

Bandpass Correction in LOFAR

John W. Romein

August 25, 2008

Abstract

This document describes how we correct the bandpass in the OnLine Application Processing pipeline (better known as "the Blue Gene/P correlator").

1 Introduction

The LOFAR pipeline contains two PolyPhase Filters (PPF) to split the spectrum into narrow-bandwidth channels.

The first PPF resides at the stations, is implemented in FPGAs, and splits the spectrum into 1024 subbands. Only the first 512 subbands are usable, since the upper 512 subbands are mirrored (subband n is the conjugate of subband $1024 - n$). Having multiple subbands allows reducing the station output datarate: only the subbands that are interesting to the astronomer are observed; the other frequency ranges are discarded.

In this paper, we assume that the 200 MHz clock is selected at the stations, which yields subbands of 195312.5 Hz wide. If the 160 MHz clock is selected, all frequencies mentioned below should be multiplied by 0.8.

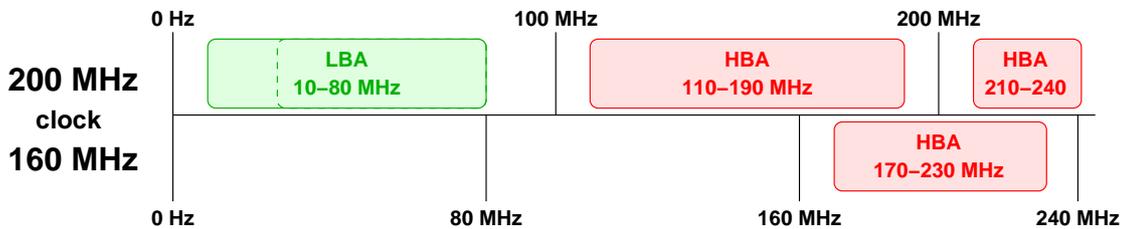


Figure 1: Observation modes.

Frequencies above 100 MHz can be observed by measuring in higher Nyquist zones. Since one cannot observe reliably at frequencies close to a multiple of half the clock rate, the 160 MHz clock must be selected for HBA observations close to 200 MHz. The observing modes are shown in Figure 1.

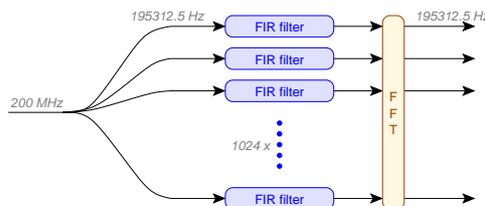


Figure 2: The first-stage PPF performs a 1024-point FFT over the output over 1024 FIR filters.

The first-stage PPF itself consists of 1024 Finite Impulse Response (FIR) filters, of which the outputs are Fourier transformed (see Figure 2), as explained below.

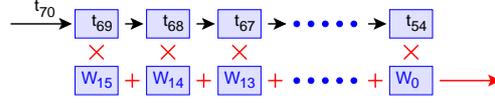


Figure 3: A 16-tap FIR filter.

A *FIR filter* is essentially a time-delay filter with a small (in our case 16) number of history buffers, i.e., *taps* (see Figure 3). Each clock tick, a sample goes in, shifting history data to the right. A weighted sum of the 16 taps goes out. We use the FIR filter as a band pass filter, i.e., a filter that attenuates high and low frequencies. The filter characteristics are determined by the weights. The PPF contains 1024 FIR filters, each with its own weight vector. The delayed stream of station samples is round-robin distributed over the FIR filters (see Figure 2), and a FFT of the FIR filter outputs yields 1024 frequency subbands, each 195312.5 Hz wide.

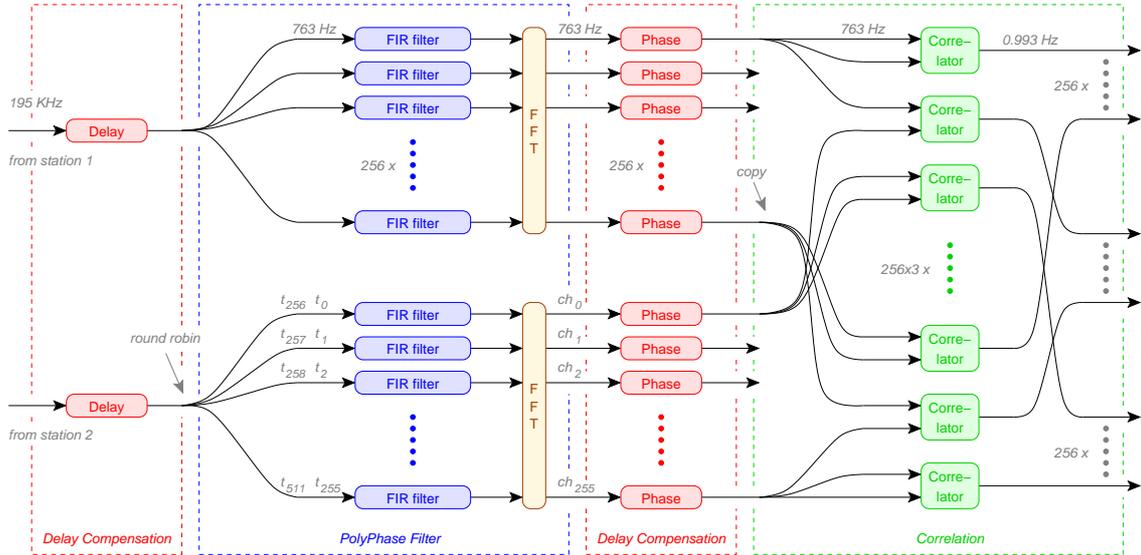


Figure 4: Signal processing on CEP for one subband.

Centrally, a second-stage PPF splits each 195312.5 Hz subband into 256 channels of 762.94 Hz wide. Having narrow frequency channels allows flagging of narrow-band RFI. This PPF runs in software on the IBM Blue Gene/P. Its place in the on-line processing chain is shown in Figure 4 and explained in more detail in [1]. Note that the correlator aggravates the bandpass in the correlations, since it multiplies the bandpasses of two PPFs.

2 The bandpass

One property of a PPF is that it is not equally sensitive to incoming signals over the entire frequency range, e.g., a 50.000 MHz input sine-wave signal results in a stronger PPF output than a 50.002 MHz input signal of the same power. This leads to problems when two consecutive PPFs are used, as is the case in LOFAR. The first-stage PPF creates subbands, each 195312.5 Hz wide, but the sensitivity over this frequency range varies, as is shown by Figure 5. This becomes visible only after applying second-stage PPF, since in the output of the first-stage PPF, there is no way to distinguish different frequencies within the same subband. However, the output of the second-stage PPF does distinguish between different frequencies

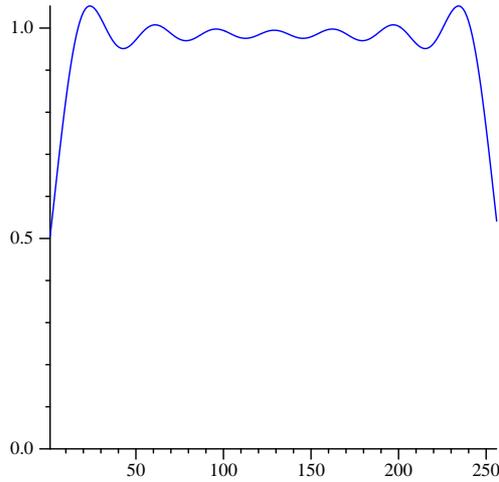


Figure 5: The bandpass of a subband.

within a subband, and the bandpass becomes immediately visible. Thus, it is the *first-stage PPF* that causes the ripple over the output data of the second-stage PPF.

The bandpass can be computed analitically and must be removed from the observed data. Logically, the best place to do this is right after the second-stage PPF, thus on the Blue Gene/P.¹ In other words, the online-processing chain should clean up the dirt caused by the stations 😊

3 Leakage

4 Shifting frequencies

Another property of a PPF is that the frequencies are shifted, as explained below. The first-stage PPF shifts frequencies by half a subband; the second-stage PPF shifts frequencies by half a channel. The shift is a result of the Fourier transform.



Figure 6: The PPF shifts the frequency by half a channel. The red numbers below indicate the channel numbers.

Figure 6 illustrates this behavior for the second-stage PPF. Here, subband 256 is shown, with a center frequency of 50 MHz. The shift causes a 50.000000 MHz signal to fall in the middle of channel 128,

¹Note that the current implementation corrects the bandpass after correlation, for efficiency reasons.

not halfway the subband between channels 127 and 128. As a consequence, half a channel at the right of the subband folds into channel 0. Channel 0 contains both power from the leftmost and the rightmost half-channel of the subband, and because of this ambiguity, the correlator always flags channel 0 as being invalid.

Something similar happens at the first-stage PPF of the station, where subband 0 contains data from 0 Hz to 97,656.25 Hz and from 199,902,343.8 Hz to 200,000,000 Hz (or higher frequencies in higher Nyquist zones). Subband 0 is thus unusable (but is unusable for other reasons as well, such as the insensitivity of the antenna elements at those frequencies).

5 Bandpass correction

In this section, we describe how we compute the bandpass from the FIR filter constants of the station PPF. The bandpass can be obtained by taking the Fourier transform of the weights of all 1024×16 FIR-filter constants in the first-stage PPF. A file called `Coeffs16384Kaiser-quant.dat` contains these weights. With matlab or octave, the correction factors can be obtained as follows:

```
f=fftshift(fft(Kaiser,262144))
m=f(131073-128:131073+127)
r=f(131073-128+256:131073+127+256)
l=f(131073-128-256:131073+127-256)
plot(2^25./sqrt(abs(m).^2+abs(l).^2+abs(r).^2))
```

f is the Fourier transform over the FIR filter constants, where the peak is shifted to the center elements (131073 – 128 to 131073 + 127). A transform length of 262144 is chosen to get 256 bandpass factors. m , r , and l , contain the subband response, the response of the subband right to the center, and the response of the subband left to the center. To compensate for channels that leak from the left and the right, their powers are added by adding the square of their values and taking the square root of the sum. (Currently, the correction is applied *after* correlation, so that the result should be squared again, or equivalently, omitting taking the square root). The reciprocal is taken to get the correct multiplication factors, as a multiplication is performed much faster than a division.

Figure 7 shows how real data is corrected. The figure shows the power of an autocorrelation in an RFI-quiet subband (300, which corresponds to a center frequency of 58.6 MHz), integrated over 30 minutes. The blue curve is not corrected and shows a (squared) bandpass; the red curve shows corrected data. The slope is caused by the receivers, that become decreasingly sensitive at higher frequencies in this subband. TODO: find out why the red curve bends at the edges.

References

- [1] J.W. Romein, P.C. Broekema, E. van Meijeren, K. van der Schaaf, and W.H. Zwart. Astronomical Real-Time Streaming Signal Processing on a Blue Gene/L Supercomputer. In *ACM Symposium on Parallel Algorithms and Architectures*, Cambridge, MA, July 2006.

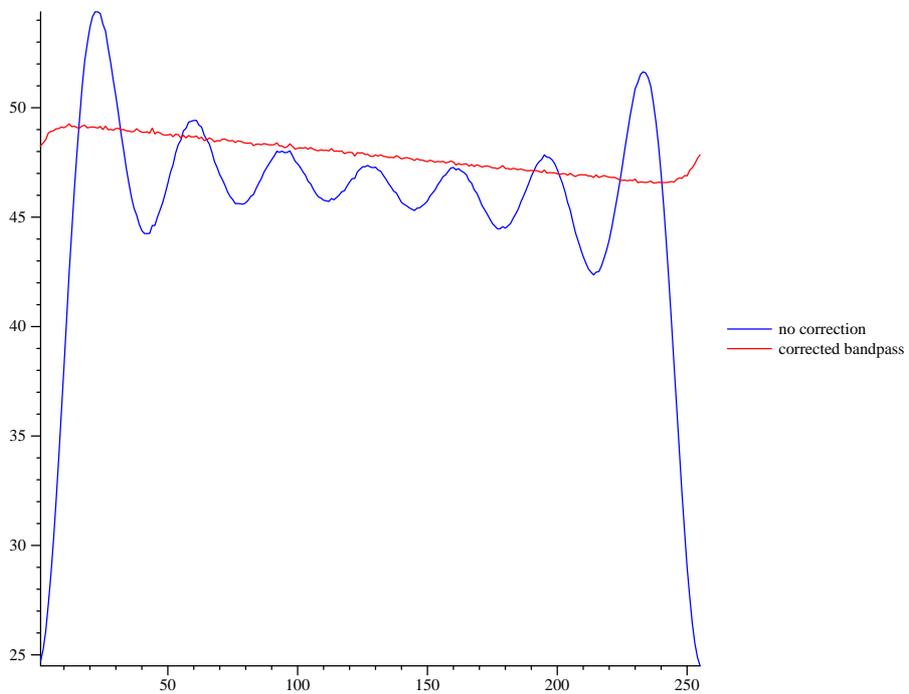


Figure 7: The bandpass correction.