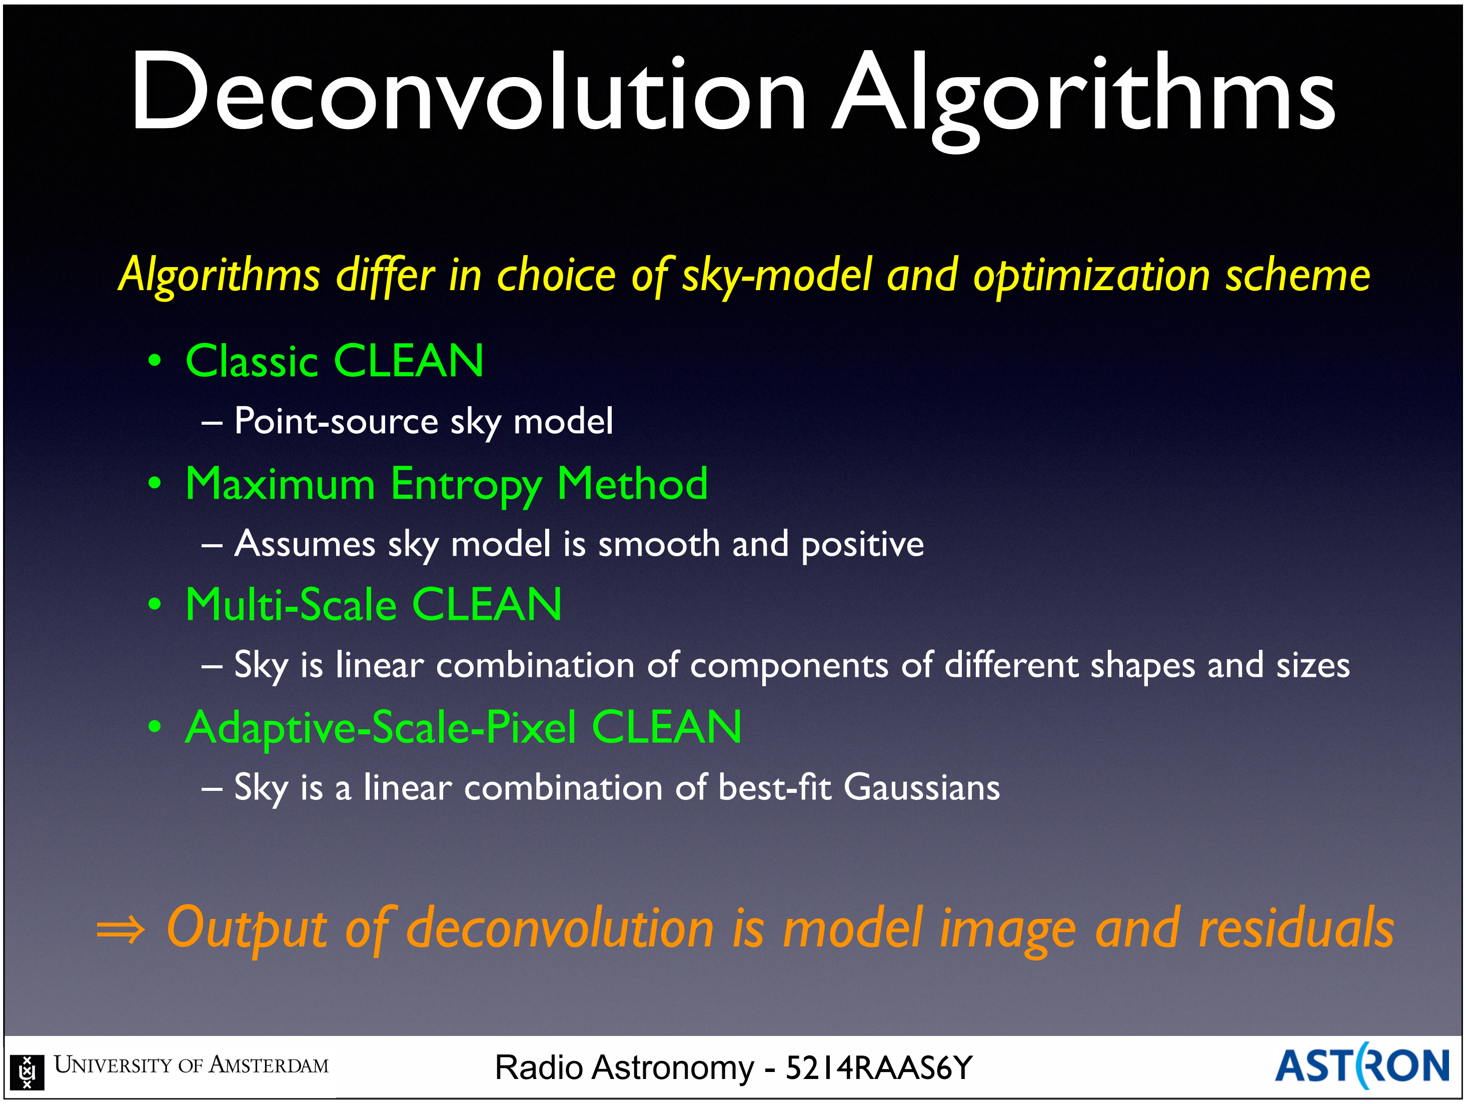

There are several common, iterative deconvolution algorithms in use currently.
The are operationally similar but differ mainly in what functions they combine to model your source.
The output of the deconvolution process in your best model image and the residuals.
Residuals are the difference between the model and the data.
Small residuals —> Good model.

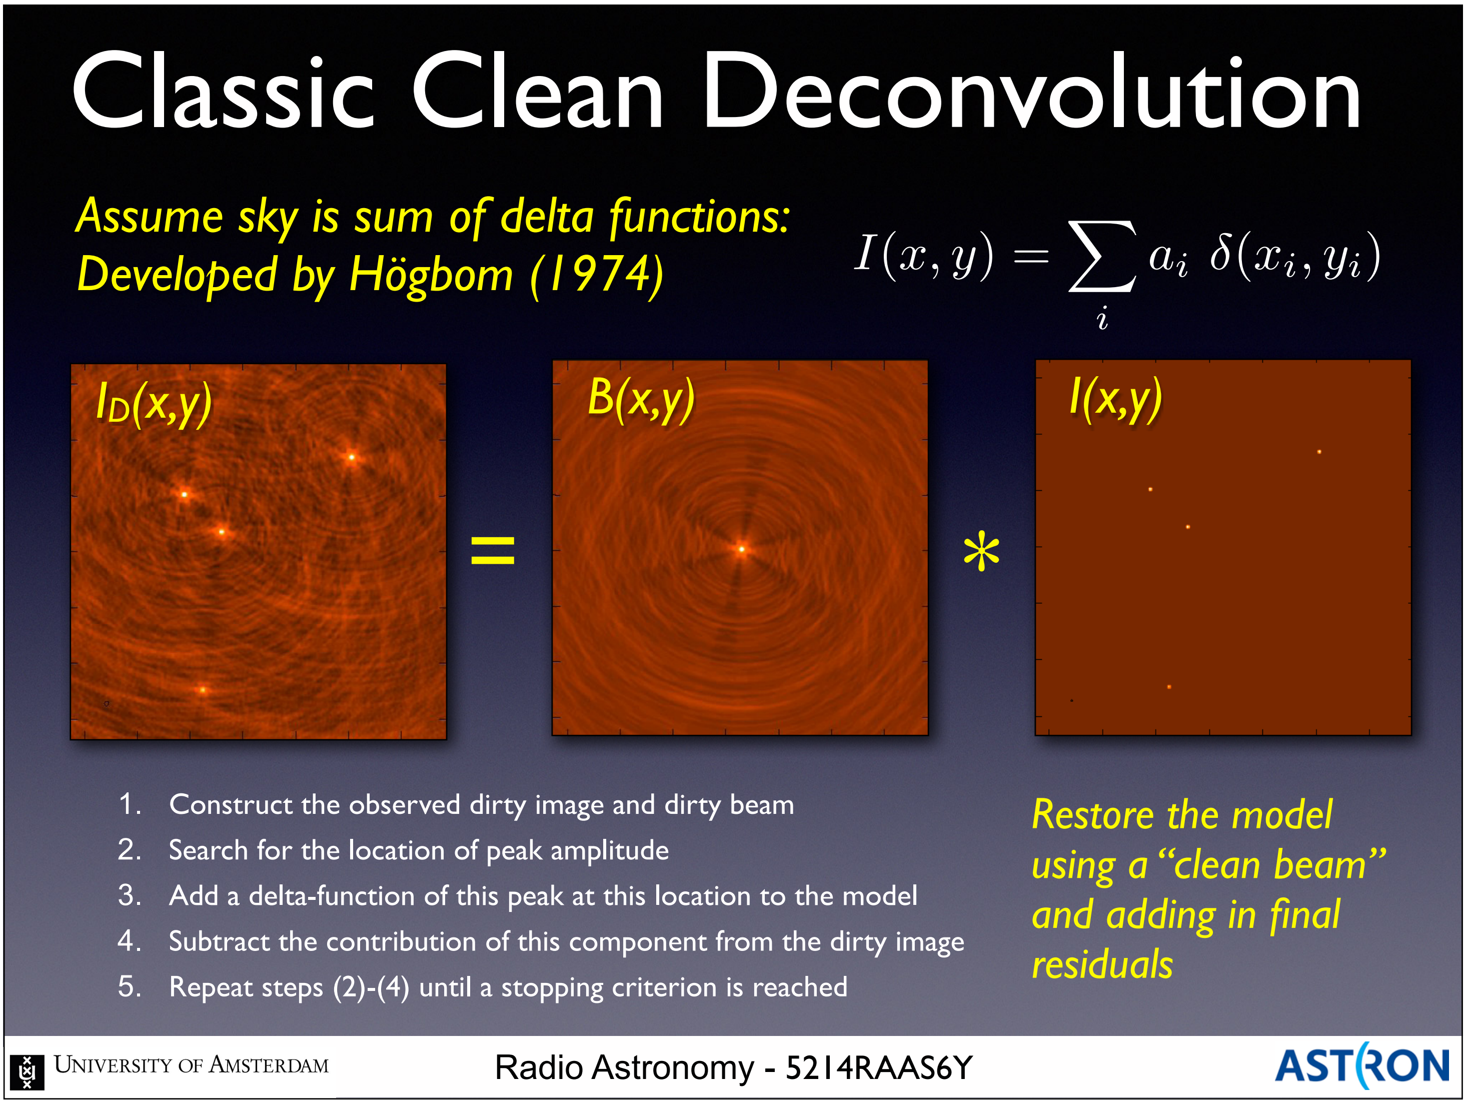

The basic CLEAN algorithm has been in use since 1974.
It assumes that your source is made up of a sum of (many potentially) delta functions.
Very simple to implement and great for point sources.
Not so great for extended, diffuse emission.

Clean in Action

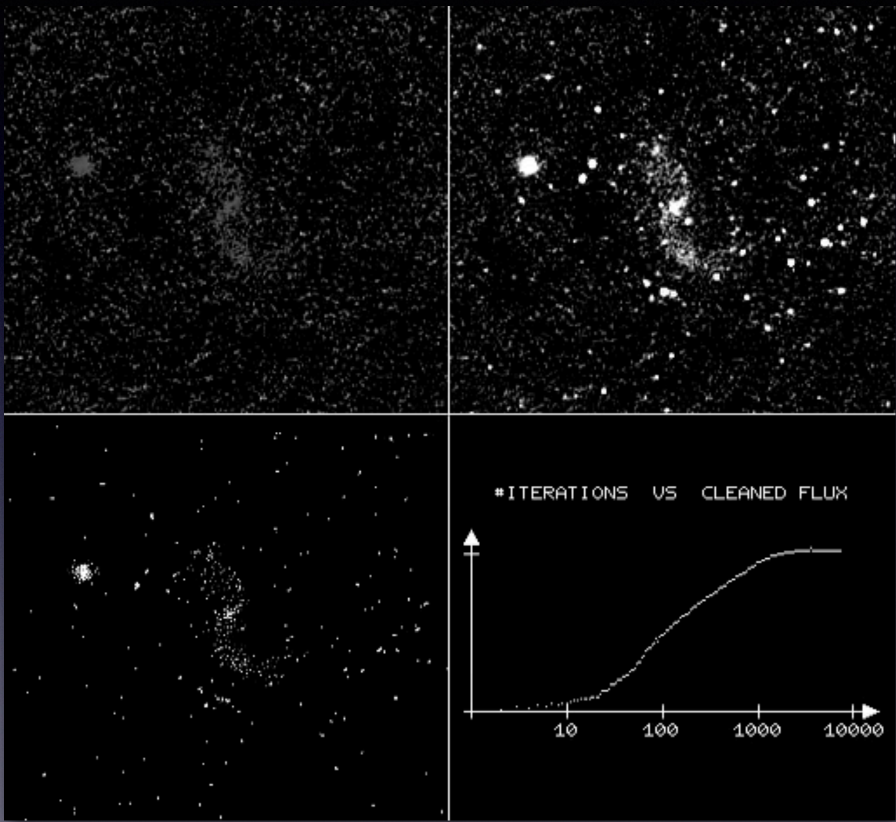

An example of the CLEAN algorithm in operation. Runs for 10,000's of iterations.
The algorithm stops cleaning when the total flux in the model stops changing significantly.
At that point, you are just cleaning noise, basically assuming peaks in the noise are real sources.

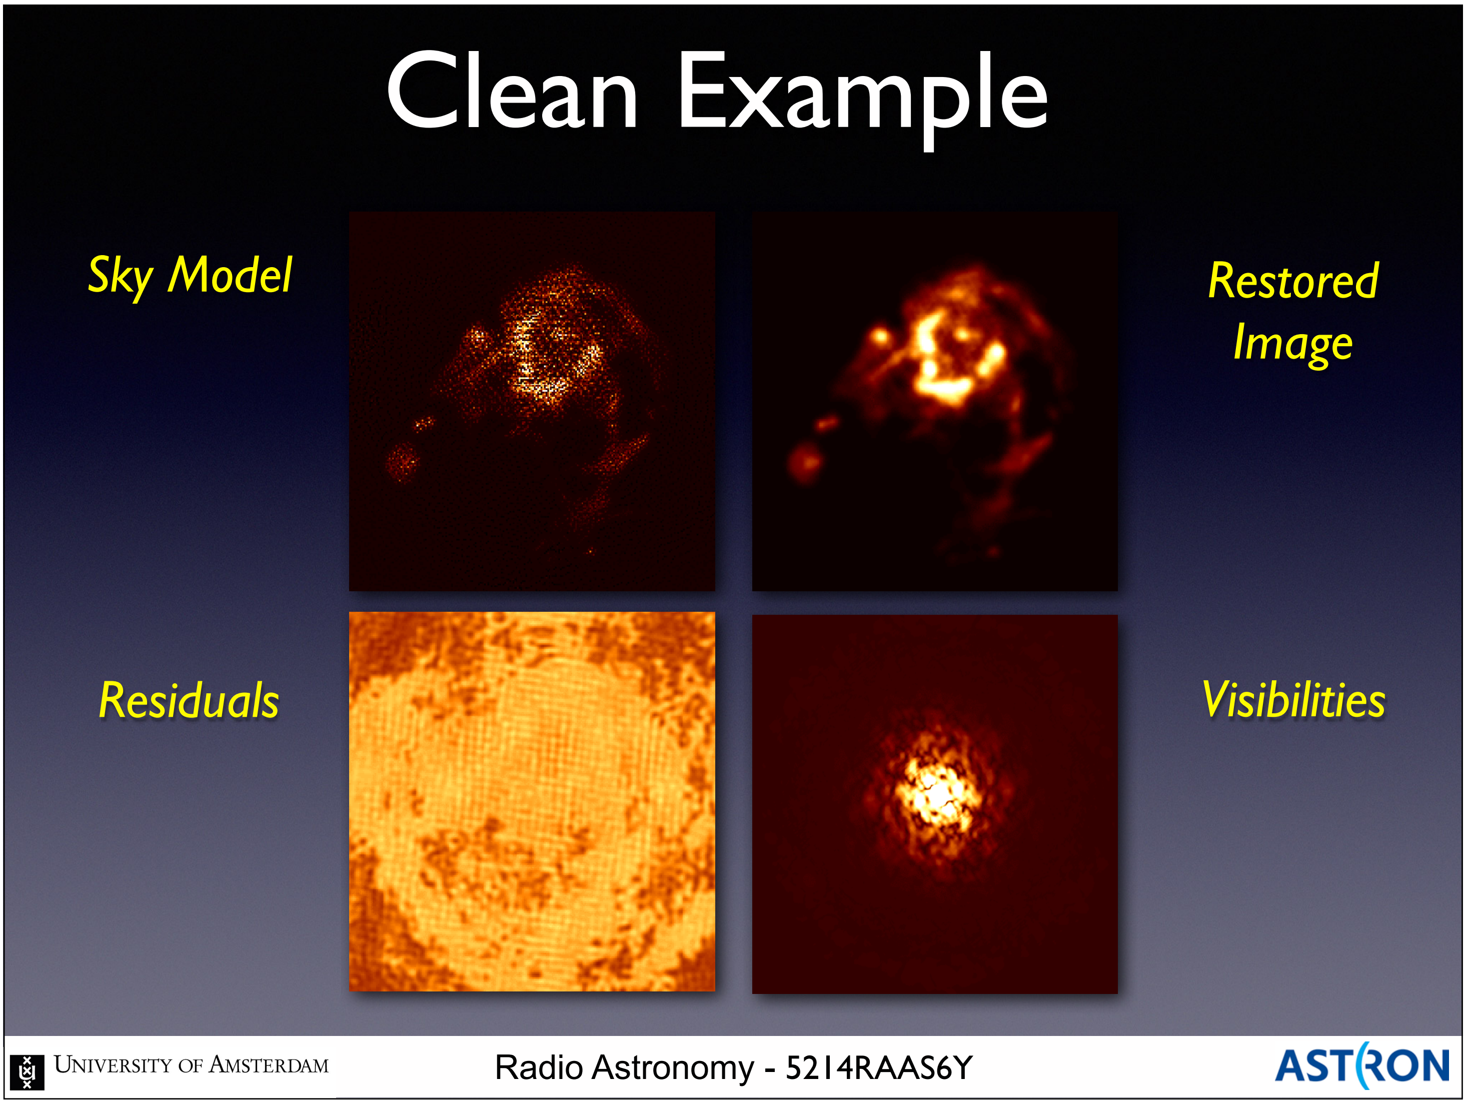

Examples of using the CLEAN algorithm on a complicated source.
The upper left shows the attempt to model the source using a large number of delta functions.
The residual map shows you what components you missed.
Ideally, it should be uniform. Notice how there are bright regions which follow the overall source shape.
The model is missing flux on this largest scale.

Adaptive Scale Pixel CLEAN

Assume sky is sum of Gaussian functions:
Bhatnagar & Cornwell (2004)

$$I(x,y) = \sum_i a_i \, e^{-\left[\frac{(x-x_i)^2}{\sigma_{xi}^2} + \frac{(y-y_i)^2}{\sigma_{yi}^2}\right]}$$

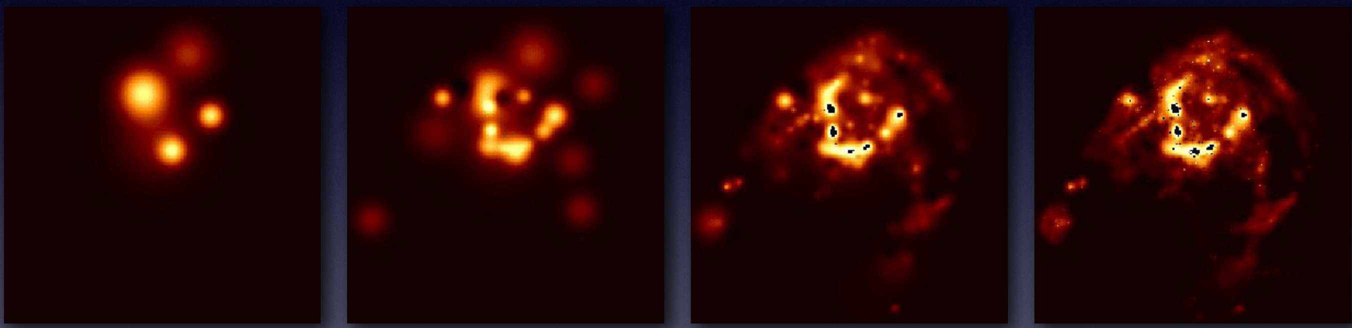

1. Calculate the dirty image, smooth to a few scales
2. Identify peak across scales to choose initial guess for new component
3. Add this new component to the list
4. Re-fit Gaussian parameters for new and old components together
5. Subtract the contribution of all updated components from the dirty image
6. Repeat steps (2)-(5) until a stopping criterion is reached

Adaptive Scale sizes leads to better image reconstruction

The Adaptive Scale Pixel CLEAN algorithm uses a sum of 2D Gaussian functions to model sources. The scales of the Gaussians can be adjusted to represent both extended and point source emission. More computationally expensive, but potentially more accurate.

ASP-Clean Example

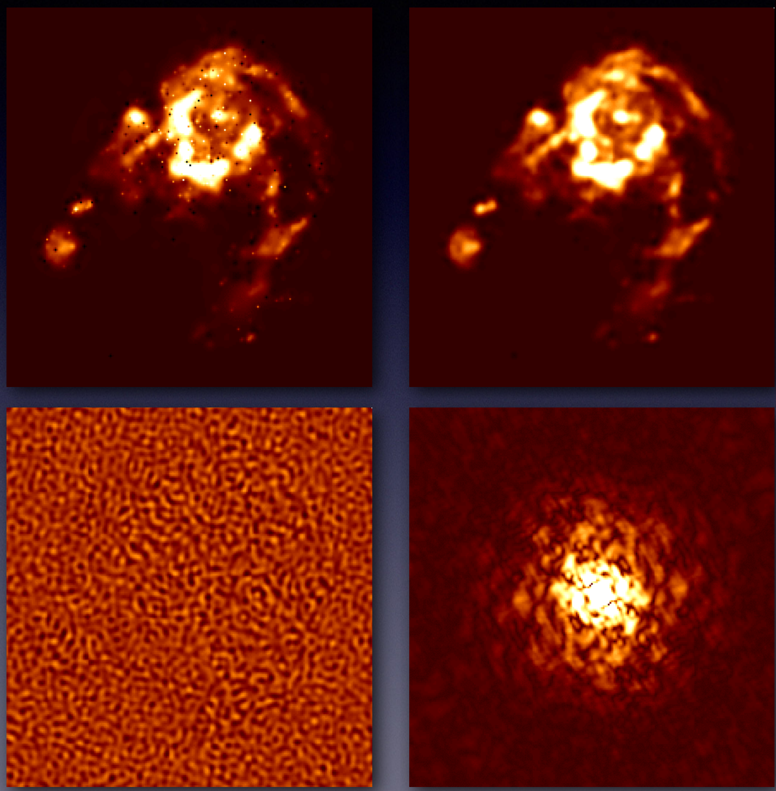

Examples of using the ASP CLEAN algorithm on the same complicated source.
Notice how the residuals are now uniform over the whole image.
We are doing a much better job at model flux from all scales in the image.

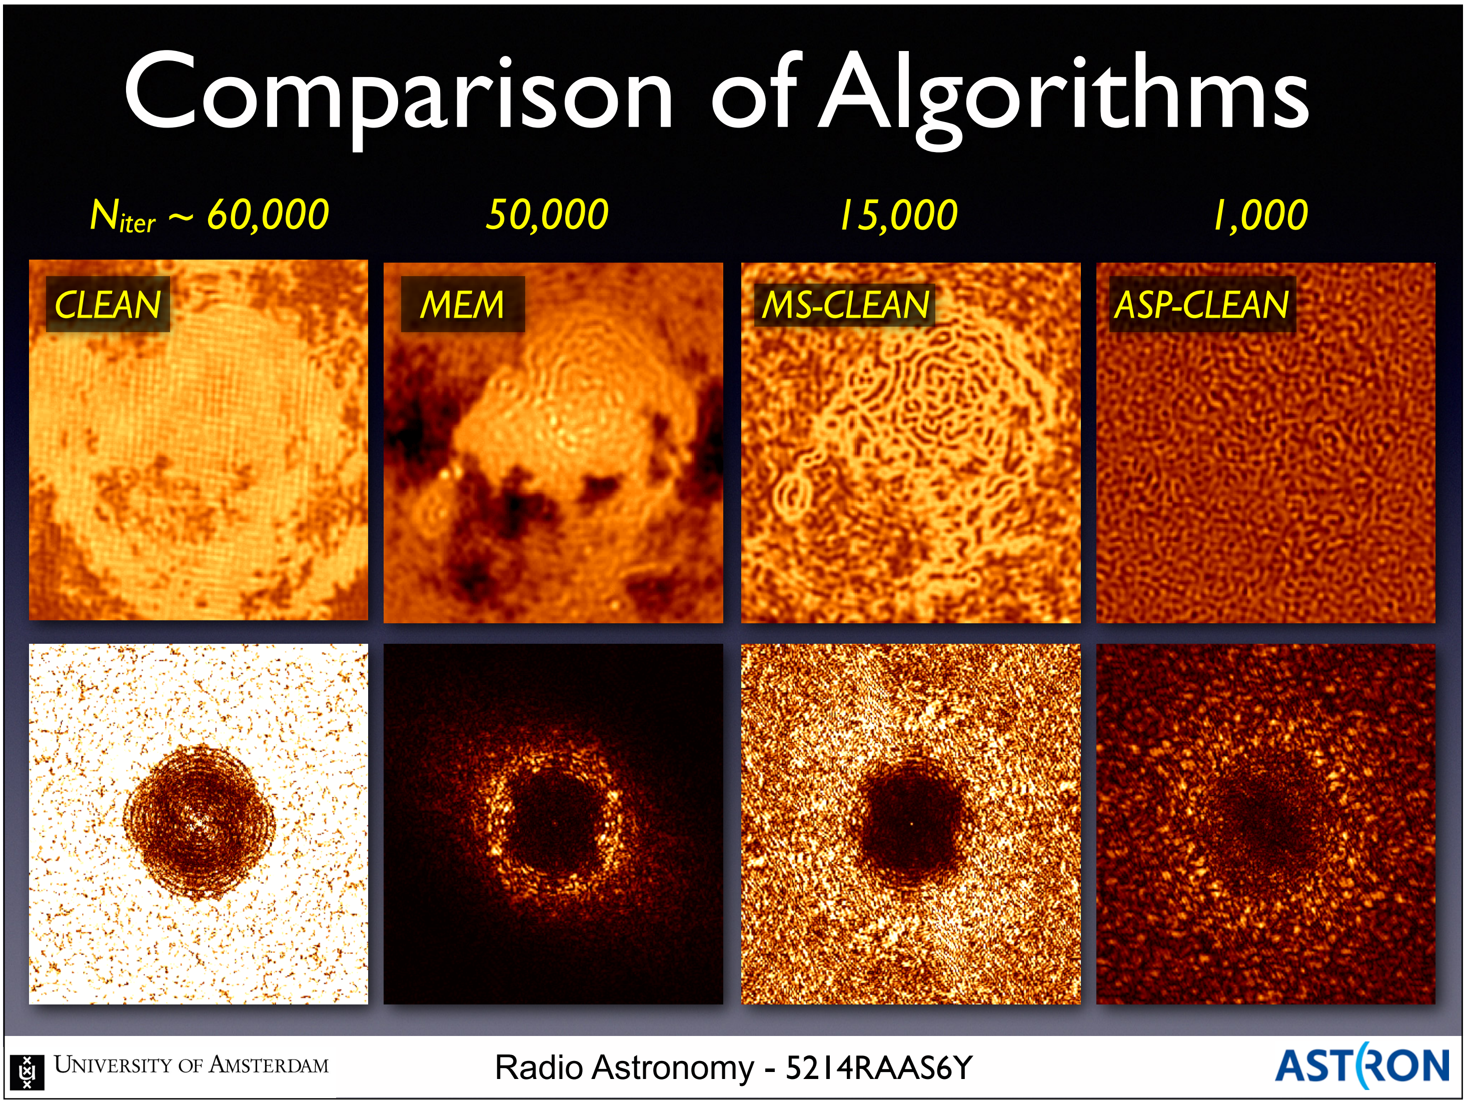

Some examples of the residuals for various CLEAN algorithms and what they cost to compute.
The images are labeled with the number of iterations required to converge to a good solution.
Choose the algorithm that suits your science!
If you have a field of point sources, the traditional CLEAN may work very well *and* be fast.

Intermission

Image Quality

Measures of Image Quality

- "Dynamic Range"
  - *Defined as ratio of peak brightness to RMS noise in a region empty of emission*
  - *Alternatively, use ratio of peak brightness to peak error (residuals)*
  - *Easy to calculate lower limit to the error in brightness in a non-empty region*
  - *Values run from DR ~10^2 - 10^6*
- "Fidelity"
  - *Difference between the calculated image and the correct image*
  - *Convenient measure of how accurately image matches true I(x,y) on sky*
  - *Need a priori knowledge of the correct image for comparison*
  - *Fidelity image = input model / difference*
  - *Similar to a SNR map*

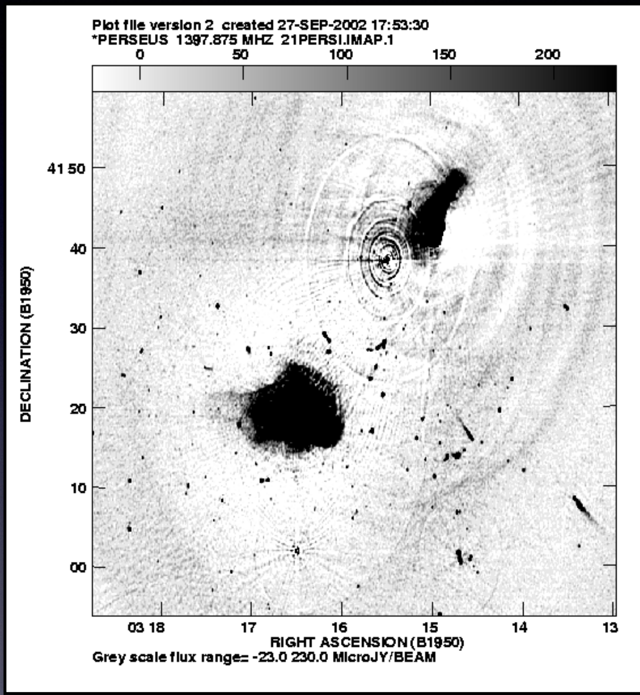

Image of the Perseus cluster showing details exposed at a dynamic range of 1,000,000:1 (de Bruyn & Brentjens 2010)

What makes an image a "good" image? When someone shows you an image, how can you tell if it can be trusted?

There are severals quantitative ways of measuring image quality.

Dynamic range is one often used metric.

Dynamic range is the ratio of the highest value in the image to the lowest, i.e. DR = max(Image)/ min(Image).

Simultaneously tells you about the noise in the image and how sensitive it is to faint structures.

Image fidelity is used less often since we don't really know the true image.

Recognizing Errors

Some Questions to ask:

Noise properties of image:

Is the rms noise about that expected from integration time?
Is the rms noise much larger near bright sources?
Are there non-random noise components (faint waves and ripples)?

Funny looking Structure:

Non-physical features; stripes, rings, symmetric or anti-symmetric
Negative features well-below a few times the rms noise
Does the image have characteristics that look like the dirty beam?

Image-making parameters:

Is the image big enough to cover all significant emission?
Is cell size too large or too small? ~4 points per beam okay
Is the resolution too high to detect most of the emission?

Things to look for when diagnosing whether an image is good or not.
The noise properties and the presence of artifacts are the most straight-forward criteria to apply.

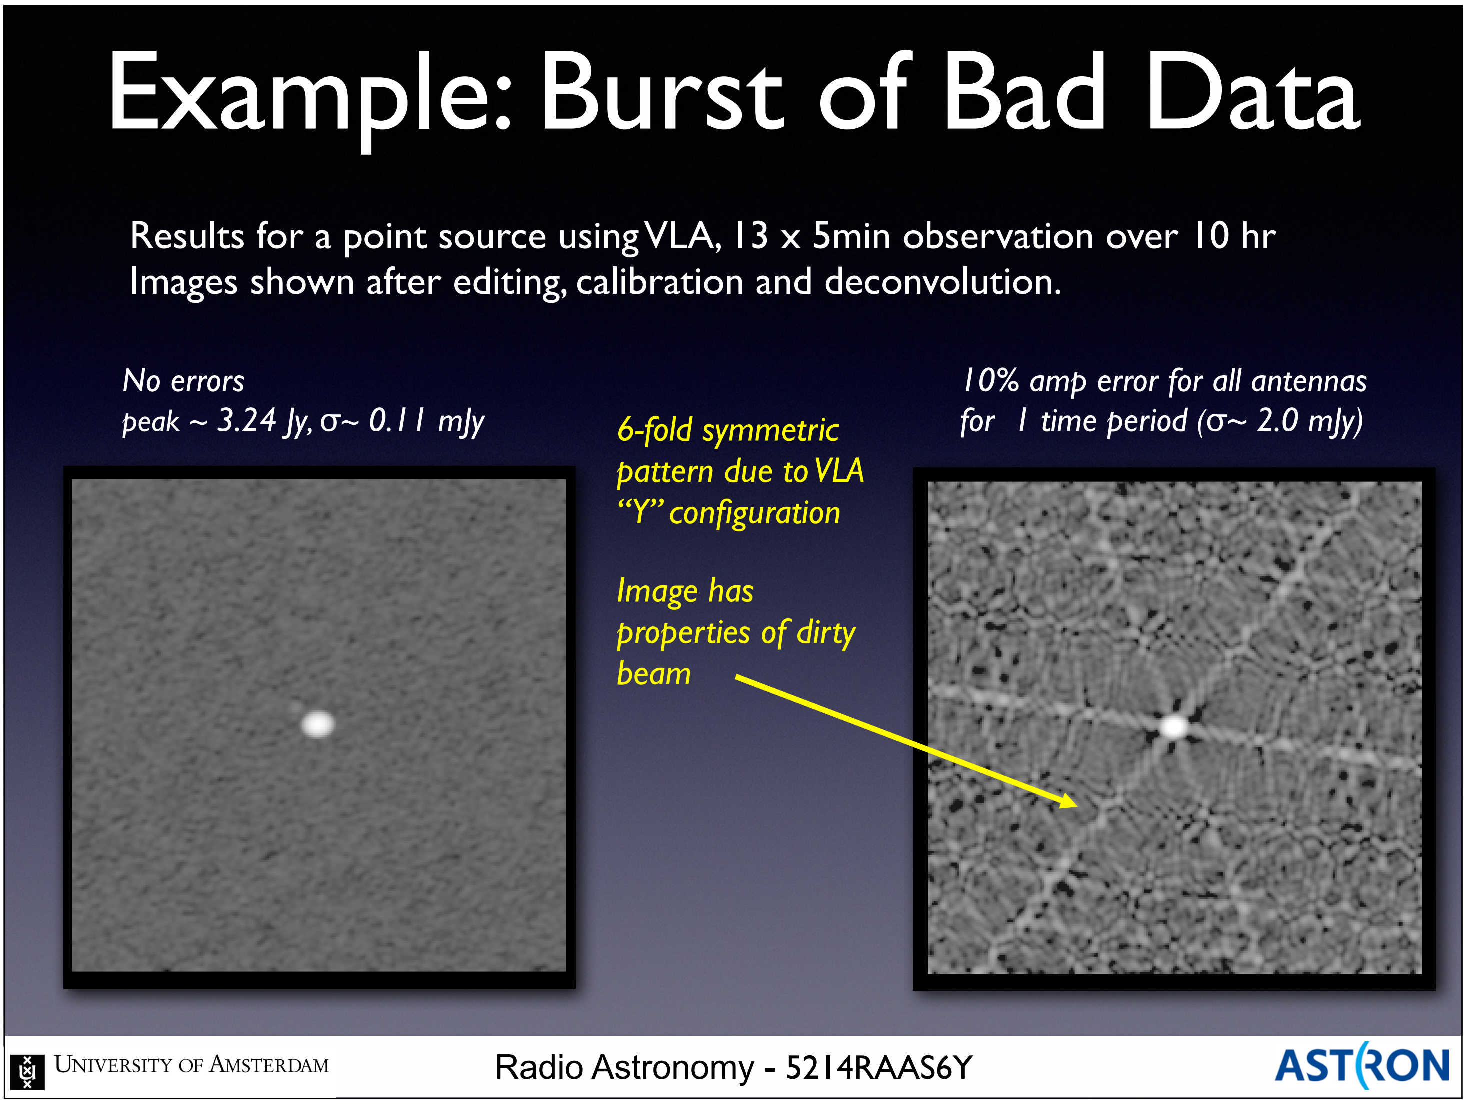

An example of an image artifact that can be caused by bad data.
In this case, a burst of bad data for a short period produces a regular, repeating artifact in the image.
Doesn't look like an astronomical source *and* looks like the VLA visibility function.
Solution —> try to recalibrate or flag that data and re-image.

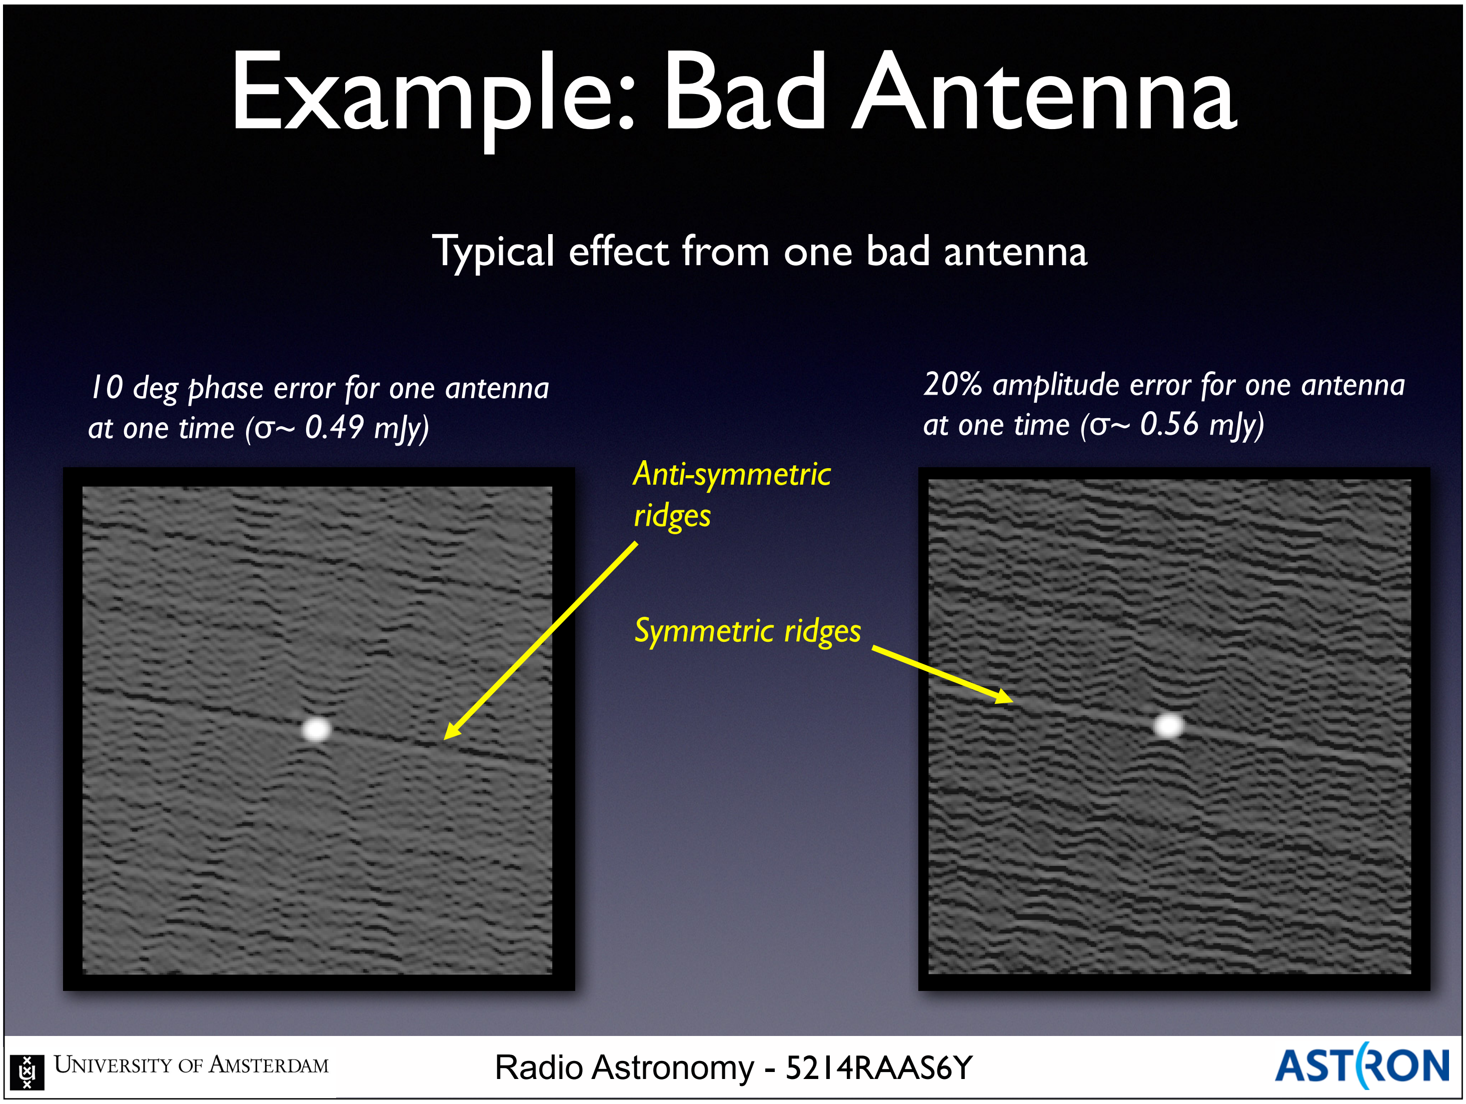

More examples of artifacts caused by bad data. Note the different artifacts caused by phase versus amplitude errors.
Solution —> try to recalibrate or flag that data and re-image.

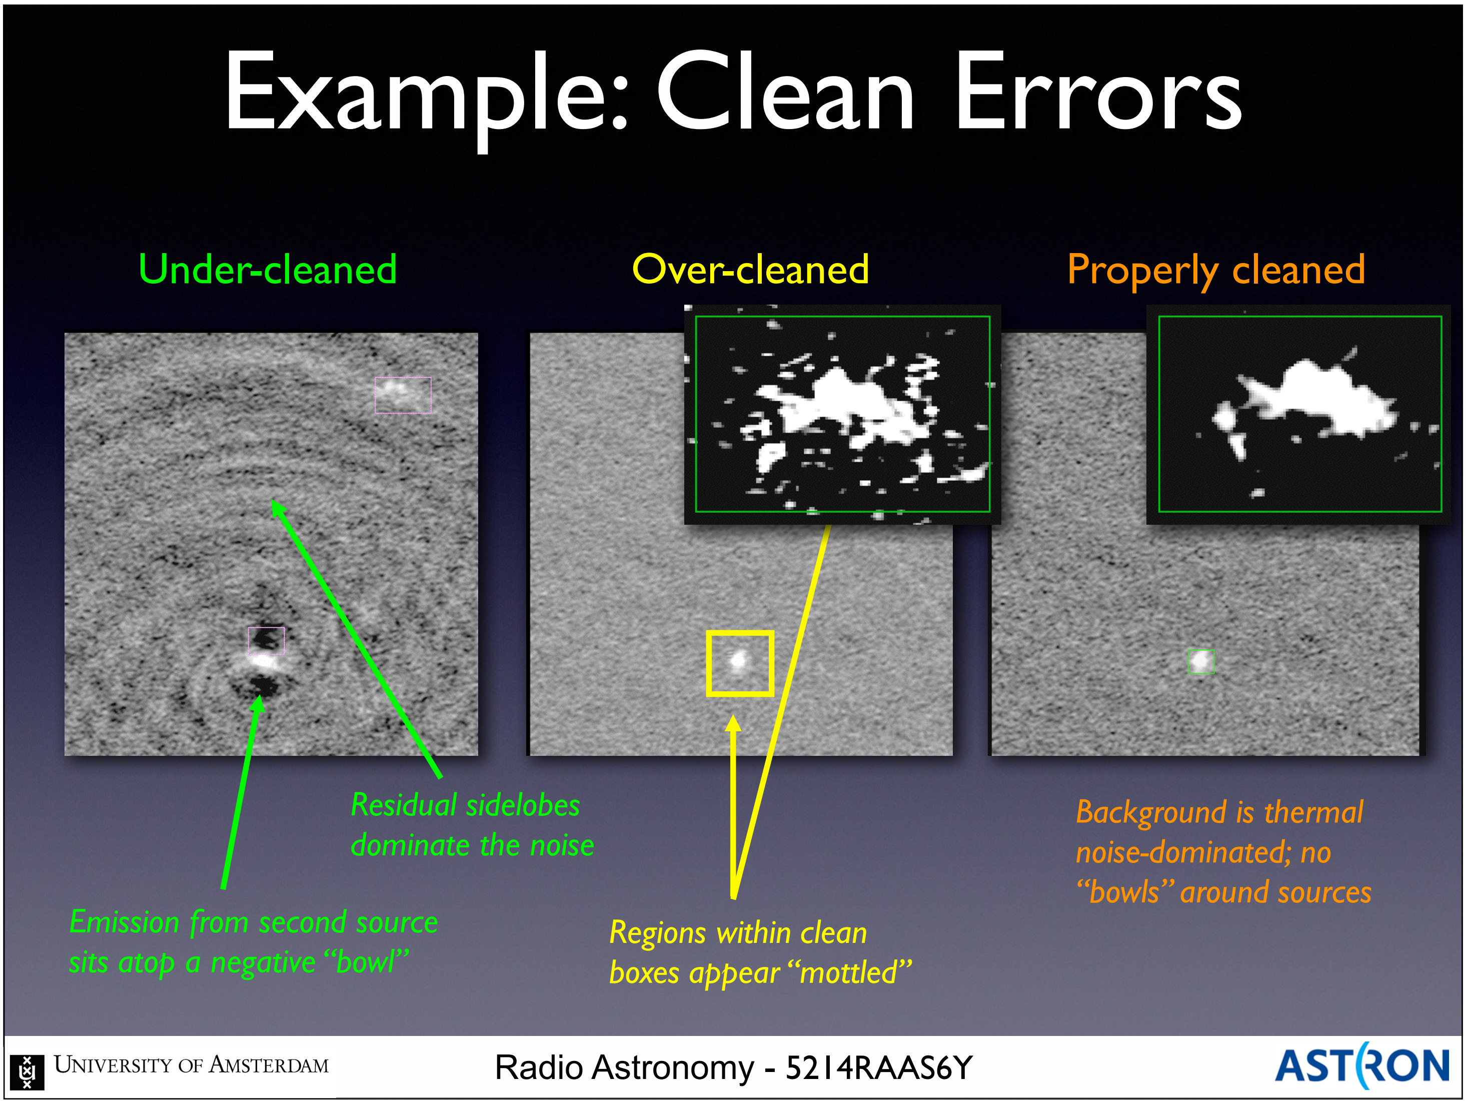

It is possible to over-use the CLEAN process.
Can produce distorted images especially for complicated sources.
Solution —> make your cleaning window smaller and re-run the CLEAN process.

Recognizing Errors

Source structure should be "reasonable", the rms image noise as expected, and the background featureless. If not:

Examine (*u*,*v*) data

Look for outliers in (u,v) data using several plotting methods.

Check calibration gains and phases for instabilities.

Look at residual data (u,v data - clean components)

Examine image plane

Do defects resemble the dirty beam?

Are defect properties related to possible data errors?

Are defects related to possible deconvolution problems?

Are other corrections/calibrations needed?

Does the field-of-view encompass all emission?

Checklist of things to examine when diagnosing your image quality.

Advanced Imaging

Wide-band Imaging

- Radio telescopes suffer from chromatic aberration ⇒ "bandwidth smearing"
- Measure visibilities in many narrowband channels to avoid bandwidth-smearing
- Construct visibilities for multiple narrowband channels, each with its own delay-tracking

  Max. channel width: $\delta\nu < \nu_0 \left(\frac{D}{b_{max}} \right)$

- Can use multi-frequency-synthesis to increase the uv-coverage used in deconvolution and image-fidelity
- Can make images at the angular-resolution allowed by the highest frequency
- Can take source spectrum into account

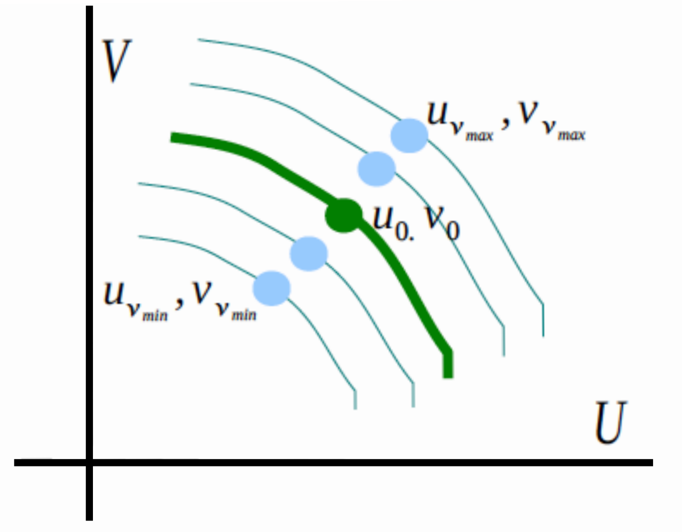

Spatial-frequency coverage changes with frequency

Modern radio telescopes can operate over a wide range of frequencies.
If the bandwidth is large enough, the properties of the telescope can vary a lot over that bandwidth.
These variations have advantages and disadvantages.
Wide-band imaging is the process of taking these variations into account to make good images.

Bandwidth Smearing

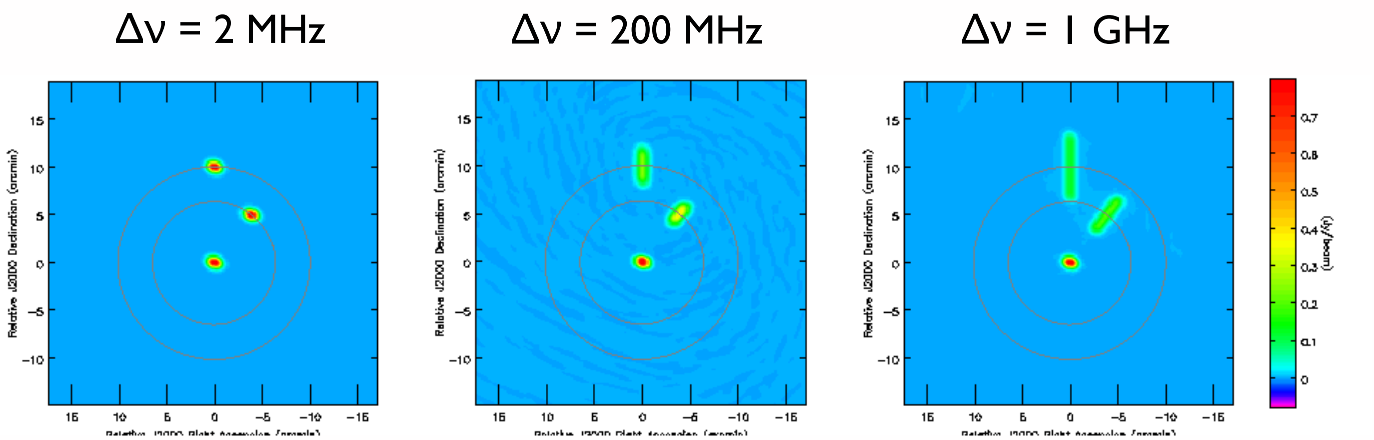

Bandwidth smearing refers to the fact that the sampling of the (u,v) changes with frequency.
If you don’t take the frequency effects into account, the resulting image will be distorted.
Sources positions will be smeared out in the image plane.
The further off the phase center the source is, the more smeared out it will be.

Multifrequency Synthesis

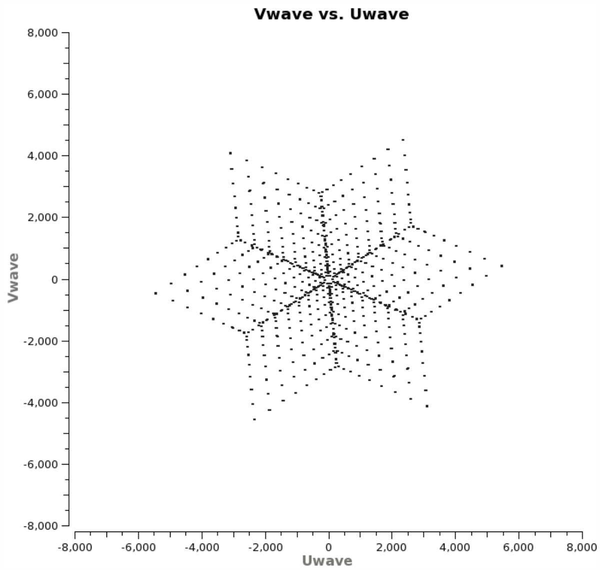

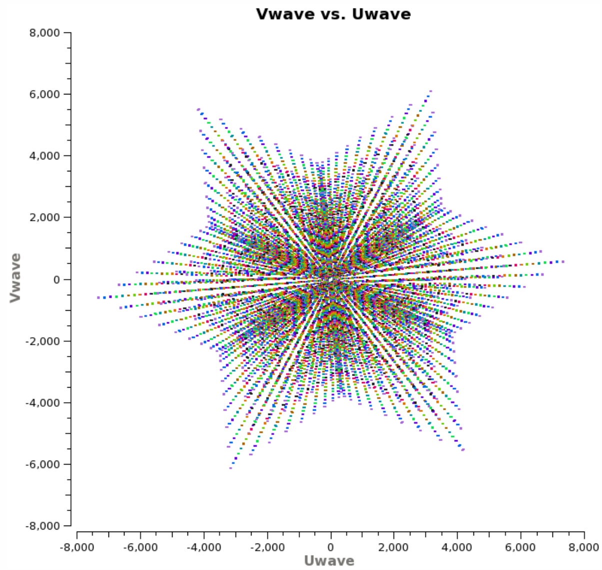

The change of position due to frequency variations in the (u,v) plane has a positive aspect as well.
At a single frequency and time, a single baseline adds a single point to the (u,v) plane.
At multiple frequencies, a single baseline adds multiple (u,v) points for a single time.
Multifrequency synthesis lets us take advantage of this behavior and get a better (u,v) sampling.

Multifrequency Synthesis

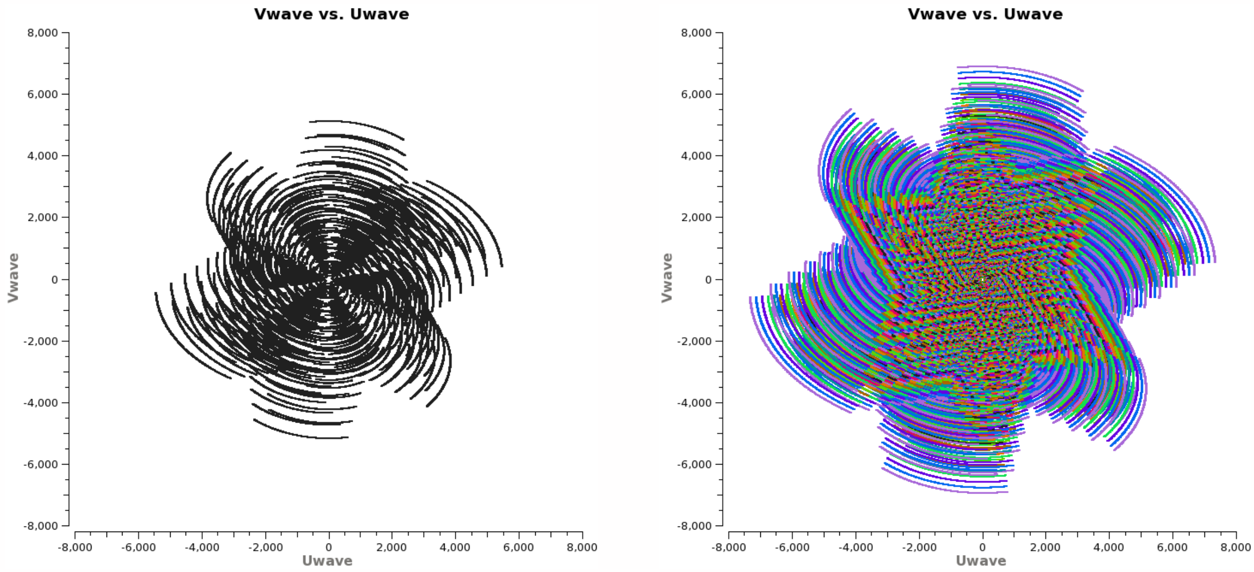

- Overlapping uv coverage ⇒ better sensitivity ⇒ $\sigma_{cont} = \frac{\sigma_{chan}}{\sqrt{N_{chan}}}$
- Increased uv filling ⇒ better imaging fidelity
- Larger spatial-frequency range ⇒ better angular resolution ⇒ $\frac{\lambda}{b_{max}}$

Add in time variation and you can get a much better sampling of the (u,v) plane.
This technique can produce much better images and is built into most standard analysis packages.

MFS Example: 3C286

DR =1600 -13000
σ=9 mJy -1 mJy
NTerms = 1

DR =10000 -17000
σ=1.0 mJy - 0.2 mJy
NTerms = 2

DR =65000 -170000
σ=0.2 mJy - 85 μJy
NTerms = 3

DR =110000 - 180000
σ=0.14 mJy - 80 μJy
NTerms = 4

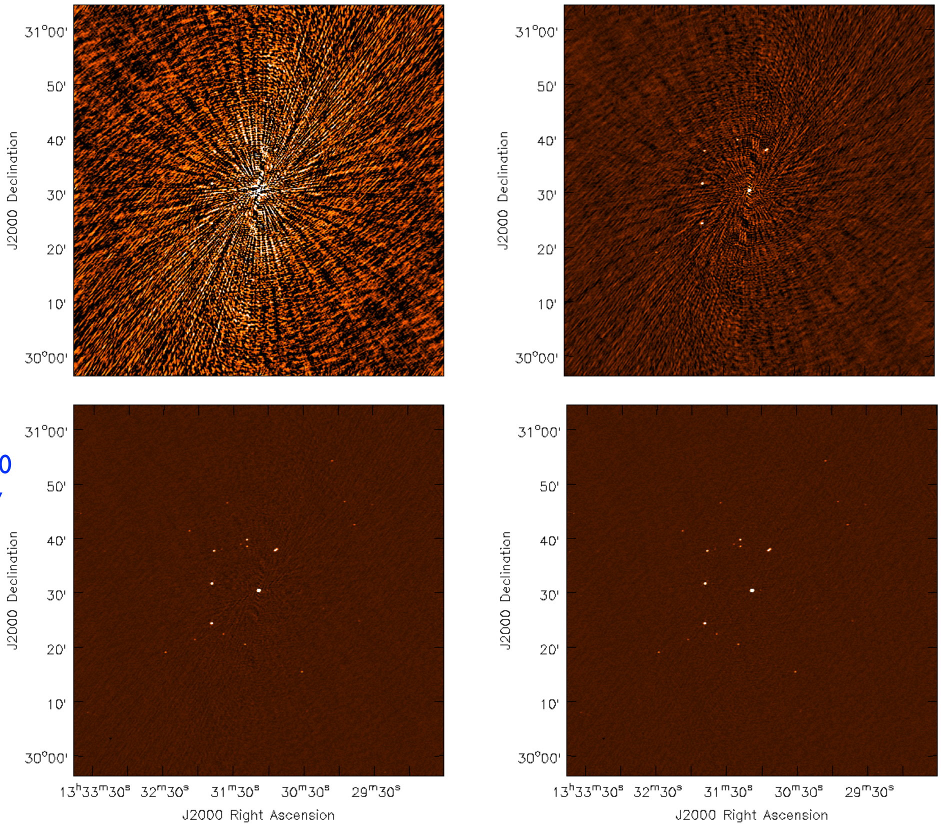

Sample images showing the potential improvement by including multiple frequency terms in the imaging.
Notice the change in the Dynamic Range (DR) of the images as more frequency terms are included.

MFS Example: Cygnus A

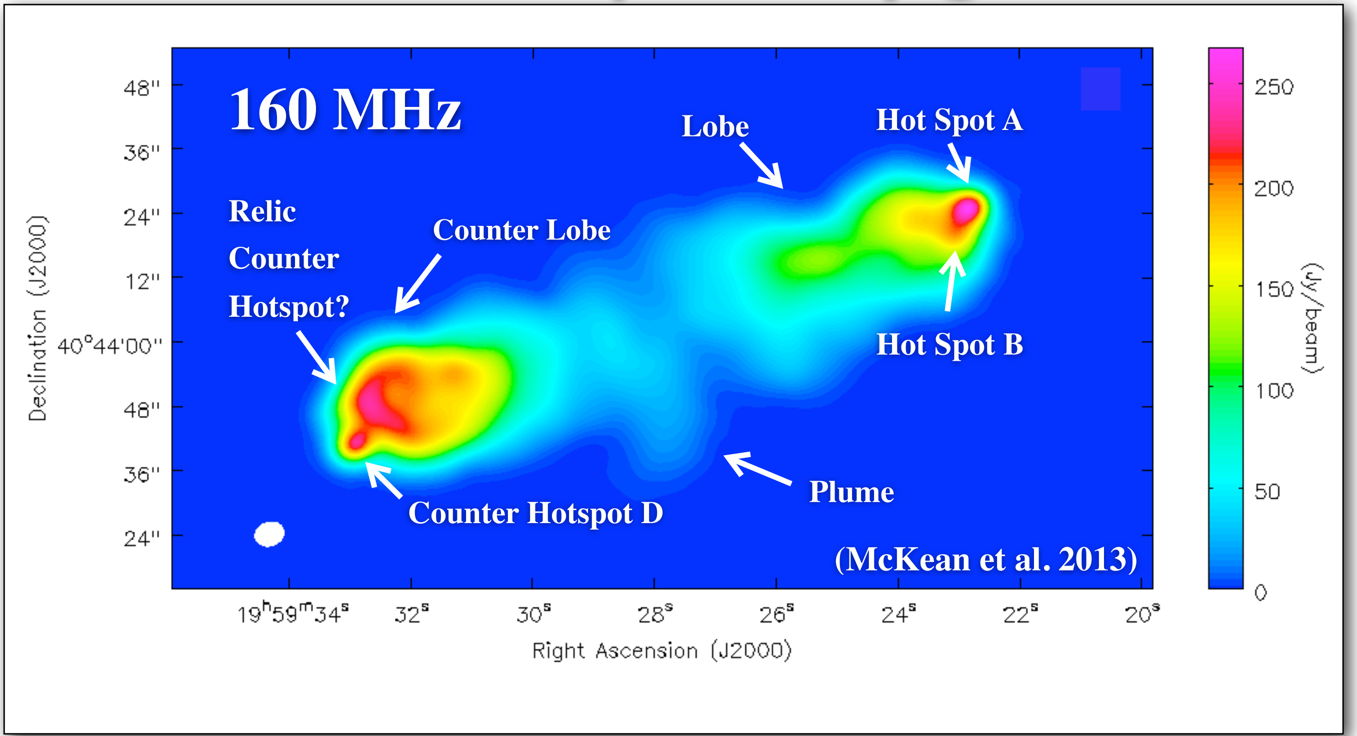

LOFAR HBA 6 hr / 110 - 182 MHz / 16 MHz
σ ~ 70 mJy / DR ~ 3000
NL baselines only, 3.0 arcsec resolution

An example from LOFAR made using multifrequency synthesis.
This image includes frequency points from 110 – 182 MHz.
Low frequency telescopes have very large frequency bandwidths, so this technique is especially useful for these telescopes.

Wide-field Imaging

Wide-band Imaging often requires wide-field imaging techniques

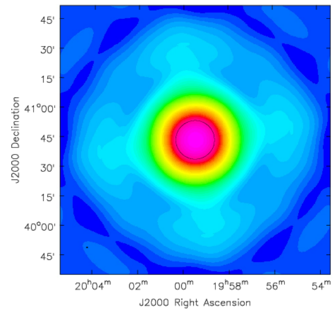

"Primary Beam" :The antenna-primary beam can introduces a time-varying spectrum in the data.

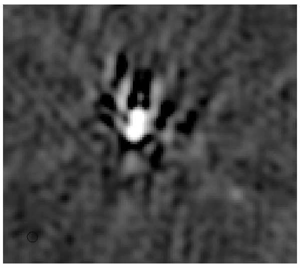

"W-term": Non-coplanar arrays also introduce a frequency-dependent instrumental effect. Narrow-band w-projection algorithm works for wide-band.

"Mosaicing": Make observations with multiple pointing and delay-tracking centers. Combine the data during (or after) image-reconstruction.

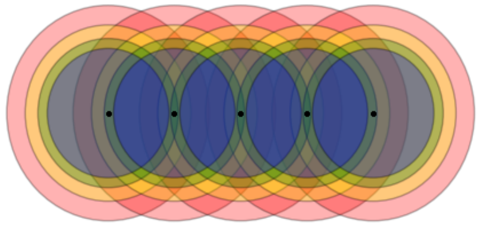

Wide-band effects refer to the frequency coverage and associated effects.
Wide-field effects refer to the variation of the telescope properties as the size of the field of view increases.
These two effects often go together, especially at low frequencies.
Wide-field imaging must take various effects into account, i.e. beam changes, non-planar effects, etc.

Wide-field Imaging

- New instruments are being built with wider fields of view (especially at lower frequencies): MeerKAT, ASKAP, Apertif, Allen Telescope Array, LOFAR
- Wide-field good for all-sky surveys and finding transients
- Traditional synthesis imaging assumes a flat sky and a visibility measurements lying on a (u,v) plane
- These approximations only hold near the phase center (implies small fields of view)
- To deal accurately with large fields of view requires more complicated algorithms (and much more computation)

Many new telescopes are built to do all-sky surveys. Wide fields of views are very useful for doing all-sky surveys, but *only* if you can correct for wide-field imaging effects.

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.

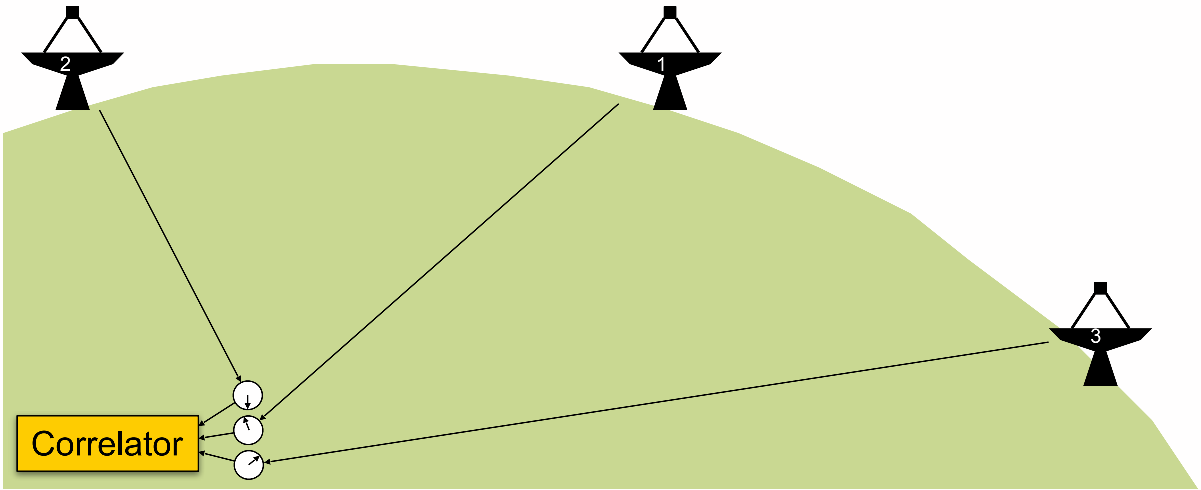

Non-coplanar Arrays

- Choose a direction of interest – this will be known as the phase centre.
- Calculate distances (in wavelengths) w_j between each jth antenna and a projection plane normal to the phase centre.
- Delay each signal $V(t)$ by $-w_j/f$ seconds.
- Signals from a source at the phase centre will then reach the correlator in phase.

Non-coplanar Arrays

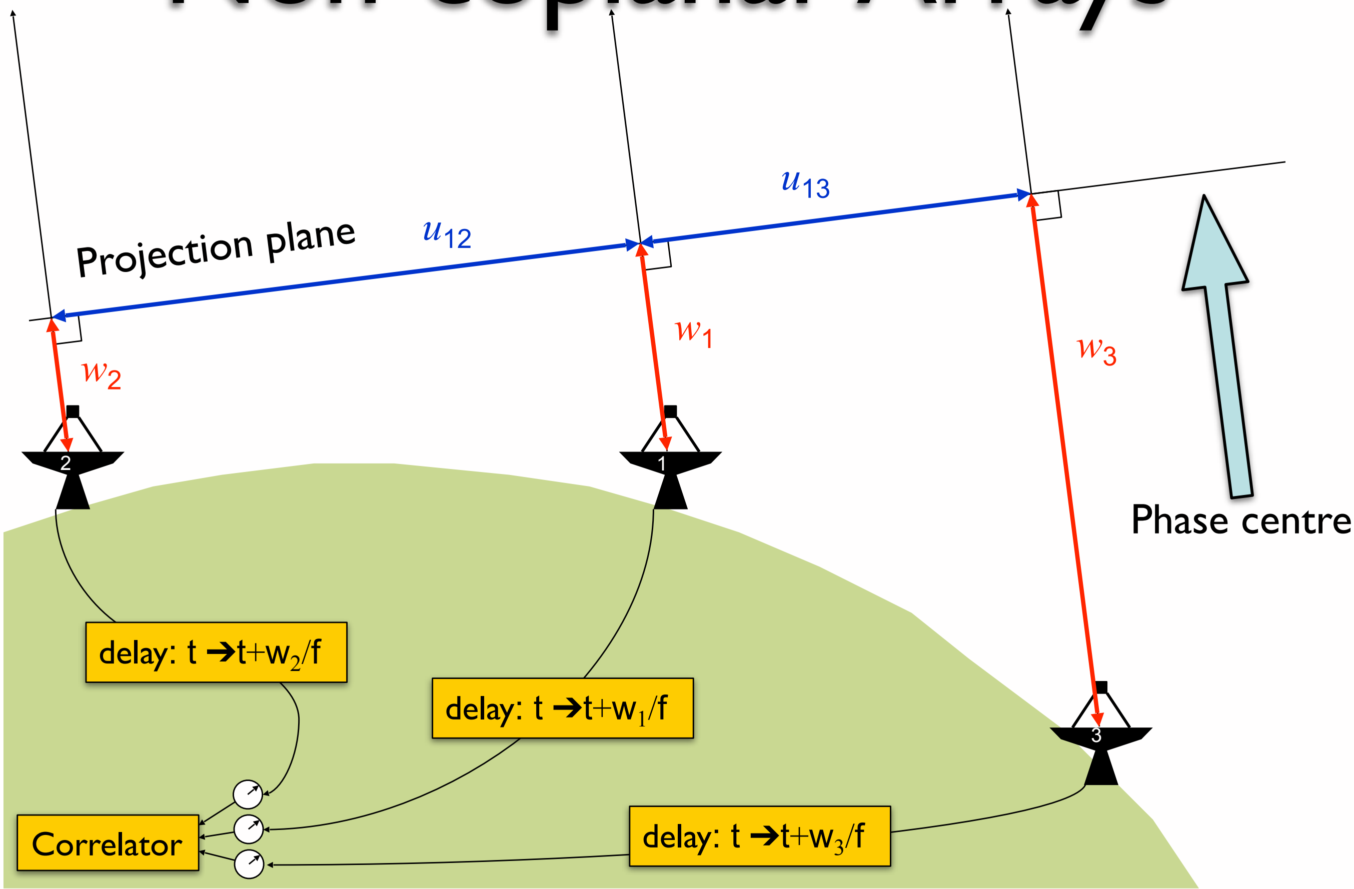

3-D Interferometers

- What if the interferometer does not measure the coherence function on a plane, but rather does it through a volume? In this case, we adopt a different coordinate system. First we write out the full expression:

$$V_\nu(\mathrm{u},\mathrm{v},\mathrm{w}) = \iint \frac{I_\nu(l,m)}{\sqrt{1-l^2-m^2}} e^{-2i\pi(\mathrm{u}l+\mathrm{v}m+\mathrm{w}n)} dldm$$

(Note that this is not a 3-D Fourier Transform).

- We orient the w-axis of the coordinate system to point to the region of interest. The u-axis point east, and the v-axis to the north.

- Then, remembering that: $n^2 = 1 - l^2 + m^2$

$$V_\nu(u,v) = \iint \frac{I_\nu(l,m)}{\sqrt{1-l^2-m^2}} e^{-2i\pi[ul+vm+w(\sqrt{1-l^2-m^2}-1)]} dldm$$

3-d to 2-d

- The expression is still not a proper Fourier transform. We can get a 2-d FT if the third term in the phase factor is sufficient small.

- The third term in the phase can be neglected if it is much less than unity:

$$w\left[1-\sqrt{1-l^2-m^2}\right]=w(1-cos\theta)\sim w\theta^2/2\ll 1$$

- This condition holds when: (angles in radians!)

$$\theta_{\max} < \sqrt{\frac{1}{2w}} \sim \sqrt{\frac{\lambda}{B}} \sim \sqrt{\theta_{syn}}$$

- If this condition is met, then the relation between the Intensity and the Visibility again becomes a 2-dimensional Fourier transform:

$$V_\nu(u,v)=\iint \frac{I_\upsilon(l,m)}{\sqrt{1-l^2-m^2}}e^{-i2\pi(ul+vm)}\,dldm$$

Non-coplanar Baselines

- Use of the 2-D transform for non-coplanar interferometer arrays (like the JVLA) always result in an error in the images.
- Formally, a 3-D transform can be constructed to handle this problem. In practice, we correct for it numerically.
- The errors increase inversely with array resolution, and quadratically with image field of view.
- For interferometers whose field-of-view is limited by the primary beam, low-frequencies are the most affected.
- Then,

$$\theta_{max} = \frac{\lambda}{D} < \sqrt{\frac{\lambda}{B}}$$

- Or, if $\frac{\lambda B}{D^2} < 1$ you've got trouble!

Example: JVLA

- The table below shows the approximate situation for the JVLA, when it is used to image its entire primary beam.
- Blue numbers show the respective primary beam FWHM
- Green numbers show situations where the 2-D approximation is safe.
- Red numbers show where the approximation fails totally.

λ	θ_{FWHM}	A	B	C	D
6 cm	9’	6’	10’	17’	31’
20 cm	30’	10’	18’	32’	56’
90 cm	135’	21’	37’	66’	118’
400 cm	600’	45’	80’	142’	253’

Table showing the JVLA’s distortion free imaging range (green), marginal zone (yellow), and danger zone (red) for different configurations and frequencies

Some numbers showing where the 2D approximation is valid for the JVLA.
Data outside this safe zone will be of less usefulness *unless* we can correct for the 3D effects.
We can, but its computationally expensive.

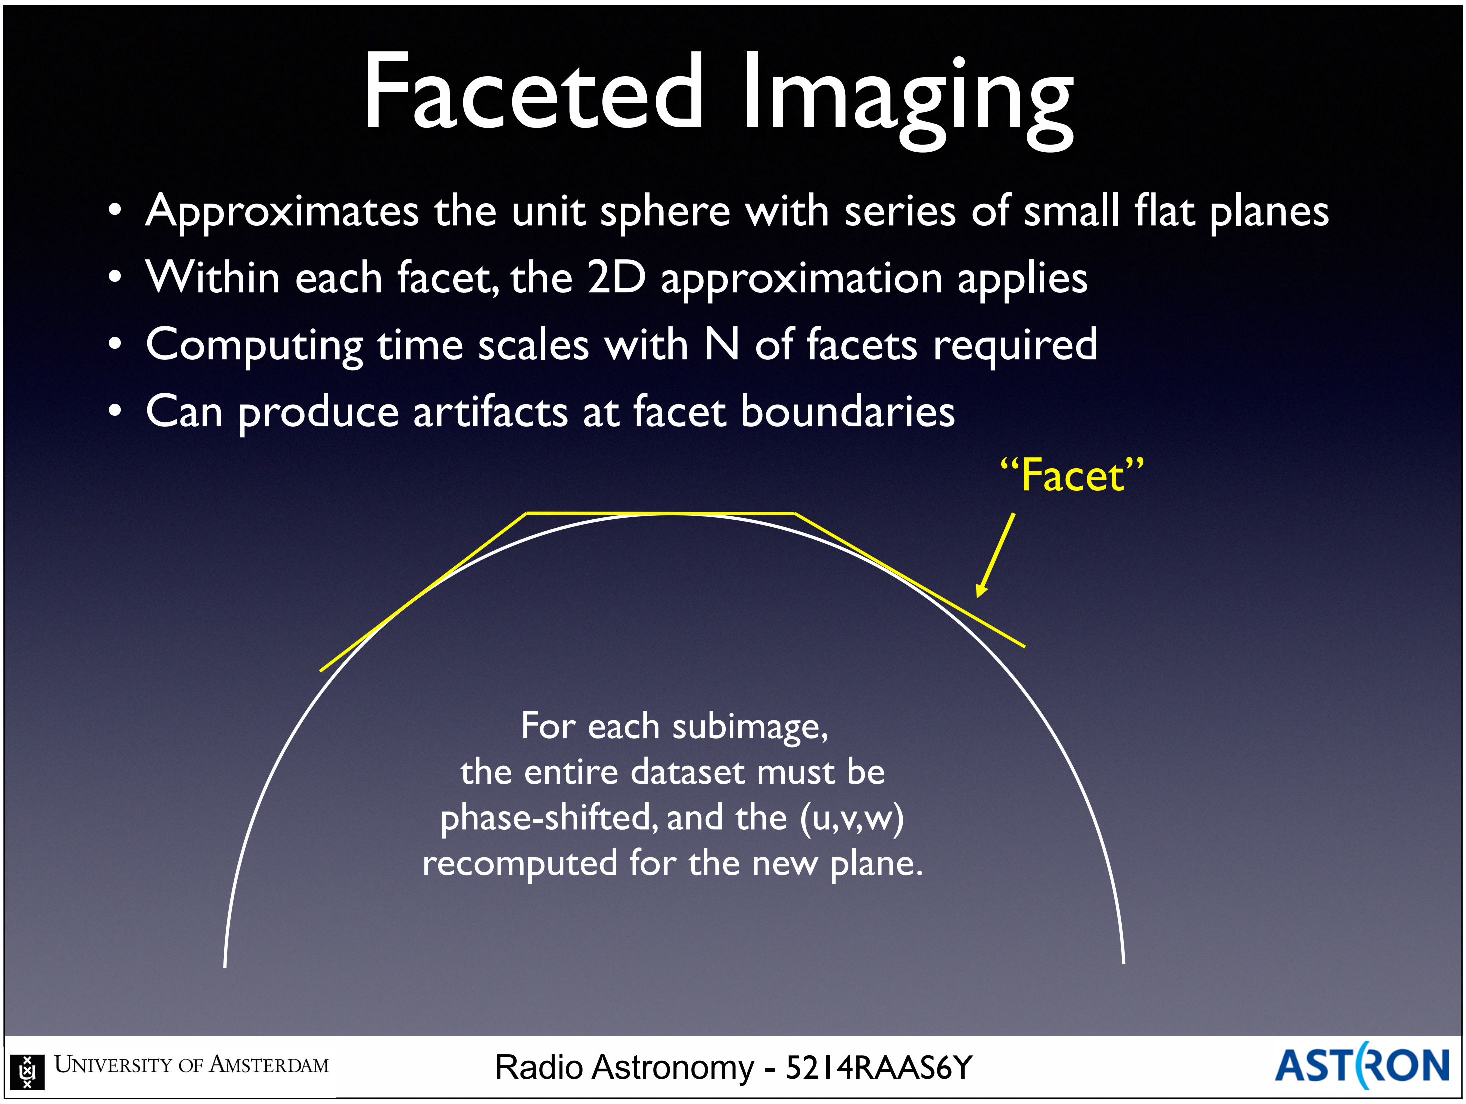

One technique to take the 3D effects into account is called "facet imaging".
Basic idea is to break the sky up into a series of tiles or facets where the 2D approximation still holds.
Each facet is imaged separately and then all the facet images are tiled together to form the final image.
Works reasonably well, but multiplies the cost of imaging by the number of facets used.
More facets —> Better final image —> More computational cost.

W-Projection

- Each visibility, at location (u,v,w) is mapped to the w=0 plane, with a phase shift proportional to the distance
- Each visibility is mapped to ALL the points lying within a cone whose full angle is the same as the field of view of the desired map (~$2\lambda/D$ for a full-field image)
- Area in the base of the cone is ~$4\lambda^2 w^2/D^2 < 4B^2/D^2$. Number of cells on the base which 'receive' this visibility is ~ $4w_0^2B^2/D^2 < 4B^4/\lambda^2D^2$

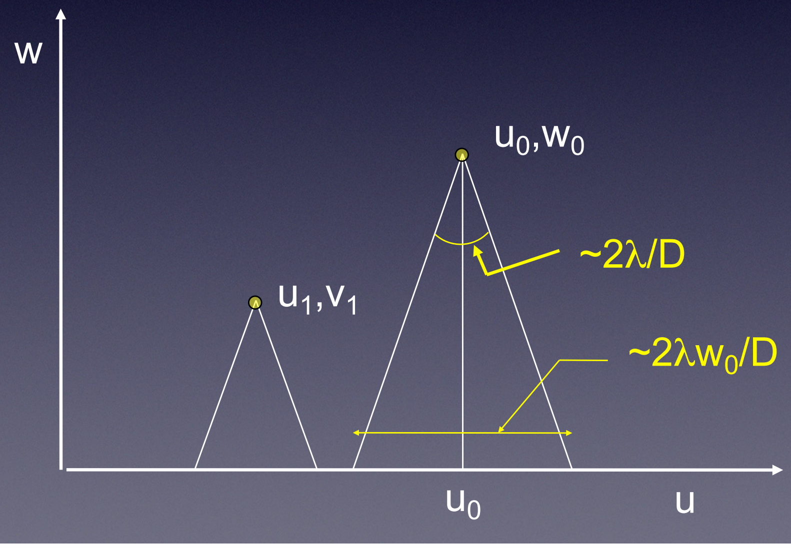

Another technique called w-projection accounts for the 3D effect by computing the image for a number of different points along the w axis.
Can produce good results, but like facet imaging increases the computing cost.
More w points —> Better final image —> More computational cost.

Without "3D" Processing

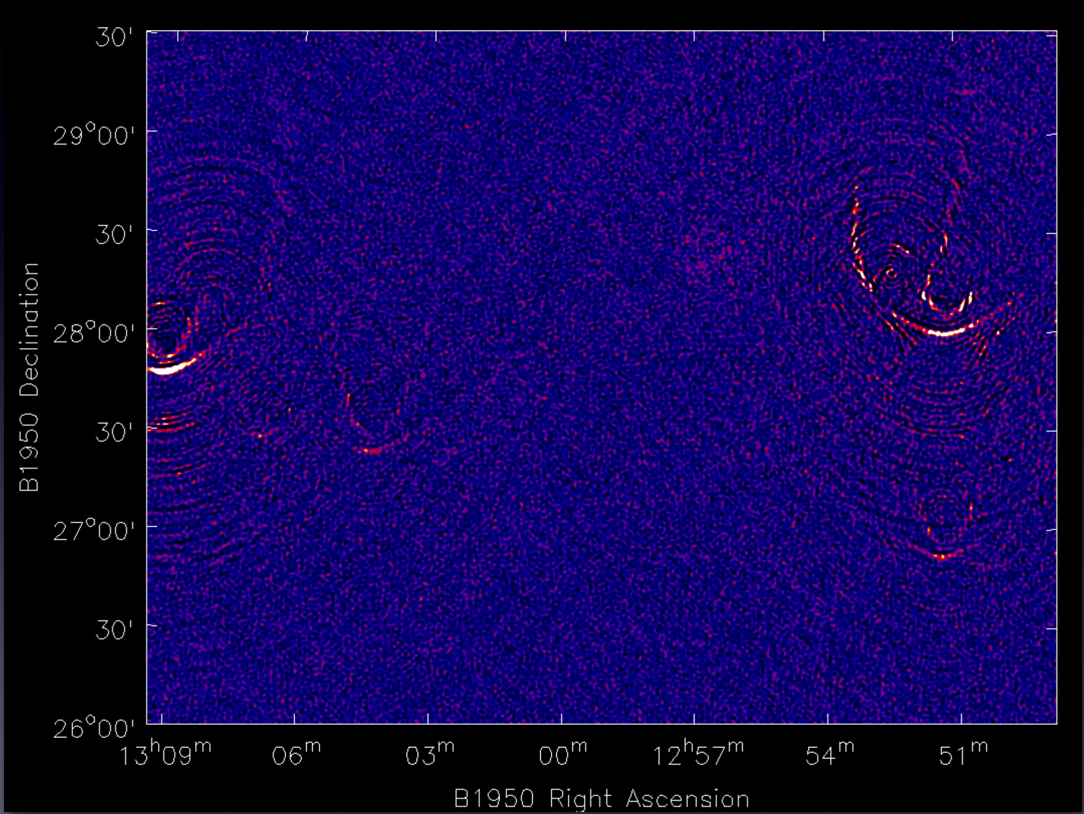

An example of a wide-field image. Notice the size of the field ~3.5 degrees.
This image has been processed assuming the 2D approximation holds.
Notice all the artifacts in the image for sources on the edge of the field of view.

With “3D” Processing

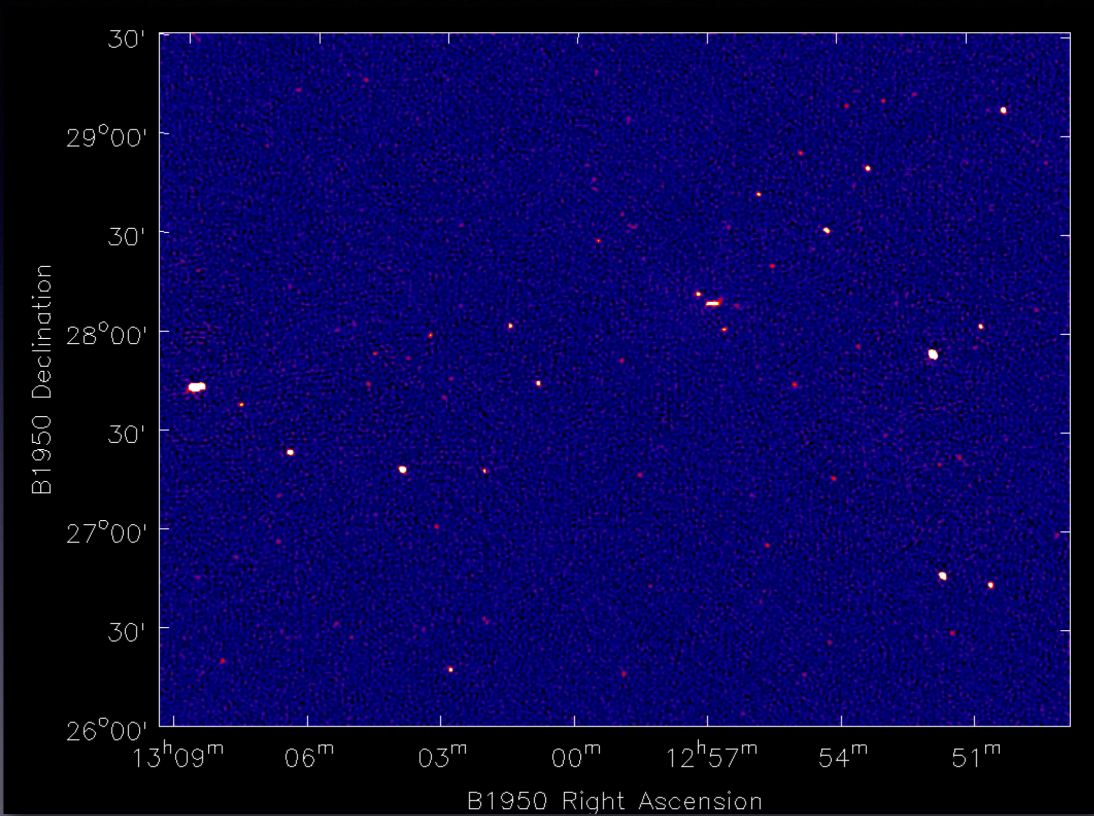

The same image including 3D effects into the imaging.
Far off-axis sources now look like point sources, as they should.

Comparison of Techniques

2D Imaging | *Facet Imaging* | *W-Projection*

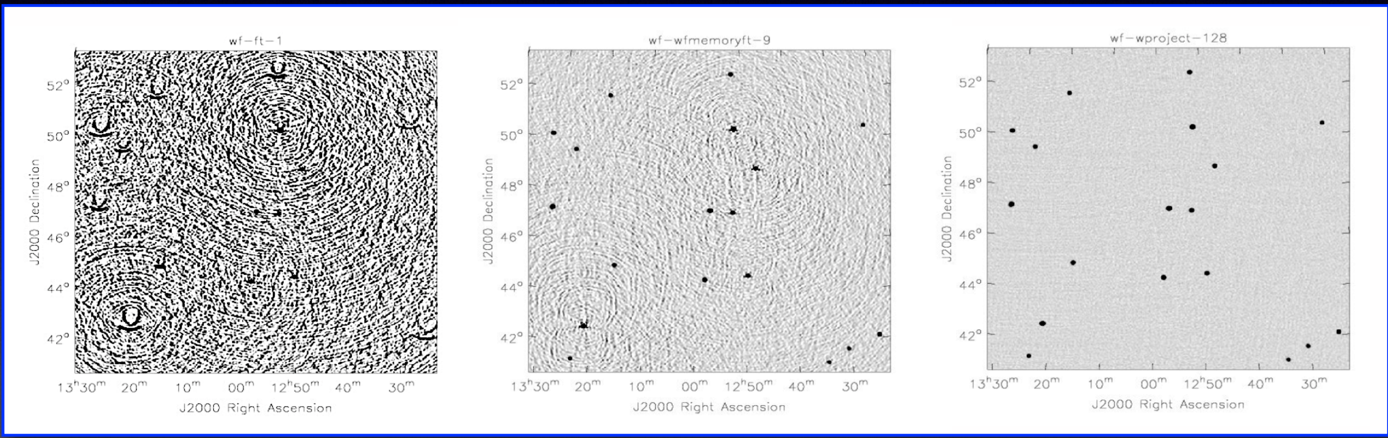

Some LOFAR images comparing the facet and w-projection techniques.
Notice the size of the field of view ~12 degrees!
LOFAR and low-frequency imaging are one of the big drivers for developing these techniques.

Mosaicing

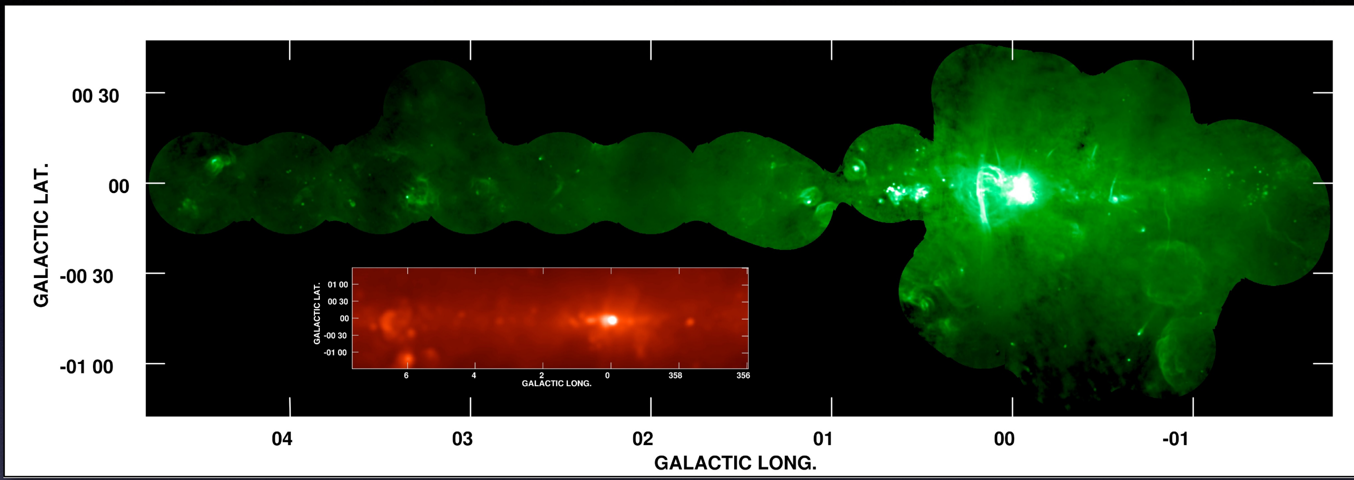

- Effects of varying primary beams must be taken into account
- Adds complexity to the deconvolution process
- Need adequate sky coverage (try to keep Nyquist sampling)
- Can also be used to add single dish data and recover zero spacings

Mosaicing Techniques

- Primary Methods
  - Linear combination of deconvolved maps
  - Joint deconvolution
  - Regridding of all visibilities before FFT

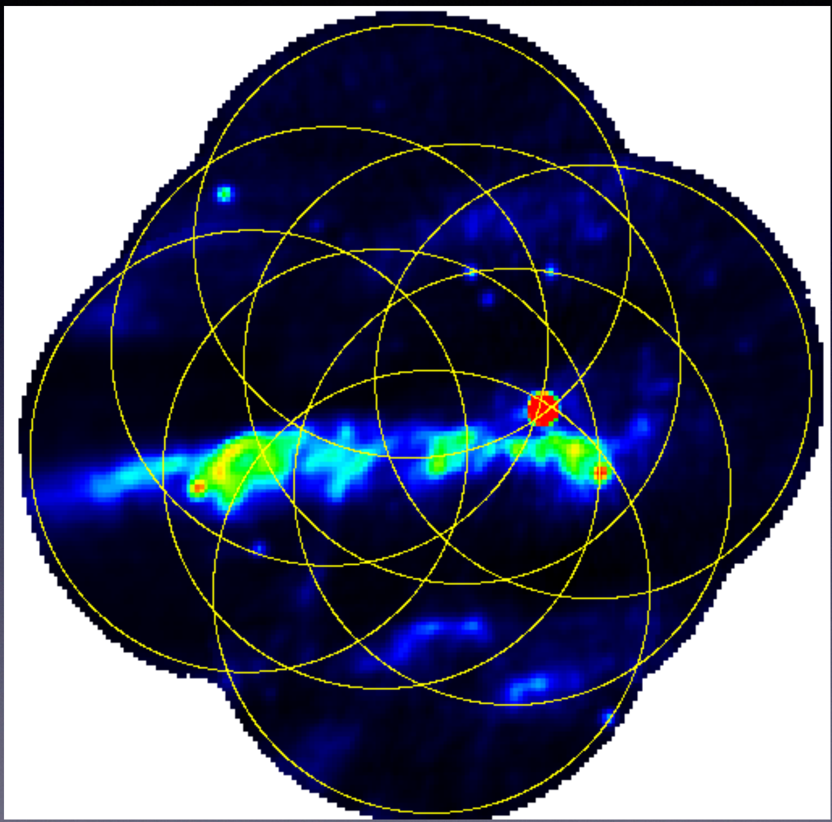

Mosaicing Techniques

- Primary Methods
  - Linear combination of deconvolved maps
  - Joint deconvolution
  - Regridding of all visibilities before FFT

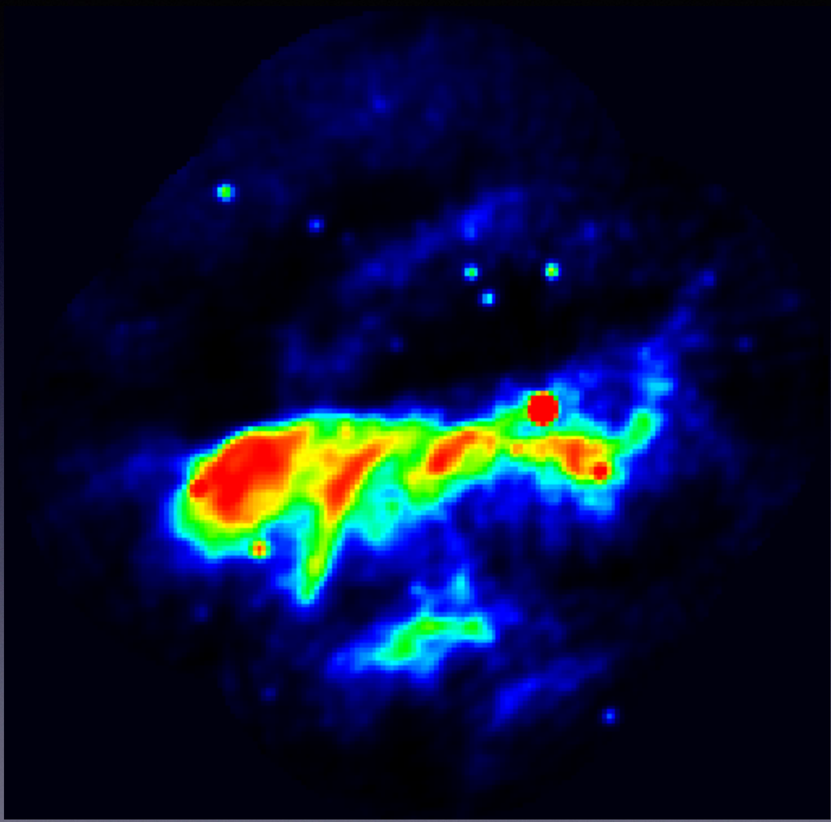

Zero Spacings

Orion Nebula

Shepherd, Maddalena, McMullin (2002)

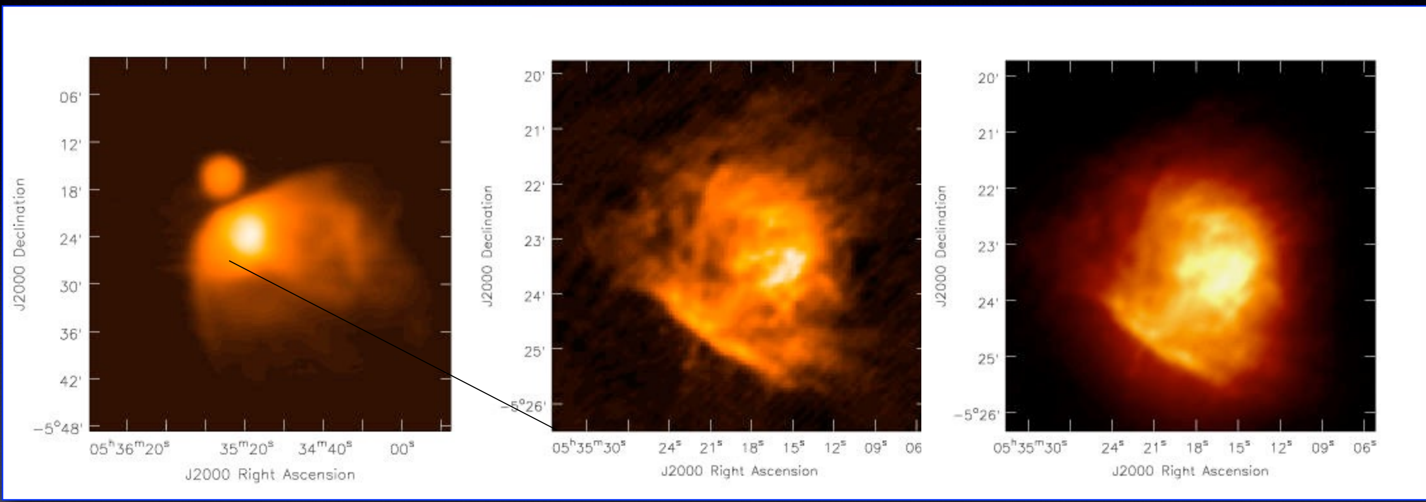

GBT map of the large field
90” resolution

VLA mosaic of central region, 9 fields
8.4” resolution

Combined GBT+VLA mosaic deconvolved with multi-scale CLEAN

Final image fidelity significantly better

Questions?

Practicum

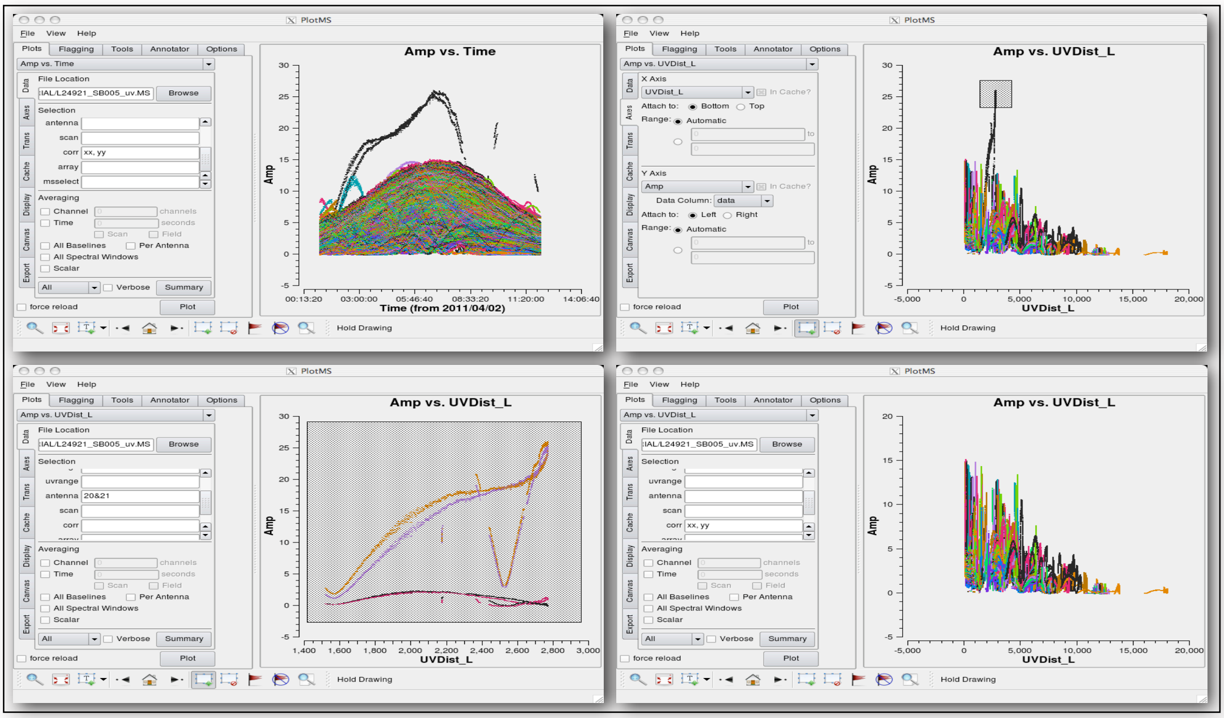

- Examine, calibrate, and image an actual radio data set

Westerbork/LOFAR Field Trip

Thursday, May 18th, 2017

Train to Hoogeven:

Depart Amsterdam Centraal: 8:08
Transfer in Almere and Zwolle
Arrive Hoogeveen: 9:53

Itinerary:

10:00 - Pick-up at Station Hoogeveen
10:30 - Presentation/tour at ASTRON
12:00 - Lunch at ASTRON
12:45 - Depart to Westerbork
13:30 - Tour Westerbork
15:00 - Depart to LOFAR
15:30 - Tour of LOFAR
17:00 - Arrive Station Hoogeveen

ASTRON Reception
+31 (0) 521 595 100