Astronomy ESFRI & Research Infrastructure Cluster

# ASTERICS - H2020 - 653477

# Technology benchmark report – D-GEX

## ASTERICS GA DELIVERABLE: D3.8

| | |
|---|---|
| Document identifier: | ASTERICS-D3.8-final.docx |
| Date: | **30-04-2017** |
| Work Package: | **WP3 OBELICS** |
| Lead Partner: | **INAF** |
| Document Status: | **Report** |
| Dissemination level: | **Public** |
| Document Link: | www.asterics2020.eu/documents/ ASTERICS-D3.8.pdf |

Abstract

The document shows results in benchmarking different computing technologies and architectures for astrophysical data analysis, focusing on performances in execution time

and power consumption. Experiments have been conducted in the mainframes of ASTRI project (INAF) and DOME (ASTRON).
Both groups have designed systems targeting the requirements of low power consumption, thus realising sort of datacenter in a box; the ASTRI team focused mainly on software (ASciSoft) and algorithms, to be run on embedded boards (Nvidia Jetson, ARM + Nvidia GPU processors) attached on ground telescopes, while ASTRON team investigating on hardware integration. Testing have been conducted giving expected results, while leaving room for improvements, both in hardware selection and integration and in software development.

# I. COPYRIGHT NOTICE



# II. DELIVERY SLIP

| | Name | Partner/WP | Date |
|---|---|---|---|
| From | Lucio Angelo Antonelli, Francesco Visconti | INAF/WP3 | 06/04/2017 |
| Author(s) | Francesco Visconti | INAF/WP3 | 06/04/2017 |
| Reviewed by | Giovanni Lamanna | LAPP/WP3 | 20/04/2017 |
| Approved by | AMST | | 30-04-2017 |

# III. DOCUMENT LOG

| Issue | Date | Comment | Author/Partner |
|---|---|---|---|
| 1 | 06/04/2017 | First Draft | Francesco Visconti/INAF |
| 2 | 19/04/2017 | Second Draft with comments | Francesco Visconti/INAF |

| | | | |
|---|---|---|---|
| 3 | 29/04/2017 | Final Version | Lucio Angelo Antonelli/INAF |

# IV. APPLICATON AREA

This document is a formal deliverable for the GA of the project, applicable to all members of the ASTERICS project, beneficiaries and third parties, as well as its collaborating projects.

# V. TERMINOLOGY

| | |
|---|---|
| GPU | Graphical Processing Unit |
| FPGA | Field Programmable Gate Array |
| ASciSoft | ASTRI Data Reconstruction and Scientific Analysis Software |
| CTA | Cherenkov Telescope Array |
| ASTRI | Astrofisica con Specchi a Tecnologia Replicante Italiana |
| DOME | Dutch government-funded project between IBM and ASTRON |
| SKA | Square Kilometre Array |

# Table of Contents

# 1 Introduction

Very high-energy (VHE, E > 50 GeV) gamma-ray astronomy is a rather young and flourishing field with high scientific potential. Gamma rays of such energies represent a unique probe for casting light on fundamental issues of galactic and extragalactic astrophysics, particle physics, fundamental physics, and cosmology. The VHE sky is currently being investigated with high collection areas by means of ground-based imaging atmospheric Cherenkov telescopes (IACTs), such as H.E.S.S., MAGIC, and VERITAS. In the last decade, with the firm detection of ~150 galactic and extragalactic sources, these arrays have demonstrated the huge physics potential at these energies as well as the maturity of the detection technique.

In order to boost the current IACTs performance and to widen the VHE science, a new generation Cherenkov telescope array has been proposed as an observatory open to the world-wide astronomical community. The full sky coverage will be assured by two CTA facilities, one in each of the Earth's hemispheres. The two arrays will make use of 100 telescopes at the southern site and 20 at the northern site of three different sizes – large (diameter 23 m); medium (diameter between 12 m and 9.5 m); small (diameter 4 m) – in order to cover the wide energy range from few tens of GeV up to few hundreds of TeV. The expected CTA sensitivity will be one order of magnitude better than current IACTs in the whole energy window and will give the possibility to fruitfully investigate several galactic, extragalactic, and fundamental physics science cases.

Within the framework of the CTA international project, the Italian National Institute for Astrophysics (INAF) is leading the ASTRI (Astrofisica con Specchi a Tecnologia Replicante Italiana) project. ASTRI is a flagship project of INAF (Italian National Institute for Astrophysics) funded by the Italian Ministry of Education, University and Research related to the next generation IACT (Imaging Atmospheric Cherenkov Telescope), within the framework of the CTA (Cherenkov Telescope Array) International Observatory. The project aims to develop and deploy a full end-to-end prototype of dual-mirror Small Size Telescope (SST-2M) to observe gamma-ray photons through the Cherenkov radiation of their atmospheric showers; the prototype will be compliant with the requirements of CTA.

IACTs are typically installed in isolated places in order to avoid luminous sky background. Such remote locations pose critical constraints on several aspects related to data processing, in particular: power consumption, data bandwidth, heat dissipation, etc.

The requirement of a fast data processing in an off-site datacenter is critical because of the costs related to data transmission, so the on-site data analysis is the only viable alternative. An effective solution to decrease the power and bandwidth required by the array facility could be to deploy an embedded module on each telescope, allowing it to carry out preliminary data reduction on its own before sending them to a central acquisition system.

Very similar technological issues are addressed by the DOME project. Within the DOME project, fundamental research is to be performed by ASTRON and IBM to inform the technology selection and roadmap for the Square Kilometer Array (SKA), especially in the areas of computing. Only by basing the overall design on architectures that are beyond state-of-the-art, on the data-centers and computing facilities of the future, will the SKA be able to handle the vast amounts of data produced by the antenna systems of the telescope.

The Dome project will take the SKA as a large and extremely challenging case study, as a carrier for fundamental research that has much broader application than just for the SKA. The research has direct relevance for the current partners, but will also lead to applications in a wide range of areas in the domain of high-tech systems.

This document summarizes the results obtained so far in investigating solution to match the issue of *green computing* in the ASTRI/DOME frameworks; green computing is to be meant as technologies to radically reduce the power needed to do computationally intensive work on extremely large amounts of data.

OBELICS is thus playing the fundamental role of benchmarking alternative computing architectures and innovating hardware solutions in such a context adapting existing software already developed or under development for ESFRI projects such as CTA and SKA with the final goal of integrating common solutions.

# 2. ASciSoft: ASTRI data reconstruction and Scientific Analysis Software

The ASTRI data reconstruction and scientific analysis software (henceforth ASciSoft) is the official package of the ASTRI project being developed for the ASTRI SST-2M prototype and mini-array data processing, in compliance with the CTA requirements [1]. The software can handle real and Monte Carlo (MC) data for both the ASTRI prototype and mini-array, and it is intended to provide all necessary algorithms and analysis tools for characterizing the scientific performance of the ASTRI SST-2M prototype and, afterwards, for carrying on the foreseen mini-array scientific program. It will be one of the first CTA data reconstruction and analysis software precursors to be developed and tested on a real-data basis.

The main purpose of ASciSoft is to reconstruct the physical characteristics of astrophysical gamma rays (and background cosmic rays) from the raw data generated by the ASTRI SST-2M prototype and mini-array. The software is composed by a set of independent modules organized in pipelines that implement all the necessary algorithms to perform the complete data reduction, from raw data to the final scientific products.

Within the OBELICS framework the code developed for ASTRI has been ported to be used in parallel computing architectures, such as multi-core CPUs and graphic accelerators (GPUs), and new hardware architectures based on low-power consumption processors (e.g. ARM). C, C++ and Python have been chosen as the programming languages for the main components of the software suite. CUDA has been used to port the most computationally demanding algorithms to GPUs.

## 2.1 Speed and low power consumption goals

Among the goals considered in the overall design, we mention the ones related to performance, in speed and low power consumption:

- On/off-site processing: ASciSoft can process data both on-site and off-site;
- On-line processing: ASciSoft can perform an on-line data reduction during data taking in order to be able to generate real-time performance and scientific monitoring alerts;
- Low-power consumption and parallel processing: ASciSoft can perform data reduction by means of low-power consumption and parallel computing processors (ARM/GPUs), besides conventional CPUs;
- Programming languages: ASciSoft is written in C++ and Python (for data processing with conventional and ARM CPUs) and CUDA (for GPU processing);
- I/O data format: ASciSoft can handle the standard Flexible Image Transport System (FITS) data format (following the NASA-OGIP standards) for input/output (I/O) operations;

# 3.Architectures

The requirement of a fast data processing in an off-site datacentre is critical because of the costs related to the data transmission; an effective solution to decrease the power and bandwidth required by the array facility could be to deploy an embedded module on each telescope, allowing it to carry out preliminary data reduction on its own before sending them to a central acquisition system.

To meet this requirement, the development has started from the very beginning using both C/C++ and CUDA, on a Nvidia TK1 board [2], a product specifically designed to meet low power consumption requirements (a brand new TX1 board [3] has been purchased less than a month ago, with very nice results obtained in porting our software).

The key design choices can be summarized in

1.Processing chunks of events

- A typical file from the camera is 55000 events, weighing 500 MB. This size has been decided in agreement with the camera team and maximizes the utilization of GPU on the Nvidia TK1 board;
- Parallelization can be made per events besides per-pixel; since events are much more than pixels, this comes handy when using GPUs, because it allows a better latency hiding.

2.Independent executable modules

- Less inter-dependency, more maintainability.

## 3.1 Hardware

The hardware currently in use, where the prototype has been developed and tested, is:

- Dual processor Intel Sandy Bridge @ 2GHz with 16 physical cores and 128 GB RAM;
- GPU Nvidia Tesla K20c (Kepler series);
- 8 disk slots of 4TB each
- Direct link and share with the ASTRI archive and storage system

Tests for low power consumption have been conducted on Nvidia Jetson TK1 board:

- Heterogeneous System-on-Chip;
- CPU: Quad-core ARM

- GPU: Kepler Architecture, 1 Streaming Multiprocessor;
- CUDA 6.5 developer toolkit;
- RAM: 2 GB shared between CPU and GPU:
- OS: Ubuntu 14.04 Linux for Tegra (L4T)
- I/O: SATA 3 Gb/s HDD

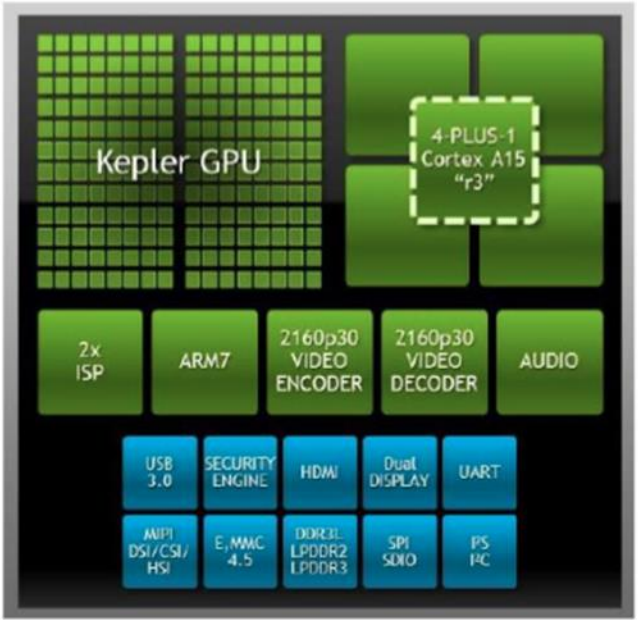


*Fig 1: Nvidia Kepler GPU architecture*


*Fig 2: Nvidia Jetson Board: it ships an ARM CPU + Nvidia TK1 processor, with Kepler architecture*

# 4.Data flow

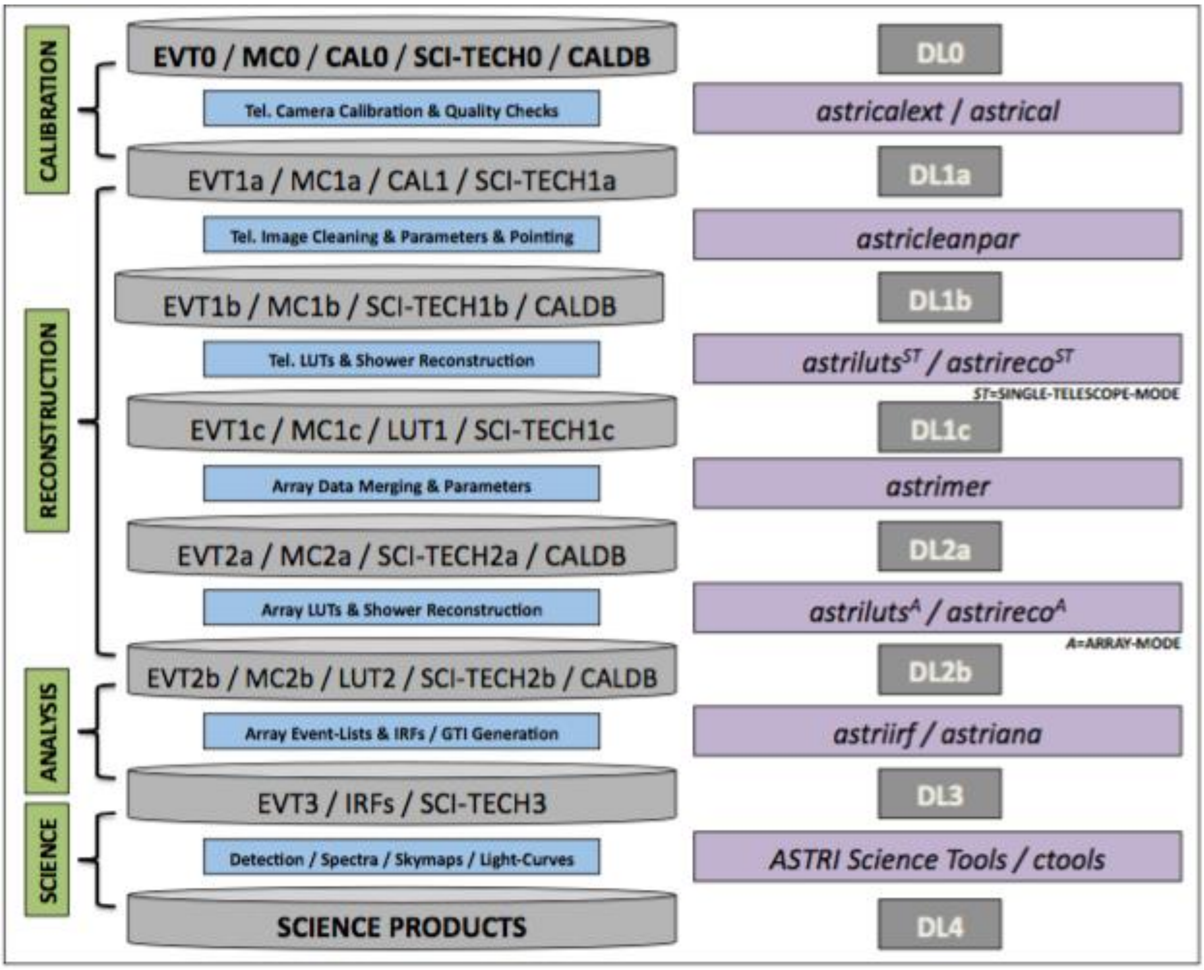


*Fig 3: Functional design layout of the ASTRI data reconstruction and scientific analysis software (A-SciSoft). Functional breakdown stages are shown as green boxes, basic software components as blue boxes, data processing executable modules and packages as purple*

*ASciSoft* functional design is the following [1]:

1. Calibration (DL0 -> DL1a): It takes DL0 data (telescope-wise) and reduces them to DL1a calibrated data;
2. Reconstruction (DL1a -> DL2b): It takes DL1a calibrated data (telescope-wise) and reduces them to DL2b fully reconstructed data (array-wise);
3. Analysis (DL2b -> DL3): It takes DL2b data and reduces them to DL3 data via optimization and selection algorithms;
4. Science (DL3 -> DL4): It takes DL3 data and generates high-level scientific products by means of the adopted Science Tools.

This data flow is compliant with the logical design and data model of CTA pipelines [4].

# 5.Tests

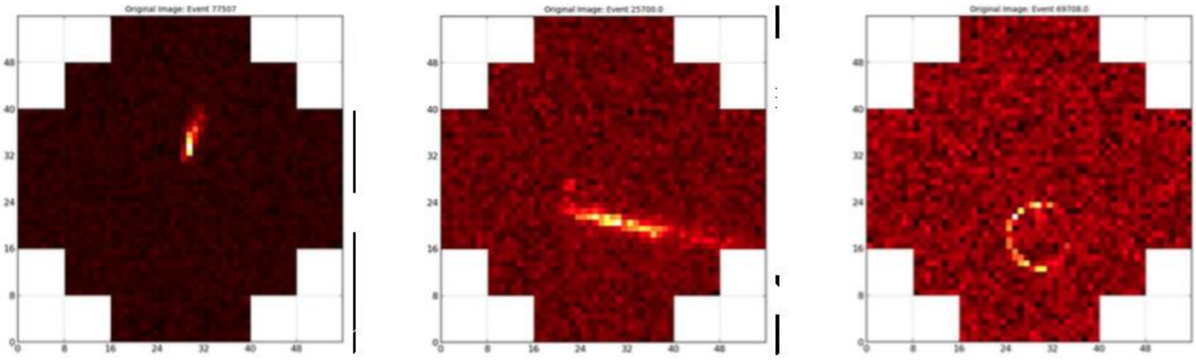

*Figure 4: Visualization of three simulated events as seen by ASTRI SST-2M camera. The camera is composed of SiliconPhoto Multiplier (SiPM) sensors with a total of 1984 connected squared pixels with angular size of 0.17 ° each. The leftmost image is a gamma-like event, the central one a proton, whereas in the rightmost image a ring produced by a muon can be seen.*

As mentioned above, ASciSoft processes chunks of 55000 events, weighing almost 500 MB.

For an input datum like this, one has:

- 110 seconds of nominal acquisition rate at 500 Hz, 550 seconds at 100 Hz;
- 81% of events survives cleaning with default settings

This is also the reference test case so far, but the whole pipeline is currently being tested on a real dataset with one crab unit of gamma events reconstructed from Monte Carlo simulated data (from SimTelArray), for a total of almost 6 hours of observation at 100 Hz (for protons) [5].

## 5.1 Low Power Unified Modules

In this section, we show some results on the given test case, with the already mentioned architecture

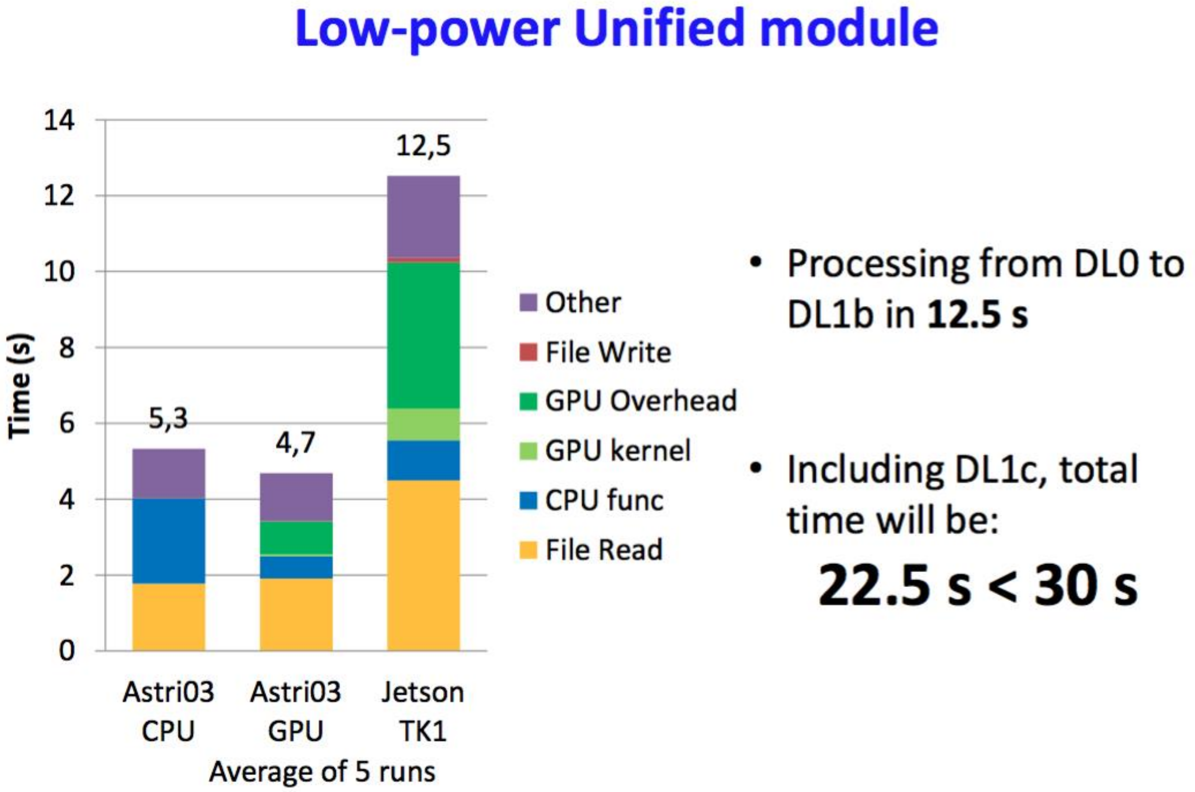


*Fig 5: comparison between time performances on Astri03 (workstation) and Jetson TK1 (low power board). Performances on the low power board are quite poor in term of execution time, but still acceptable*

In the above figure, a comparison of the unified *astrical* + *astricleanpar* modules performances on different architectures can be seen.

On the workstation (Astri03), the GPU version performs better than the CPU one, as expected - the gain in execution time is small, and this needs a further investigation - while on the Jetson low power board the overhead of GPU is far more present: in particular, data transfer between host and device memory as well as I/O operations are clear bottlenecks.

Preliminary results with the new TX1 model show a further improvement to 8.5 s Unified Module wall time execution.

Nevertheless, regarding power consumption following good results were achieved:

| | Workstation | Jetson TK1 | Improvement |
|---|---|---|---|
| Time to solution | 6,5 s | 22,5 s | 0,28 x |
| **Energy to solution** | **1430 J** | **225 J** | **6,3 x** |
| Energy delay product | 9295 Js | 5062 Js | |
| Event/s | 7690 evt/s | 2200 evt/s | |
| Energy/event | 29,6 mJ | 4,5 mJ/evt | |

## 5.2 Tests on an ARM + GPU workstation: ARKA by E4

ASciSoft has been run on an ARKA workstation equipped with Octa-core ARM APM X-Gene1 CPU and an NVIDIA K40 obtaining for the unified module an improvement in execution time of about 1 seconds with respect to our workstation, thus demonstrating the feasibility of a full Cherenkov analysis on an ARM+GPU server. Unfortunately, no power consumption measurements have been possible.

# 6.Ongoing and future activities

Currently, we have released an alpha version of ASciSoft for internal testing, which will be followed by a beta release available to the whole ASTRI team; packaging has been released via conda, a tool coming with the open platform from Continuum Analytics [6].

While giving support and maintenance for ASciSoft, which was specifically developed for the ASTRI prototype and mini array, we planned a roadmap which includes following main activities:

- Benchmark of Machine Learning methods for event reconstruction on CPU, GPU and Jetson TX1.
- Ongoing developments by GPU manufacturers on fast and energy efficient libraries for neural network deployment (e.g. Nvidia TensorRT) are a field to explore in order to exploit low-power devices for deep learning image inference in gamma-ray astronomy [7].
- Test of frameworks specific to Machine Learning and Deep Learning.
- Hillas parameterization on GPU.

We have some preliminary results regarding gamma hadron classification, where we compared the behavior of a Random Forest method towards a deep neural network with two hidden layers.

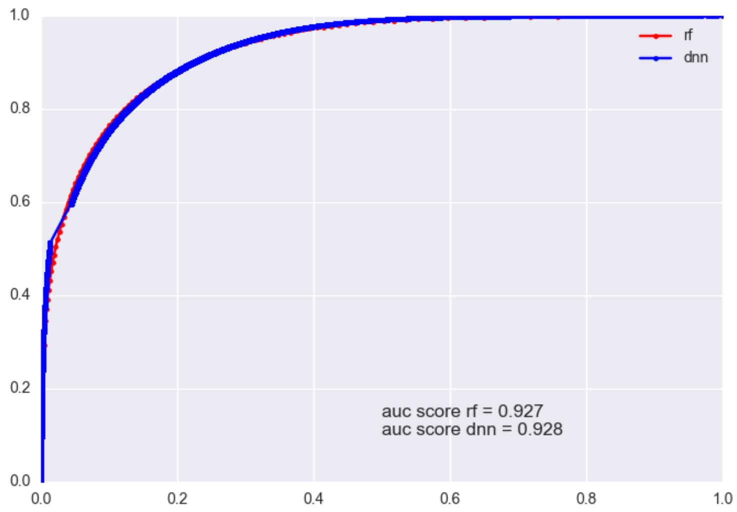


*Fig 6: ROC curve for gamma acceptance: Random Forest vs Deep NN with two hidden layers*

Figure 6 shows the ROC curve with the fraction of true positive gammas on *y* axis and the fraction of false positive gammas on *x* axis. Data are referred to a Monte Carlo dataset, we predicted the probabilities for the events collected by eight telescopes in the simulated array.

It can be seen that the two methods are somewhat equivalent, a direct measure being the score, namely the area under the curves. Similar results have been obtained with a shallow neural network, with an only hidden layer.

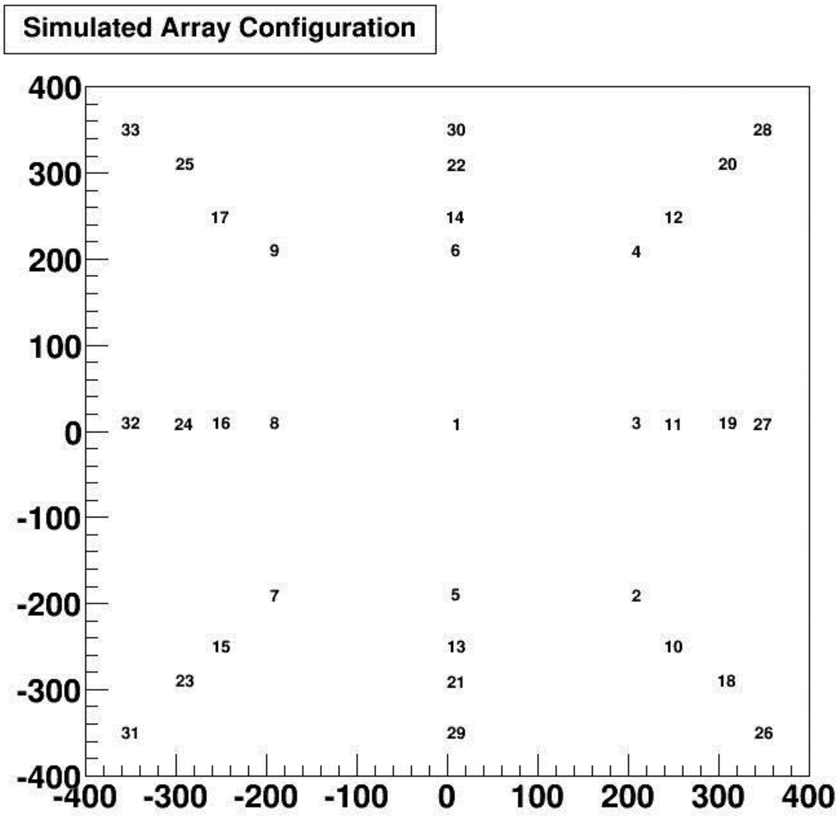


*Fig 7: Simulated ASTRI array configuration*

## 6.1 Simulations

The development of the ASciSoft requires the simulation of events as closer as possible to the expected data that will be collected by the ASTRI-SST-2M telescope prototype and later by the ASTRI mini array. The development of the atmospheric showers initiated by cosmic particles hitting the Earth atmosphere, the emission of the Cherenkov light by the ultra-relativistic particles in such showers and its propagation through the atmosphere up to the telescopes need to be carefully simulated by Monte Carlo technique. The standard CTA package for the simulation of the shower development in the atmosphere (CORSIKA version 6.99 [8]) has been used along with the CORSIKA IACT/ATMO add-on package [9] which provides a more customizable simulation of the atmospheric transmission, a flexible interface for any arbitrary configuration of Cherenkov detectors and a machine-independent output format for the Cherenkov photons reaching the assumed positions of the Cherenkov detectors. This package allows also to reuse each event several times by randomizing the shower core. In this way, the amount of CPU needed to simulate the required number of events can be reduced considerably while introducing a negligible correlation between events.

Both showers initiated by primary gammas, which constitute the expected signal, and by background particles, like protons, electrons and heavier nuclei, need to be simulated to develop and test the ASciSoft package as well as to estimate properly the expected performance of the different configurations of telescopes under study.

It's worthwhile to notice that thanks to the very high rejection of hadronic background achieved with the imaging atmospheric imaging technique, huge samples of simulated hadronic events are needed to achieve statistically significant estimates of the performance and to carefully optimize the analysis chain. More than $10^8$ cosmic ray induced atmospheric showers for each telescope configuration are needed to properly estimate the array sensitivity, energy and angular resolution requiring extensive computing needs in term of both disk space and CPU power. The simulation of such amounts of needed atmospheric showers has been realized exploiting the GRID technology [10].

Showers produced by gamma primaries from a point-like source at 20° zenith angle and diffused protons at the same zenith angle within a cone with 6°radius have been simulated. The main characteristics of the simulated atmospheric showers are summarized in the following table:

| Primary particle | Min Energy [TeV] | Max Energy [TeV] | Maximum Impact Distance [m] | Maximum Offset Angle [deg] | Simulated Events |
|---|---|---|---|---|---|
| Gamma | 0.1 | 330 | 1200 | 0 | 24.95E6 |
| Proton | 0.1 | 600 | 2000 | 6 | 699.8E6 |

No showers induced by heavier nuclei nor by electrons have been generated so far. The energy spectral index used for both primaries is −2.0 (rescaled to −2.62 for gammas and to −2.7 for protons in the analysis step) to equally distribute the CPU time over the entire energy range. The telescope response has been simulated with the sim_telarray package [11] which is the simulation code adopted by the CTA consortium. Such package carefully simulates photon propagation through the optical system of the telescopes taking into account shadowing elements, reflections by primary and secondary mirror, transmission through the camera cover and photon conversion in the photon-detectors. The sim_telarray code simulates also the electronic signal produced by the photon-detectors and the propagation of such signal through the electronic chain. The trigger logic adopted to select events is carefully simulated too. The simulation of the electronic chain has been customized to simulate properly ASTRI front-end electronics which is quite different from that of any other Cherenkov telescope.

## 6.2 Low-power architecture work

This section provides some highlights concerning low-power architectures R&D in which ASTRON is involved. It concerns work that has been conducted in the recent past; it also includes ongoing work. The research is not only aimed at platform level, it also includes somewhat higher-level system design considerations, as power minimization can best be tackled starting from a system level perspective. The activities reported here are mainly based on the specific needs of SKA project and its precursors. Nevertheless, adopted solutions are well suited also for the needs of other ESFRI projects in which low-power consumptions hardware is requested for massive on-site and real time data reduction and analysis. In particular, as shown in the previous chapters, CTA is one of the ESFRI projects having a very similar challenging approach to a massive data rate acquisition and real time analysis.

### *a. The Dome project*

At the ASTRON & IBM Centre for Exascale Technology (www.dome-exascale.nl) in Dwingeloo and in the IBM-ZRL laboratories, we are looking into low-power architectures aimed at addressing the computing challenges the Square Kilometre Array radio telescope (SKA) faces. This work is conducted in the Dome project, where we consider this challenge at different levels, including algorithms, data transport, data storage and pre-fetching, (holistic) architecture exploration. References can be found on the Dome web page, and concerning low-power architectures we would like to mention the following project topics.

### *b. Accelerators in the Dome project*

We have compared a wide range of accelerators in terms of performance, energy efficiency, and programming effort. This was done by means of implementing a signal-processing pipeline that processes data from a radio telescope. We optimized the algorithms from this pipeline on GPUs from AMD and NVIDIA, the Intel Xeon Phi, a Digital Signal Processor (DSP) from Texas Instruments, and a dual Intel Xeon CPU as reference platform. The algorithms have quite different compute and memory access characteristics, hence we have assessed various subsystems of the accelerators. Results, showing for example that GPUs yield the best performance and energy efficiency, can be found in [8]. It also turns out that the differences between the (energy) efficiencies of the accelerators are large. More specific, results on accelerated imaging by implementing an efficient image domain gridder can be found in [21]. In a follow-up project (funded by the Netherlands eScience Center), we will explore Intel's OpenCL/FPGA framework. And finally, in parallel we are also looking into using the Compaan framework (www.compaandesign.com/) to reduce programming effort.

### *c. Dome microservers*

Within the Dome project (mainly at ZRL) the team has developed the microserver system aimed at creating a “datacentre in a box” or a “micro datacentre”. Very low power operation is achieved by integrating the electronics into a very small form factor, by placing the different electronic components such a memory close together, and also by using “hot” water to cool the system. Publications describing the concept and prototypes include [15] [16] [17]. We have created microservers based on PowerPC architectures, FPGAs, and recently, ARM.

### *d. Dome ASICS*

We have looked into a power efficient ASICs approach for the basic building blocks of the radio astronomical signal chain. In [20] a scalable and versatile architecture is presented, that uses only three chip architectures, and that in principle can be used as a building block for SKA phase one and two.

### *e. Dome “holistic” architecture modelling and exploration*

We have created an analytic model that takes as inputs a parametric microarchitecture-independent characterization of the target workload, and a hardware configuration of the core and the memory hierarchy. It yields an estimation of processor-core performance. We have compared our performance estimates with measurements on an Intel® Xeon® system. We conclude that our model enables fast design space exploration and represents a first step towards an analytic exascale system model [14]. A higher-level architecture model for the SKA signal processing challenges is presented in [13].

### *f. AARTFAAC*

The Amsterdam-ASTRON Radio Transients Facility And Analysis Center (AARTFAAC) all sky monitor is a sensitive, real time transient detector based on the Low Frequency Array (LOFAR). It generates images of the low frequency radio sky with a time cadence of a few seconds. We implemented a heterogeneous, hierarchical design to manage the 120 Gbps raw data rate, and large scale computing to produce real-time images with minimum latency. In [18] we present the implementation of the instrumentation, its performance, and scalability.

### *g. Accelerators in the DAS project*

The Dutch Advanced School for Computing and Imaging has built five generations of a distributed system for computer science research. The system has supported years of award-winning research, underlining the benefits of investing in a smaller-scale, tailored design. The cluster at ASTRON was used extensively to prototype applications that exploit new accelerator technologies. These applications, which included two GPU correlators, a beam former, an imager, and a pulsar pipeline, are now used in production by the LOFAR radio telescope. A summary of the DAS work can be found in [12].

# 7. Conclusions

CTA and SKA and their precursors are, among the ESFRI projects with the most challenging instrumentations in terms of huge amount of data to be elaborated in real time and directly on site in order to drastically reduce both transfer rate and storage dimension. New computing technologies and innovative architectures alternative to traditional x86 CPUs or similar can be largely adopted in order to perform fast data elaboration with a lower power consumption. The goal of OBELICS is to identify and benchmark innovative computing architectures adopting as use cases the requirement of CTA and SKA and their precursors in order to extend in a second step this test also to other ESFRI projects. In this part of the project we performed many tests on innovative hardware with low-power consumption within ASTRI for CTA and DOME for SKA projects. We developed dedicated data analysis software or adapted existing software in order to perform tests and benchmarks on new architectures also in collaboration with hardware producers. We also used simulated data and or data provided by precursors already in operation. Obtained results show that the use of innovative computing architectures is effectively able to reduce massive data amount in short time and with a very low (more than a factor of 10 less) power consumption.

# 8.References


1. Saverio Lombardi et al, *Proc. SPIE* 9913, Software and Cyberinfrastructure for Astronomy IV, 991315 (July 26, 2016); doi:10.1117/12.2231656
2. Nvidia TK1: http://www.nvidia.com/object/jetson-tk1-embedded-dev-kit.html ;
3. Nvidia TX1: http://www.nvidia.com/object/embedded-systems-dev-kits-modules.html
4. Contreras, J. L., et al., "Data model issues in the Cherenkov Telescope Array project," Proc. 34th ICRC, (2015).
5. Michele Mastropietro et al, *Proc. SPIE* 9913, Software and Cyberinfrastructure for Astronomy IV, 99133V (July 26, 2016); doi:10.1117/12.2232003;
6. https://www.continuum.io/
7. Nvidia whitepaper: *http://www.nvidia.com/content/tegra/embedded-systems/pdf/jetson_tx1_whitepaper.pdf*
8. D. Heck et al., Forschungszentrum Karlsruhe Report No.FZKA 6019 (1998).
9. https://www.mpi-hd.mpg.de/hfm/bernlohr/iact-atmo/iact_refman/index.html
10. L.Arrabito et al. Journal of Physics Conference Series, 396(3) (2012) 032007
11. K. Bernlöhr, Simulation of Imaging Atmospheric Cherenkov Telescopes with CORSIKA and sim_telarray, Astroparticle Physics 30 (2008) 149–158. [arXiv:0808.2253].
12. H. Bal and D. Epema and C. de Laat and R. van Nieuwpoort and J. Romein and F. Seinstra and C. Snoek and H. Wijshoff, "A Medium-Scale Distributed System for Computer Science Research: Infrastructure for the Long Term," in Computer, vol. 49, no. 5, pp. 54-63, May 2016. doi: 10.1109/MC.2016.127
13. R. Jongerius, S. Wijnholds, R. Nijboer and H. Corporaal, "An End-to-End Computing Model for the Square Kilometre Array," in Computer, vol. 47, no. 9, pp. 48-54, Sept. 2014., doi: 10.1109/MC.2014.235
14. R. Jongerius, G. Mariani, A. Anghel, G. Dittmann, E. Vermij and H. Corporaal, "Analytic processor model for fast design-space exploration," 2015 33rd IEEE International Conference on Computer Design (ICCD), New York, NY, 2015, pp. 411-414. doi: 10.1109/ICCD.2015.7357136
15. R. P. Luijten, M. Cossale, R. Clauberg and A. Doering, "Power measurements and cooling of the DOME 28nm 1.8GHz 24-thread ppc64 µServer compute node," 2015 International Conference on IC Design & Technology (ICICDT), Leuven, 2015, pp. 1-4. doi: 10.1109/ICICDT.2015.7165919
16. R. Luijten, D. Pham, R. Clauberg, M. Cossale, H. N. Nguyen and M. Pandya, "4.4 Energy-efficient microserver based on a 12-core 1.8GHz 188K-CoreMark 28nm bulk CMOS 64b SoC for big-data applications with 159GB/S/L memory bandwidth system

density," 2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers, San Francisco, CA, 2015, pp. 1-3.

17. R. P. Luijten, "Microserver system use of optical technologies," 2016 Optical Fiber Communications Conference and Exhibition (OFC), Anaheim, CA, 2016, pp. 1-3.
18. Peeyush Prasad, Folkert Huizinga, Eric Kooistra, Daniel van der Schuur, Andre Gunst, John Romein, Mark Kuiack, Gijs Molenaar, Antonia Rowlinson, John D. Swinbank and Ralph A.M.J Wijers, "The AARTFAAC All Sky Monitor: System Design and Implementation", Journal of Astronomical Instrumentation, Volume 5, Number 4, December 2016
19. J. W. Romein, "A Comparison of Accelerator Architectures for Radio-Astronomical Signal-Processing Algorithms," 2016 45th International Conference on Parallel Processing (ICPP), Philadelphia, PA, 2016, pp. 484-489. doi: 10.1109/ICPP.2016.
20. M. L. Schmatz et al., "Scalable, efficient ASICS for the square kilometre array: From A/D conversion to central correlation," 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Florence, 2014, pp. 7505-7509, doi: 10.1109/ICASSP.2014.6855059
21. Bram Veenboer, Matthias Petschow, and John W. Romein, "Image-Domain Gridding on Graphics Processors", IEEE International Parallel and Distributed Processing Symposium (IPDPS'17), Orlando, FL, May 29-June 2, 2017