Development of a GPU-based wide band receiver at the Nançay radio telescope

J.Borsenberger, P.Colom, J.M. Martin, G.Kenfack

Observatoire de Paris, LESIA-CNRS, GEPI-CNRS, USN-CNRS

Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique

Galaxies Étoiles Physique et Instrumentation

MOTIVATION

The main reasons to enlarge the bandwidth is to rise the sensitivity toward continuum sources and to widen the spectral coverage in view of searching for the HI and OH lines in redshifted galaxies.

Another need is to achieve a much greater temporal resolution than the present correlator used at the NRT, which is too slow compared to some RFIs, e.g. radars.

Our digital receiver is aimed at registering either the waveform in a sustained manner, either the real-time computed power spectra. While the former task does not need the GPU processors, the latter one is not performed in the FPGA-based ROACH board, but rather in GPUs in order to benefit of the high precision implied by using real*4 float numbers.

The use of floating point frees from the constrains usually encountered in FPGA FFT, which involve shifts that become difficult to handle when increasing frequency resolution. The architecture of GPU allows to go to very high resolution, while keeping wide band capability.

RECEIVER CONFIGURATION

The first task assigned to the FPGA is to tag the output packets so that they can securely be dated. In fact there is a tiny packet loss using the standard 10GbE, even using jumbo frames, so we need a tag to assign the received data in their proper place in the analysis window.
To diminish the number of lost packets, an effort has been needed to reduce file-system activity, and to properly assign interrupts on specified cores.

We use the four 10GbE outputs to switch each polarization in two back-end PCs (see schema), allowing for forthcoming improvements in RFI mitigation field. The size of the data chunk must be at least twice the window analysis, currently this frame size is set to 2^26

The back-end PCs use shared memory to store these frames, together with ancillary data, and status flags. This allows for the use of any number of GPU without re-programming. The limit is set by the number of cores available to control GPUs. One of these cores is dedicated to receive data on the 10GbE and store them in memory.

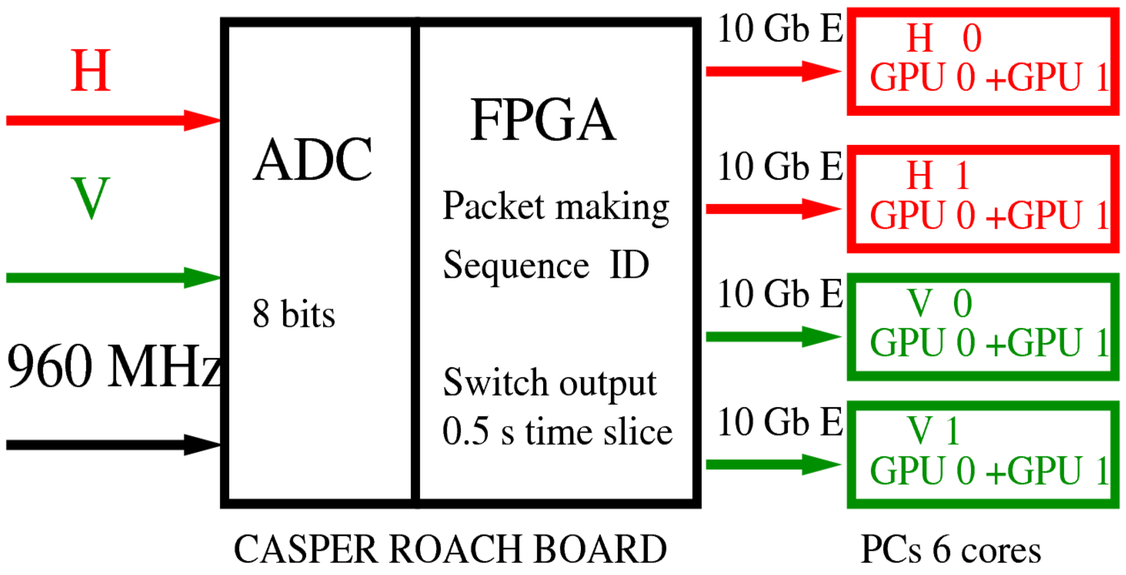

Schematic view of the new digital receiver. Two entries on the left are receiving the horizontal and vertical linear polarisation in the intermediate frequency band of the NRT, filtered with about 500 MHz bandwidth. Also a 960 MHz clock signal is sent to the Analog to Digital Converters (ADC).

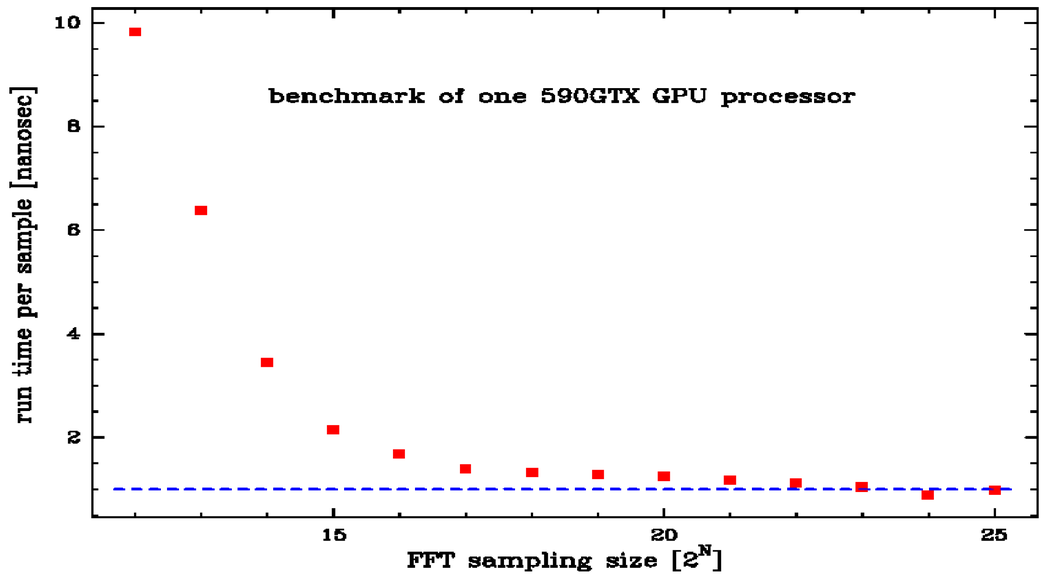

Benchmark of a GPU processor. Note the run time par sample is not at all proportional to the exponent N, as it should be for an FFT in a CPU processor. Here, the optimum size is in the range 20-25, that is for one million samples and more.

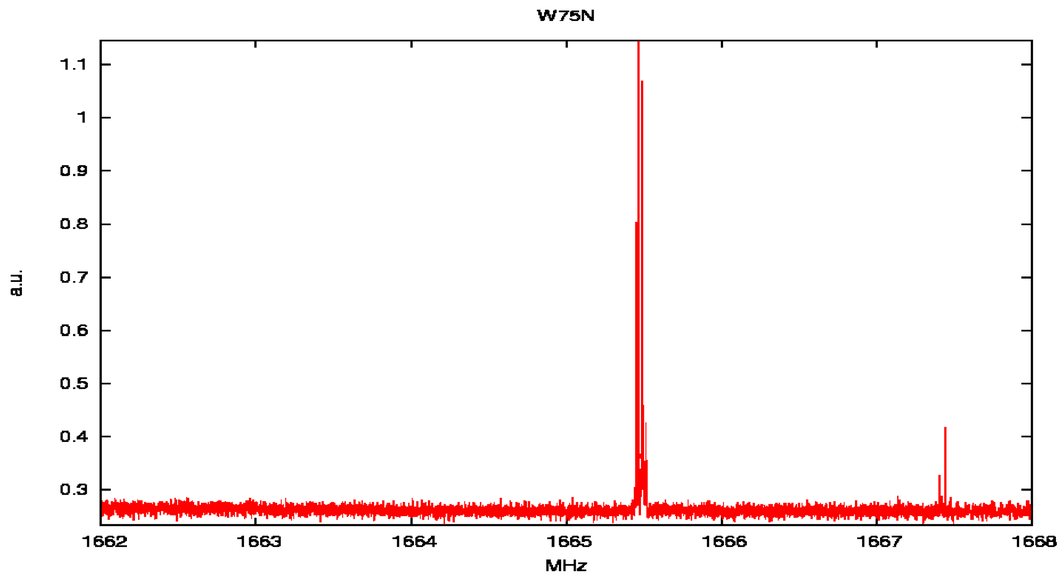

Zoom of a spectrum obtained toward W75N, an OH galactic maser. The two main lines of the OH ground state are seen. No baseline subtraction has been applied.

BENCHMARK

As regard to real-time aspects, it was essential to measure the execution time of power spectra calculations per sample, as a function of the
FFT size.

Hence, we present here a benchmark of the 590GTX GPU (NVIDIA), which gives very similar results as compared to the 275GTX et 285GTX processors.

The execution times per sample are given in aside figure, as a function of the exponent N involved in the FFT sampling size.

As the sampling clock is close to 1 GHz, we also have shown the limit of 1 nanosec (blue dashed line) to which every result has to be compared.
This time comprises the FFT, the squared coefficients, and a power accumulation calculations.

It should be noted that, according to the number of operations involved, we should observe a log(size) behaviour. On the contrary there is a decay of the per sample time. The architecture of the GPU parallelism, used by the CUDA toolkit, is very probably responsible of this behaviour.

Also worth to mention is that we have gained a factor ~2 over these times by implementing an FFT algorithm that takes every second data set in the imaginary part, hence handling two real data sets in one complex numbers FFT calculation.

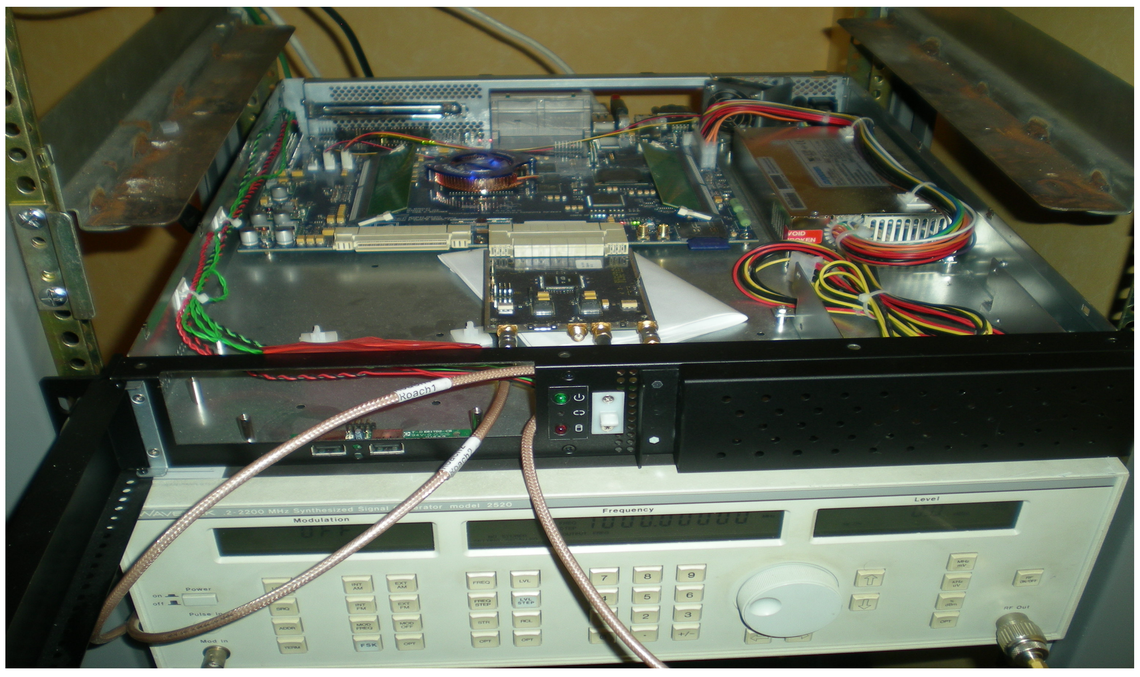

View of ADC, ROACH board, clock, and the two signal feeds