Phased Array Spectral Voltage mode of the GMRT Wideband Backend

Version 1.0

March 5, 2022

Visweshwar Ram Marthi, Sanjay Kudale & Nilesh Raskar

NCRA • TIFR

NATIONAL CENTRE FOR RADIO ASTROPHYSICS
TATA INSTITUTE OF FUNDAMENTAL RESEARCH
Post Bag 3, Ganeshkhind, Pune - 411 007.

Contents

1 Background

The GMRT Wideband Backend (GWB) provides a beamformer mode of operation, used for time-domain astronomy. Typically, most observations of pulsars and fast radio bursts (FRBs) utilize the beamformer. The beamformer combines the signals from all the antennas selected in the GMRT Array Combiner (GAC) in one of three modes available: (a) the phased array (PA) mode where the voltages from the antennas are added after applying the relative phase offset corrections, (b) the incoherent array (IA) mode where the voltages from the antennas are first detected (i.e. squared and converted to power) before being co-added and (c) a coherent dedispersion pipeline (CDP), which is identical to the PA mode except that the voltage in each channel is coherently dedispersed before detection. Each of these three modes has its own utility and are sometimes used concurrently.

This document describes the implementation of the new Phased Array Spectral Voltage (PASV) mode of the beamformer and spells out the Standard Operating Procedure (SOP) for running this mode during astronomical observations.

2 The PASV Mode

The GWB FX correlator takes the delay-corrected antenna voltage timestreams of the two polarizations and produces a complex voltage spectrum every $2Nt_s$ s via a $2N$-point Fast Fourier Transform (FFT), where N is the number of channels and t_s is the Nyquist sampling time. The FX correlator picks up the antenna voltage spectra, applies the fringe correction and the instrument phase solutions. These spectra are now available for the X part of the FX correlator, namely the cross-multiply-accumulate (XMAC) engine. The interferometer, PA and IA beamformer kernels then run independently. This intermediate product is available in a shared memory location for the PA beamformer for further processing, like co-adding the voltage spectra, FFT integration and square-law detection. The PASV kernel has been introduced at this point, where it accesses the co-added voltage spectra from the shared memory and writes the two polarizations, one each to a disk as .raw files in a time-frequency channel (TF) order. Figure 1 shows a high-level schematic view of the GWB FX correlator and the feeding points for the interferometer as well as the PA and IA beamformers.

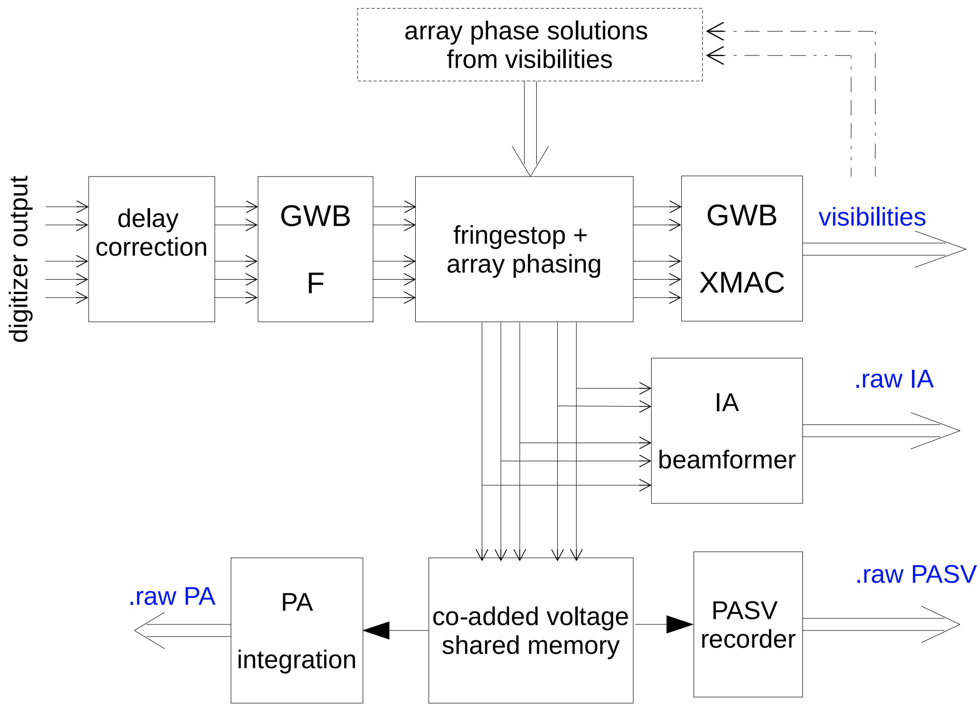

Figure 1: A high-level schematic view of the architecture of the GWB beamformer and the PASV block in relation to it. The CD pipeline is not shown here in the interest of clarity.

3 Available Modes

The GWB FX correlator can be run in 100-MHz, 200-MHz or 400-MHz mode: this is necessitated by the fact that in all bands other than Band-5, a choice of 100 MHz or 200 MHz of RF bandwidth is available. In band-5, a 400-MHz wide band is accessible in addition to the 200-MHz bandwidth and the sub-band mode which gives a choice of four sub-bands with a 120-MHz bandwidth. The choice of the bandwidth dictates the sampling rate and hence the configuration of the GWB FX correlator. At present, PASV recording in the 100-MHz , 200-MHz and 400-MHz modes of the GWB correlator have been tested, with additional flexibility in the final bandwidth being provided by the Digital Down Converter (DDC) in the 100-MHz mode. Specifically, it should be noted that within the 100-MHz band, DDC can be combined with the DDC Local Oscillator (DDCLO) to mimic the narrowband mode available for spectral line observations. This can be used with downsampling factors of 2 onwards. At the time of writing this document, the PASV mode has not been tested explicitly in the 120-MHz sub-band mode in Band-5.

The modes currently tested and released for astronomical observations are enumerated below:

1. 400-MHz/4-bit
2. 200-MHz/8-bit
3. 200-MHz/4-bit
4. 100-MHz/8-bit
5. 100-MHz/4-bit
6. 50-MHz/8-bit
7. 50-MHz/4-bit
8. 25-MHz/8-bit
9. 25-MHz/4-bit

	25-MHz	50-MHz	100-MHz	200-MHz	400-MHz
4-bit	✓	✓	✓	✓	✓
8-bit	✓	✓	✓	✓	X

Table 1: Bandwidth and bit-depth options available for PASV recording.

The 400-MHz, 200-MHz and 100-MHz modes are standalone modes derived from the respective GWB settings, while the 50- and 25-MHz modes are obtained from the 100-MHz mode through DDC. Besides, it is also possible to deploy the DDCLO to shift the target band by an arbitrary amount within the RF bandwidth.

The 400-MHz/4-bit and the 200-MHz/8-bit modes are slightly different in their implementation detail from the rest: the recording program puts out the individual polarizations streams on two different nodes. In terms of the node-level throughput, they are identical to the 2 100-MHz/8-bit or 2 200-MHz/4-bit polarization streams going to two disks on the same host node. At the time of writing the current version of this document, the 400-MHz/4-bit and 200-MHz/8-bit modes are not part of the Status Document of GTAC-42 Announcement of Opportunity: they have been subsequently implemented and tested.

4 Implementation

The PASV mode data are available in a shared memory location for the PA beamformer kernel. This shared memory location is accessed through a standalone program, written using the C language. The shared memory block size is determined through various parameters of the GPU header: typically is it 64/128/256 MB. One MB here is exactly 1024 × 1024 bytes, i.e. 1048576 bytes. Since the FFT length is always an integer power of 2, it is always guaranteed that there are always an exact integer number of spectra in each shared memory block region. The shared memory blocks are polarization-interleaved, i.e. two adjacent shared memory blocks together constitute data with identical timestamps from the two polarizations. These two blocks are processed by two independent threads and written to two independent disks. While the 25-MHz and 50-MHz modes can both be written to a system of RAID disks, the higher throughput of the 100- and 200-MHz modes constrains the data to be recorded on non-volatile memory express (NVMe) devices. The current implementation loops through four shared memory blocks in a round-robin fashion to provide for sufficient latency between the write-to-disk and fill-shared-memory-buffer operations.

The implementation of the PASV mode undertakes no manipulation of the data at any stage. It merely reads the data from the shared memory, ascribes to each block a verified timestamp and writes it to disk. The statistics of the data are therefore completely attributable solely to the upstream settings such as the RF and LO system, the digitizer and the GWB. For brevity and consistency, the statistics of the data are not described here.

5 Auxiliary Files

5.1 Timestamp information

The recording program, in addition, also saves in an auxiliary binary file a block-level timestamp, referenced to the starting time boundary of each block: the timestamp refers to the first sample in each shared memory block identically for both the polarizations. The binary timestamp file is saved with a `.ts` extension. The structure that holds the timestamp is defined as

```
typedef struct
{ struct timeval timestamp;
  double blk_nano;
} BeamTimeType;
```

The `timeval` structure holds the `UNIX` time in seconds and microseconds, and the `blk_nano` element holds the nanoseconds elapsed since the last microsecond.

	400-MHz	200-MHz	100-MHz	50-MHz	25-MHz
8 bits	NA	3200 Mbps	1600 Mbps	800 Mbps	400 Mpbs
4 bits	3200 Mbps	1600 Mbps	800 Mbps	400 Mbps	200 Mbps

Table 2: Data throughput numbers for the different modes. Modes not implemented are marked as `NA`.

5.2 Record log

The recording program also writes a `.log` file. The `.log` file is an ASCII file with information on the current record from the correlator and the record number being accessed in the shared memory block. In addition, it gives information on the association between the correlator record number and the shared memory block identifier.

6 Data Format and Throughput

The two polarizations go into a file each. For each polarization:

```
Loop : Time
{ Loop : Channel
  { Real | Imaginary }
}
```

6.1 8-bit

The channel-wise real and imaginary parts of the complex spectrum for each polarization are interleaved, and the spectra are time ordered.

```
RR : [[real-8bits | imaginary-8bits] X NCHAN ] X NTIME
LL : [[real-8bits | imaginary-8bits] X NCHAN ] X NTIME
```

6.2 4-bit

The channel-wise real and imaginary parts of the complex spectrum for each polarization are packed into a single byte: the upper nibble represents the real part (extracted from the 4 MSbits) and the lower nibble represents the imaginary part (extracted from the 4 MSbits). The spectra are time ordered.

```
RR : [[real-4MSbits-nibble 1 | imaginary-4MSbits-nibble 2] X NCHAN ] X NTIME
LL : [[real-4MSbits-nibble 1 | imaginary-4MSbits-nibble 2] X NCHAN ] X NTIME
```

6.3 Throughput

The data throughput per polarization is a pertinent quantity as it has a bearing upon both the volume of the data as well as the duration for which such data can be recorded. The data throughput can be calculated for a given mode configuration, based on the parameters set for the GWB FX correlator that can be found in `GPU.hdr`, as well as the bit depth of the final PASV file. Given
B - bandwidth in MHz
b - bit depth
the data throughput per polarization is calculated as $r = 2Bb$ million bits per second (Mbps). For example, the 100-MHz, 8-bit mode produces data at the rate of $r = 1600$ Mbps or 200 million bytes per second. Currently there is a policy limit of ~ 60 minutes of data at $r = 1600$ Mbps for GTAC and other astronomical observations, but this will likely see an upward revision as more NVMe device capacity becomes available on the beam host nodes. In principle, each 2TB NVMe device can support ~ 160 minutes of data assuming a 90% net capacity. Table 2 summarizes the data rates for the different modes.

At the time of writing this document, the NVMe capacity available in the four beam host nodes is listed below:

- `gwbh7  :  2TB × 1`
- `gwbh8  :  2TB × 2`
- `gwbh9  :  2TB × 1`
- `gwbh10 :  2TB × 1`

7 Offline Diagnostics and Analysis Tools

Offline tools for diagnosis and analysis are available for the PASV data. The software are available in all the gwb host nodes under `/common-h10/gpuuser/PASV/`:

- PA Total Intensity beamformer: `/common-h10/gpuuser/PASV/PASV2PATI`
  Tool for converting PASV to PA Total Intensity (PATI) spectra with user-defined FFT integration. The output is in the same format as the GWB PATI output and can be analysed using standard tools like `gptool`.

- PA Full Stokes beamformer: `/common-h10/gpuuser/PASV/PASV2PAFS`
  Tool for converting PASV to PA Full Stokes (PAFS) spectra with user-defined FFT integration. Again, it is exactly the same format as the GWB Full Stokes beam. Note that the auto- and cross-polar products are available in the .raw output files, and not the Stokes products.

- Offline CD: `/common-h10/gpuuser/PASV/PASV2CDP`
  An offline version of the Coherent Dedispersion Pipeline (CDP) is available for running on the PASV data.

- PPS detector: `/common-h10/gpuuser/PASV/PASV2PPS`
  An offline tool for reading the timestamp information from the accompanying `.ts` file is available. In addition, a tool exists for detecting the 1PPS instance(s) within each shared memory block of PASV data. The size of the shared memory block depends on the .hdr settings.

- VDIF converter: `/common-h10/gpuuser/PASV/PASV2VDIF`
  An offline tool for converting the PASV data to a 2-thread timestream VDIF is available. This tool converts the full bandwidth into either 8-bit or 2-bit VDIF for downstream analysis using standard pulsar processing software. Tools for VDIF conversion for VLBI with VLBI-friendly bandwidths are also available, but described elsewhere (see Internal Technical Report, *GMRT baseband signal processing for VLBI*).

8 Quality-check results from pulsar observations

Bright pulsars were observed in all the above modes. All modes have been tested with ~ 1 hour of continuous recording on NVMe disks. No data loss has been detected. Figure 2 shows example folded profiles from a 100-MHz/8-bit and a 200-MHz/8-bit run on PSR B0329+54 at Band-4 and Band-3 respectively. Figure 3 shows the polarized profiles of PSR B0329+54 obtained by processing PASV data.

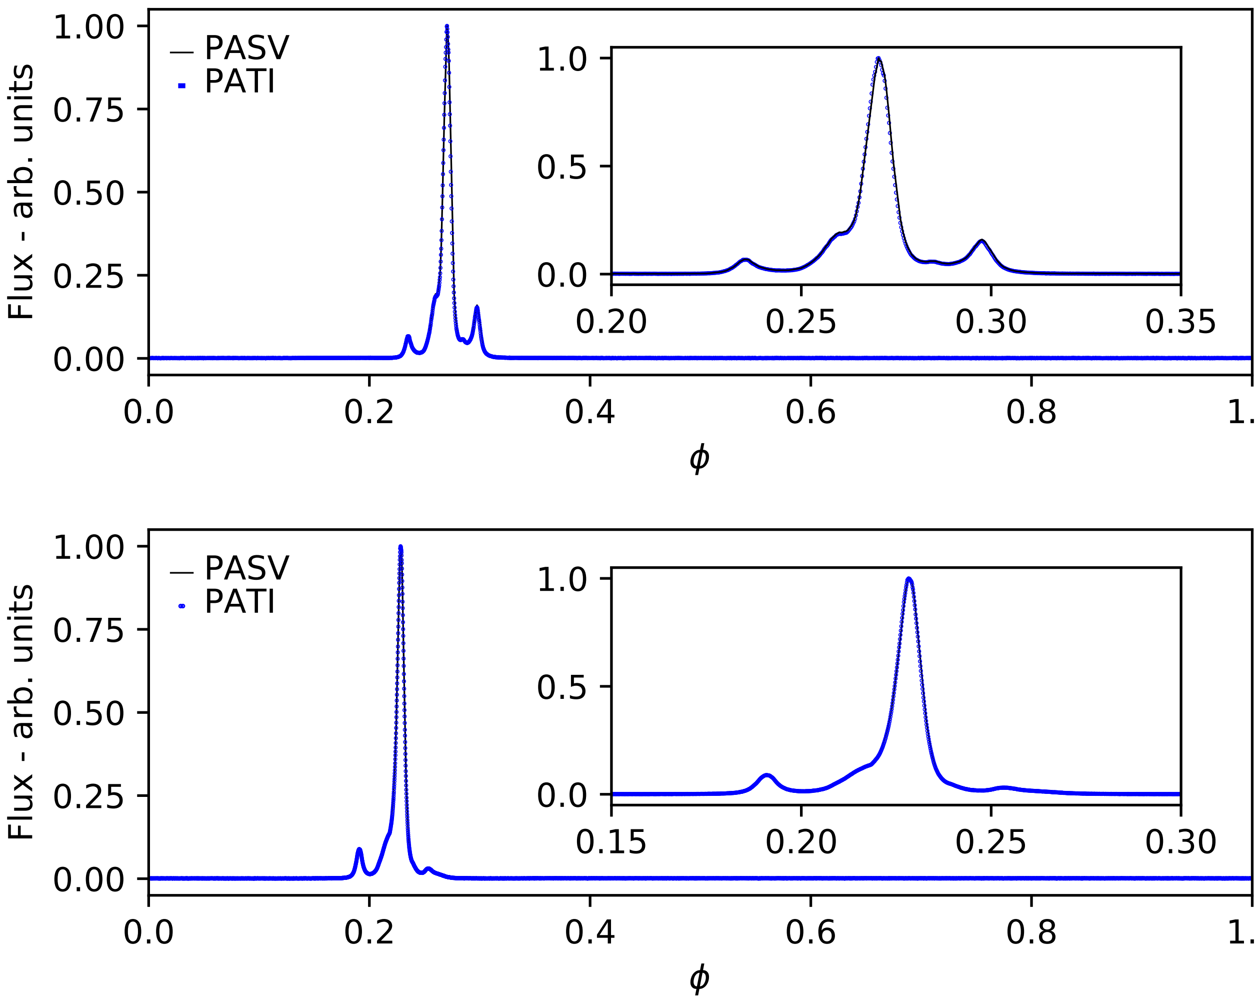

Figure 2: Comparison of the folded profiles of PSR B0329+54 from simultaneous recording of PA Total Intensity (PATI) and post-processed PASV data. *Top:* Band-4 100-MHz/8-bit mode. *Bottom*: Band-3 200-MHz/8-bit mode. Inset in each panel shows a magnified view of the pulse.

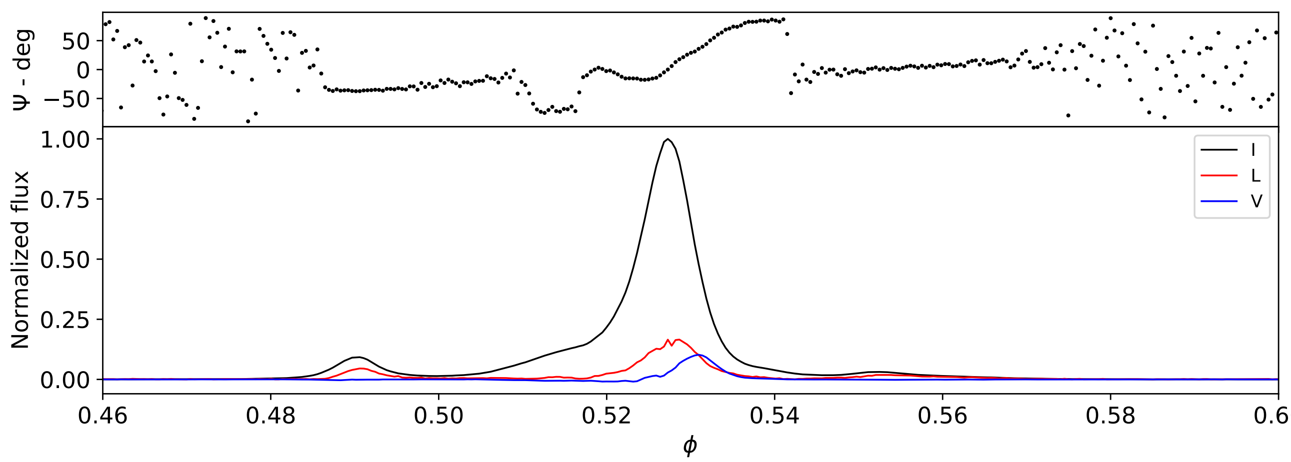

Figure 3: The total intensity (*black*) linearly (*red*) and circularly (*blue*) polarized profiles and position angle of B0329+54 over the pulse-on phase at Band-3 generated from the same PASV data shown in the bottom panel in Figure 2.

9 Standard Operating Procedure for running the PASV mode

This Standard Operating Procedure (SOP) outlines the steps to invoke the PASV mode recording. As a prerequisite, the GWB has to be configured with valid settings and the correlator has to be running. For the PASV mode to be recorded on beam-`x`, the CD beamformer has to be running in that beam, i.e. beam-`x` has to be invoked in the GWB GUI in the Voltage mode. However, note that the coherent dedispersion and recording themselves should not be running for the PASV mode to be recorded. For the 200-MHz mode, only one CD beam can be invoked, whereas in the 100-MHz mode, two CD beams can be invoked. Where detailed SOPs are available (e.g. starting and running the GWB), details are excluded in this SOP.

1. Configure and start the GWB acquisition in the required configuration of `RELEASED` or `TRIAL` mode.

2. Start the GWB scan.

3. On the TGC machine terminal, enter the recording command with the necessary arguments. This command can also be used in the observation command file after scan start commands:

   ```
   pasv.pl <prjcode> <path1> <path2> <obs-duration> <source> <frequency>

   Example :  pasv.pl TEST nvme0 nvme1 300 3C48 1450
   ```

   where

   `prjcode` : project code
   `path1` : the disk where the 130-POL PASV baseband data file is to be written
   `path2` : the disk where the 175-POL PASV baseband data file is to be written
   `obs-duration` : the observing duration in seconds
   `source` : the source name
   `frequency` : the LO frequency (or channel 0 frequency) in MHz

   Note that under the given disk name, the full path is generated internally using the combination of prjcode, disk-name and date.

   E.g. if command entered is: `pasv.pl TEST nvme0 nvme1 300 3C48 1450`

   the script generates a directory with full path as:

   `/nvme0/gpuuser/TEST_nvme0_10feb2022`

   where it records the following files:

   ```
   3C48_PASV_P2_bm2_2048_8_100_1450_10feb22_154201.raw
   3C48_PASV_P1_bm2_2048_8_100_1450_10feb22_154201.raw
   bm2pasv_10feb22_154201.log
   3C48_PASV_P1_bm2_2048_8_100_1450_10feb22_154201.raw.ts
   ```

4. The `.log` file is an ASCII file with information on the current record from the correlator and the record number being accessed in the shared memory block. In addition, it gives information on the association between the correlator record number and the shared memory block identifier.

5. After completion of observing time given in step 3, stop the GWB scan.