Infrastructure	Description	Contact people	Data rate and volume	Data formats and documentation
SKA: Square Kilometer Array https://www.skatelescope.org/	The largest **radio** telescope array (South Africa and Australia) made up by dishes (high frequency) and antennas (low frequency). **ESFRI project**.	Signal and data transport: Tyler Bourke <t.bourke@skatelescope.org>, Rodrigo Olguin <r.olguin@skatelescope.org>, Jeff Wagg <j.wagg@skatelescope.org> Science data processor: **Miles Deegan** <m.deegan@skatelescope.org>, Juande Santander-Vela <j.santandervela@skatelescope.org>	On-site data: 2 TB/s (SKA1-mid) and 1 TB/s (SKA1-low) Archive data: 100 PB/yr	To be determined.
CTA: Cherenkov Telescope Array https://www.cta-observatory.org/	Two arrays (La Palma and Chile) of Cherenkov telescopes for very-high energy **gamma ray** astronomy. **ESFRI project**.	General Data Management: **Giovanni Lamanna** <lamanna@lapp.in2p3.fr> Data model : **José Luis Contreras** <contrera@gae.ucm.es>	Total data rate: 10 GB/s Reduced raw data volume: 4 PB/year	The low level data format is to be determined. High level data will be provided to the scientific community in FITS format.
KM3NeT: Km^3 Neutrino Telescope http://www.km3net.org/	Future European deep-sea infrastructure in the Mediterrenean Sea hosting a **neutrino** telescope. **ESFRI project**.	Software and computing system: **Kay Graf** <Kay.Graf@physik.uni-erlangen.de> On-line DAQ and readout system: Tomaso Chiarusi <tommaso.chiarusi@bo.infn.it>	Total raw data rate: 25 GB/s. Reduced raw data rate: 1.2 MB/s.	Data stored in “flat” files, each containing a run and only metadata stored in the experiment database. The software Jpp is used for data analysis.
E-ELT: European Extremely Large Telescope https://www.eso.org/sci/facilities/eelt/	Large ground-based **optical/near-infrared** telescope in Chile. **ESFRI project**.	E-ELT Science contact ESO Quality Control and Data Processing Group: **Reinhard Hanuschik** <rhanusch@eso.org>		Image-based format.
LIGO/EGO: Laser Interferometer Gravitational Wave Observatory / European Gravitational Observatory http://www.ligo.org/ https://www.ego-gw.it/	Two LIGO multi-kilometer-scale **gravitational wave** detectors (Washington and Lousiana) using laser interferometry, and the giant Michelson interferometer Virgo (Pisa).	Virgo data analysis software: **Marie Anne Bizouard** <mabizoua@lal.in2p3.fr>	Advanced LIGO raw data rate: 20 MB/s. Advanced VIRGO raw data rate: 400 MB/s. Advanced VIRGO reduced data rate: 24 MB/s.	Common Data Frame Format for Interferometric Gravitational Wave Detectors Open data may also be provided as HDF5 files and in VOEvent format.
EVN: European VLBI Network http://www.evlbi.org/	Large scale **radio** astronomical facility using Very Long Baseline Interferometry (VLBI). The **e-EVN** programme is an introduction of the e-VLBI technique to create a VLBI network based on broadband data transmission and real-time data processing and it is **pathfinder for SKA**.	Web page of the Joint Institute for VLBI in Europe (JIVE)	e-VLBI data rate: 128 MB/s each telescope (x 16)	Image-based format (FITS files). In e-VLBI, data is streamed directly from the telescope to the correlator and correlated in real time.
LOFAR: The low frequency array http://www.lofar.org/	**Radio** interfermotric array of dipole antenna stations (Netherlands, Germany, France, UK, Sweden…) for astronomical observations at frequencies below 250 MHz. **Pathfinder for SKA.**	Data products and management: Roberto Pizzo <pizzo@astron.nl> Science Support: **Luciano Cerrigone** <cerrigone@astron.nl>	Observational data: 10 GB/s Archive data: 6 PB/yr	Beam formed observations recorded in HDF5 and processed data stored in binary files with standard formats in the pulsar community. Interferometric observations produce measurement sets in the standard format of the CASA package.

EUCLID http://www.euclid-ec.org/ http://sci.esa.int/euclid/	Satellite mission to map the geometry of the **dark Universe** (redshifts ~2)	Euclid Consortium Lead: **Yannick Mellier** <mellier@iap.fr> Euclid data processing: Fabio Pasian <fabio.pasian@inaf.it>	Data rate (4 h periods): 106 GB/day (7 MB/s)	Image-based format. Public data will be distributed in a common format (e.g., FITS and Virtual Observatory).
LSST: Large Synoptic Survey Telescope http://www.lsst.org/	**Optical** telescope placed in Chile that can identify changing or moving objects.	Data products: **Mario Juric** <mjuric@lsst.org>	Data rate: 15 TB/night (400 MB/s)	Image-like structure (e.g., FITS).
H.E.S.S.: High Energy Stereoscopic System https://www.mpi-hd.mpg.de/hfm/HESS/	Array of atmospheric Cherenkov telescopes located in Namibia for the study of cosmic very-high energy **gamma rays**. **Pathfinder for CTA**.	DAQ, software, telecommunications: **Christian Stegmann** <christian.stegmann@desy.de>	Data rate: 46 MB/s	Analysis framework SASH (Storage & Analysis Software at H.E.S.S.) relying on ROOT.
MAGIC: Major Atmospheric Gamma-Ray Imaging Cherenkov Telescopes https://magic.mpp.mpg.de/	Two imaging Cherenkov telescopes in La Palma for the study of cosmic very-high energy **gamma rays**. **Pathfinder for CTA**.	Software coordinator: **Julian Sitarek** <jsitarek@uni.lodz.pl>	Data rate: 100 MB/s	Proprietary format for low-level data and ROOT for higher levels. Public results are in FITS format.
IceCube https://icecube.wisc.edu/	Particle detector at the South Pole to search for astrophysical **neutrinos**. It encompasses a cubic kilometer of ice. **Pathfinder for KM3NET**.	Computing and data management: **Gonzalo Merino** <gonzalo.merino@icecube.wisc.edu>, Data management: James Bellinger <james.bellinger@icecube.wisc.edu>	Data rate (detector): 1 TB/day (12 MB/s) Data rate (satellite transmission): 100 GB/day (1.2 MB/s) Public datasets: 10 TB/year	The analysis framework is proprietary software called IceTray. Public datasets: each event is a python class, and the binary file is then a stream of serialized event objects.
ANTARES: Astronomy with a Neutrino Telescope and Abyss environmental RESearch http://antares.in2p3.fr/	Large area water Cherenkov detector in the deep Mediterranean Sea (coast of Toulon, France) for the detection of muons from high-energy astrophysical **neutrinos**. **Pathfinder for KM3NET**.	Data acquisition system: **Mieke Bouwhuis** <mieke.bouwhuis@nikhef.nl>	Raw data rate: 0.3 – 10 GB/s Filtered data rate: 30 – 100 MB/s	Filtered data are written in ROOT format.

Abstract

This working document is part of the task 3.2 D-GEX (Data Generation and information eXtraction) in OBELICS. Its goal is to give an introductory overview of ESFRI projects and related pathfinders with regard to the data formats that they (will) use as well as to the data flow and volume in these experiments. Together with this report, the above table provides at-a-glance summary including contact people in charge of data management in each facility. Subsequently, as one of the deliverables of the task, we will write a Data Format Survey Report, where the different data formats will be classified and compared with each other with the final aim to propose more general standards that may benefit to these and future projects.

1. SKA

1.1. Introduction

The Square Kilometre Array (SKA) is an ambitious project to build a radio telescope that will enable breakthrough science and discoveries not possible with current facilities. Built over two sites in Australia and Africa it will, when both phases are complete (SKA1 and SKA2), provide over a million square metres of collecting area through many thousands of connected radio telescopes. Constructed in two phases: SKA1 is being designed now; SKA2 is planned to follow.

The SKA radio telescopes will provide continuous frequency coverage from 50 MHz (6 m wavelength) to 20 GHz (1.5 cm wavelength). SKA1 will cover most of this frequency range, while greater sensitivity at all frequencies and fast surveying is being added in SKA2. A project of this scale has perforce been developed from the very beginning as an international partnership that will draw on the scientific, technological, industrial and financial resources of its members. What will emerge is a new international observatory.

1.2. Description of SKA1

Figure 1.1 shows the major SKA Observatory entities: SKA1-low in Australia, SKA1-mid in South Africa and the SKA Global Headquarters in the UK. The thick flow-lines show the uni-directional transport of large amounts of digitised data from the receptors to the central signal processing facilities on the sites, and from the central signal processing facilities to the Science Data Processing Centres and Archives. The thin dash-dot lines show the bi-directional transport of system monitor and control data.

The SKA1-mid telescope will be a mixed array of 133 15-m SKA1 dishes and 64 13.5-m diameter dishes from the MeerKAT telescope. The antennas will be arranged in a moderately compact core with a diameter of ~1 km, a further 2-dimensional array of randomly placed dishes out to ~3 km radius, thinning at the edges. Three spiral arms will extend to a radius of ~80 km from the centre. The SKA1-dish array will be built essentially in the same location as the MeerKAT array, and the array can be expanded to a much larger SKA2 array from that location.

The SKA1-low telescope receptors will consist of an array of ~131,000 log-periodic dual-polarised antenna elements. The antenna elements will be grouped into ~512 stations, whose antennas will be beam-formed. Many of the antennas will be arranged in a very compact configuration (the 'core') with a diameter of ~1 km. The rest of the elements will be arranged in stations a few 10s of metres in diameter. The stations will be distributed over a 40-km radius region lying within Boolardy Station, most likely organised into spiral arms with a high degree of randomisation.

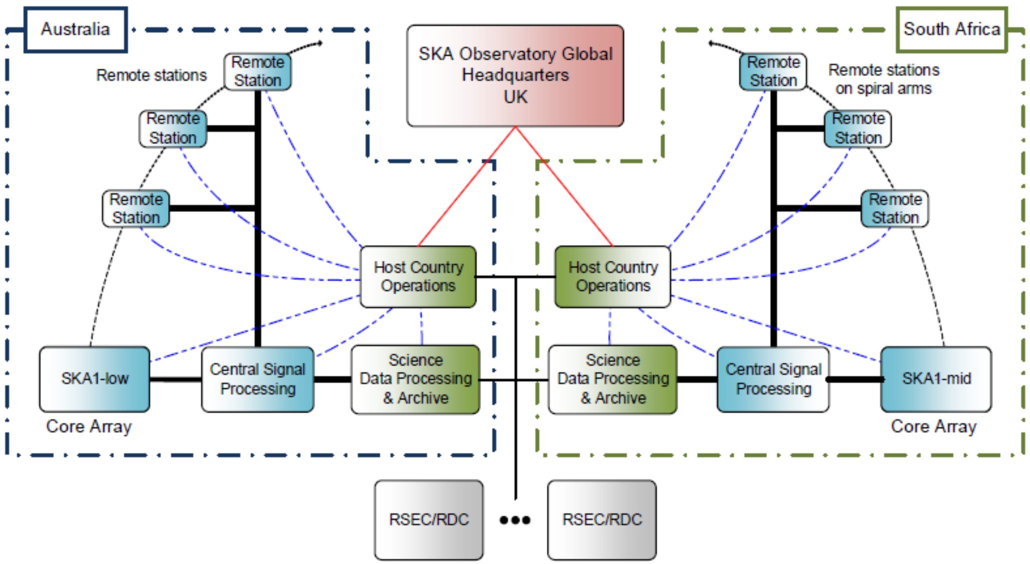

Figure 1.1. A schematic diagram of the SCA Observatory.

Signals from the dishes of the SKA1-mid or from the beamformers of the stations of the SKA1-low will be transported to the corresponding Central Signal Processing Facility, where they will be channelized and cross-correlated with each other. Output data from the correlator will be transported to the

Science Data Processing Centre of the site. The signals will also be combined into a large number of array beams, the outputs of which will then be distributed to specialised pulsar search equipment. Pulsar candidates will be sent to the Science Data Processing Centre for further analysis.
The Science Data Processor is envisaged to be a supercomputing facility with an attached or nearby archive to store science-ready data. The science data processor is where calibration of the data takes place, images of sky brightness are formed, and further analysis of time-domain effects are carried out.
The Archives will store the outputs from the Science Data Processors in the site countries, where they will be kept for an indefinite time. The RSEC/RDCs are the facilities where it is expected that actual science analysis will take place and/or science data will be stored, as well as engineering design work for upgrades and future developments.

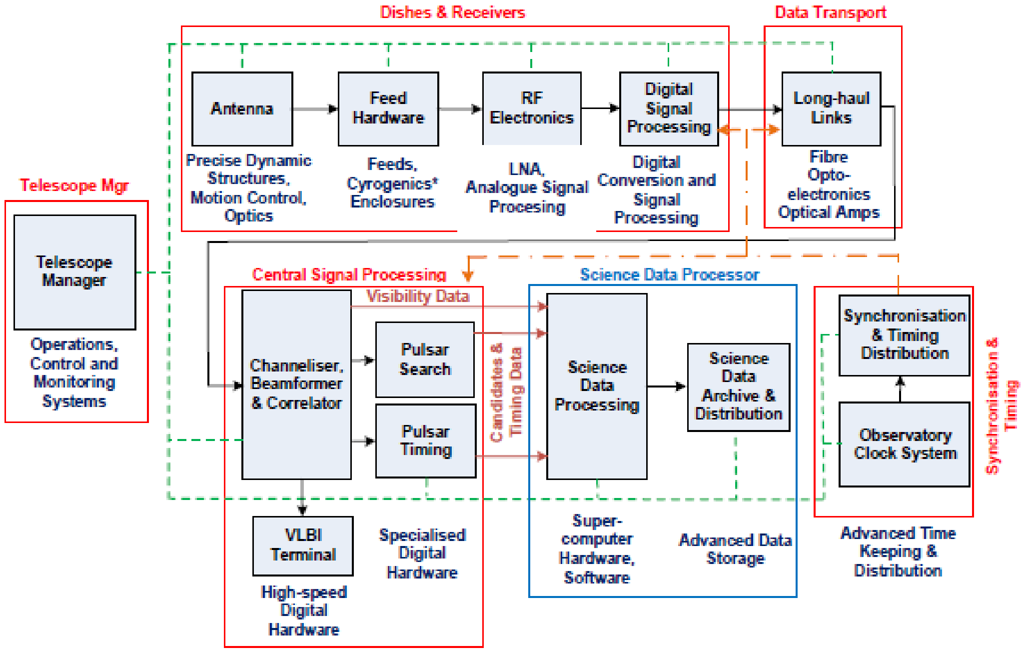

Figure 1.2. Major components of SKA1-mid through the signal flow.

1.3. Data Transport

For both SKA1-mid and SKA1-low, 'Data transport' refers to transport of output data from dishes or stations to the correlator/beamformer subsystems. Data volume is very steady and needs to be transported in only one direction.
The on-site data volume for SKA1-mid is very large, ~12 Tb/s from the antennas in the core and ~4 Tb/s from antennas in the arms. All dishes will receive 100 Gb/s Ethernet capability through to the correlator interface. The data rate required from the correlator-beamformer to the science data processor is ~12 Tb/s.
The on-site data rates for SKA1-low are: ~7 Tb/s from the antennas in the core and ~0.7 Tb/s from antennas in the arms. The SKA1-low outer stations, which will be within 80 km of the processor building, will require a 20 GB/s connection; assuming 45 outer stations the total is 0.9 TB/s. These will be provided by two 10 Gb/s Ethernet links. The data rate required from the correlator-beamformer to the science data processor is ~4 Tb/s.

Bibliography

[1] https://www.skatelescope.org/
[2] SKA1 System Baseline v2. Available at https://www.skatelescope.org/wp-content/uploads/2014/03/SKA-TEL-SKO-0000308_SKA1_System_Baseline_v2_DescriptionRev01-part-1-signed.pdf

2. CTA

2.1. Introduction

The Cherenkov Telescope Array (CTA) is an observatory for very-high-energy gamma-ray astronomy, which will provide observers with data on astrophysical objects over a very wide range of energies. The CTA Observatory (CTAO) will operate two arrays of Cherenkov telescopes, one in each hemisphere. The Northern array will be hosted on the European Southern Observatory (ESO) Paranal grounds, in Chile, and the Southern one at the Instituto de Astrofísica de Canarias (IAC), Roque de los Muchachos Observatory in La Palma, Spain. CTA is intended to provide a service to a wide scientific community.
The design goal is a factor of ten improvement in sensitivity in the currently accessible energy domain of about 100 GeV to some 10 TeV and for an extension of the accessible energy range to well below 100 GeV and up to more than 300 TeV. This ambitious aim can only be achieved with a combination of telescopes of different sizes: large ones for the lowest energies, medium ones for the core energy range and many small ones for the highest energies. To achieve a substantially improved sensitivity at the highest energies, CTA requires a collection area of the order of 10 km^2 which means spreading numerous telescopes over a large area.

2.2. Data Model

Data will be processed on the observatory sites by both a real time and a delayed analysis chains to generate science alerts and monitor the instrument. Afterwards it will be transmitted to the off site data centers for the final analysis and storage. These centers will provide data access services to the scientists and technical personnel of the observatory.

Figure 2.1 depicts the main path and rate of data within the CTAO. On each CTA site, the data rates are based on the event rates from Cherenkov and night-sky-background triggers registered by the telescope array (calculated from Monte-Carlo simulations and site measurements), given a particular trigger scheme. It is assumed that the full-waveform signal is kept only for selected pixels in a camera while for the other pixels only basic information is stored. The nominal data rates result in 5.38 GB/s for CTA south and 3.18 GB/s from CTA north. They also include 20% calibration data and 10 MB/s of device monitoring and control data for each site, for a resulting total data rate of about 8.56 GB/s. A data volume reduction system will cut the nominal data rate by a factor of 10. Raw data and device control data are transferred from CTA sites to the CTAO Data Management Centre (CTAO-DMC) along the network. Data are then distributed over four data-centres (DC) participating in the processing and archive of data.

The estimated annual (reduced) raw data volume is 4 PB plus 4 PB of data products. The total volume to be managed by the CTAO Archive is of the order of 27 PB/year, when all data-set versions and backup replicas are considered. This will correspond to a permanent archive of the order of more than 400 PB in 2031.

Off-site, in each data-centre DCi, the data production is comprised of a series of processing steps that transform archived (reduced) raw data DL0 (Data Level 0) to calibrated camera data (DL1), then to reconstructed shower parameters such as energy, direction, and particle ID (DL2), and finally to high-level observatory products comprised of selected gamma-like events, instrument response tables, and technical data (DL3). DL3 data will have a total volume of about 2% of the DL0 data volume and guaranteed access will be provided in the CTA archive to basic users. The science tools are then used either automatically or by users to produce DL4 (e.g. spectra, sky-maps). Finally (DL5) legacy observatory data, such as CTA survey sky maps or the CTA source catalogue will be produced.

For the low level data (DL0), it is assumed that the images from each telescope will be kept in separate files together with some calibration and technical information needed for their first processing. Three options are being considered and prototyped right now for the file format. One of them is based on Google protobuffs protocol and compressed FITS, another one in the PACKETLIB format used in space missions and a third one is an extension of the format presently used for the Monte Carlo data by the H.E.S.S and CTA collaborations.

High level data (DL3-DL5) comprises the data levels that the CTA observatory will provide to guest observers and the scientific community in general. They must be provided in open, self documented formats. The observatory requirement is to use the FITS format and to ensure the integration of CTA high-level data within the Virtual Observatory infrastructure.

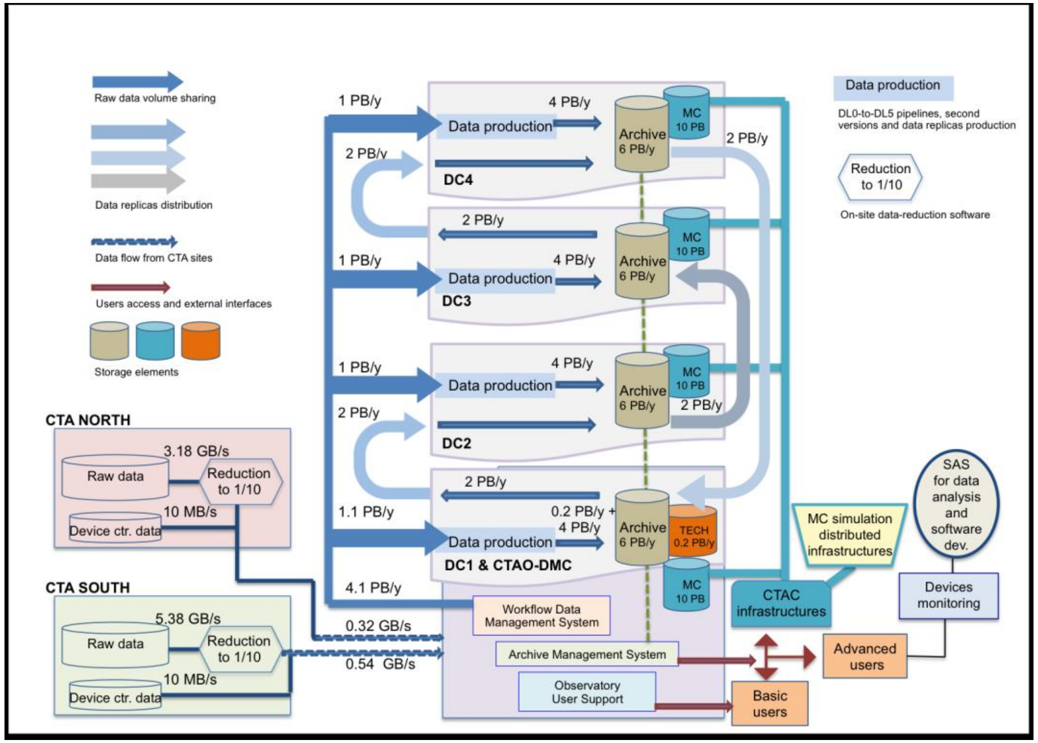

Figure 2.2. Data volume management of the CTA Observaroty.

Bibliography

[1] https://www.cta-observatory.org/
[2] B.S. Acharya et al., Astropart. Phys. 43 (2013) 3.
[3] J.L. Contreras et al., PoS(ICRC2015)960.
[4] Data Management Technical Design Report.

3. KM3NeT

3.1. Introduction

The KM3NeT Observatory is a large scale neutrino telescope to be built in the deep waters of the Mediterranean Sea. With four to five cubic kilometres instrumented with thousands of optical sensors, KM3NeT will be the largest and most sensitive neutrino detector in the energy range of 1–10000 TeV. It will be capable of neutrino astronomy with unprecedented accuracy. Being situated in the Northern Hemisphere it will be particularly suited to the investigation of high energy neutrinos from our Galaxy. In addition, it will provide innovative science opportunities in Earth and Sea Science.

3.2. Infrastructure Description

The deep sea infrastructure consists of a neutrino telescope and a network of nodes for marine and earth science investigations. The neutrino telescope occupies an area of several square kilometres of the seabed and the marine and earth science nodes are located far enough to avoid interference with the neutrino telescope but close enough to make use of a common deep sea cable network. Next, the neutrino telescope is described.

The neutrino telescope will consist of a three-dimensional array of optical modules supported by vertical structures anchored on the seafloor and connected to a seabed cable network for power distribution and data transmission. These structures together with their optical modules are referred to as detection units.

A prototype of a Digital Optical Module (DOM) consists of a glass sphere housing 31 photomultipliers and it has already been integrated in the ANTARES detector for testing and validation. For these tests, the data are transported the 45 km to shore through the ANTARES 1 Gb/s multiplexed optical link that uses reflective modulation in the DOM.

In KM3NeT, the expected data rate of the telescope will be roughly 25 GB/s, assuming 64 bits per recorded photon. This data rate to shore can be accommodated on a number of optical fibres using dense wavelength division multiplexing (DWDM) techniques. A backup solution to the data transmission scheme within a detection unit is a fibre-optic daisy-chain concept. A shore station will house the power supplies, the lasers that will drive the fibre optic network, and will also host the data acquisition system that will implement data filtering and storage, before transferring them to remote nodes for analysis. The overall readout system includes the submarine infrastructure, a shore station and various computing centres around Europe, together with external systems such as the gamma ray coordination network (GCN).

3.3. Data Flow

The total data rate exceeds that of any data storage capacity by several orders of magnitude. Hence, the raw data have to be filtered. The rare neutrino (muon) signal can be discriminated from the random background utilising the time-position correlations produced by the traversing particle. The main challenge is the real-time filtering of the neutrino signal from the continuous random background. The total data rate from the submarine infrastructure should be reduced by a factor 10,000 to less than 10 Mb/s, to be able to transfer the data to the various computing centres in real-time and to store on a permanent medium. Therefore, a high-bandwidth link of 1 Gb/s to the various European computing centres is required. This internet link allows for remote operation of the infrastructure.

The data is gathered as a snapshot or time slice. Each of these is processed in an individual processor within a large computer cluster. A snapshot covers a time span of 10-100 ms. Thereby, a single processing unit acts on a snapshot of the whole telescope and is able to build events based on triggering and selection algorithms for the whole telescope.

The output from each processing unit is the raw data of a time slice that passed event selection criteria together with a set of event description data, called meta-data. The event meta-data for time slices which did not pass any selection may also be kept. For a fraction of these time slices the whole raw data may be archived to study biases and inefficiencies in the processing. The data base is also stored at the permanent data storage facilities and will be updated for each configuration change and after production of new calibration data.

The use of Grid concepts is envisaged, as used by experiments at CERN, with a form of decentralisation for sharing computing resources in different tiers with Europe-wide (worldwide) distribution. All data from the shore station have to be copied to other remote facilities. A guiding principle is to ensure two copies of all data to be secure against accidental loss. A remote facility will provide permanent storage with access to primary processed versions of the data and the capability for reprocessing.

Three types of data will have to be stored on persistent media:

- The PMT acquisition data constitutes the main bulk of all the data produced by the experiment. It is organized in “runs”, each run being associated with the run conditions in which the data have been acquired. Past experience has shown that

such large amounts of data are handled with better performance over “flat” files, possibly using specialized high performance file I/O software frameworks such as ROOT I/O. For this kind of data, only the corresponding metadata should be stored in the experiment database, using potentially complex relations between the different elements of data description.
- Control and configuration data is typically stored on a relational database because of the complex inter-data relationships. This kind of data describes the hardware setup and the values of all the parameters with which this setup has been configured before producing its acquisition data. These data must be time tagged and associated with the corresponding acquisition data in a non modifiable way. At the same time, it has to be structured in a way that makes it intelligible to those who operate the detector and those who analyse data. Configurations are defined as sets of tree structures, automatically mapped over Oracle or MySQL database tables and automatically support configuration history.

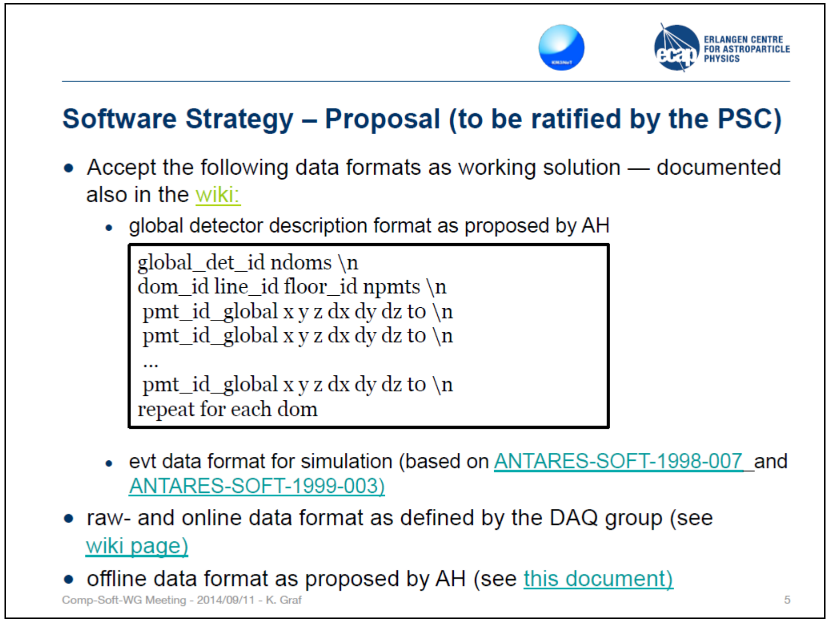

Figure 3.1. Slide taken from the presentation of Kay Graf in the KM3NeT Computing and Software WG Meeting in 2014.

- Instrumentation acquisition data, as acoustic positioning data. These data can represent an important amount of data and are therefore stored in flat files for raw data and database storage for metadata. The data from low data volume instruments are stored directly in the database.

3.4. Data Formats

KM3NeT will develop a data processing chain with well-defined interfaces (mainly i/o and configuration data formats) at each processing step. One official software package per step will be maintained (with alternatives for crosschecks), which can read/write the official data formats directly. Simulation and processing software from ANTARES and IceCube (and corresponding data formats) is being adapted for KM3NeT. The distinct types of data for which formats are being developed are identified in figure 3.1.

Bibliography

[1] http://www.km3net.org/home.php
[2] KM3NeT Technical Design Report. Available at http://km3net.org/TDR/TDRKM3NeT
[3] S. Adrián-Martínez et al., Eur. Phys. J. C 74 (2014) 3056
[4] K. Graf, KM3NeT, Computing and Software WG Meeting (2014)

4. E-ELT

4.1. Introduction

The European Extremely Large Telescope (E-ELT) is a revolutionary scientific project for a 39.3 m segmented optical telescope that will gather 13 times more light than the largest optical telescopes existing today. The E-ELT will have an adaptive optics, providing images 16 times sharper than those from the Hubble Space Telescope. The E-ELT will vastly advance astrophysical knowledge by enabling detailed studies of planets around other stars, the first galaxies in the Universe, super-massive black holes, and the nature of the Universe's dark sector.
The E-ELT is currently under construction on Cerro Armazones as part of the European Southern Observatory (ESO). The telescope will be operated from, and as part of, the existing Paranal site, where is located the Very Large Telescope (VLT).

4.2. Data Management

Fundamentally, the E-ELT will be operated by the La Silla Paranal Observatory (LPO) and Data Management and Operations (DMO) Divisions within the Directorate of Operations. In particular, the data management shall be common to other facilities of the LPO.

ESO is committed to the long-term preservation and accessibility of all the scientific data obtained with the E-ELT, their associated calibrations and ancillary information, as well as processed data products. The raw scientific data, including calibrations, shall arrive at the ESO Science Archive Facility not later than 1 h after they have been obtained, and will be made available to the worldwide community at the end of a proprietary period to be defined by the science policy document. The Science Archive Facility will also store pipeline-processed data products and higher level, science-ready data products. These include calibrated images, spectra, and catalogues. Science-ready data products will contain metadata compliant with Virtual Observatory standards and be made available to Virtual Observatory tools through publication in the appropriate registries.
All data generated by the instruments shall be in a standard format, which shall be deemed to be the most useful for the astronomical community at the time of first light.

Bibliography

[1] https://www.eso.org/sci/facilities/eelt/
[2] The E-ELT construction proposal. Available at https://www.eso.org/public/archives/books/pdf/book_0046.pdf

5. LIGO/EGO

5.1. Introduction

By the early 2000s, a set of initial interferometric detectors of gravitational waves (GWs) was completed, including TAMA 300 in Japan, GEO 600 in Germany, the Laser Interferometer Gravitational-Wave Observatory (LIGO) in the United States, and Virgo in Italy. Combinations of these detectors made joint observations from 2002 through 2011 while evolving into a global network. In 2015, Advanced LIGO became the first of a significantly more sensitive network of advanced detectors to begin observations. Advanced Virgo, KAGRA (in Japan), and a possible third LIGO detector in India will extend the network
On September 14, 2015 at 09:50:45 UTC the two Advanced LIGO detectors simultaneously observed a transient gravitational-wave signal. This is the first direct detection of gravitational waves and the first observation of a binary black hole merger.

5.2. Advanced LIGO

From its inception, the LIGO project was planned to involve the development and operation of multiple generations of increasingly sensitive gravitational wave detectors. The three Initial LIGO detectors were three specialized Michelson interferometers: the 4 km-long H1 and 2 km-long H2 detectors housed by an observatory on the Hanford site in Washington State, and the 4 km-long L1 detector housed by an observatory in Livingston Parish, Louisiana. The detectors included a resonant Fabry–Perot optical cavity in each arm and power recycling, resulting in an increase of the power by a factor of ~8000 with respect to a simple Michelson, and were designed to be sensitive to GWs in the frequency band 40–7000 Hz and capable of detecting a GW strain amplitude as small as 10^{-21}. They operated at design sensitivity in a continuous data-taking mode from November 2005 to September 2007. Subsequently, upgrades were made to the detectors leading to Enhanced LIGO, which collected data from July 2009 to October 2010.
The Advanced LIGO detectors are the second generation of interferometers, which are configured like Initial LIGO with the addition of a signal recycling mirror. Each observatory hosts one Advanced LIGO interferometer with 4 km long arms. A third interferometer is consigned for a planned later installation at a new site in India. All three interferometers are intended to be identical in design and expected performance. Compared to Initial LIGO, Advanced LIGO is designed to provide a factor of 10 increase in strain sensitivity over a broad frequency band, and to extend the low end of the band to 10 Hz. As the probed volume of the universe scales as the cube of the strain sensitivity, this represents an enormous increase (of order 10^3x) in the number of potential astrophysical sources detectable by these instruments. At design sensitivity, Advanced LIGO is likely to detect dozens of compact binary coalescence sources per year.
Each Advanced LIGO interferometer record duplicate copies of the data, which will be written to independent file systems in different buildings at each Observatory. The data rates are about 10 MB/s per interferometer for an aggregate rate of about 1 PB/year. The raw LIGO data is combined with auxiliary measurements and models to build a time series representing the GW strain signal. In addition, there are a large number of auxiliary instrumental and environmental monitoring channels that are ingested (only about 1% of data is in the GW strain channel).
A minimum of three full archival copies of the LIGO data will be stored in at least three different geographic locations – Caltech, Hanford, Livingston – to provide access to the LIGO Scientific Collaboration and other users with

sufficient redundancy. In addition it is important to provide catalogs of what has been archived and where it is located. This applies to both the raw data and all derived data products.
Currently, LIGO data is available only to the LIGO Scientific Collaboration. The plan for the future calls for a phased approach to the release of LIGO data moving from the detection or Discovery era of today (phase 1) to the epoch of routine detections (phase 2). During this phase 2 the entire body of GW data will be openly released to the broader research community.

5.3. Advanced Virgo

The Virgo detector, located in Cascina, near Pisa is Michelson interferometer having two 3 km-long perpendicular arms with Fabry-Perot resonant cavities and a power recycling cavity. Construction was achieved by CNRS and INFN, which created the European Gravitational Observatory (EGO) Consortium for the running and maintaining of the instrument and the site infrastructures. Then, the Collaboration has been joined by Dutch, Hungarian and Polish groups. Virgo performed its first scientific data taking run in 2007.
Similarly to the LIGO interferometers, this detector is being upgraded in two steps: Virgo+ and Advanced Virgo. The goal of the Advanced Virgo project is to start taking data in an early operation phase in 2016. Then, for a late operation (tentative date is 2018), a signal recycling mirror will be introduced and the laser power will be increased up to 200 W. The foreseen changes should result in an increase of the sensitivity of one order of magnitude in the whole detection band. The accessible universe volume should then be increased by a factor of 1000, compared to the Virgo design.
Advanced Virgo data include interferometer sensing and control signals, monitoring signals from all the detector subsystems and environmental-monitoring signals. They are acquired through the fast-data network (~400 MB/s) and merged in different streams at the front-end of the data collection pipeline. Then, data are reduced (decimation, compression, image processing, etc) and provided to several online processes which enrich the streams with additional computed channels. At the end of the data-collection tree, different data streams are stored on disk, with the total data flow being 24 MB/s.
The data are kept on circular buffers at the Virgo site, with a depth of ~6 months for the raw data. During science runs all the data are transferred to computing centers for permanent storage and offline data analysis.

5.4. Data Format

Fortunately the data format has been in place since 1997 as the consequence of a LIGO-Virgo agreement. It is the “Common Data Frame Format for Interferometric Gravitational Wave Detectors”, also called an IGWD Frame file. This expresses multi-channel time and frequency series, which can have different sampling rates for different channels. Each Frame file has thousands of channels, one of which represents the GW strain, with all the other channels used for environment and instrument monitoring.
Both Virgo and LIGO plan to make data publicly available. The open data infrastructure will be able to deliver data files that are easier to use for most open data users, for example HDF5 in addition to the existing Frame format. In addition, it is proposed to begin a program that uses the Virtual Observatory VOEvent Transport Protocol to communicate real-time alerts to follow-up observers.

Bibliography

[1] http://www.ligo.org/index.php
[2] https://www.advancedligo.mit.edu/
[3] https://www.ego-gw.it/
[4] http://www.virgo-gw.eu/
[5] B.P. Abbott et al., Rep. Prog. Phys. 72 (2009) 076901.
[6] T. Accadia et al., JINST 7 (2012) P03012.
[7] F. Acernese et al., Class. Quantum Grav. 32 (2015) 024001.
[8] J. AAsi et al., Class. Quantum Grav. 32 (2015) 074001.
[9] LIGO Data Management Plant. Available at https://dcc.ligo.org/LIGO-M1000066/public
[10] Common Data Frame Format for Interferometric Gravitational Wave Detectors. Available at http://dcc.ligo.org/cgi-bin/DocDB/RetrieveFile?docid=329

6. EVN

6.1. Introduction

The European VLBI Network (EVN) is a very-long-baseline interferometric (VLBI) array of radio telescopes spread throughout Europe, Asia and South Africa. EVN is the most sensitive VLBI array in the world, thanks to the collection of extremely large telescopes that contribute to the network, and

conducts unique, high resolution, radio astronomical observations of cosmic radio sources.
As the EVN telescopes observe the same cosmic radio source simultaneously, the data are recorded on high capacity magnetic tapes, and these are later replayed and combined at a special purpose data processor, often referred to as the "correlator". The EVN software correlator at JIVE correlates essentially all of the EVN observations and about half of the global VLBI observations. The Correlator can input data recorded at rates from 32 to 1024 Mbps per station, up to a maximum of 16 stations at a single pass.

6.2. e-EVN

The EVN is conducting a development programme called e-EVN, which aims at creating a VLBI network based on modern technologies, including broadband data transmission and real-time data processing. The key element of the e-EVN programme is an introduction of the e-VLBI technique which enables real-time data transfer from remote radio telescopes to the central processing facility via optical fibre cables, as opposed to the "traditional" implementations of VLBI in which the data are first recorded at the telescopes on tapes or discs and then physically delivered to the processor. In parallel to establishing the network connectivity to EVN telescopes, the e-EVN programme includes relevant hardware and software upgrades of the EVN Data Processor at JIVE, enabling this facility to correlate e-VLBI streams in real time, and development of UniBoard, a generic high-performance FPGA-based computing platform for next generation radio astronomy. The e-EVN programme is officially recognised as a SKA Pathfinder for its continuing contribution to the SKA's requirements for new methods of high rate data transport and processing.

Bibliography

[1] http://www.evlbi.org/
[2] http://www.jive.eu/e-vlbi
[3] http://www.nexpres.eu/

7. LOFAR

7.1. Introduction

LOFAR, the LOw-Frequency ARray, is a new-generation radio interferometer constructed in the north of the Netherlands and across Europe. LOFAR's capabilities are revolutionising the astronomical capabilities in the 10-240 MHz range. It is one of the first radio observatories to feature automated processing pipelines to deliver fully calibrated science products to its user community. LOFAR's new capabilities, techniques and modus operandi make it an important pathfinder for the Square Kilometre Array (SKA).
LOFAR is being developed by a consortium of knowledge institutes, universities and industrial parties, led by ASTRON. The LOFAR facilities are jointly operated by the International LOFAR Telescope (ILT) foundation, as an observatory open to the global astronomical community. The infrastructure also has applications in Geophysics and Agriculture.

7.2. System Overview

There are two distinct antenna types: the Low Band Antenna (LBA) operates between 10 and 90 MHz and the High Band Antenna (HBA) between 110 and 250 MHz. These "sensors" are organised in aperture array stations containing local computing resources to perform beam-forming. In the Netherlands, a total of 40 LOFAR stations are being deployed with an additional 8 international stations currently built throughout Europe. The densely sampled, 2-km-wide, core hosts 24 stations and is located ~30 km from ASTRON's headquarters in Dwingeloo.
For the typical sampling rate of 200 MHz, the raw data-rate generated by the entire LOFAR array is 13 Tbit/s, far too much to transport in total. Even utilizing beam-forming at the station level, the long range data transport rates over the array are of order 150 Gbit/s requiring dedicated fibre networks. The datastreams from all LOFAR stations are sent via this high-speed fiber network infrastructure to a central processing (CEP) facility located in Groningen in the north of the Netherlands.
An overview of the data processing is shown in figure 7.1. At the computing center of the University of Groningen, data from all stations are aligned, combined, and further processed "online" using a flexible IBM Blue Gene/P (BG/P) supercomputer. When in full operation, LOFAR can produce observational data at rates up to 80 Gbit/s. After pre-processing, raw data products are written to a storage cluster for additional "offline" processing. This cluster currently hosts 2 Pbyte of working storage. Once on the storage cluster, a variety of reduction pipelines are then used to further process the data into the relevant scientific data products depending on the specific type of observation. The final scientific data products are transferred to the LOFAR long-term archive (LTA) for cataloging and distribution to the community. The archive of LOFAR science data products is expected to grow by up to 5 Pbyte per year.

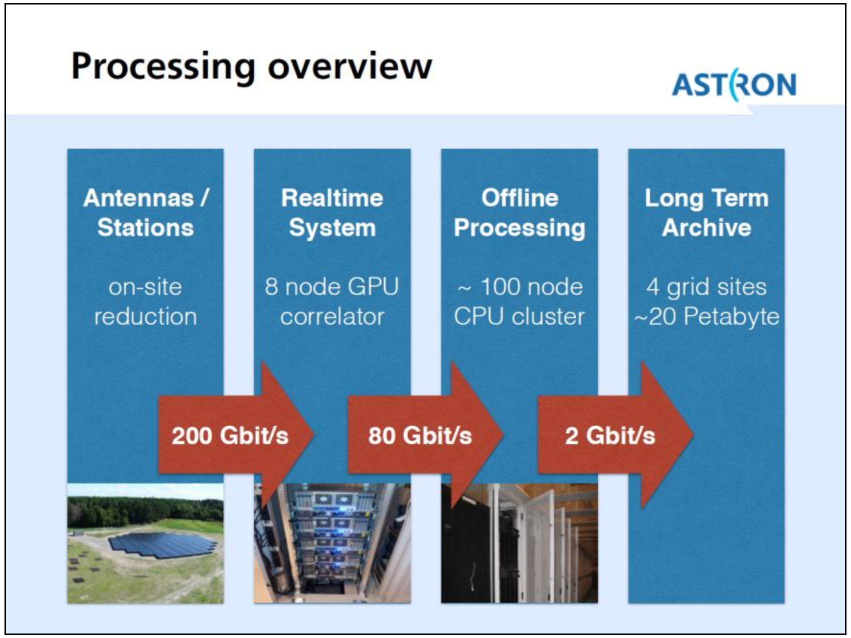

Figure 7.1. Slide taken from the presentation of Tammo Jan Keijkema and Sebastiaan van der Tol in the OBELICS f2f Meeting in 2016.

The network connecting LOFAR to the LTA sites in Groningen, Amsterdam and Jülich, Germany consists of light-path connections that are shared with LOFAR station connections and with the European eVLBI network (e-EVN). The current 10 Gbit/s bandwidth between the sites is sufficient for regular one-time LTA data transports but to allow transparent processing within the LTA it may grow to 60–80 Gbit/s bandwidth in the future.

7.3. Data Formats

In the interferometric imaging mode, a first pre-processing pipeline includes flagging of the data for interference and averaging. It produces so-called measurement sets (MSs), which are the standard format for the CASA package. These pre-processed visibility data sets are a large fraction of the archived data. Data are further processed in the Imaging Pipeline and final image products are made available via the LTA.

In the pulsar processing, the raw beam-formed data written by BG/P are stored on the LOFAR offline processing cluster and LTA in the HDF5 format. Several processing pipelines exit for the different science cases. The most advanced of these pipelines is the standard pulsar pipeline, “Pulp”, which is currently implemented within a python-based framework that executes the various processing steps.

Several conversion tools have been developed to convert these HDF5 data into other formats, e.g., PSRFITS, PSRCHIVE, PRESTO and SIGPROC. The long-term goal is to adapt these packages to all natively read HDF5, using the LOFAR Data Access Layer (DAL) for interpreting the HDF5 files. This adaptation has already been successfully done with the well-known program DSPSR.

Bibliography

[1] http://www.lofar.org/
[2] http://www.astron.nl/radio-observatory/radio-observatory
[3] M.P. van Harlem et al., A&A 556 (2013) A2. Available at http://dx.doi.org/10.1051/0004-6361/201220873
[4] L. Cerrigone, private communication.
[5] T.J. Dijkema and S. van der Tol, OBELICS f2f Meeting (2016).

8. EUCLID

8.1. Introduction

Euclid is an ESA mission from the European Space Agency (ESA) to map the geometry of the dark Universe. The mission will investigate the distance-redshift relationship and the evolution of cosmic structures by measuring shapes and redshifts of galaxies and clusters of galaxies out to redshifts ~2, or equivalently to a look-back time of 10 billion years. In this way, Euclid will cover the entire period over which dark energy played a significant role in accelerating the expansion. The launch date is planned for 2020.

8.2. Data Flow

During its lifetime the Euclid mission will deliver an unprecedented large volume of data for a space mission: about 850 Gbit of compressed data per day, where the daily telecommanding and communications period will be 4 hours. Lossless compression will be applied, with compression rates of the order 2 to 3. This data stream will be handled by the Euclid ground segment.

The complete survey represents hundreds of thousands images and several tens of Petabytes of data. Dedicated teams from the scientific consortia will process data from the instruments during all phases of the mission through the

Instrument Operation Centres and the Science Data Centres. The data handling system includes a common archive, the Euclid Mission Archive, which is accessible by all priviledged parties in the project. Once calibrated and qualified the data products will be stored in the Euclid Legacy Archive. This archive will be accessible by the astronomical community at large.
The scientific products are categorised in three data levels. The first data level consists of the raw decompressed telemetry frames, the second level consists of calibrated data with instrument signatures removed, and the third level consists of extracted scientific information such as catalogues. In addition, Euclid provides quick-release data, Level Q, representing transient products suitable for most purposes in astronomy, except for the core cosmology objectives.
Processed data included in the Euclid Legacy Archive will be provided in a format compatible with all of the international wide-spread standards (e.g., FITS and the Virtual Observatory) and accessible through public e-infrastructures and web services.

Bibliography

[1] http://sci.esa.int/euclid/
[2] http://www.euclid-ec.org/
[3] Euclid Definition Study Report. Available at http://arxiv.org/abs/1110.3193

9. LSST

9.1. Introduction

The Large Synoptic Survey Telescope (LSST) will be an automated astronomical survey system that will survey approximately 10000 deg^2 of the sky every few nights in six optical bands from 320 to 1050 nm. Over the planned 10-year baseline survey, it will uniformly and repeatedly image about 18000 deg^2 of the sky over 800 times. These data will be used to explore a wide range of astrophysical questions, ranging from discovering "killer" asteroids to examining the nature of Dark Energy. The LSST survey system consists of a large-aperture, wide-field, ground-based telescope currently being constructed on the El Peñón peak of Cerro Pachón in the Chilean Andes, a 3.2 gigapixel camera, and a data management system.

9.2. The LSST Data Management System

The rapid cadence and scale of the LSST observing program will produce approximately 15 TB per night of raw imaging data. These data will be automatically reduced to scientifically useful images and catalogs by the LSST Data Management (DM) system. The catalogs will be stored and offered to the users as relational databases which they will be able to query.
The DM system processes in real time the incoming stream of exposures by archiving raw images, generating alerts and updating catalogs ("Level 1" data products), and reprocesses annually the accumulated survey data to form a static, self-consistent Data Release ("Level 2" data products). The DM system also facilitates the creation of added-value ("Level 3") data products by providing the users with suitable software, application programming interfaces and computing infrastructure. 10% of the processing and storage capabilities will be devoted to this issue.
Over the ten years of LSST survey operations and 11 data releases, this processing will result in a cumulative processed data size approaching 500 PB for imaging, and over 50 PB for the catalog databases. The contents of the final data release alone are expected to be around 70 PB. Therefore, only the two most recent data releases will be kept on the fast storage and with catalogs loaded into the database. Small samples of data from older releases will also be kept for data quality monitoring and assessment, whereas their full contents will be archived to mass storage (tapes). Archived releases will be made available as bulk downloads in some common format (e.g., FITS binary tables).
Raw exposures, processed visit images, calibration frames and co-added images will be available as FITS files. The formats in which the image characterization data (e.g., backgrounds and point spread function) are served will depend on the final adopted algorithm in consultation with the scientific community. Alerts will be issued in VOEvent format or some other format that is broadly accepted and used by the community.

Bibliography

[1] http://www.lsst.org/
[2] Data Products Definition Document. Available at
https://docushare.lsstcorp.org/docushare/dsweb/Get/LSE-163/
[3] M. Jurić et al., The LSST Data Management System. Available at
http://arxiv.org/abs/1512.07914

10. H.E.S.S.

10.1. Introduction

The High Energy Stereoscopic System (H.E.S.S.) is an array of 5 Imaging Air Cherenkov Telescopes (IACTs) located in the Khomas Highland of Namibia. It is dedicated to investigate cosmic gamma rays in the energy range from 10s of GeV to 10s of TeV. The instrument allows scientists to explore gamma-ray sources with intensities at a level of a few thousandths of the flux of the Crab nebula (the brightest steady source of gamma rays in the sky).
The first of the four telescopes of Phase I of the H.E.S.S. project went into operation in summer 2002; all four were operational in December 2003, and were officially inaugurated on September 28, 2004. A much larger fifth telescope - H.E.S.S. II - is operational since July 2012, extending the energy coverage towards lower energies (from about 100 GeV to about 30 GeV) and improving the overall sensitivity by a factor of ~2 in the central energy range.

10.2. Data Flow

The maximum hardware event readout rate is about 900 Hz with an event size of 4.5 kB for each of the H.E.S.S. I telescopes and about 3.4 kHz with an event size of 10 kB for the fifth telescope. This leads to maximum data rates of approx. 50 MB/s for the primary scientific data during routine operation. To account for other hardware devices on-site as well as some additional capacity for tests, the DAQ system is required to process at least 80 MB/s, thereby having sufficient throughput for network connections and storage facilities.
The same ROOT-based data format, as decided by the H.E.S.S. Collaboration, is for both on-line and off-line analysis software as well as for storage. This approach allows high level analysis algorithms and visualisation tools to be used on a DAQ level. Therefore, the raw data that are sent from the Cherenkov detectors have to be converted into a common data format at a very early stage, which requires additional computing power. Moreover, the common data format allows data that are identical to those recorded by the DAQ to be easily simulated.
The H.E.S.S. array is situated in a remote location where no cheap, reliable and fast internet connection is available. Therefore, the data cannot be streamed to the different institutes of the H.E.S.S. Collaboration. Instead, the data have to be shipped to the institutes in regular intervals via magnetic tapes. The DAQ provides the resources to store the data on-site for up to three months until the shipped data has been verified in Europe. Every month H.E.S.S. Phase I takes 420 GB of data and for H.E.S.S. Phase II this value is around11 TB, while, the total available disk space on-site is about 60 TB.

Bibliography

[1] https://www.mpi-hd.mpg.de/hfm/HESS/
[2] A. Balzer et al., Astropart. Phys. 54 (2015) 67.

11. MAGIC

11.1. Introduction

The MAGIC (Major Atmospheric Gamma-ray Imaging Cherenkov) Telescopes are located on the Canary Island La Palma, Spain. The system consists of two 17 m diameter telescopes operating in stereoscopic mode and designed to study the universe and discover new Gamma-ray sources in the energy range from 50 GeV to 5 TeV.
The instrument MAGIC-I has been operational since 2004. In 2009, the second telescope MAGIC-II was constructed and commissioned, and the system has been successfully running in the stereoscopic mode ever since. During 2011 and 2012, the telescopes underwent a major upgrade, aimed at homogenization and improvement of the performance of both instruments.

11.2. Data Flow

The digitized signals from the telescopes are stored by a data acquisition system capable to process a sustained data rate of ~100MB/s (at 1 kHz trigger rate), the data volume to store being of the order of several Terabytes per night. Immediately after the data is taken, thanks to an independent network for the storage, an Online Analysis starts on the data.
MAGIC raw data consist of digitized pulses, recorded for every event and every pixel, with amplitudes expressed in ADC units. Originally stored in binary form, these data are translated to ROOT format by means of the merging and preprocessing program.
Scientific MAGIC results are made available through a Virtual Observatory server, and the release of public MAGIC results in FITS format.

Bibliography

[1] https://magic.mpp.mpg.de/
[2] J.A. Coarasa et al., The Data Acquisition of the MAGIC II Telescope, 29th ICRC (2005)

12. IceCube

12.1. Introduction

IceCube, the South Pole neutrino observatory, is a cubic-kilometer particle detector made of Antarctic ice and located near the Amundsen-Scott South Pole Station. Its main scientific goal is to map the high-energy neutrino sky, which is expected to include both a diffuse neutrino flux and point sources.
The in-ice component of IceCube consists of 5,160 digital optical modules (DOMs), each with a ten-inch photomultiplier tube and associated electronics. The DOMs are attached to 86 vertical “strings” arrayed over a cubic kilometer from 1,450 meters to 2,450 meters depth. The strings are deployed on a hexagonal grid with 125 meters spacing and hold 60 DOMs each; except for 8 strings at the center of the array, which were deployed more compactly. This denser configuration forms the DeepCore subdetector, which lowers the neutrino energy threshold from 100 TeV down to 10 GeV. The pilot project AMANDA is contained within the IceCube detector, but was decommissioned in 2009.
A surface array, IceTop, was also deployed. It consists of 81 stations located on top of the same number of IceCube strings. Each station has two tanks, each equipped with two downward facing DOMs. IceTop, built as a veto and calibration detector for IceCube, also detects air showers from primary cosmic rays in the 300 TeV to 1 EeV energy range.
The IceCube Neutrino Observatory was built under a National Science Foundation (NSF) Major Research Equipment and Facilities Construction grant, with assistance from partner funding agencies around the world. The University of Wisconsin–Madison is the lead institution, coordinating data-taking and Maintenance and Operations activities. The international IceCube Collaboration, with more than 40 institutions worldwide, is responsible for the scientific research program.

12.2. Data Flow

The IceCube sensors collect light, which is subsequently digitized and time stamped inside each DOM. Data is sent to the computers in the IceCube Lab on the surface through a cable network enabling data rates up to ~900kb/s for the most remote DOMs. The IceCube Lab receives raw data collected from IceCube, ARA and DM-Ice. One terabyte of data comes off the detectors daily and is filtered down to 105 gigabytes for satellite transmission to the North.
The first level of reconstruction and filtering of IceCube data happens in near real-time at the South Pole station. The raw data are written to disk and categorized into data sets for satellite transmission to the Northern hemisphere and/or archiving to tape by a system called SPADE (for South Pole Archival and Data Exchange). Raw unfiltered data is recorded on one set of tapes, and all other data – filtered data from the detector, monitoring, calibration, or any other ad hoc data – on another set of tapes. Potentially more than one copy of a given dataset can be made. SPADE creates tape archives in parallel with transferring the data, so as to avoid delaying data transfer. The data is prepared for transfer and sent to the Northern Hemisphere by one of three transfer methods: high-volume bulk transfer over TDRSS satellite, low-volume transfer by scp, or transfer as email attachment.
SPADE also begins the process of cataloging IceCube data by producing a metadata file for each data file it receives. This XML file follows a metadata specification called "DIF Plus". The "DIF" portion of the XML is metadata in Directory Interchange Format (DIF), which is the format required by the Office of Polar Programs for all experiments funded by the National Science Foundation, and is required for cataloging data in the Antarctic Master Directory. The "Plus" portion of the XML is a collection of fields that have been defined specifically for IceCube, to aid in the storage of data in the data warehouse and later searching of the contents of the warehouse. SPADE tars and gzips the data and "DIF Plus" metadata files, transferring them together to the Northern Hemisphere.
In-depth analysis of selected physically interesting events is carried out in the North. Reconstruction and filtering modules are provided by the physics working groups in the collaboration and are written in the IceTray analysis framework.
It is distinguished between the level 1 data, which is the filtered stream coming from the South Pole, and level 2 cuts providing the basic reconstruction of upgoing muons and good reconstructions of downgoing muons. Further levels of data filters are set by working groups within project collaborations

Bibliography

[1] https://icecube.wisc.edu/
[2] F. Halzen and S. Klein, Rev. Sci. Instrum. 81 (2010) 081101.
[3] R. Abbasi et al., NIM A 601 (2009) 294.
[4] IceCube Open Data Project
[5] http://wipac.wisc.edu/science/computing

13. ANTARES

13.1. Introduction

The ANTARES detector, located 40 km off shore from Toulon at 2475 m depth, was completed on 29 May 2008, making it the largest neutrino telescope in the northern hemisphere and the first to operate in the deep sea. The ANTARES telescope extends in a significant way the reach of neutrino astronomy in a complementary region of the Universe to the South Pole experiments, in particular the central region of the local galaxy. Furthermore, due to its location in the deep sea, the infrastructure provides opportunities for innovative measurements in Earth and sea sciences.

The detector consists of an array of approximately 1000 photomultiplier tubes in 12 vertical strings, spread over an area of about 0.1 km^2 and with an active height of about 350 metres. The basic unit of the detector is the optical module (OM) housing a photomultiplier tube. The OMs are grouped together in “storeys” of three modules and interconnected via an electro-mechanical cable.

Another project benefiting from the deep sea infrastructure is an R&D system called AMADEUS of hydrophones which investigates the detection of ultra-high energy neutrinos (>100 PeV) using the sound produced by their interaction in water.

13.2. Data Flow

The Data Acquisition (DAQ) system is based on the “all-data-to-shore” concept. In this mode, all signals from the PMTs that pass a preset threshold (typically 0.3 Single PhotoElectron (SPE)) are digitized and sent to shore, where they are processed in real-time by a farm of commodity PCs. The raw data are packed offshore as arrays of hits of predefined time frame duration of about 100 ms. These data packets amounts to 60–200 kB and the total data flow ranges from a couple of Gb/s to several tens of Gb/s, depending on the level of the submarine bioluminescent activity.

The data are then sent to shore in such a way that the data collected for the full detector for the same time frame are sent to a single data filter process in the onshore data processing system. There, the data flow is reduced, on average, by a factor of about 10,000. The filtered data are written to disk in ROOT format by a central data writing process and sent via the commercial fiber optic network to be stored remotely at a computer centre in Lyon. The count rate information of every PMT is stored together with the physics data.

Bibliography

[1] http://antares.in2p3.fr/
[2] M. Ageron et al., NIM A 656 (2011) 11.
[3] J.A. Aguilar et al., NIM A 570 (2007) 107.