EVLA Memo 114

Image Pixelization and Dynamic Range

W. D. Cotton, Juan M. Uson
NRAO [1]

Abstract

This study investigates some of the effects of representing the sky by a rectangular grid of pixels on the dynamic range of images derived from radio interferometer measurements. In particular, the effects of image pixelization coupled to the CLEAN deconvolution representation of the sky as a set of discrete delta functions can limit the dynamic range obtained when representing bright emission not confined to pixels on the grid. Another effect which can limit dynamic range is a failure of the approximation that either the sky has no curvature over the field of view or that the antenna array is strictly coplanar. In this case, the image derived by the usual 2D Fourier transform techniques will be increasingly defocused away from the image tangent point. A technique to reduce these effects is described and examples from an implementation in the Obit package (http://www.cv.nrao.edu/∼bcotton/Obit.html) are given.

[1]The National Radio Astronomy Observatory (NRAO) is operated by Associated Universities Inc. under cooperative agreement with the National Science Foundation.

Contents

1.1 Introduction

With the new generation of high sensitivity interferometers to come on-line in the next few years, EVLA, ALMA, LOFAR, wide–field imaging will be necessary to achieve the sensitivity possible with these instruments. The problem is especially acute at lower frequencies where every field of view will contain several relatively bright sources at any time. The sensitivity of instruments such as the EVLA at lower frequencies may be compromised much of the time by artifacts of the bright sources in the field if these artifacts are not reduced to an acceptable level. This memo describes artifacts arising from using pixellated images to describe the sky and describes a technique for reducing them. All data manipulations discussed in this report used the Obit package (http://www.cv.nrao.edu/∼bcotton/Obit.html); plots were generated in AIPS.

1.2 Effects of Pixellated Images

It is generally convenient to represent the sky seen by an imaging interferometer as a set of pixel values on a rectangular grid. This is a good match to the widely used CLEAN deconvolution technique which represents the sky as a set of delta functions located the centers of cells of such a grid. A commonly used measure of the quality of an image is its "dynamic range", generally defined to be the ratio of the brightest pixel in an image to the RMS pixel-to-pixel fluctuation in areas away from emission.

One limitation of the pixelization technique is that emission not confined to points on the grid cannot be represented exactly and the CLEAN technique will approximate such structure by a potentially infinite series of alternating positive and negative delta functions. This is a problem particularly when the image contains very bright, unresolved emission as is common in the radio sky. A combination of the limited support of the actual representation and the effects of the finite

precision of digital computers will limit the dynamic range obtained by introducing artifacts in this image. This effect is easily understood for a point–like source located between pixels. This effect has long been recognized as a problem [2], [6] and has led to the technique of centering the point–like emission on a grid cell. In the past, this effect has not been a particularly serious problem as the bright source was usually the object under investigation and centering it in a pixel was straightforward. This is not the case for surveys or future observations where the source(s) of interest may be faint sources in the presence of multiple, much stronger sources and the locations on the imaging grid are not easily controlled. A simple adjustment was made to the imaging grid in the NVSS survey [3] to deal with this problem.

A second limitation of the pixelization technique is that the rectangular grid is flat whereas the sky is not (see [7]). In the case of an array confined to a plane, such as the E-W linear WSRT, a projection of the sky is possible which avoids this problem; but in the general case of a 3D array, this is not possible. This effect is referred to in the following as the coplanarity problem. If the field of interest is small enough, the curvature of the sky can be insignificant. Several techniques have been developed to deal with the more general problem. Some of these are:

- Full 3-D imaging
  Interferometric measurements are made in visibility space described by the coordinate set (u,v,w). These measurements can be convolved onto a 3D grid and Fourier transformed into a 3D image. The celestial sphere is a spherical surface in this 3D image. A 3D deconvolution followed by projection onto a plane is possible but in practice this is sufficiently expensive in computing resources that it is not used. See [7] for details.

- Fly's Eye
  The curved surface of the celestial sphere can be approximated by a mosaic of facets, each tangent to the celestial sphere and of sufficiently limited extent that the error introduced is insignificant. However, the errors increase quadratically with the distance of a cell from the corresponding tangent point and this can still limit the dynamic range. See "Polyhedron imaging" in [7] for details.

- W projection
  It is possible to correct for the diffractive effects on the wavefront as it propagates from the antenna closest to the source to the farthest on each baseline. This correction is made in the convolution of the visibility data onto the grid prior to Fourier transformation. This allows a single, flat 2D grid to represent the curved sky. See [4] for details.

In the following, only adaptations of the Fly's Eye technique are considered.

1.3 Wide–field Imaging with Fly's Eye and Catalog of Sources

The technique used for wide–field imaging in the following testing is as follows. The imaging in Obit manipulates image mosaics using class ObitImageMosaic. A field of view is specified by the radius from the pointing center to be completely imaged. The data is examined and the cell spacing (if not specified) is picked on the basis of the longest baseline in the data and the size of a facet "undistorted" by coplanarity effects determined from the maximum w in the selected data. A "Fly's Eye" tessellation of round regions in a hexagonal pattern is then defined which fully covers the field of view and a set of square images enclosing these circular regions defines a mosaic.

In general, the imaged field of view does not enclose all of the sky to which the array elements have significant gain, so there is the possibility of sources visible outside of the fully imaged field of view. To include such sources, facets are added to the mosaic centered on the positions of outlying sources from a catalog (currently a stripped–down version of the NVSS) which are deemed to have an apparent brightness above a given threshold based on a given spectral index and a model of the antenna pattern. At high frequencies, a catalog could be generated from the WMAP point sources [1].

Following the "OVERLAP 2" method of AIPS, all facets in the mosaic are formed and a quality measure based on both peak brightness and extended emission is used to determine which facet is to be CLEANed first. Components are selected from this facet and then subtracted from the visibility data. The next highest quality measure facet is re-imaged and if it is still the highest quality facet, it is CLEANed, otherwise, the next highest facet is re-imaged and tested, etc. After the CLEAN is finished, components are (optionally) restored to the facet from which they were subtracted as well as to any overlapping facet in which their position appears. After restoration of the subtracted components, all facets are projected ("flattened") onto a grid covering the specified field of view.

1.4 autoCenter Technique

In order to minimize pixelization related artifacts using the Fly's Eye technique, it is desirable to locate each strong, point–like source centered on the pixel that is at the tangent point on the facet containing it. It is also necessary to restrict CLEAN from assigning any components in any overlapping regions of other facets.

An initial CLEAN is used to determine which, objects in the field, if any, have sufficiently bright emission to warrant being centered in a special facet. The CLEAN needs to be deep enough that an accurate measure of the centroid of the source can be made to center it on a pixel. The implementation in the Obit package CLEANs to a factor 0.1 of the minimum value specified for special treatment. Also, since the residuals are not needed, it is not necessary to derive a full set of residual images at the end of the initial CLEAN. The decision to enter this mode is based on the highest pixel value in the CLEAN-able portions of the initial dirty images.

Peak brightnesses are determined from the sum of the CLEAN components within a given radius (cells within 62.5% of the PSF in current implementation) of each component derived from each facet. If this sum exceeds the threshold level, then the centroid of the peak is determined from the CLEAN components by a moment analysis. A facet, currently 96×96 pixels centered at the derived centroid, is added to the current imaging mosaic. It is possible that the same source may appear in several overlapping facets in the initial CLEAN so it is necessary to check that a given strong source has not already been added, a coincidence of less than two cells is considered the same source. If a source is already within 0.5 pixel of the tangent pixel of the enclosing facet, a new facet is not created, but the facet is recentered (if not already within 0.01 cell) and is considered in subsequent recentering operations.

In order to ensure that CLEAN assigns no components to the re-centered source appearing in overlapping facets of the initial mosaic, each of the facets is examined to see if it contains the position of the source to be re-centered. In any facet in which the re-centered source appears in the CLEAN-able region, a round "unbox", currently of radius 33 pixels, is added to that facet centered on the position of the re-centered source. This size is slightly smaller than the size of

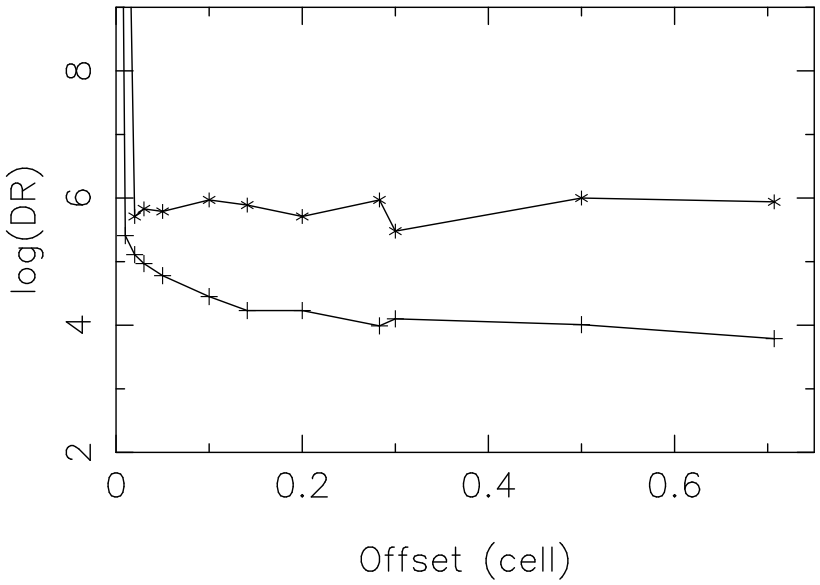

Figure 1.1: Dynamic Range(DR) obtained on simulated, noiseless data as a function of fractional pixel offset. A combination of offsets on a one and both axes are included. The "+" symbols represent CLEANing without using autoCenter and the "*" represent the results of using autoCenter.

initial cleanable region in the the new facet added, a radius of 38 pixels, to ensure that no pixel will be excluded a priori from a component search. These choices are somewhat arbitrary but seem to work well. An "unbox" is like a normal CLEAN window except that any pixel inside of it will not be considered as a location for CLEAN components even if inside of another regular CLEAN box, i.e. the CLEAN process ignores any pixels inside of an unbox. Pixels inside of unboxes are also excluded from statistical measures such as maximum, minimum and RMS. In Obit CLEAN tasks, the autoCenter function is invoked by setting the parameter autoCen to the threshold pixel value for re-centering. If any recentering operations are required, the initial CLEAN is repeated.

Because the accuracy of the determination of the centroid of a source to be re-centered is adversely affected by the limited CLEAN and the very effects this technique is trying to correct, some iteration may be in order. Images in which high dynamic range is desired also generally benefit from one or more iterations of self–calibration. At the beginning of each imaging cycle, the centroid of the previous iteration can be checked to see if it is sufficiently close to the tangent pixel. In the implementation in Obit, the centroid from the previous CLEAN is checked at the beginning of a CLEAN "major cycle". If the first moment of CLEAN components within 1.5 pixels of the center pixel is more than 0.01 pixel off, then the image is re-centered and all components from the previous cycle are discarded. This allows an iterative refinement of the centroid of the peak and significantly improves the results over a single estimate.

1.5 Examples

The following sections give two sets of examples using this technique. The first involves using simulated data and the second, actual VLA observations of 3C84.

Table 1.1: Offset Source Dynamic Ranges

Offset(pixels)	DR[1]	DR_corr[2]
0	$> 1.0 \times 10^{23}$	
0.01	256×10^3	
0.02	130×10^3	515×10^3
0.03	94.0×10^3	671×10^3
0.05	59.6×10^3	618×10^3
0.1	28.5×10^3	923×10^3
0.141	17.0×10^3	775×10^3
0.2	17.1×10^3	516×10^3
0.283	9.7×10^3	934×10^3
0.3	12.5×10^3	300×10^3
0.5	10.3×10^3	998×10^3
0.707	6.15×10^3	879×10^3
10	65.4×10^3	89100×10^3
20	16.4×10^3	95170×10^3
30	7.34×10^3	2853×10^3
50	2.70×10^3	2057×10^3
100	0.74×10^3	548×10^3

Notes:
[1] Dynamic range without autoCenter
[2] Dynamic range with autoCenter

1.5.1 Simulated Data

Simulated data have the advantage that their properties are known which makes it simple to separate artifacts from source structure. In these tests, the model was a single 1 Jy point source with no noise or other corruptions added. In order to test the effects of fractional pixel offsets, the source was at the phase center of the data and the data imaged with a series of fractional pixel shifts of the center of the image. The artificial data set was derived by replacing the observed data from a 74 MHz VLA data set by model values. This data set consisted of 12 frequency channels of 122 kHz bandwidth to reduce bandwidth smearing. The tests were performed in the Obit package using task Imager. The CLEAN used the visibility–based technique [5] and proceeded for 1000 iterations with a loop gain of 0.1 and a CLEAN window of radius 10 cells centered on the peak. The model visibility computation used the "DFT" method which is more accurate than the "GRID" method as implemented in Obit (and AIPS). Several iterations of phase self–calibration were done in order to refine the estimate of the centroid; the self–calibration otherwise has little effect as there were no errors added to the model. The cell spacing used was 0.25 of the fitted Gaussian restoring beam so the image was reasonably over–sampled. The components removed by the CLEAN were not restored and the RMS of the pixel values in the final residual image were used to derive the dynamic range. In this memo, dynamic range is defined as the ratio of peak value in the image to the off–source RMS. Each of the tests were then repeated turning on the autoCenter mode of Imager and the corresponding dynamic range redetermined. A combination of offsets on a single and on both axes are included. The results are shown in Table 1.1 and Figure 1.1.

A second set of tests explored the effects of non-coplanarity by inserting a series of large, whole–

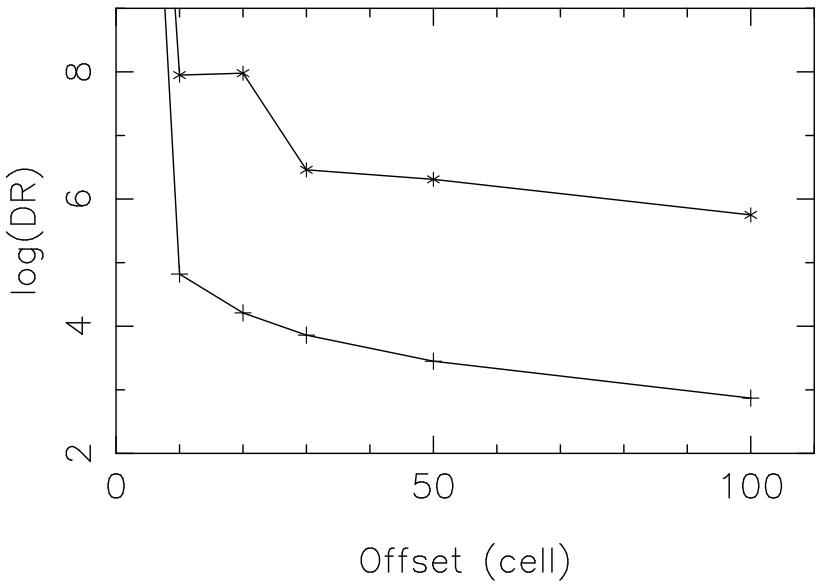

Figure 1.2: Dynamic Range(DR) obtained on simulated, noiseless data as a function of large, whole pixel offset offsets to test non co-planarity effects. The "+" symbols represent CLEANing without using autoCenter and the "*" represent the results of using autoCenter.

pixel offsets by offsetting the position of the point model and using a procedure like the one for small pixel offsets. The results are shown in Table 1.1 and Figure 1.2. The Obit routine to suggest the undistorted field size gives a radius of 95 cells (given the 4 cells per beam used).

While these tests are not exhaustive, it is clear that fractional pixel offsets of bright point–like sources can limit the dynamic range to of order 10^4 and non-coplanar effects can limit the dynamic range to of order 10^3 even with a moderately conservative limit on facet size. Applying the autoCenter technique improved the dynamic range by typically a factor of 50 for the fractional pixel offset tests and typically a factor of 350 in the large pixel offset tests.

Figure 1.1 shows a fair amount of scatter in the corrected dynamic range achieved. This is presumed to be the result of residual errors in the centroiding; corrections less than 0.01 of a pixel were not made. All corrected dynamic ranges were substantially better than that obtained using a 0.01 pixel offset without correction. At very high dynamic ranges, the accuracy of the centroiding needs to be exceedingly precise. This is possible in this test as there are no systematic errors and the source is exactly a point. In the real sky, resolution may be a problem even for a source whose size is a very small fraction of the PSF; position errors of even less than 0.0025 of the PSF seem to be capable of limiting dynamic range so resolution on similar scales might be a problem. In Figure 1.2 the efficacy of the correction seems to diminish with increasing pixel offset. This may be as much a problem in the simulated data as in the imaging; the longer position shifts needed to model a source with a substantial offset from the pointing center will result in loss of numerical precision.

Noncoplanar effects coupling to the pixelization can limit significantly the dynamic range achieved. A source observed 20% of the way to the edge of its imaging facet suffers comparable dynamic range loss to a source observed 0.2 cells from the closest pixel. This strongly suggests that imaging using the Fly's Eye technique needs to be applied with bright sources at the center of a facet (tangent point) and not merely on a pixel.

Table 1.2: 3C84 Dynamic Range

Image	Peak Jy	Near RMS Jy	Far RMS Jy	Near DR	Far DR
center[1]	24.6	0.0020	0.00122	12289	20145
half[2]	12.0	0.00239	0.00121	5029	9933
half/auto[3]	11.6	0.000991	0.00055	11745	21163

Notes:
[1] 3C84 observed at pointing center
[2] 3C84 observed at half power
[3] 3C84 observed at half power using autoCenter

A coarser grid spacing will likely lead to more serious errors than presented here.

1.5.2 Actual Data

A test using real data and wide-field imaging was made using the VLA at 1.4 GHz and observations of 3C84 (peak = 24.6 Jy). 3C84 was observed both at the pointing center and at the half power point of the beam. The observations were made in spectral mode with 15 × 390 kHz channels. The data were bandpass calibrated in addition to the amplitude and phase calibration and the edge channels were excluded from processing. Wide–field 3671×3671 pixel images were made of the primary beam of the antennas using the Fly's Eye technique and a 37 facet mosaic. In all cases, Imager applied amplitude and phase self–calibration to optimize the dynamic range. The data set with 3C84 at the half power of the antenna power pattern was imaged both with and without the autoCenter technique. In these images, 3C84 was half–way to the edge of its facet (coplanarity limit) and 0.4 of a cell from the nearest grid cell in RA and 0.3 of a cell in Declination. Sampling was 4 pixels per beam. Contour plots of a portion of the images around 3C84 are shown in Figure 1.3.

It is immediately obvious from Figure 1.3 that the autoCenter technique helps improve to the dynamic range of the image with 3C84 well away from the pointing center. This is explored further in Table 1.2 which gives the relevant image statistics. The RMS was determined in 601×601 pixel windows either centered on 3C84 ("Near RMS") or far from 3C84 ("Far RMS"). The near RMS values were determined using a histogram analysis.

As can be seen from Table 1.2, using the autoCenter technique doubles the effective dynamic range of the image, even in regions far from the obvious artifacts caused by 3C84 and does even better near the source. The autoCenter image has comparable dynamic range to the observation with 3C84 on axis. Especially in the neighborhood of 3C84, it is clear that other systematics (e.g. bandpass mismatches, pointing errors, beam squint) are limiting the dynamic range.

1.6 Discussion

In the preceding, we have demonstrated that the dynamic range obtained imaging interferometric observations can be adversely affected by pixelization of the images and non-coplanarity effects. Image pixelization effects can limit dynamic range to of order 10^4 even for point sources and non-coplanar effects can limit the dynamic range to of order 10^3 even with a moderately conservative

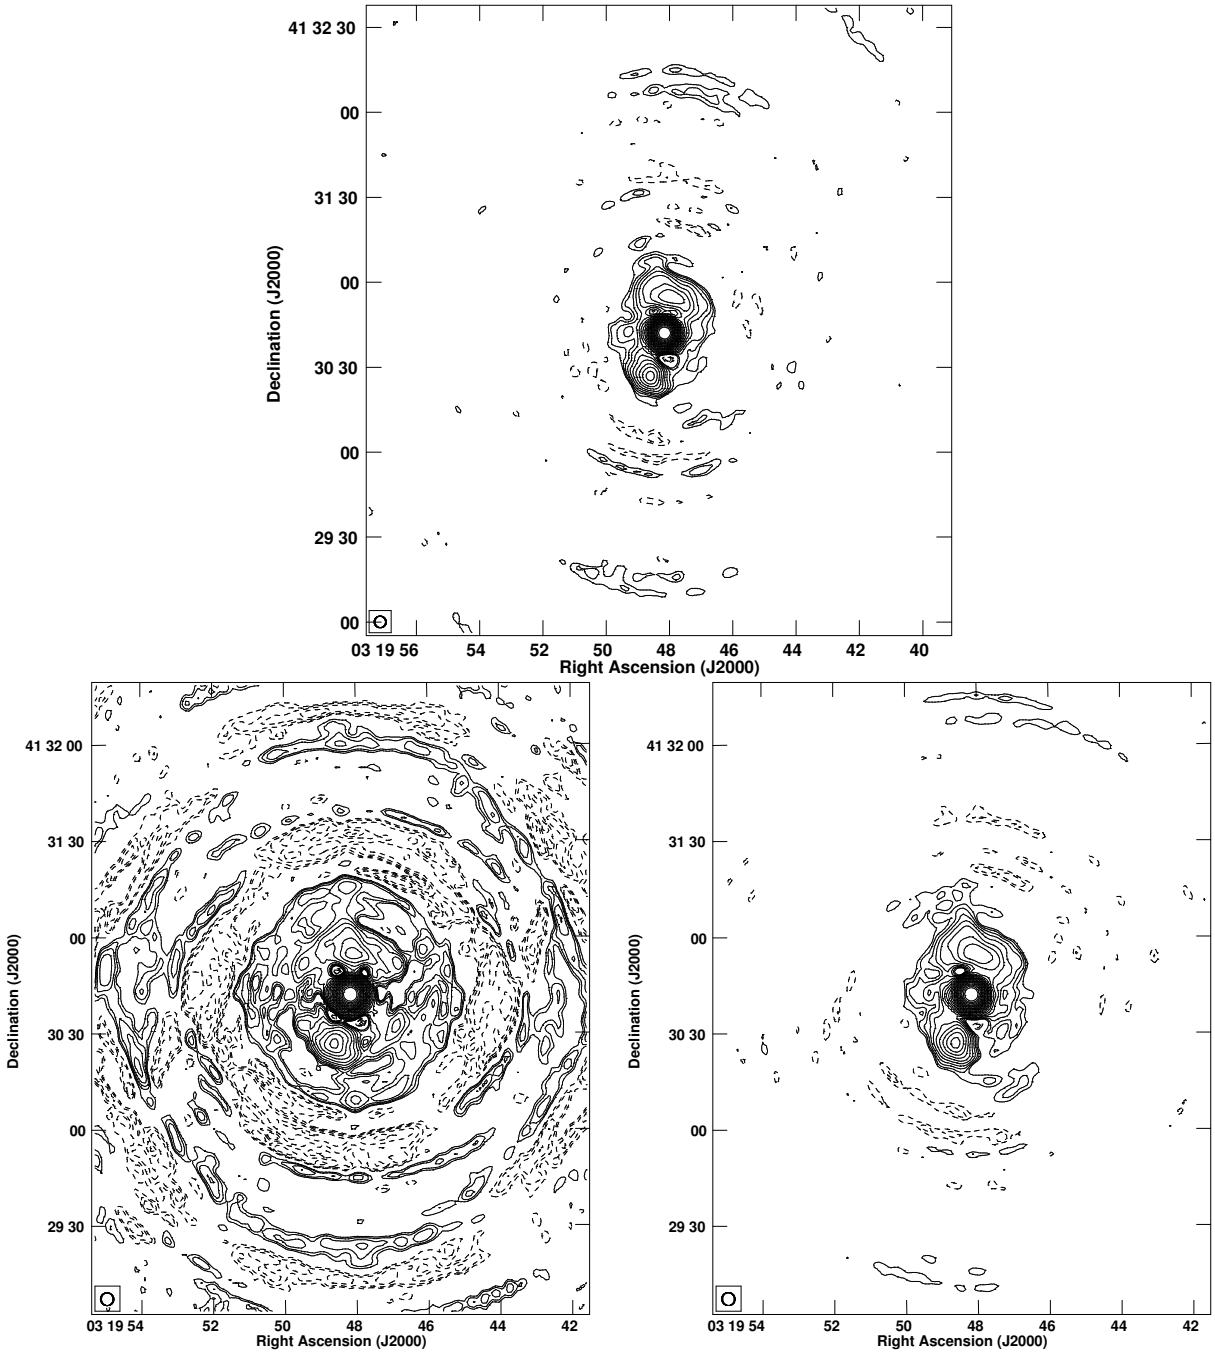

Figure 1.3: 3C84 with contour interval of powers of $\sqrt{2}$. The same contouring relative to the peak in the image is used for all plots.
Top: 3C84 observed at the pointing center. The most negative contour (dashed) is at -8 mJy/beam
Bottom Left: Portion of wide-field image with 3C84 at the half power of the antenna pattern. The most negative contour (dashed) is at -11 mJy/beam
Bottom Right: Like Left except using the autoCenter technique. The most negative contour (dashed) is at -4 mJy/beam

limit on facet size. The higher dynamic range needed for the EVLA and LOFAR, where wide fields of view with numerous bright sources will be common, need improved techniques to get around these limits.

The autoCenter technique presented here has shown factors of 50 to 500 improvement in simulated data with no noise or systematic errors. An improvement of over a factor of 2 was achieved in real observations of the bright source 3C84, even far from the obvious artifacts from the source. It is also likely that using the "DFT" rather than the less accurate "GRID" method of computing model visibilities will result in improved dynamic range.

Fractional pixel corrections are more difficult in a single image w-projection method. A simple shift can center a single source, also varying the pixel spacing could center two sources and a rotation could add a third but centering larger numbers of sources will not be possible. The introduction of separate, parallel "w–projection" grids to accomodate such a strong sources would seem to reduce the benefits of the "w–projection" technique. [2]

The technique described here can be applied to an arbitrary number of point–like sources. The tests presented here suggest that high dynamic range imaging of bright extended sources needs a better basis function than delta functions on grid cells as used by CLEAN.

1.7 Acknowledgments

The authors would like to thank Huib Intema of Leiden University for stimulating discussions on this subject.

[2] Alternatively, a finite number of tangent facets could be used to supplement the "w–projection" grid in a manner similar to the implementation discussed here. Again, this would seem to reduce the benefits of the "w–projection" technique.

Bibliography

[1] Bennett and WMAP team. WMAP Point Source Catalog. *NASA Web site*, http://lambda.gsfc.nasa.gov/product/map/dr2/ptsrc_catalog_get.cfm.

[2] D. S. Briggs and T. J. Cornwell. An Alternative Interpretation for the Physical Basis of CLEAN. In D. M. Worrall, C. Biemesderfer, and J. Barnes, editors, *Astronomical Data Analysis Software and Systems I*, volume 25 of *Astronomical Society of the Pacific Conference Series*, pages 170–+, 1992.

[3] J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick. The NRAO VLA Sky Survey. *Astron. J.*, 115:1693–1716, May 1998.

[4] T. J. Cornwell, K. Golap, and S. Bhatnagar. W Projection: A New Algorithm for Wide Field Imaging with Radio Synthesis Arrays. In P. Shopbell, M. Britton, and R. Ebert, editors, *Astronomical Data Analysis Software and Systems XIV*, volume 347 of *Astronomical Society of the Pacific Conference Series*, pages 86–+, December 2005.

[5] W. D. Cotton. Special Problems in Imaging. In G. B. Taylor, C. L. Carilli, and R. A. Perley, editors, *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, pages 357–+, 1999.

[6] R. A. Perley. High Dynamic Range Imaging. In G. B. Taylor, C. L. Carilli, and R. A. Perley, editors, *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, pages 275–+, 1999.

[7] R. A. Perley. Imaging with Non-Coplanar Arrays. In G. B. Taylor, C. L. Carilli, and R. A. Perley, editors, *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, pages 383–+, 1999.