



RFI Mitigation with BETA

Keith Bannister, Greg Hellbourg, Aidan Hotan

ASKAP Commissioning and Early Science Memo 012

March 15, 2016

CSIRO Astronomy and Space Science
Cnr. Vimiera and Pembroke Roads
PO Box 76, Epping, NSW 1710, AUSTRALIA Telephone : +61 2 9372 4100
Fax : +61 2 9372 4310

Copyright and disclaimer

© 2016 CSIRO To the extent permitted by law, all rights are reserved and no part of this publication covered by copyright may be reproduced or copied in any form or by any means except with the written permission of CSIRO.

Important disclaimer

CSIRO advises that the information contained in this publication comprises general statements based on scientific research. The reader is advised and needs to be aware that such information may be incomplete or unable to be used in any specific situation. No reliance or actions must therefore be made on that information without seeking prior expert professional, scientific and technical advice. To the extent permitted by law, CSIRO (including its employees and consultants) excludes all liability to any person for any consequences, including but not limited to all losses, damages, costs, expenses and any other compensation, arising directly or indirectly from using this publication (in part or in whole) and any information or material contained in it.

Contents

Summary	1
1 Introduction	2
2 Satellite Radio Frequency Interference	2
3 Mitigating RFI with PAF-based projection algorithms	3
4 Estimators for the RFI subspace dimensionality	6
5 Real-time Implementation	6
6 Visibility Processing	8
7 Results	8
7.1 Looking directly at a GPS satellite	8
7.2 Observing PKS B1934–63.	8
8 Conclusions and further work	16
9 Acknowledgements	17

Summary

The Boolardy Engineering Test Array (BETA) is a 6 antenna interferometer, with each antenna having a 188-element phased array feed at the prime focus. Here we attempt to mitigate the affect of radio frequency interference (RFI) by updating the beam weights using a projection algorithm applied to the measured single-dish array covariance matrix. We observe an astronomical source at 1225 MHz, which contains substantial interference from the GPS L2 band, and compare a number of different projection algorithms by looking at the affect on correlated visibilities, and on the resulting images. The RFI projection algorithms perform remarkably well, yielding a promising future for this technique.

Contents

1 Introduction

The Boolardy Engineering Test Array (BETA) is a 6-antenna prototype of the Australian Square Kilometre Array Pathfinder (ASKAP) (Hotan et al., 2014). Each 12 m dish is equipped with a Phased Array Feed (PAF) comprising 188 elements in two orthogonal linear polarisations, arranged in a roughly circular pattern at the prime focus. Beamforming hardware at each antenna forms up to 9 dual-polarisation beams by calculating the weighted sum of the elements. Typically, the programmable weights are determined offline from specially targeted observations of a reference source and a blank noise field, and uploaded to the beam former at the beginning of each observing run. The voltage stream from each beam is correlated with the same beam from all other antennas to form the interferometric visibilities.

In addition to forming beams, the beamformers can produce Array Covariance Matrices (ACMs), which are the complex cross correlation between all pairs of the 188 elements in a given antenna. An ACM is formed for 64 1 MHz channels spread throughout the 304 MHz observing band (i.e. with a gap of 4–5 MHz between each ACM), and dumped on a programmable dump time, which is typically 1 second of integration downloaded every 2 seconds. Only every 4th voltage sample is used in calculating the ACM and longer integrations can be used if needed.

2 Satellite Radio Frequency Interference

In spite of its location at an extremely RFI quiet site in the Murchison Shire, Western Australia, BETA is still subject to strong interference from satellites, and occasional terrestrial interference during ducting events. Global navigation systems (GNSS) such as the GPS (United States), Galileo (Europe), GLONASS (Russia) and Beidou (China) all transmit in the BETA bands and can interfere strongly with the reception of astronomical signals. The frequency plan of a number of these systems is shown in Figure 2.1.

The systems themselves are at different levels of maturity. GPS is well established, with 33 operational satellites. GLONASS somewhat fewer, and Galileo and Beidou fewer still. But the number of satellites is only expected to increase, making satellite RFI an increasingly severe problem.

The orbital parameters of all (known) GNSS satellites are measured and publicly released by NORAD¹. The satellites themselves transmit their own orbital parameters,

¹<https://celestrak.com>

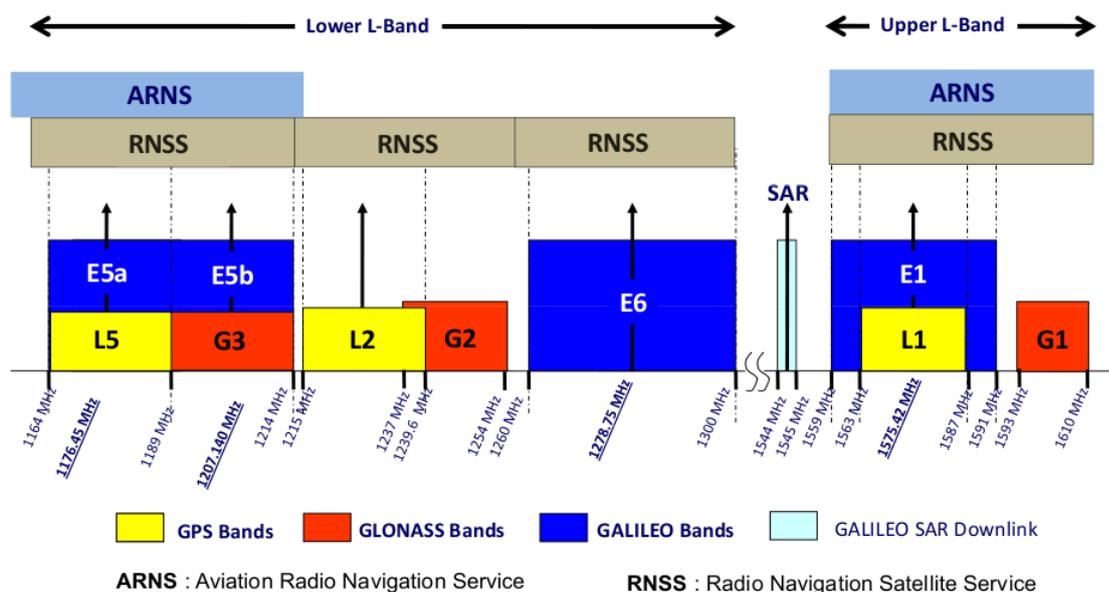


Figure 2.1: Frequency plan for GPS, GLONASS and Galileo systems. Credit: www.navipedia.net

but the data format is system specific and requires a receiver on site. For this analysis, the homogeneous NORAD data is the easiest to obtain and use, in spite of its somewhat lower accuracy. The position of each satellite can be easily computed by downloading the two line elements (TLEs) and using a standard package such as PYEPHEM to compute the azimuth and elevation of a satellite at any time. Figure 2.2 illustrates that there is a large number of GPS satellites above the ASKAP horizon during a 12 hr observation of PKS B1934–63 described hereafter. There are multiple passes of a GPS satellite within 10 degrees of the target (Table 2.1).

3 Mitigating RFI with PAF-based projection algorithms

For the investigations described in this memo we used a projection algorithm using the ACM for a single dish. The single-dish approach was chosen as it was the easiest to implement given the constraints of the BETA hardware and software. The projection algorithm takes the initial weight vector w_0 , and computes a new weight vector w_t nulling the contribution of a satellite estimated through the singular value decomposition of the ACM Hellboug (2014). It has the advantage of not requiring any communication between antennas. The projection algorithms operate as follows.

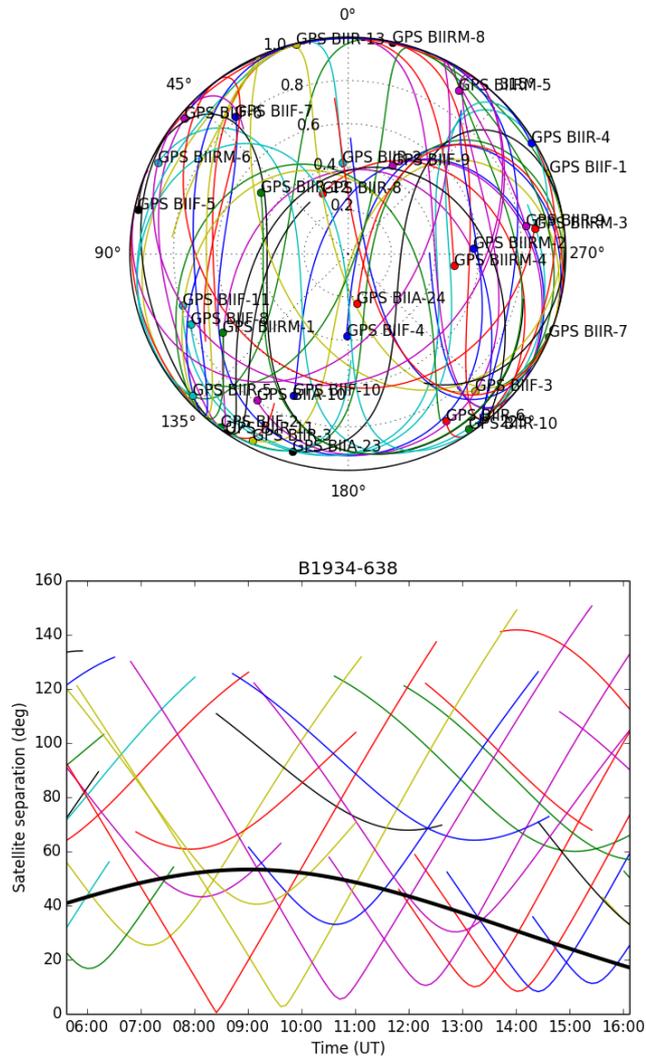


Figure 2.2: Prediction of the GPS constellation for a 12 hr period starting 2015-11-04 05:36:52 UT. Top panel: trajectory of each satellite in azimuth and elevation from BETA. The circles mark the position of the satellite at the start time. Bottom panel: Separation of satellite from the position of PKS B1934-63. The elevation of PKS B1934-63 is shown as a thick black line.

Table 2.1: Close passes to PKS B1934–63 of a GPS satellites for the 12 hrs after 2015-11-04 05:36:52 UT.

Satellite	Time of minimum distance (UT)	Minimum Distance (deg)
GPS BIIRM-1	06:02	16.7
GPS BIIRM-3	08:24	0.4
GPS BIIF-1	09:38	2.4
GPS BIIRM-5	10:43	5.4
GPS BIIF-9	12:18	10.5
GPS BIIR-8	13:13	10.0
GPS BIIA-24	14:02	8.3
GPS BIIF-4	14:24	8.2
GPS BIIF-10	15:25	11.2

Given the ACM at time t , \mathbf{R}_t , compute the singular values (\mathbf{s}_t) and singular vectors (\mathbf{U}_t) with the singular value decomposition:

$$\{\mathbf{U}_t, \mathbf{s}_t\} = \text{SVD}(\mathbf{R}_t) \quad (3.1)$$

where we assume SVD returns singular values in decreasing order in the vector \mathbf{s}_t , and the corresponding singular vectors as column vectors of the matrix \mathbf{U}_t . The next step is to form \mathbf{A}_{est} , which is a non-square matrix whose vectors span the vector space occupied by the RFI. The simplest method is to assume (or estimate) the dimensionality of the RFI subspace, n , in a single ACM snapshot, then take the n singular vectors associated with the n dominant singular values.

$$\mathbf{A}_{\text{est}} = \mathbf{U}_t[:, \mathbf{0} : n] \quad (3.2)$$

Alternatively, \mathbf{A}_{est} can be made of the dominant singular vectors from the previous m snapshots:

$$\mathbf{A}_{\text{est}} = [\mathbf{U}_t[:, \mathbf{0}], \mathbf{U}_{t-1}[:, \mathbf{0}], \dots, \mathbf{U}_{t-m-1}[:, \mathbf{0}]] \quad (3.3)$$

Once \mathbf{A}_{est} is generated, and providing $\mathbf{A}_{\text{est}}^H \mathbf{A}_{\text{est}}$ is full-rank, then the orthogonal projection matrix is calculated according to:

$$\mathbf{P}_{\text{orthog}} = \mathbf{I} - \mathbf{A}_{\text{est}} (\mathbf{A}_{\text{est}}^H \mathbf{A}_{\text{est}})^{-1} \mathbf{A}_{\text{est}}^H \quad (3.4)$$

where $(\cdot)^H$ indicates the Hermitian transpose and \mathbf{I} the identity matrix. Equation 3.4 is only invertible if the columns of \mathbf{A}_{est} are independent. This is strictly true if it is

formed from the columns of \mathbf{U}_t , as the SVD is an orthogonal decomposition. However, if \mathbf{A}_{est} is formed from Equation 3.3, then it may be poorly conditioned. In that case, columns can be iteratively removed from \mathbf{A}_{est} until it is well conditioned. The oblique projector Hellbourg et al. (2012) can be calculated according to:

$$\mathbf{P}_{\text{oblique}} = \mathbf{w}_0(\mathbf{w}_0^H \mathbf{P}_{\text{orthog}} \mathbf{w}_0)^{-1} \mathbf{w}_0^H \mathbf{P}_{\text{orthog}} \quad (3.5)$$

Finally, the new weights can be calculated as:

$$\mathbf{w}_{t,\text{orthog}} = \mathbf{P}_{\text{orthog}} \mathbf{w}_0 \quad (3.6)$$

$$\mathbf{w}_{t,\text{oblique}} = \mathbf{P}_{\text{oblique}}^H \mathbf{w}_0 \quad (3.7)$$

4 Estimators for the RFI subspace dimensionality

We used 2 different estimators to determine the number of interferers.

In Method A, we took the logarithm (base 10) of the singular values, and fit a 3rd order polynomial between the 10th and 83rd singular value. The number of interferers n_A was estimated as the number of singular values that exceeded the fit by 0.1.

In Method B, we took the logarithm (base 10) of the singular values, and calculated the ratio of each singular value with the next smaller singular value. The number of interferers n_B was estimated as the number of these ratios which exceeded 0.1.

5 Real-time Implementation

We created a set of weights in the usual way, with observations of the Sun and max S/N beam forming (Hotan et al., 2014). We processed the standard weights file ² with a script ³ to make a new weights file ⁴ that had the boresight weights (beam 0) for all beams. Therefore, the initial state where all beams reported exactly the same data. On loading, the script loaded the initial weights from the weights file.

The procedure described in Sections 3 and 4 is implemented as an OSL script ⁵. This script subscribes to ACMs from the beam former of a single antenna, computes a new set of weights, and uploads the weights to that antenna. We run the script 5 times ⁶ once for each of the 5 (working) BETA antennas.

²201511040243_6128_9beams

³\$ACES/tools/copyweights.py

⁴201511040243_6128_b0

⁵osl_a_rfi_mitigation.py

⁶setting the antenna of interest using the -a command line switch

The script was configured to integrate for 2 s and only send every second integration (i.e. a 50% duty cycle). We found that every 2nd ACM received by the script contained (essentially) zeros. We therefore had a 2s integration every 8s. The script was configured to download ACMs from card 15, FPGA C, which corresponds to a sky frequency of 1225 MHz, when observing at the standard band 2 frequency of 1119.5 MHz. That sky frequency is in the GPS L2 band.

The script spawned a separate thread for processing ACMs. When the script received an ACM, it immediately pushed the raw ACM data onto a queue. The processing thread popped the ACM data off the queue, calculated new weights, and pushed the weights onto a weights queue. The main thread watched the weights queue and uploaded weights whenever new weights were on the queue. Low level software copies identical weights across the 5 MHz bandwidth that is processed by the card and FPGA of interest. Thus, while the ACM and weights are calculated for a 1 MHz channel at 1225 MHz, the weights are uploaded to 5 channels from 1225-1230 MHz. We expect there to be a loss of RFI suppression for channels far away from the measured ACM frequency. We do not consider these additional channels any further.

The ACM processing thread dealt with each polarisation independently. The cross polar sections of the ACMs were not used. For each polarisation, the input ACM was obtained (containing 94 elements), and the new weights were calculated. Both the initial weights and the calculated weights had zeros for the cross polar terms. We calculated weights for each beam using the same initial weights, so that the algorithms could be compared directly. The algorithms parametrization for each beam are described in table 5.1

Table 5.1: List of mitigation algorithms applied to each beam. n_A and n_B refer to automatic estimates of the number of emitters, as described in Section 4

Beam	Algorithm
B0	Max S/N beam (no mitigation)
B1	Orthogonal, snapshot $n = 1$
B2	Oblique, snapshot $n = 1$
B3	Orthogonal, historical, $m = 10$, $n = 1$
B4	Oblique, historical, $m = 10$, $n = 1$
B5	Orthogonal, snapshot, $n = n_A$
B6	Oblique, snapshot, $n = n_A$
B7	Orthogonal, snapshot, $n = n_B$
B8	Oblique, snapshot, $n = n_B$

6 Visibility Processing

We recorded visibility data for all beams as a single CASA (McMullin et al., 2007) measurement set. First we split the measurement sets into nine, one for each beam. Memory errors in the correlator occasionally produce visibilities with extremely high amplitudes. Therefore, we flagged any visibility with a raw amplitude $> 10^6$, which does a reasonable job of flagging visibilities that are corrupted by memory errors. After flagging, we converted visibility data to MIRIAD format Sault et al. (1995) for further processing.

In MIRIAD we channel-average the data by a factor of 54, yielding a 1 MHz frequency resolution. We performed bandpass calibration using a model of PKS-B1934–63 (Reynolds, 1994). We chose a relatively RFI-free period between 10:00 UT and 10:10 UT during the observation to derive the bandpass solution. In order to compare the effects of flagging to unflagged data, we made a copy of the unflagged data set, and flagged all visibilities with a Stokes I amplitude > 50 Jy. We performed gain and phase calibration with an interval of 1 minute using a relatively RFI-free 16 MHz band centered at 1012 MHz. In both cases we used AK06 as the reference antenna. To create images, we inverted the visibilities with a pixel scale 5 times smaller than the synthesized beam, and image size of 512 pixels, and Briggs robust weighting of zero. Each image was cleaned with 2000 iterations and restored with the restoring beam.

7 Results

7.1 Looking directly at a GPS satellite

To test the system in real-time, we pointed the antennas directly at a GPS satellite. Figure 7.1 shows the real-time plots produced by the RFI mitigation script running on an antenna. Such plotting enables us to directly monitor the RFI mitigation performance while the observation is in progress. The spectrum obtained from real-time monitoring is shown in Figure 7.2. This spectrum clearly shows the spectrum of the GPS L2 signal, among others, and the effect of the various mitigation algorithms taking some RFI power out of the correlation amplitude.

7.2 Observing PKS B1934–63.

We observed PKS B1934–63 for about 10 hrs starting at 2015-11-04 05:36:52 UT. The visibility amplitudes for the unmitigated beam (beam 0) as a function of time, for all baselines is shown in Figure 7.3. As predicted from the ephemeris modelling earlier, there is a close pass of a GPS satellite around 08:26 UT, as well as a number of other, weaker passes. The visibility amplitude varies with baseline, probably because the satellite is moving so fast, it is smeared on the longer baselines.

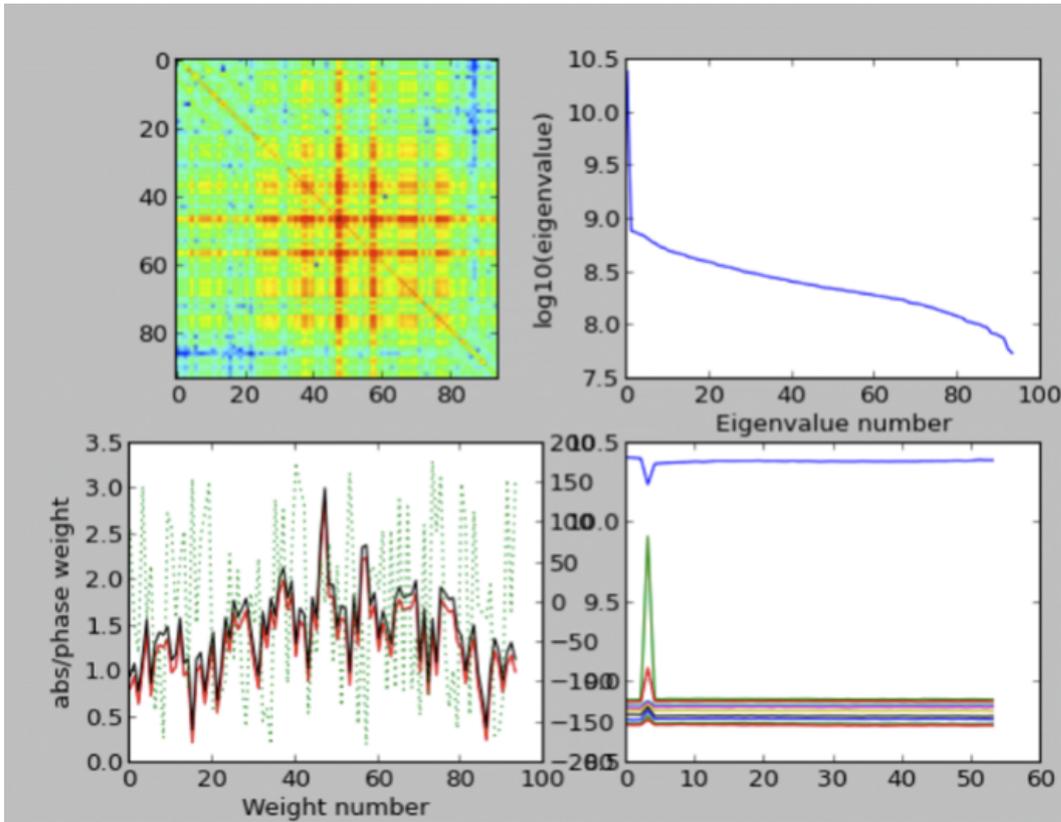


Figure 7.1: Real-time plots produced by the RFI mitigation script. Top left: ACM amplitude. The colour scale is an arbitrary log scale. Top right: Log of the singular values. Bottom left: beam 0 weight amplitude (red), beam 1 weight amplitude (black) beam 1 weight phase (green). Bottom right: Amplitude of the largest 10 singular values as a function of integration number.

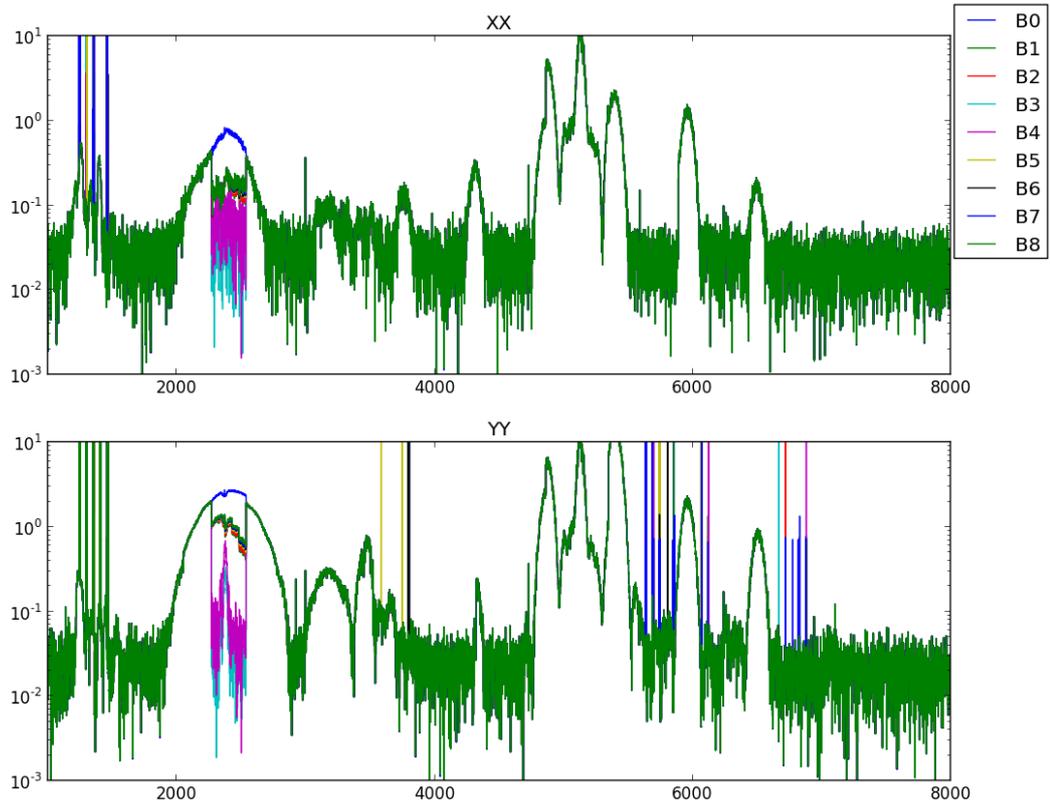


Figure 7.2: Cross correlation amplitude between AK01-AK06 for 2 polarisations and all beams. The frequency resolution is 18 kHz. The mitigation is active in the region near channel 2500.

Figure 7.4 shows the visibility amplitude and phase for a particular baseline. Notably, the close pass of GPS BIIR-3 at 08:26 UT (shown in detail in Figure 7.5) produces a reduction in the visibility amplitude to around 0 Jy in the mitigated beam (B3), while at the same time, the unmitigated beam reaches amplitudes of around 2×10^4 Jy. We suspect this reduction is because the satellite came so close to the target (predicted distance 0.4 deg at 08:24 UT) that the projection algorithm nulls out the target as well as the satellite signal. Qualitatively, the mitigation performs well for the other passes. For example, the pass near 09:45 produces amplitudes near 4000 Jy and substantial scatter in phase in B0 while there is only marginally increased scatter in the amplitudes and phases in B3 at the same time. The increased amplitude in B3 at around 14:15 is probably due to the fact that there were 2 satellites at around 10 deg from the target at that time (making closest approaches at 14:02 and 14:24 respectively). Given the B3 algorithm only mitigates the primary singular vectors, only a small part of a combination of the multiple satellites has been nulled. Although it is notable that the auto-detecting algorithms (B5–B8) fared no better at this time. There are 3 phase jumps of a few minutes duration during the observations (at around 5:50, 6:15 and 13:40). The origin of these phase jumps is unknown, but we suspect either RFI in the band on which we did the phase calibration, or hardware errors.

Images of the data show that the interference mitigation performs extremely well. Figure 7.6 shows images of the 1 MHz channel centred at 1225 MHz for the flagged and unflagged data, and for all algorithms. Noise and dynamic range measurements are shown in Table 7.1. The RFI is so poor in the unflagged B0 data that PKS B1934–63 is not at all visible. However, good dynamic range is achieved for B3 and B4. Flagging improves the dynamic range of almost all algorithms, but the best dynamic range is still achieved by the B3 and B4 algorithms. No doubt the unmitigated data could be improved by more careful flagging, but this is true for all the data sets, and ultimately, the unmitigated dataset will always be poorer as more flagging reduces *uv*-coverage and integration time.

Comparing to the theoretical and measured RMS away from RFI, if we assume a System Equivalent Flux Density of 2500 Jy, a 10 hr observation should have a naturally weighted image RMS of 2 mJy in a 1 MHz channel. In a 1 MHz channel near 1012 MHz, our imaging achieves an RMS of 12 mJy. While these figures are not perfectly comparable (the SEFD is slightly better at 1012 MHz than at 1225 MHz), it nonetheless shows that the RFI mitigation is not perfect. No observations approach the ideal dynamic range (assuming the thermal noise limits the dynamic range) but this is most likely due to poor *uv*-coverage.

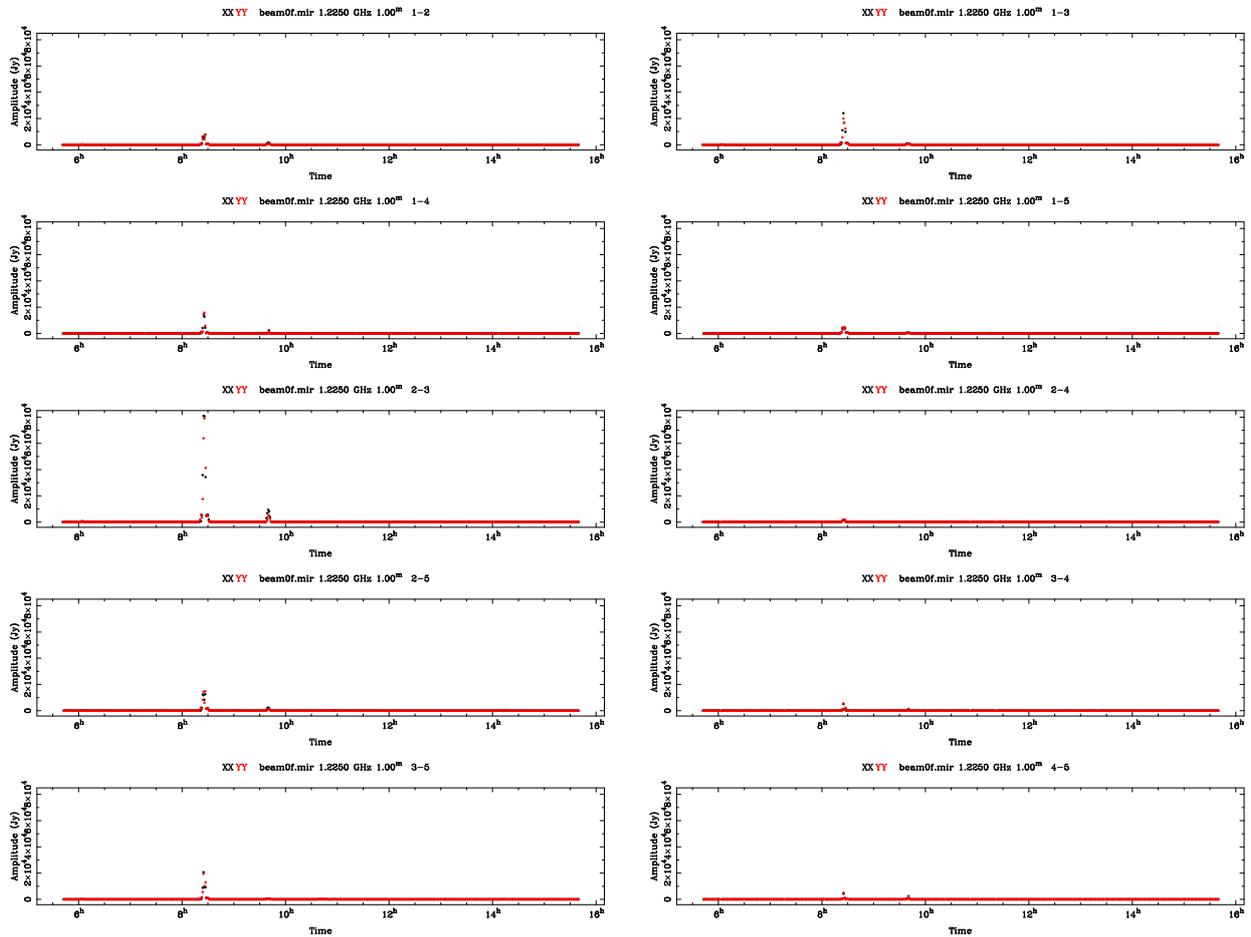


Figure 7.3: Calibrated visibility amplitudes for the unmitigated beam as a function of time while observing PKS B1934–63, at 1225 MHz. The data have been averaged to 1 minute intervals. A few close passes of GPS satellites are visible. Their amplitude is somewhat reduced by the averaging interval. The differing amplitudes on the different baselines is clearly visible, and probably due to time average smearing on the longer baselines.

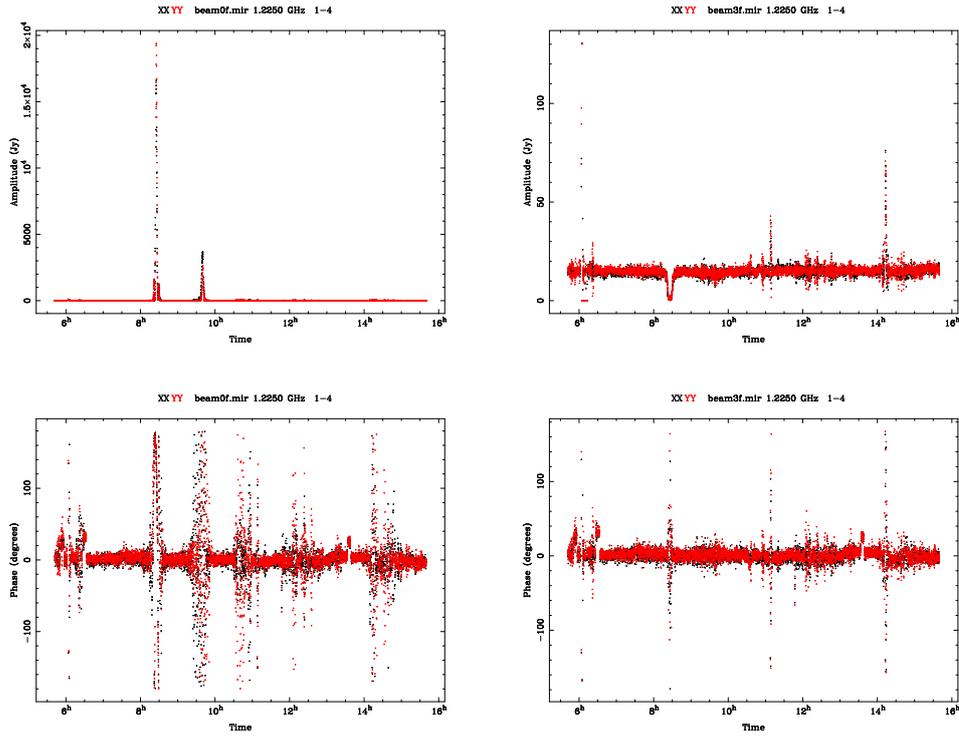


Figure 7.4: Visibility amplitude (top row) and phase (bottom row) of the 1 MHz channel centred at 1225 MHz on the AK06-AK01 baseline while observing PKS B1934–63. Left column: B0, right column: B3. The dip in B3 amplitude at 08:26 UT is because the satellite passes within $\simeq 0.4$ degree of PKS B1934–63 (the predicted pass time was 08:24). The pass at around 09:45 UT shows substantial interference in B0 but is well mitigated by B3.

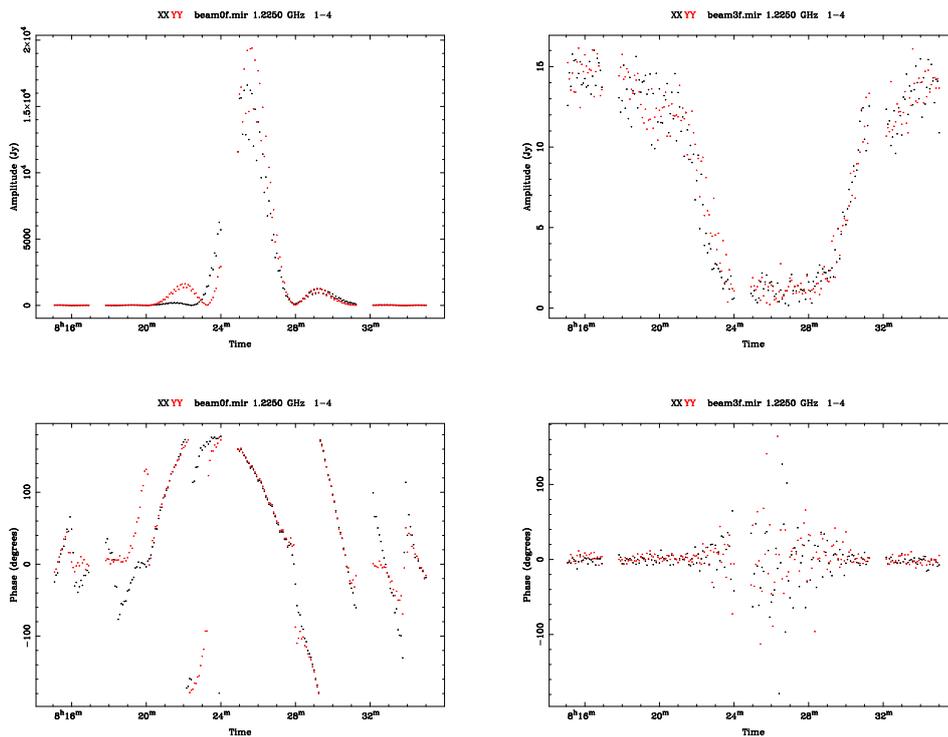


Figure 7.5: Same as Figure 7.4 but zoomed into the bright pass of GPS BIIRM-3, which was predicted to pass within 0.4 deg at 08:24 UT. These data are at the native resolution of 5 seconds.

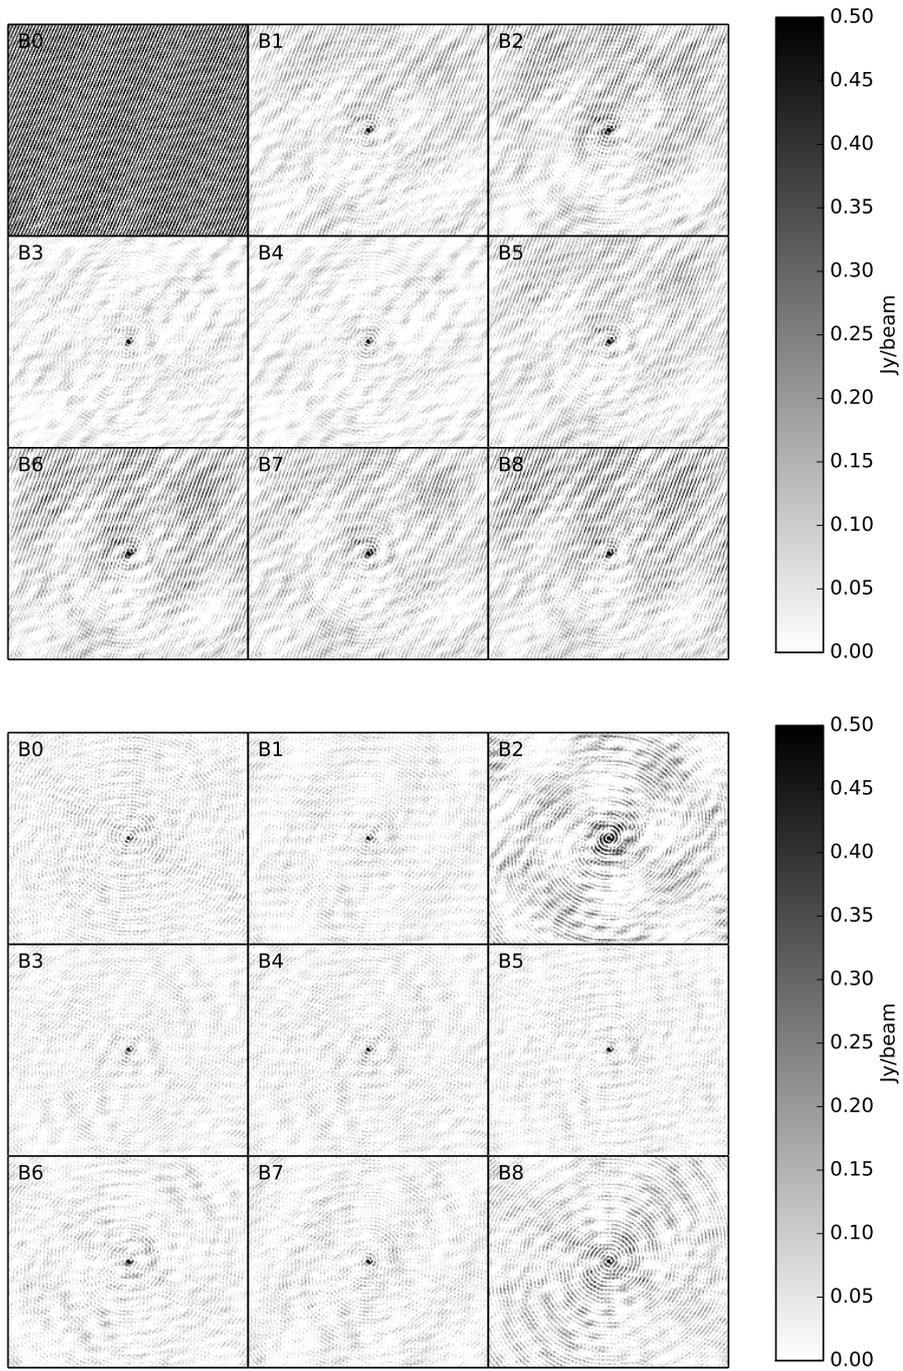


Figure 7.6: Images of PKS B1934–63 in the 1 MHz channel centred at 1225 MHz. Top 9 panels: Unflagged data. Bottom 9 panels: automatically flagged data.

Table 7.1: Statistics of images produced based on various mitigation methods, and with the flagged and unflagged data in a 1 MHz channel at 1225 MHz (in the GPS L2 band), as well as an RFI-free channel at 1012 MHz. The ideal numbers assume an SEFD of 2500 Jy and natural weighting.

Beam	Unflagged			Flagged			
	Peak Jy/bm	RMS Jy/bm	Dynamic Range :1	Peak Jy/bm	RMS Jy/bm	Dynamic range :1	Flagged %
B0	19.8	9.220	2	14.3	0.070	205	16
B1	14.7	0.144	102	14.4	0.068	210	9
B2	13.1	0.197	66	13.1	0.065	202	9
B3	14.4	0.059	245	14.4	0.054	268	2
B4	14.3	0.058	248	14.3	0.050	283	2
B5	14.8	0.153	96	14.5	0.065	222	10
B6	15.0	0.237	63	14.7	0.063	232	10
B7	15.0	0.164	91	14.6	0.068	214	11
B8	14.9	0.245	60	14.5	0.064	225	11
B0 (1012 MHz)	14.6	0.012	1130	14.8	0.093	1585	1
Ideal	15.1	0.002	7500				

8 Conclusions and further work

We have shown that PAF-based projection algorithms work extremely well in mitigating RFI from GPS satellites. While limited flagging is some of use in removing the same RFI, the best dynamic range is achieved by a combination of PAF-based RFI mitigation, and flagging.

It is clear that the mitigation is not perfect, and there are a number of theoretical and practical improvements that could be made. On the theory side, a more thorough investigation of the algorithms is required. Clearly the historical algorithms outperform the snapshot ones. Improving on the historical algorithm, including an estimate of the number of interferers, as well as the history, would be useful in the case of multiple interferers. No doubt our RFI rejection suffers as there is a substantial time delay between receiving an ACM and new weights coming into effect. In the case of sufficiently strong interferers, one could imagine tracking the spatial signature vectors of the RFI (e.g. with a Kalman filter) and predicting future vectors, which could be added to a_{set} to project out.

Practically, there are many improvements to make. The ability to mitigate a contiguous swath of spectrum, rather than 1 MHz in 5 is required for any real-world astronomy application. Improving the duty-cycle would also increase the S/N of the ACM and

the turn-around time, both of which should improve the interference rejection.

References

- Hellbourg, G. (2014). *Radio Frequency Interference spatial processing for modern radio telescopes*. Ph.D. thesis, Université d'Orléans.
- Hellbourg, G., et al. (2012). In *Statistical Signal Processing Workshop (SSP), 2012 IEEE*, pp. 93–96.
- Hotan, A. W., et al. (2014). *PASA*, **31** e041.
- McMullin, J. P., et al. (2007). In R. A. Shaw, F. Hill, D. J. Bell, eds., *Astronomical Data Analysis Software and Systems XVI, Astronomical Society of the Pacific Conference Series*, vol. 376, p. 127. [ADS].
- Reynolds, J. (1994). A revised flux scale for the AT compact array.
- Sault, R. J., et al. (1995). In R. A. Shaw, H. E. Payne, J. J. E. Hayes, eds., *Astronomical Data Analysis Software and Systems IV, Astronomical Society of the Pacific Conference Series*, vol. 77, p. 433. [ADS].

9 Acknowledgements

KB would like to thank Craig Haskins for providing the skeleton RFI mitigation script.

CONTACT US

t 1300 363 400
+61 3 9545 2176
e enquiries@csiro.au
w www.csiro.au

YOUR CSIRO

Australia is founding its future on science and innovation. Its national science agency, CSIRO, is a powerhouse of ideas, technologies and skills for building prosperity, growth, health and sustainability. It serves governments, industries, business and communities across the nation.

FOR FURTHER INFORMATION

CSIRO Astronomy and Space Science

Keith Bannister
t +61 2 9372 4295
e keith.bannister@csiro.au
w [Astronomy and Space Science](#)

CSIRO Astronomy and Space Science

Greg Hellbourg
t +61 2 9372 4148
e greg.hellbourg@csiro.au
w [Astronomy and Space Science](#)