Calibration & Imaging at low frequencies

André Offringa (ASTRON)

SALF IV Sydney, Australia, 2017-12-14

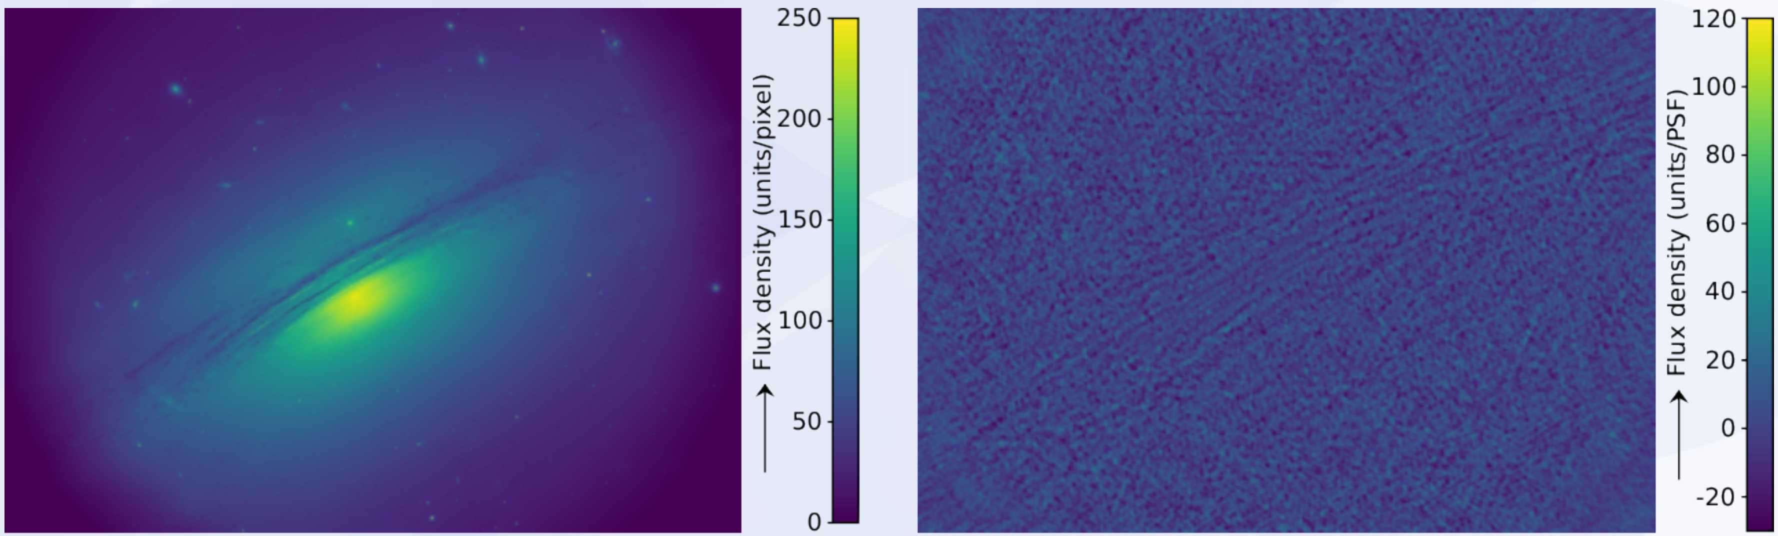

(a) Reconstructed model image

(b) Residual image (σ=8.6 units/PSF)

Automatic scale-dependent masking applied on the UGC12591 test-set.

European Research Council
Established by the European Commission
Supporting top researchers from anywhere in the world

Dynamical Statistical Compression (Dysco)

- Transparent *lossy* compression of visibilities
- Compresses by a factor of 4 with <1% noise increase
- Being rolled out to LOFAR archive

Published in Offringa (2016)

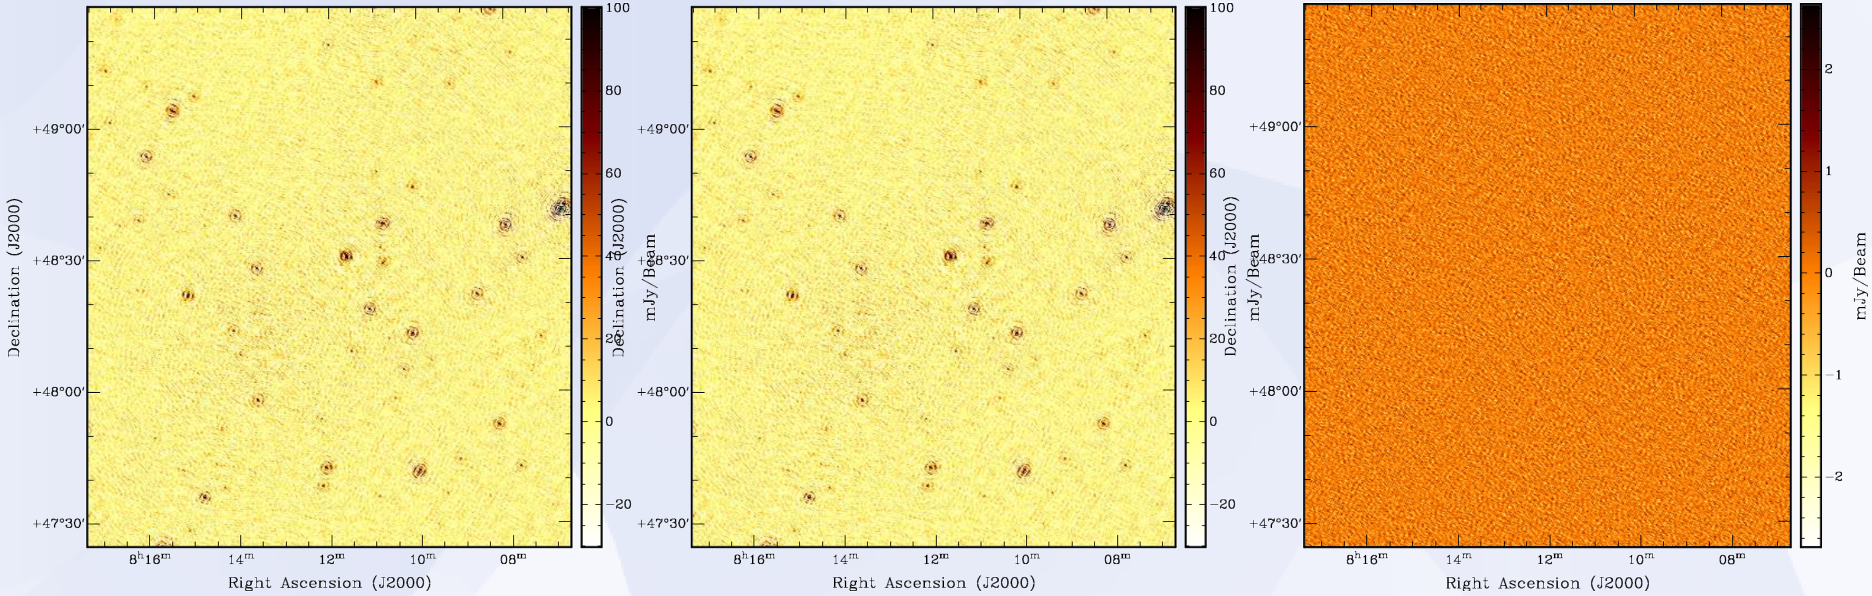

Uncompressed

4x compressed
(8 bit compression)

Difference
(Unstructured noise)
Rms of 400 microJy

Dynamical Statistical Compression (Dysco)

- Transparent *lossy* compression of visibilities
- Compre… noise increase
- Being r…

…shed in Offringa (2016)

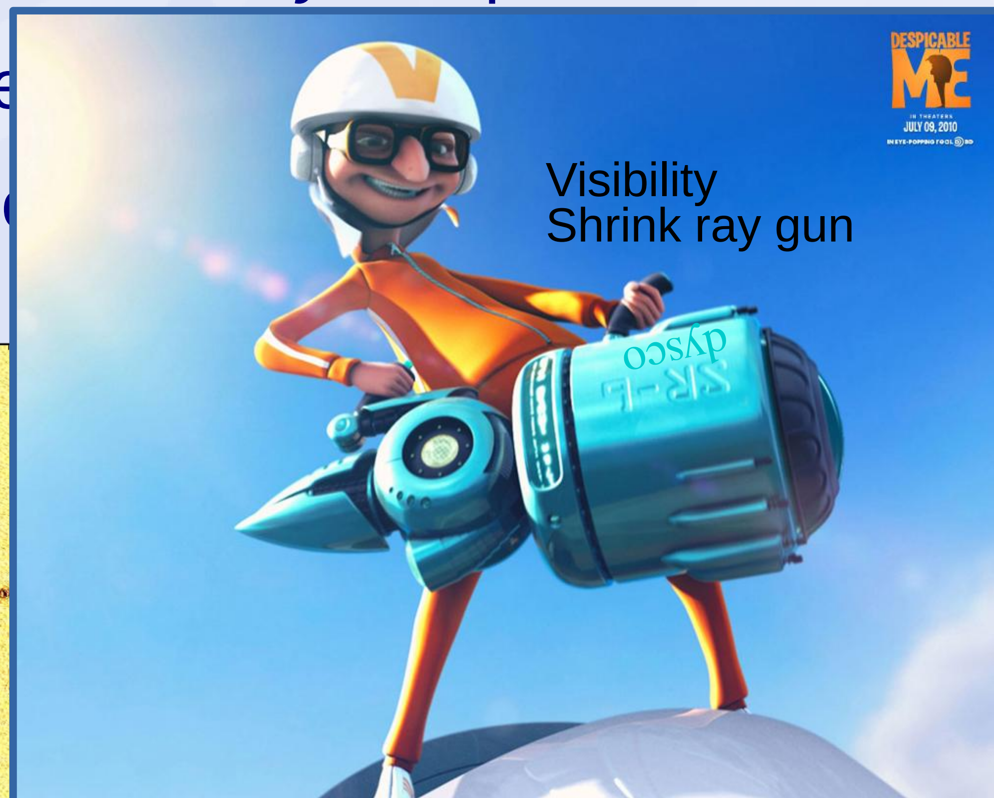

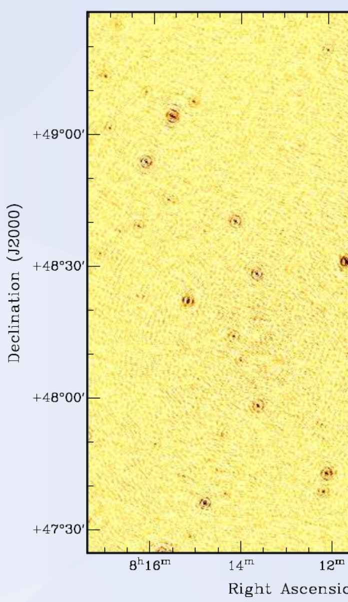
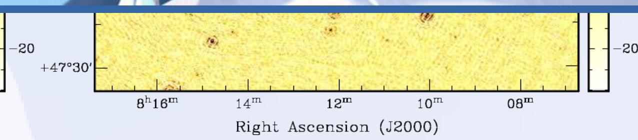
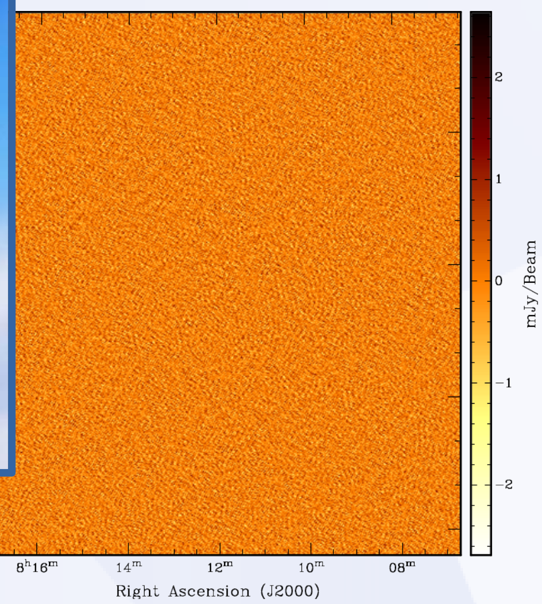

Uncompressed

4x compressed
(8 bit compression)

Difference
(Unstructured noise)
Rms of 400 microJy

Challenges in low-frequency imaging

- Large FOV
  - Large w-values
  - Harder to deconvolve
- Large fractional bandwidth
  - Requires multi-frequency approaches
- Large data volumes
- LF beams time dependent & difficult to model
- Calibration errors are higher
- Requires direction-dependent cal

Multi-frequency deconvolution

- Common approach in MF deconvolution is imaging / predicting "frequency derivative" images ("nterms>1", the Sault & Wieringa (1994) method).

That results in:

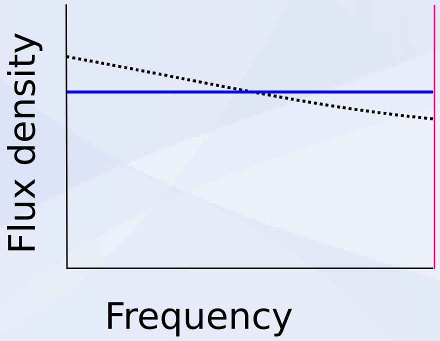

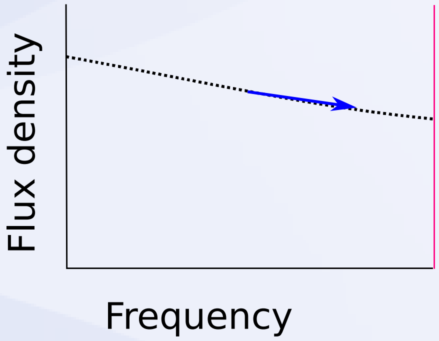

Instead, WSClean splits the bandwidth and creates separate images for each part:

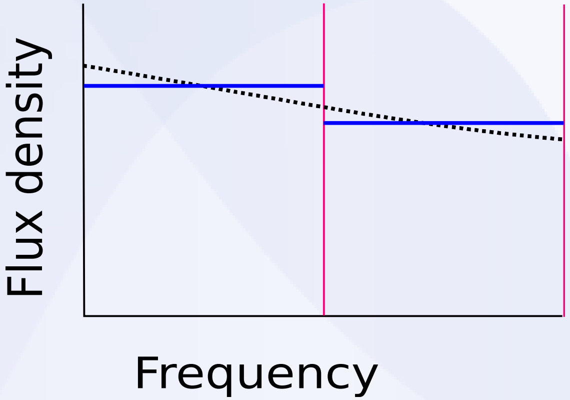

(Similar strategy is used by B. Cotton's OBIT)

Multi-frequency deconvolution

- Common approach in MF deconvolution is imaging / ... derivative" images ("nterms>1", the Sault &

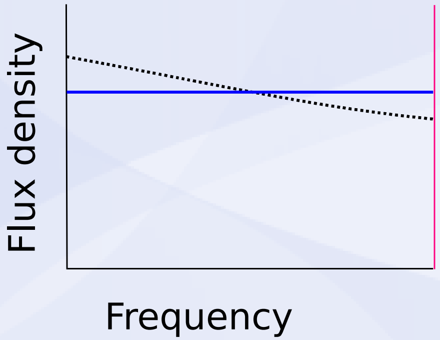

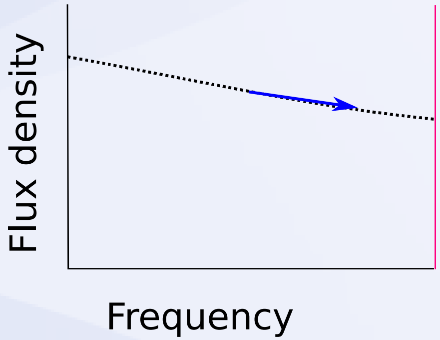

- Of course, these contain the same information

  (they can be converted from one to the other)

- Algorithm to clean second option is simpler.

...an splits the ... and creates ...r each part:

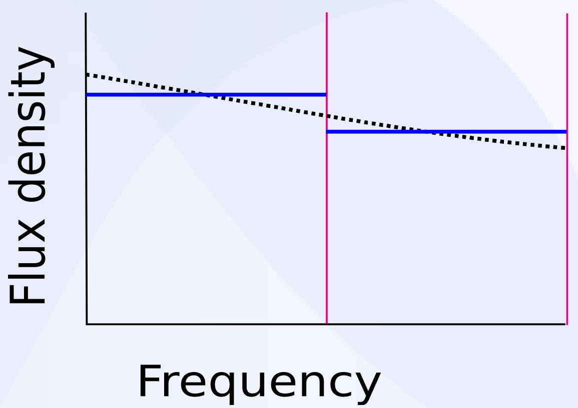

(S... on's OBIT)

Multi-scale kernel

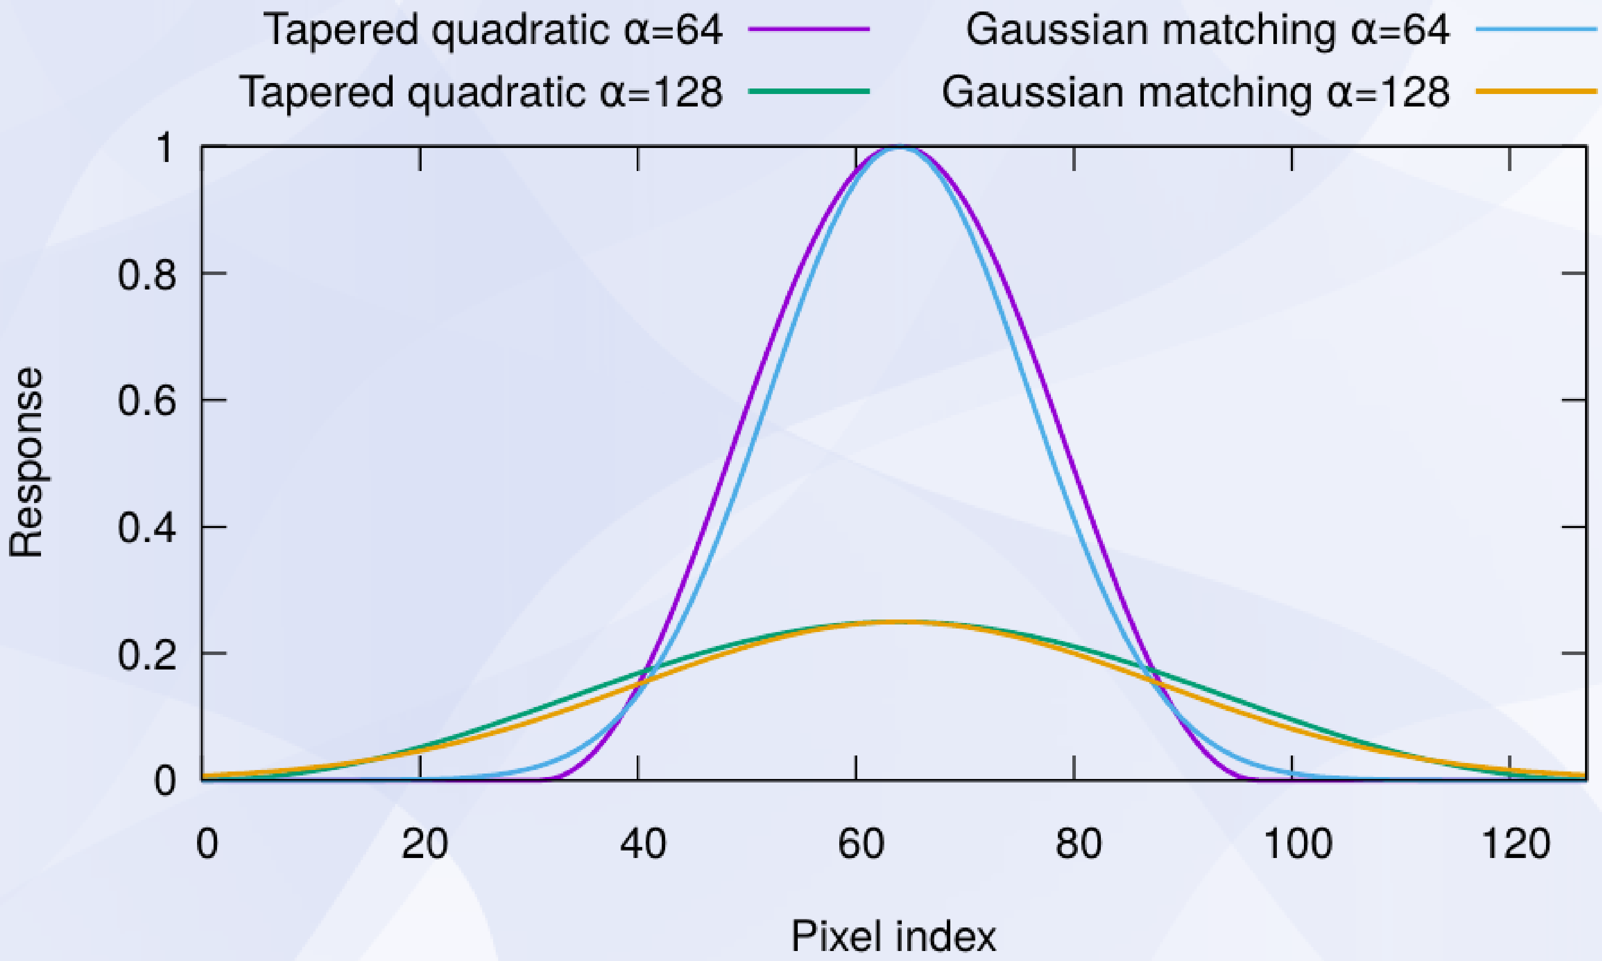

Figure 1. Shape functions for scales $\alpha = 64$ pixels and $\alpha = 128$ pixels.

Fast multi-scale deconvolution

- In Cornwell’s (2008) multi-scale method, the appropriate scale is determined every minor iteration
- Cornwell’s algorithm can be sped up by keeping the scale fixed “for a while”
- This is the algorithm implemented in WSClean

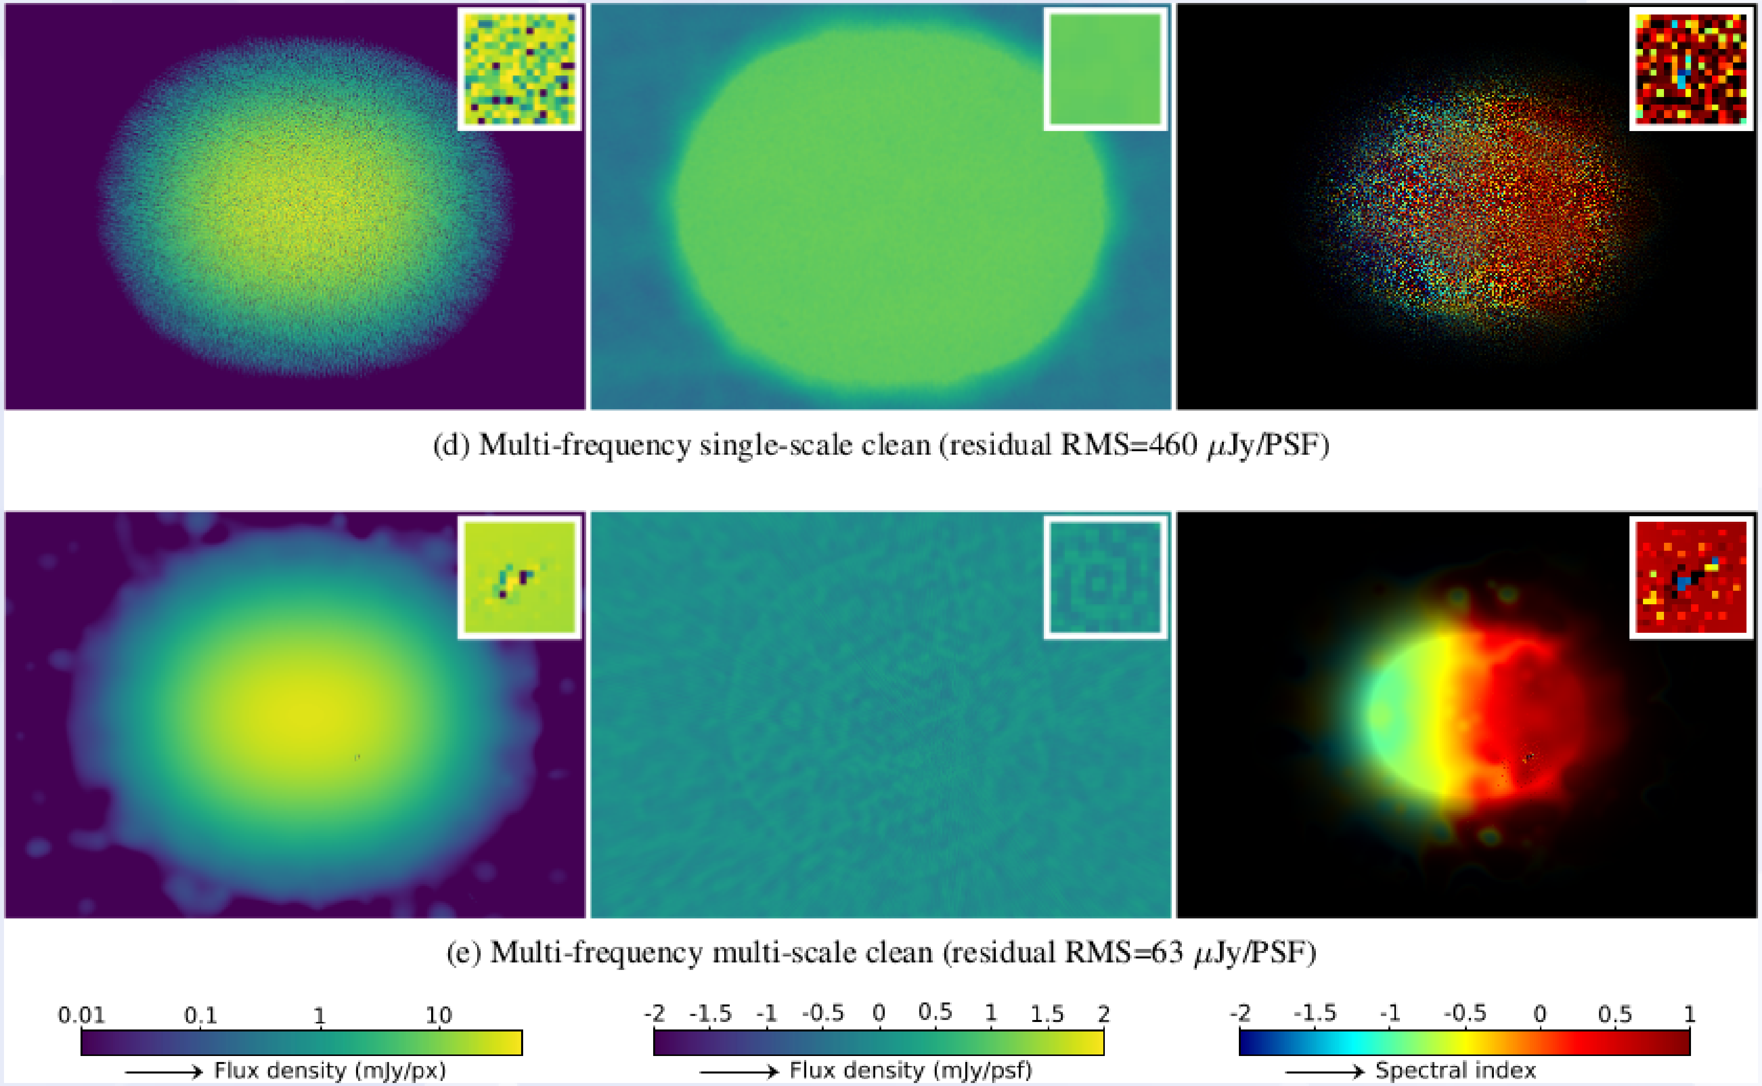

(d) Multi-frequency single-scale clean (residual RMS=460 μJy/PSF)

(e) Multi-frequency multi-scale clean (residual RMS=63 μJy/PSF)

- Comparison of WSClean MF single scale and multi-scale cleaning
- Simulated bandwidth of 30 MHz at 150 MHz.
- MWA layout, 2 min snapshot

Offringa and Smirnov (2017)

Deconvolution performance

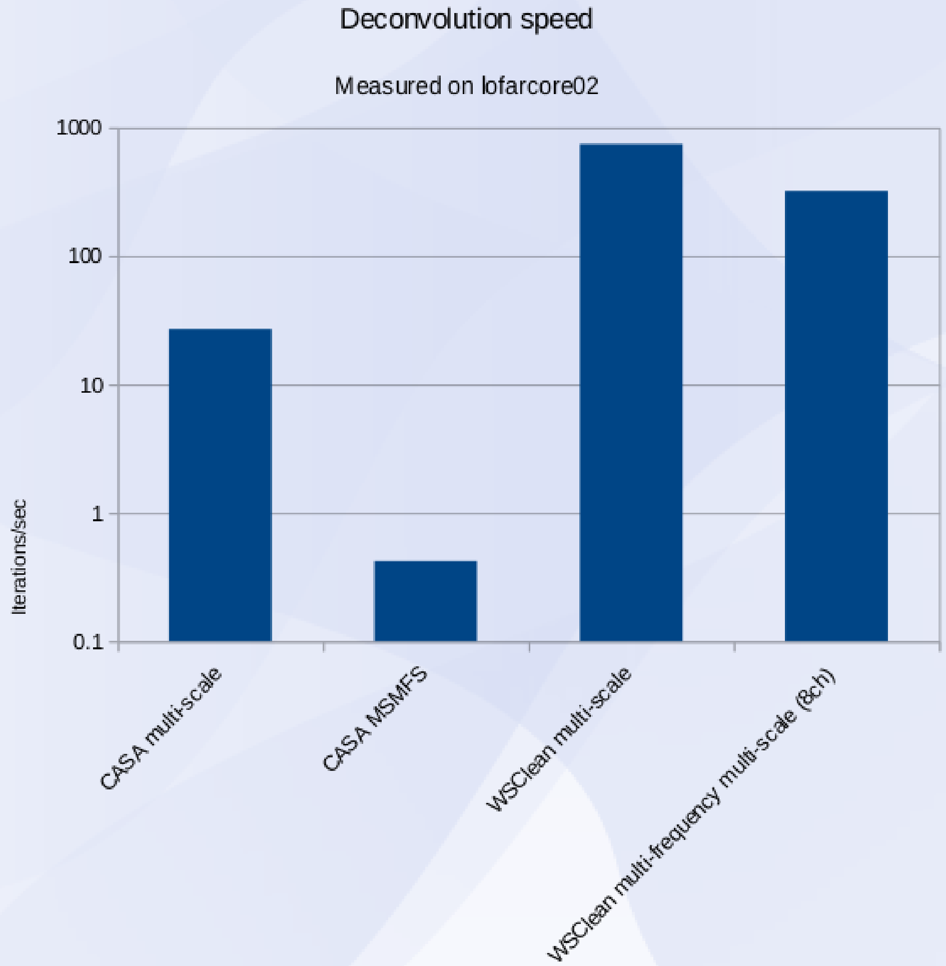

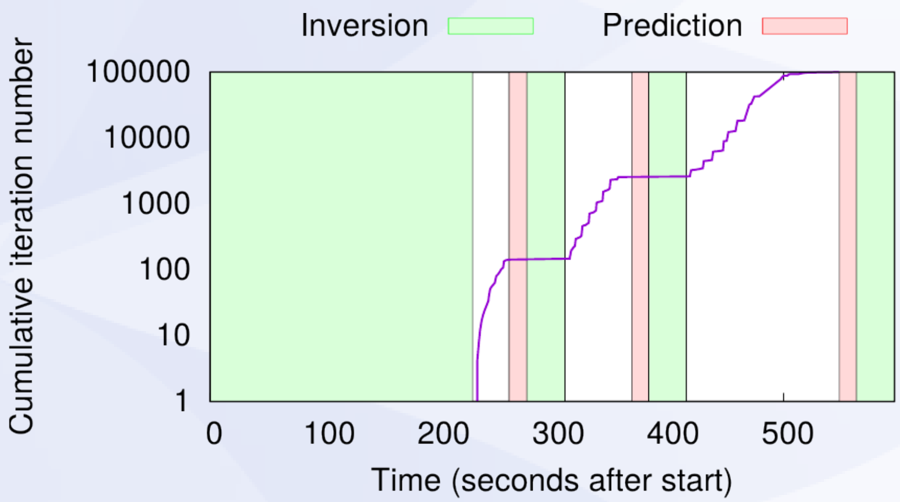

Figure 12. Example of the progression over time when using the new multi-scale clean algorithm on a 2048×2048 image.

Offringa and Smirnov (2017)

Local RMS cleaning

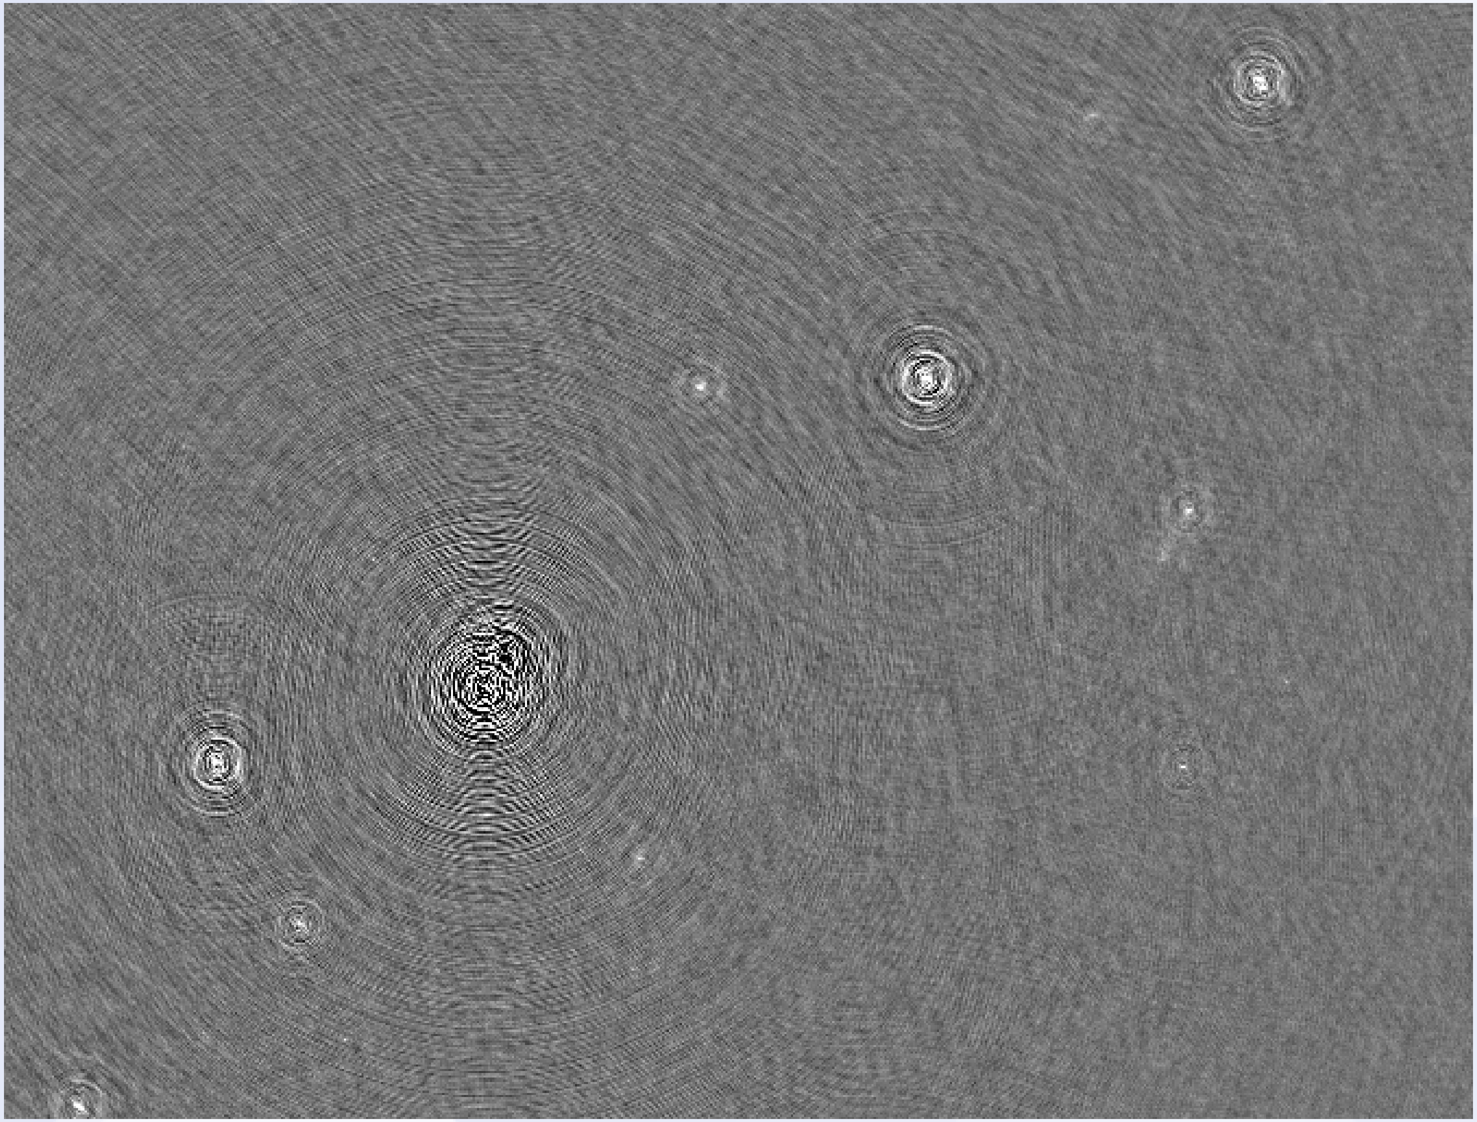

Local RMS cleaning

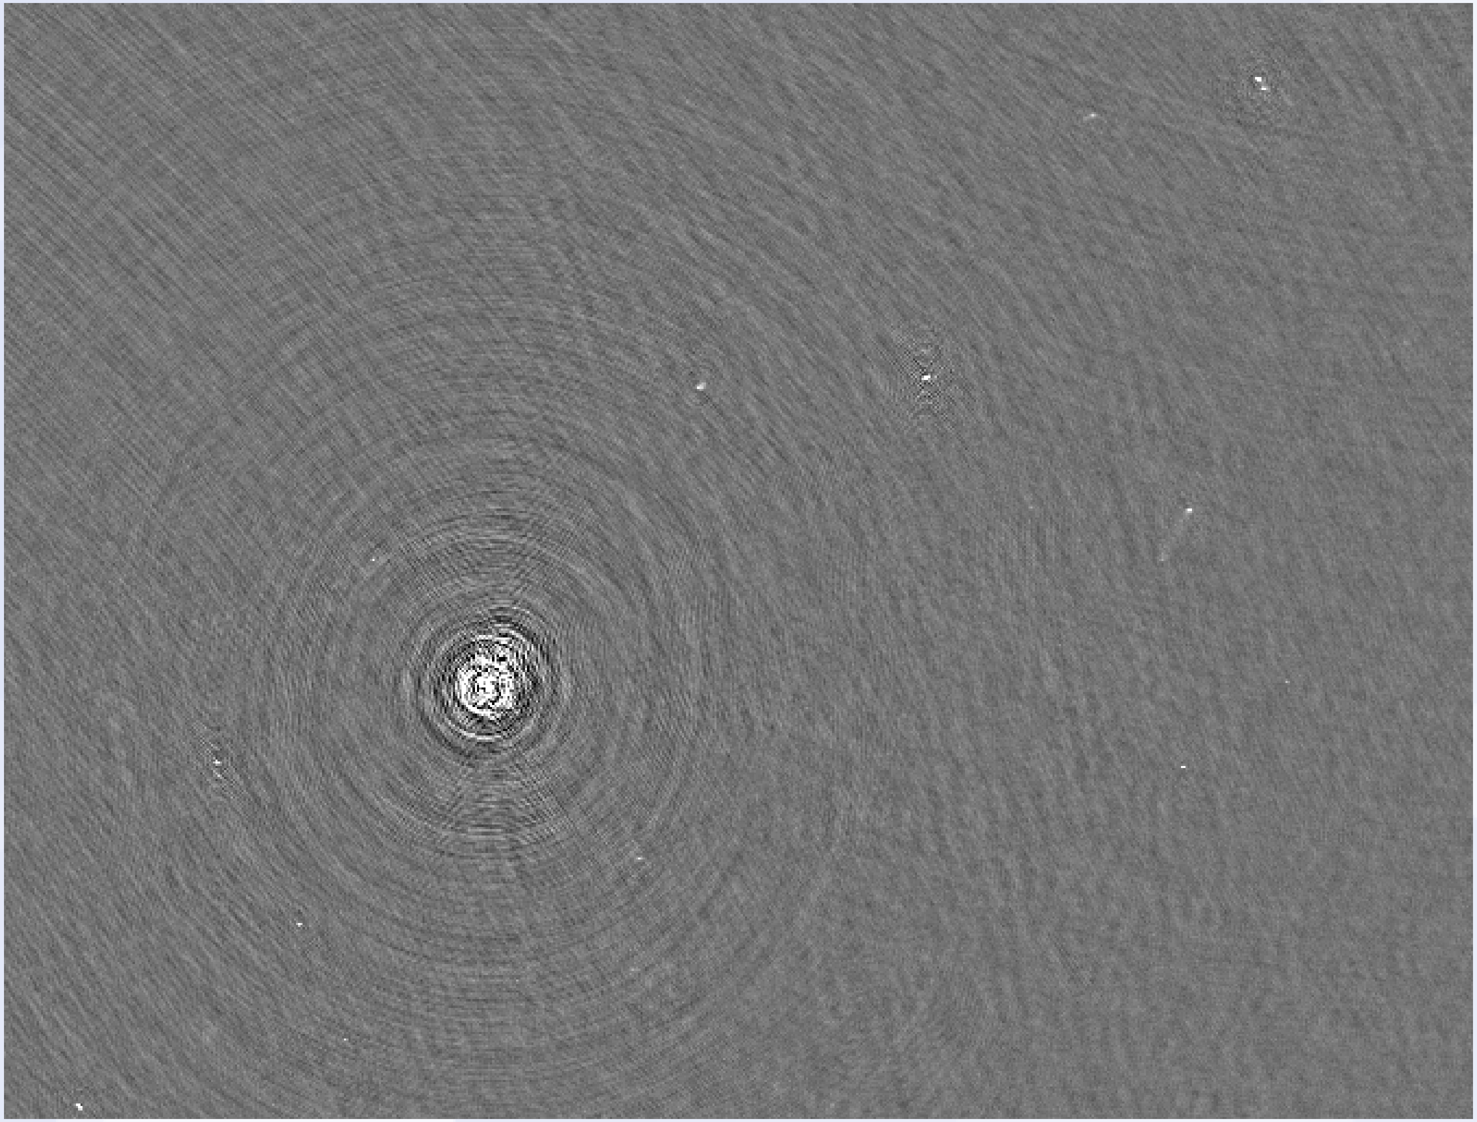

Polarized cleaning

- Standard iquv imaging: minimize sum pol^2
  - Available in CASA, WSClean, ...
- WSClean supports some RM cleaning methods
  - E.g., sum-over-squared Q/U pol & freq cleaning
- Since 2.6, also “linked polarization” cleaning
  - Base cleaning of subset of pols on others
  - E.g. search components in XX/YY, also remove from XY/YX.

Automatic scale-dependent masking

- Normal cleaning requires manual threshold tweaking, manual masking, etc…
- Masking is hard when structures are diffuse
- Move towards non-interactive, fully automatic cleaning
- “Automatic scale-dependent masking” :
  - For each scale, a mask is accumulated
  - Clean normal to 3-5σ, continue to 0.5σ with a scale-dependent mask. In one run.

Automatic masking

- Threshold is relative to RMS estimate
- RMS estimate can be “local” when RMS is expected to change over the image
  (avoids picking up calibration errors)
- Avoids interaction & somewhat-arbitrary selection of features, etc.
- Allows deeper & more stable cleaning of complex structures. Limits clean bias.
- Can be done in multi-frequency mode

Auto-masking on point sources

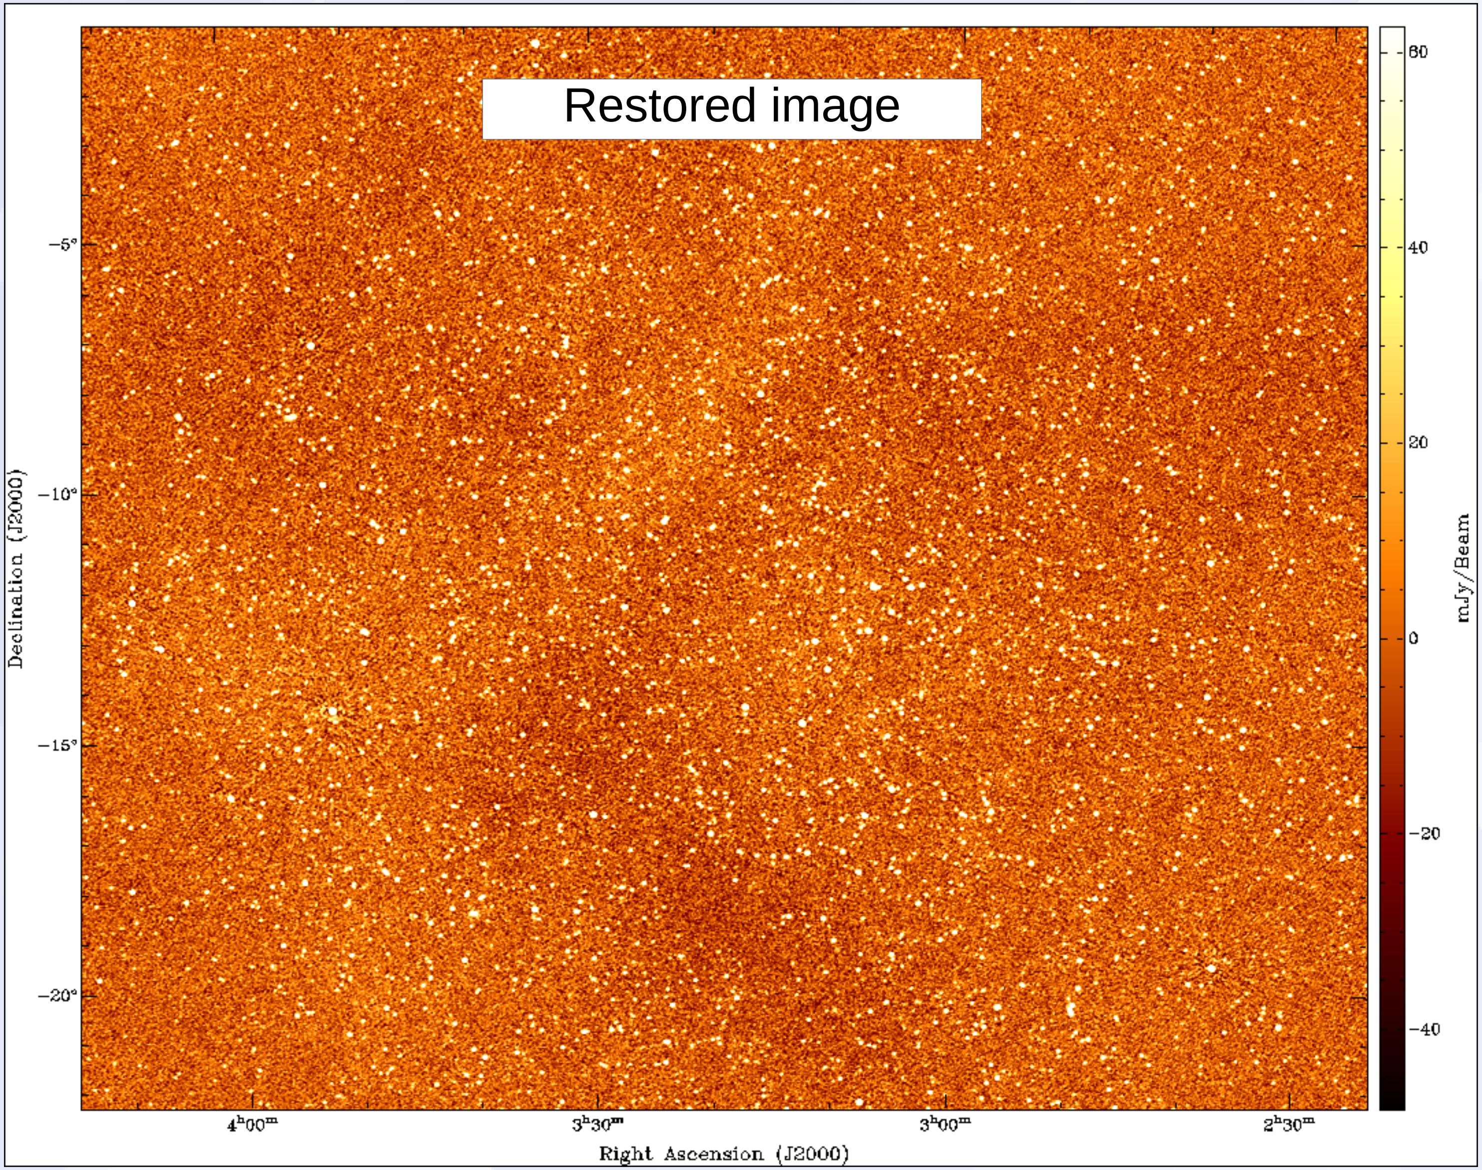

From data by T. Franzen

Auto-masking on point sources

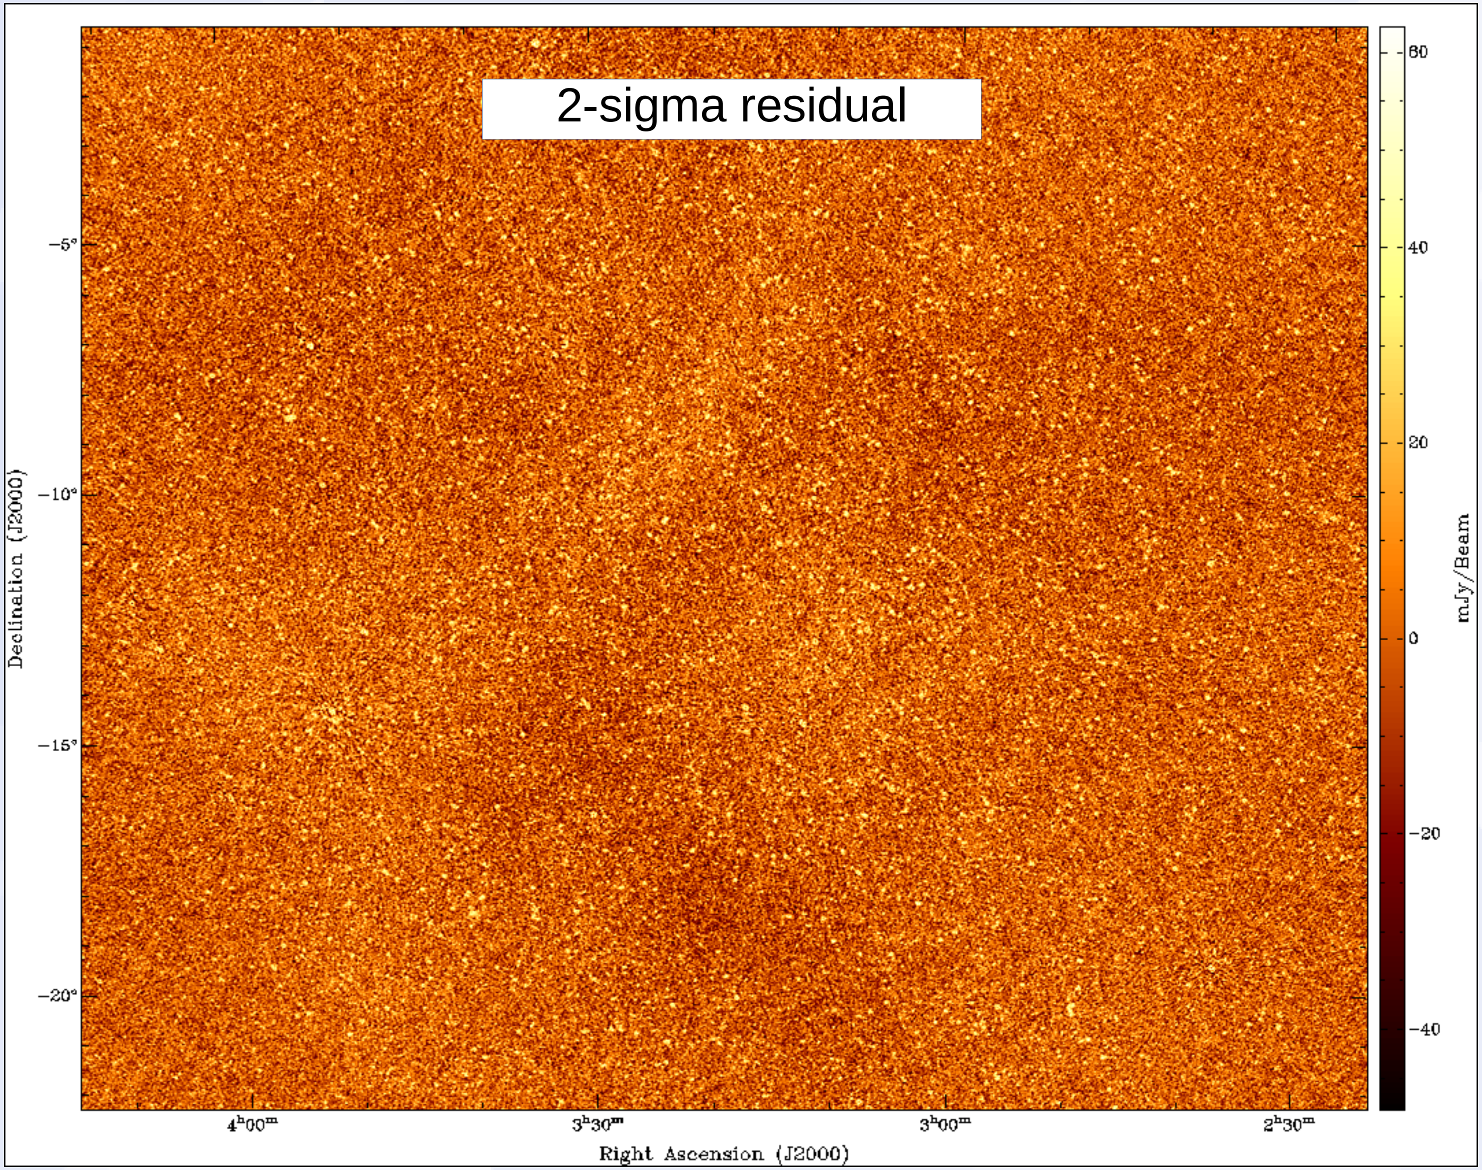

From data by T. Franzen

Auto-masking on point sources

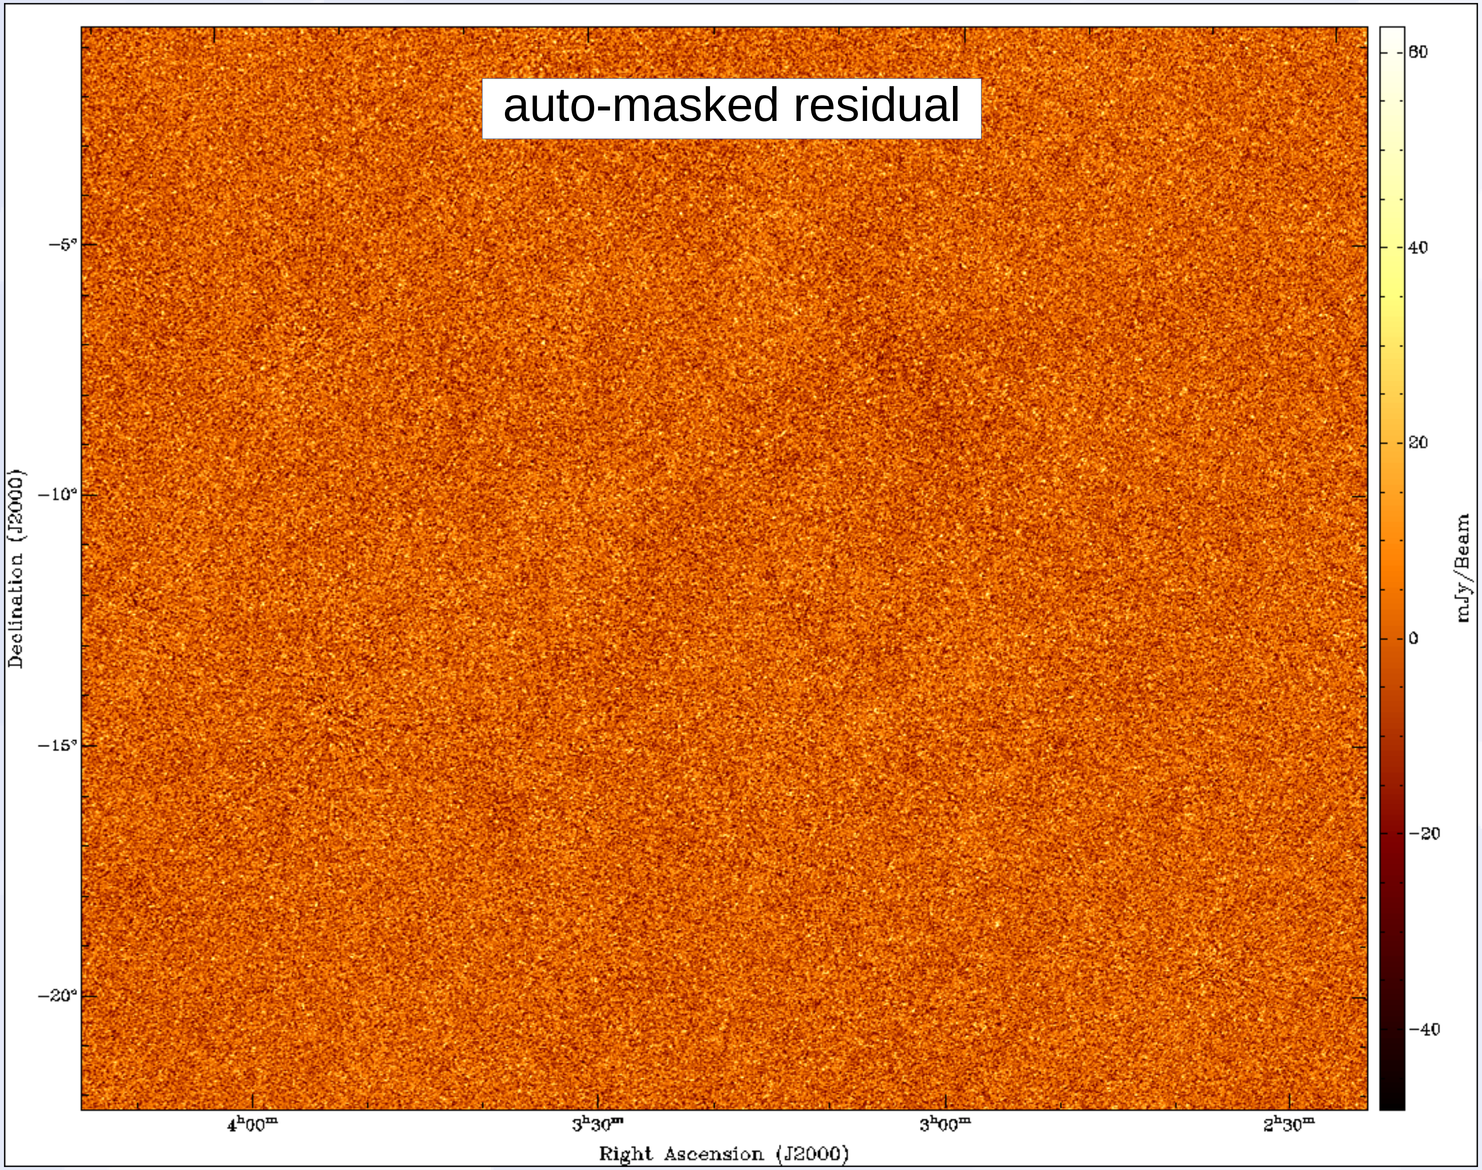

From data by T. Franzen

Automasking VLBI example

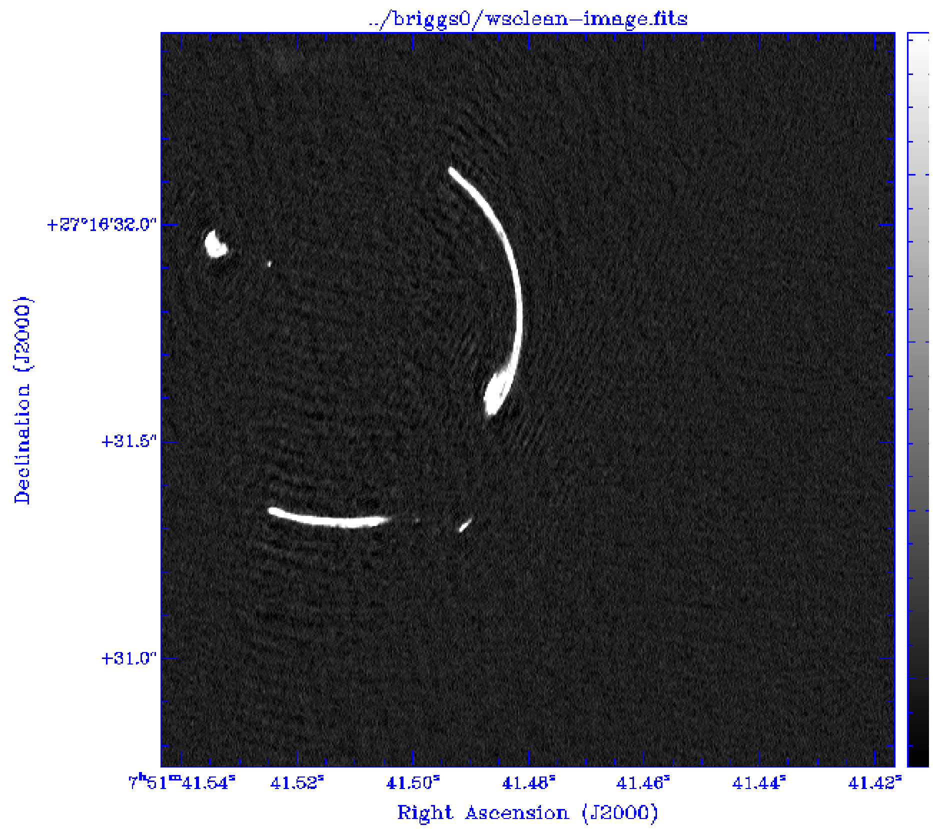

Restored

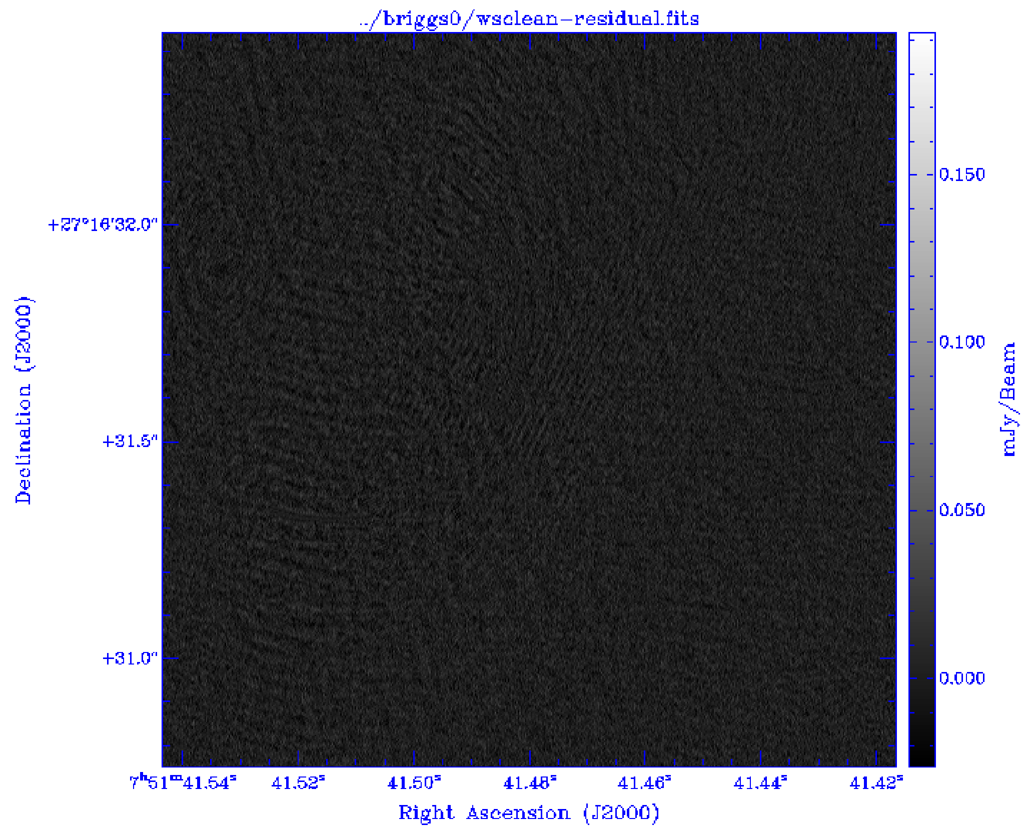

Residual

Data by J. P. McKean
and C. Spingola

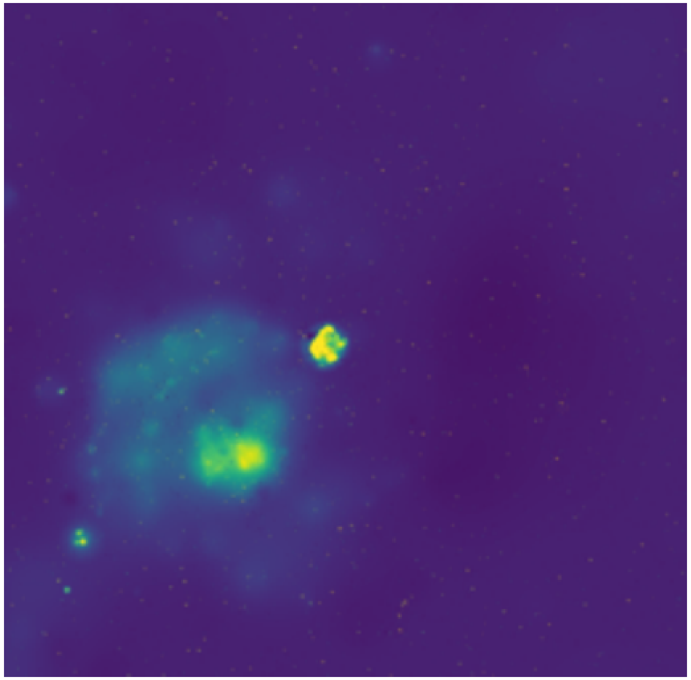

(a) Multi-scale model image without masking

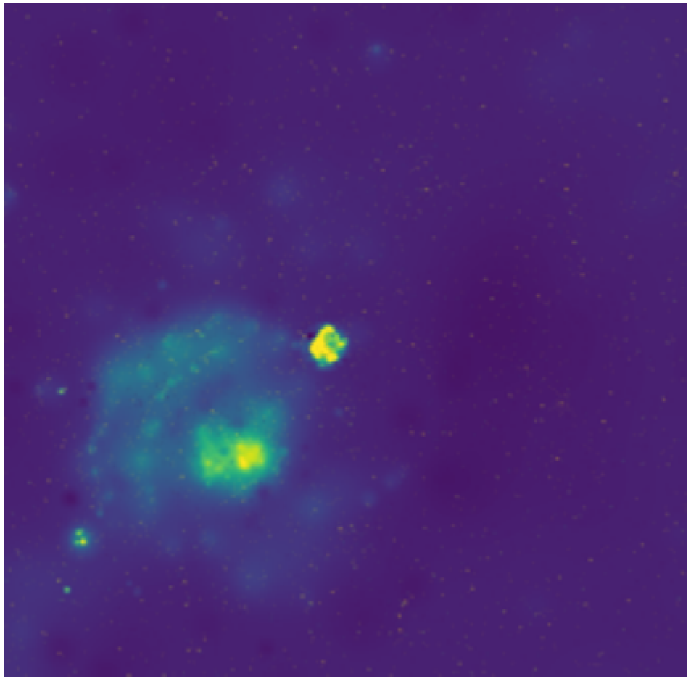

(b) Multi-scale model image with automatic masking

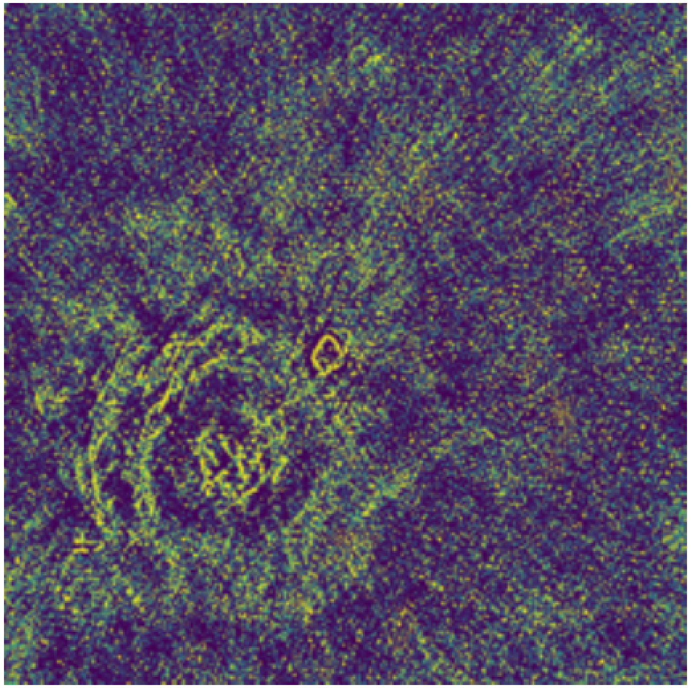

(c) Multi-scale residual without masking (rms=50 mJy/B)

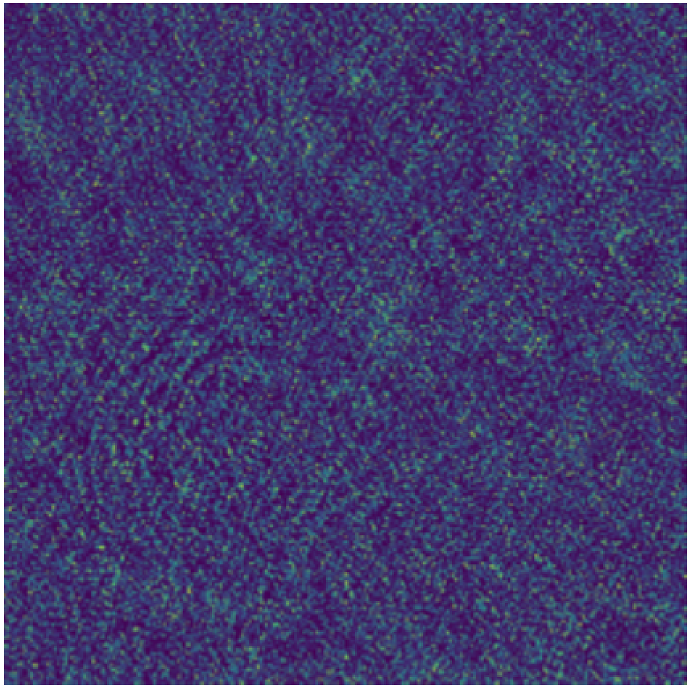

(d) Multi-scale residual with automatic masking (rms=38 mJy/B)

Offringa and Smirnov (2017)

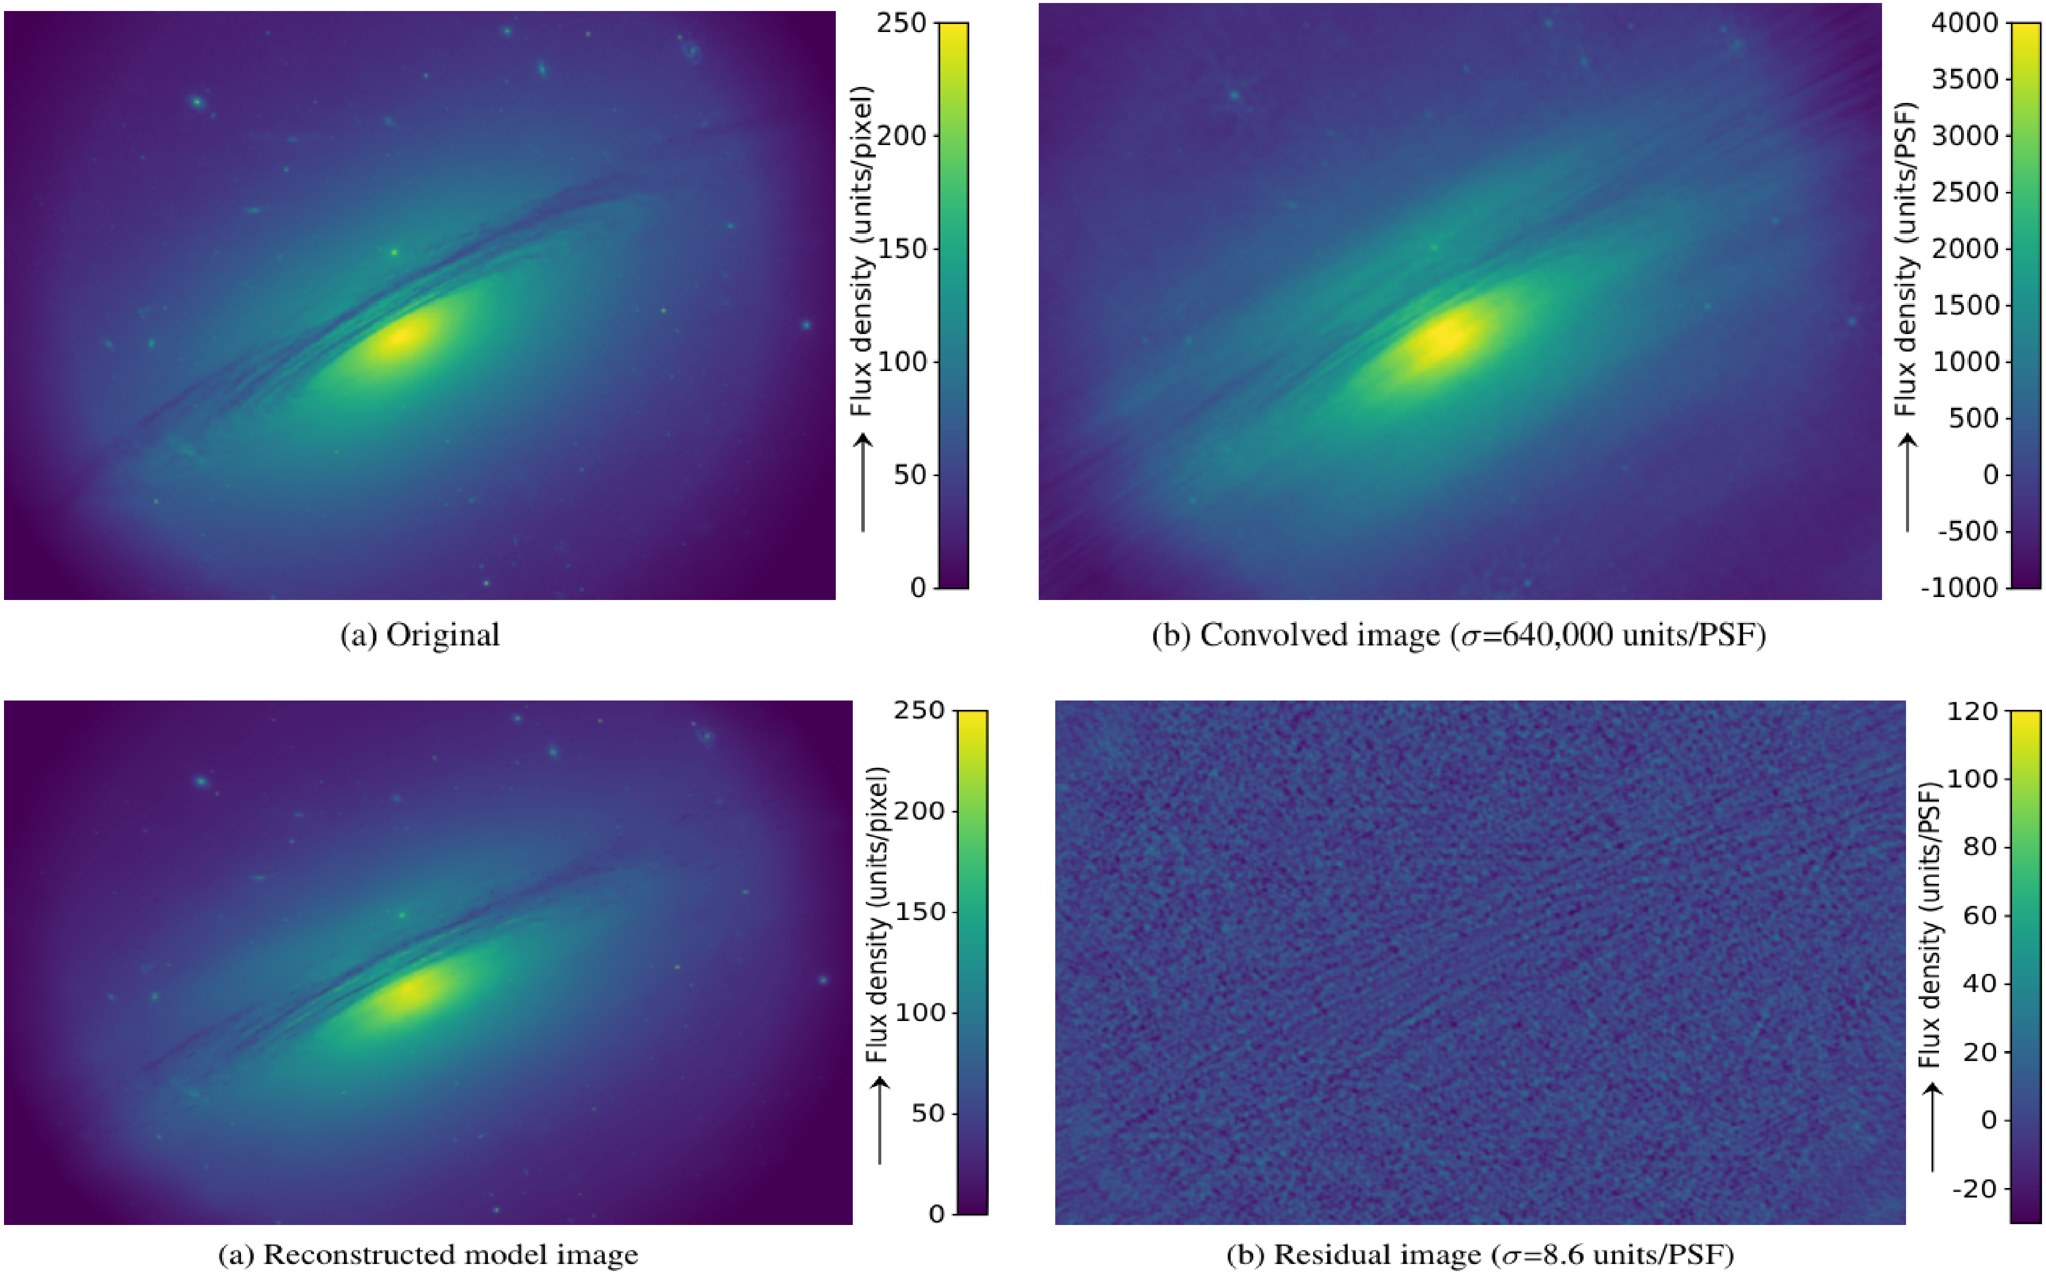

Figure 9. Automatic scale-dependent masking applied on the UGC12591 test-set.

Offringa and Smirnov (2017)

Image Doman Gridding (IDG)

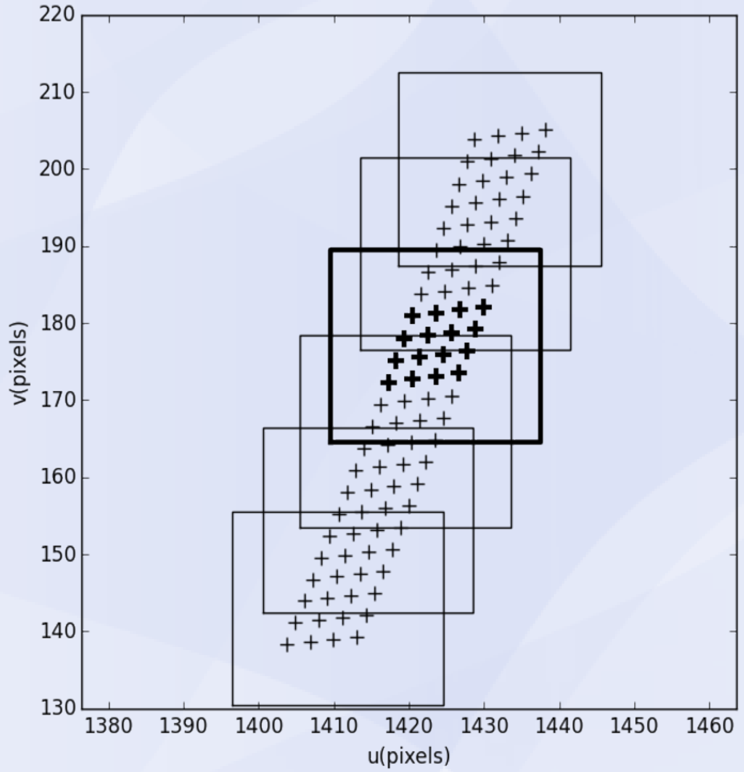

(a)

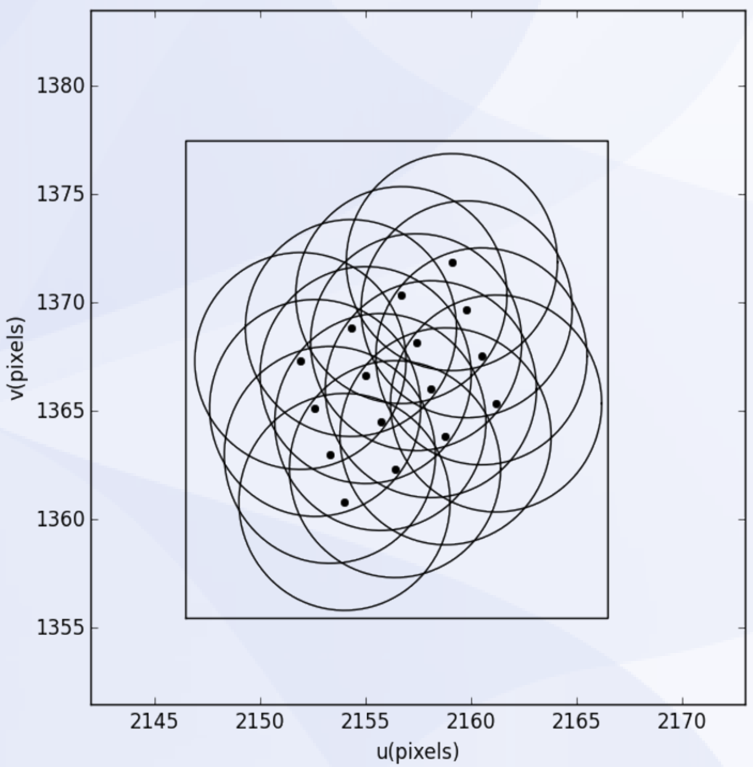

(b)

(a) uv track for a single baseline and multiple channels. The boxes indicate the position of the subgrids. The bold box correspond to the bold samples. (b) A single subgrid (box) encompassesing all affected pixels in the uv grid. The support of the convolution function is indicated by the circles around the samples.

Van der Tol, Veenboer & Offringa (to be submitted)

Image Doman Gridding (IDG)

- Compared to normal gridding, IDG does (on first order) not change the amount of operations to be performed
- However, parallelizes extremely well on GPUs
- W & A-term (beam/ionosphere) correction “for free”
- Results in very high gridding accuracy:

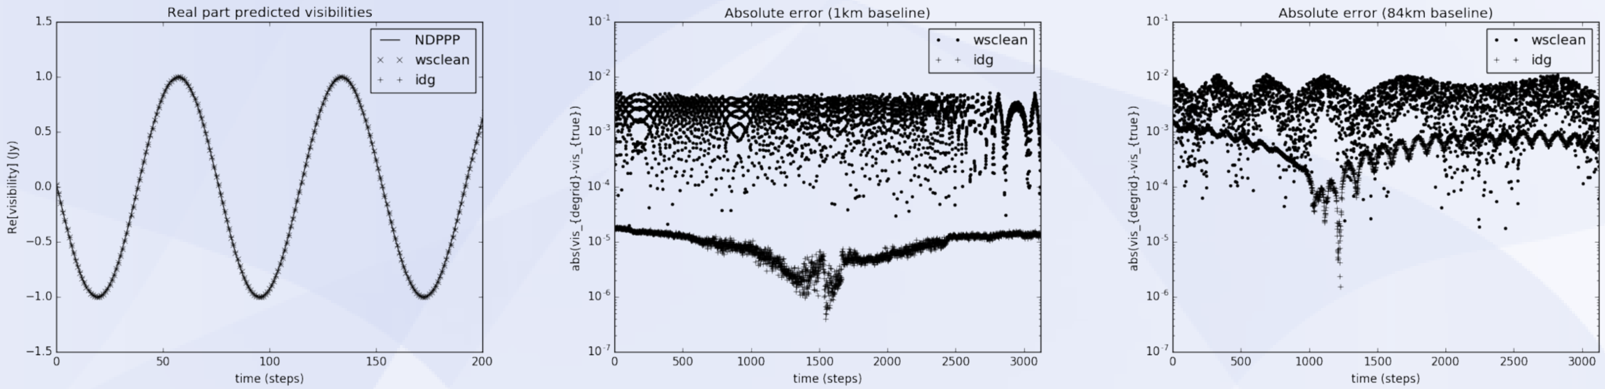

Left: visibilities for a point source as predicted by direct evaluation of the ME, and degridding by the classical gridder and image domain gridder. The visibilities are too close together to distinguish in this graph. The plot and the middle and on the right show the absolute value of the difference between direct evaluation and degridding for a short (1km) and a long (84km) baseline. On the short baseline the image domain gridder rms error of 1.03×10^{-5} Jy is about 242 times lower than the classical gridder rms error of 2.51×10^{-3} Jy. On the long baseline the image domain gridder rms error of 7.10×10^{-4} Jy is about 7 times lower than the classical gridder error of 4.78×10^{-3} Jy.

Van der Tol, Veenboer & Offringa (to be submitted)

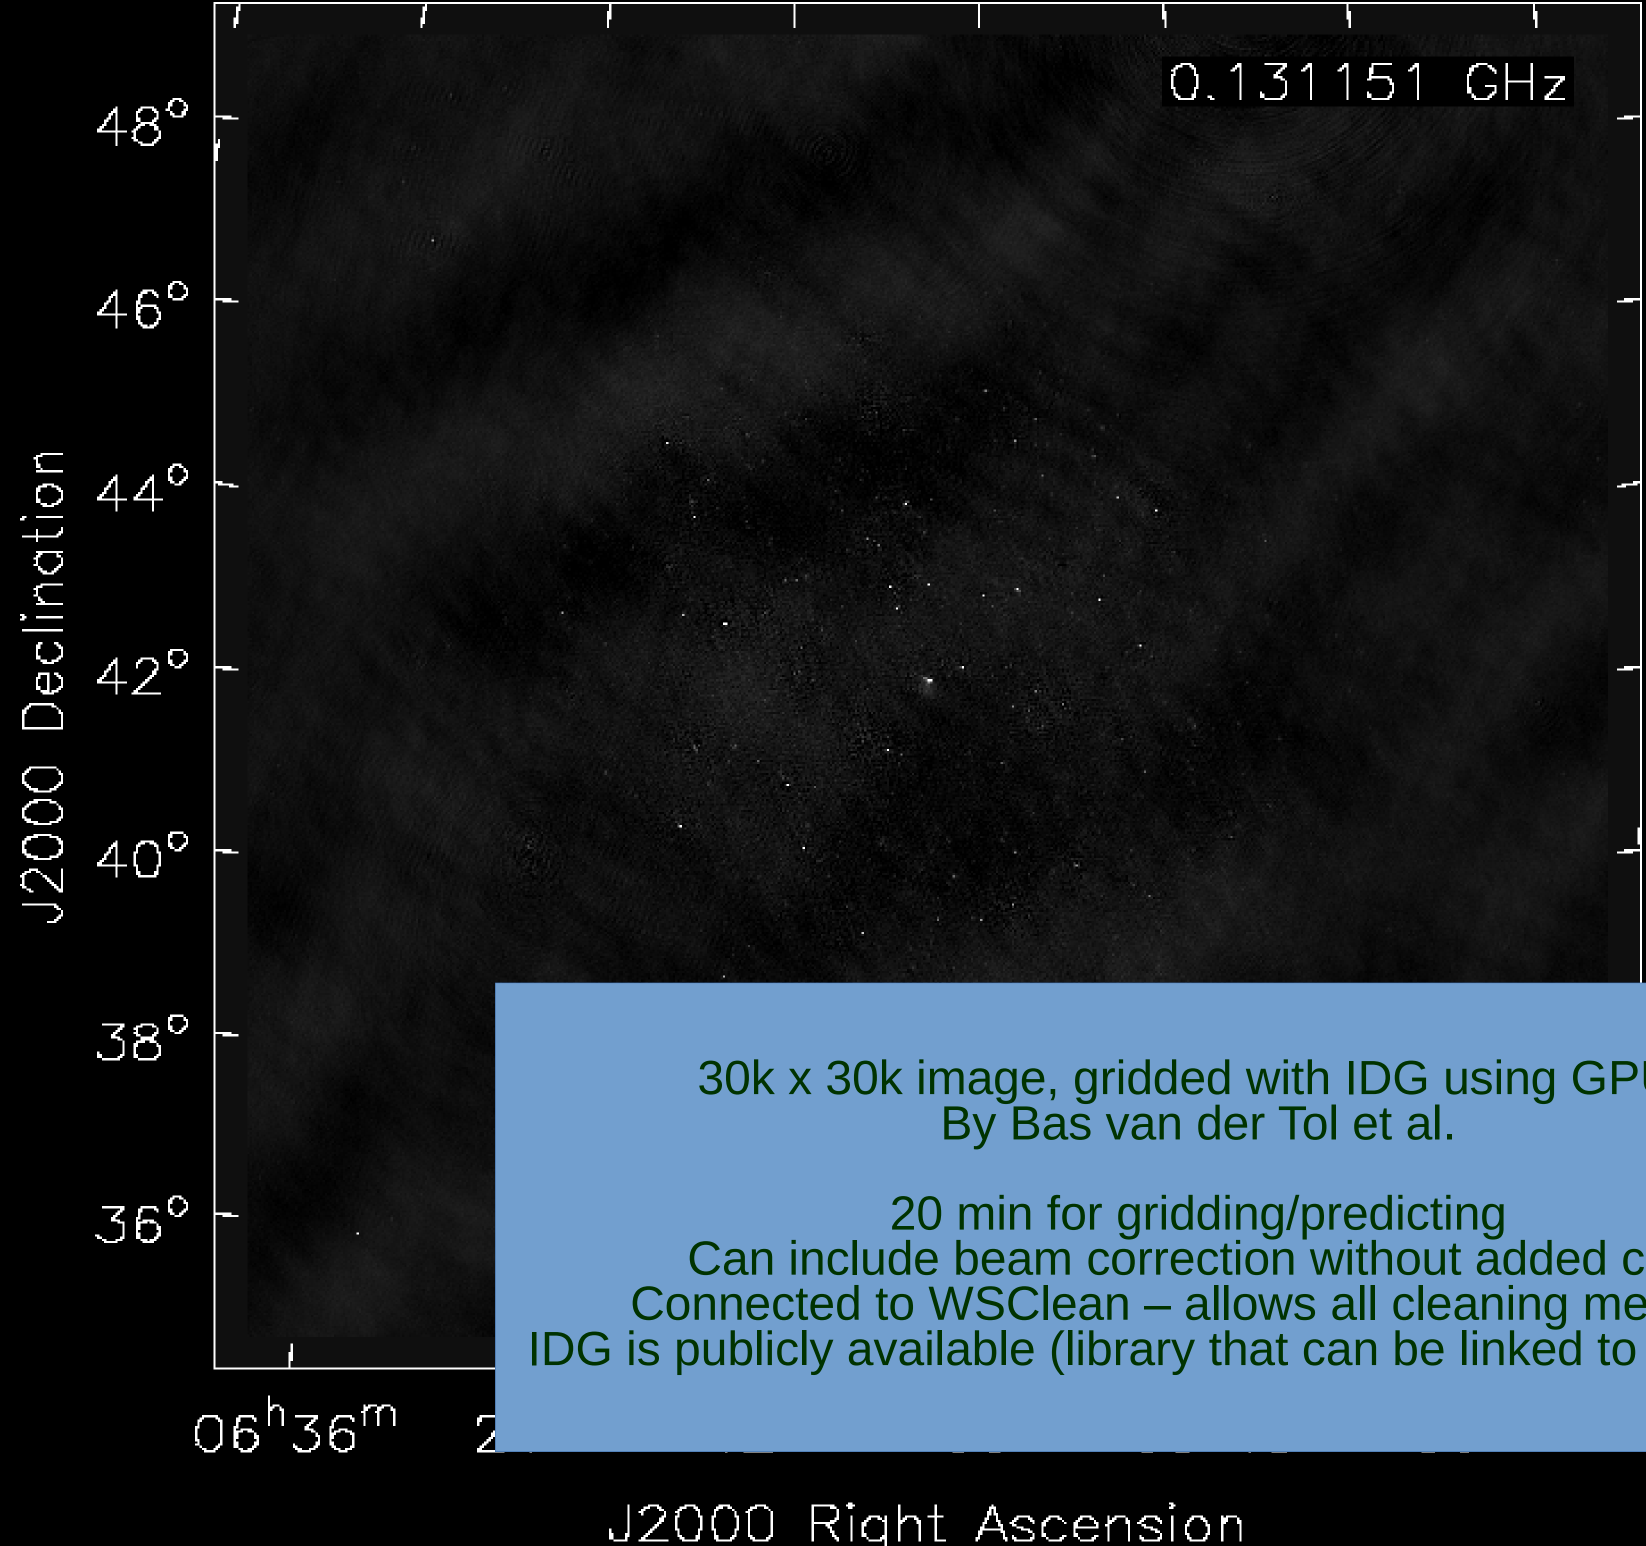

30k x 30k image, gridded with IDG using GPUs
By Bas van der Tol et al.

20 min for gridding/predicting
Can include beam correction without added cost
Connected to WSClean – allows all cleaning methods
IDG is publicly available (library that can be linked to WSClean)

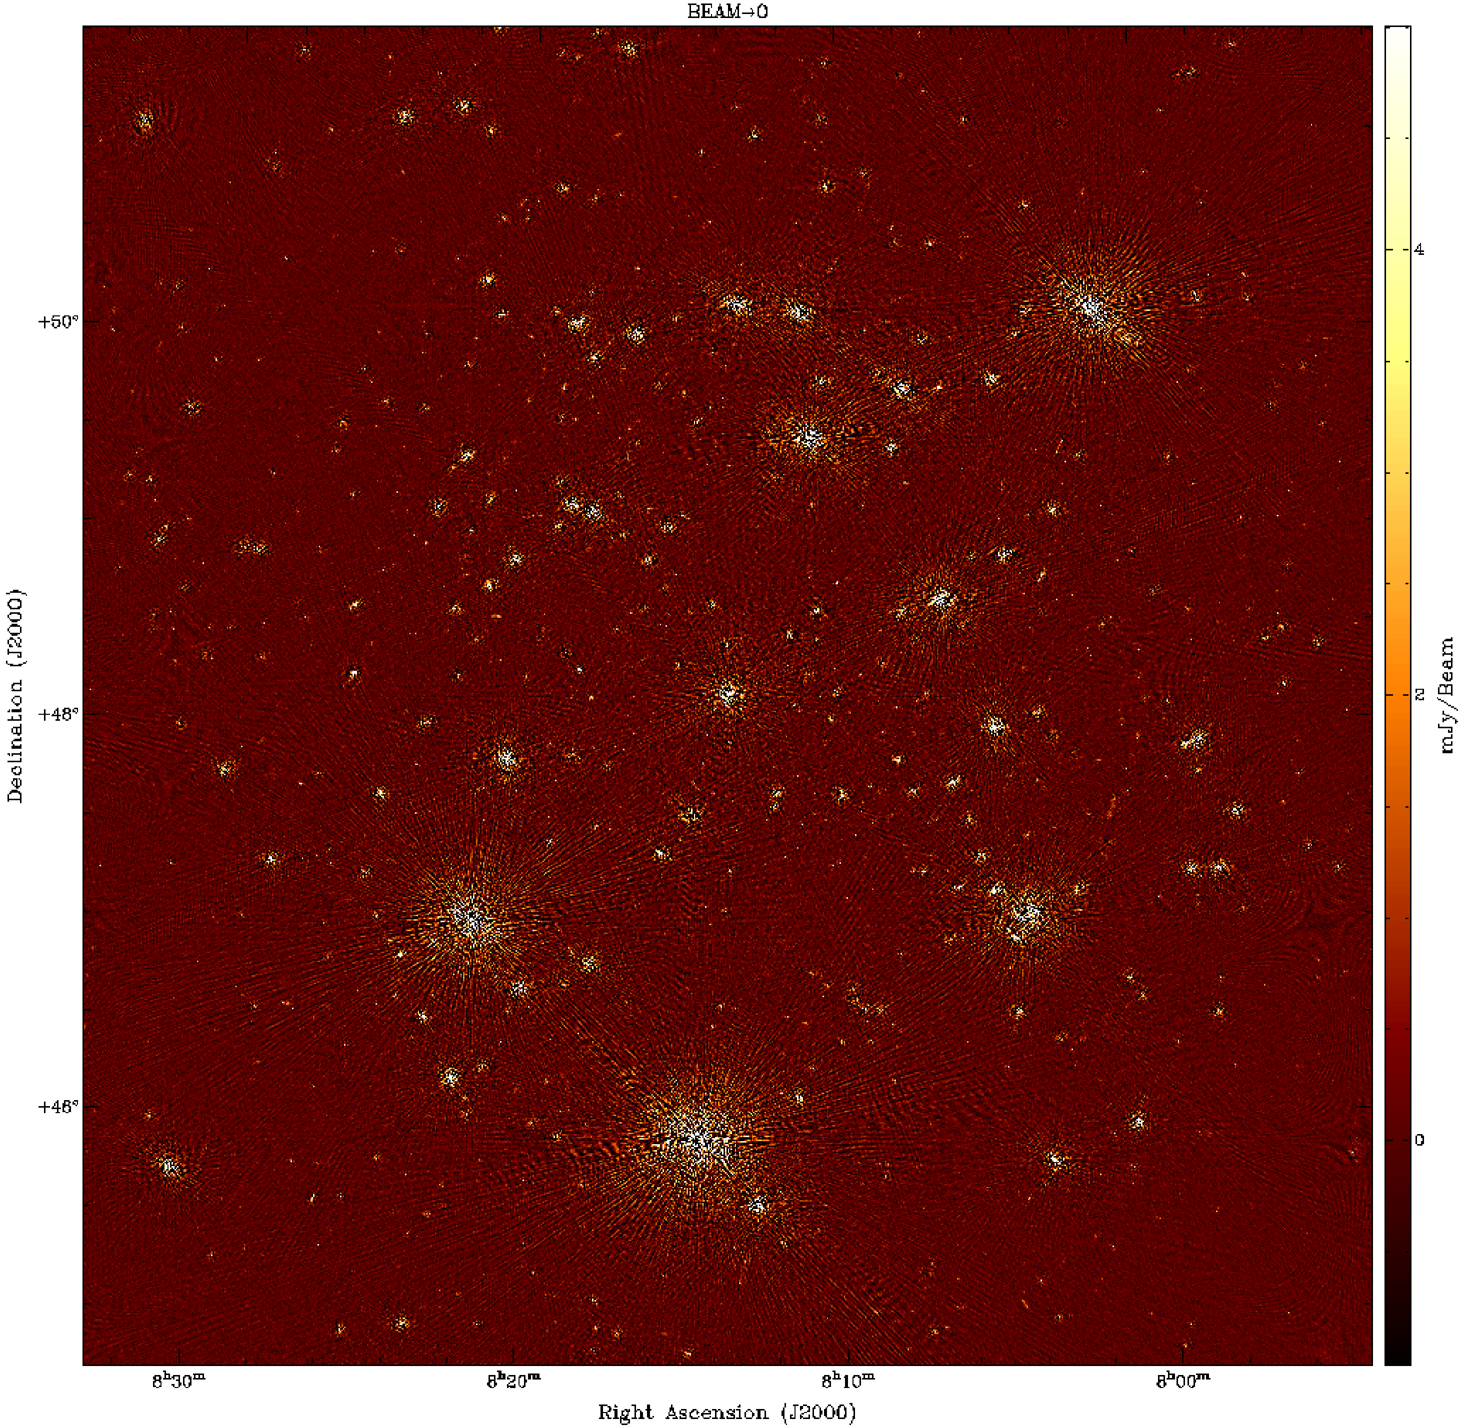

14k x 14k image

(7° x 7°, about up to first null)

LOFAR, 20 MHz 6 h

Gridding with IDG on GPU

250 µJy noise

IDG 0.2 + WSClean 2.5

(Both are publicly available)

Fully multi-scale multi-frequency cleaned IDG 14k x 14k result

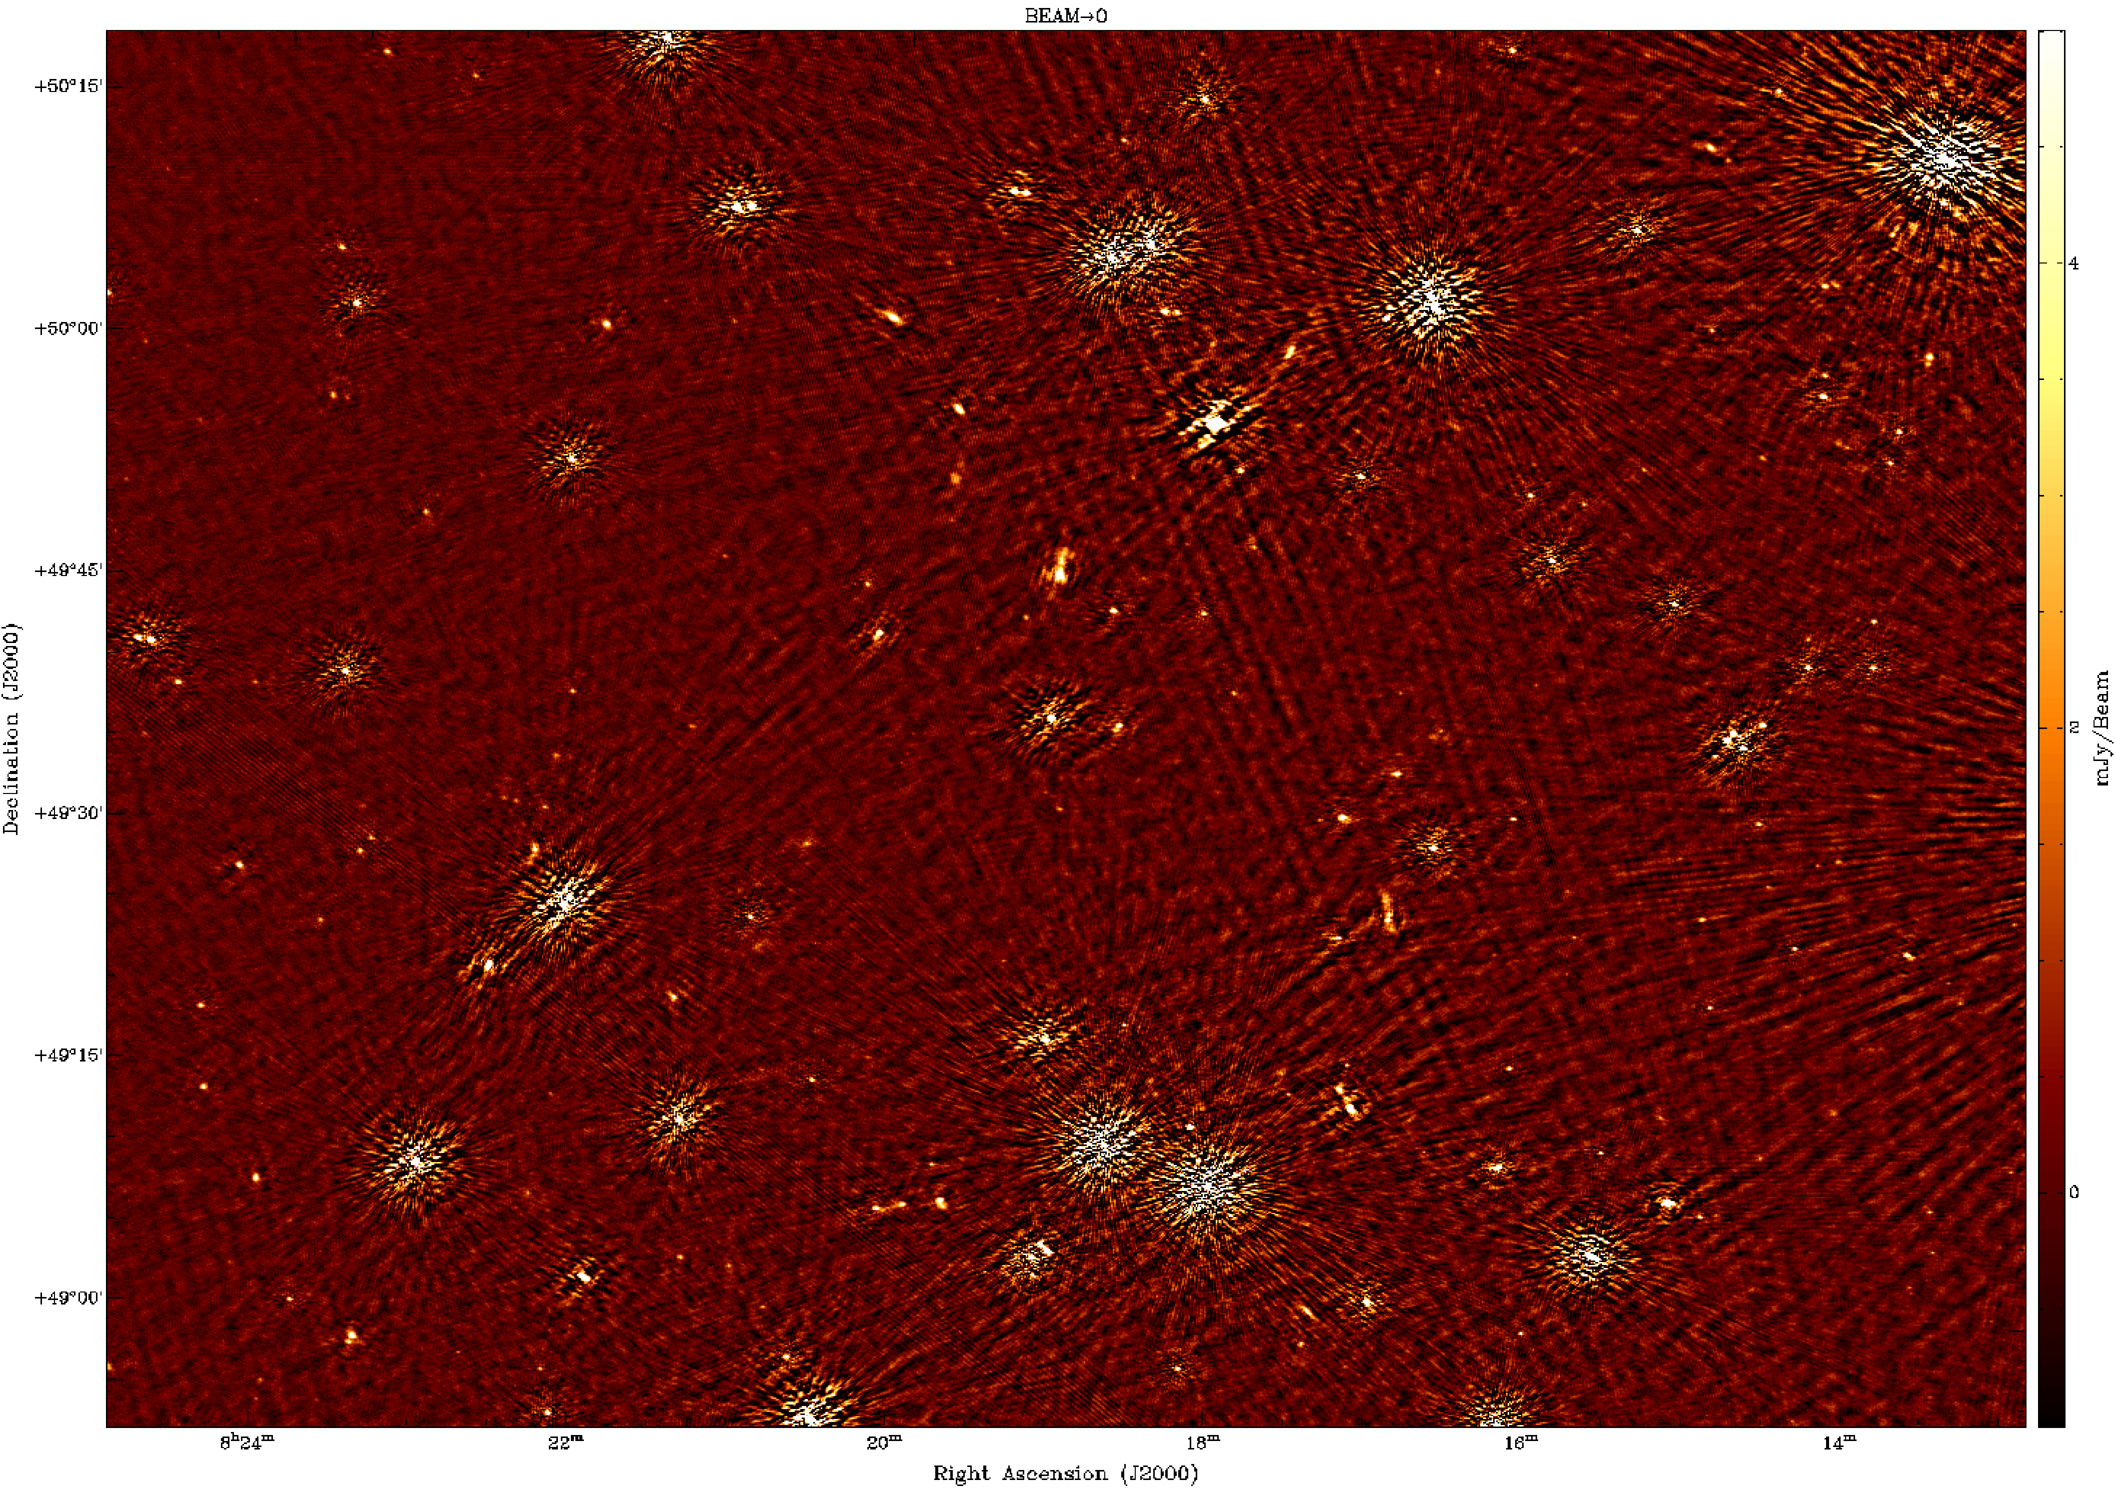

Zoom in to 2° x 1.5°

Implementation of IDG

1) Connect IDG to WSClean
2) Apply beam corrections during gridding
3) Apply DD ionospheric corrections

Implementation of IDG

1) Connect IDG to WSClean

2) Apply beam corrections during gridding

3) Apply DD ionospheric corrections

Connection to WSClean finished: all cleaning modes are supported.

(work by Van der Tol, Offringa, Veenboer, Dijkema and others)

Applying a-term with IDG

Next step: apply LOFAR beam

- Working & to be released in next WSClean version
- Applies full-Jones antenna beam in forward and backward imaging step
- No extra computational cost(!)

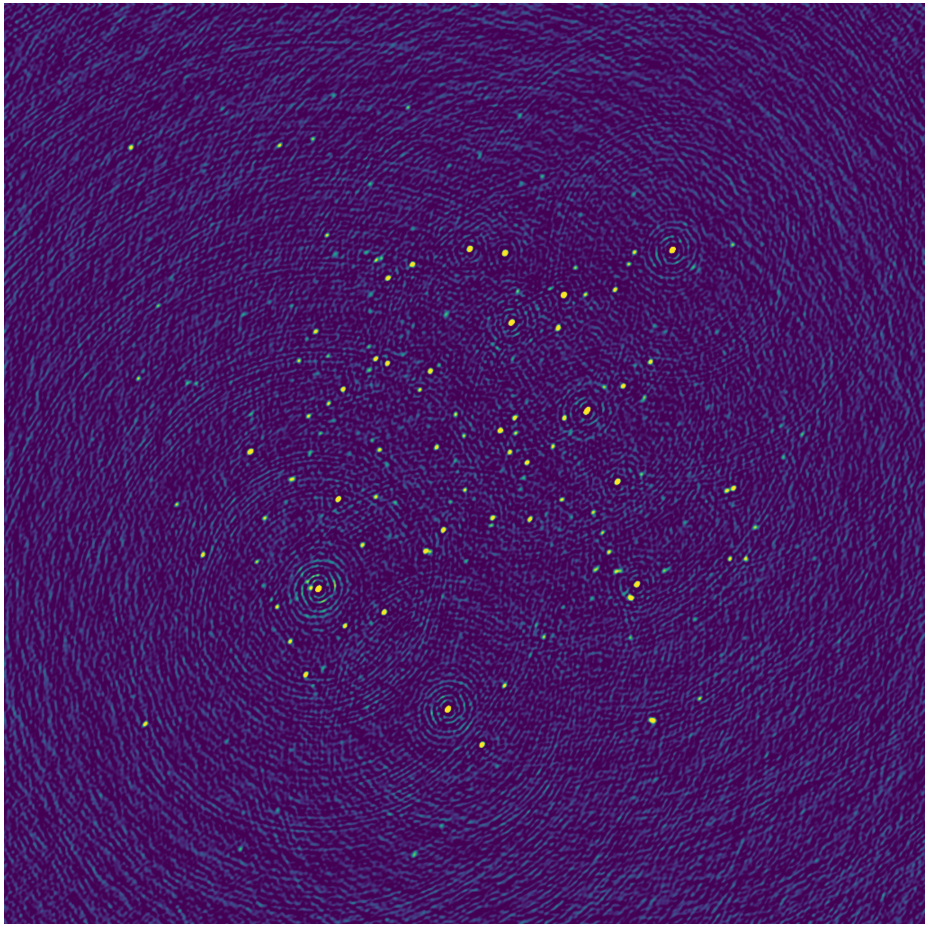

Normal imaging with w-stacking gridder (no beam)

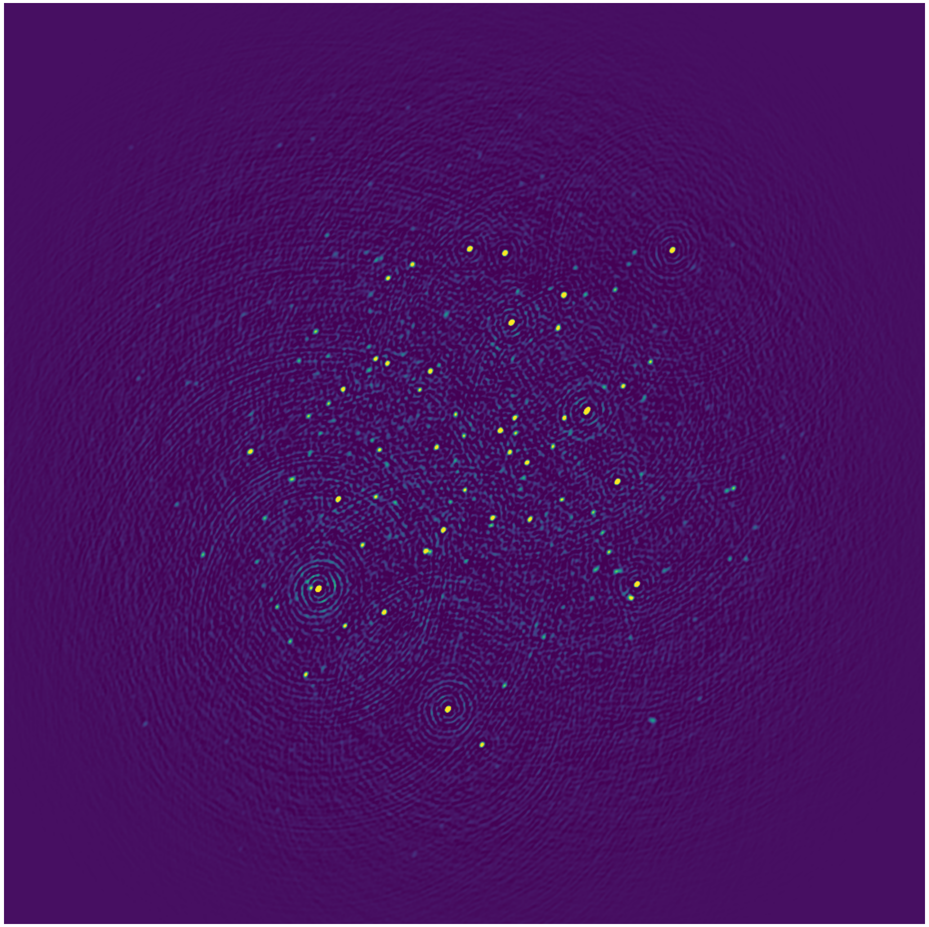

LOFAR beam applied during imaging stage Producing “optimally weighted” image

Applying a-term with IDG

Next step: apply LOFAR beam

Working & to be released in next WSClean version

…ward imaging step

Snapshot of the LOFAR beam for the 48 stations:

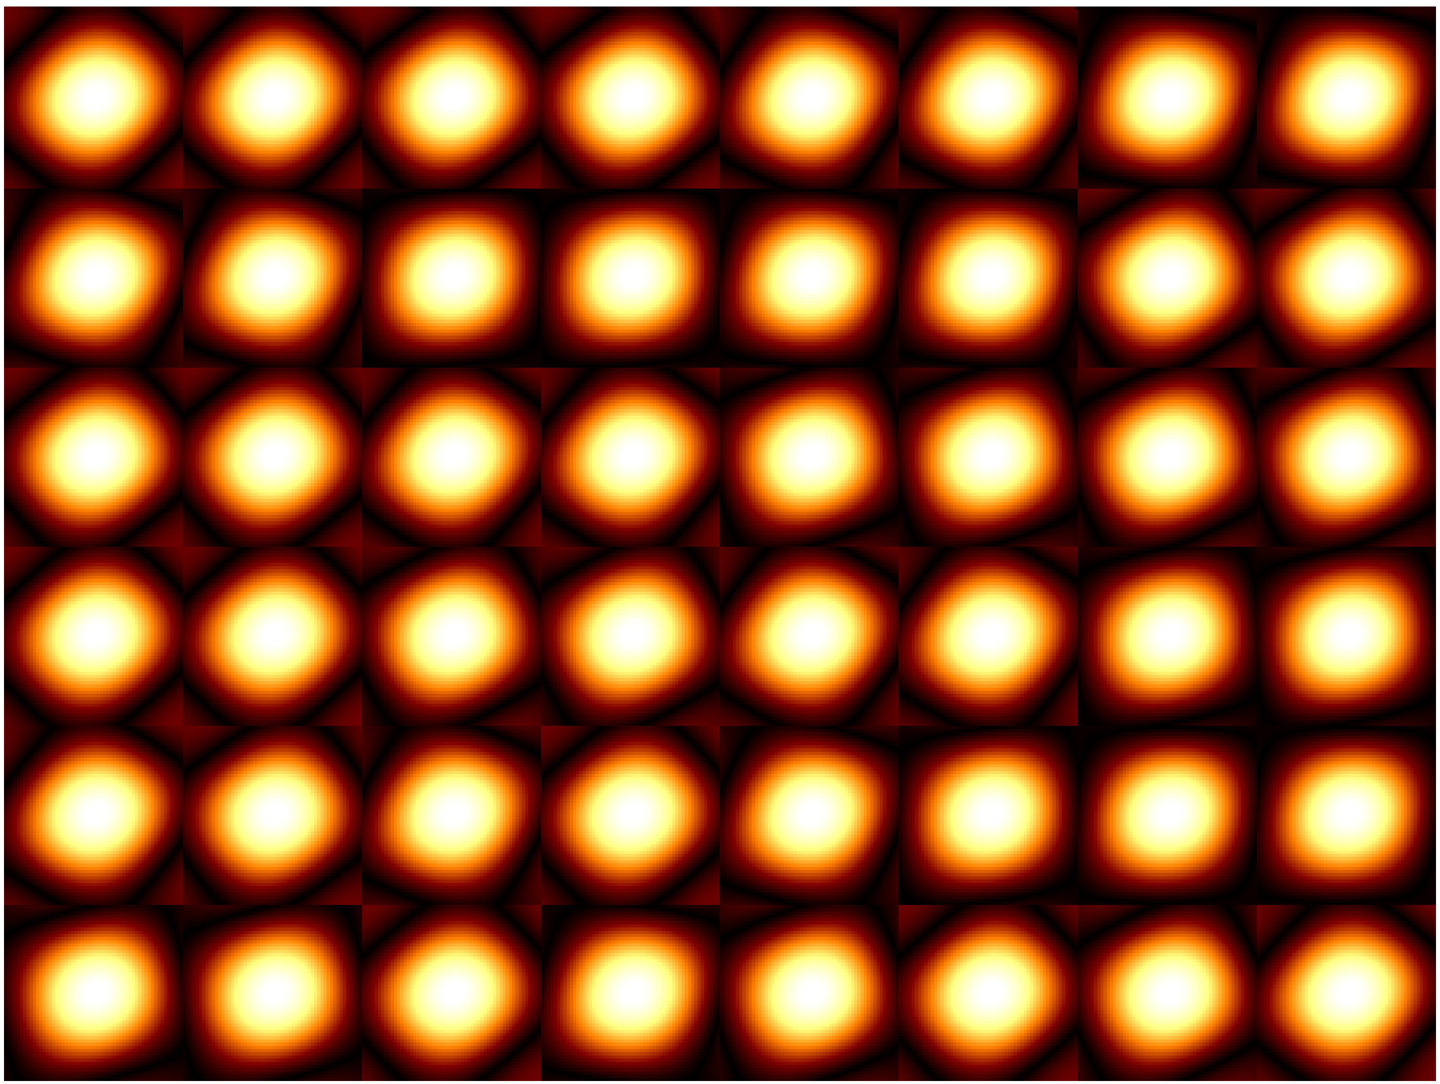

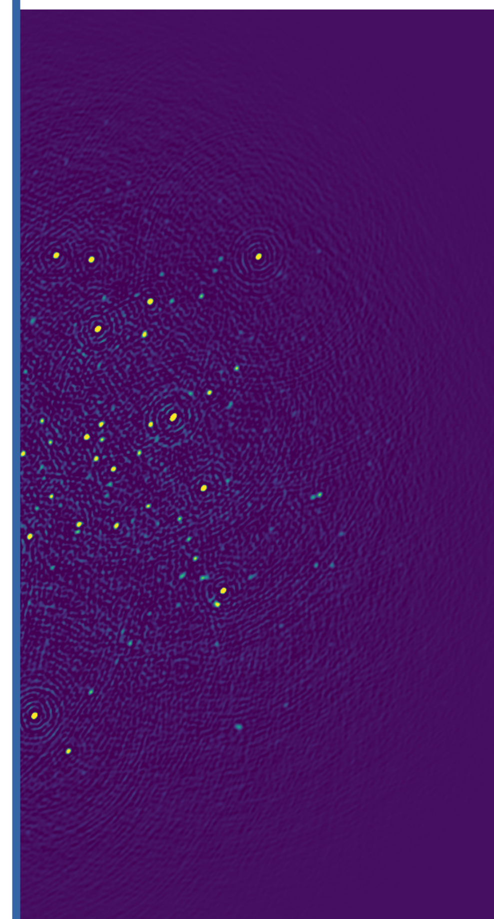

(no beam)

…ed during imaging stage

…ally weighted" image

Applying a-term with IDG

Next step: apply LOFAR beam

Working & to be released in next WSClean version

Snapshot of the LOFAR beam for the 48 stations:

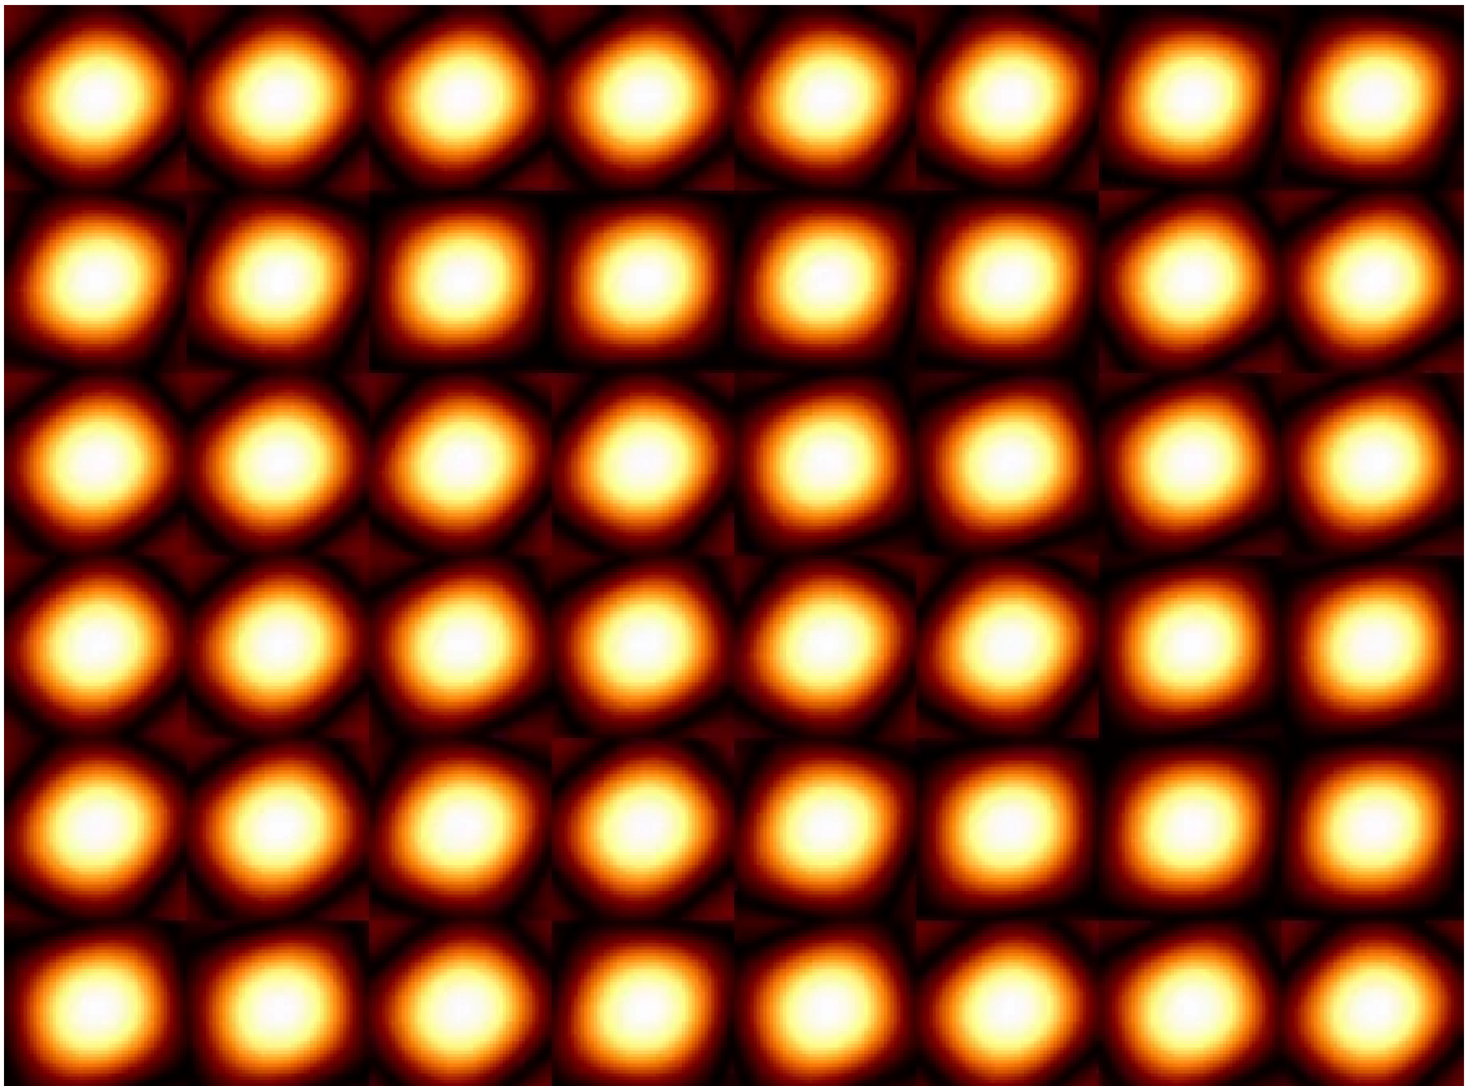

...ward imaging step

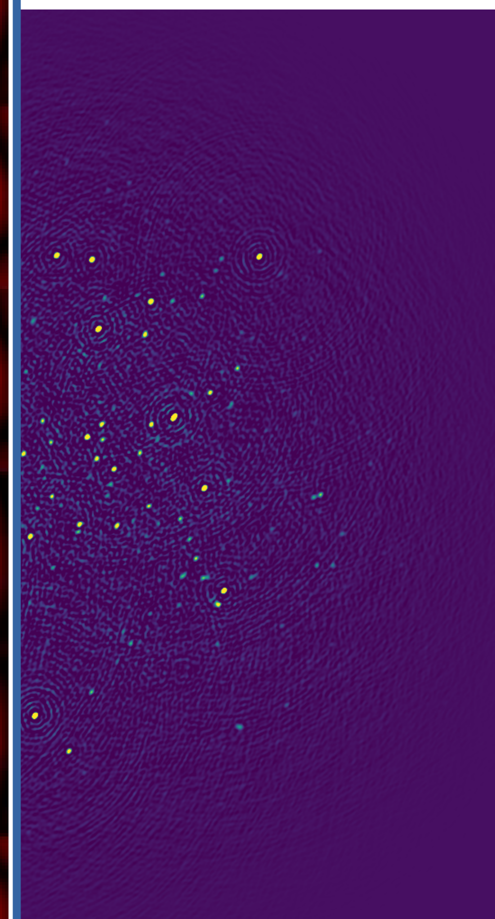

(no beam)

...ed during imaging stage

Producing "optimally weighted" image

Implementation of IDG

1) Connect IDG to WSClean

2) Apply beam corrections during gridding

3) Apply DD ionospheric corrections

Parallel cleaning

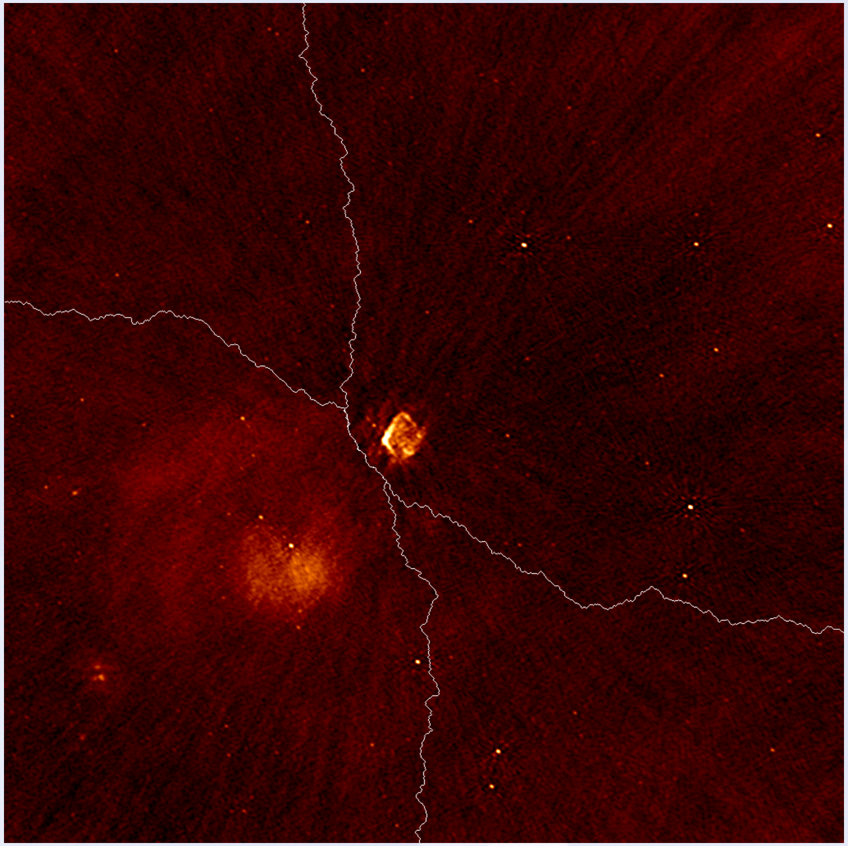

Division of a (small) dirty images into 4 independent areas. Bounding boxes are send to deconvolution algorithm, edges are applied with a mask.

- IDG makes it computationally possible to make 30k x 30k images
- Computational bottleneck has (again) been moved to deconvolution
- However, big images can easily be subdivided and cleaned independently
- Implemented in WSClean by using Dijkstra's algorithm (with constraints)

Parallel cleaning

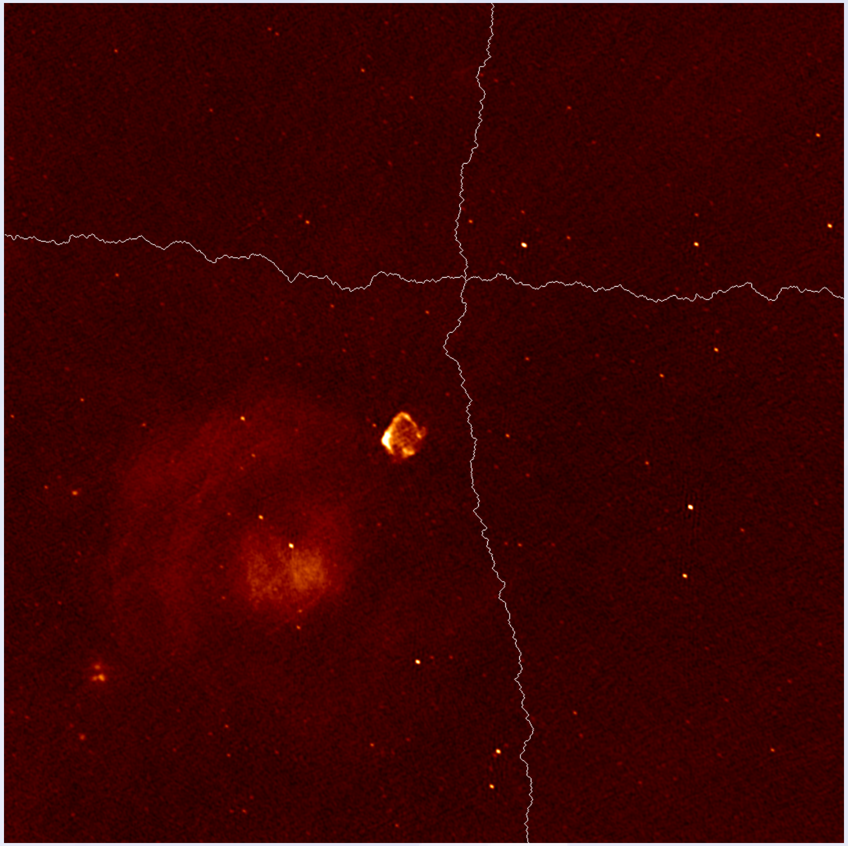

Division is recalculated each major iteration.
This shows the division during the final major iteration.

- IDG makes it computationally possible to make 30k x 30k images
- Computational bottleneck has (again) been moved to deconvolution
- However, big images can easily be subdivided and cleaned independently
- Implemented in WSClean by using Dijkstra's algorithm (with constraints)

Direction-dependent calibration

Options for direction-dependent calibration

- Several packages can perform DDE calibration:
  - MWA's real-time system (Mitchel et al., 2008)
  - SageCal (Yatawatta et al., 2009)
  - Ionpeel (Offringa et al., 2016)
  - Killms (Smirnov & Tasse, 2015)
  - Factor (Van Weeren et al., 2016)
  - SPAM (Intema, 2014)
  - OBIT (Cotton, 2008)
  - [..?]
  - → DPPP… (Offringa et al. in prep)

Issues with current DDE pipelines

The large degree of freedom in DDE calibration causes several issues:

- LOFAR LBA (30-80 MHz) calibration
  - S/N ratio very low
  - No current pipeline can produce (good) DD solutions
- Diffuse low-frequency imaging
  - Current pipelines calibrate diffuse structures out
- EoR imaging
  - Need to avoid suppression of EoR signals
  - Frequency stability important
- Deep HBA imaging
  - Requires faster solution interval
  - Interpolated TEC screens

Issues with current DDE pipelines

The large degree of freedom in DDE calibration causes several issues:

- LOFAR LBA (30-80 MHz) calibration
  - S/N ratio very low
  - No ... D solutions
- Diffuse low-
  - Cu ... es out
- EoR imagi
  - Ne ... s
  - Frequency stability important
- Deep HBA imaging
  - Requires faster solution interval
  - Interpolated TEC screens

Solution:

Introduce more constraints into the calibration

Constrained DDE calibration approach

- DPPP is the “Default Pre-processing Pipeline” – written for LOFAR (but works also for other arrays)
  - Largely written by G. van Diepen and T. J. Dijkema
- Good starting point: DPPP already has a fast prediction implementation

1) A DDE algorithm was implemented in the DPPP software
   - Base of algorithm is a multi-directional version of alternating least squares (see Smirnov & Tasse 2015)

2) Constrain solutions:
   - Added generic constraint ‘hook’ into DDE algorithm
   - Implemented TEC frequency slope, spatial smoothness on sky and/or on ground, etc.

Are constraints inside calibration necessary?

Or is fitting the constraint after calibration also an option?

We found that in the low S/N regime, calibration does not properly converge without extra constraints inside the calibration:

Calibrated with TEC constrained Inside the solver

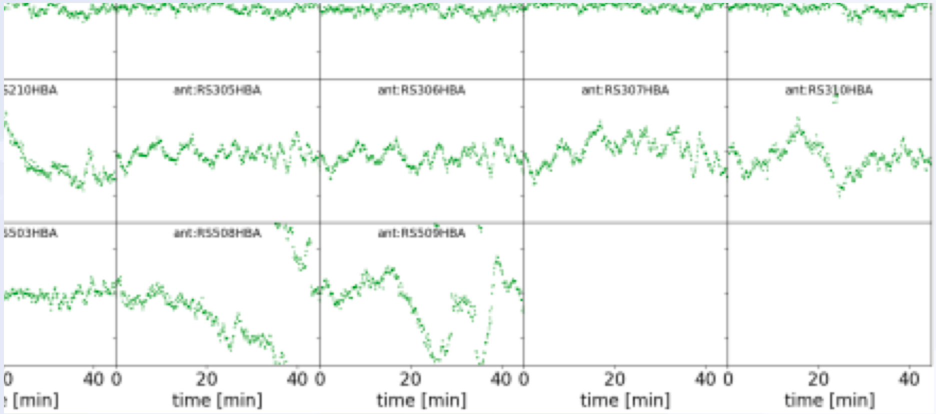

Calibrated per channel, TEC fitted afterwards

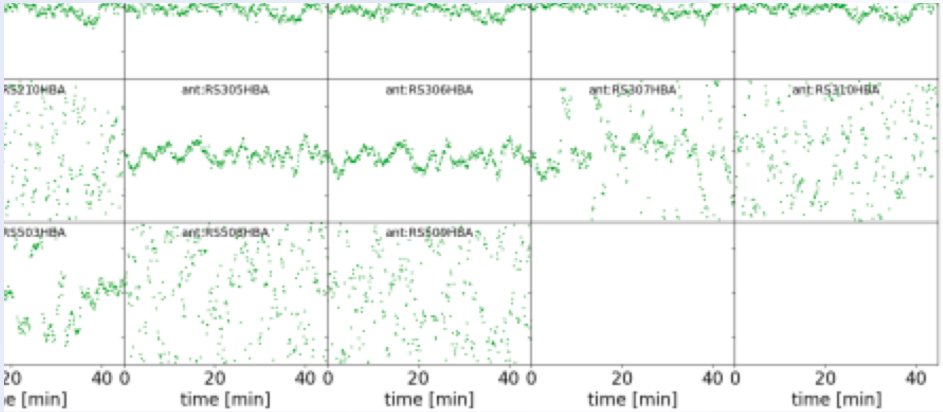

Plots by R. van Weeren

Applying TEC constraint

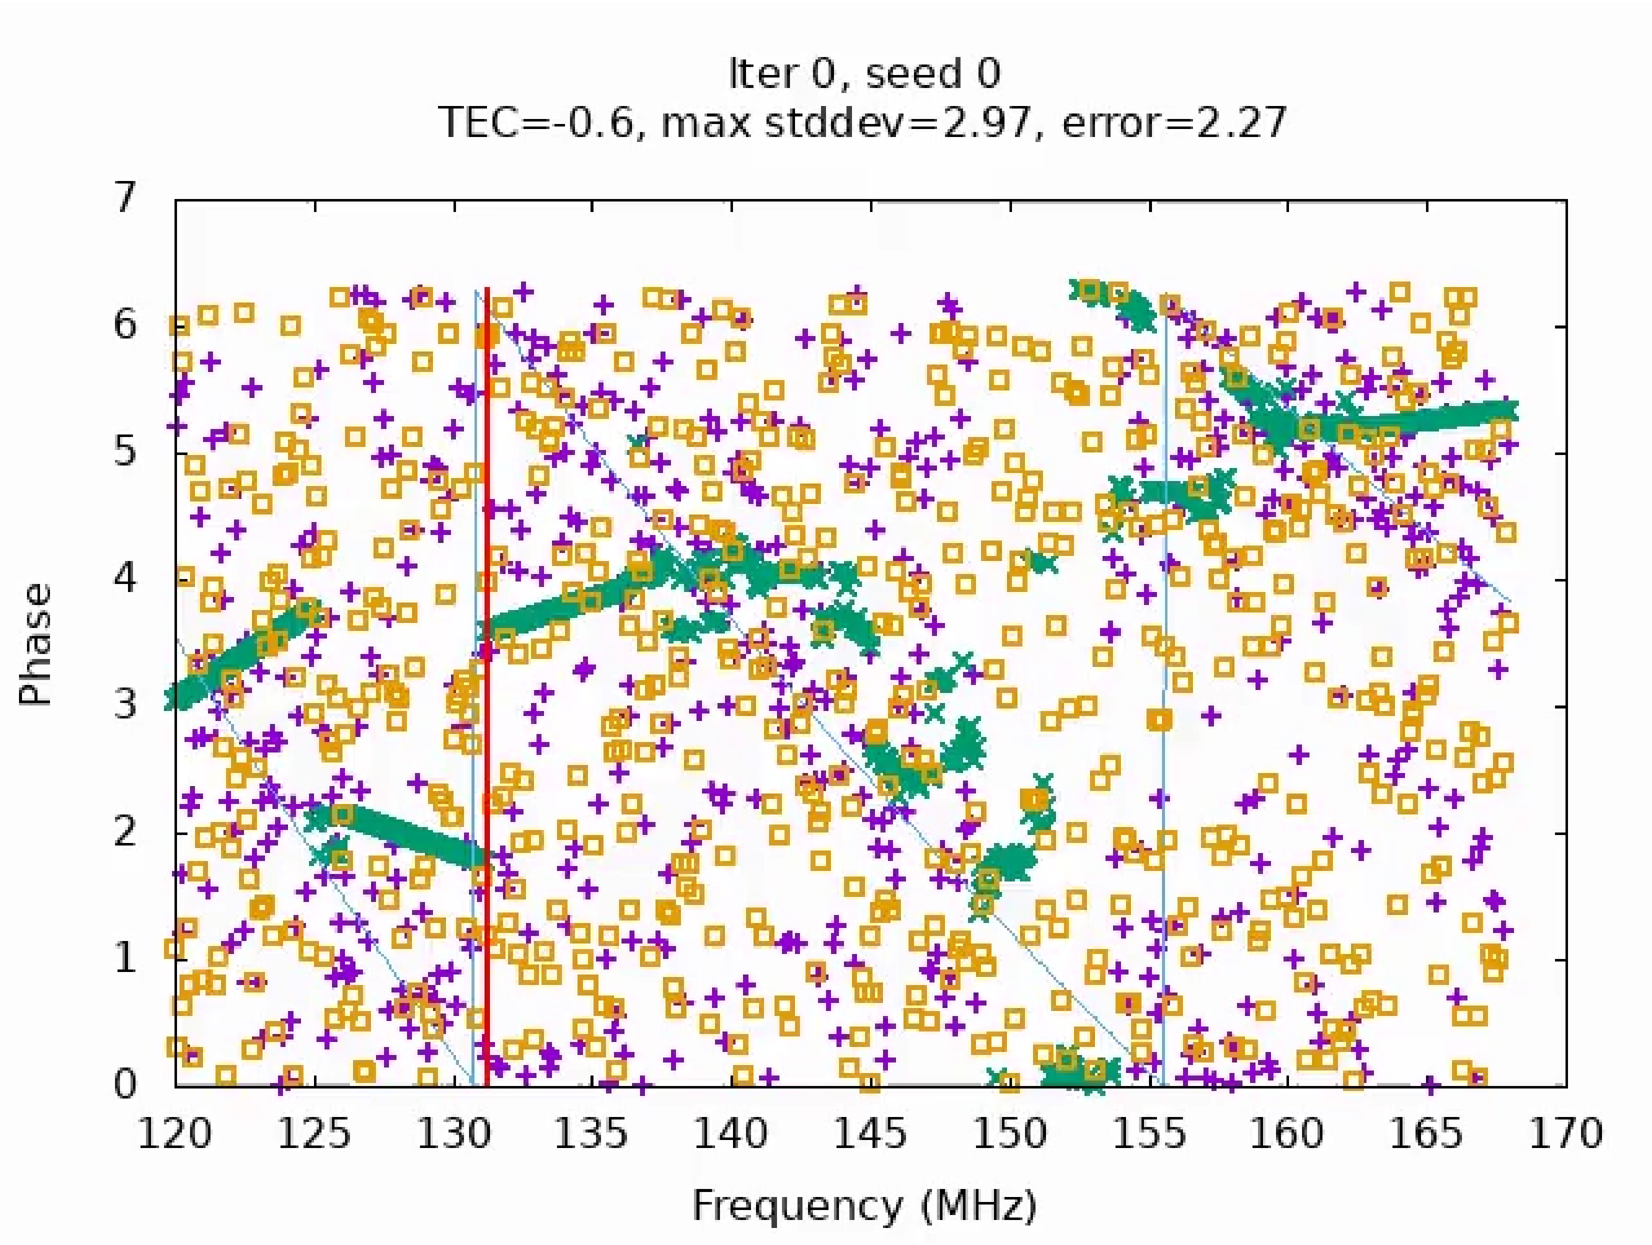

Applying TEC constraint

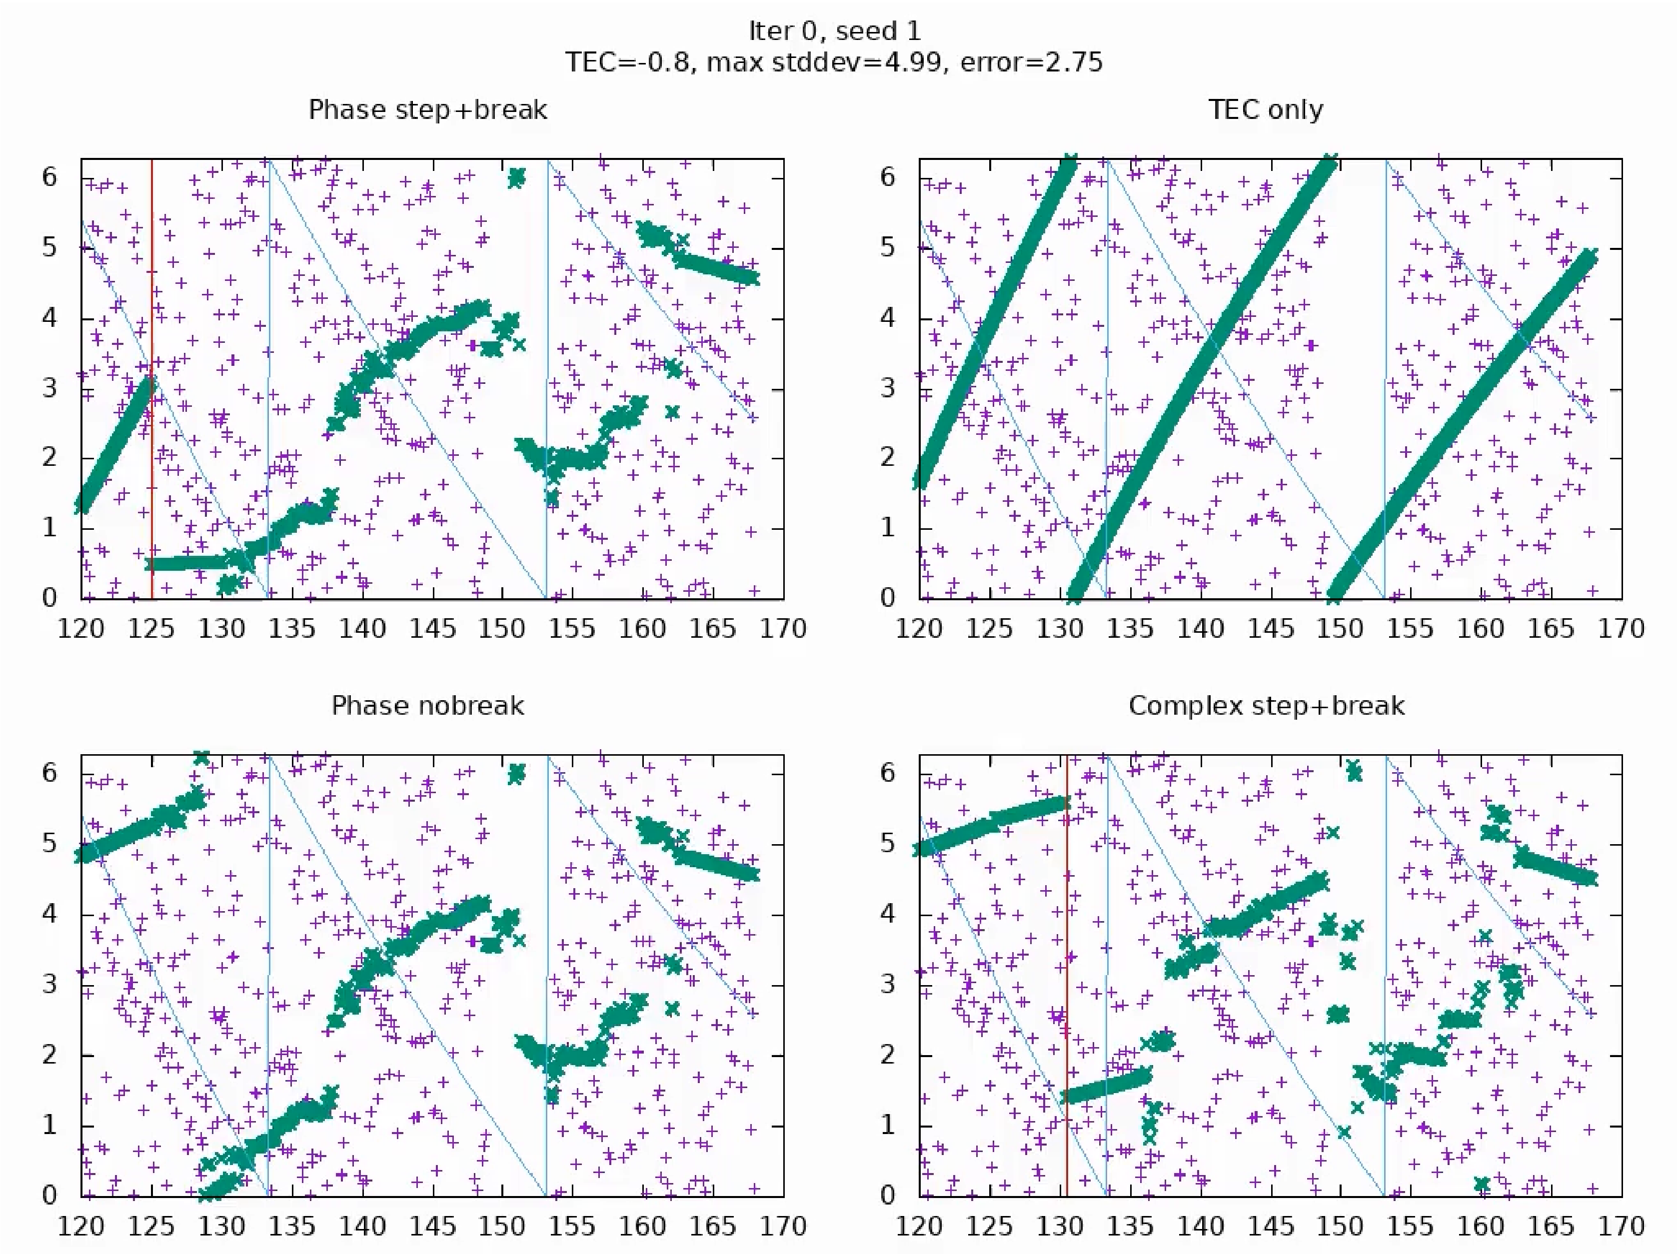

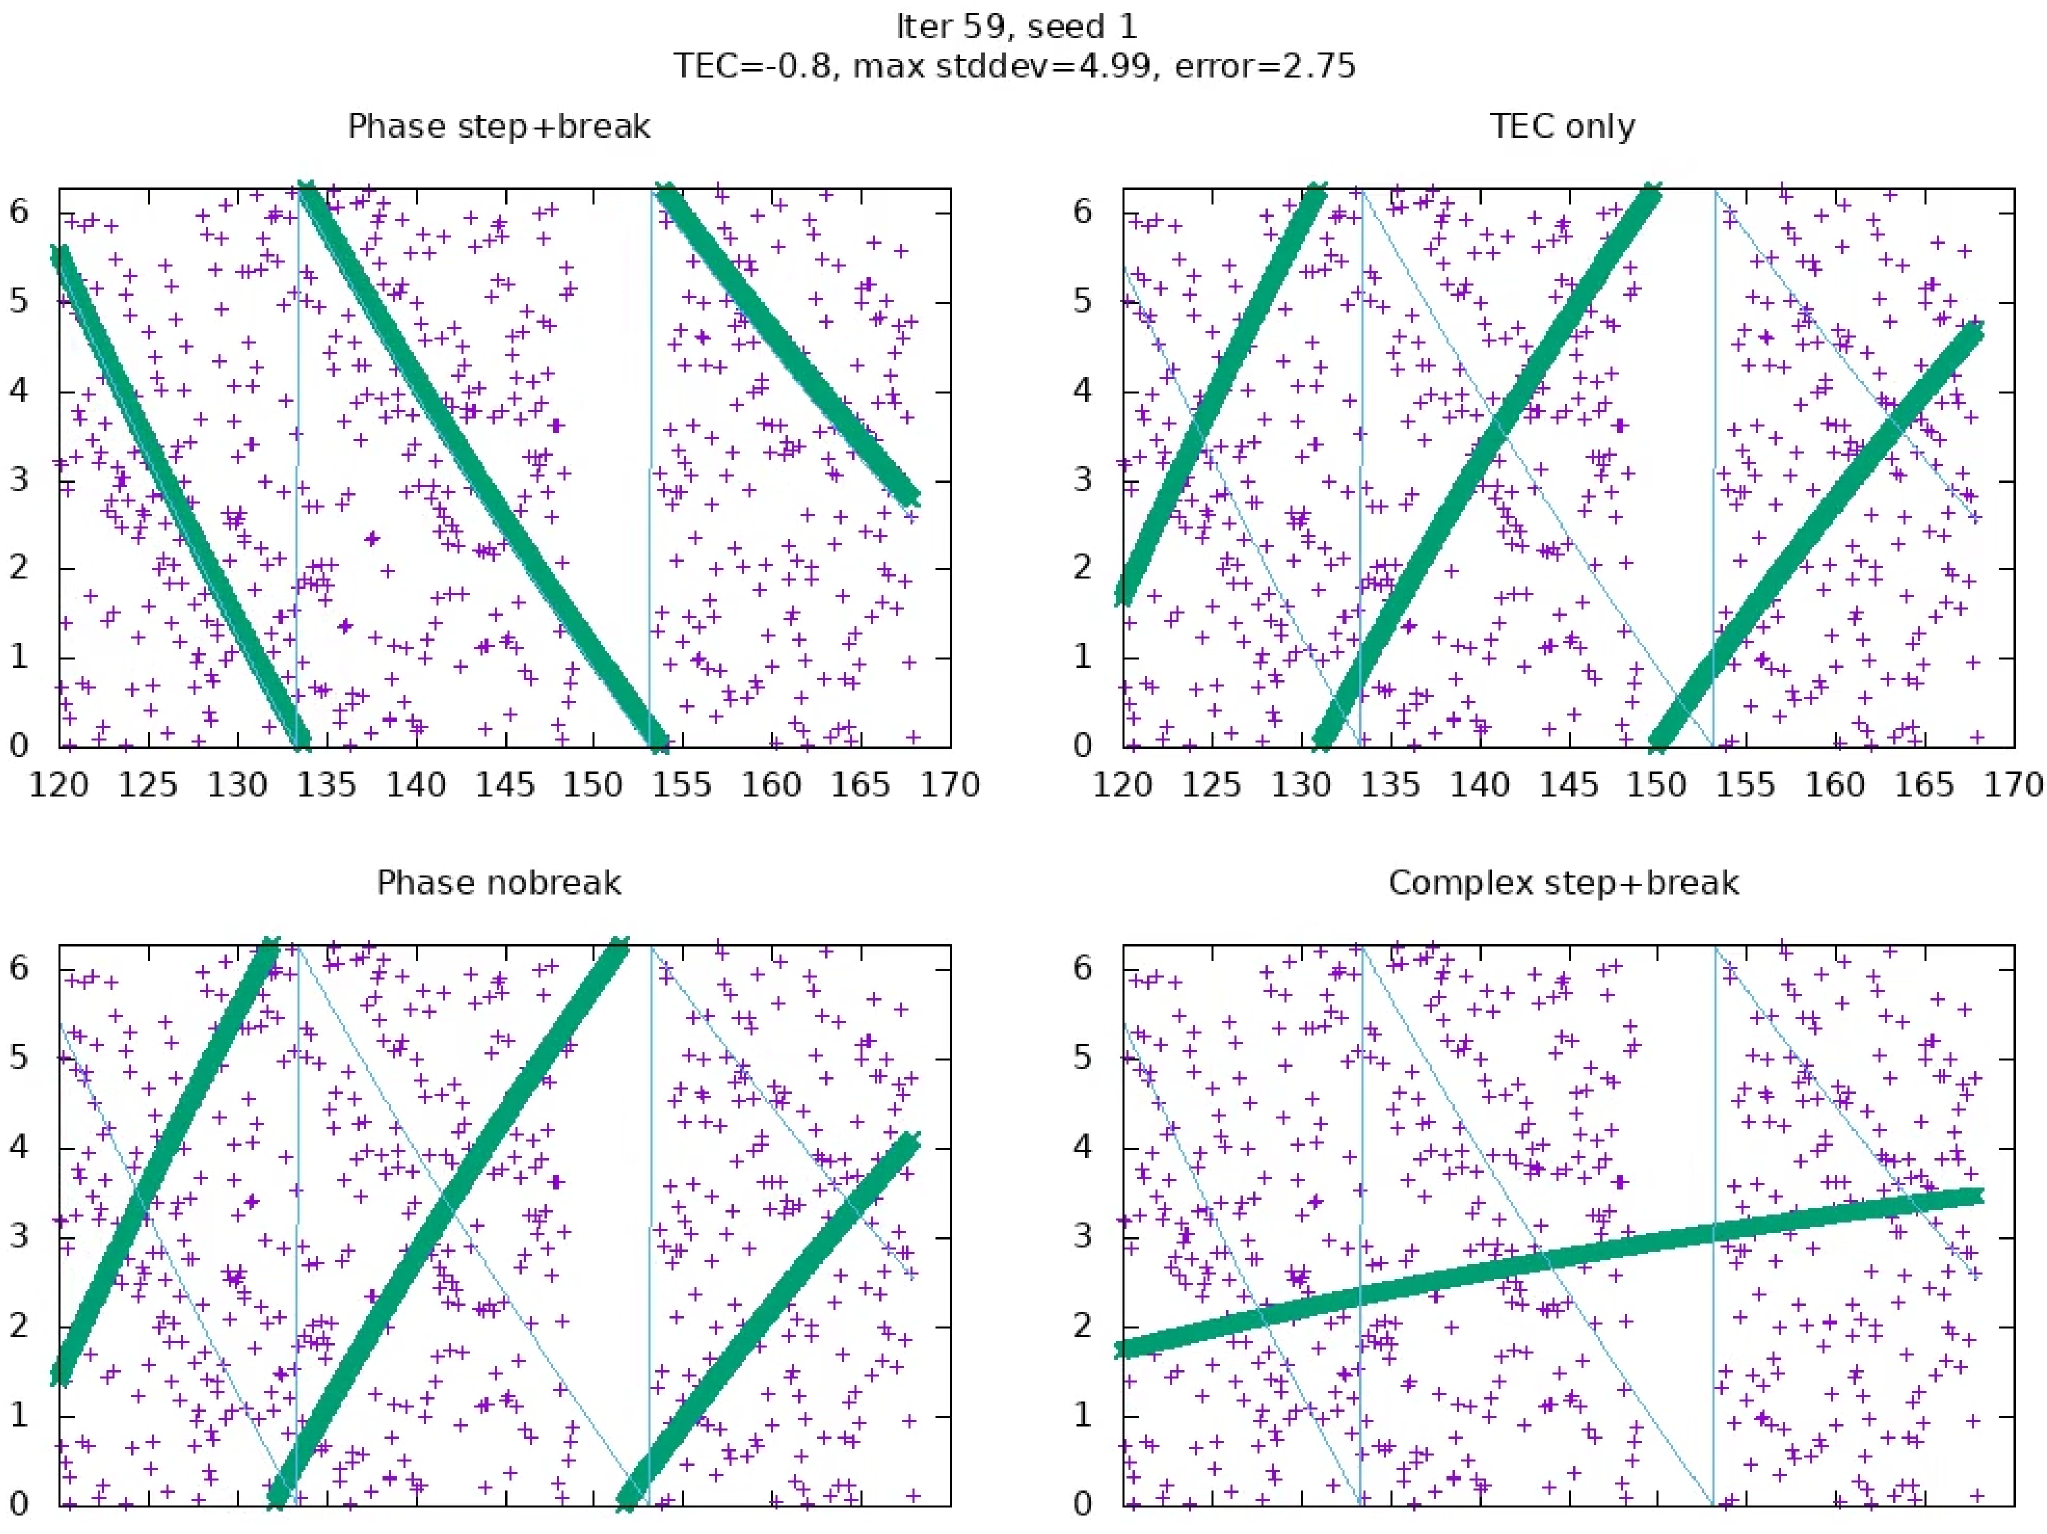
Iter 59, seed 1
TEC=-0.8, max stddev=4.99, error=2.75
Phase step+break
TEC only
Phase nobreak
Complex step+break
0
1
2
3
4
5
6
120
125
130
135
140
145
150
155
160
165
170

TEC solving result

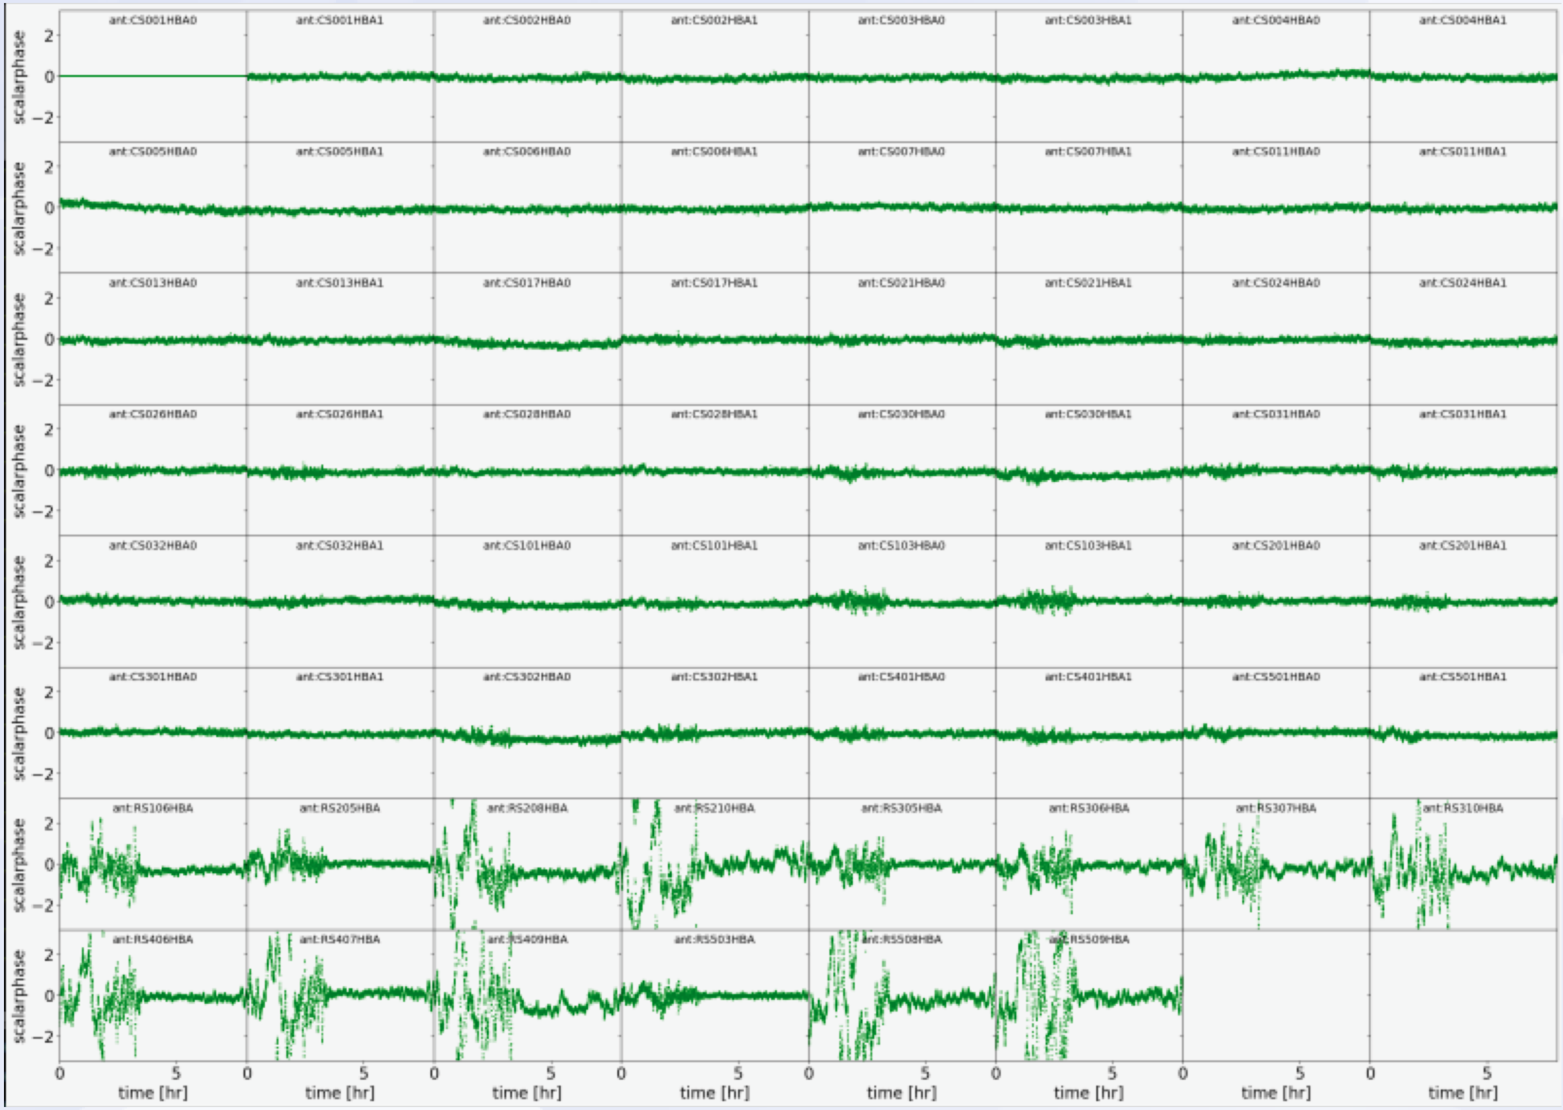

Plots by R. van Weeren

TEC solving result

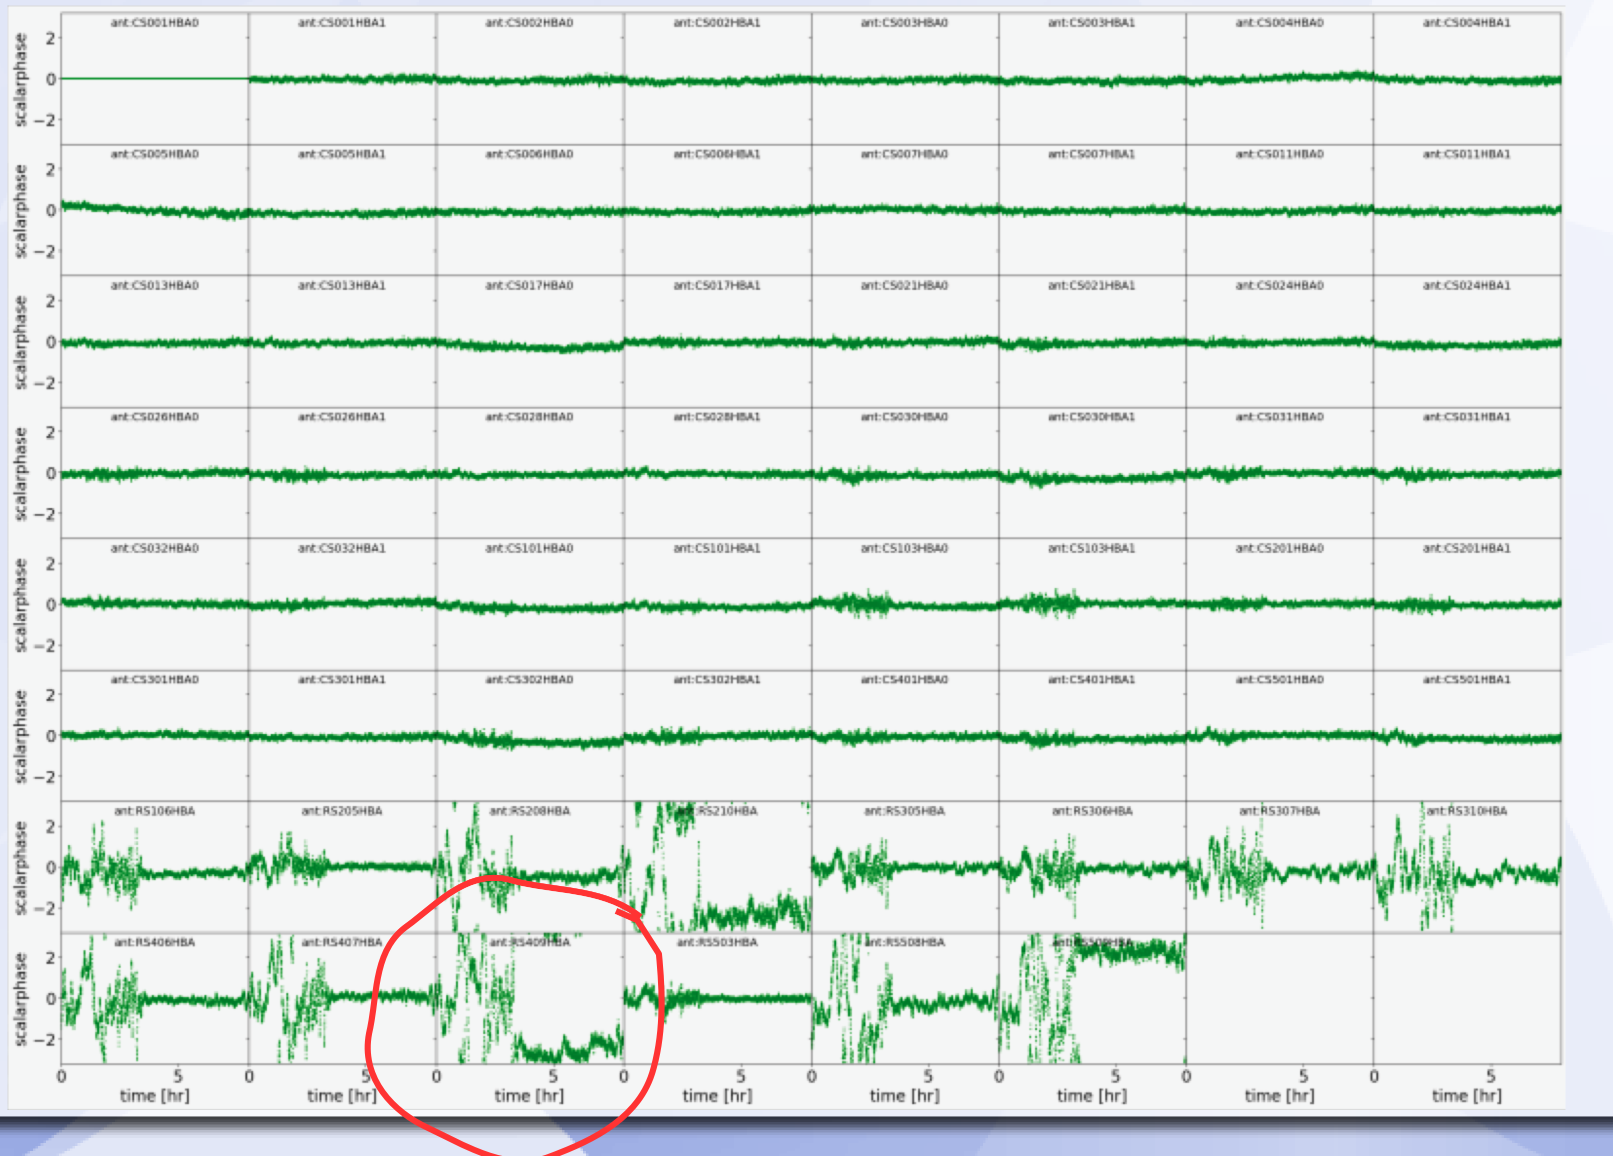

Result of constraining too aggressively during solve

Result of TEC constraint

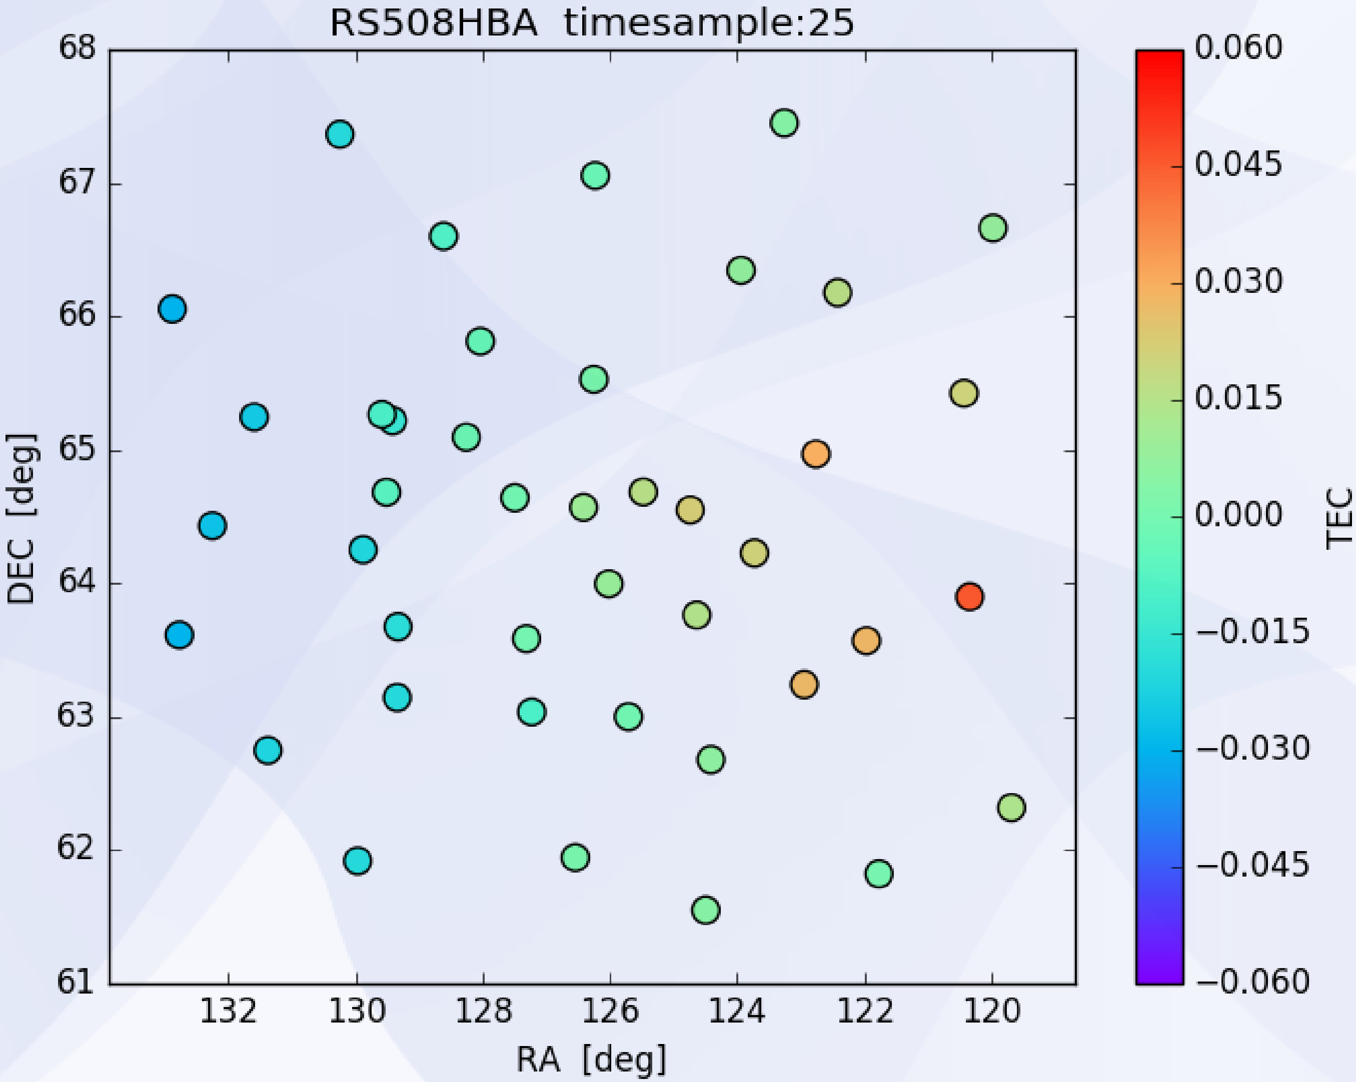

Plot by R. van Weeren, including work by F. de Gasperin, M. Mevius, B. van der Tol et al.

Result of TEC constraint

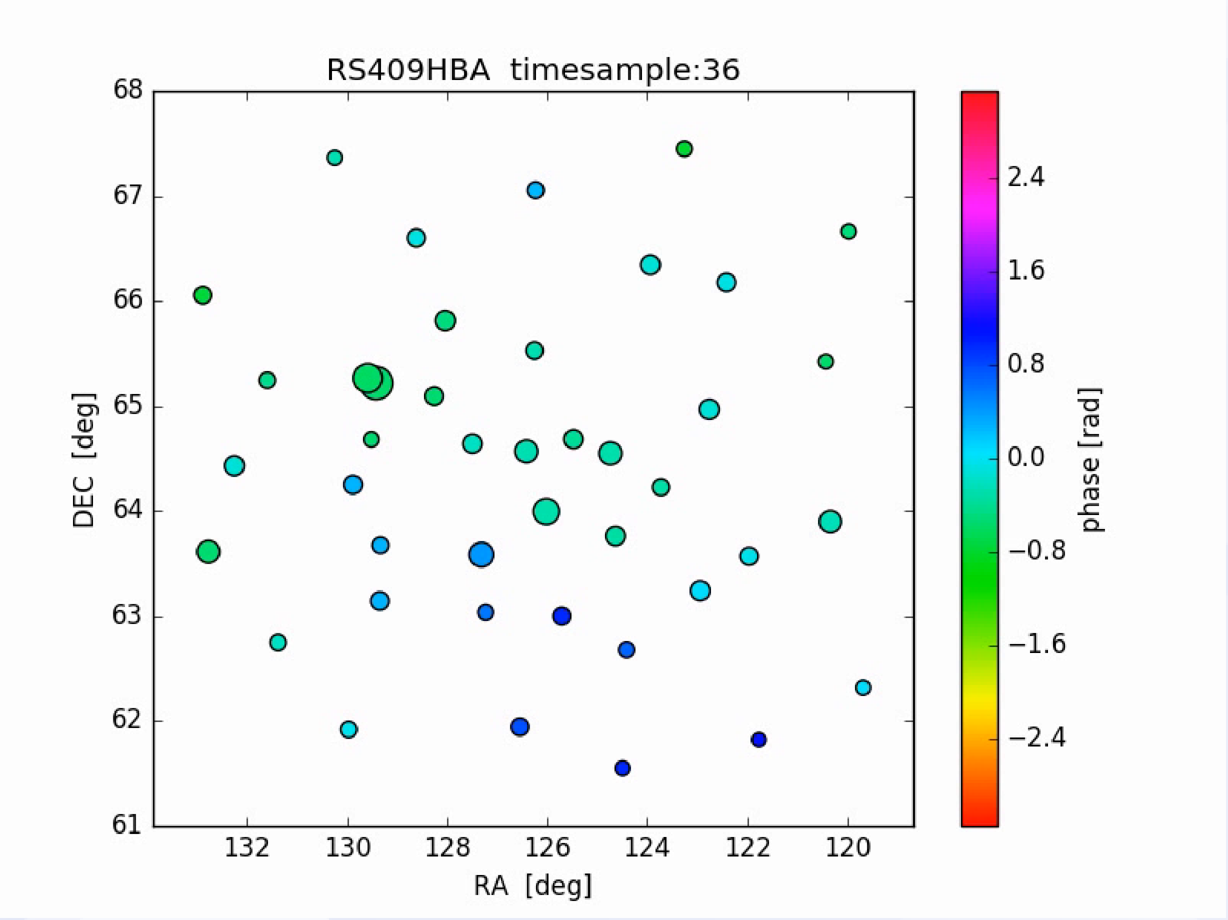

TEC solver

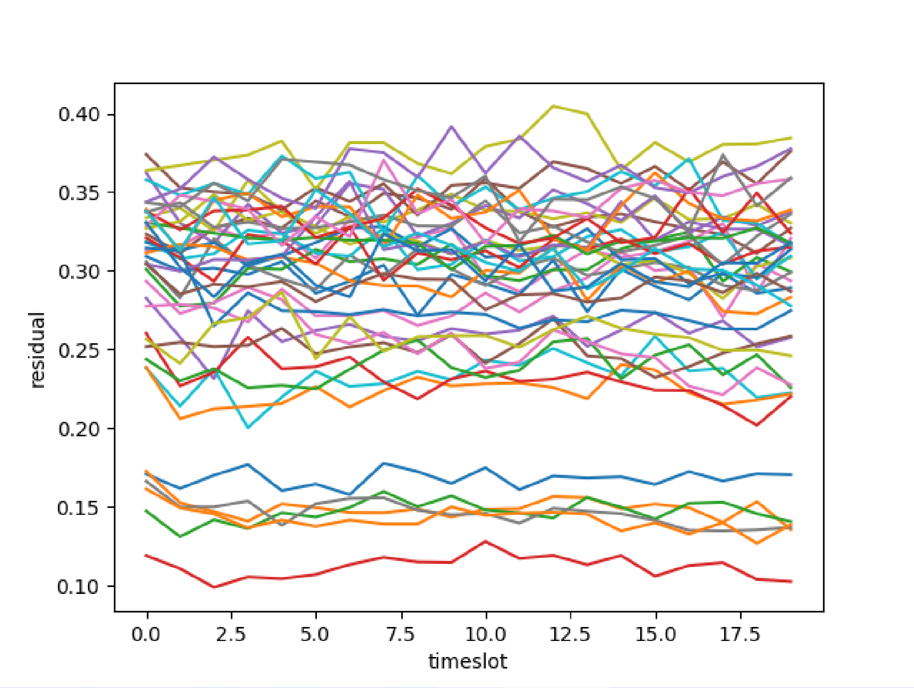

- Plot: fitting error for station RS509

  (difference of TEC fit with final solver step)

- Each line shows the error for one direction
- Error dominated by signal to noise

Modeling with WSClean

- WSClean (since 2.4) can directly output a beam corrected calibration model
- Consists of point sources, Gaussians and spectral information
- Directly readable by DPPP (T.J. Dijkema)
  - Allows DD calibration with WSClean + DPPP
- Local RMS method reduces false components
- Future goal: use IDG for prediction

Summary

- Low-frequency calibration & imaging very much still in development
- WSClean provides many new features:
  - Fast gridding & deconvolution
    - Ideal for LOFAR high res imaging & MWA phase 2
  - (LOFAR) beam correction
  - Fast multi-scale, multi-frequency (joined channels) clean
  - Fully automated (masked) cleaning
- Constrained multi-directional TEC solver in DPPP
  - State of the art algorithm, generic platform for any constraint
  - TEC solving is hard, but we can now do this

WSClean: http://wsclean.sourceforge.net/
Offringa et al. 2014
Offringa & Smirnov 2017

DPPP is part of the public LOFAR software