SKA MEMO 60, 24 May 2005

Real Time Imaging

Melvyn Wright

Radio Astronomy Laboratory, University of California, Berkeley, CA, 94720

ABSTRACT

In this paper, we propose to integrate the imaging process with the correlator hardware in order to handle the high data rates and imaging problems for radio telescope arrays with large numbers of antennas and large fields of view. We use FX correlators and beam formers with a high data bandwidth into computer clusters to support a flexible programming environment. The correlation function is computed with narrow frequency channels and short integration times so that images can be formed over a large field of view. Images can be made simultaneously in multiple regions within the field of view by integrating the output from the correlators at multiple phase centers on targets of interest, calibration sources, and sources whose sidelobes will confuse the regions of interest. Calibration is made in close to real time using a model of the sky brightness distribution. The derived calibration parameters are fed back into the imagers and beam formers. Images are made simultaneously for multiple phase centers using an FFT algorithm in restricted fields of view. Sidelobes from sources outside each of the regions imaged are minimized by subtracting the model from the *uv* data before imaging. The regions imaged are used to update and improve the a-priori model, which becomes the final calibrated image by the time the observations are complete.

1. Introduction

The Square Kilometer Array is a next generation radio telescope which will be able to form simultaneous images in multiple regions within the field of view. The current SKA science requirements (D.L.Jones, 2003), specify imaging multiple regions within the FoV ($\sim$ 1 degree at 1.4 GHz) with an image fidelity 10^4 between 0.5 and 25 GHz. The input bandwidth $\sim$ 25% at observing frequencies below 16 GHz, and 4 GHz above 16 GHz. Each band is to have $\sim 10^5$ spectral channels with a minimum accumulation interval 0.5s. The

images should have at least 10^5 beam areas at the maximum angular resolution. The SKA is to be configured with 20% of the collecting area within 1 km, 50% within 6 km, and 75% within 300 km, with maximum baselines of $\sim$ 3000 km.

In this paper, we explore how these impressive requirements can be met using an array with a large number of small antennas by integrating the calibration and imaging into the data acquisition process. There are several reasons for pursuing this approach.

First, we recognize that the end user is primarily interested in astronomy, and less interested in details of the data acquisition and data processing which tend to preoccupy the radio astronomer specialist. The SKA is a powerful instrument, and will only achieve its full potential if it can easily be used by non radio astronomers. The SKA should produce final, calibrated images as its normal output.

The data processing poses significant problems for arrays of 1000's of antennas. Aperture synthesis imaging in radio astronomy has been developed with arrays of $\sim$ 10 to 100 antennas. In the current data processing paradigm, the cross correlations are formed in on-line, custom designed correlator hardware. The correlations and calibration data are written to datasets which are transferred to off-line computers where the calibrations are applied and images are made using data reduction packages. There is a severe mismatch between the data rates in the on-line correlator hardware and those supported by the off-line processing which can typically handle only a few percent of the data rates large correlators are capable of producing. This can be resolved by integrating the calibration and imaging into the data acquisition process.

At decimeter wavelengths there are problems with non coplanar array geometry and non isoplanicity of the atmosphere. The whole field of view (FoV) must be imaged in order to remove the sidelobes of all the sources within primary beam pattern of the antenna, and a 2D FFT cannot be used with a large FoV. For high dynamic range it will be necessary to subtract strong sources in the sidelobes of the primary beam pattern, which may have poorly determined and time variable antenna gain. Off-line data processing needed to deconvolve the sidelobes of sources in a large field using the current algorithms is very expensive.

Sophisticated image processing algorithms have been developed to self-calibrate the measured cross-correlation function using images of the sky brightness, and to remove sidelobes of the synthesised beam and confusing sources (see e.g. Thomson, Moran and Swenson, 2001 (TMS); Cornwell and Perley, 1992). These algorithms have been very successful, but are time consuming for the end-user and require a level of expertize which many astronomers do not wish to acquire in order to do their science.

The delayed calibration and analysis of the data limit the science which can be done.

Variable sources, targets of opportunity, and RFI are more easily handled as the data are being acquired. In this paper we propose to implement calibration and imaging in close to real time in order to reduce the burden of expert data reduction on the end user, and to make best use of both telescope and human resources.

2. Pointing, Delay, and Phase Centers.

The central thesis of this paper is the application of reconfigurable digital signal processing to independently control pointing, delay and phase centers for wide field imaging. That these three centers are used ambiguously should not be surprising, since they lie at the root of the diversity of radio telescope designs - from an array of detectors sampling the wavefront, to a monolith bringing the wavefront into phase at a single point. We define our usage of these terms:

The pointing center is defined by the peak amplitude in the primary beam pattern - the illumination pattern of the sky brightness distribution by the antennas. A parabolic reflector brings the wavefront into phase at the receiver feed, and the pointing center is determined by the location of the feed and geometry of the antenna. Using an array receiver at the focus of each antenna provides multiple pointing centers. For a phased array telescope, the pointing center is determined by the fixed or adjustable phasing of the individual elements. For an SKA "station beam", the pointing center is determined by the geometry and phasing of arrays of small diameter antennas. The primary beam pattern which results from a cross correlation is the product of the voltage pattern of one antenna and the complex conjugate of the other for each antenna pair, and is complex valued if the two voltage patterns are not identical.

The delay center is the direction in the sky for which the propagation delays of the signals from each antenna are equal, and is determined by the geometry of the array of telescopes, and by cable, fiber, or digital delays from each antenna. Different back end instruments, correlators and beamformers, may have different delay centers.

The phase center is the direction in the sky for which the signals from each antenna have the same phase. Different back ends may have different phase centers. The delay center is the default phase center. However, for a cross correlation receiver, the signals from other directions in the sky can be brought into phase by adjusting the phase of the cross correlation. This, of course, is the power of a correlator and allows images to be made by phasing the correlator outputs over a grid of positions.

This paper carries three interrelated themes:

1) Using reconfigurable backend data processing we can adjust the pointing, delay and phase centers independently and simultaneously for multiple regions in the sky. Multiple delay centers within each pointing center, and multiple phase centers within each delay center enables us to observe multiple targets over a wide field of view simultaneously.

2) Calibration, imaging, and deconvolving the sidelobes of sources outside the fields of interest are intimately related, and are best handled in close to real time, rather than using off-line data processing.

3) High performance digital signal processing enables us to handle high data rates in parallel, and to make images in close to real time.

3. System Architecture and Data Processing Model.

A hybrid solution using beam formation and correlators provides a flexible development path for imaging large fields of view. Phased array beams can be formed anywhere in the sky by adding the voltages in phase from all, or any subset of antennas. The sidelobes structure of each beam depends on the array geometry, the source direction, and the amplitude and phase weighting of the signals from each antenna. Beam formation is appropriate for analyzing signals from discrete radio sources such as pulsars, SETI targets and RFI sources. Beam formation allows us to channel the collecting area of large arrays of telescopes into expensive back end analysis engines. Direct imaging using beam formation is appropriate for compact sources, but is currently too expensive for imaging large fields.

Correlators provide a familiar and versatile mechanism for integrating the data in the uv domain. We have the advantage of over 40 years of development for calibration and imaging using uv data. Large correlators are needed to correlate 1000's of antennas, and the high data rates are increased by forming correlation functions.

In this section we trace the data flow through an imaging system using correlators and beam formers. Figure 1 shows the overall system.

3.1. Digitize and packetize

The total bandwidth of signals from N antennas is $N \times B \times N_{pol}$, where B is the analog bandwidth and N_{pol} the number of polarizations from each antenna.

The data for each antenna are digitized and packetized to exploit commercially available hardware and protocols. The total data bandwidth is $N \times 2B \times N_{pol} \times N_{bits}$, e.g. for $N =$

1000, $B = 1$ GHz, $N_{pol} = 2$, and $N_{bits} = 8$, the total data bandwidth is 4 10^{12} bytes/s.

3.2. Channelize

The input bandwidth must be channelized, to provide spectral resolution, to facilitate interference rejection, and to reduce bandwidth smearing by using multifrequency synthesis (MFS). The science and RFI requirements for a large number of frequency channels favor an FX architecture for the correlators. Excellent separation of frequency channels can be obtained using an 8 tap polyphase filter. After the frequency transform, the total data bandwidth is $N \times 2B/N_{chan} \times N_{chan} \times N_{pol} \times N_{bits}$ where N_{chan} is the number of spectral channels. The data bandwidth may be reduced by sampling the narrow band channels with fewer bits after interference rejection and by selecting appropriate bits in each frequency channel so that RFI can be characterized using post correlation techniques. E.g. using 8-bit sampling of the wide bandwidth analog signal from each antenna, and 4-bit sampling after channelization, the data bandwidth is reduced by a factor 2. After the frequency transform, the data can be processed in parallel, reducing the data rate in each frequency channel by a factor N_{chan}.

3.3. Switches

The data in each packet are uniquely identified and self contained so that they can be routed as desired to multiple asynchronous data processing engines without ancillary information. The data packets should contain the metadata needed to calibrate and make images for multiple regions. Commercial routers can be used to distribute the data to beam formers and correlators. There may be multiple routers in the system to allocate and reuse the hardware for IF processors, beam formers, correlators, and integrators in different projects. Flexible routing allows the data processing hardware to be upgraded, repaired, and reprogrammed with minimum interruption to array operations

4. Beam Formation and Correlation Function Imaging

An array forms a beam by adding up the voltages, V_k, in phase from each antenna, k, and multiplying by the complex conjugate to get the power

$$I(s) =< [\Sigma V_k(r) \exp(2\pi i\ \nu t + 2\pi i/\lambda\ r.s)] \times [\Sigma V_k(r) \exp(2\pi i\ \nu t + 2\pi i/\lambda\ r.s)]^* >,$$

where s is the sky position vector, r is the antenna position vector, and <> denotes a time average. An image can be made by forming beams in a grid of source positions across the region of interest. This direct image formation solves a number of problems for large field imaging, but is computationally expensive (see SKA memo 46).

The above equation can be re-written as the sum of the powers measured at each antenna and the cross products between all antenna pairs:

$$I(s) =< \Sigma V_k(r_k)V_k^*(r_k) > + < \Sigma V_j(r_j)V_k^*(r_k)\exp(2\pi i/\lambda\ (r_j - r_k).s) >$$

Information about the brightness distribution is contained in the cross products, which are the output of a cross correlator. Images are formed from the Fourier transform of the sampled correlation function. (e.g., see TMS and references therein). For full polarization, we must form all 4 cross products of the 2 polarizations from each antenna.

After correlation there are $N(N+1)/2$ auto and crosscorrelations, and $N_{pol} = 4$, polarization products. The required sampling rate is reduced from the input bandwidth, B to the rate of change of the cross correlations; the fringe rate. For a siderial source, the fringe rate is $(r_j - r_k)/\lambda\ .sdot$ Hz, where $sdot$ is the earth rotation rate, 7.27 10^{-5} radian/s. Fringe rates for the ATA and SKA are given in Table 1. Although the required sample rate is proportional to the baseline length, the data are usually sampled at the Nyquist rate for the longest baseline.

In order to correlate the signals from a siderial source anywhere in the entire sky, the data bandwidth from the correlator is:

$$N(N+1)/2 \times N_{pol} \times N_{chan} \times N_{bits} \times 2\ sdot \times D_{max}/\lambda$$

e.g. for $N = 1000$, $N_{chan} = 10^5$, $N_{pol} = 4$, $D_{max} = 1000\ km$, $\lambda = 1\ cm$, and $N_{bits} = 2\times16$ (complex data), the total data bandwidth would be $\sim 10^{16}$ bytes/s. The data rate can be greatly reduced with more appropriate integration and sampling.

4.1. IF processing

The IF processor corrects for the delay and phase at the delay tracking center. The delay tracking center is usually set at, or near to the pointing center. For an array of fixed antennas, such as a dipole array with a wide FoV, the delay and phase tracking center may be the zenith. The signals from each antenna are processed to compensate for the propagation delays introduced by array geometry, atmosphere, and electronics. Signals from sources in different directions have different geometric delays, $r_j.s$, and may have different propagation

delays through the atmosphere (troposphere at short wavelengths and ionosphere at long wavelengths). In order to form a beam, the signals from each antenna must be added in phase. For a correlator with sufficiently narrow frequency channels, no delay and phase tracking may be needed before the signals are correlated. Phase errors, which reduce the forward gain and increase the sidelobes for beam formation, produce a phase shift $2\pi i/\lambda$ $(delay_j - delay_k)$ in the cross correlation of antennas j and k, which can be calibrated. The relative delays and phases of the antennas used for beam formation are conveniently measured using a correlator (see Figure 1). The IF processor may also produce in-phase and quadrature phase inputs for beamformers and correlators, and demodulate phase switching patterns which were used to remove cross-talk between antennas or polarization channels.

4.2. Correlators

The correlation function is computed with narrow frequency channels and short integration times so that images can be formed over a large field of view. Correlations of signals from different directions produce a phase gradient across the bandwidth. In order to correlate signals from sources over a large field of view, the channel width must be much less than c/D_{max} to avoid bandwidth smearing, where c is the velocity of light and D_{max} is the maximum delay within the field of view (e.g. see TMS). In order to cross correlate the signals from any direction, D_{max} is the maximum antenna separation. This limiting value of the channel width is listed in Table 1.

The length of the sampled data stream must be sufficient to define the channel width. E.g. for an 8 tap polyphase filter the channel width must be greater than 8 $\times$ the sample rate. This sets a upper limit on the sample rate, and a lower limit on the channel width. i.e. D_{max}/λ $sdot << \Delta\nu << c/D_{max}$. If the fringe rate, $Fringe << c/D_{max}$, we can correlate signals from anywhere in the siderial sky using one delay center (see Table 1).

Thus, for short baselines, $D_{max} << \sqrt{(}c\ \lambda/sdot)$, we can image sources anywhere in the sky with a spectral line correlator. For large antenna separations at short wavelengths we cannot image the whole field of view with a single correlator, but must use different delay centers in restricted fields of view. A large FoV can be imaged using multiple delay tracking centers and correlators.

4.3. Integration at Multiple Phase Centers

Sampling the correlator at the fringe rate allows us to make images over a wide field of view, including targets of interest, calibration sources, and sources whose sidelobes confuse the regions of interest. We can form simultaneous images in multiple regions within the field of view by integrating the output from the correlators at multiple phase centers. The data from each correlator is multiplied by a phase factor, $\exp(2\pi i/\lambda\ r.s_o)$, where $r = (r_j - r_k)$ is the baseline vector, and s_o is the phase center in each region of interest.

Within a restricted field of view, the required sample rate is reduced to the range of fringe rates within the FoV. The regions which can be imaged are limited by the primary beam, the isoplanatic patch sizes, and limits to the image size imposed by non coplanar baselines, $\sim\ sqrt(\lambda/D_{max})$. The primary beam $FWHM\ \sim\ \lambda/D_{ant}$, can be imaged by making $N_f \sim \lambda D_{max}/D_{ant}^2$ tangent plane images using a 2D FFT, or N_f image planes using a 3D FFT. (references in TMS). Table 1 lists λ/D_{ant}, $sqrt(\lambda/D_{max})$, N_f, and the Nyquist sample rate within the primary beam $FWHM$ at the longest baseline, 2 D_{max}/D_{ant} $sdot$ Hz. e.g. for the ATA, with a 6m diameter antenna and a maximum antenna separation 1 km, the FoV is defined by the primary beam $FWHM$, $\sim$ 17 arcmin at $\lambda = 3$ cm, and by non coplanar baselines at $\lambda = 1$ m. The 10 degree primary beam at $\lambda = 1$ m can be imaged with $\sim$ 28 subfields if these are each within a single isoplanatic patch. For the SKA, with 12 m diameter antennas at $\lambda = 1$ cm, the primary beam $FWHM$ is only $\sim$ 3 arcmin, and mosaics at multiple pointing and phase centers, array receivers, or on-the fly mapping are required to image larger regions.

Phase tracking and integrating the data after the correlation requires complex multiplies by a phase factor for each antenna pair and frequency channel, but the data rate is reduced to the range of fringe rates within the restricted field of view. The data bandwidth for imaging the full primary beam $FWHM$ is:

$N(N-1)/2 \times N_{pol} \times N_{chan} \times N_{bits} \times 2\ sdot \times D_{max}/D_{ant}$

e.g. for $N = 1000$, $N_{chan} = 10^5$, $N_{pol} = 4$, $D_{ant} = 12\ m$, and $N_{bits} = 2 \times 16$ (complex data), the total data bandwidth is 2 10^{10} bytes/s for 1 km baselines, and 2 10^{13} bytes/s for 1000 km baselines.

4.4. Calibration

An a-priori model of the sky brightness distribution is used for calibration and imaging. In the standard observing paradigm, strong compact sources are used as primary calibrators,

and self-calibration is used to improve the calibration during the off-line imaging process.

The model visibility is calculated as

$V'_{j,k} = exp(2\pi i/\lambda\ r.s_0) \times \Sigma(I \times A \times B \times P \times G \times \exp(2\pi i/\lambda\ r.(s - s_o),$

where $I(s, \nu, p)$ is the model image, $A(s, \nu, p)$ is the primary beam response, $B(\nu)$ is the instrumental bandpass, $P(s, \nu, p)$ is the polarization calibration, and $G(time, s_0)$ is the gain. $r = (r_j - r_k)$ is the baseline vector, and s, ν, and p are the position, frequency and polarization. s_o is the phase center for each region of interest.

Each of the calibrations, A, B, P, and G are complex valued functions, and can be decomposed into antenna dependent components. The primary beam response is the product of the antenna voltage patterns for each correlation and will modulate the phase as well as the amplitude of the primary beam response to sources away from the pointing center. For a phased array station beam, atmospheric fluctuations make the primary beam response time variable. Even for a clean voltage pattern with low level sidelobes from a single antenna, the complex sidelobe pattern will vary with time due to pointing errors. Over a wide field of view the atmospheric calibration will vary due to non isoplanicity.

Our approach to these problems is to separately calibrate the data for each phase center. We can identify regions which have bright emission from a-priori images of the sky brightness, and image only regions which are of interest or contain sources whose sidelobes corrupt the regions of interest. Confusing sources may be in the sidelobes of the primary beam, or in different isoplanatic regions.

The data from the correlators are integrated at each phase center by multiplying by the phase factor $\exp(2\pi i/\lambda\ r.s_o)$. The slowly varying parts of A, B, P, and G can be characterized, and applied to the model. After correcting the model by these best estimates, the model is used in a self calibration algorithm to determine the antenna gains $g(t, s_0)$, as a function of time for each phase center which contains suitable sources. These antenna gains include time variations and errors in the measured values of A, B, P, and G as well as the atmospheric amplitude and phase fluctuations at each phase center.

The gains are not totally independent for each phase center. The SNR of the gains can be improved by developing a global model of the gain variations across the array as a function of time and frequency. When the source model is a spectral line source, multiple frequency channels are used simultaneously to determine the antenna gains using a 3D image model. Observations at multiple frequencies can be used to separate the gains into tropospheric and ionospheric delays. The data streams must be delayed by a calibration interval so that the gains can be averaged and interpolated before being applied to the data stream.

The self calibration algorithm proceeds by minimizing $\chi^2 = < \Sigma[V \times g_i g_j - V']^2 \ / \ \sigma_{ij}^2 >$, where the summation is over all antenna pairs (i, j), frequency channels with known source structure, and time intervals in which the antenna gains are to be defined. χ^2 can be accumulated in distributed processors associated with each correlation engine, and passed to a calibration solver which computes the gains g_i for each time interval and phase center. These antenna gains can be checked for consistency with the expected temporal and spatial variations of the global model, and used to calibrate the data streams into beam formers and imagers.

4.5. Sidelobe subtraction

Sidelobes of sources outside the regions of interest confuse the images, and must be subtracted to achieve high fidelity. The usual approach pioneered and developed by Cornwell and Perley (1992), is to subtract the confusing sources when deconvolving the images. The algorithm proceeds iteratively, deconvolving tangent plane subfields, and re-computing the image when the residual noise in the deconvolved subfields is comparable to the sidelobe level from sources outside each subfield.

Perley and Clark (2003) have developed scaling relations for the image processing requirements for large field imaging. For imaging the full FoV, they show that the processing cost scales as $\sim \lambda^{1.7} N_{sta} D_{max}^3 / D_{ant}^6$, where N_{sta} is the number of antennas per station. If the antennas are clustered in stations the effective uv coverage is reduced, the synthesised beam has higher sidelobe levels, and the tangent plane images must be re-computed more often. For an array with a constant total collecting area, the cost scales as $\sim \lambda^{1.7} N_{sta} (N \ D_{max})^3$. Cornwell (2004), taking into account the effect of non-coplanar baselines, derives a data processing cost equation which scales as $\sim \lambda N^2 D_{max}^3 / D_{ant}^4$, which for an array with a constant total collecting area, scales as $\sim \lambda D_{max}^3 / D_{ant}^8$. As noted by Perley and Clark, these scaling relations only apply to full field of view imaging at the highest resolution. In many cases, the SKA will image regions larger than λ / D_{ant} at lower resolution, or regions smaller than λ / D_{ant} at the highest resolution, and these scaling relations will not apply. (Lonsdale, Doeleman, & Oberoi, 2004; Wright, 2004).

In any case, extrapolating the existing deconvolution algorithms for off-line data processing of large fields of view is too expensive for a large N array. Building a smaller number of larger antennas does not solve the problem for the SKA since sources larger than the primary beamwidth must be mosaiced, and the image fidelity is degraded (see for example, Cornwell, Holdaway, & Uson, 1993).

In this paper, we propose to remove the sidelobes of confusing sources by subtracting the source model from the calibrated *uv* data stream. The subtraction can be made for each region of interest and frequency channel in distributed processors associated with each correlation engine, but includes the response from the whole sky model, especially of course the strong sources. In order to subtract the a-priori model from the *uv* data stream we must determine the instrumental response across the whole sky. This calibration includes the primary beam response $A(s, \nu, p)$, the instrumental bandpass $B(\nu)$, polarization calibration $P(s, \nu, p)$, and the gain $G(time, s_0)$, where s, ν, and p are the position, frequency and polarization, and s_0 is the phase tracking center.

4.6. RFI subtraction

RFI must be subtracted from the *uv* data stream before it is passed to the imaging engine and beam formers. RFI presents a special case in several ways. RFI sources may be stationary, or moving across the sky at a non-siderial rate. A correlator can be used to locate and characterize RFI as a function of time, frequency and polarization. The SNR can be improved by pointing some of the antennas or beam formers at the RFI sources. Correlators allocated to measuring RFI may need to sample the signal at high data rates.

4.7. Imaging Engine

Images in different frequency channels are processed in parallel in a distributed architecture. The images are formed from the calibrated *uv* data stream from which an a-priori source model has been subtracted, and are therefore difference images from the model image. These difference images are used to update the model, including not only the regions of interest, but also improving the accuracy of the sources whose sidelobes are being subtracted. As the observations proceed, both the model image and the calibration are improved. The process converges when the difference images approach the noise level and the model image is consistent with the *uv* data streams. Indeed, this criteria could be used to define when the observations are complete.

For a small field of view a 2D FFT can be used to image the region around each phase center. The maximum image size for a 2D FFT scales as D_{max}/λ, $\sim 10^8$ beam areas on a 1000 km baseline at λ 1 cm. Deconvolution is minimized by the excellent *uv* coverage and low sidelobe level of the large N array design. e.g. for the ATA with 350 antennas, the sidelobe level $\sim$ 0.3 %. In many cases, deconvolution in the image plane may not be

needed at all, since the model image and sidelobes of confusing sources have been subtracted from the *uv* data. In addition, the difference images may be limited by atmospheric noise or instrumental errors which must be removed from the *uv* data and can not be removed by deconvolving the synthesized beam.

The imaging engine must also make images using all the frequency channels. Spectral line images can be made for individual frequency channels, averaged into the desired frequency or velocity intervals. Wideband, MFS imaging requires averaging over frequency channels. The a-priori model used in the calibration is updated at intervals when the difference from the best current model is significant.

Variable sources are detected as intermittent sources which are inconsistent with the current model. We should also accumulate a χ^2 image to help identify pixels where there are time variable sources or RFI sources. In some cases we may want to keep a time series of difference images and the model images used for the calibration.

We view imaging as a dynamic process which can be guided in real time by observers inspecting the convergence of the model image and the χ^2 image. As the observations proceed, the phase centers can be moved to image regions where more data are needed to define the science goals, either regions of interest, or sources whose sidelobes are confusing, or new sources which are discovered in the imaging process. Isoplanatic patches may vary during the observations requiring more, or fewer phase centers to adequately determine the calibration across the FoV.

4.8. Archive

The archive serves as the data base for the observations, calibrations and instrument status during the observations. The *uv* data streams from each phase center are saved in the archive along with the metadata. Since earlier data has not been calibrated, and sidelobes have not been subtracted using the final model image, the images can be improved by replaying the *uv* data, repeating the calibration using the improved model. Thus, we envision that the *uv* data stream should be saved, at least until the final calibrated images have been made.

The *uv* data streams from the archive can be replayed though the imaging system or a copy of the same system. The best available model of the sky and phase screen across the array resulting from the completed observations can be used to improve the calibration - especially of the early data which was acquired with the initial a-priori model.

5. Proof of concept

This paper presents a somewhat radical departure from the conventional model of off-line data reduction. We made some tests to offer some assurance - to the author, as well as to readers, that it might possibly work. In this section we describe some off-line tests using model data simulations and some real data from a 4-antenna dipole array at 175 MHz. The off-line procedures are simple **csh** scripts which control MIRIAD tasks. The results from four scripts are presented here.

5.1. Change phase tracking center.

The first script is a test of processing the *uv* data at multiple phase centers. We added new code to the **uvcal** task to move the phase center around in a large field of view as required for integrating the data at multiple phase centers.

The script uses the existing MIRIAD tasks **uvgen** and **uvmodel** to generate uv-data for a point source at the phase center. We move the phase center to other positions in a large field of view and add a point source at each phase center. We then move the phase center around and subtract each source from the *uv* data, and inspect the results in the *uv* and in the image plane using a standard grid and 2D FFT algorithm at each phase center.

At each phase center, the J2000 RA, DEC position is precessed to the current epoch, and the spectral line *uv* data is phase rotated to the new position. The phase rotation is performed after any calibrations derived for the previous phase center have been applied.

This process worked as expected leaving residuals at a level consistent with the floating point arithmetic. The source subtraction is perfect for source subtraction at the phase center, and degrades slowly with added noise in a field of view $\sim sqrt(\lambda/D_{max})$. This script is really just a test of the software.

5.2. Source models at 2 different phase centers.

The second script is a test of simultaneously deconvolving multiple regions of interest (or removing sidelobes from confusing sources, depending on one's point of view).

We generated source models, distributed within a region $\sim sqrt(\lambda/D_{max})$ around 2 phase centers at the positions of Cygnus A and Cas A. We verified that we can move the phase center and subtract both sources from the *uv* data. The wide field image including

both sources is distorted by a 2D FFT due to non coplanar baselines, but each source is correctly imaged when moved to the phase center.

The image fidelity depends on the complexity of the source distributions, the signal to noise ratio and the *uv* coverage. At each phase center we deconvolved the images within a FoV $\sim sqrt(\lambda/D_{max})$ (see Table 1). The speed and quality of the deconvolution depends on the complexity of the model image and the *uv* sampling. We used a source model with up to 10 Gaussian components, and different levels of added thermal noise. We tried imaging with a 4-antenna dipole array, 15-antenna CARMA configuration, and the ATA with 42, 98 and 350 antennas.

In the off-line simulation the multiple regions are handled sequentially. Each source was moved to the phase center, imaged and deconvolved, and subtracted from the *uv* data. A few iterations suffice to remove the strongest source components down to the sidelobe level of the synthesised beam which goes approximately as $1/N$, ~ 0.3 % for ATA-350.

The simulations show that we can deconvolve multiple regions by subtracting source models from the *uv* data. Our goal is to develop efficient data handling of *uv* data streams to calibrate and image multiple regions of interest over a wide field of view. Deconvolution and subtracting a source model can be made using existing software in floating point processors.

The off-line data processing is fast enough to handle quite large datasets, processing spectral line data at 10-20,000 *uv* data points per second (1-3 Mbytes/s) on a modest sized desktop (see figures 2-4 in SKA memo 46). Source subtraction from the *uv* data can be made for each region of interest and frequency channel in distributed processors associated with each correlator and phase center integration.

5.3. Epoch of reionization experiment

The third script used real data from a 4-antenna dipole array which is a prototype for a large N array to study the epoch of reionization (Backer and Bradley, 2005)

Data were obtained using a prototype FX correlator for the ATA. An RF bandwidth 150-200 MHz was sampled in 513 spectral channels at 0.1 Hz, sufficient to image sources within the entire field of view (see Table 1). The amplitude was calibrated by dividing the *uv* data by the square root of the autocorrelations for each pair of antennas. The *uv* data were phase rotated to the positions of the strongest sources, Cygnus A and Cas A, and instrumental delays and antenna positions fitted over a limited HA range. The data at each phase center were self-calibrated using a 30 min time interval to remove slow instrumental

and ionospheric phase drifts. Figures 2 and 3 show our first images of Cygnus A and Cas A. No correction for the antenna primary beam response or polarization has been applied, and residual sidelobes of other sources are clearly present. The measured peak to RMS ~ 100, and the differences of spectral channels are at a level $\sim 10^{-3}$ which gives us some hope of achieving much higher image fidelity when more calibrations are applied to the data.

5.4. Extended source structures

Sources larger than $\sim \lambda/D_{ant}$ must use mosaics of multiple pointing centers. From the perspectives of this paper, each pointing center is also a phase tracking center, but need not be a delay tracking center if $\Delta\nu << c/D_{max}$.

Imaging with the ATA has been simulated for a range of source sizes and declinations using the MIRIAD software. Figure 4 shows a mosaic MFS image of Cas A scaled up 40 times the real size, imaged with the ATA with 350 antennas. We generated *uv* data for a snaphot observation using a VLA image of Cas A as source model. We used a mosaic observation in a hexagonal pointing pattern sampled at the Nyquist interval $\lambda/(2\ D_{ant})$ at an observing frequency 1.42 GHz. A 1 min integration at each of 7 pointing positions generated 1221500 *uv* data samples with 20 frequency channels. Thermal noise appropriate for the ATA was added to the *uv* data. The *uv* data was Fourier transformed to make images and deconvolved using Maximum Entropy (MEM) algorithms. The original image model was convolved to the same resolution by a Gaussian restoring beam and subtracted from the deconvolved image. The residual imaging errors are characterized by the total recovered flux, peak flux density and the RMS residuals. The results are pipelined into tables. The mosaicing process, the simulated images and the residual images are displayed on the terminal, and the images are saved on disk for further analysis. For more details see ATA memo 52.

Generating the *uv* data for a 7-pointing mosaic takes $\sim$ 60s. Fourier transform to a 1585 x 1510 mosaic MFS image and 2049 x 2049 x 7-pointing beam set takes only $\sim$ 47s, but the Maximum Entropy (MEM) deconvolution with 200 iterations takes over 1 hour to converge. Fortunately, because of the low sidelobe level for the ATA synthesised beam, such deconvolution will only be needed for bright sources. Most images of low brightness emission will be limited by thermal noise and will not require deconvolution. The performance of MIRIAD running on a modern PC appears adequate for the calibration and imaging to keep up with the data acquisition for the ATA, but the MEM deconvolution does not. In this paper, we propose to make difference images from which the sidelobes of confusing sources have been subtracted from the *uv* data, and no deconvolution is needed in the image plane.

6. Data Structures

The tests described above are not comprehensive, but most aspects of our proposed on-line imaging can be tested and developed using existing off-line software with simulated model data and real data from the ATA and EoR telescopes. The MIRIAD *uv* data structure which we are using for the simulations is quite close to the block floating format used in the hardware. The *uv* data consists of a stream of spectral line data stored as scaled 16-bit integers. The metadata are stored as a stream of named variables and values. Source names, frequencies, pointing centers, and phase tracking centers are variables which can change throughout the *uv* data stream.

MIRIAD calibration tasks produce or use calibration tables and parameters which are stored in the data structure. The history of observation and data reduction, including the steps and parameters used in observing and reducing the data are stored in the data structure. The data structure is quite flexible. Other entities, such as WVR data, and a copy of the observing script and parameters from Hat Creek observations have been easily added when needed. The ATA correlator directly writes MIRIAD *uv* data.

7. Hardware prototype

We plan to build a hardware prototype imager which can be tested on the ATA and other wide field telescope arrays. We are working with the Berkeley Wireless Research Center (BWRC) to develop a multipurpose computing platform for radio telescope DSP applications (Brodersen etal, 2004). Our interest is motivated by the very fast growth rate of DSP using FPGA devices which enables a wide range of radio astronomy applications. The growth rate far outstrips our ability to implement the latest technology in a cost effective and timely manner.

Currently, radio astronomers use a specialized instrument design with separate printed circuit boards for each subsystem, with dedicated functionality, custom interconnects, memory interfaces, clock distributions and packaging, and programming. The design cycle typically has taken 5 to 10 years from concept to production usage, so by the time the instruments are put in the field, they are many years behind the current capabilities.

The BWRC is a partnership of UC Berkeley and industry associates. One of the current developments is the BEE project to develop a multipurpose DSP platform using FPGA technology. Typical applications are to analyze and process signals from analog devices, which includes radio telescopes. Four BEE1 processing units are currently in operation. The development time for different applications has ranged from 1 to 6 FTE months, far less

that the many FTE years which radio astronomers have experienced with custom built DSP devices.

Now in its 2nd generation, the BEE2 board is designed to handle 250 10^9 CMAC/s (complex multiply-add per sec), with memory bandwidth 12 GB/s and 360 Gb/s I/O bandwidth. The board is well suited to a wide range of radio astronomy DSP applications including spectrometers, wideband correlators, and arrays with large numbers of antennas or receivers.

The system design philosophy treats BEE2 boards as a modular DSP resource with a flexible interconnect architecture which allows reconfiguration of the computing resources for multiple applications. The programming model uses a system generation library with BEE2 specific hardware abstractions which allow the application programmer to focus on the application rather than the details of the hardware. The system design and programming model together allow the application software to survive by using a technology independent design flow.

With the BWRC, we have developed a unified DSP architecture for radio astronomy applications which enables multiple, reconfigurable, backend data processing, including correlators and beam formers. The BEE2 design is complete and 2 prototype boards have been tested and used to demonstrate a 256 million channel spectrometer on the ATA at Hat Creek. The project timetable includes manufacturing 10 more BEE2 boards, and building a 4000 channel, 32-antenna correlator in 2005, which can be used on ATA, CARMA, and Epoch of Reionization experiments.

Other radio astronomy applications of BEE2 include image formation, RFI mitigation, and adaptive optics for large antennas and interferometer arrays. Digital images can be made in close to real time at the telescope. The phase across the wavefront can be digitally adjusted to correct for atmospheric turbulence or RFI. In many ways, the BEE project provides a programmable "digital lens" for radio astronomy.

8. Conclusions

Imaging a wide field of view with arrays of large numbers of antennas presents a number of interesting problems. Sources of non thermal noise may limit the capabilities of the SKA. For high dynamic range, sidelobes of sources outside the regions of interest must be subtracted or the SKA will be confusion limited. In order to achieve the sensitivity and image fidelity goals, the calibration needs to match the thermal noise.

Calibration and imaging are intimately related. We use a model of the sky brightness distribution to determine the instrumental and atmospheric calibrations. A model is used to subtract sidelobes from sources outside the regions of interest. The calibration varies across the sky due to uncertainties and time variations in the primary beam, instrumental polarization and non-isoplanicity. We must measure the gain variations in the directions of sources whose sidelobes corrupt the regions of interest. The calibration can be improved by using a global model of the sky brightness and gains. RFI and time variable sources must be identified and measured so they can be correctly separated and subtracted from the uv data.

Extrapolating the existing deconvolution algorithms for off-line data processing of large fields of view is very expensive for a large N array. Building a smaller number of larger antennas is not a good solution for the SKA since sources larger than the primary beamwidth must be mosaiced, and the image fidelity is degraded by pointing errors and by the sparse uv coverage.

Off-line data processing is time consuming and requires a level of expertize which many astronomers do not want to learn in order to do their science using radio telescopes. Delayed calibration and analysis of the data also limit the science which can be done. Variable sources, targets of opportunity, and RFI are more easily handled as the data are being acquired. The large-N design for the SKA does not fit well within the constraints of the current paradigm using custom designed correlator hardware and off-line data reduction.

In this paper we propose to implement calibration and imaging in close to real time using FX correlators and beam formers with a high data bandwidth into computer clusters. The mismatch between the data rates in the on-line correlator hardware and those supported by off-line processing is resolved by integrating the calibration and imaging with the data acquisition process. Calibration and imaging are handled with the real time feedback of the antenna gains needed for beam formers and RFI suppression.

Images can be made simultaneously for multiple regions within the field of view by integrating the output from the correlators at multiple phase centers centered on targets of interest, calibration sources, and sources whose sidelobes confuse the regions of interest. The

regions imaged are used to update and improve the a-priori model, which becomes the final calibrated image by the time the observations are complete.

The data processing requirements for an array with ~ 4000 antennas can be satisfied, but the the development time for custom designs is very long and lags significantly behind Moore's law. We propose to use modular DSP boards with a flexible interconnect architecture which allows reconfiguration of computing resources for IF processors, beam formers, correlators, and integrators for multiple projects. The system design and programming model together allow radio astronomy to implement the latest technology in a cost effective and timely manner.

9. Acknowledgements

I would like to thank the many colleagues and students who have made comments, or asked difficult questions during the course of this work. In particular, I appreciate the students who participated in my imaging seminar, Don Backer for the opportunity to work on real data and develop algorithms for the EoR array, Dan Werthimer for his enthusiasm and introducing me to the BWRC, Udaya Shankar for a careful reading of the manuscript, Gerry Harp for turning my pencil drawing into figure 1, and Chen Chang and the BEE group at BWRC for designing hardware and software tools which enable us to build reconfigurable data processing for radio astronomy.

10. References

Backer and Bradley, 2005, NRAO newsletter January 2005, no. 102, p27.

Brodersen, B., Chang, C., Wawrznek, J., Werthimer, D., & Wright, M., 2004 "BEE2: A Multi-Purpose Computing Platform for Radio Telescope Signal Processing Applications" http://bwrc.eecs.berkeley.edu/Research/BEE/BEE2/presentations/BEE2_ska2004_poster.pdf

Cornwell, T.J. & Perley, R.A., 1992, "Radio-Interferometric Imaging of Very Large Fields", A&A 261, 353

Cornwell, T.J., Holdaway, M.A. & Uson, J.M., 1993, A&A 271, 697, "Radio-interferometric imaging of very large objects: implications for array design",

Cornwell, T.J., Golap, K., & Bhatnagar, S., 2003, "W projection: a new algorithm for non-coplanar baselines", EVLA memo 67

Cornwell, T.J., 2004, "EVLA and SKA computing costs for wide field imaging (Revised)" EVLA memo 77

Jones, D.L., 2003, "SKA Science Requirements", SKA memo 45, version 6, 1 December 2003 (D.L.Jones, 16 December 2003, ISSC11-9.4)

Lonsdale, C.J., Doeleman, S.S., & Oberoi, D., 2004, "Imaging Strategies and Postprocessing Computing Costs for Large-N SKA Designs", SKA memo 54, July 2004.

Perley, Rick, & Clark, Barry, 2003, "Scaling Relations for Interferometric Post-Processing", EVLA memo 63

Thompson, A. R., Moran, J. M. & Swenson, G. W., 2001, "Interferometry and synthesis in radio astronomy", 2nd ed. New York : Wiley, 2001.

Wright, M.C.H., 2002, "A model for the SKA", SKA memo 16, March 2002.

Wright, M.C.H., 2002b, "Allen Telescope Array Imaging", BIMA memo 92; ATA memo 52.

Wright, M.C.H., 2004, "SKA Imaging", SKA memo 46, Feb 2004.

Table 1: Limits on the Field of View and Sampling Rates for the ATA and SKA. D_{max} is the maximum antenna separation. D_{ant} is the antenna diameter. λ/D_{ant} is the primary beam FWHM. $sqrt(\lambda/D_{max})$ is the limit to the image size imposed by non-coplanar baselines. $N_f \sim \lambda D_{max}/D_{ant}^2$ is the number of image regions needed, using a 2D FFT, to cover the primary beam $FWHM$. c/D_{max} is an upper limit to the channel width to avoid bandwidth smearing on the longest baseline without delay tracking. $Fringe$ is the maximum rate of change of the cross correlation, and $Nyquist$ is the sample rate within the primary beam $FWHM$.

λ [m]	D_{max} [km]	D_{ant} [m]	λ/D_{max} [arcsec]	λ/D_{ant} [arcmin]	$sqrt(\lambda/D_{max})$ [arcmin]	N_f	c/D_{max} [kHz]	$Fringe$ [Hz]	$Nyquist$ [Hz]
1.00	1	6	206.26	572.9	108.7	27.8	300	0.07	0.02
0.21	1	6	43.32	120.3	49.8	5.8	300	0.35	0.02
0.03	1	6	6.19	17.2	18.8	0.8	300	2.42	0.02
1.00	1	12	206.26	286.5	108.7	6.9	300	0.07	0.01
0.10	1	12	20.63	28.6	34.4	0.7	300	0.73	0.01
0.01	1	12	2.06	2.9	10.9	0.1	300	7.27	0.01
1.00	10	12	20.63	286.5	34.4	69.4	30	0.73	0.12
0.10	10	12	2.06	28.6	10.9	6.9	30	7.27	0.12
0.01	10	12	0.21	2.9	3.4	0.7	30	72.7	0.12
1.00	100	12	2.06	286.5	10.9	694.4	3	7.27	1.21
0.10	100	12	0.21	28.6	3.4	69.4	3	72.7	1.21
0.01	100	12	0.02	2.9	1.1	6.9	3	727.	1.21
1.00	1000	12	0.206	286.5	3.4	6944.	0.3	72.7	12.12
0.10	1000	12	0.021	28.6	1.1	694.4	0.3	727.	12.12
0.01	1000	12	0.002	2.9	0.3	69.4	0.3	7270.	12.12
1.00	1000	25	0.206	137.5	3.4	1600.	0.3	72.7	5.82
0.10	1000	25	0.021	13.8	1.1	160.	0.3	727.	5.82
0.01	1000	25	0.002	1.4	0.3	16.	0.3	7270.	5.82

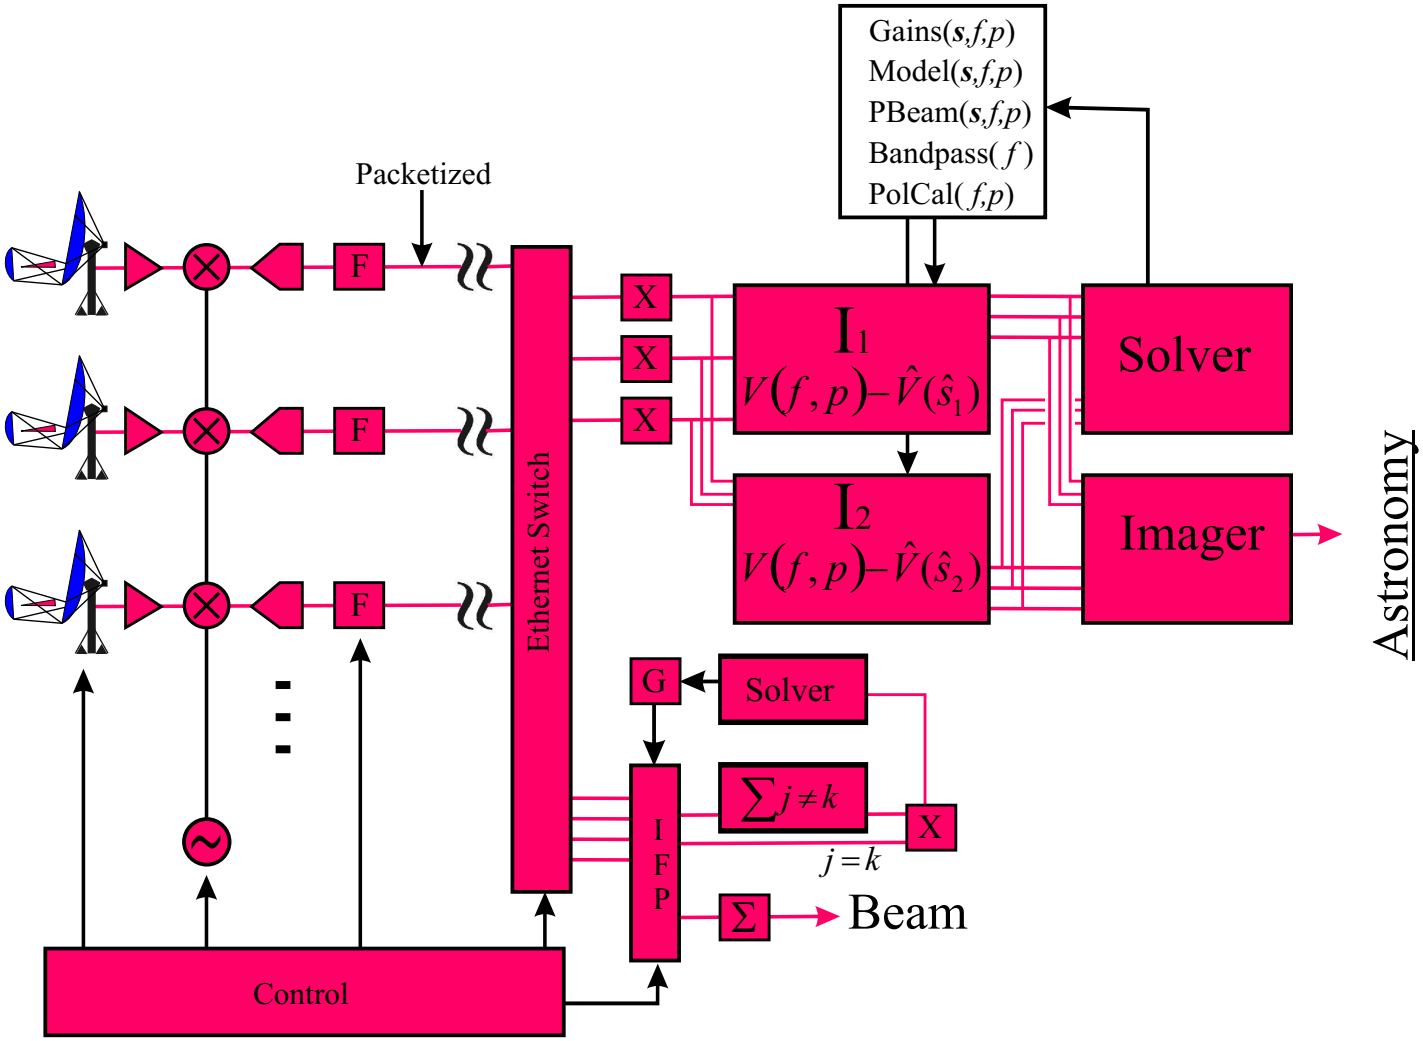

Fig. 1.— Data flow from telescopes to images. See text

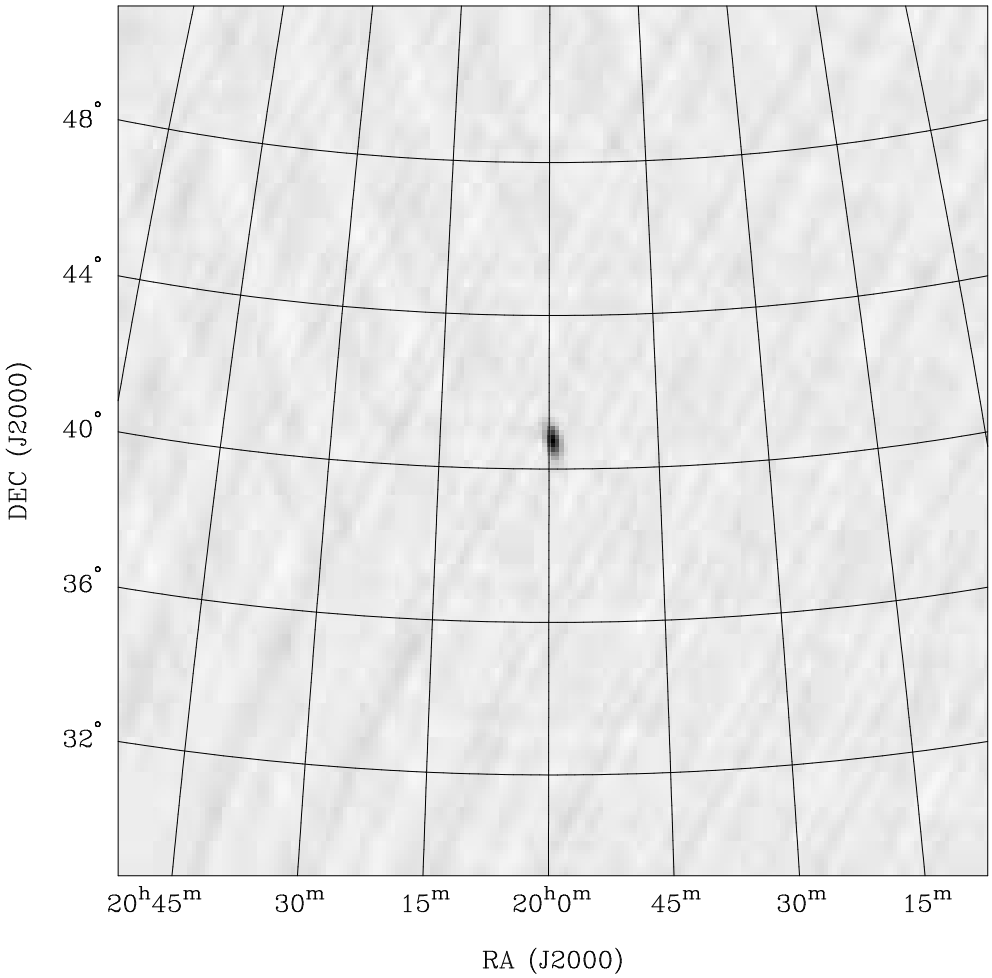

Fig. 2.— Cygnus Image at 175 MHz from 4-dipole array

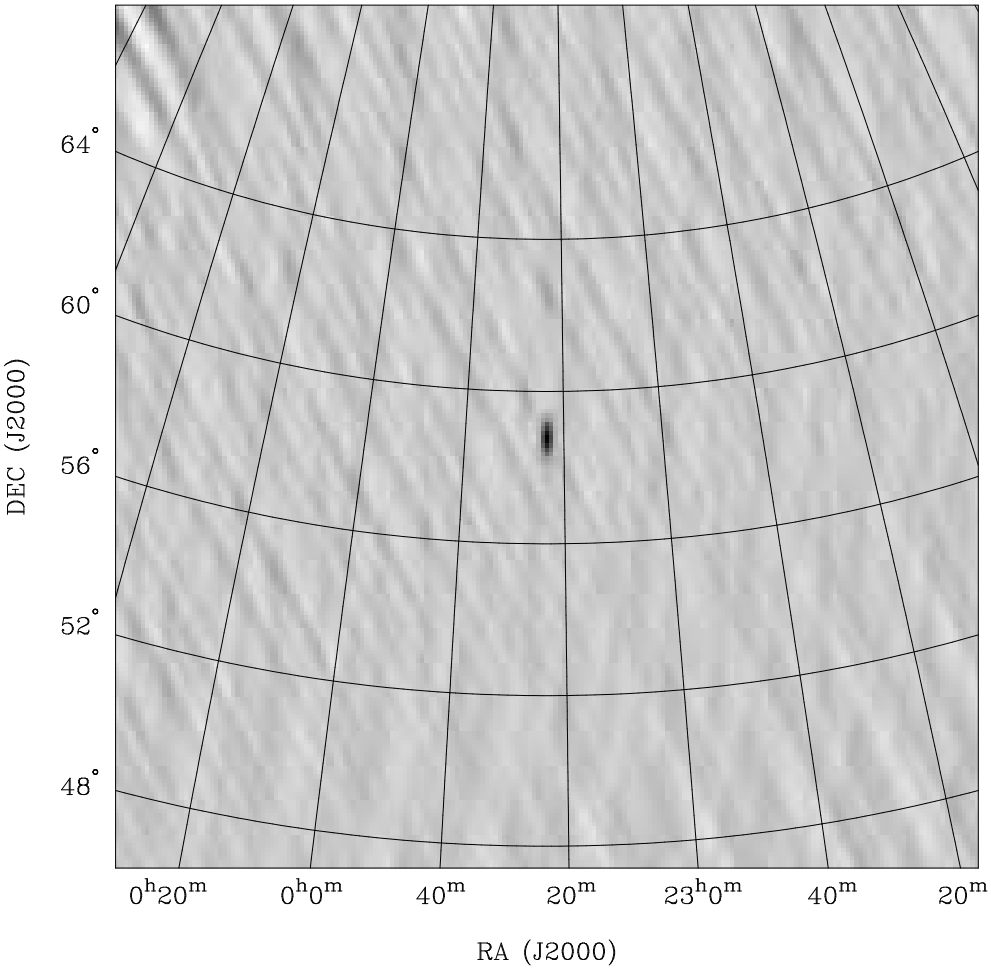

Fig. 3.— Cas A Image at 175 MHz from 4-dipole array

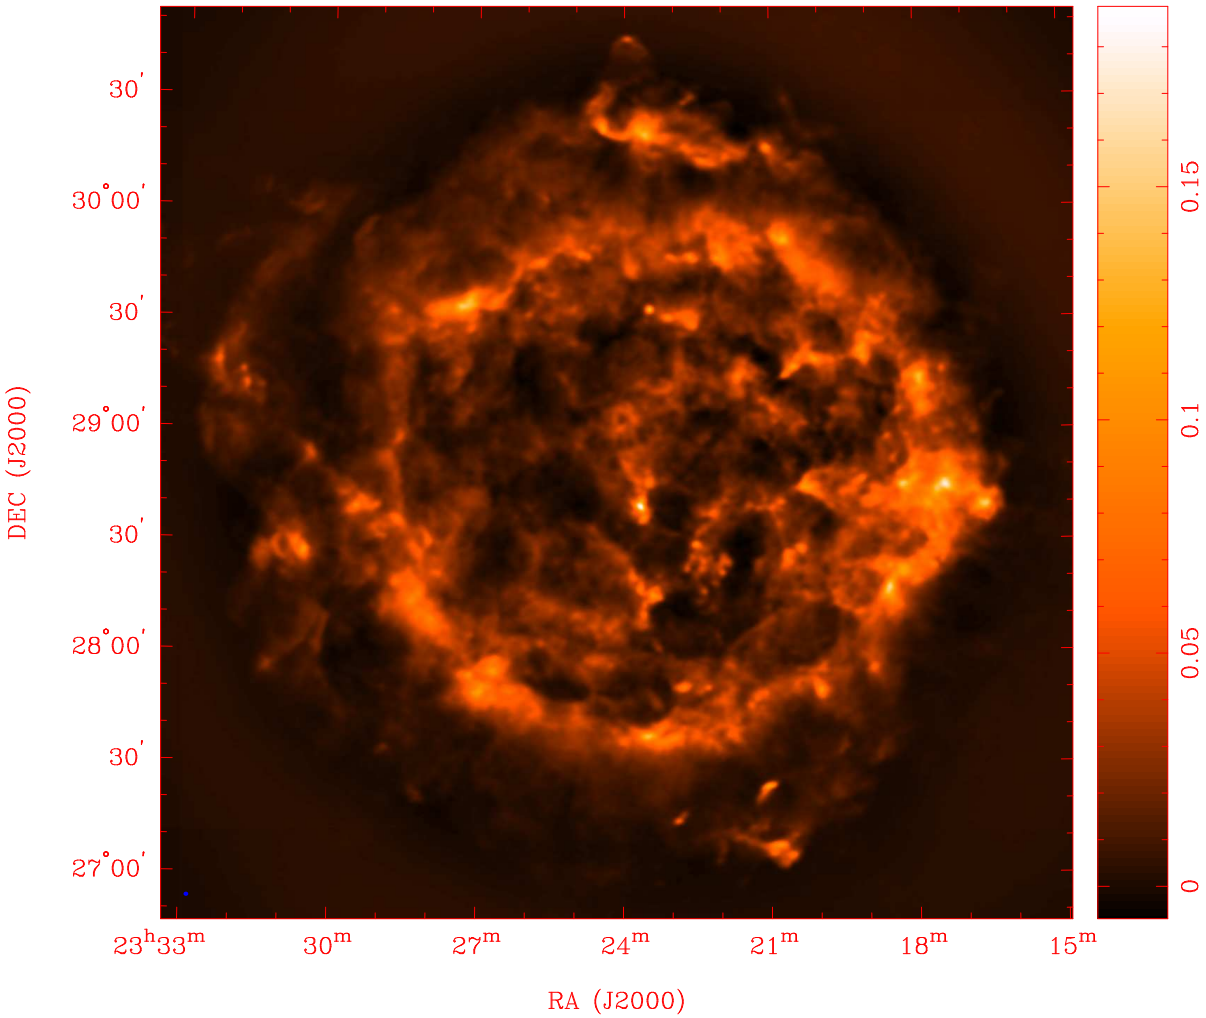

Fig. 4.— A mosaic MFS Image of CASA scaled up 40 times real size, imaged with the ATA at 1.42 GHz. The image is 4 degrees on a side. The synthesised beam FWHM, 77 x 78 arcsec is shown in the lower left corner.