GFORC:
A fast expandable and cheap computer cluster devoted to Geant4

Memmo Federici[1], Bruno Martino[2], Simone Lotti[1]

[1]INAF- IASF Istituto di Astrofisica Spaziale e Fisica Cosmica Roma, [2] CNR – IASI Istituto di Analisi dei Sistemi ed Informatica

email contact: memmo.federici@iasf-roma.inaf.it ; bruno.martino@iasi.cnr.it ; simone.lotti.@iasf-roma.inaf.it

Abstract

For the development of new X-rays detectors as well as for simulation of interactions between particles and matter, very often the scientific community makes his choice on Geant4. This is a very versatile and powerful tool, but some applications can require substantial computing resources. In particular simulations in which is necessary to use a large number of particles can take an extremely long time (sometimes many weeks). To expedite the simulation process we have developed a Cluster: GFORC who currently consists of 4 PCs for a total of 16 GB of RAM 32 cores 6 TB of data storage. The Cluster is easily expandable and thanks to the adopted design criteria achieves an excellent compromise between performance and cost. The management software of Cluster, allows to divide the single instance of the simulation on the available processors, creating a simulation speed similar to that obtained on "parallel systems". Furthermore the use of Geant4 on GFORC is facilitated by a series of custom graphical user interfaces. The Cluster improves the speed of the simulation used as a test, about a factor of 20 compared to the times obtained by a modern PC.

The priorities in GFORC cluster designing were chosen by low cost and high performance computing. To achieve this goal it was decided to use only conventional and readily available components. The individual nodes that make up the cluster are of a commercial personal computer, then low-cost hardware with good quality, but conventional.

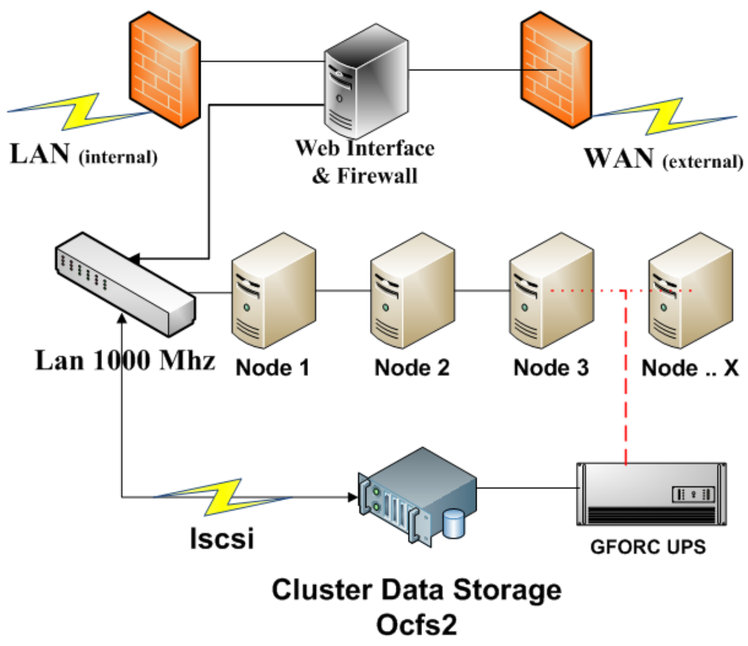

Software Strategy

GFORC is a calculation tool that is able to split the simulation program in so many programs running simultaneously on the cores of the allocated nodes. Thanks to this mechanism, the result obtained in terms of speed is comparable to the speed of calculation made on parallel platforms.

The great advantage of GFORC is the ability to reuse software written in serial with results very similar to implementations on parallel platforms.

To efficiently use the features of the calculation in clustering and to facilitate their use have been product a series of software procedures that implement the following functions:

•Allocation of resources

• Launch of specific tools

• Synchronization instances subjected to the calculation tail

Thanks to the strategies adopted in developing the GFORC software, the obtained performance computing, are a factor 20 above the ones obtained with a workstations with single core.

It is important to underline that the cost of the current structure of GFORC is only about 4 times greater than the cost of a single workstation used as a comparison.

The flow chart at right shows one of the main procedures for managing cluster: the login procedure. This is primarily designed to allocate the nodes of the cluster and set the environment variables required for simulations. Also is shown one of the key features of GFORC which is the ability to resume, controlled by the users, active sessions still in calculation. This essentially allows users to run multiple programs, simulations, simultaneously. Thanks to this function the user can choose at next login, which of the sessions calculation still running wants to resume. This function also allows the user to disconnect its connection to the Cluster in all those cases in which the simulation is very long.

Graphical User Interface (GUI)

The using of GFORC is facilitated by the implementation of some graphical user interfaces. Furthermore through these interfaces, users can interact with the simulation programs by changing the input parameters and saving the settings for subsequent reuse. Below are shown (as an example) some of the graphical interfaces available.

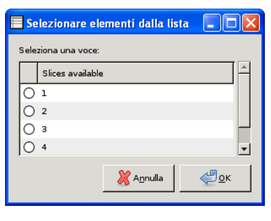

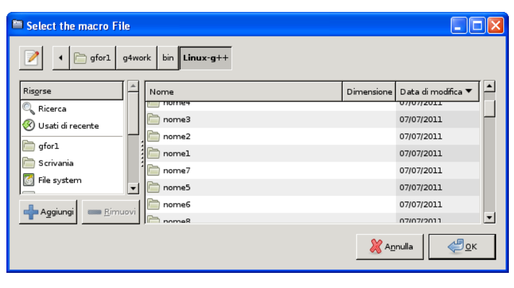

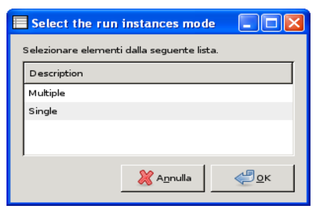

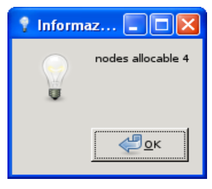

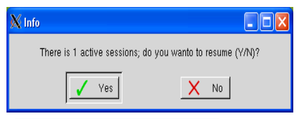

The diagram shows the current configuration of GFORC. It is at present composed of 4 nodes, each equipped with Intel *"I7"* Eight-Core processor, 4 GB of DDR3 RAM by 1333 Mhz, and 500 GB of Hard Disk. Overall GFORC develops with its 16 GB of RAM and 32 cores computing power close to 100 Giga Flops (100 x 10^9 Floating point Operations Per Second). The cluster is managed by the resource manager SLURM (Simple Linux Utility for Resource Management). Simulations data results reside on a server with NAS RAID-6 storage capacity of 6 TB. The DATA STORAGE file system is realized by OCFS2 (Oracle Cluster file system). This is a clustered file system which is shared by being simultaneously mounted on each node of Cluster. The OCFS2 is used by nodes in the cluster through the ISCSI transport protocol (Internet Small Computer System Interface) which allows the use of the existing TCP/IP infrastructure and produces a great gain in speed communication and data exchange between nodes. This choice is extremely performant compared to the use of a file system like ext3 or ext4 by a NFS (Network File System protocol). Thanks to the choices made the speed of read/write on the disk that containing the users area *"home"* has increased by an order of magnitude. Access to GFORC is via a secure connection in the *"ssh X"* protocol which allows users to export graphical windows. The Cluster has an individually controlled power supplies (UPS) controlled via WEB. The power of 10kVA, can guarantee in the absence of mains power, an endurance of several hours for a load of up to 60 nodes. Is currently studying a system of user access through a web interface.

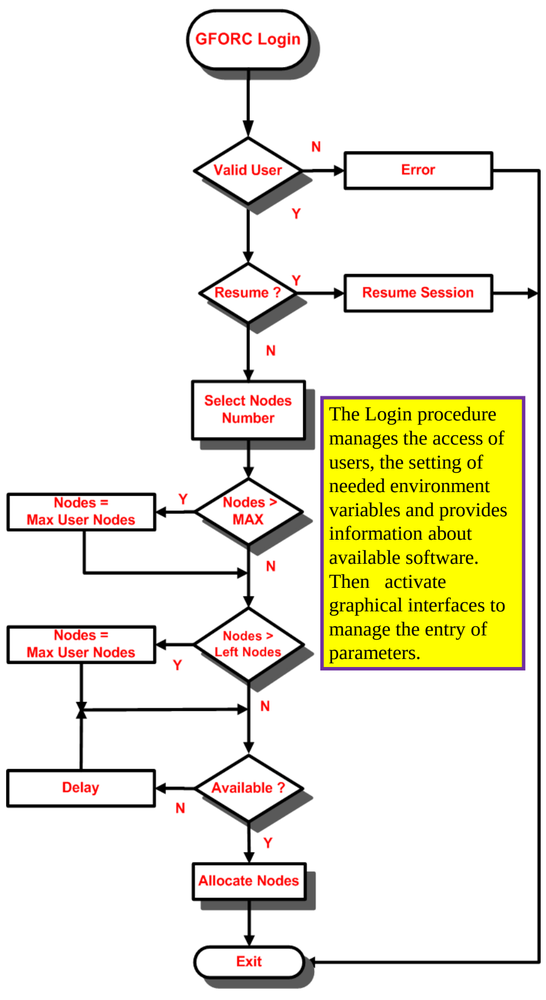

The Login procedure manages the access of users, the setting of needed environment variables and provides information about available software. Then activate graphical interfaces to manage the entry of parameters.

Simulation Results

We present here the results, obtained with GFORC, of a study on the impact of L2 cosmic particles with galactic (CGR) and solar origin on the X-ray Microcalorimeter Spectrometer (XMS) onboard of the ATHENA mission (ESA). We used a Monte Carlo based simulator to investigate the rejection efficiency of the anticoincidence system and the impact of residual background on the detector performances. Similar GEANT4 simulations have already proven to able to reproduce the background measured by the XRS microcalorimeter flown on Suzaku (Perinati et al. - Journal of Low Temperature Detectors, accepted paper).

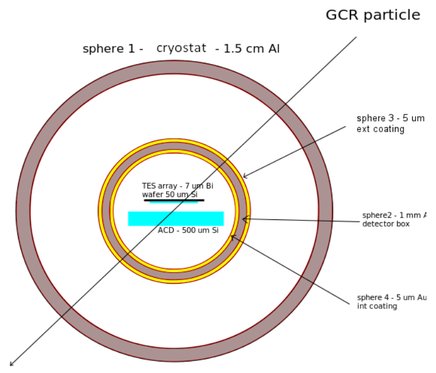

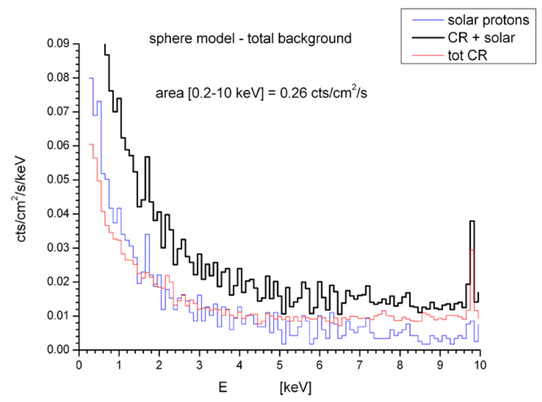

We initially prepared a simple geometrical model sphere model where we assumed a simplified geometry composed by 4 concentric spheres. This is a standard approach, often used in background simulations for space applications (H.Murakami, 2006). Simulations on the sphere model has been used to check the right response of the simulators, requiring low time machine computation, and has been run on a PC with standard performaces.

The spectrum of unvetoed events (100 eV bin), with contributions from different input components highlighted, obtained with the sphere model

GFORC allowed us to exploit a factor by 20 higher in computational power, so we decided to prepare a second model – detailed model – where the cryostat is modeled in great detail, based on the ATHENA specifics.

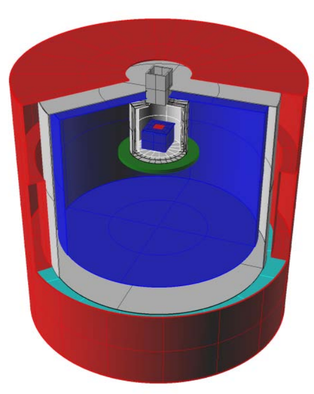

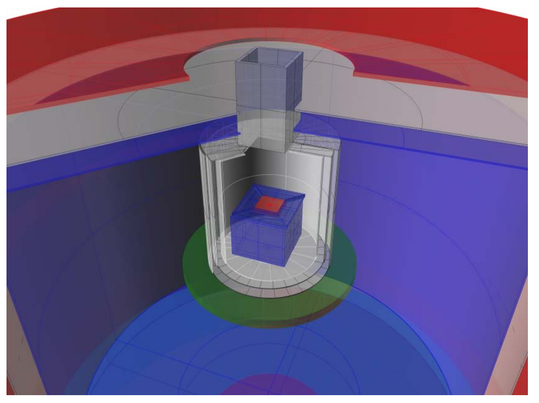

The Geant4 detailed model of the cryostat:
Left - Three Al external thermal shields are implemented (red, gray and blue).
Right – the innermost part of the cryostat: the last stage cooler plate (green), the internal shielding composed of three grey layers (Cryoperm-Aluminum-Niobium from outside to inside). Inside the last shielding the detectors (red), and their supports (blue).

Simulations with the detailed model allowed to:

• identify the geometrical origin of the background

• characterize correctly the weight of each background component

and therefore to find the appropriate solutions for background reduction.

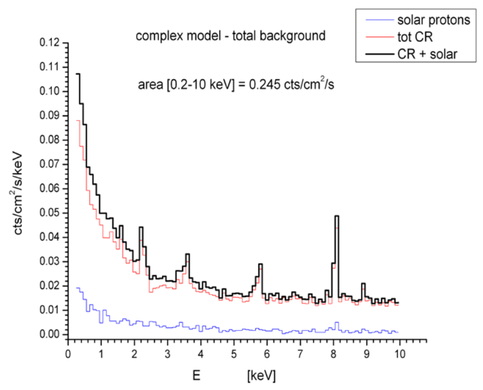

The spectrum of unvetoed events (100 eV bin) with contributions from different input components highlighted, obtained with the detailed model.

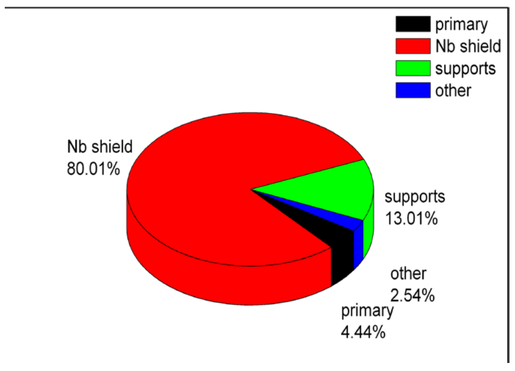

The geometrical origin of unrejected particles. Note that a new component of the background has arisen from the insertion of the supports.

Thanks to the results obtained with GFORC the Italian ATHENA consortium have started to analyze potential area of improvement for the background. Providing a passive layer for electron shielding can allow to significantly reduce the background level. Preliminary simulations have predicted a background drop of about a factor 50% with a 250 μm kapton liner in the Nb shield. Another option under investigation, that promises to reduce even more substantially the residual background, employs a thin filter very close to the array surface to stop low-energy electrons.