

FLAGCAL: A flagging and calibration package for radio interferometric data

Jayanti Prasad · Jayaram Chengalur

Received: date / Accepted: date

Abstract We describe a flagging and calibration pipeline intended for making quick look images from GMRT data. The package identifies and flags corrupted visibilities, computes calibration solutions and interpolates these onto the target source. These flagged calibrated visibilities can be directly imaged using any standard imaging package. The pipeline is written in “C” with the most compute intensive algorithms being parallelized using OpenMP.

Keywords Flagging · Calibration · Synthesis imaging

1 Introduction

Radio interferometric data taken at the GMRT has traditionally been analyzed interactively using the AIPS data package (but see also Sirothia (2009)). This becomes cumbersome when large data sets need to be analyzed, or when the data analysis has to be done in quasi real time. We describe here a “C” based program which calibrates GMRT data as well as flags data affected by interference or by instrumental problems. This package was developed largely in the context of two ongoing programs at the GMRT viz. (a) a search for transient radio sources for which quasi real time data analysis is needed and (b) a search for HI emission at high redshifts for which large volumes of data have to be analyzed. The package is designed to give a quick look at the data for the first program (in order to determine whether a given burst of radio emission can be localized in the sky or not) and to do the first pass of flagging and calibration for the second program. However, since it operates on files in

Jayanti Prasad
IUCAA, Postbag 4, Ganeshkhind, Pune University Campus, Pune 411007, India
E-mail: jayanti@iucaa.ernet.in

Jayaram Chengalur
NCRA-TIFR, Postbag3, Ganeshkhind, Pune University Campus, Pune 411007, India E-mail: chengalur@ncra.tifr.res.in

a standard format (viz. FITS) and is relatively flexible, it is expected to be useful for a larger range of problems than it was specifically designed for.

Radio interferometers measure “visibilities”, i.e. the Fourier transform of the sky brightness distribution. Descriptions of the processes required to convert this information into an image of the sky can be found in Thompson et al. (2001); Taylor et al. (1999); Chengalur et al. (2003). Here we focus only on the initial stages of this process, viz. that of identifying and flagging corrupted data, and subsequently determining and correcting for the antenna based complex gains.

Data at the GMRT could be corrupted for a variety of reasons, e.g. instrumental failure, radio frequency interference, ionospheric scintillations etc. A wide range of interfering signals could be present, including interference that is narrow band and persistent (e.g. from local digital equipment), interference that is broad band but bursty in time (e.g. from satellites, aircraft radar etc.) low level broad band interference that is persistent in time (e.g. from power transmission lines). A diverse range of approaches have been taken for identifying interfering signals in radio astronomical data (see e.g. Chengalur (1996); Kanekar & Chengalur (1997); Urvashi (2003); Briggs & Kocz (2005); Kocz et al. (2010); Offringa et al. (2010); Paciga et al. (2011)). Here we identify corrupted data by assuming that the true visibilities should be continuous in the time-frequency plane. We identify discontinuities using robust estimators of the underlying statistics of the visibilities. This approach is well suited to finding RFI that is localised in time or frequency, but not low level persistent RFI. We also identify corrupted calibrator visibilities by requiring that the visibility phase for calibrators be stable over time. We note that if the calibrator is resolved at some baselines or if there is confusing structure in the field of view, the requirement that the phase be stable will not be satisfied. However, we find in practice that at the GMRT operating frequencies and the generally used calibrators, the visibility phase is stable enough for this algorithm to work. Often data corruption affects entire subsets of data. For example, there may be narrow line RFI in one channel, or a few contiguous channels. Similarly, a particular baseline may have corrupted data because of a correlator problem. We identify such subsets of corrupted data by making a second pass through the data. Channels, baselines or antennas for which more than a specified fraction of the data has been identified as corrupted in the first pass are identified and then completely flagged out. We note that general purpose packages like AIPS and CASA also have tools that flag data based on the assumption that the true visibilities should be continuous in the time frequency plane. While the algorithms that are used there are similar in spirit to those used here, `flagcal` consistently uses robust estimators and the most compute intensive calculations are parallelized. It is also worth noting that a wrong choice of parameters in algorithms that are based on smoothness in the time-frequency plane could lead to over flagging of the data.

Calculation of the robust statistics as well as of the calibration solutions is computationally intensive. Fortunately both calculations can be trivially parallelized over time or frequency. Hence, wherever possible, these calculations

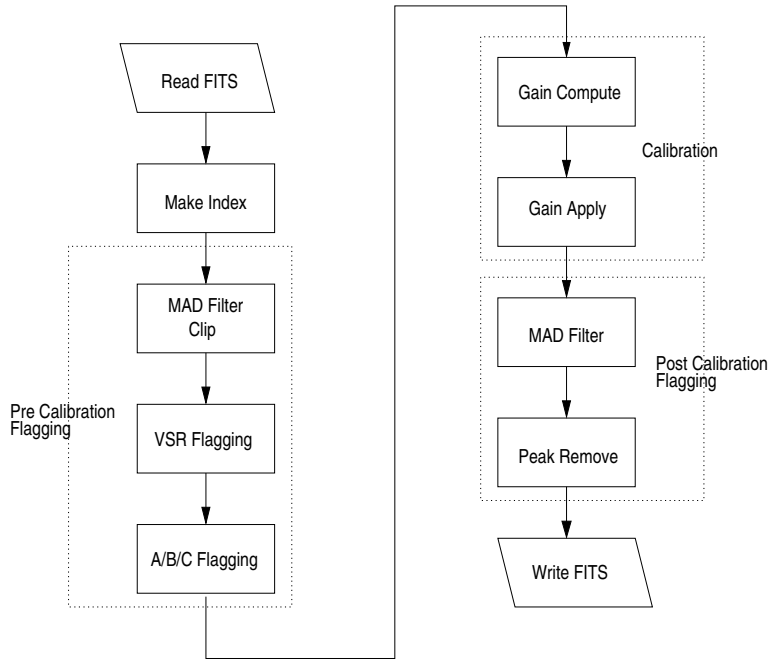


Fig. 1 Flow chart for the `flagcal` pipeline. See the text for more details.

are computed in parallel using OpenMP pragmas. This gives substantial speed ups when running the pipeline on the now standard multi-core workstations.

Below we give some details on the different algorithms used in the program as well as present sample results and images made using data processed through the pipeline.

2 Details of the Pipeline

The flowchart for the `flagcal` pipeline is shown in Fig. 1. The main stages of the pipeline are detailed below. The program consists of a number of routines some of which take several user defined parameters (see Table 2). These parameters can be specified in a plain text configuration file which is read by the program at run time. The configuration file also allows the user to specify which of the routines are to be run on a given data set, allowing a great deal of flexibility while executing the program.

2.1 Indexing and Pre Calibration Flagging

GMRT data is available as a random group UVFITS file. Accordingly, the `flagcal` program takes a random group UVFITS file as input and its output

is a calibrated and flagged random group UVFITS file. The first stage of the pipeline reads the input random group fits file and creates an index table for the file. The index table carries information on the start and stop record, as well as the source ID and source type (i.e. flux calibrator, phase calibrator or target source) for each scan. By a scan we mean a continuous observation of a given source, with fixed receiver and backend settings. A GMRT data file typically consists of a series of scans on different (e.g. flux calibrator, phase calibrator, target source) sources. The visibility data themselves are stored internally in the form of a row-major multi dimensional array. The size of a typical GMRT data is less than 4 GB. The current algorithm hence keeps all of the indexed data in memory. However, for large files (greater than 8GB) the algorithms can be easily modified to process the data scan by scan. This would keep the memory requirements modest while also minimising the amount of disk I/O. Each element of the array is a `complex3` (i.e. consisting of real, imaginary, and weight fields) structure. Data with negative weight are regarded as being flagged. A typical visibility data point for baseline b , stokes parameter s , frequency channel c and record r is indexed as $V[b + n_b(s + n_s(c + n_c \times r))]$, or $V_r(c, s, b)$ where n_b , n_s and n_c are the total number of baselines, stokes, and channels respectively.

The next stage of the pipeline does pre-calibration flagging. Several types of flagging are possible. The first is a simple thresholding or clipping, where all visibilities whose amplitudes lie above a user defined threshold are clipped. If the user has some prior information on corrupted baselines, antennas or channels, this can be passed to the pipeline in the form of a flag file. Clipping and initial flagging (on the basis of the input flag file) is followed by two stages of MAD filtering¹. A MAD filter flags all visibilities whose amplitudes differ from the median amplitude by more than a user defined threshold times the median absolute deviation (MAD). This filter has been chosen because of its robustness (i.e. the median and the median absolute deviation are robust to the presence of outliers, unlike the mean and the standard deviation). The two MAD filter steps before calibration are a Global Mad Filter and a Pre Mad filter. The Global Mad Filter flags the data points for which the visibility amplitude is discrepant as compared to the the global median and MAD. These statistical parameters are computed over all baselines, channels and stokes parameters present in the data file for that given source. For constructing the Global Mad Filter, the visibility data array A_i where i varies from 0 to $n_r \times n_b \times n_c \times n_s$ is split into n_r sub-arrays, one for each time sample. The median M_1^j and mad M_2^j are computed for each sub array². For each

¹ The median absolute deviation or MAD, which we refer to as M_2 in this paper, is defined by the median of the absolute deviation around the median value.

² For an array $\{x\} = (x_1, x_2, x_3, \dots)$

$$M_1 = \text{median}(x_1, x_2, x_3, \dots)$$

and

$$M_2 = \text{median}(|x_1 - M_1|, |x_2 - M_1|, |x_3 - M_1|, \dots)$$

source we collect the different m_1^j and m_2^j to construct the global median M_1 and MAD M_2 for that source. All visibilities for source k which satisfy

$$\text{abs}(|V_r(c, s, b)| - M_1^k) > \text{thrs_gmad} \times M_2^k \quad (1)$$

are flagged, where `thrs_gmad` is a user defined threshold (see Table 2). Since the Global Mad Filter takes into account all the data points available for a source, it is expected to be very robust.

In contrast to global mad filtering, pre mad filtering is done only on the basis of visibility data of the specific channel, stokes parameter and scan. This filtering does not combine data from different channels, baselines or stokes parameters i.e., it works only in time domain. In the Pre-Mad Filter all visibilities that satisfy the following condition are flagged.

$$\text{abs}(|V_r(c, s, b)| - M_1^k(c, s, b)) > \text{thrs_pmad} \times M_2^k(c, s, b) \quad (2)$$

where `thrs_pmad` is a user defined threshold value and where $M_1^k(c, s, b)$ and $M_2^k(c, s, b)$ are the median and mad values of the visibilities which are computed separately for every source k , channel c , baseline b and stokes s .

2.1.1 VSR and ABC flagging

For observations of an unresolved point source calibrator, one would expect that, in the absence of RFI or instrumental problems, the phase of the visibility for any given stokes, baseline and channel would vary slowly with time. Consequently, for sufficiently small stretches of time, (typically 1-2 min at the GMRT) one would expect that the ratio of the amplitude of the vector sum of the visibilities to the scalar sum of the visibilities (the ‘‘vector to scalar ratio’’, VSR) would be close to unity. In this stage of the pipeline, the calibrator visibilities are broken up into blocks (of size `nflag` decided by the user) and for each block the ratio $\mathcal{R}^j(c, s, b)$ define below is computed.

$$\mathcal{R}^j(c, s, b) = \frac{\left| \sum_{r=r_{start}^j}^{t_{end}^j} \mathbf{V}_r(c, s, b) w_r(c, s, b) \right|}{\sum_{r=r_{start}^j}^{r_{end}^j} |\mathbf{V}_r(c, s, b) w_r(c, s, b)|} \quad (3)$$

Where $\mathbf{V}_r(c, s, b)$ and $w_r(c, s, b)$ are the complex visibility and weight respectively for the flagging block j which starts at record r_{start}^j and ends at record r_{end}^j . A flagging block is considered good only if it has a VSR $\mathcal{R}^j(c, s, b)$ that is greater than a user defined threshold `thrs_vsr`. Clearly, the expected value of $\mathcal{R}^j(c, s, b)$ is unity for the case when the phase of the visibility remains constant and is zero when the phase varies randomly between zero and 2π . We note that this test is applied only on the calibrator scans, and not on the target source scans.




Fig. 2 Plots of the ratio of the modulus of the vector sum to the scalar sum (VSR, see Eqn. 3) of the visibility for the calibrator scans. For good data the VSR is expected to be 1.0. All data points for which the VSR is smaller than a user given threshold (0.95 in the above case) are flagged and are shown by cross (\times) symbol. Data points which are not flagged are shown by plus ($+$) symbol.

In Fig. 2 is shown the VSR for a typical frequency channels for few baselines. Data blocks for which the VSR is less than a user settable threshold (0.95 in this case) are flagged. The default value of `nflag` is 2min. This value was iteratively determined after running the program on several different GMRT data sets.

RFI or instrumental problems often affect a finite subset of the data, for example, a particular channel may be bad, or a particular baseline or antenna. In order to identify such subsets of corrupted data, the VSR $\mathcal{R}^j(c, s, b)$ is marginalized over two of its indices to obtain a set of normalized “measures of the badness” \mathcal{A} , \mathcal{B} and \mathcal{C} for all the antennas, baselines and channels respectively. Specifically, for each calibrator scan i we compute:

$$\mathcal{A}_l^i = \frac{1}{N_A} \sum_{c=1}^{n_c} \sum_{s=1}^{n_s} \sum_{b=1}^{n_b} \mathcal{F}^i(c, s, b) [\delta_D(p-l) + \delta_D(q-l)] \quad (4)$$

$$\mathcal{B}_b^i = \frac{1}{N_B} \sum_{c=1}^{n_c} \sum_{s=1}^{n_s} \mathcal{F}^i(c, s, b) \quad (5)$$

$$\mathcal{C}_c^i = \frac{1}{N_C} \sum_{b=1}^{n_b} \sum_{s=1}^{n_s} \mathcal{F}^i(c, s, b) \quad (6)$$

where the function δ_D is unity when its argument is zero, otherwise it is zero, p, q are the indices of the antennas that comprise the baseline b . We normalise \mathcal{A}, \mathcal{B} and \mathcal{C} with $N_A = n_c \times n_s \times (n_a - 1)$, $N_B = n_c \times n_s$, and $N_C = n_b \times n_s$ respectively so that their values are constrained to lie between zero and one, i.e. they represent the fraction of data that has been identified as corrupted in the earlier pass through the data. $\mathcal{F}^i(c, s, b)$ takes values of either 0 or 1. It is set to 0 only if one of the two conditions below are satisfied.

1. The number of “good blocks” (i.e. blocks for which the value of \mathcal{R} is greater than the user defined threshold `thrs_vsr`) is larger than a user defined threshold value `thrs_ngood`. Note that `thrs_ngood` is in units of the flagging block length `nflag`. Multiplying `thrs_ngood` by the time duration of the flagging block length would give the corresponding time interval.
2. The length of the longest contiguous time interval in the scan for which all the flagging blocks are “good” (i.e. as defined above) is greater than a user defined threshold value `thrs_block`. As above, this is in units of the flagging block length.

The idea behind the second test is that sometimes the problem causing the data to be corrupted (e.g. hardware failure) occurs partway through the scan. In that case there would be a contiguous stretch of time for which the data quality is good. This can be distinguished from an intermittent problem because if the problem is intermittent the probability that nonetheless `thrs_block` contiguous blocks remain unflagged is small. Note that we carry out the test (2) on a scan only when the test (1) fails.

As mentioned above, we apply two passes of flagging for the calibrator scans, the first one based on \mathcal{R} (where individual flagging blocks are flagged) and the second based on \mathcal{A}, \mathcal{B} and \mathcal{C} (where entire antennas, baselines or channels are flagged). Antennas for which more than `thrs_ant` of the data is already flagged, baselines for which more than `thrs_base` of the data is already flagged and channels for which more than `thrs_chan` of the data is already flagged are flagged in this second pass. Note that these thresholds represent fractions of the total data for that antenna, baseline or channel.

For the target source, the assumption that the visibility phase varies slowly with time need not be satisfied. However, it is often the case that a data subset (e.g. a channel or a baseline) that is corrupted during the calibrator scan, is also corrupted during the target source scan. Hence the user is allowed to interpolate the flags from the calibrator scans on to the target source scans. Specifically for the target source scans i , the values of $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are computed using

$$X^i = \max(X^{i-1}, X^{i+1}) \quad (7)$$

where X can be \mathcal{A}, \mathcal{B} or \mathcal{C} . The above interpolation scheme assumes that there is a calibrator scan before and after every target source scan, which is typical for most GMRT observations. In case the target source scan is not bracketed by calibrator scans, the interpolation is done from the nearest calibrator scan.




Fig. 3 This figure shows the measure of badness for antennas \mathcal{A} for different calibrator scans (shown by different line styles). The value of \mathcal{A} is close to one for antennas which are "bad" and it is close to zero for antennas which are "good".




Fig. 4 This figure shows the measure of badness for different channels \mathcal{C} for different calibrator scans (shown by different line styles). The value of \mathcal{C} is close to one for channels which are "bad" and it is close to zero for channels which are "good".

2.2 Gain Computation and Calibration

2.2.1 Gain Computation

In situations where baseline based errors can be ignored, the relation between the observed visibility V_{ij} and the true visibility \tilde{V}_{ij} can be written as:

$$V_{ij} = g_i g_j^* \tilde{V}_{ij} \quad (8)$$

where g_i is complex *gain* (i.e. the overall gain, including instrumental and ionospheric contributions) of antenna i . The goal of a calibration program is to use the observed visibilities V_{ij} to determine the gains g_i (which are expected to vary with frequency, stokes parameter and time) and hence the true visibilities \tilde{V}_{ij} . In “ordinary calibration”, which is the scheme implemented here, it is assumed that the g_i , , do not vary much with time. This allows one to compute g_i using the observations of the calibrators (for which the true visibilities are known) and interpolate these on to the visibilities of target source.

As is well known, Eqn. 8 describes an over constrained system since one has a total of $n_a(n_a - 1)/2$ complex visibilities V_{ij} from which one needs to determine n_a complex gains. In the `flagcal` implementation an iterative method of solving for g_i by least square minimization (see e.g. Bhatnagar (2001)) is followed. The estimated complex gain of g_i^n at iteration n is given by:

$$g_i^n = g_i^{n-1} + \mathbf{alpha} \times \left[\frac{\sum_{j \neq i}^{n_a} X_{ij} g_j^{n-1} w_{ij}}{\sum_{j \neq i} |g_j^{n-1}|^2 w_{ij}} - g_i^{n-1} \right] \quad (9)$$

where $X_{ij} = V_{ij}/\tilde{V}_{i,j}$. We use the following convergence criteria for the above iterative solutions

$$\mathit{Max} \left(\frac{|g_i^n - g_i^{n-1}|}{|g_i^{n-1}|} \right) < \mathbf{eps} \quad (10)$$

where the maximum is computed over all antennas. As formulated above the solutions g_i are independent for each frequency channel, stokes and record, and this algorithm can hence be trivially parallelized over any of them. In the `flagcal` implementation the calculation is parallelized over frequency channels. The solution is computed after averaging the data over time; the length of the solution interval `nsol` and the maximum number of iterations `niter` can be set by the user. The typical solution interval is 2m and convergence is typically obtained in less than 10 iterations. The computed gains compare well with those computed independently by tasks in the AIPS package. The calibration is estimated to be accurate at about the 10% level. The residual errors are due to a combination of systematic errors (for e.g. any elevation dependent gains are not corrected for) and do not reflect the precision of the calculations.

2.3 Post Calibration Flagging

After calibration, the visibilities on the calibrators should have identical amplitudes and zero phase, within the noise. Data which is slightly corrupted will hence be easier to detect than before calibration. We hence allow for a stage of post calibration MAD filtering, in which all visibilities $V_r(c, s, b)$ for which

$$\text{abs}(|V_r(c, s, b)| - M_1(c, s, r)) > \text{thrs_qmad} \times M_2(c, s, r) \quad (11)$$

are flagged. `thrs_qmad` is a user defined threshold, and $M_1(c, s, r)$ and $M_2(c, s, r)$ are the median and mad for record r , channel c , stokes s and baseline b . The intention is that the calibration step can then be repeated, leading to more robust gain solutions, although this is currently not yet implemented in `flagcal`.

2.3.1 SmoothSubtractThreshold (SST) Flagging

As discussed in Sec. 2.1.1 the expected visibility for the target source is not known a priori, and hence the flagging algorithms suggested there may not be appropriate for target source visibilities. However, even though one does not know the expected visibilities for the target source, it is reasonable to assume that the visibilities do not vary very rapidly with either time or frequency. One could then try and identify corrupted data by either looking for deviant points after polynomial fitting to the visibilities, or smoothing the data (see e.g. Chengalur (1996); Kanekar & Chengalur (1997); Urvashi (2003); Offringa et al. (2010)). In `flagcal` the latter approach is followed. The data is smoothed over one dimension (either time or frequency) and then visibilities that differ from the smoothed value by more than a threshold amount are flagged. This can be repeated for `nriter_X` iterations (where X can be time t or frequency f) with the smoothing size increasing with each iteration. Specifically, for a set of n visibilities x_i one first computes the smoothed visibilities over a window of length $2w_j + 1$, viz.

$$X_i^j = \begin{cases} \frac{1}{2w_j+1} \sum_{j=0}^{2w_j+1} x_{i+j} & \text{for } 0 < i < w_j \\ \frac{1}{2w_j+1} \sum_{j=i-w_j}^{i+w_j} x_{i+j} & \text{for } w_j < i < n - w_j \\ \frac{1}{2w_j+1} \sum_{j=-(2w_j+1)}^0 x_{i+j} & \text{for } n - w_j < i < n \end{cases} \quad (12)$$

then subtracts the smoothed data from the original

$$Y_i^j = x_i - X_i^j \quad (13)$$

Data points for which

$$\text{abs}(Y_i^j - M_1) > \text{thrs_peak_X}^j \times M_2 \quad (14)$$

Module name	1	2	4	8
Global Mad filter	205.68	108.08	66.93	41.45
Pre-Mad filter	48.49	25.66	15.23	9.17
Gain computation	81.56	53.58	31.51	23.45
Peak in time	218.94	118.76	59.06	28.64
Peak in frequency	202.61	101.94	55.39	27.89
Post Mad filter	45.79	27.84	20.01	16.88
Total time	803.07	435.86	248.15	147.48

Table 1 This table shows the time taken in seconds by various modules for 1, 2, 4 and 8 threads (column 2, 3, 4 and 5 respectively). The last three rows show the total time taken (without considering the time take by read/write and other serial sections). For timing we used a visibility data set with 128 frequency channels, 2 stokes, 435 baselines and 1506 time samples (16 second sampling). The file size is ~ 2 GB and it takes ~ 73 s to read in the file from disk.

where M_1 and M_2 are the median and median of absolute deviation for the series X_i are flagged. The user can choose a range of window sizes w_j , and thresholds T^j and the number of iterations $nriter_X$ (see table 2 with X replaced by t for time and f for frequency), over which this process can be repeated. As before this operation is parallelized using OpenMP pragmas.

The width of the smoothing window w^j and threshold $thrs_peak_X^j$ at iteration j is given by :

$$w^j = (dwin_X)^j \times win_min_X \quad (15)$$

and

$$thrs_peak_X^j = \left(\frac{1}{dpX} \right)^j thrs_peak_X \quad (16)$$

where descriptions of win_min_X , $dwin_X$, dpX and $thrs_peak_X$ are given in Table 2 (with X replaced by t or f) with their typical values.

We note that running this procedure on the target source could result in genuine emission from very strong transients being filtered out along with undesired RFI. In such situations, it would be better to disable this flagging. In the more general case, where one is searching for transients that are too faint to show up in the visibility of a single channel of a single baseline, this algorithm could be used.

3 Results

The `flagcal` programs have been tested several GMRT data sets. The program works well at removing RFI and calibrating for all the GMRT bands with minimal tuning of user input parameters. The default values of the parameters have been set after running the programs on various data sets and examining what data was flagged. The fact that many of the parameters are expressed in terms of the underlying robust statistics of the data also reduces the necessity to fine tune them for each data set. In Fig. 5 is shown a time




Fig. 5 [A] The top panel shows a time-frequency gray plot for a baseline for a typical GMRT observation at 157 MHz. The vertical stripes at constant intervals in the figure are the scans for phase calibrator. The presence of a large number of structures shows that the baseline is heavily corrupted due to RFI. [B] The bottom panel shows the same data after processing through `flagcal`. Note that the RFI has been flagged out, and the intensity scale also has become narrower. The wedge on right side shows the intensity scale in Jy.




Fig. 6 A 157 MHz image obtained after processing the data with `flagcal`. A snapshot of the processed visibilities from which this image has been made is shown in Fig. 5. Note that the only calibration applied is that computed by `flagcal`. The image was produced using the AIPS task `IMAGR` and has been `CLEANed`, but not self calibrated.

frequency plot of visibilities on a given baseline taken from a GMRT 150 MHz observation. This data set has been selected for illustration because at GMRT the 150 MHz band is generally worst affected by interference. Panel [a] shows the raw visibilities, while Panel [b] shows the visibilities after being processed through `flagcal`. As can be seen, the strong RFI visible in Panel [a] has been removed and the calibration of the data also reduces the variation in the intensity. About 30% of the data has been flagged out, note that this also includes data from antennas that had problems with their servo systems etc.

An image made with this processed data is shown in Fig. 6. The image was made using tasks in the AIPS package. No further calibration or flagging was done in AIPS. Instead the data was imaged and cleaned using the task `IMAGR`. The `UVRANGE` was set to 0-10 kL and the `UVTAPER` to 8 kL resulting in a resolution of $30'' \times 27''$. Given that the pipeline is intended only for a quick look at the data we do not do more detailed cross checks. We postpone these for a later stage as the pipeline becomes more mature.

Table 2 lists the parameters available to the user in tuning the performance of the `flagcal` program. The execution times for the various modules on an AMD Opteron (2.6GHz, 8 GB RAM) with dual processors, each with 4 cores is given in Table 1.

In summary, we present the initial results from `flagcal` a pipeline aimed at doing initial flagging and calibration of GMRT data. The aim of this version of the pipeline was to pre process the visibility data sufficiently in order to allow the user to make a quick look image of the data. We demonstrate this using a GMRT 150 MHz data set, a frequency band where RFI problems are generally severe.

Acknowledgements We are grateful to the referees of this paper for their careful reading of the draft, and for their several useful comments, which have substantially improved the readability of this paper. This work was partially supported the Australia-India Strategic Research Fund.

References

- Bhatnagar, S. 2001, PhD thesis, Pune University
- Briggs, F. H., & Kocz, J. 2005, *Radio Science*, 40, RS5S02
- Chengalur, J. N. 1996, Reduction of WSRT Compound Interferometer data, Tech. Rep. 453, NFRA
- Chengalur, J. N., Gupta, Y. G., & Dwarakanath, K. S. 2003, *Low Frequency Radion Astronomy*, 2nd ed. (NCRA)
- Kanekar, N., & Chengalur, J. N. 1997, *MNRAS*, 292, 831
- Kocz, J., Briggs, F. H., & Reynolds, J. 2010, *AJ*, 140, 2086
- Offringa, A. R., de Bruyn, A. G., Biehl, M., Zaroubi, S., Bernardi, G., & Pandey, V. N. 2010, *MNRAS*, 405, 155
- Paciga, G., et al. 2011, *MNRAS*, 413, 1174
- Sirothia, S. K. 2009, PhD thesis, TIFR
- Taylor, G. B., Carilli, C. L., & R. A. Perley, R. A., eds. 1999, *Synthesis Imaging in Radio Astronomy II*, Vol. 180 (ASP Conference Series)
- Thompson, A. R., Moran, J. M., & Swenson, G. W. J. 2001, *Interferometry and Synthesis in Radio Astronomy*, 2nd Edition (Wiley)
- Urvashi, R. V. 2003, Automatic RFI identification and flagging, Tech. Rep. R00202, NCRA-TIFR

Modules and their parameters		
Module name	Parameter	Description of parameters (with reference)
Gain solution	<code>nsol(4)</code> <code>niter(50)</code> <code>alpha(0.1)</code> <code>eps(0.01)</code> <code>refant(1)</code>	number of record used for gain solution (see §2.2.1) number of iterations for gain solution (see §2.2.1) a coefficients (see Eqn (9)) used for checking convergence (see Eqn (10)) reference antenna (which has phase zero)
Clip	<code>thrs_clip(1000.0)</code>	threshold used for clipping
Global mad filter	<code>thrs_gmad (9)</code>	threshold used for filtering (see Eqn (1))
Pre mad filter	<code>thrs_pmad (9)</code>	threshold used for filtering (see Eqn (2))
VSR/ABC flagging	<code>nflag(6)</code> <code>thrs_vsr (0.95)</code> <code>thrs_ngood(0.7)</code> <code>thrs_block(0.5)</code> <code>thrs_ant (0.3)</code> <code>thrs_base(0.3)</code> <code>thrs_chan(0.3)</code>	number of samples used for VSR (see §2.1.1 para 1) threshold for VSR (see Eqn (2.1.1)) threshold for bad block (see point 1 in §2.1.1) threshold for bad block (see point 2 in §2.1.1) threshold for bad antennas (see Eqn (2.1.1)) threshold for bad baseline (see Eqn (2.1.1)) threshold for bad channels (see Eqn (2.1.1))
Post mad filter	<code>thrs_qmad(9)</code>	threshold for post mad filtering (see Eqn (11))
SST flagging in time	<code>nriter_t(4)</code> <code>thrs_peak_t(9.0)</code> <code>dpt(0.7)</code> <code>win_min_t(2)</code> <code>dwin_t(2)</code>	number of SST iterations (see §2.3.1) starting threshold for SST (see Eqn (16)) factor for decreasing SST threshold (see Eqn (16)) starting smoothing window for SST (see Eqn (15)) factor for increasing SST window (see Eqn (15))
SST flagging in frequency	<code>nriter_f(4)</code> <code>thrs_peak_f(9.0)</code> <code>dpf(0.7)</code> <code>win_min_f(2)</code> <code>dwin_f(2)</code>	number of SST iterations (see §2.3.1) starting threshold for SST (see Eqn (16)) factor for decreasing SST threshold (see Eqn (16)) starting smoothing window for SST (see Eqn (15)) factor for increasing SST window (see Eqn (15))

Table 2 The Flagging and Calibration pipeline FLAGCAL consists of modules which are used for calibration and flagging. Every module has its set of parameters and the values of parameters control the effectiveness and accuracy of the module. In most cases the default (typical) values of the parameters can be used, however, in some case the tuning of parameters may be required. In the above table the first column gives the names of modules, the second and the third columns list the parameters and give their description respectively. We also give the typical values of the parameters in the second column (in bracket) and the reference in the text where these parameters are discussed. The names of the parameters in the above table are identical to that are given in the code. Note that the typical length of a GMRT records is 16 seconds.