

# Processing Methodologies

The standardised methods for processing I-LOFAR data products

- Single Pulse Source Observations
- LuMP Processing
- PRESTO Timing
- Timing With Tempo2 (Empty)
- Processing Non-Pulse-Based Observations
- Getting TOA Measurements from Single Pulses

# Single Pulse Source Observations

FRBs and RRats are highly transient by their nature. As a result, we could see a single peak in a 10 hour observation. While keeping the raw voltages on hand can help with re-processing observations to avoid missing features due to issues in the current methodology, we no longer have the storage space on the UCC processing nodes or NUIG1.

Proposed methodology for processing single pulse observations observations:

- Process the observation with **CDMT**. Produce both a **0-DM and N-DM** filterbank at the nominal time resolution (currently 655.36us, 16x **ts**, 8x **chan**)
- Perform **RFI detection** on the 0DM filterbank
  - We currently do not have a strong methodology here, apart from bandpass analysis and rfind for detecting DM=0 features
- Search this output after an **8-bit decimation from digifil**
  - Log the heimdall commands used, RFI channels flagged
  - Investigate the optimal scale timescale for digifil (-l, default 10s) for FRBs; some are expected to last up-to or over 1 second at our frequencies due to scattering.
- After a search is complete, **archive the CDMT filterbanks**
  - **Digifil: 2x ts** for further space savings if needed, mostly a layover from previous 8x tsch
    - -l 0 : No scale changes, raw 2:1 conversion
    - -b-32 : Float32 output, no change from raw filterbank
    - -t 2 : Down sample to 655us resolution
  - **Compressed** with zstandard: Further 10-20% compound storage saved

There are a few ways that this methodology could be changed to make the resulting filterbanks easier to search + store, or improve SNR

- Future changes
  - **Chop bandwidth?** Top 5MHz / Bottom 7 are Nyquist suppressed + RFI contaminated
    - Removing these could save us 15% of storage and speed up processing as searching the last 10MHz introduces an addition delay of 25 seconds @ R3's DM
    - No easy way to do this with the current voltage extraction/processing method, would need to be after the filterbanks are formed
  - Investigate having CDMT **split filterbanks** every N samples
    - Consider overlap requirements to not miss signals on the boundaries

- Duplicated data, but higher theoretical SNR when we can include more channels by more selective RFI flagging
- Or just find a decent RFI flagging algo...

We note that for RRats, we do not recommend forming a ODM filterbank as those sources often do not need validation as they should be bright enough to be obvious with/without coherent dedispersion.

| Step                             | Method                     | Storage Used | Product  | Overall on Disk       |
|----------------------------------|----------------------------|--------------|----------|-----------------------|
| <b>Generate Voltages</b>         | Observer                   | 1            | 1        |                       |
| <b>Compressed</b>                | zstandard, Olaf's recorder | ~0.6-0.8     | 0.6      | 0.6                   |
| <b>CDMT</b>                      | -a -b 16 -d 0,DM,2         | 0.125        | 0.125    | 0.725                 |
| <b>Digifil (Search)</b>          | -b 8 -l <DECIDE>           | 0.03125      | 0.03125  | 0.75625               |
| <b>Cleanup: Digifil (search)</b> | rm                         | -0.03125     | -0.03125 | 0.725                 |
| <b>Compress CDMT (compress)</b>  | zstandard                  | ~0.1         | 0.1      | 0.825                 |
| <b>Cleanup: Voltages, CDMT</b>   | rm                         | -0.6 - 0.125 | -0.725   | 0.1                   |
| <b>Overall</b>                   |                            |              |          | <b>~100 GB/obs-hr</b> |

# LuMP Processing

An alternative recording method for some observations is with the LuMP software from MPIfRA. It has been used in the past for coordinated observations with FR606 and at the request of observers from the UK and Poland.

Any DSPSR sub-program (dspsr, digifil, digifits) can be used to process a LuMP observation, but each port/process (if using a multi-processed recording mode) must be processed separately and then combined (fils: filmerge, fits/ar: psradd).

So as an example, to process with digifil you may choose to process a set of observations using the command

```
for file in *.raw; do
  digifil -b-32 -F <NUMCHAN>:D -l 0 -c -set site="I-LOFAR" $file -o $file".fil"
done

filmerge *.raw.fil
```

To perform coherent dedispersion (`-F <NUMCHAN>:D`) for a known pulsar target (inside LuMP metadata), without any bandpass/temporal offsets (`-l 0 -c`), producing a 32-bit output (`-b-32`) filterbank.

Many issues arise with modern versions of DSPSR when processing raw data, not limited to the dedispersion kernel failing, the default filterbank failing, misaligned folds when directly processing with DSPSR, etc. As a result we use a modified version of the workflow presented above for processing a typical LuMP observation.

```
baseName=$1

# Process the raw data with digifil. Perform 8x channelisation, 2x time scrunching (tsamp = 81us)
# Fake machine to COBALT as sigproc's filmerge will refuse to merge fils if the header is FAKE
for fil in *.00.raw; do digifil -b-32 -l 0 -c $fil -set machine=COBALT -set site=I-LOFAR -t 2 -F 328:1 -o $fil".fil" &
echo "hi" ; done; wait;

for fil in *.01.raw; do digifil -b-32 -l 0 -c $fil -set machine=COBALT -set site=I-LOFAR -t 2 -F 328:1 -o $fil".fil" &
echo "hi" ; done; wait;

for fil in *.02.raw; do digifil -b-32 -l 0 -c $fil -set machine=COBALT -set site=I-LOFAR -t 2 -F 320:1 -o $fil".fil" &
echo "hi"; done; wait;

# Each port should have the same number of samples and starting MJD; merge each of them
```

```
filmerge ./udp16130*raw.fil -o "./udp16130_"$baseName".fil"
filmerge ./udp16131*raw.fil -o "./udp16131_"$baseName".fil"
filmerge ./udp16132*raw.fil -o "./udp16132_"$baseName".fil"
filmerge ./udp16133*raw.fil -o "./udp16133_"$baseName".fil"

for fil in udp*"$baseName".fil; do digifil -b 8 $fil -o $fil"_8bit.fil"; done

# Fold the data, 1024 bins, ~3 secnd integration (change turns as needed)
for fil in *_8bit.fil; do dspsr -turns 4 -nsub 512 -t 4 -b 1024 -skz -skzn 4 -k lelfrHBA -O $fil"_fold" $fil; done

# Attempt to combine the data. This will not work 90% of the time due to packet loss, but worth trying.
psradd -R *.ar -f $baseName".ar"
```

# PRESTO Timing

PRESTO can be used for generating timing files for use with tempo(2).

To start, a standard `psrcat` command should be run, though to use the output archives for timing the `nosearch` flag must be used, as a result you will need a well-timed target (good entry in psrcat) or an existing ephemeris file on hand for the folding.

Once you have a `psrcat` pfd generated, you can use the `psrcat2tim` script to generate TOA .tims to process with tempo(2).

# Timing With Tempo2

## (Empty)

# Processing Non-Pulse-Based Observations

The backend used for CDMT is also available in a CLI, `lofar_udp_extractor`, which is installed on the [Docker containers](#) available on the REALTA nodes.

This guide assumes you have a UDP recording (compressed or uncompressed) from Olaf Wucknitz's VLBI recording program (standard for observing with I-LOFAR) and will explain the standard operating modes, and workarounds for issues with the `lofar_udp_extractor` program. [The full, up to date documentation for the CLI can be found here.](#)

## Standard Usage

```
lofar_udp_extractor \  
[-i /path/to/raw/udp_1613%d.TIMESTAMP.zst \  
-o /output/file/location \  
-p <procMode>
```

This sets up the program to take a compressed ZST file, starting at port 16130 and iterating up to port 16133, outputting to the provided location in a set processing mode. Some processing modes have multiple outputs, and will require '%d' to be in the output name as a result. The most useful processing modes are

| Mode ID | Output (Stokes) | Tsamp (us) | Outputs |
|---------|-----------------|------------|---------|
| 100     | I               | 5.12       | 1       |
| 104     | I               | 81.92      | 1       |
| 150     | I, Q, U, V      | 5.12       | 4       |
| 154     | I, Q, U, V      | 81.92      | 4       |
| 160     | I, V            | 5.12       | 2       |
| 164     | I, V            | 81.92      | 2       |

Modes 150+ are only available in more recent versions, and may error out of the docker containers have not been updated recently.

There are several other useful flags for processing data, such as `-u <num>` which will change the number of ports of data processed in a given run, `-t YYYY-MM-DDTHH:MM:SS -s <num>` or `-e <files>` can be used to extract a specific chunk of time, or specify a file with several time stamps and extraction duration (with the requirement that these regions do not overlap).

The `-s <flag>` flag passes flags to `mockHeader` which generates a sigproc-compatible header of metadata about the observation. This can make handling Stokes data easier later on, through the use of `sigpyproc` for loading and manipulating data, though as of right now it is not possible to set a per-subband frequency as is needed for mode357, so a dummy `fch1` (central top frequency) and `foff` (frequency offset between channels) should be used instead.

As an example, during a processing run on 29/10/20 of some Solar Mode357 data, the following command was used.

```
olaf_udp_extractor \  
-H /mnt/ucc1_recording/data/sun/20201028_sun357/20201028090300Sun357/udp_1613%d.ucc1.2020-10-28T09:05:00.000.zst \  
-o ./2020-10-28T09:05:00_mode357_StokesVector%d.fil \  
-p 164 \  
-a "-fch1 200 -fo -0.1953125 -tel 1916 -source sun_357"
```

## Known Issues and Workarounds

When recording starts later than the supplied start time, Olaf's recorder may pick up stale packets in the UDP cache and record them at the start of your observation. This will manifest itself as a **segfault when trying to process the start of an observation**, as the program will run into issues attempting to align the first packet on each port. As a workaround, use the `-t YYYY-MM-DDTHH:MM:SS` flag to set a start time shortly after the actual data begins recording, at which point the software will be able to accurately align the packets as needed,

# Getting TOA Measurements from Single Pulses

This page describes the process to get a TOA measurement for a single pulse, assuming

- You know the rough TOA of the pulse
- The input data is a Sigproc Filterbank
- DSPSR and PSRCHIVE (with GUI) are available

Many steps of this process are automated on REALTA using this python script[gist].

<getting the .ar>

## Generating a Noise-Free Model

We will use the `baas` tool to generate a noise free model, which will then be used for cross-correlation or other analysis methods to determine the pulse TOA. Choose your brightest or most characteristic pulse and begin the fitting process by running

```
baas [H] # Interactive fitting
[!-d /xwin # Visual GUI of choice
<input .ar> # Input profile to use as a reference
```

Once loaded in, focus on the pulse itself by pressing `z` to set the left limit of a zoom, and left click to select the right limit. Then, left click on the left and right edges of the pulse to set the phase limits of the pulse, you will then be able to select the peak of the pulse vertically.

Once you have a rough model in the view, you can press `u` to iteratively update the model to the data, continue to update the model until you believe a good fit of the amplitude and position of the pulse has been achieved and the residuals of the region (red lines) are similar to the noise floor.

You can then quit by pressing `q`, this will save the model to disk as 3 files, `model.dat` (the model we generated), `model.ar` (an archive profile containing the shape of the model) and `model.txt` (an ASCII copy of the model)

We will be using the `model.txt` file for determining the pulse TOAs.

# Determining Pulse TOAs using the Noise-Free Model

Now that we have our archives and model, we can use `pat` to determine the pulse TOAs. We typically perform this using the following command,

```
pat [-f tempo2 \ # Output in the tempo2 format
[-A PIS \ # Generate cross correlations using the Parabolic interpolation method, chosen for the it's performance
on a test dataset from J2215+45
-F \ # Sum across frequencies before determining TOA
-m <paas model>.m \ # Model generated by paas in the previous section
-s <paas profile>.std \ # Archive generated by paas in the previous section
<input archives>.ar > <output filename>.tim

# Optional flags, you may need to remove -m for these
[-t \ # Plot the profile, template and residuals
-K /xwin \ # Using an xwindow
```

The output timing file can be used for analysis in tempo2.