The Astrophysical Journal, 864:160 (9pp), 2018 September 10 https://doi.org/10.3847/1538-4357/aad815

A Post-correlation Beamformer for Time-domain Studies of Pulsars and Transients

Jayanta Roy[1], Jayaram N. Chengalur[1], and Ue-Li Pen[2,3,4,5]
[1] National Centre for Radio Astrophysics, Tata Institute of Fundamental Research, Pune 411 007, India
[2] Canadian Institute for Theoretical Astrophysics (CITA), 60 St. George Street, Toronto, M5S 3H8, Canada
[3] Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada
[4] Canadian Institute for Advanced Research, CIFAR Program in Gravitation and Cosmology, Toronto, ON M5G 1Z8, Canada
[5] Perimeter Institute for Theoretical Physics, Waterloo, ON N2L 2Y5, Canada
Received 2018 June 5; revised 2018 July 21; accepted 2018 August 2; published 2018 September 12

Abstract

We present a detailed analysis of post-correlation (PC) beamforming (i.e., beamforming which involves only phased sums of the correlation of the voltages of different antennas in an array), and compare it with the traditionally used incoherent and phased beamforming techniques. Using data from the GMRT we show that PC beam formation results in a manyfold increase in the signal-to-noise for periodic signals from pulsars and reductions, of several orders of magnitude, in the number of false triggers from single-pulse events like fast radio bursts (FRBs). This difference arises primarily because the PC beam contains less red noise, as well as less radio frequency interference. The PC beam can also be more easily calibrated than the incoherent or phased array beams. We also discuss two different modes of PC beam formation: (1) by subtracting the incoherent beam from the coherent beam and (2) by phased addition of the visibilities. The computational costs for both these beam formation techniques, as well as their suitability for studies of pulsars and FRBs, are discussed. The techniques discussed here should be of interest for all upcoming surveys with interferometric arrays. Finally, we describe a time-domain survey with the GMRT using the PC beam formation as a case study. We find that PC beamforming will improve the current GMRT time-domain survey sensitivity by ∼2 times for pulsars with periods of few hundreds of milliseconds and by many-folds for even slower pulsars, making it one of the most sensitive surveys for pulsars and FRBs at low and mid radio frequencies.

Key words: instrumentation: interferometers – pulsars: general – techniques: interferometric – techniques: radar astronomy

1. Introduction

Despite over five decades of pulsar surveys there are only ∼2600 pulsars that have been discovered so far.[6] This is 5% or less of the total Galactic population of pulsars (estimates of the Galactic population range from 40,000 to 90,000 objects, see, e.g., Lorimer 2008) and a large population of pulsars remains to be discovered by current and future surveys. Even though there has been an accelerating rate of discovery over the last decade, this has not been uniform across the entire parameter space occupied by pulsars. For example, even though the population of known Galactic field millisecond pulsars (MSPs) has increased approximately fourfold over the last decade,[7] there has only been a ∼40% increase in the number of known slow pulsars (i.e., pulsars with period $P > 30$ ms). This very modest increase in the number of known slow pulsars is particularly unfortunate, since the already known population of relatively slow pulsars contains several interesting objects, such as double neutron stars (∼15 known, Tauris et al. 2017), which enable tests of strong field gravity, energetic young pulsars with significant spin-down noise, normal pulsars showing intermittency, drifting, and nulling, probing hitherto unknown emission physics, magnetars with extraordinarily high magnetic fields (∼29 known[8]), and ultra-slow pulsars (only 2 known) with periods >10 s that graze the theoretical death-line. One of the major reasons for the relatively slow increase in the number of such known pulsars is that the detection of these objects via periodicity searches can be severely affected by both instrumental red noise and radio frequency interference (RFI). Both of these phenomena particularly reduce the search sensitivity at the low-frequency end of the power spectrum of the detected time-series, which is where the signal from these objects is strongest.

In addition to pulsars, the population of time-domain radio transients consists of rotating radio transients (RRATs; McLaughlin et al. 2006; 112 known[9]) and fast radio bursts (FRBs) (Lorimer et al. 2007, Thornton et al. 2013; 33 known[10]). All the FRBs discovered to date are single events (except for one repeating FRB) of millisecond duration with dispersion measure (DM) values generally higher than the possible Galactic contribution. The non-repeating nature of these sources warrants real-time time-domain detections aided by simultaneous millisecond timescale imaging to localize these events in order to maximize the science returns. RRATs show occasional flashes of dispersed radio bursts of typically a few milliseconds duration. The cause of their sporadic emission as well as their connection to other neutron star populations are not fully understood. Detection of a large number of FRBs and RRATs is essential in order for us to gain a better understanding of the nature of these sources. However, detection of such single-pulse events with millisecond duration in dedispersed time-series data is severely hindered by the presence of RFI.

[6] http://www.atnf.csiro.au/people/pulsar/psrcat/
[7] http://astro.phys.wvu.edu/GalacticMSPs/GalacticMSPs.txt
[8] http://www.physics.mcgill.ca/~pulsar/magnetar/main.html
[9] http://astro.phys.wvu.edu/rratalog/
[10] http://www.frbcat.org/

Time-domain surveys are generally sensitivity-limited, hence surveys with more sensitive instruments should lead to a higher discovery rate. Many of the existing as well as future high sensitivity radio telescopes are interferometric arrays. Planned surveys with telescopes, like MeerKAT (e.g., TRAPUM[11]) and SKA Phase1 (e.g., Levin et al. 2017), also need to optimally combine signals from many small telescopes (i.e., do "beam formation"). The GMRT was one of the first interferometric instruments to be systematically used for pulsar searches. The high discovery rate of the GMRT High Resolution Southern Sky (GHRSS[12]; Bhattacharyya et al. 2016; Bhattacharyya 2017), as well as the Fermi-directed survey (Bhattacharyya et al. 2013), demonstrate the capabilities of the GMRT for low-frequency pulsar searches. The recent upgrade of the GMRT allowing much larger instantaneous bandwidths (uGMRT; Gupta et al. 2017) brings a significant increase in its theoretical survey sensitivity for pulsars and FRBs at low and mid radio frequencies. With the uGMRT, Phase2 of the GHRSS survey (Roy 2018) is expected to achieve a sensitivity better than all existing and ongoing off-galactic plane surveys. Most of the existing and planned surveys, however, use one of the two traditionally used methods of beam formation, incoherent array (IA) or phased array (PA) beams, which are described in more detail below. In this paper we explore the possibility of significantly improving the observed time-domain sensitivity using yet another kind of beam formation, post-correlation (PC) beamforming. We show that in this kind of beamforming, the contribution of instrumental red noise to the power spectrum is significantly reduced, thus greatly improving the sensitivity toward low and mid spin frequency pulsars. We also show that PC beam formation can be used to significantly reduce the effect of RFI, thus improving the time-domain sensitivity for periodicity and single-pulse search. Both of these factors lead to reduction of the number of false detections by several orders of magnitude. This not only allows one to lower the candidate detection threshold (i.e., probe fainter flux levels) but also greatly eases the problem of carrying out on-the-fly imaging and other follow-up of these events to maximize the science returns.

2. Beam Formation with Antenna Arrays

2.1. IA and PA Beam Formation

Two commonly used beam formation techniques for antenna arrays are IA beam formation and PA beam formation. For example, the GMRT backends (GSB; Roy et al. 2010 or GWB; Reddy et al. 2017) allow one to form both these type of beams by making per spectral channel combinations of the delay- and fringe-corrected signals from different antennas. The IA beam is formed by summing together the squares of the individual antenna voltages, i.e., it adds together the signal powers. Mathematically

$$P_{\rm IA} = \sum_{i=0}^{N-1} |V_i|^2, \tag{1}$$

where $P_{\rm IA}$ is the IA beam signal and V_i are the voltages from the individual antennas. This kind of combination leads to a wide field of view (but at reduced sensitivity compared to a phased combination) and is useful for blind searches (such as, for e.g., the GHRSS survey). The coherent or PA beam is produced by summing together the voltages (after phasing them appropriately so that the beam points to the direction of interest) and then squaring the resultant sum. Mathematically,

$$P_{\rm PA} = |\sum_{i=0}^{N-1} V_i e^{-i\phi_i}|^2, \tag{2}$$

where $P_{\rm PA}$ is the PA beam signal and V_i are the delay- and fringe-corrected voltages from the individual antennas, and ϕ_i is the phase introduced in antenna i in order to steer the beam toward the desired direction. The PA beam has higher sensitivity than the IA beam. The signal-to-noise ratios (S/Ns) for observations of a single pulse of flux density S located at the pointing center of a dual polarized array for the IA and PA beam are

$$({\rm S/N})_{\rm IA} = \frac{GS\sqrt{2N_{\rm a}\Delta\nu\tau}}{T_{\rm sys}}, \tag{3}$$

$$({\rm S/N})_{\rm PA} = \frac{GSN_{\rm a}\sqrt{2\Delta\nu\tau}}{T_{\rm sys}}, \tag{4}$$

where G is the gain of a single telescope, $N_{\rm a}$ is the number of antennas used for beam formation, $\Delta\nu$ is the instantaneous observing bandwidth, τ is the integration time, and $T_{\rm sys}$ is the total system noise. These expressions assume that the sky noise is small compared to the receiver noise of the antennas. The sensitivities of the IA and PA beams under different scenarios are discussed in detail in Kudale & Chengalur (2017).

The IA beam is not only less sensitive that the PA beam, it is also more vulnerable to instrumental gain fluctuations and RFI. This is because the IA beam is the sum of the auto-correlations of the individual antennas. Since most of the terms in the PA beam correspond to cross correlations between antennas, it has some immunity to RFI (which gets decorrelated by the delay tracking/fringe rotation operations), as well as to fluctuations in the instrumental gains. We illustrate this by showing in Figure 1 the dedispersed time-series for PSR J2144−3933 from simultaneous IA and PA observations using the GMRT. As can be seen, fewer RFI bursts are seen in the PA beam as compared to the IA beam. The PA beam noise properties in general appear better to be than those of the IA beam; one can see individual single pulses in the dedispersed *PRESTO* (Ransom et al. 2002) output, while these pulses are lost in the noise of the IA beam. Still, further improvement in the noise properties can be seen in the PC beam output (the lowest panel in the figure). We discuss this in more detail below.

2.2. Post-correlation Beam Formation

Post-correlation (PC) beam formation (e.g., Kudale & Chengalur 2017), conceptually consists of forming the desired beam not by combining the individual antenna voltages, but rather by combining the (suitably phased) visibilities from the different baselines in the array. Effectively, this eliminates the auto-correlation terms from the PA beam. According to the radiometer Equation (3) and (4), for an array with $N_{\rm a}$ elements, in situations where sky noise is negligible (i.e., $T_{\rm sky} \ll T_{\rm rec}$), the IA beam sensitivity

[11] http://www.trapum.org/

[12] http://www.ncra.tifr.res.in/ncra/research/research-at-ncra-tifr/research-areas/pulsarSurveys/GHRSS

scales as $\sqrt{N_a}$, whereas PA beam sensitivity scales as N_a. Following Equation (29) of Kudale & Chengalur (2017), PC beam sensitivity scales as $\sqrt{N_a(N_a-1)}$. For a telescope like the GMRT with $N_a = 30$, the theoretical degradation of the sensitivity for the PC beam compared to the PA beam is $<2\%$. The reduction in sensitivity arises from the non-inclusion of the N_a auto-correlation terms. However, in practice, since the auto-correlation signals are the ones that are most affected by instrumental gain fluctuations and RFI, one could in fact (as can be seen in Figure 1) get a significant improvement in the S/N when using PC beam formation instead of PA beam formation. In addition, the PC beam is also easier to calibrate. Basically, as far as calibration is concerned, since the PC beam consists only of visibility data, and assuming that visibilities are also computed in parallel (as is the case for the GMRT) the beam can be calibrated using exactly the same techniques as standard interferometric imaging calibration. This holds also for polarimetric calibration. Since both the IA and the PA beams contain auto-correlation terms, proper calibration of these beams involves calibration of the system temperature. We note that calibration of the PC beam could be done in real time, in situations where the visibility data is also output. For example, Kudale & Chengalur 2017) demonstrate for the GMRT that it is possible to apply in real time phases obtained via in-field self calibration to keep the PA beam phased.

Although the name implies that PC beam formation has to be done after correlating the antenna voltages, the beam can in fact be operationally produced in two different ways by (1) subtracting the IA beam from the PA beam; this effectively removes all the auto-correlation data that is contained in the PA beam, or (2) by a phased addition of the cross visibilities. Mathematically, we could do either

$$P_{\rm PC} = |\sum_{i=0}^{N} V_i e^{-i\phi_i}|^2 - \sum_{i=0}^{N} |V_i|^2 \quad (5)$$

when using the antenna voltages directly, or

$$P_{\rm PC} = 2 \times Re\left[\sum_{i=0}^{N-1}\sum_{j=i+1}^{N-1} V_{ij} e^{-i\phi_{ij}}\right] \quad (6)$$

when using the visibilities. Here, $P_{\rm PC}$ is the PC beam signal, V_i and ϕ_i are delay- and fringe- corrected voltages and the beam steering phase of the ith antenna, V_{ij} and ϕ_{ij} are the raw visibility and beam steering phase for the baseline between the ith and jth antenna. We also show in Figure 2, schematic block diagrams of these two ways of forming the PC beams. In Section 4 we compare the computational costs of these two forms of beam formation.

3. Comparison of Different Beam Formation Schemes

We use a number of data sets to compare these different beam formation schemes. The first set of data is based on *SIGPROC*[13] *filterbank* data from uGMRT GWB backend observations in the 300–500 MHz band of the the slow ($P \sim 8.5$ s) pulsar PSR J2144−3933. Data from the IA and PA beams formed in real time using the GWB were recorded, and the PC beam was formed offline using the difference between the PA and IA beams as described above (Equation (5)). Simulated pulsar signals were injected into the IA and PA *filterbank* data files using the *inject_pulsar* routine of the *SIGPROC* pulsar package. A total of 12 data sets (each for IA, PA, and PA−IA) were generated in this way, where the difference between the data sets is the period of the injected pulsar signal, this varies from 25 ms to 128 s. The original data also of course contains the signal for PSR 2144−3933, making for a total of 13 data sets from the uGMRT data. In addition we also used data from GMRT GSB backend observations at 607 MHz of PSR J2144−3933. The Nyquist-sampled antenna voltages were recorded on disk, and all beam formation as well as correlation was done offline. These data sets allow us to compare the performance of the different beamforming schemes with the exact same input data. The GSB data set also allows us to compare the two different ways of PC beam formation discussed above.

In Figure 3 we show the low-frequency end of the power spectra after de-dispersion for the three different beamforming modes using the uGMRT data. As can be seen, the power spectra for the PA and IA beams are essentially the same, since, as mentioned above, this part of the power spectrum is dominated by the instrumental red noise and the RFI that is contained in the auto-correlation spectra. Consistent with this, the PC beam, which does not contain auto-correlation data, has significantly lower noise. This "de-reddening" of the power spectrum should greatly ease the problem of detecting slow pulsars. Indeed, one can see that in the PC beam, the signal from the 8.5 s pulsar is detectable from the first harmonic onward. For the IA or PA beam on the other hand, only harmonics beyond the ∼60th harmonic are visible in the power spectrum. Figure 4 shows the folded profiles of PSR J2144−3933 for these IA, PA, and PC beam data. A systematic and significant improvement in the S/N is clearly visible even to the eye as one goes from the IA beam to the PA beam and PC beam. The PC beam's S/N is ∼5–6 times better than that of the PA beam. This clearly shows the dramatic improvement in the detectability of slow pulsars when the noise and systematics contained in the auto-correlation spectra are eliminated. We note that the beams were formed using all of the input data, i.e., there has been no effort at RFI mitigation. We discuss below specific advantages that the PC beam offers as far as targeted removal of RFI is concerned.

In Figure 5 we show the ratio of the S/Ns of the PC and PA beams as a function of the pulse period. This plot was generated using the data for the simulated pulsars as well as the data for PSR J2144−3933. In all cases the PC beam has a higher S/N than the PA beam. The PC beam S/N is about 10% better than that of the PA beam for a spin period of 25 ms; this difference reaches factors of 5−6 for spin periods of ∼10 s. Beyond spin periods of ∼10 s, the increase in the S/N is not as large, but it is still as much as a factor of ∼3 for spin periods as long as 100 s, (i.e., spin frequency ∼0.01 Hz). This is due to the fact that red noise in the PC beam also goes up below 0.1 Hz, as can be seen in Figure 3.

The two methods of PC beam formation presented in Section 2.2 are mathematically equivalent. One might imagine then that all that distinguishes these two methods is their respective computational costs. We discuss this issue in Section 4 below. However, there is one further way in which these two methods are different: in the possibilities that they offer for identifying and removing RFI. When the PC beam is formed as the PA−IA beam, one can only flag out data at the

[13] https://github.com/SixByNine/sigproc/

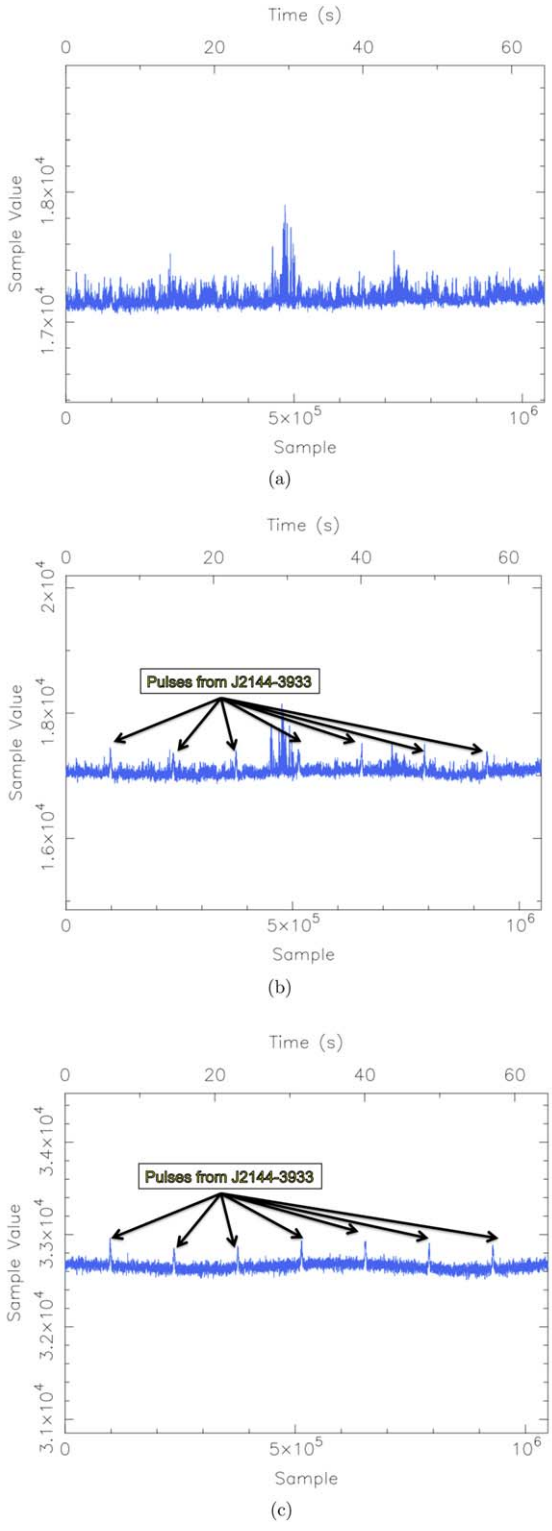

Figure 1. Dedispersed time-series for PSR J2144–3933 from simultaneous observations of the (a) IA beam and (b) PA beam are shown here. The plots were generated using the *PRESTO* software tools. The plot shows the mean value computed using moving average of 8 time samples. The *y*-axis scale is different for the different panels. Fewer RFI bursts are seen in the PA beam compared to the IA beam. Also, individual single pulses (as marked in the plot) are visible for the PA beam while these pulses are lost in the noise of the IA beam. (c) Dedispersed time-series for PSR J2144−3933 from the post-correlation beam. As can be seen, there is a significant improvement in the immunity against RFI, and the individual pulses can be clearly seen.

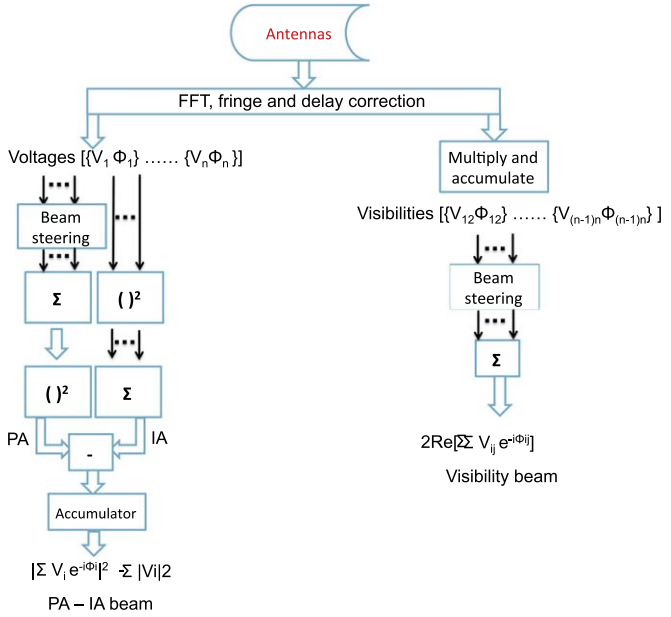

Figure 2. Schematic block diagram for post-correlation beam formation. The left branch shows PA–IA beam formation and the right branch shows beam formation using visibilities. Note that in PA–IA beam formation the beam steering has to be done at the FFT block level, while in visibility-based beam formation the beam steering is done after accumulation.

granularity of an antenna. When forming the PC beam from the visibilities, one can flag out data at the granularity of baselines. This is particularly useful in arrays that contain antennas at a range of separations. Often data from the short baselines contain significantly more RFI than the data from long baselines. Since nearby antennas also have baselines with more distant antennas, this could allow one to greatly eliminate the RFI while retaining much of the raw sensitivity. We show in Figure 6 that short timescale (i.e., few seconds) RFI bursts present in the PC (visibility based) beam can be removed by flagging out the data from all the baselines shorter than ∼450 meters (i.e., 3%–4% of the total GMRT baselines). As shown in the figure, these RFI bursts generate pseudo pulse-like features in the folded profile of PSR J2144−3933; flagging the short baseline very effectively mitigates the problem. We note that the flagging done here was "blind," i.e., short baselines were flagged, without looking at the data quality on these baselines. In principle one could use flagging algorithms (such as FLAGCAL, Prasad & Chengalur 2012) to automatically identify and flag only those baselines that actually do have RFI.

So far we have been comparing the characteristics of the IA, PA, and PC beams in relation to detecting pulsars. Another class of pulsed signals that is of great interest currently are transients such as FRBs, which emit single pulses. While observing with an interferometric array, one can save the visibilities for candidate events, so that one can also image the field in order to localize any confirmed sources (see, e.g., Bhat et al. 2013). As discussed in detail in Bhat et al. (2013), in such searches, it is important to reduce false positives as much as possible, in order to minimize the amount of data that have to be saved and processed. Since the PC beam contains far less RFI than the IA and PA beam, one would expect that the number of false positives in the PC beam would also be less than that for the other beamforming modes. Figure 7 shows the

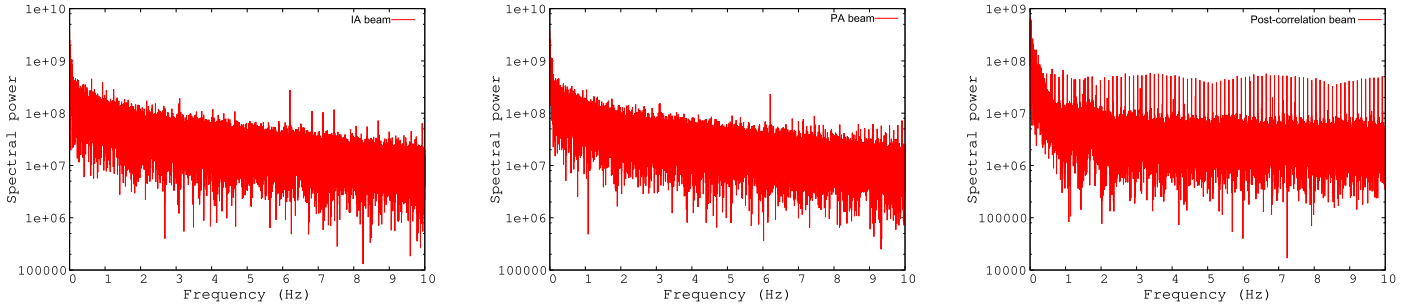

Figure 3. Power spectra for the IA, PA, and PC beams for PSR J2144−3933. The harmonics of the pulsar are spaced at 0.12 Hz. As can be seen, the power spectra for the PA and IA beams are essentially the same. However, there is an order of magnitude reduction in red noise for the PC beam. This enables the detection of low-order harmonics, which are completely buried in the noise for the IA and PA beams. See the text for more details

number of candidates detected from IA, PA, and PC beams for simulated FRB events with various DMs injected in the same uGMRT 300–500 MHz band data as discussed above. The signals were injected using the same *inject_pulsar* routine, but with the pulse period being much larger than the duration of the data (i.e., 60 s). The PC beam is formed as PA−IA. There are 8 FRB events simulated at DMs of 10, 20, 50, 100, 200, 500, 1000 and 2000 pc cm^{-3}. *PRESTO*-based single-pulse searches were performed for all three beams over a range of DMs (indicated by the error bars). As can be seen from the figure, the number of triggers from the PC beam is almost two orders of magnitude lower than that of the IA beam, even at a DM as high as 2000 pc cm^{-3}. The number of false positives in the PC beam data is also a factor of ∼5 less than that found in the PA beam data. Interestingly, over the full FRB DM search space (i.e., 250–2600 pc cm^{-3}), the candidate detection rate is almost constant for the PC beam. The percentage of true positives at the highest (2000 pc cm^{-3}) DM values of the simulations for IA, PA, and PC beams are 0.0004, 0.8, and 5, respectively. We note that this plot was generated for candidates detected above a threshold of 5σ in order to make the uGMRT 300–500 MHz PC beam sensitive enough to detect all the known FRBs (ignoring frequency dependent scattering and spectral steepening). At this threshold the uGMRT IA beam detects only 30% of the known FRBs. Raising the threshold to 10σ generates very few triggers from the PC beam for this data, whereas the IA beam continues to be equally corrupted. The recently detected FRBs are all at the lower end of the FRB flux distribution; all of these will be completely missed at the sensitivity offered by the IA beam. As is the case for FRBs, the PC beam data would also contain far fewer false positives in searches for other transients such as RRATs (most of which have DM $<$ 300 pc cm^{-3}). Both manual as well as automated searches for RRATs in the IA beam data would be swamped by the large number of false positives. Ways of overcoming this problem by forming multiple incoherent sub-array beams and using coincidence filtering are discussed in Bhat et al. (2013). However, splitting antennas in sub-arrays significantly reduces the survey sensitivity. Another major difference between the IA and PC beam is of course the field of view. In blind surveys one would like to have as large a field of view as possible, in which case PC beam formation is not competitive, unless one is able to form multiple beams. In the next section we detail the computational cost involved in forming multiple PC beams.

4. Computational Requirements

As discussed above, there are two different ways of forming the PC beam. The first is via the difference of the PA and IA beams, while the second is via a phased addition of the visibilities. While the PC beams formed in these two ways is mathematically equivalent, we also saw that operationally the visibility route might have some advantage because of the better opportunities it provides for flagging data affected by RFI. Here, we take a look at the difference in the amount of computation required to make the PC beam in these two ways. To start with, we note that correlators require a fan-out of the data, i.e., in order to correlate the data from one antenna with all other antennas, one needs multiple copies of the data stream. On the other hand beam formation operates on the data stream from each antenna independently, except in the final addition stage. This would lead to differences in architecture. Here, we do not look at this in detail, but instead focus only on the number of computations required to make the PC beam in these two different ways.

We start by defining the parameters needed to determine the required computation, with the assumed value of the parameter for the GMRT (where relevant) given in parenthesis. Let the total number of beams to be formed (each with an independent phase center) be N_B. The total number of antennas is $N_{\rm a}$ (30), the bandwidth of operation is B (200 MHz), the number of time samples in a given FFT block is N_f (4096 for 2048 spectral channels), and the number of FFT blocks per integration is N_b. In terms of these parameters the channel resolution is $\Delta\nu = \frac{B}{N_f}$ and the integration time $\Delta\tau = \frac{N_b N_f}{2B}$.

The total computational load (in number of operations per second) for PA−IA beam formation for one integration is

$$5N_{\rm a}N_bN_f\log N_f + N_BN_{\rm a}N_bN_f + N_BN_f(N_{\rm a}N_b + N_b - 1) + N_bN_{\rm a}N_f + (N_{\rm a} - 1)(N_b - 1)N_f + N_BN_f, \quad (7)$$

and consists of the following components:

1. $5N_aN_bN_f\log N_f$ for FFT.
2. $N_BN_aN_bN_f$ for fringe and fractional-delay corrections, as well as beam steering.
3. $N_BN_f(N_aN_b + N_b - 1)$ for PA beam formation, including addition, squaring, and integration.
4. $N_bN_aN_f + (N_a - 1)(N_b - 1)N_f$ for IA beam formation, including squaring, addition, and integration.
5. N_BN_f for the PA–IA operation.

We note that for PA−IA beam formation the phase corrections for beam steering need to done before antenna addition, which requires working at the FFT resolution. However, since the maximum fringe rate of the GMRT is ±5 Hz (Chengalur 1998), the maximum possible delay change even over a period as large as 1 ms is much smaller

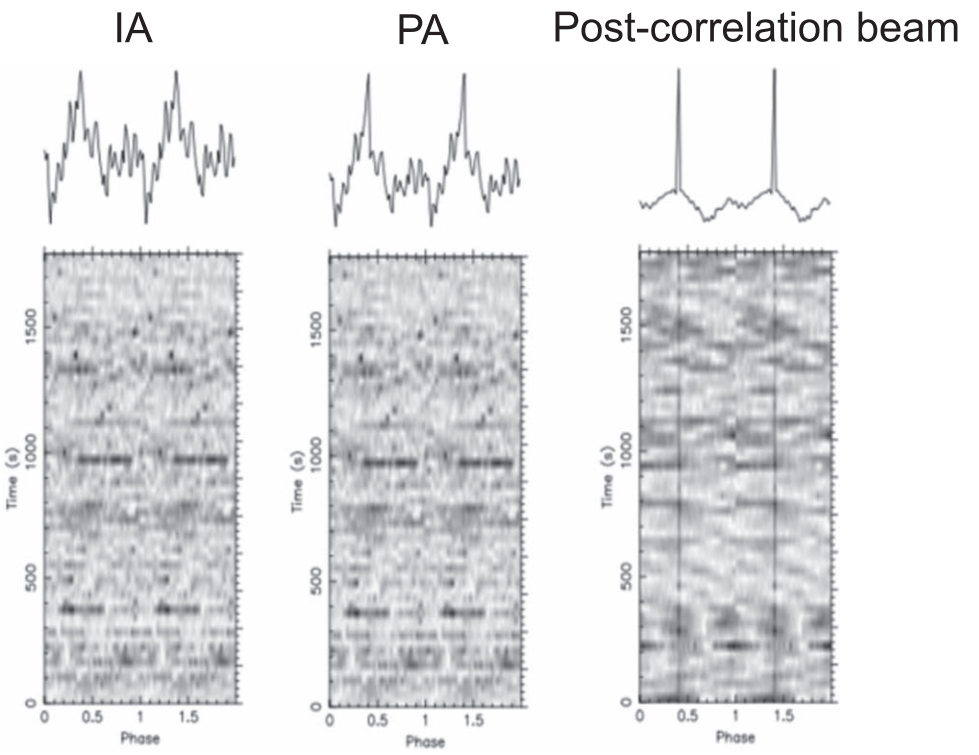

Figure 4. Folded profiles of PSR J2144−3933 observed with the uGMRT 300−500 MHz bands.

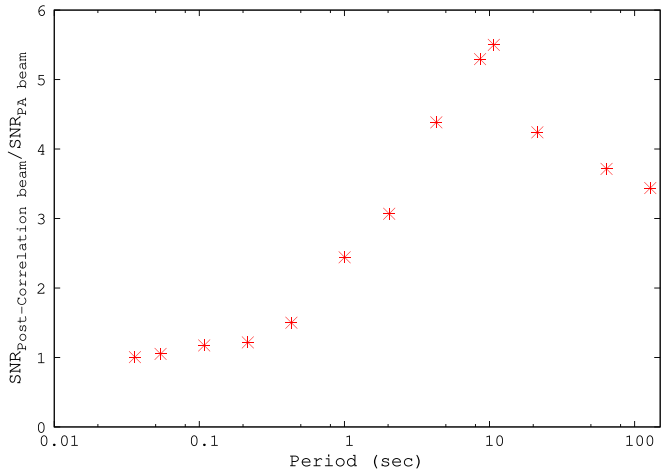

Figure 5. Sensitivity improvement with a post-correlation beam, compared to the PA beam, as a function of pulse period.

than the Nyquist-sampling resolution. This means that for PC beam formation from the visibilities, we can perform the differential beam steering after the visibilities have been computed. The total computation load (in number of operations per second) for visibility-based PC beam formation for one integration is

$$5N_aN_bN_f\log N_f + N_aN_bN_f + N_bN_f\frac{N_a(N_a-1)}{2} + (N_b-1)N_f\frac{N_a(N_a-1)}{2} + N_BN_f\frac{N_a(N_a-1)}{2} + N_BN_f\left[\frac{N_a(N_a-1)}{2}-1\right] \quad (8)$$

and consists of the following components:

1. $5N_aN_bN_f\log N_f$ for FFT.
2. $N_aN_bN_f$ for fringe and fractional-delay corrections at the pointing center (common to all beams).
3. $N_bN_f\frac{N_a(N_a-1)}{2} + (N_b-1)N_f\frac{N_a(N_a-1)}{2}$ for correlation, including multiplications and additions.
4. $N_BN_f\frac{N_a(N_a-1)}{2}$ for phase corrections required for steering the individual beams.
5. $N_BN_f[\frac{N_a(N_a-1)}{2}-1]$ for visibility addition for the beam formation.

For the given GMRT configurations with 1600 beams, in PA–IA-based PC beam formation (i.e., Equation (7)), terms 2 (fringe, fractional delay, and beam steering) and 3 (PA) dominate equally and are at least 20 times higher than any other terms. Whereas for visibility-based PC beam formation (i.e., Equation (8)), terms 4 (beam steering) and 5 (visibility addition) dominate, but they are only an order of magnitude higher than the next most dominant term, term 2 (FFT). However, the contributions of terms 4 and 5 of Equation (8) increase manyfold compared to the other terms, as beams are formed at high time resolutions. Considering these, one would expect that PA−IA beam formation is computationally cheaper for a small number of high time resolution beams, while the visibility-based beam formation is computationally cheaper for a large number of beams at low time resolution. The crossover point would depend on the total number of elements. We show in Figure 8 a comparison of these two computational loads as a function of the total number of beams formed, and the time resolution for a GMRT-like array of 30 antennas (upper panel), as well as an SKA Phase1 Mid like array of 256 antennas. For the GMRT array we use 200 MHz instantaneous bandwidth with 2048 spectral channels at 163.84 μs (upper right) and 1.31 ms (upper left) time resolution. For the SKA Phase1 Mid array (Levin et al. 2017) we use 300 MHz instantaneous bandwidth with 4096 spectral channels at 64 μs (lower right)

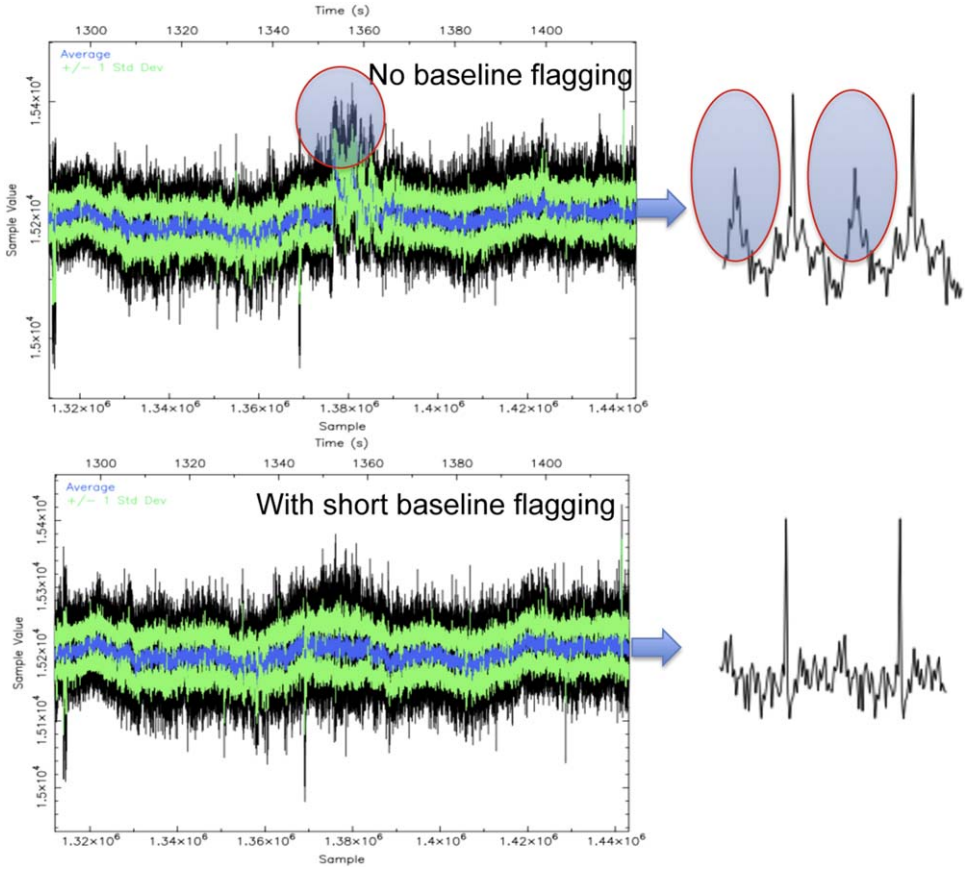

Figure 6. Visibility-based post-correlation beam formation. The upper panel shows the time-series when one uses all the data, while the lower panel shows the time-series when the short baselines (which tend to be those most affected by RFI) are flagged. The plots on the side show the corresponding folded profiles. A dramatic decrease in the systematics can be seen when one flags the baselines with RFI. See the text for more details.

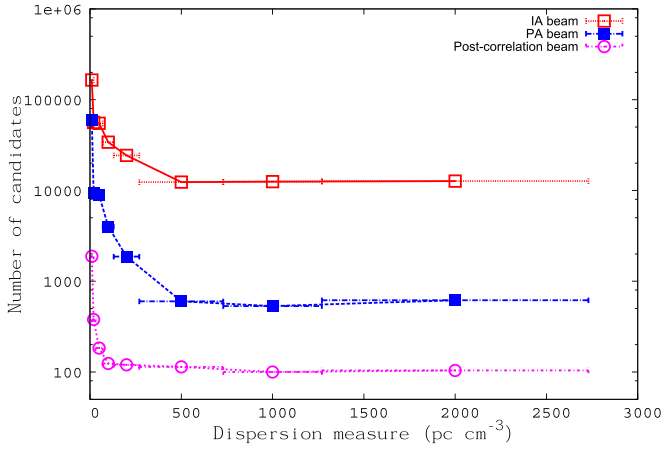

Figure 7. Number of candidates detected from IA, PA, and post-correlation beams for simulated FRB events with various DMs for data of 60 s duration. There are 8 FRB events simulated at DMs of 10, 20, 50, 100, 200, 500, 1000, and 2000 pc cm^{-3}. The DM ranges used in each step for searches are shown by horizontal bars around each central DM. This detection rate is for a threshold of 5σ. The number of (false) detections from the IA beam is almost two orders of magnitude higher than the post-correlation beam, even at a DM of 2000 pc cm^{-3}.

and 2.048 ms (lower left) time resolution. The figures clearly bring out the broad trends expected for time resolution and number of beams. For the GMRT, visibility beam formation is economical compared to PA−IA beam formation for time resolutions $\geqslant$163.84 μs and for $\geqslant$10 beams. For the SKA Phase1 Mid array, visibility beam formation is economical compared to PA−IA beam formation for time resolutions $\geqslant$2.048 ms and for $\geqslant$800 beams. A configuration with a small number of high-time resolution beams would be useful in searches for pulsars (especially MSPs) via targeted observations of globular clusters (GCs). GCs are the most likely hosts of exotic binary systems, like MSP main-sequence binaries, highly eccentric binaries, MSPs in evolutionary phases like Redback and Black Widow, and MSP black hole binaries, which may not form via normal stellar evolution in the disk. The multiple beams should be sufficient to cover the expected sky area within which MSPs are expelled from the center but that are still within the cluster tidal radius. A moderate number of high time resolution beams offers an opportunity to greatly increase the pulsar timing efficiency in arrays where the individual elements have a large field of view, by allowing simultaneous observations of multiple pulsars (Stappers et al. 2018). A large number of lower time resolution beams (as would be cheaper via the visibility route) would be useful in blind searches for all but the fastest pulsars.

5. Case Study of a Proposed GMRT Survey

The improvements seen in time-domain processing using PC beam formation, aided by the enhanced sensitivity of the uGMRT for the GHRSS Phase2 survey, provide the motivation to develop a time-domain survey with a PC beamformer. We compute here the estimated parameters for such a survey. As a benchmark, we consider the uGMRT 300–500 MHz band with

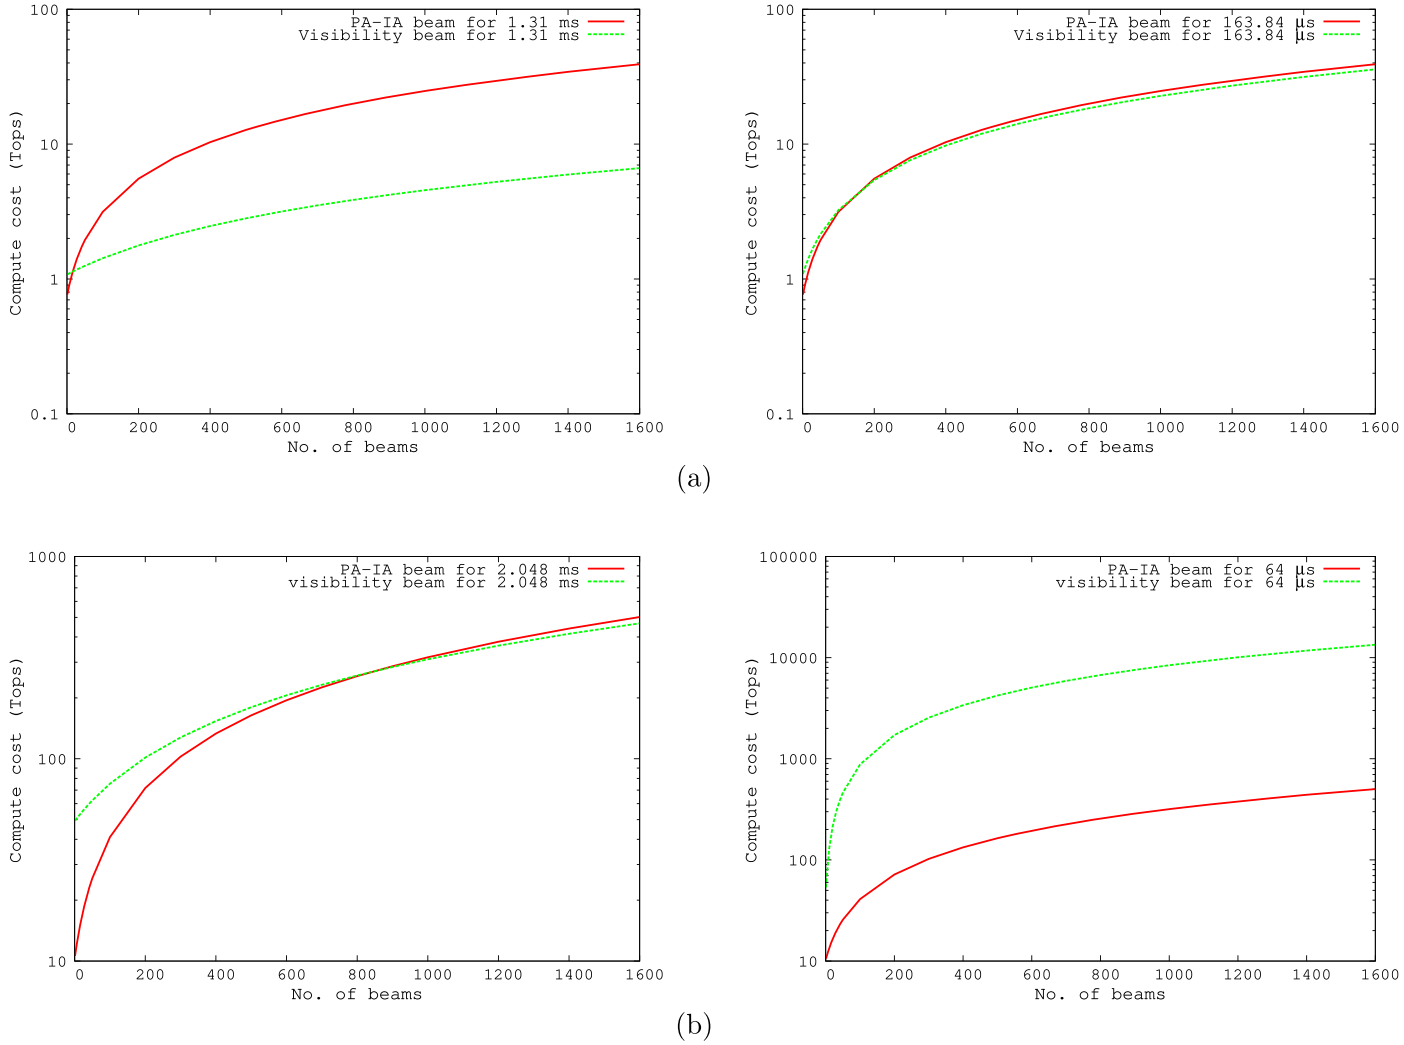

Figure 8. Comparison of compute costs for two ways of forming post-correlation beams for the GMRT (a) and the SKA Phase1 Mid array (b). For the GMRT, a 1.31 ms (left panel) and 163.84 μs (right panel) output time resolution is plotted, while for SKA Phase1 Mid, a 2 ms (left panel) and 64 μs (right panel) out time resolution is plotted. PA–IA beams are marked in red and visibility beams are marked in green. The costs for both beam formation modes are similar: 163.84 μs for the GMRT and 2 ms for the SKA Phase1 Mid.

30 antennas, 200 MHz bandwidth, 2048 spectral channels, and visibility beam formation at a 1 ms time resolution, with about 128 beams covering a $\sim 10'$ field of view. We note that covering the entire field of view with PC beams would require $\sim$1600 beams. We estimate a survey sensitivity of $\sim$0.1 mJy at 400 MHz (considering the radiometer Equation (4) with a 2% loss for ignoring the auto-correlation power), for a 10σ detection for a 10% duty cycle, a PC beam gain of 7 K Jy^{-1} for 200 MHz bandwidth, 10 minutes of dwell time, and a system temperature of 106 K. We also calculate a sensitivity of 0.05 Jy as the 5σ detection limit for 5 ms transient millisecond bursts, which would correspond to weak scattering (Thornton et al. 2013).

Figure 9 shows the components required for such a time-domain survey, with the PC beamformer as specified above. The required components are shown in four different colors. Visibilities computed in the uGMRT backend (GWB; marked in blue) at 1 ms time resolution are transferred to the the PC beamformer nodes (marked in orange) with an aggregate data rate of $\sim$3 GB s^{-1}. We aim to implement in-field phasing (Kudale & Chengalur 2017) using a sky model derived from

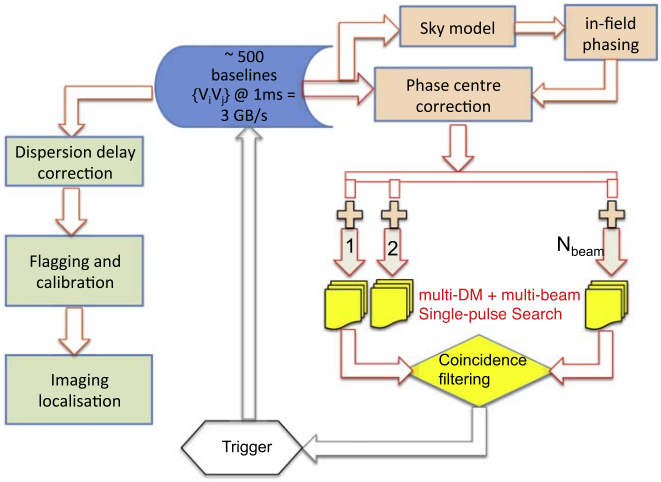

Figure 9. Proposed multibeam time-domain survey with visibility beam formation. There are four functional modules, each colored differently and running on different compute hardware. Please refer to the text for further details.

the time-averaged visibilities in order to improve the coherence in phasing up to a baseline length of several kilometers. This optimizes the GMRT PA sensitivity beyond a central compact core (most current PA observations use only the antennas in the central square). In addition to deriving the phasing model, a baseline-based flag masking the bad baselines will also be generated in real time from these time-averaged visibilities. Coherent additions of these visibilities will result in 128 such visibility beams. The multi-DM search for single pulses (colored in yellow) on each of these visibility beams would need to be executed on a separate FRB cluster, followed by coincidence filtering to remove spurious events (Bhat et al. 2013). We also propose recording these 128 beams with a 1 ms time resolution, giving a total data rate of 200 MB s^{-1} into a disk for a quasi-real-time search for pulsars using the same cluster. We note that the proposed 1 ms time resolution is sufficient to detect double neutron star systems, young pulsars, and normal pulsars, as well as objects like radio magnetars. Visibility buffers corresponding to candidate single-pulse events will be recorded at a 1 ms time resolution covering the full DM sweep time-range. For an event at a DM of 2000 pc cm^{-3}, the total DM sweep time over 200 MHz band in uGMRT 300–500 MHz band is ∼50 s, which results in a 40 GB buffer size on each of the PC beam nodes. This means one can easily hold few buffers for accommodating the pipeline delay and flush them into storage based on the real-time triggers. These visibilities will be processed through the processing blocks (marked in green) for millisecond imaging localization at quasi-real-time. This block includes removal of dispersion delay, followed by a flagging and calibration pipeline and snapshot imaging. We note that part of this imaging pipeline for localizing time-domain events has already been demonstrated for the GHRSS Phase1 survey (Bhattacharyya et al. 2016).

6. Summary

In this paper, we demonstrate that use of a PC beamformer for a radio interferometric array results in a manyfold increase of the detection significance of time-domain events, compared to the conventional incoherent and coherent array beamformer. This increase in sensitivity is driven by the lower red noise and RFI contamination of the PC beam. Post-correlation beam formation also allows one to use standard interferometric calibration techniques for calibrating the beam. We compare two different modes of PC beam formation: (1) PA–IA beam formation, which does not require computation of the visibilities; and (2) visibility beam formation, where the beam is formed from the computed visibilities. We also show that the PA–IA beam formation is computationally economical for a small number of high time resolution beams. At low time resolutions, the visibility-based beam formation is computationally cheaper. Visibility-based beam formation also allows for better control in flagging/suppressing RFI. For a multi-element feed system (e.g., Parkes multibeam system) or for a PA feed, the PC beam can also be used to subtract RFIs (correlated within the feed elements) from a feed element response (Kocz et al. 2010). These new beamforming techniques could significantly improve the sensitivity of time-domain studies with both existing (e.g., uGMRT, JVLA) and upcoming (e.g., CHIME, Amiri et al. 2018; OWFA, Subrahmanya et al. 2017) radio interferometric arrays. As a specific example, we have presented a proposed time-domain survey with the uGMRT.

We thank the computer groups at GMRT and NCRA. We thank Mr. Harshavardhan Reddy of GMRT for insightful discussions on the data-rate issues in the GWB. We thank Mr. Sanjay Kudale of GMRT for trying out in-field phasing in generating post-correlation beams. Ue-Li Pen acknowledges NSERC for support. We also acknowledge the support of telescope operators during our test observations, which also involves data intensive baseband recording observations. The GMRT is run by the National Centre for Radio Astrophysics of the Tata Institute of Fundamental Research.

ORCID iDs

Jayanta Roy https://orcid.org/0000-0002-2892-8025
Jayaram N. Chengalur https://orcid.org/0000-0002-0269-1154

References

Amiri, M., Bandura, K., Berger, P., et al. 2018, ApJ, 863, 48
Bhat, N. D. R., Chengalur, J. N., Cox, P. J., et al. 2013, ApJS, 206, 2
Bhattacharyya, B. 2017, in Proc. IAU Symp. 337, Pulsar Astrophysics the Next Fifty Years, ed. P. Weltevrede et al. (Cambridge: Cambridge Univ. Press), 17
Bhattacharyya, B., Cooper, S., Malenta, M., et al. 2016, ApJ, 817, 130
Bhattacharyya, B., Roy, J., Ray, P. S., et al. 2013, ApJL, 773, 12
Chengalur, J. N. 1998, Delay Tracking and Fringe Stopping for the GMRT Correlator, NCRA Tec. Rep., http://www.ncra.tifr.res.in/library/
Gupta, Y., Ajithkumar, B., Kale, H. S., et al. 2017, CSci, 113, 4
Kocz, J., Briggs, F. H., & Reynolds, J. 2010, AJ, 140, 2086
Kudale, S., & Chengalur, J. N. 2017, ExA, 44, 97
Levin, L., Armour, W., & Baffa, C. 2017, in Proc. IAU Symp. 337, Pulsar Astrophysics the Next Fifty Years, ed. P. Weltevrede et al. (Cambridge: Cambridge Univ. Press), 171
Lorimer, D. R. 2008, LRR, 11, 8
Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Sci, 318, 777
Lorimer, D. R., & Kramer, M. 2004, Handbook of Pulsar Astronomy, Vol. 4 (Cambridge: Cambridge Univ. Press)
McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, Natur, 439, 817
Prasad, J., & Chengalur, J. N. 2012, ExA, 33, 157
Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, AJ, 124, 1788
Reddy, S. H., et al. 2017, JAI, 6, 1641011-336
Roy, J. 2013, PhD thesis, National Centre for Radio Astrophysics, Tata Institute of Fundamental Research
Roy, J. 2018, in Proc. IAU Symp. 337, Pulsar Astrophysics the Next Fifty Years, ed. P. Weltevrede et al. (Cambridge: Cambridge Univ. Press), 183
Roy, J., Gupta, Y., Pen, U.-L., et al. 2010, ExA, 28, 55
Roy, J., Ray, P. S., Bhattacharyya, B., et al. 2015, ApJL, 800, 12
Stappers, B. W., Keane, E. F., Kramer, M., Possenti, A., & Stairs, I. H. 2018, RSPTA, 376, 20170293
Subrahmanya, C. R., Manoharan, P. K., & Chengalur, J. N. 2017, JApA, 38, 10S
Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, ApJ, 846, 170
Thornton, D., Stappers, B., Bailes, M., et al. 2013, Sci, 341, 53