



Enabling Grids for E-scienceE

Porting of Astronomy & Astrophysics and Quantum Chemistry Applications

IFCA - CSIC
Santander

www.eu-egee.org



- **CSIC: Spanish National Research Council**
 - Largest public entity in Spain dedicated to research
 - Multidisciplinary (<http://www.csic.es>) :
 - Physics, Chemistry and Materials Science
 - Biology, Biomedicine and food sciences
 - Environmental Sciences.
 - CSIC has 124 research institutes distributed across Spain
 - The **Institute of Physics of Cantabria-IFCA** and the **Institute of Corpuscular Physics-IFIC** are involved in EGEE-III
 - Grid applications activities at IFCA (Santander)
 - Tier-2 of CMS experiment
 - Astrophysics and Environmental Sciences support
 - Grid applications activities at IFIC (Valencia)
 - Tier-2 of ATLAS experiment
 - Medical imaging applications

- **Support to MPI and i2glogin**

CSIC as coordinator of inteugrid has developed a strategy of support for **parallel MPI apps** and for **interactive job submission**

- **Middleware deployment**
 - *mpi-start*
 - *i2glogin*
- **Infrastructure support**
 - Low latency intranet cluster: Infiniband
- **Documentation for user and applications developers (wiki)**


- **Support to our user communities**

- **Astrophysics (AA cluster)**
 - Planck mission
 - User induction (XMM group)
- **Fusion**
 - SA1 coordinators in EUFORIA (FP7 Jan 08 – Dec 10)
- **Environmental research**
 - Modelling of coastal landscapes
 - Grid instrumentation: networks of sensors in the framework of DORII (FP7 Feb 08 – June 10)
- **Medical Applications**
- **Quantum Chemistry**
- **Support to our HEP community (CMS and ATLAS Tier-2)**


- **Single Frequency:** we have ported to the GRID a code to detect point sources using the Mexican Hat Wavelet 2. We have used the GRID to test it with realistic simulation of Planck. Its performance is very good and it has been installed in the Planck Data Processing pipelines for the High Frequency and Low Frequency instrument and it is being used during operations.
- **Multidimensional parameter space tests.** Hundreds of test simulations needed.

Parameter Space: Image Size (7 Sq. Deg – 14 Sq. Deg.)


- Pixel Size (Planck nominal sizes)
 - Overlap between patches (1 - 5 Deg.)
 - Border (5-15 pixels)
 - Matching Radius (FWHM)
- In addition, we have tested this tool in data from NASA's WMAP satellite producing a new catalogue of extragalactic point sources at 23, 33, 41, 61 GHz, with 25% more objects than the WMAP one. These results have been recently published MNRAS, 2009, 392, 733.
 - **Status:** we continue with the testing and development of the code with even more realistic simulations and with Planck data.




+



**Filtering
MHW2**




→



- **Characteristics:**
 - Healpix Pixelization Scheme Libraries: HEALPIX
 - FITS Format I/O library: CFITSIO
 - Fortran 90
 - Static Compilation: 2MB executable
 - Not yet parallel
- **Requirements:**
 - Full Planck resolution images: 200 MB each (9 Planck Frequencies) (to be copied from the local Storage Element)
 - 1 GB RAM
 - 2 MB executable: to be copied from the local Storage Element.
 - 10 Kb parameters file: to be included in the input sandbox
- **Execution time:**
 - Aprox. 90 minutes in 1 3.65 GHz Intel Xeon CPU depending on the Freq.
- **Output:** Catalogue of objects (position in the sky, flux estimation, error, etc.)
 - very small (a few MB's)

- Multi-Frequency:** We have ported to the GRID a new code to detect point sources using multi-frequency information. This new algorithm is giving very good results and we are investigating it further in the GRID using LFI Planck realistic simulations. The results have been published MNRAS, 2009, 394, 510.


Results for full-sky 30, 44, 70 and 100 GHz PSM simulations



- We have ported to the GRID a tool to detect clusters of galaxies using the SZ effect, an interaction between the photons from the CMB and the hot ionized gas inside the clusters. This interaction produces a decrement in the CMB signal below 217 GHz and an increment above, the frequency dependence is very well known and can be used to design a filter known as Matched Multifilter to preserve only those signal in the maps between 30 and 857 GHz with this frequency law.
- We have extensively tested this algorithm using the GRID and is the best example so far to show the power of the GRID. The analysis of 9 Planck maps at full resolution using a single CPU could last 1 month, but we have managed to do this analysis in 1 day. This is also an ongoing work and a production version of this code has been installed in the Planck DPC.
- We plan to publish the results from a comparison exercise within Planck working groups in the next months.



- **We are porting to the GRID three algorithms to detect non-gaussianity in maps of the CMB.**
- **These algorithms are:**
 - Spherical Mexican Hat Wavelet
 - Minkowski functionals
 - Scalars on the sphere
- **Some of this codes have been ported to the Data Processing Center in Trieste and they can be used to check the products of Planck, for example, the CMB map.**
- **To do a proper non-gaussianity analysis 10.000's of gaussian simulations have to produced and the systematics of the experiment properly understood.**
- **To do this kind of simulations we will use the GRID with an ambitious simulation plan using the LevelS, a Planck simulation software that we plan to port to the GRID.**

- **SIESTA (Spanish Initiative for Electronic Simulations with Thousand of Atoms)**
 - Method and implementation for electronic structure calculations
 - Ab initio molecular dynamic simulations of molecules and solids using Density Functional Theory (DFT)
 - Also able to compute molecular mechanics (similar to GROMACS)
- **CPU intensive task**
 - But designed to scale linearly
 - Fortran 90 code
 - Parallel and serial versions available
 - Input and Output files relatively small




<http://www.icmab.es/siesta/>

- **SIESTA dependencies:**
 - BLAS & LAPACK (serial) BLACS & ScaLAPACK (parallel)
 - We are developing an automated way of obtaining and compiling the software
- **We developed helper scripts that create JDL, sandboxes and submits to grid on behalf of the user**
- **Initial Tests of scalability:**



- **MPI support in EGEE is a longstanding issue**
 - Older WMS have hardcoded MPICH support
 - Sites have to use workarounds to this hardcoded startup
- **Configuration for MPI sites:**
 - Sites should follow: http://egee-uig.web.cern.ch/egee-uig/production_pages/MPIJobs.html
- **MPI-START initiates the jobs:**
 - Allows user to avoid low-level details
 - Allows sites to customize their MPI support

- Specifies a unique interface to the upper layer in the middleware to describe MPI jobs
- Support basic file distributions
- Implemented as portable shell scripts
- Extensible via user hooks and plugins at the site level
- User specifies the job characteristics using env variables
- Initially developed by HLRS in int.eu.grd, now maintained by CSIC



► The script:

```
[imain179@i2g-ce01 ~]$ cat test2mpistart.sh
#!/bin/sh
# This is a script to show how mpi-start is called

# Set environment variables needed by mpi-start
export I2G_MPI_APPLICATION=/bin/hostname
export I2G_MPI_APPLICATION_ARGS=
export I2G_MPI_TYPE=openmpi
export I2G_MPI_FLAVOUR=openmpi
export I2G_MPI_PRECOMMAND=time

# Execute mpi-start
$I2G_MPI_START
```

► The StdOut:

```
[imain179@i2g-ce01 ~]$ cat test2mpistart.sh.o114486
Scientific Linux CERN SLC release 4.5 (Beryllium)
Scientific Linux CERN SLC release 4.5 (Beryllium)
lflip30.lip.pt
lflip31.lip.pt
```


► The StdErr:

```
[lflip31] /home/imain179 > cat test2mpistart.sh.e114486
Scientific Linux CERN SLC release 4.5 (Beryllium)
Scientific Linux CERN SLC release 4.5 (Beryllium)
real  0m0.731s
user  0m0.021s
sys   0m0.013s
```


- **MPI commands are transparent to the user**
 - No explicit mpiexec/mpirun instruction

- **Sites tested:**
 - Sites publishing **MPICH** or **OPEN-MPI** tag (or a variant of these)
 - and **MPI-START** tag
 - Used one of the queues available at the site (the shortest one if available)
 - VOs: cms, biomed & fusion
- **Job submitted:**
 - Simple MPI job calculates π
 - Job requesting 2 processes
 - Used MPI-START hooks for compilation before the execution
 - Assume MPI-START is available at the default location
`$I2G_MPI_START`

MPI Support Biomed 2008




MPI Support Biomed 2009




- **Support for MPI is increasing in the sites**

- New flavours supported
- MPI-START available at the sites
- Less errors:




- Other Vos also getting more resources:

MPI support CMS



MPI support Fusion VO



- Good success rates for the jobs sent:

Sites without errors

Open MPI	82%
MPICH	50%

Sites without errors

Open MPI	100%
MPICH	83%

- **Some sites publish MPI support but they do not have MPI-START**
 - Usually this is a legacy configuration of the site
 - 33% of the sites publishing MPI support have this problem
- **Frequent errors in sites:**
 - Wrong installation of the MPI environment:
 - mpiexec for PBS but having LSF scheduler
 - Open MPI mpiexec for mpich or viceversa
 - Site information does not match site configuration
 - Lack of tools (e.g. Compilers)
 - Site errors:
 - Abort the jobs without apparent reason
 - Not allocating the CPUs requested
- **Use adequate requirements:**
 - `Requirements = Member("MPI-START",
other.GlueHostApplicationSoftwareRunTimeEnvironment) && Member ("OPENMPI",
other.GlueHostApplicationSoftwareRunTimeEnvironment);`
 - `Requirements = Member("MPI-START",
other.GlueHostApplicationSoftwareRunTimeEnvironment) && Member ("MPICH",
other.GlueHostApplicationSoftwareRunTimeEnvironment);`

- **Astronomy & Astrophysics**
 - Marcos López (caniego@ifca.unican.es)
 - also on Thursday at 11:00 in the A&A Applications in Grids session
- **Computational Chemistry:**
 - Iñaki Silanes (isilanes@ifca.unican.es)
- **MPI:**
 - Enol Fernández (enolfc@ifca.unican.es)
 - <https://twiki.cern.ch/twiki/bin/view/EGEE/MpiTools>