Mem. S.A.It. Vol. 89, 366

Memorie della

High resolution low cost radio spectrometer for astrophysical observations

Roberto Lulli[1,2]

[1] National Institute for Astrophysics (INAF), Radio Astronomy Institute (IRA), Via Fiorentina 3508/B, I-40059 Medicina (BO), Italy
[2] University of Camerino, School of Advanced Studies (SAS), via Camillo Lili 55, I-62032 Camerino (MC), Italy, e-mail: roberto.lulli@unicam.it

Abstract. We realized a high resolution, low cost spectrometer that can be also used in the radio field SETI (Search for Extra-Terrestrial Intelligence Cocconi & Morrison 1959) researches. In its last implementations, SETI instrumentation needs a high resolution (of the order of 1 Hz), real-time or almost real-time spectrum analysis that must be performed over the entire receiver bandwidth.
Till now, this kind of spectrometers were developed dividing up the bandwidth into contiguous blocks analysed separately by single nodes of clustered systems. Although very effective, these systems are also very expensive not only for the number of computing nodes required but mainly for the data communication tools among the nodes.
Recent technological advances make possible designing efficient and inexpensive data acquisition and processing system for radio signal analysis based on COTS (Component Off-The-Shelf) parts like multi-core CPUs (Central Processing Units), many-core GPUs (Graphics Processing Unit) and FPGAs (Field Programmable Gates Array).
Over the last four years we carried out many pilot projects based on high performance PC (Personal Computer) that produced very promising results. We present here the current state of effort, the results we obtained and the future work.

Key words. SETI – SERENDIP – spectrometer – FFT – GPU

1. Introduction

On April 8, 1960, Frank Drake spent several weeks listening for extraterrestrial signals pointing a 26-meter radio telescope (National Radio Astronomy Observatory in Green Bank) to two nearby stars. This was known as Project Ozma (Drake 1961) and it was the first modern SETI research. In the 1970s, NASA started to be interested in the SETI program and then, at the end of the 80ies, such a program became a reality. Observations began in 1992 on the 500th anniversary of Columbus' arrival in the New World. In 1993 the US Congress cancelled the program. At that point the SETI Institute, which was conceived in 1984 to facilitate scientific research and educational programs related to life in the universe, decided to base its activities on private funding to continue at least part of the NASA SETI Program (Phoenix Project). The first Phoenix observations took place at the 64-meter Parkes radio telescope in Australia.

Since 1960 there have been more than 98 SETI projects around the world. Nowadays, SETI research activities are focused not only on radio spectrum but also in the optical one usually including the near UV too (Siemion et al. 2010).

This observation windows have been chosen because they are the only ones in which the atmosphere is clear enough to allow effective observations from the Earth.

The current reference scientific and technological tools for SETI observations are those developed within the six "Serendip" projects; these tools are currently installed in various radio telescope both in the USA (Hat Creek, Green Bank, Arecibo, Ohio State, ATNF) and in the rest of the world (Medicina station, Italy and Parkes, Australia).

In the years, the evolution of the Serendip spectrometers was possible thanks to the latest electronic and computing technologies available on the market.

The first two spectrometers Serendip I (1979) (Bowyer et al. 1983) and Serendip II (1986) (Werthimer et al. 1984), (Bowyer et al. 1988) were designed in a different way from the subsequent ones, namely, Serendip III (1992), Serendip IV (1997) (Werthimer et al. 1995), Serendip V[1] (2005) (Siemion et al. 2010) and Serendip VI[2] (2014) that have quite similar architectures and characteristics consistent with following guidelines:

1. modular development: radio signal analysis is divided among different computational units linked to each other by high speed communication systems;
2. Hz order resolution;
3. real time or almost real time signal analysis.

In fact, even though using the best available technology, the latest Serendip instruments have never allowed a high frequency resolution analysis of the entire bandwidth with a single computing unit. In the years, the increase in computing capabilities of the integrated circuits used for the analysis never compensated the increase in the bandwidth of the receivers. Consequently the bandwidth coming from the receiver has to be suitably divided among different computing units.

In order to obtain a complete analysis, the results given by different units must be linked to each other. This poses a big problem because the bandwidth partition cannot be sharp (the transfers function of the filter cannot be perfectly rectangular). This means that, in order to perform correct calculations near the edges, each computing unit must deal with a portion of the bandwidth assigned to the nearby units. This makes the system structure even more complex and reduces its efficiency when the number of elaboration units increases from some units to several dozens. Costs related to the filters bank that splits the incoming bandwidth among the computing units and the master unit that calculates the final results increase as well: e.g. the architecture of the sixth Serendip is composed by an analog beam-splitter, two FPGA based multi-channel acquisition units, five high-end servers with GPUs linked to each other by means of a high-speed Ethernet switch. The overall cost of this processing system can be estimated in the order of tens of thousands of dollars.

The idea behind this research project is building a new kind of spectrometer with a cost of € 10,000, completely enclosed in a single processing unit (high performance PC) containing the data acquisition system, the computing system, the control and the display systems. This system does not need components that are usually necessary in multi-unit (clustered) systems, like the data communication appliances and the master unit. In this way we save costs related to such components.

The entire bandwidth is analysed by the system itself (without dividing it among different computing units) taking into account that the spectrometer performances are limited by the capabilities of less performing data processing chain module. This implies that such spectrometer needs a more correctly balanced design to maximize performances and reduce

[1] http://casper.berkeley.edu/wiki/SETI_Spectrometer

[2] https://casper.berkeley.edu/wiki/SETI/FRB_Spectrometer_(SERENDIP_VI)

costs. Such design could allow to build an instrument capable of analysing most part or the entire bandwidth coming from the current INAF radio telescopes receivers.

Another goal of this project is the identification and, hence, the achievement of the performance limit of our spectrometer. The pilot projects completed so far show that it is possible to analyse 100 MHz bandwidth with 1 Hz resolution in real-time. The last phase of the project expects to expand the analysed bandwidth above 100 MHz up to a maximum of 600 MHz.

2. Previous experiences and pilot projects

The first project of a high performance spectrometer at the INAF Radio Station of Medicina (BO) was started in 2012. The system was based on a high-end PC with a DAQ (Data AcQuisition) board hosted on PCI Express and capable of sampling at 1 Gsample/s with an 8-bit resolution. Spectrum analysis and post-processing computation were performed by a CPU (Intel i7 LGA 2011 four core, 3.4 Ghz) that also managed system control and data visualization. The algorithm used for spectrum analysis was the FFT (Fast Fourier Transform) implemented by FFTW library (Fastest Fourier Transform in the West) (Frigo et al. 1997), (Frigo et al. 1998), (Rajovic et al. 2016), that, on the available hardware, reached a maximum of 4 millions of channels at variable resolutions. The spectrometer was also able to reach the Hz resolution, but not in real-time because it was necessary to make a substantial decimation of the sampled signal due to the impossibility of the acquisition board to work at data rates lower than 380 Msample/s.

The spectrometer worked properly but its performances were limited by the microprocessor capabilities and had a cost of about € 14,000 of which the great part (two thirds) was devoted to the acquisition board. Moreover, the system did not have a high scalability and reliability mainly because the DAQ board was supported by a very small set of operating systems and an absent-minded update was enough to put the system out of order.

Building up on this earlier experience, we tried to improve the quality of the system increasing both the reliability and the computation power. On the basis of other international attempts (Kondo et al. 2010), we thought to reach the latter goal by inserting into the PC accelerators for massive parallel computation. So the system architecture has been revised by assigning spectrum analysis to the accelerators and leaving to the CPU only tasks involving system control and data visualization.

Among the many accelerators currently available on the market (GPUs by NVIDIA, or AMD, Intel Xeon Phi many core systems and FPGAs), we chose NVIDIA GPUs for the following reasons: the lower cost, the immediate availability and the possibility to use the NVIDIA software library cuFFT dedicated to FFT computation (NVIDIA Corporation, CUDA Programming Guide 2009-2017), (NVIDIA Corporation, CUDA Toolkit and SDK 2009-2017), (NVIDIA Corporation, CUDA CUFFT Library 2009-2017). The path towards the design of the new high performance spectrometer evolved through two pilot projects devised to validate the adopted technological strategies. The first pilot project was carried out in 2015 using old hardware already present in the laboratory (Fig. 1). The basic system was a PC with a first generation Intel i5 dual-core processor as CPU and a Picoscope digital oscilloscope with USB 2.0 interface, bandwidth up to 25 MHz and 12-bit resolution as DAQ.

The major limit of this spectrometer was the fact that DAQ could not acquire single data sets bigger than 256 ksample. In these circumstances and using the FFTW library, the CPU was capable to analyze the maximum bandwidth of 25 MHz with a resolution of 200 Hz and a real-time output of 24 spectra per second. This instrument was mainly used as a test bed for the software development and also for testing the capabilities of NVIDIA GPUs available in lab.

In the autumn of 2016 a second pilot project was carried out with the goal of building a spectrometer with real-time analysis ca-

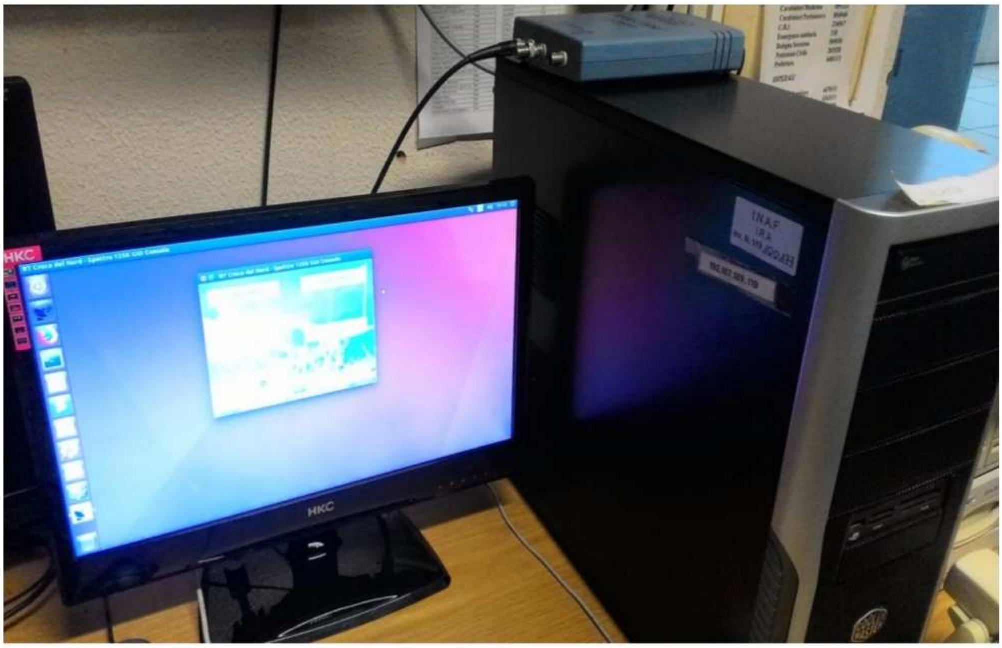

Fig. 1. Photo of the first pilot project spectrometer composed by a PC that host on the top of case the USB Oscilloscope.

pabilities at 1 Hz resolution (Fig. 2). As DAQ we used a relatively obsolete board with PCI interface and maximum sampling rate of 20 Msample/s borrowed from the University of Camerino. This board was installed on a PC with a single core CPU (due to software driver limitations) that also supports PCI Express bus to host a NVIDIA Quadro 5000 GPU. This system has proven to be able to meet the project specifications analyzing the entire bandwidth of 10 MHz in real-time at 1 Hz resolution, but the tests highlighted some incompatibility issues between the hardware and the software which made the spectrometer unstable and therefore not reliable. Despite that problem, this system validated our technological strategy so we could move on to design the new high performance spectrometer.

A few months later the end of the second pilot project, following some recent developments in the supercomputing field, we had the idea to solve the incompatibility issues by replacing the NVIDIA GPU with an array of low-cost calculation boards based on ARM processors. The use of ARM processors in HPC (High Performance Computing) applications has been the subject of several studies in both academic and industrial environments. Those studies were focused on energy consumption optimization rather than on performance. The main research project in this area is called Mont-Blanc[3] and has been funded by the European Union under the program Horizon 2020. At present, array of two different boards for educational applications are being evaluated, one based on Raspberry Compute Module 3[4] (Fig. 3) and another based on UDOO QUAD[5] (Fig. 4). We expect to achieve results of this new experiment by the spring of 2019.

[3] http://montblanc-project.eu/project

[4] https://www.raspberrypi.org/products/compute-module-3/

[5] https://www.udoo.org/udoo-dual-and-quad/

Fig. 2. Photo of the internals of second pilot project spectrometer. From the bottom the first add on board is the DAQ and the second is the NVIDIA Quadro 5000 GPU.

3. Project architecture

In the autumn of 2017 we carried out the design of the new spectrometer based on the experience accumulated thus far. Following the architecture block diagram of Fig. 5, the analog signal from the radio telescope antenna is converted into digital data by DAQ board that produces at least a data set per second; these data sets are analysed in real-time (or almost in real-time, in case of more complex computations) allowing to reach a 1 Hz resolution.

Data produced by the acquisition board is transferred to the PC RAM in DMA mode freeing the processor from transfer management. The RAM is the real core element of the architectural block diagram where the data exchanged by the other blocks must mandatory pass.

In fact, acquisition and processing of signals with a bandwidth greater than 50 MHz produces a such amount of data that, in a very short time, would saturate the SATA bus interfaces of any kind of mechanical or solid state hard disks. On the market there are also solid state hard disks with PCI Express interfaces whose performances are, at least in theory, good enough for our purposes. But, they are very expensive and the number of their write cycles is too low to make our spectrometer a reliable system. So, at present, for this application, it is impossible to use any kind of mass storage as computation support memory.

The solution we found is using the PC RAM as a storage area for both the acquired data and the results of computation. We store in the mass memory only the results that are significant and relevant for the application pur-

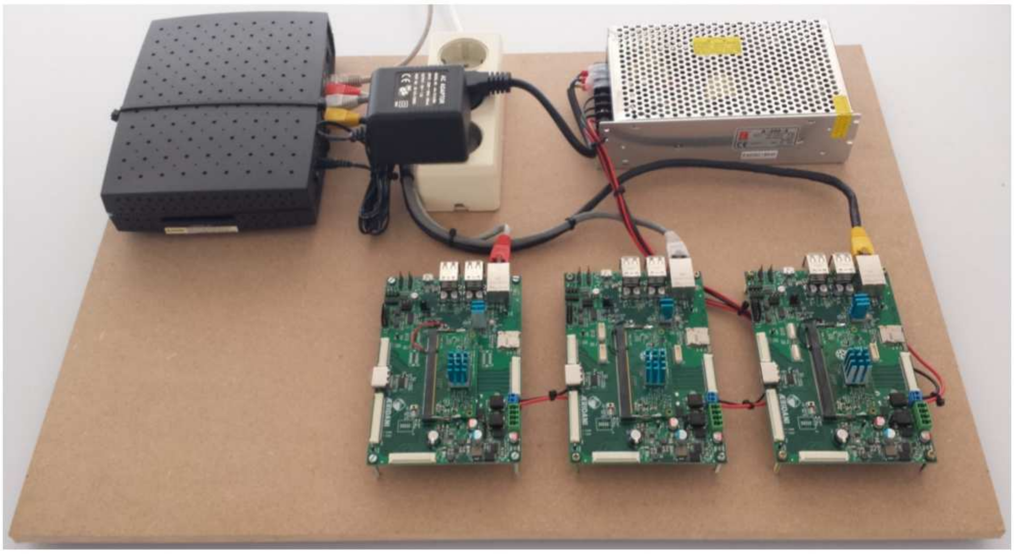

Fig. 3. Photo of a 3 Raspberry Pi Compute Module 3 array.

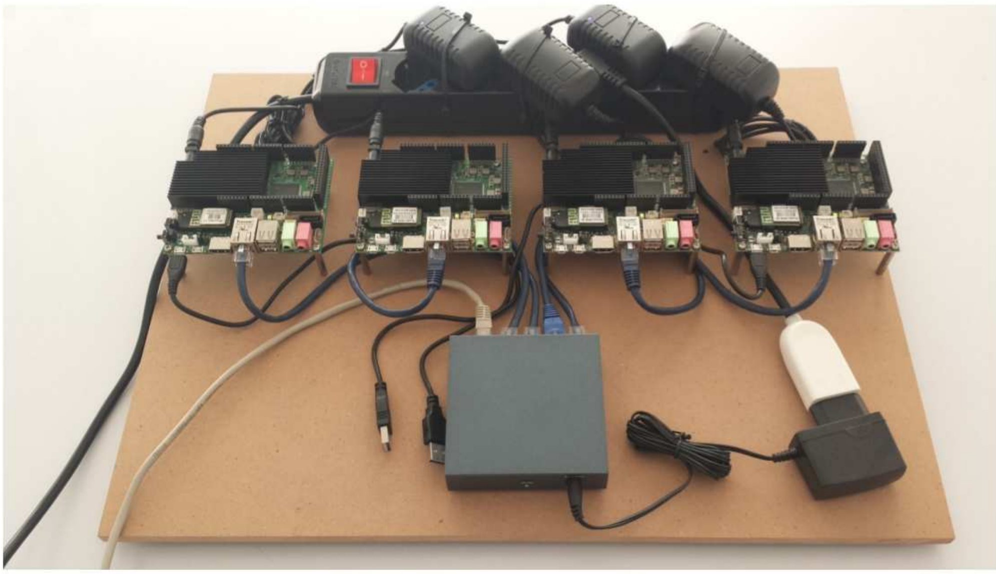

Fig. 4. Photo of a 4 UDOO QUAD board array.

poses; we think that these data should not be more than 5% of acquired ones.

All the calculations necessary for the processing result will be carried out by two accelerators which are used, not in parallel, but in alternate mode. Each accelerator will exclusively manage a single data set produced by the acquisition board and will therefore have two seconds of time to analyse them completely.

Both the acquisition board and the accelerators are connected to the PC via a PCI Express bus which, at the moment, is the only kind of

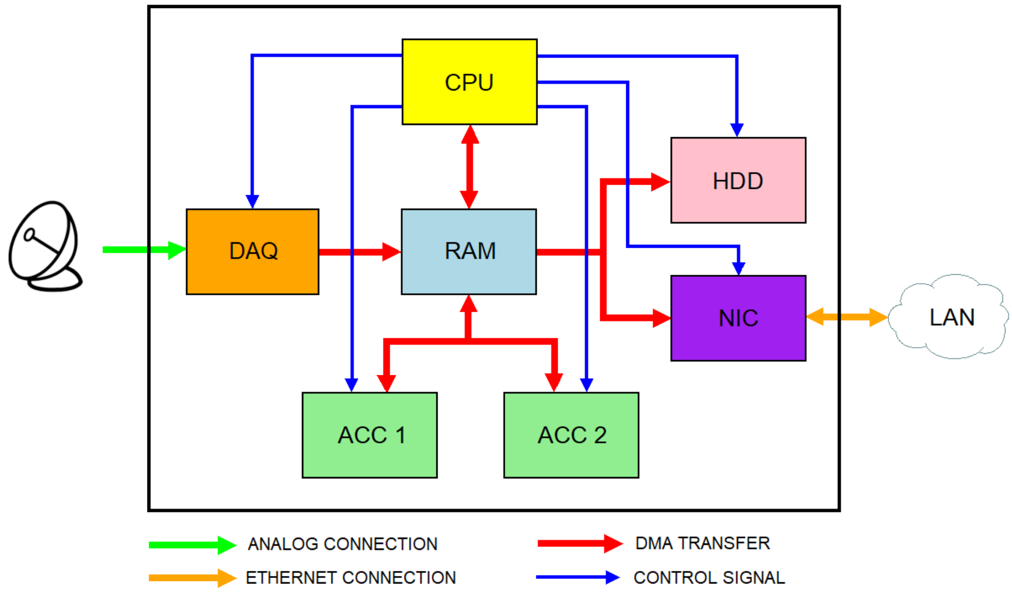

Fig. 5. Block diagram of new spectrometer architecture.

bus capable of transferring the expected data streams in an acceptable time. The use of the accelerators in alternate mode also allows us to reduce the data transfer time for the benefit of data computing time.

The CPU is instead used for the following aims: system control, diagnostics, user interface management, and filtering of significant results. Since the CPU has to accomplish all these tasks at the same time, we expect to use at least a 4-core processor with a clock frequency greater than 3 GHz.

Finally, from the data processing results we extract significant data to be stored on mass storage or directly transferred to the outside via a network interface. The data can be stored and transferred in different formats chosen by the user, including the one compatible with the data reading and display software Scan7.

4. Criteria for the choice of the spectrometer components

One of the most important parts of the spectrometer design has been the choice of its internal components. This was strongly conditioned by two key parameters, namely performance and cost.

For what concerns the data acquisition board, after the tests carried out during the first pilot project, we decided not to use an USB oscilloscope for the design of the new spectrometer. Although the cost is very competitive, the USB oscilloscope used in the first pilot project is currently capable of acquiring a maximum bandwidth of 25 MHz. New models developed by the Picoscope[6] company, are able to reach bandwidth of several GHz but the USB 3.0 connection in continuous data streaming mode (the mode we need for real-time analysis) limits the maximum bandwidth to 60 MHz. This bandwidth is too narrow for the requirements of actual INAF radio telescopes but can be useful for smaller or older radio telescopes usually used in educational or amateur astronomy fields. The design of such a particular instrument is left for future work since it might be of help in building spectrometers with even lower costs (€ 3,000).

[6] https://www.picotech.com/oscilloscope/3000/usb3-oscilloscope-logic-analyzer

So for the new spectrometer we chose a DAQ board with a PCI Express interface whose characteristics allow to solve the problems related to USB 3.0 bottleneck. From market researches, we was able to identify a board that acquires signals data with a maximum bandwidth up to 600 MHz and with a 8-bit resolution (about 48 dB of dynamic range). It works in continuous streaming with optimized data transfer (this last task is performed in background without overloading the CPU and the PCI Express bus). This board can works with internal and external clocks whose frequency can vary from a few tens of kHz up to 1.25 GHz. It has support for different operating systems; moreover there is a technical support (provided by the producer) that replies with competence and promptness. The cost of this DAQ is about half of that of the first spectrometer project but having better performances.

All types of accelerators currently available on the market have been evaluated: NVIDIA and AMD GPUs, Intel Xeon Phi many-core systems and FPGA based accelerators. An in-depth analysis shown that most of them have integrated circuits of a comparable integration scale and are able to provide the same performances in terms of clock frequency and power consumption. Also the peak performances declared by the producers are aligned to the same order of magnitude (from 1 to 10 Tflops). However, from the literature[7] (Teodoro et al. 2014), (Liu et al. 2015), and from our tests, we discovered that this peak performances are only theoretic and they are reachable only when the hardware architecture is optimised for the implemented algorithm. Otherwise, the performance degrades even by several orders of magnitude.

Irrespective of the used technologies and of the specific model, professional accelerators cost is up to € 10.000 each. So, to limit costs, we decided to use no-high-end accelerators whose architectures is optimized for the spectrum analysis algorithm we need to implement.

[7] `http://www.nvidia.com/object/justthefacts.html`

In particular, we chose a high-end COTS gaming NVIDIA TITAN Xp GPU (Fig.6). The architecture of the NVIDIA GPU boards is especially optimized for the FFT (Fast Fourier Transform) algorithm; moreover these boards support a high-performance parallel version of the FFTW. These boards use state-of-the-art GPU technology circuits with almost the highest performance available and have a very competitive cost because they were developed specifically for the consumer market. On the other hand, they are not so reliable as professional boards and provide less features. With the first pilot project spectrometer we tested two GPU boards: a Quadro 5000 (Fermi series GPU for professional market) and a GeForce GTX 680 (newer Kepler series GPU but for high-end gaming market) which pointed out a bandwidth real-time analysis of 100 MHz (Quadro 5000) and almost 160 MHz (GeForce GTX 680) respectively.

project where the obsolete DAQ board compatibility imposed also the use of an old GPU. In this situation the Quadro 5000 was able to perform a real-time analysis of a signal with a bandwidth of 10 MHz at 1 Hz resolution.

NVIDIA cuFFT library documentation analysis and subsequent tests showed that bandwidth limitations mainly depends on GPU onboard memory: the same tests where carried out on a newer board, a GeForce GTX 970 (Maxwell series) with the same onboard memory as GeForce GTX 680 and the results showed roughly same bandwidth analysis capabilities (but with different computation times).

For the final project we decided to use a pair of NVIDIA GPUs TITAN Xp Pascal series equipped with 12 GB of memory each. We were expecting that this memory size allows to analyse signals with a bandwidth triple with respect to boards equipped with 4 GB of memory like the GeForce GTX 680 and GTX 970. Tests carried out in past few months showed that the assumption was true because Titan XP was able to analyse in real-time signals of 450 MHz of bandwidth.

The graph in Fig. 7 shows an overall comparison of the performance of the four GPUs

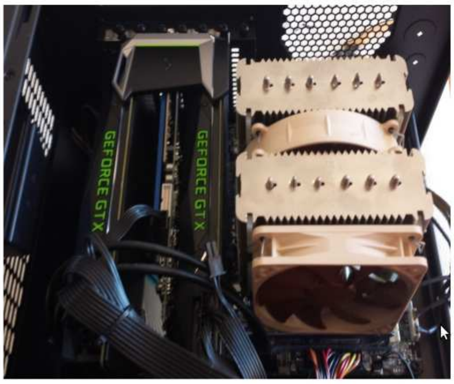

Fig. 6. Photos of the externals (on the left) and the internals (on the right) of the new spectrometer's prototype. In the right photo form the left, the second GPU, the DAQ board, the first GPU and the CPU's cooler system.

analyzed. Along the X axis it is represented the size of the data set provided as input to the algorithm for the calculation of the FFT. Being the input data set composed of real numbers, the size of such data set must be double than the size of the output one: e. g. to obtain a spectrum analysis of a signal with a bandwidth of 100 MHz at a resolution of 1 Hz, it is necessary to provide an input data set of 200 Msample.

Along the Y axis the computational times of FFT algorithm are shown. They are always below 350 ms, so considering a data set acquisition per second, it is always possible to carry out the analysis of the signal in real-time. The irregular trend of the four tracks is typical of the algorithm that allows to obtain lower calculation times if the size of the data set can be factored with fewer prime numbers of smaller value (2, 3, 5 and 7): the best possible performance is obtained when the size of the data set it is a power of 2.

Actually the cost of NVIDIA TITAN Xp GPUs is less than € 1,500, making it possible to set the entire system at a cost of less than € 10,000.

Finally if a low cost gaming GPU with 16 GB of memory is available in the future it is likely that this will allow analysis of up to 600 MHz of bandwidth and take full advantage of the capabilities of the DAQ board.

5. Software and user interface

In designing the spectrometer's software we followed the same guidelines as in the hardware design: costs and performances.

To limit costs, we decided to use open-source and free software together with ad-hoc developed software. Among the many available operating systems (remind that our spectrometer is PC-based) we use Ubuntu 16.04 because it is directly supported by the SDK (Software Development Kit) of both the GPU and the DAQ board.

The application is mostly implemented in C/C++, with parts regarding the GPU written in CUDA (Compute Unified Device Architecture). In effect the great part of standard computation acceleration libraries have a Fortran-oriented interface because originally written in this language, so we could also choose to use this language for the application. But the use of C/C++ is preferable due to an easier interaction with the operating system and with drivers and SDKs of DAQ and GPUs. The application is multi-threaded. We have planned a number of different threads that perform the following tasks: data acquisition, GPU elaboration, selection of meaningful results and their storage in the mass memory.

The application has also a modular structure, especially for what concerns software packages that implement the elaboration pro-

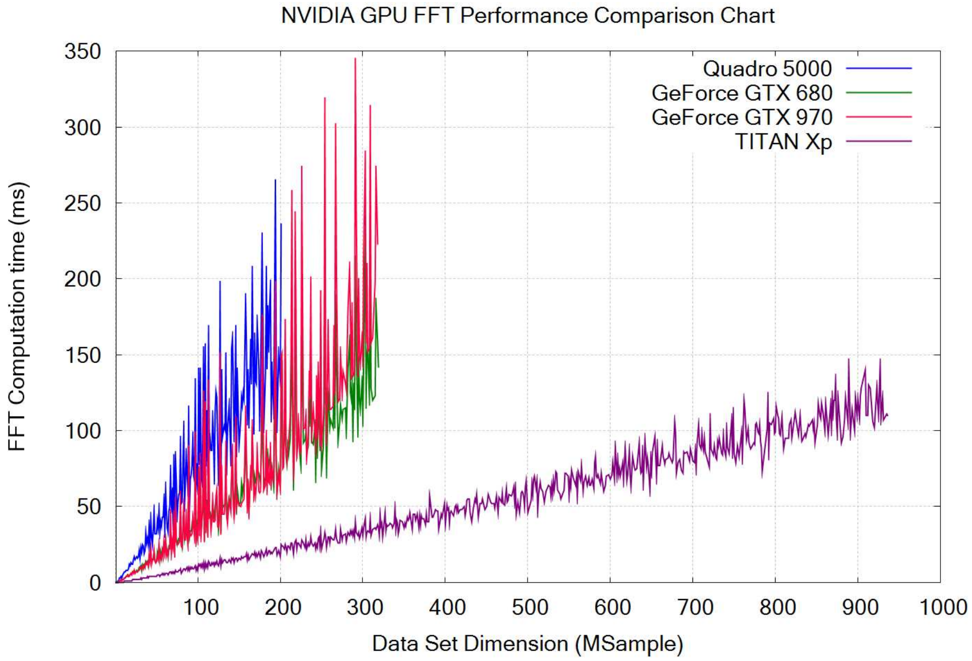

Fig. 7. Comparison chart of the performance in FFT computation time of different NVIDIA GPU boards with increasing dimension of the data set (one test every million of points).

cess and the interfaces of spectrometer hardware components. Over time, these components could also change, replaced for instance by more powerful GPUs and data acquisition boards with a wider bandwidth.

The main part of the application performs the following processing chain:

1. data acquisition;
2. transfer of the acquired data to RAM;
3. synchronization with the GPUs and transmission of the data to be processed;
4. FFT execution (or other spectrum analysis algorithms);
5. boxcar (or other algorithms) execution to reject the filter shape from the spectrum;
6. threshold computation (with average, moving average or other algorithms);
7. result formatting based on the requested output format.

The processing chain can be also built dynamically and may even include modules not currently planned. For instance, we are considering the implementation of the following functionality:

a. auto-optimization of the application based on hardware components;
b. dynamic construction of the processing chain;
c. configuration of acquisition and elaboration parameters;
d. selection of the output data format;
e. monitoring and diagnostics of the system status.

The application will also include functions for visualizing the processing results, status error and error messages. The use of these features can also take place in graphical mode through a user friendly front-end written in Python. The choice of Python for the development of the front-end was motivated by the portability and standardization of the language as well as its ability to easily interface with standard C/C++ applications.

6. Conclusion and future work

The modularity of the processing chain and the chance given to the user to personalize such a chain, allow the spectrometer to become a very flexible instrument especially if the number of modules and, hence, the available functionalities constantly increases.

We are already planning to add other software modules for spectrum analysis that collaborate with or replace the FFT, like those dedicated to KLT (Karhunen-Love transform), to the Wavelets or to Agnostic Entropy analysis. Other modules could integrate the current post-processing with different equalisation modes and signal extraction with respect to noise. Finally, other modules could concern data reduction and formatting.

As well known, other algorithms have better performances if we execute them with other kind of accelerators different from NVIDIA GPU simply because these accelerators are best suited for them. For this reason, we need to consider the possibility of integrating different accelerators to replace the current ones or to provide a joint use of both. Some tests have already been carried out, but finding suitable solutions would require more work in order to properly evaluate performances and to reduce hardware and software compatibility issues.

In case the wideband DAQ was replaced by a narrower bandwidth DAQ equipped with more channels, it would be possible to perform integrated multi-channel analysis on arrays of antennas such as the N/S arm of the Medicine station. This would imply a different software processing chain to preserve the effectiveness of the hardware architecture, but a priori this change seems to be technically possible and economically feasible.

Moreover, by having several spectrometers connected to similar but distant antennas (Medicina, Noto, SRT, etc.) and having an adequate bandwidth for data exchange, it is possible to use the spectrometers to carry out VLBI analysis (Very Long Baseline Interferometry). As for multichannel analysis, even this analysis would require a modification to the software processing chain, but it seems technically possible and economically convenient.

Finally, if the analyzable bandwidth of the spectrometer is too narrow for receivers currently installed at the radiotelescopes or in the case of simultaneous use of multiple receivers in different bands, one might even think to replicate the Serendip model and then build a modular system by connecting multiple spectrometers in parallel.

In conclusion, we believe that a quite inexpensive instrument, open at software level, scalable and updatable over time, set up the way to new research with reduced funding and allows the consolidation of those currently in progress.

References

Bowyer, S., et al. 1983, Icarus, 53, 147

Bowyer, S., Werthimer, D. and Lindsay, V. 1988, The Berkeley Piggyback SETI Program: Serendip II, in Bioastronomy: The Next Steps, G. Marx ed. (Kluwer, Dordrecht), 363

Cocconi, G. & Morrison, P. 1959, Nature, 184, 844

Drake, F. 1961, Physics Today, 14, 40

Frigo, M. & Johnson, S.G. 1997, The Fastest Fourier Transform in the West, Technical Report MIT-LCS-TR-728

Frigo, M. & Johnson, S.G. 1998, FFTW: An Adaptive Software Architecture for the FFT, in the Proceedings of the 1998 IEEE ICASSP, 3, 1381

Frigo, M. & Johnson, S.G. 2005, The Design and Implementation of FFTW3, in Proceedings of the IEEE, 93, 216

Kondo, H., et al. 2010, A Multi-GPU Spectrometer System for Real-time Wide Bandwidth Radio Signal Analysis, in International Symposium on Parallel and Distributed Processing with Applications (ISPA), 594

Liu, T. , Xu, S.G., Carothers, C.D. 2015, Annals of Nuclear Energy, 82, 230

NVIDIA CUDA Programming Guide version 2.3 and subsequent revisions (2009 - 2017)

NVIDIA CUDA Toolkit and SDK version 2.3 and subsequent revisions (2009 - 2017)

NVIDIA CUDA CUFFT Library version 2.3 and subsequent revisions (2009 - 2017)

Rajovic, N., et al. 2016, The Mont-Blanc prototype: an alternative approach for HPC systems, in SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (IEEE), 444

Siemion, A.P.V., et al. 2010, Current and Nascent SETI Instruments, Astrobiology Science Conf., 2010

Teodoro, G., et al. 2014, Comparative Performance Analysis of Intel Xeon Phi, GPU, and CPU: A Case Study from Microscopy Image Analysis, in 2014 IEEE 28th International Parallel and Distributed Processing Sumposium, 1063

Werthimer, D., Tarter, J.C. and Bowyer, S. 1984, The Serendip II Design, in The Search for Extraterrestrial Life: Recent Developments, ed. M. D. Papagiannis (D. Reidel, Dordrecht), IAU Symp., 112, 421

Werthimer, D., et al. 1995, The Berkeley SETI program: SERENDIP III and IV instrumentation, in Progress in the Search for Extraterrestrial Life, ed. G. Seth Shostak (ASP, San Francisco), ASP Conf. Ser., 74, 293