VLB ARRAY MEMO No. 697A

REVISION

High Dynamic Range Imaging With the VLBA

D.S. Briggs
R.J. Davis
J.E. Conway
R.C. Walker
National Radio Astronomy Observatory *

July 23, 1994

1 Introduction

The VLBA is being tested for high dynamic range imaging using a complete track on DA193. All the data have been fringed, calibrated and edited by Craig Walker. Here we report on attempts made to image just one of these IF channels. In principle this data set had the possibility at being imaged at $\gtrsim$28,000:1 dynamic range. This is well beyond current limits achieved in VLBI. Furthermore the VLBA has the potential in terms of signal to noise to exceed dynamic ranges of 150,000:1. At this level we are concerned about non-closing errors in the correlator and station hardware, self calibration errors, and deconvolution errors. Our current work achieves a measured dynamic range of 27,700 and extremely good agreement with the a priori thermal noise. The best images show no sign of any hardware effect or processing instability. Significant artifacts are seen in the CLEAN only deconvolution image, but several other techniques are investigated with successful results.

2 Initial Calibration

The experiment consisted of full tracks on DA193 at C band, for a total observation of 17 hours. Data were recorded for 13 minutes out of every 18. The observation was single

*The National Radio Astronomy Observatory (NRAO) is operated by Associated Universities, Inc., under contract with the National Science Foundation.

polarization, RR, and after bandpass calibration the effective bandwidth was 1.75 MHz.

The initial calibration was done in AIPS. The amplitude calibration was done using the a priori gains and system temperatures. Since the overall scaling factor was not known, it was adjusted to force the total flux density of DA193 to match that measured in a special pointing run made a few days before. Large gain adjustments were required for Pie Town (pointing problem) and Hancock (snow). The data were edited based on a priori edit information provided by the VLBA system and based on system temperatures and known antenna problems. Once all of these edits were done, and low weight points were removed, no further editing was required. Fringe fitting was done with `FRING` and a bandpass calibration was done after fringe fitting. The bandpass calibration was done after fringe fitting so that proper amplitude corrections would be made, to the fully calibrated data, for amplitude losses due to the phase slopes caused by delay offsets. Finally the data were averaged in frequency, omitting the end channels, and in time. No baseline dependent (closure) calibrations were made.

3 CLEAN Deconvolution

Initial attempts to image IF 3 were made with CLEAN, strongly favoring only positive components. This produced images with clear rings of emission around the source and covering the entire map. Allowing negative components solves this problem and produces the maps of figure 1.

The extended halo with twin banana like structure is similar to the dirty beam and is probably non-physical. At this stage we were worried about the possibility that this could have been caused by closure problems. To this end we processed IF 7 through the same software route. This IF channel has been through completely different hardware on its route to the workstation disk. The two IFs passed through different filters at the stations, different samplers, tracks on the tape, decoding at the correlator, delay lines, phase rotation, correlator chips and so forth. However, within mapping errors which are bound to be a few times noise given the selfcal correction, these images are identical. Subsequent images were all made from IF 7.

Examination of the CLEAN components shows 4 alternate rings of positive and negative CLEAN components. Studies and simulations, [1], have shown this to be due solely to the CLEAN deconvolution algorithm. Tight CLEAN windows actually make the solution somewhat worse. The peak of the extended emission in these images is roughly 0.1%. The dynamic range is 9000:1, worse than the theoretical of 14500:1. It was clear that a better means of deconvolution is required, especially when the VLBA is used at the full recording rate of 128 MBits/s instead of the effective 3.5 Mbits/s used in this experiment. Several additional deconvolution/calibration strategies were explored, all of which were successful.

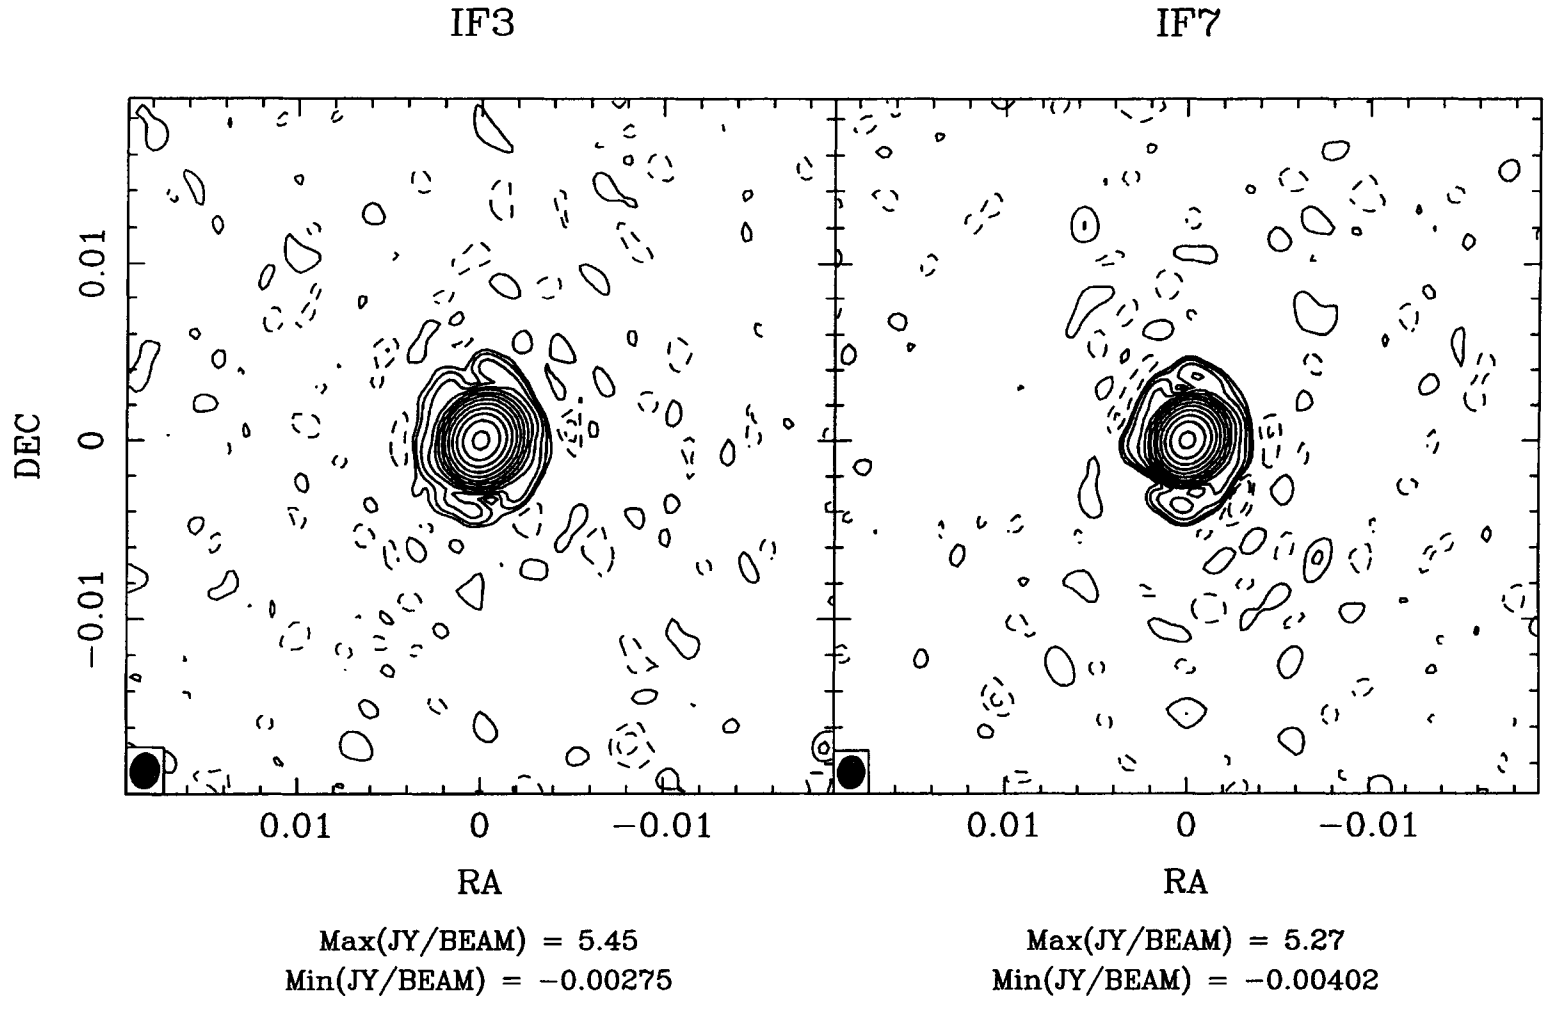

Figure 1: CLEAN images of two different IFs produced in `difmap` by RJD. Peak to offsource RMS is 9000 and 9400 for IFs 3 and 7 respectively.

4 Imaging Presentation and Metrics

In situations where there is a significant mismatch between the main lobe of the beam and its fitted Gaussian, the issue of the residual flux scale is important and directly effects the meaning of "dynamic range." The images presented here have all had the residual noise cleaned to convergence. This practice does distort the noise histogram somewhat, and it also causes the noise RMS to deviate from that calculated with the standard expressions. (This is not surprising, since the units of these expressions are strictly Janskies per dirty beam, and the whole point of cleaning to convergence is to to change units to Janskies per restoring beam.) However, so far as is possible the restored model and noise are on a common scale and this is our justification for quoting dynamic ranges based on the cleaned noise images. Dynamic range in our usage is the peak in the image over the offsource rms in the cleaned noise images. By contrast, our noise comparison with the a priori sensitivity calculations are based on the off source RMS in the images without noise cleaning. For the naturally weighted fitted beam of 2.7×2.3 milliarcseconds, the noise cleaned offsource RMS drops by 19% compared to the image where only the source is cleaned. For the uniformly weighted beam of 1.8×1.5 mas it drops by 28%. For the super resolved restoring beam of 1.0×1.0 mas, it drops by 47%. Thus while our definition of dynamic range is self consistent, measured values can exceed the nominal values derived by simply dividing peak source strength by thermal noise. Our theoretical dynamic ranges have been corrected for the noise rescaling effect.

All images presented here are contoured logarithmically, with a factor of 3 between contours. Unless stated otherwise in the caption, the lowest positive contour is .02% of the peak.

5 Model Fit with CLEAN Tidying

John Conway has fitted the bulk of the emission with a two component elliptical Gaussian model. The residuals from the modelling processes were then processed with CLEAN. This image was then used as the basis for one iteration of self calibration and remapped with the same model. The result is Fig 3 which shows no artifacts in the noise and a smaller extended region around the source. It seems likely that further work with self calibration might improve the image still, but this has not been done. This image is given in the left half of figure 2. The dynamic range is 10500:1.

6 NNLS

Dan Briggs has brought into operation a new deconvolver based on on Non Negative Least Squares matrix inversion. The current implementation exists in the package SDE as the task svdconv, and takes FITS format dirty map and beam as input. The algorithm is a quite straightforward application of a rather complicated preexisting constrained linear algebra algorithm. The task forms the matrix connecting an arbitrary region in the dirty map with another arbitrary region in the component plane, and solves the resulting algebraic equations with the NNLS algorithm from [2]. The algorithm and variations will be presented in more detail in his dissertation, [1], but the initial results are very encouraging. NNLS deconvolution is somewhat slower than existing algorithms for compact sources, and very much slower for extended objects. Memory is the limiting factor to the source size that can be deconvolved. One copy of the beam matrix must be held in memory, of size $N_{data}\,N_{flux}$, where N_{data} is the number of pixels in the dirty map used as input to the algorithm, and N_{flux} is the number of pixels where the algorithm is allowed to place flux. It is important that this fit into physical memory, as once page swapping sets in performance degrades dramatically, and problems that were just practical become quite impractical. Running time is roughly proportional to $N_{data}\,N_{flux}^2$, and also varies with the SNR of the data, with higher quality data taking longer to process. Currently a map with approximately 6000 pixels of significant emission and high SNR can be deconvolved in several hours on an IBM RS/6000.

The strength of the algorithm is obviously related to the positivity constraint, without which the problem is extremely ill posed. Since the NNLS algorithm is a direct solution, it does not suffer the problem of a spatially iterative algorithm like CLEAN, where an initial error is made at the start of the algorithm before the support constraints have a chance to become effective. This initial error produces the ring like pattern noted earlier which is never subsequently removed. As NNLS is a direct solver, the support constraints enter the solution in a globally optimal way, which appears to be quite important for self calibration. NNLS shares the positivity constraint with the Maximum Entropy method, but unlike MEM it has no entropy (smoothness) criterion. This actually seems to be an advantage for NNLS when deconvolving compact structure. Noiseless simulations on both Gaussian and disk sources show NNLS to be much better than either MEM or CLEAN for compact objects less than a few beamwidths across. The deconvolution improvement degrades with increasing source size, but even at the computational limit of $\sim$6000 pixels, it is still the best deconvolution algorithm studied.

The right panel of figure 2 was made from data calibrated entirely in SDE using the NNLS deconvolver in the hybrid mapping loop. User input was minimal, with a single loose window around the entire region containing probable flux. The entire procedure was automated in a shell script, and the solution converged after 5 phase-only self calibration

iterations, and an additional 25 amplitude+phase iteration. Note the striking similarly of this image to that of the model fit + CLEAN procedure. These two images calibrated and produced in completely different packages nearly convinced us that the extended emission might be real.

However, when a human entered the loop and placed a much tighter window on the allowed flux, the solution improved still further. Five more iterations of self calibration with a tight window yields the left image in figure 4. A naturally weighted map made from the same calibrated data set is also given. The extended halo is completely gone. The dynamic range is 13,700 for uniform weighting and 27,700 for natural weighting!

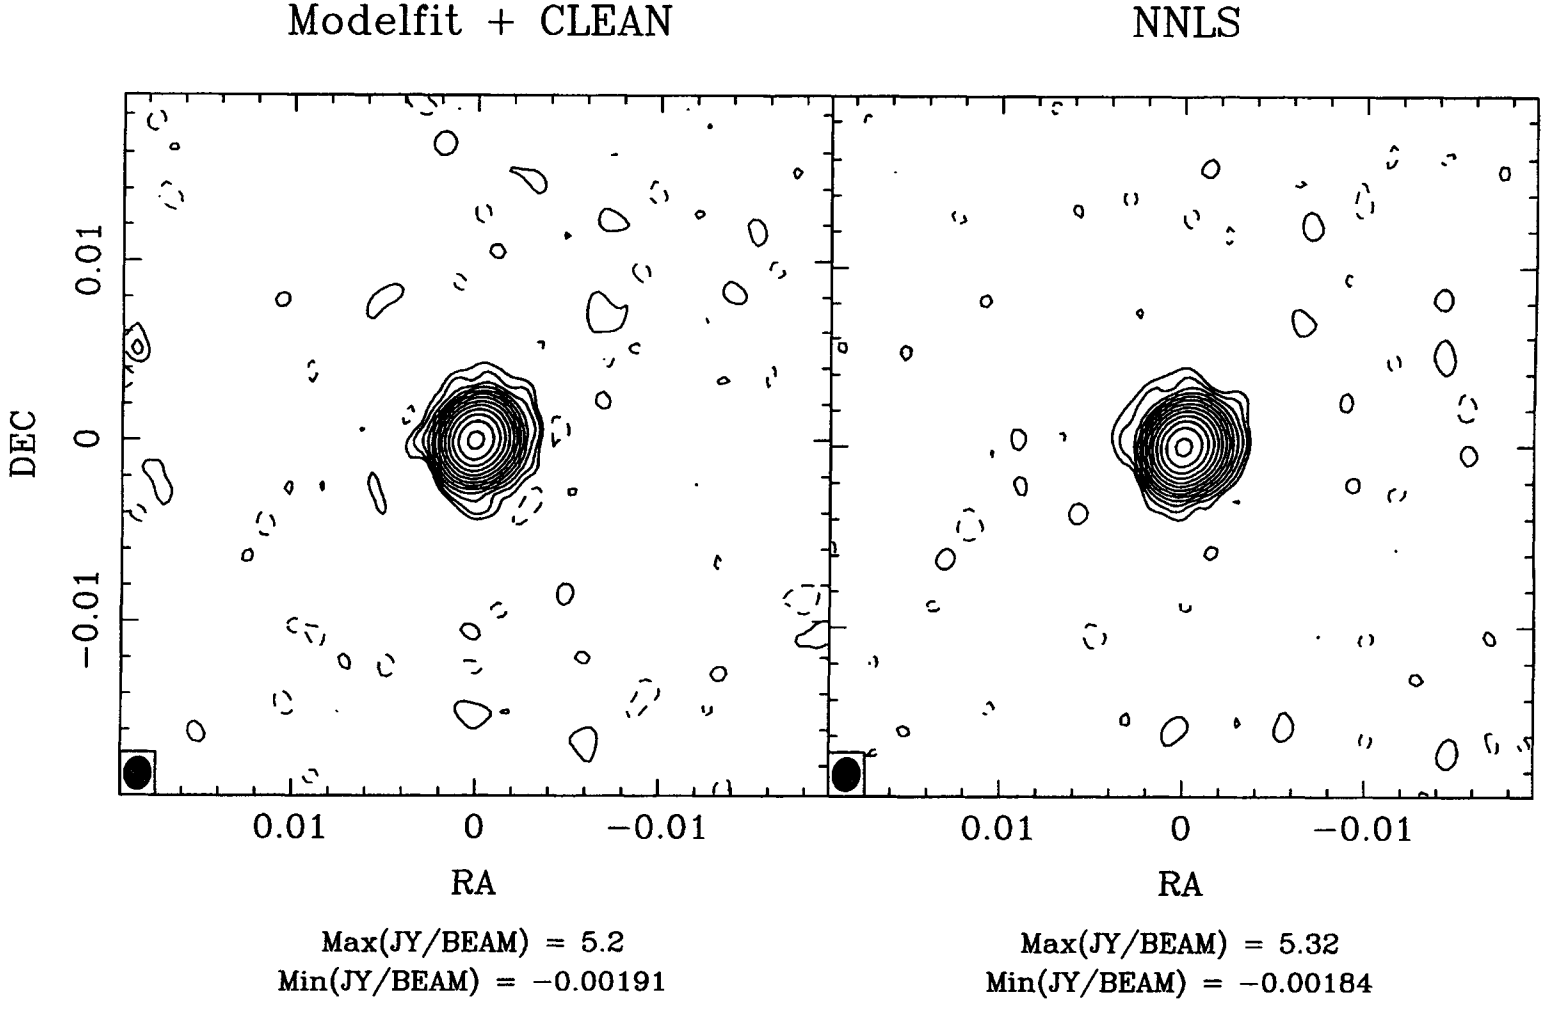

Figure 2: The left panel is JEC's image produced and calibrated in **difmap** via model fitting and cleaning the residual flux. The right panel is DSB's NNLS image produced and calibrated within SDE. Notice the extreme similarity in the (erroneous) extended structure. Dynamic range for the two images is 10,500 and 11,400 respectively.

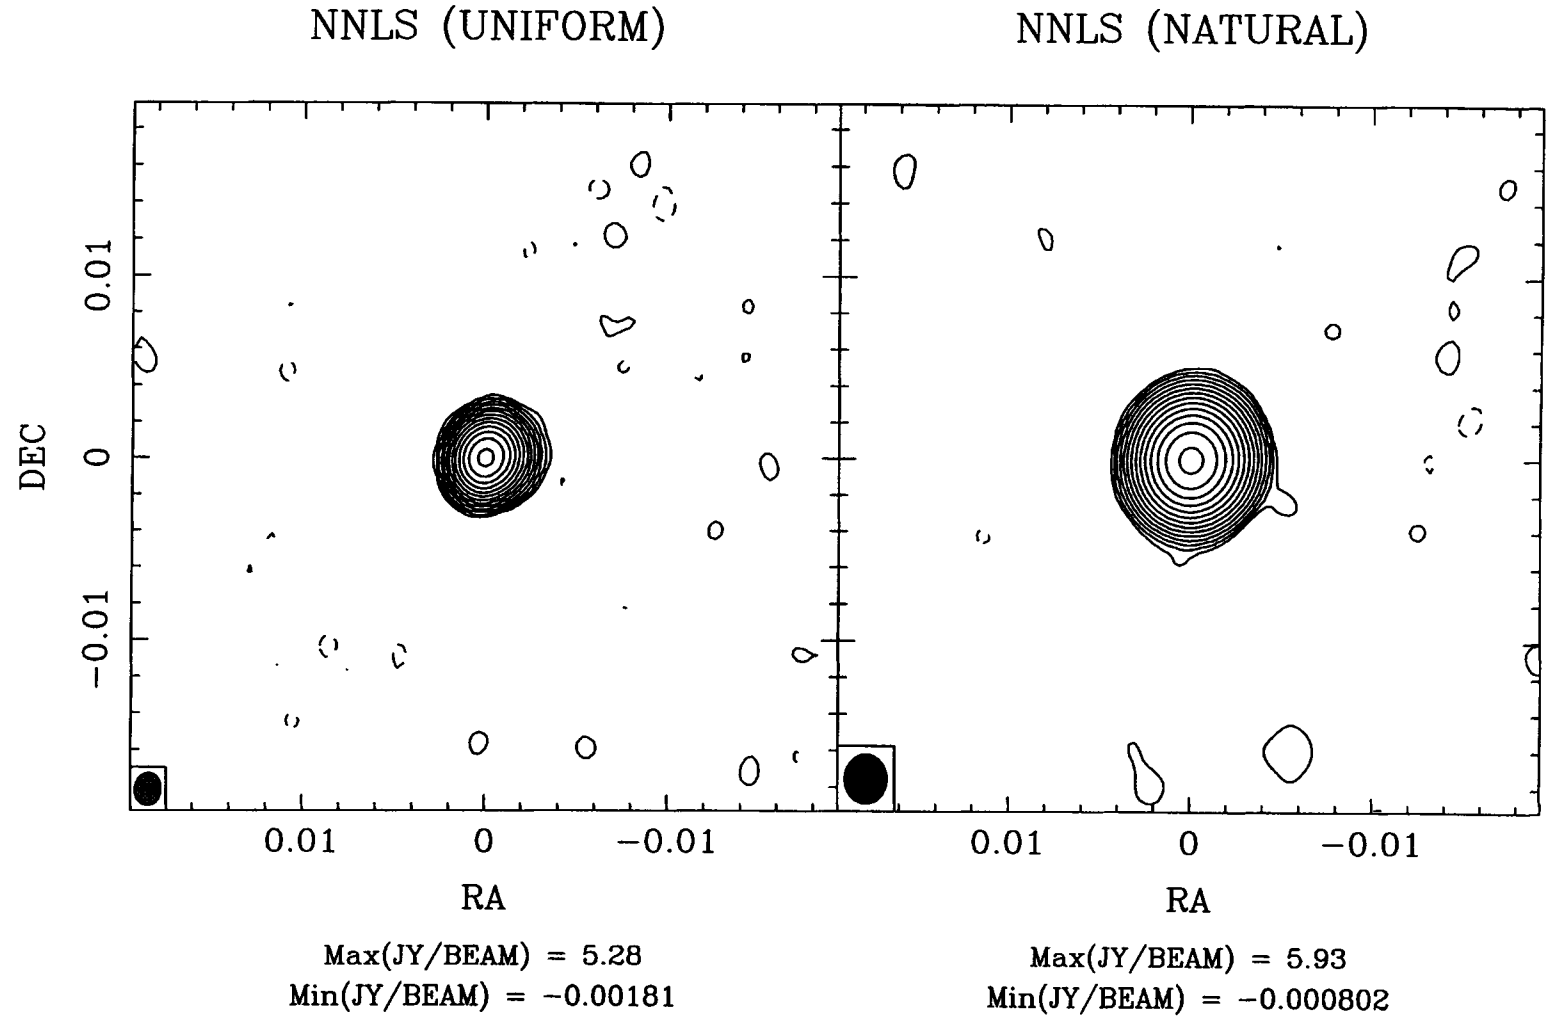

Figure 3: DSB's best NNLS images, with uniform and natural weighting. The diffiference between this figure and the right panel of figure 2 is that extremely tight windows were used in the final stages of self calibration. The naturally weighted image is contoured another factor of 2 deeper than other images presented, so that the lowest positive contour is .01% of the peak. The uniform dynamic range is 13,700 and the natural dynamic range is 27,700!

7 Noise Comparison

From the VLBA Observational Status Summary of 23-June-1994, the zenith Source Equivalent Flux Density for C band is 296 Jy. For a full track observation the effective SEFD will be degraded somewhat by elevation dependent effects, but we will use 296 to derive a best case lower limit. Equation 3 of the status summary gives the following standard expression for a single polarization naturally weighted map plane dispersion.

$$\Delta I_m = \frac{1}{\eta_S} \frac{SEFD}{\sqrt{2\,\Delta\nu\, t_{int} N_{base}}}$$

For this observation $\Delta\nu = 1.75 \times 10^6$. The data set contains 19989 visibilities, the great majority of which are 60 second integrations. We replace $t_{int}\, N_{base}$ with $\Delta t\, N_{vis}$. The factor of $1/\eta_S$ accounts for system efficiency losses and is not yet well measured for the VLBA. In the same, "lower limit" spirit, we use $1/\eta_s = 1.57$ which is appropriate for a perfect one bit correlator at Nyquist sampling. The numbers yield an expected $\Delta I_m = .227$ mJy.

Software also exists in SDE to allow the estimation of thermal noise directly from the measured visibility data. The visibility averaging task can solve for the proportionality constant between the weights and inverse variance from the scatter in finite averages. After subtracting off the transformed NNLS component model, leaving mostly thermal noise, it seems to work quite well in this case. The SDE mapper takes proper account of weights and given this constant can calculate both the expected RMS thermal noise in the naturally weighted map and also the degradation caused by using non-natural weights.

The averaging task produces the following proportionality constants as a function of integration time. 150 seconds: .003042. 300 seconds: .003035. 600 seconds: .003042. The point here being that it is well behaved at the one part in several hundred level. The noise behavior is close to Gaussian.

Using the 300 second figure, the expected thermal RMS noise in the maps are .259 and .508 mJy for natural and uniform weighting respectively. The measured offsource RMS from the uncleaned final NNLS images are .264 and .538 mJy.

The conclusion we draw from this is that the NNLS images are very close to the thermal RMS in terms of the noise actually present in the data. In agreement with visual inspection of the images, neither the deconvolution nor the calibration nor non-closing errors (all of which should vary slowly on the scale of 300 seconds) are significantly contributing to the off source rms. In addition, the measurements are matched exactly by noise limited observations, theoretical correlator efficiency, and a mean SEFD of 340 Jy. Considering the known problems at Pie Town and Hancock, plus the effect of non-zenith observations, this is extremely good agreement between the measured thermal noise and a priori noise estimates.

8 Super Resolution

In previous super resolution work on SN1987A, [3], the deconvolver of choice was MEM. At the low signal-to-noise ratio of the SN1987A data, pure deconvolution error due to incomplete sampling was completely dominated by thermal noise and the choice of algorithm was made on other criteria than high precision. MEM extrapolates well in the Fourier domain and has the property that image features are made only as sharp as mandated by the data. In fact, since agreement between the data and the model is enforced in a statistical sense only, the MEM model image[1] will always be smoother than the true image. Related to this is that the model image has an effective resolution greater than that imposed by the pixel size, and the transform of the model drops towards zero at the highest spatial frequencies. The net effect of these properties is that one can view the MEM model as an image (without residuals) that has a spatially variable resolution as high as warranted by the data. Features present in this image are guaranteed to be mandated by the data though whether by the astrophysical source or by the thermal noise fluctuations is ambiguous. Conventional astronomical usage convolves the MEM model image with a restoring beam and adds the residuals back to the image, but the MEM algorithm itself limits the maximum resolution attempted in the reconstruction, no matter what restoring beam is selected.

By contrast, while NNLS does a superb job of extrapolation in the Fourier domain near the envelope of the measured data, the spectral power does not necessarily roll off at the highest spatial frequencies. At some resolution the solution must become spurious. Direct examination of the NNLS model image in simulations shows decidedly non-physical structure where flux in several adjacent pixels has been collapsed into one pixel of greater value surrounded by several zeros. With NNLS the onus is on the user to select a restoring beam that provides an acceptable degree of superresolution without producing obvious artifacts. Unfortunately, the restored image RMS and maximum negative as a function of restoring beam often gives little guidance as to image fidelity and a degree of arbitrary judgement must be exercised.

By itself, the CLEAN algorithm combines the worst features of the other two algorithms when used for superresolution. The high spatial frequency reconstruction is not generally as good as that of NNLS, but it requires a similarly delicate choice of optimum restoring beam. However it can be used in conjunction with MEM or model fitting. MEM suffers from the problem that even when the RMS residuals approaches the noise, there can still be considerable beam-like structure left in the residuals. Model fitting to the thermal noise limit is simply very difficult for all but the simplest sources. In both of these cases, CLEAN can be used to remove any residual structure from the residuals.

[1] Actually, it is the entropy (smoothness) in the ratio between the model image and the best a priori estimate of the image that is maximized. Usually the latter is simply a constant value.

Provided that the residual structure is small, this does not destroy the superresolution properties of the solution, and the MEM + CLEAN approach can be used in situations where computational limits preclude the use of NNLS.

In figure 4 we present images superresolved from the nominal 2.0 × 1.6 mas to 1.0 × 1.0 mas. These are produced from the same calibrated data as in figure 3, the best NNLS calibration. The MEM algorithm was allowed to place flux in a fairly loose box around the source, though the details of the window made relatively little difference. The East–West extension of the source traces nicely the jet seen at higher frequencies. Great care must be exercised in interpreting super resolved images astrophysically, but the structure here at least seem plausible. A NNLS image superresolved to 0.6 mas was rejected for showing obvious artifacts.

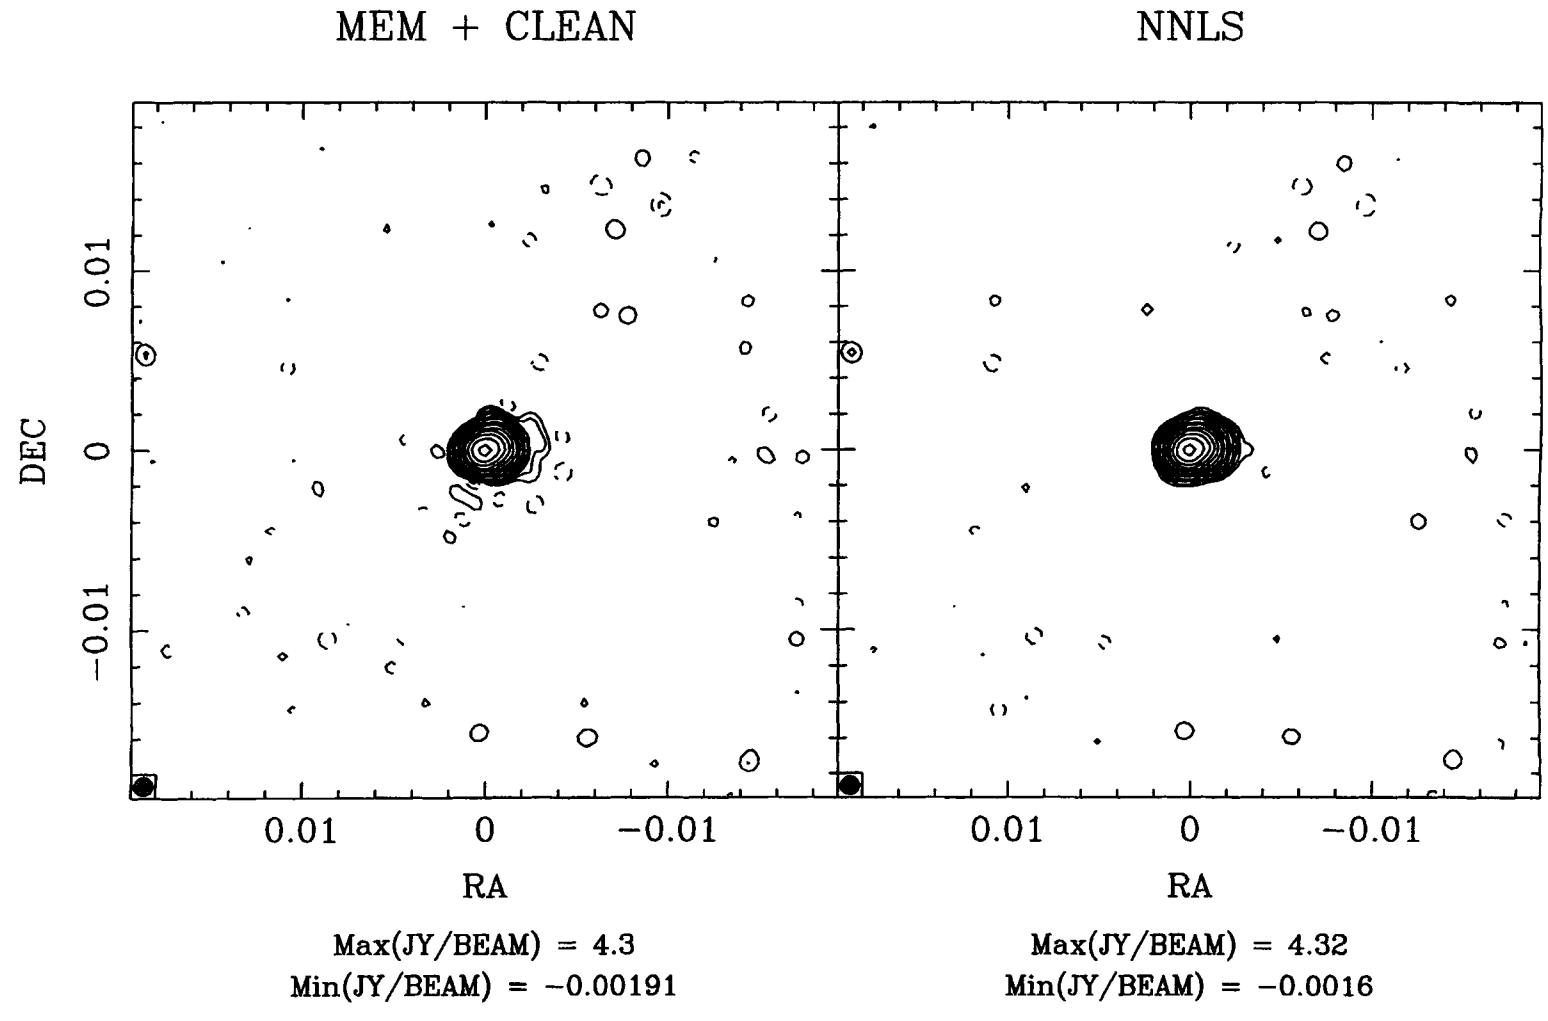

Figure 4: Superresolution. The left panel is produced with combination of MEM and CLEAN, the right with NNLS.

9 Self Calibration

Clearly the interaction between the deconvolution step and self calibration is critical in the high dynamic range limit. Both CLEAN and MEM produce much better images on the best calibrated NNLS data set than they do on data calibrated via other methods. MEM + CLEAN on the best data set produces an image comparable to NNLS, though either alone is somewhat worse.

NNLS seems to have several properties that are highly desirable in the self calibration loop. With very little user input and loose windows it converged to a solution as good as that of the best other methods. This property allowed it to be used in an open loop shellscript where many selfcal iterations were performed automatically without user interaction. Once converged to this intermediate solution, examination of the model components provided clear delineation of the source support, and selection of a tighter flux window was straight forward. The tighter window quickly drove the iteration to the final solution.

NNLS has a clear advantage over MEM in self calibration in that it is not biased against compact structure. NNLS has an advantage over CLEAN in that it makes efficient use of the support constraints. With CLEAN, the characteristic errors have already been introduced into the solution before the support constraints make any difference in the iteration. NNLS encodes the support information in the structure of the beam matrix and this affects the solution in a global manner.

The majority of the calibration can also be performed with lower precision deconvolvers for speed reasons. The calibration required 35 iterations to converge to the highest quality solution with NNLS, using first the loose and then the tight flux window. The calibration still required nearly the same 35 iterations to converge, restarting from the beginning and using the tight window throughout. Until the last few iterations, the intermediate solutions were nearly identical. Similar experiments varying the self calibration integration time yielded similar insensitivity to parameters until the last few iterations. The final calibration allowed independent amplitude and phase solutions for every 60 second visibility average. When the calibration was attempted with MEM as the deconvolver, intermediate solutions were very similar to that of NNLS until the MEM solution ceased to improve at about 15 iterations. Highest quality deconvolution seems necessary only at the end, and the final calibration with NNLS seems stable to different intermediate routes to the solution.

10 Conclusions and Further Work

- No obvious artifacts at all in the best images, and good agreement with a priori thermal noise. No evidence for closure errors at the 27,700:1 level! There appear

to be several workable approaches to high fidelity imaging with the VLBA.

- NNLS is very promising. It is the best deconvolver when it can be applied, and it has the advantage being very forgiving about windowing. This is an excellent algorithm in the self calibration cycle.
- Model fitting + CLEAN is a workable imaging approach and will likely drive self calibration well.
- CLEAN & MEM individually produce image artifacts even on the best calibrated data set, but the combination is very good and not limited by source size. This does not drive self calibration well.
- It is a bit disturbing that the first NNLS attempt & the Model fitting + CLEAN are so similar. Be conservative when interpreting such image features.
- NNLS seems to superresolve very nicely in this case – but more work would be needed before publishing it as astrophysics!
- Further testing continues on a 128 Mbits/s data set which will enable the VLBA to be tested at 150,000:1 dynamic range.

References

[1] D. S. Briggs. *High Fidelity Deconvolution of Moderately Resolved Sources.* PhD thesis, New Mexico Institute of Mining and Technology, Socorro, New Mexico, 1994. in preparation.

[2] Charles L. Lawson and Richard J. Hanson. *Solving Least Squares Problems.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1974.

[3] D. S. Briggs. Superresolution of SN 1987A: A Comparative Deconvolution Study. In Robert J. Hanisch and Richard L. White, editors, *The Restoration of HST Image and Spectra – II*, pages 250–256, 1994, Space Telescope Science Institute, Baltimore, MD