

Data processing activities at the MAGIC site

David Fidalgo^{*1}, Miguel Nieves-Rosillo¹, Ana Babic², Jose-Luis Contreras¹, Michele Doro³, Nikola Godinovic⁴, Dario Hrupec⁵, Alejandro Lorca⁶, Abelardo Moralejo⁷, Konstancja Satalecka¹ and Martin Will⁸ for the MAGIC[†] collaboration

¹UCM, Madrid, Spain

²FER, UniZg, Zagreb, Croatia

³INFN, Padova, Italy

⁴FESB, Split, Croatia

⁵IRB, Zagreb, Croatia

⁶ESAC, Madrid, Spain

⁷IFAE, Barcelona, Spain

⁸IAC, Tenerife, Spain

MAGIC is a system of two imaging atmospheric Cherenkov telescopes located on the Canary Island of La Palma. The fast processing of the data at the observation site plays an essential part in the operation of the telescopes and has continuously improved since the beginning of the experiment. The on-site computing can be divided into three major contributions: the MAGIC online analysis (MOLA), providing preliminary real time analysis results; the on-site analysis (OSA), providing final data products at the end of each observation night; and the Data Check (DC), a daily check on the performance of the telescope's subsystems and the quality control of the data observed during the previous night. We present the status of the system, including the latest upgrades and details on its performance.

*The 34th International Cosmic Ray Conference,
30 July – 6 August, 2015
The Hague, The Netherlands*

*dfidalgo@gae.ucm.es

†<https://magic.mpp.mpg.de>

1. Introduction

This paper describes the current status of the computing pipelines operating at the MAGIC site. Section 2 provides an overview of the computing infrastructure available at the MAGIC site in the Observatorio del Roque de los Muchachos (ORM) and the organization of the data flow across the different systems. The next three sections are devoted to the three pipelines operating there: MOLA, OSA and DataCheck.

2. Computing infrastructure and data flow

The computing system at the MAGIC site consists of a cluster of computers linked by internal networks and accessible from outside via a public web server. *Data Acquisition* (DAQ) and operations are carried out in the DAQ and subsystem servers, while data processing takes place in the analysis cluster. The later is composed of a set of HPC servers running the same operating system (Scientific Linux Cern 6.3) and sharing both network and disk access.

The storage capacity of the whole cluster is provided by four RAID systems and local disks for temporary storage. Each RAID system has a different capacity, configuration and purpose. Volumes devoted to DAQ (RAID1 and RAID2 for the two MAGIC telescopes M1 and M2, respectively), use the XFS file system and are handled by the DAQ machines. Shared volumes (RAID3 and RAID4) store the compressed raw data, software and user data. They are connected by means of a fibre-channel dedicated network, and use a GFS2-formatted file system accessible to every




Figure 1: Data flow scheme at the MAGIC site. SuperArehucas (1) is the central control for the MAGIC telescopes, controlling the data taking and merging the information of all important subsystems into report files. The copy/compress activities (3, 4) are done simultaneously, while the others follow approximately the sequential order indicated by the number inside the circle. The output of the on-site analysis (5) are low-level analysis products with different key names (*_D_*, *_Y_* etc.). The transfer of the data to the MAGIC data center (7) and the deletion (8) are triggered by the MySQL database that keeps track of the daily activities on-site. The processes discussed in this work are MOLA, OSA and the DataCheck.

machine of the cluster. The decoupling of the two storage arrays is important since it allows the on-site activities (like the on-site analysis and the data check) to act independently from the data acquisition and to maintain legacy systems used by the DAQ machines isolated from the more modern analysis cluster.

Data created at the telescopes (~ 100 GB per telescope per hour of observation) are processed in a well-defined chain, comprising different *activities* according to the logical actions and subsystems involved. The responsibility for each activity relies in a group which takes care of the full development and deployment of the services and tools required for a correct processing. The data flow of the activities carried out on-site is sketched in Fig. 1.

3. The online analysis

The *MAGIC online analysis* (MOLA) is a multi-thread C++ program and part of the *MAGIC Analysis and Reconstruction Software* (MARS [2]). It is used to obtain on-the-fly estimates of the gamma-ray flux from sources in the field of view of the telescope during data taking.

The program runs simultaneously with the data acquisition software on an independent computer and receives the event information of each telescope (M1 and M2) from the two DAQs through TCP/IP by means of ASCII files. This is handled by two independent *reading* threads which then go on to perform the non-stereo analysis steps using MARS algorithms with some simplifications in the signal extraction and image cleaning in order to gain speed. At the same time, an *analyzing* thread is appointed to match the events from the two telescopes and performs the stereoscopic reconstruction as well as the background suppression. The same thread also produces high-level analysis results like gamma-ray excess sky-maps and light curves of the measured gamma-ray flux in real time only a few seconds after the data taking started. However, being automatically generated with some simplifications, these results must be considered preliminary.

Currently the sensitivity of the MOLA analysis above ~ 200 GeV is estimated to be 1.0% of the Crab Nebula flux in 50 hours of observation, compared to 0.66% for the offline standard analysis with MARS. MOLA provides the operators of the telescope with all the necessary information to judge the behavior of the observed astrophysical source, together with early diagnostic information about the data itself¹.

4. The On-Site Analysis

The aim of the *on-site analysis* (OSA) is to provide low and intermediate level analysis products to the collaboration on the day after the observation night. OSA essentially takes care of the calibration of the data, the cleaning and parameterization of the images and merges the data from the two MAGIC telescopes. Those steps reduce the file sizes by a factor of ~ 200 compared to the compressed raw data and are performed in parallel with the raw data transfer to the MAGIC data center located at the *Port d'Informació Científica* (PIC) in Barcelona. OSA consists of the *MAGIC Analysis and Reconstruction Software* (MARS [2]), a set of python scripts, cron jobs and a PBS/Torque resource manager. Its objective is achieved by a high degree of parallelization when processing the data and by starting its operation already during the observation night.

¹For a more detailed description on the mode of operation of MOLA see [1]

4.1 OSA and MARS

MARS is a set of C++ classes based on the well known ROOT package from CERN² and allows the analysis of MAGIC data by means of compiled macros called MARS executables. OSA interacts with those executables via sub-processes, which are spawned within the python scripts, passing them arguments and evaluating their return codes. The main MARS executables used by OSA are the following:

- *sorcerer*: takes as input the compressed raw data and calibrates it.
- *merpp*: adds the central control reports to the calibrated files.
- *star*: performs the cleaning and the parametrization of the images.
- *superstar*: combines the M1 and M2 *star* files, containing image parameters, and performs the stereo reconstruction.
- *melibea*: estimates the event properties, e.g. energy, direction and particle kind.

In general, collaboration members start their analyses with *star* or *superstar* files provided by OSA since running *melibea* requires input generated with the help of Monte Carlo simulations, which on the other hand depend on the instrument setup and analysis goals. Another reason is that OSA does not make any kind of data selection regarding the quality of the data. This is usually done by the analyzer at the *star* or *superstar* level. *Melibea* files are produced to allow for a quick but preliminary off-line analysis of targets of opportunity.

4.2 OSA workflow

Every hour, the copy/compress process (see Fig. 1) looks for new raw data on the immediate (DAQ) storage and copies it together with the corresponding report files to the on-site shared

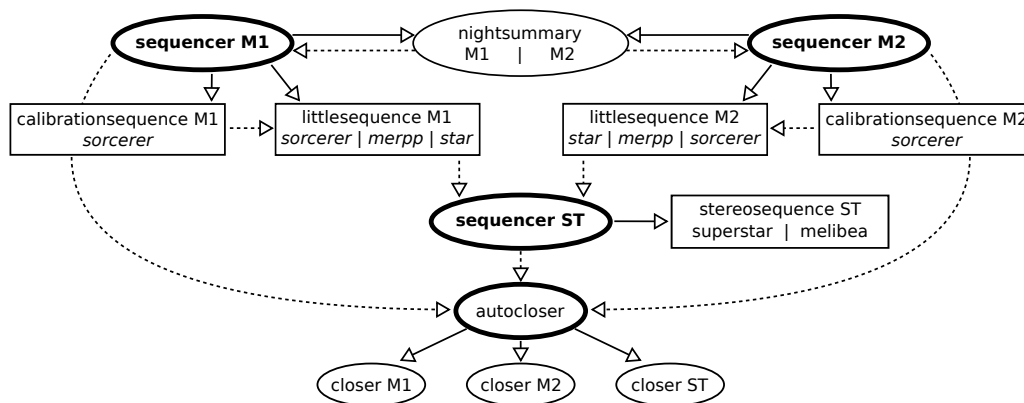


Figure 2: A sketch of the OSA workflow. Ellipses indicate scripts run in the user interface (bold ones are initiated by cron jobs) while scripts inside boxes are executed anywhere in one of the working nodes. Solid lines represent job submissions or process calls while dashed lines illustrate dependencies. The MARS executables called by the respective python scripts are written in italic.

²<https://root.cern.ch>

storage, classifying it by observation nights. With the same frequency and after a short delay, the on-site analysis chain is initiated by cron jobs that execute the so-called *sequencer* (see Fig. 2). This script, with the help of the *nightsummary* script, checks for available raw data in the shared storage system and gathers information from the report files for a given observation night. The *sequencer* then goes on to create a list of sequences to be analyzed, establishes relationships between them, making the execution of some of them dependent on the successful completion of others, and sends them to the Torque queue system. Three types of sequences are built, depending on the content of the input files:

- *Calibration*: to process the calibration runs taken for each source before starting the observation.
- *Data*: to process a data run of one telescope, which by default lasts 20 minutes.
- *Stereo*: to merge and process the output files of the *Data* sequences for each telescope sharing the same data run.

While the command line arguments M1 and M2 build *Calibration* and *Data* sequences for the corresponding telescope, ST results in a list of *Stereo* sequences. The sequences are processed in the working nodes of the computer cluster and handled by different scripts according to their type (*calibrationsequence*, *littlesequence* and *stereosequence*, respectively). The Torque system balances the charge among the nodes.

The standard output of the *sequencer* is a human readable table displaying the characteristics, status and progress of each sequence. This table is automatically interpreted by the *autocloser* script that is executed by a cron job on the half-hour. When it detects that the analysis has finished for all the sequences of M1, M2 or ST, the *autocloser* calls the *closer* script. This *closes the day* for the corresponding data (M1, M2, or ST), triggering the transfer of the respective analysis products to the MAGIC data center through an entry in the MySQL database. The output of the *sequencer* is also copied periodically to a web page for visual inspection by the OSA team.

4.3 OSA performance

The first version of OSA was installed in 2005 and consisted of a small set of shell scripts called by cron jobs. Since then it underwent several updates/changes with a significant remodeling

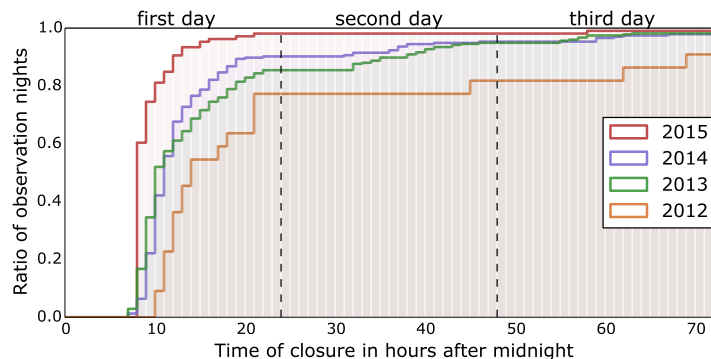


Figure 3: OSA performance for closing M1 for the last 4 years of operation. The y-axis shows the ratio of nights closed before x hours after midnight. Statistics could only be gathered from 2012 on.

during 2012 when the whole system was transferred to python and the storage system was separated into the immediate storage and the on-site shared storage system. This decoupling allows OSA to start the analysis chain already during the night after the first data is available without disturbing the rest of the data taking.

At the end of 2014, a change in the *nightsummary* script and some helper modules permitted the automatic analysis of some non-standard observations. Additionally, the introduction of the *autocloser* script automatized the interpretation of the *sequencer* table and the *closing of the day*, tasks that formerly had to be taken care of manually. Both modifications led to a higher degree of automatization and significantly decreased the workload for the OSA staff.

The performance of OSA over the last 4 years measured by the time needed to process the data of a single day is shown in Fig. 3. As for 2015, OSA completes the analysis of M1 and M2 for $\sim 85\%$ of the observation nights within the first 12 hours after midnight, while for the stereo analysis it is 77%. About 94% of the observation nights are completely analyzed and closed within the first 24 hours. For the rest of the nights, OSA encounters major problems that require a deeper intervention by the OSA staff and are normally fixed within three days.

Together with a stable transfer of the analysis products to the MAGIC data center, where it becomes accessible to the general collaboration, OSA provides the possibility of a fast offline analysis and therefore allows for an adaptive scheduling of targets of opportunity like flaring sources.

5. The Data Check

The Data Check is a series of tasks for monitoring the performance of the telescope subsystems, with the aim of ensuring the safety of the hardware and quality of the data. It comprises of scripts (shell, python, php, etc.) and MARS executables which process the data collected during the night. A scheduled operator is in charge of checking the output and appropriately reacting to the automated warnings. The Data Check ensures timely repair and anticipation of possible failures, leading to the reliable and stable performance of the telescopes.

With respect to the timeliness of the response, there are three basic levels of activities. In parallel to the data-taking, the *iScream* online monitoring tool provides the telescope operators with critical information about 23 parameters from the Central Control (CC) software of the telescopes. Each morning following the data-taking, Daily Check scripts process all the collected subsystem data and generate detailed reports, which form the basis of a summary report by the Daily Check operator. At the end of the daily activities, the long-term monitoring data is generated and stored in a dedicated database, allowing analysis of parameter correlations and long-term trends.

5.1 Online Monitoring

The main purpose of the online monitoring is to alert the telescope operators about conditions which are outside of the hardware safety limits, or have the potential to severely degrade the quality of data.

The *iScream* program is based on the time series data logging and graphing system *RRDtool*. It receives hardware and environment-related data from the CC software, processes them to obtain basic statistics (mean, maximum and minimum values) and stores the consolidated data in different databases: hourly, daily, weekly, monthly, and yearly. With these values, *iScream* creates online

monitoring plots. If values go outside the predefined limits, within one minute warnings and alerts are issued in form of changed background color, pop-up and voice messages [3].

5.2 Daily Check

Detailed reports on subsystems are generated and checked daily. The purpose is to ensure that the majority of the problems are detected and investigated before the start of observations of the following night. A detailed description of these activities is given in [4].

5.2.1 Reports

During the data-taking, the CC and Data Acquisition software write reports every 10 s and about once per minute, respectively. At the end of the night, these reports are first processed into a ROOT file format. A series of MARS macros then process these files and create graphical outputs in pdf and root format. More than one hundred parameters are monitored this way, 24 of which are monitored for each of the 1039 pixels in each of the two cameras.

After the OSA scripts are finished, data check macros create reports on the night's calibration, pedestal and data runs. For each of those runs, a detailed report is created with the charge distributions and timing information for the calibration and interleaved calibration. From the data runs, muon information is extracted, and is used as basis for the estimation of the effective point spread function of the event reconstruction.

5.2.2 Automated and Operator Checks

About 40 of the monitored parameters are checked by automated scripts. Checks are performed against various criteria: a) Fixed predefined limits are applied to the majority of the parameters; b) Some parameters, e. g. the temperature of the receiver electronics or the individual pixel temperatures, are instead required to stay stable during the entire data-taking; c) Dynamically determined limits take into account the variation of atmospheric conditions; d) In a camera with 1039 pixels, even in the case of average and RMS values being within fixed predefined limits, clusters of pixels can significantly deviate from the mean. Several camera plots are automatically checked for such clusters, which would point to failures in the electronics chain.

The duty of the operator in charge of the daily check is to examine all the relevant plots published on a dedicated web page, and particularly those plots where automated warnings have been generated. In the latter cases, the operator consults all related plots which can help diagnose the problem, and alerts the person(s) responsible for the particular telescope subsystem. The task is usually completed by 1 p.m. CET the day following the observations. In most cases, this is timely for the responsible person(s) to investigate the problem and organize intervention with the on-site telescope staff.

5.2.3 Analyzer Summary Reports

In addition to the extensive plots used for monitoring the hardware and data taking conditions, a light-weight graphical file, the “superplot” is produced. This file is particularly useful as basis for the data quality selection for analyzers, as it will provide for each observed source: telescope pointing, light level as measured by direct current of the pixels, discriminator thresholds, telescope

coincidence rates re-scaled to eliminate the zenith angle-dependence, atmospheric transmission measured by LIDAR and infra-red pyrometer.

5.3 Long-Term Monitoring

Following the Daily Check activities, 78 of the analyzed parameters are further processed by calculating basic nightly statistics (mean, median, minimum, maximum, RMS) which are stored in a separate *MySQL* database. The database covers the period since the last major upgrade in 2012.

The web interface for the database is based on the *Twitter Bootstrap 2* framework. It allows the selection of two parameters, and for each of the two parameters one of the five statistical measures. The user can also select a date range for the plots. Once submitted, the query is returned in a few seconds, with five plots pre-generated for the query: histogram for each of the two parameters, time series for each of the parameters, and a correlation plot for the two parameters. All the plots are available for download in either pdf or root format, and the time series data on which the plots are based can also be viewed or downloaded.

In addition to this basic functionality, the user can interactively explore the time series for the two parameters. These plots are based on the *Dygraphs* charting library and allow click-and-drag zooming in the graph on either axis, displaying the data for any data point by hovering, and displaying a running average plot for the number of days chosen by the user.

The purpose of this tool is to enable investigation of long-term trends and correlation of parameters. It is being used to predict potential hardware failures, determine parameter limits, and gather deeper understanding of the detector.

6. Conclusions

Despite the high data volume generated by the telescopes, an important part of the MAGIC processing is performed on-site. MOLA and Date Check allow a prompt feedback both for physics analyses and quality checks, while OSA provides high quality calibrated and pre-processed data which in fact is used in most of the articles published by the MAGIC collaboration. The scheme implemented should be easily scalable and applicable to next generation observatories like the Cherenkov Telescope Array, CTA.

References

- [1] D. Tescaro et al., Proc. of the 33rd ICRC, Rio de Janeiro, Id. 122, 2013.
- [2] R. Zanin et al., Proc. of 33rd ICRC, Rio de Janeiro, Id. 773, 2013.
- [3] T. Terzic, PhD Thesis, 2015.
- [4] N. Godinovic et al., Proc. of the 33rd ICRC, Rio de Janeiro, Id. 687, 2013.

Acknowledgements

We would like to thank the Instituto de Astrofísica de Canarias for the excellent working conditions at the Observatorio del Roque de los Muchachos in La Palma. The support of the German BMBF and MPG, the Italian INFN, the Swiss National Fund SNF, and the ERDF funds under the Spanish MICINN is gratefully acknowledged. This work was also supported by the CPAN CSD2007-00042 and MultiDark CSD2009-00064 projects of the Spanish Consolider-Ingenio 2010 programme, by grant 268740 of the Academy of Finland, by the Croatian Science Foundation (HrZZ) Project 09/176, by the DFG Collaborative Research Centers SFB823/C4 and SFB876/C3, by the Polish MNiSzw grant 745/N-HESS-MAGIC/2010/0 and by the Ramón y Cajal program from the Spanish MICINN.