PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Conceptual design of the data handling system for the European Solar Telescope

Ermolli, Ilaria, Cauzzi, Gianna, Collados, Manuel, Paletou, Frederic, Reardon, Kevin, et al.

Ilaria Ermolli, Gianna Cauzzi, Manuel Collados, Frederic Paletou, Kevin Reardon, Jean Aboudarham, Roberto Cirami, Rosario Cosentino, Dario Del Moro, Paolo Di Marcantonio, Fabrizio Giorgi, Martine Lafon, Ermanno Pietropaolo, Paolo Romano, "Conceptual design of the data handling system for the European Solar Telescope," Proc. SPIE 8448, Observatory Operations: Strategies, Processes, and Systems IV, 84481S (13 September 2012); doi: 10.1117/12.926109

SPIE.

Event: SPIE Astronomical Telescopes + Instrumentation, 2012, Amsterdam, Netherlands

Conceptual design of the data handling system for the European Solar Telescope

Ilaria Ermolli[a], Gianna Cauzzi[b], Manuel Collados[c], Frederic Paletou[d], Kevin Reardon[b], Jean Aboudarham[e], Roberto Cirami[f], Rosario Cosentino[g], Dario Del Moro[h], Paolo Di Marcantonio[f], Fabrizio Giorgi[a] , Martine Lafon[d] , Ermanno Pietropaolo[i], Paolo Romano[g]

[a]INAF Osservatorio Astronomico di Roma, via Frascati 33, 00040 Monte Porzio Catone, Italy; [b]INAF Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, 50125 Firenze, Italy; [c]Instituto de Astrofisica de Canarias, via Lactea, 38205 La Laguna, Tenerife, Spain; [d]Universit ́e de Toulouse, CNRS, 14 av. E. Belin, 31400 Toulouse, France;[e]Observatoire de Paris, 5 Place Jules Janssen, 92190 Meudon, France;[f]INAF Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, 34131 Trieste, Italy; [g]INAF Osservatorio Astrofisico di Catania, via S. Sofia 78, 95123 Catania, Italy; [h]Dipartimento di Fisica, Universita' Tor Vergata, 00100 Roma, Italy;[h]Dipartimento di Fisica, Universita' di L'Aquila, L'Aquila, Italy.

ABSTRACT

We present an overview of the conceptual design of the data handling unit of the ECS, the Control System for the European Solar Telescope (EST). We will focus on describing the critical requirements for this unit resulting from the overall design of the telescope, together with its architecture and the results of the feasibility analysis carried out to date.

Keywords: solar telescope, EST, telescope control, control software

1. INTRODUCTION

The EST (European Solar Telescope) is a 4-meter class telescope that will be optimized for the study of the magnetic field and dynamics from the deep photosphere to the upper chromosphere of the Sun (Collados et al. 2010 [1,2]). This will require diagnostics of the thermal, dynamic and magnetic properties of the plasma over many scale heights, by using multiple wavelength imaging, spectroscopy and spectro-polarimetry. To achieve these goals, EST will specialize in high spatial and temporal resolution using instruments that can efficiently produce two-dimensional spectral information. EST will improve by a considerable factor the presently achieved spatial resolution. In addition, the operation of several narrow-band tunable visible and near-infrared imaging instruments together with grating spectrographs, all with polarimetric capabilities, and large format broad-band imagers, will allow the simultaneous observations of photospheric and chromospheric layers and the study of the temporal evolution of the three-dimensional structure of solar magnetic fields. These aspects will make EST a unique infrastructure.

The EST project is promoted by the European Association for Solar Telescopes (EAST), which is a consortium with 15 research institution members from 15 European countries with the aim, among others, of undertaking the development of the telescope in the Canary Island (Spain). The conceptual design study, aiming to demonstrate the scientific, technical, and financial feasibility of the project, started on February 2008 with funding by the European Commission through an FP-7 Collaborative Project and completed in June 2011. The study involved 29 partners and 9 collaboration institutions from 14 countries and covered all key aspects of EST design, from optical configuration of the telescope and instruments, to dome and auxiliary buildings needed to operate the facility.

Observatory Operations: Strategies, Processes, and Systems IV, edited by
Alison B. Peck, Robert L. Seaman, Fernando Comeron, Proc. of SPIE Vol. 8448,
84481S · © 2012 SPIE · CCC code: 0277-786X/12/$18 · doi: 10.1117/12.926109

The operation of the EST infrastructure, as well as the supervision as a whole of all its sub- systems, shall be carried out by an integrated control system, which will be characterized by a distributed, object-oriented architecture and a common software employed in the entire system. The EST control system (ECS, hereafter) shall also provide efficient management of the data and metadata produced by the facility, and their transmission from each sub-system to a real-time repository, as well as to users and to temporary and permanent archives.

The ECS shall consist of the EST control facilities and other elements, specifically the common software, the safety and monitoring facilities, and a sub-set of data analysis applications. It shall allow EST users to manage the flux of information generated by the EST facility. The term User encompasses five categories: engineers, technical operators, staff astronomers, responsible for the operation of the EST, visiting astronomers. The flux of information generated by the facility shall consist of science data, metadata, and monitoring data.

The ECS shall consist of the four main blocks shown in Figure 1. A description of ECS requirements and architecture is given by Di Marcantonio et al. (2012) [3].

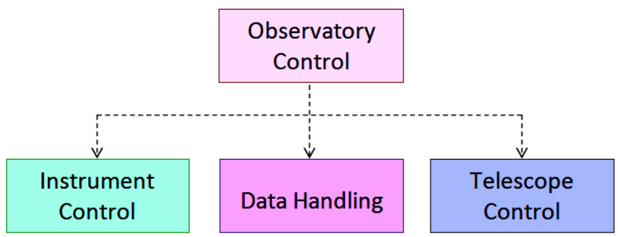

Figure 1. Main blocks of ECS.

The partition of the ECS into four main blocks derives from the topological, functional and management reasons. These four blocks have to work in a coordinated manner though they are to a certain degree independent from each other. In particular, the Observatory, Instrument, and Telescope Control blocks are required to enable EST to execute various modes of both operation (classically scheduled, queue observing service, engineering) and observation (day-time solar observation, day-time non-solar observation, night-time observation), by various users (engineers, technical operators, staff astronomers, visiting astronomers, responsible of EST operation). This implies control, monitoring, and operation of all the various main parts of the EST. These blocks are required to enable EST also to point and track a position in various coordinate systems (none, sidereal, heliocentric, heliographic) and to perform observations with different instrument setups, also in the framework of simultaneous and coordinated campaigns carried out with various instruments and other telescopes. On the other hand, the Data handling block (Data Handling System, DHS, hereafter) of the ECS shall provide for the recording of all the data and metadata acquired by the telescope and instruments, for the access and display of the data needed to facilitate telescope operations, and for the transfer of the data produced away from the telescope to intermediate and then to long-term data archiving facilities.

The four main blocks of ECS shall consist of 16 sub-systems, hereafter referred to as high-level sub-systems of ECS. These sub-system shall comprise, e.g.:

- The Operation Repository Sub-system, which shall allow to store all the data generated by the EST operation.
- The Observing Program Management Subsystem that shall provide EST users with the tools needed for the creation, modification, and submission of observing proposals. It shall allow the EST end users (Principal Investigator, PI) to plan observations, as well as to monitor the status of granted observations.
- The Database Management Sub-system, which shall provide the EST end users with the access to the scientific and operational data generated by the EST through remote internet access, or through the transfer of physical media if necessary. All these data will become public after a proprietary period during which the PI will have exclusive right to the data.
- The Data Processing Sub-system that shall allow to process both scientific and engineering data, to perform Quality Control of science data, and to archive the data managed. The DPS shall provide a common framework for the processing of the science data, including quick-look analysis and reduction. This system shall be a distributed facility,

since some of the data processing tasks might be executed at different processing nodes to improve the system performance. The instrument teams shall have to provide the specific reduction templates of science data, i.e. the set of rules to be followed and codes to be applied for the reduction of the science data taken with a given instrument.

2. DATA HANDLING SYSTEM REQUIREMENTS

The EST will have a main instrumentation station at the Coudé focus with three types of instruments, each one composed of different channels to observe different wavelengths: broad-band imagers, narrow-band tunable filter spectropolarimeters and grating spectropolarimeters. In the baseline design the various channels include 32 detectors. The science goals of EST and technical solutions identified to date point to use detectors with 4000x4000 pixels, 2 bytes/pixel depth and reading 100 frames per second. The data flux estimated from each instrument is 18.9 GB/s to 24.6 GB/s, depending on the instrument; the overall data flux from the three types of instruments exceeds 80 GB/s. Twelve hours of EST operation at maximum data rate translate in a data volume produced of over 3.5 PB per day. The data store at the telescope shall have a total capacity for several days of consecutive observations of order of 25 PB per week. In addition to the scientific data, many other details of the facility operation shall also need to be recorded. The data store at the telescope should never become a limitation to the infrastructure operation and amount of high-quality data that can be obtained at the telescope. Besides, the connection from the detectors an control computers to the data store shall occur through a dedicated layer of the control network, to allow for a highly efficient point-to-point transmission of the camera data stream.

As introduced above, it may be necessary to have some subset of the acquired data processed at the telescope prior to moving it off the summit The near–real–time reduction of a small subset of the data for data evaluation will need only relative modest, off–the–shelf computing resources. On the other hand, bulk data–volume reduction will come in two forms, data compression and image reconstruction. However, the computing power required for such processing may not be reasonable at the telescope even by the time EST becomes operational, primarily because of the power and cooling constraints at the summit facility. An alternative exists, however, in massively–multi–core Graphics Processing Units (GPU). Existing GPU, such as the ATI Radeon 5970, have 960 cores on a single PCI package and can achieve 1 teraFLOP/sec, after taking into account the bottlenecks imposed by the low memory bandwidth per operation on this type of hardware.

The facilities dedicated to data handling at the telescope and other sites shall also allow for the re-processing and re-use of old data; in addition, they shall serve as dispenser of Science-ready data (VO-compliant) and processing tools. They shall also allow for the publication of data into VO and the interoperability with databases produced by other facilities. The display of the data generated by the EST shall provide near real-time feedback to the operators and users about the telescope operations. The data display may have to combine the raw data stored on disk with metadata extracted from the repository in order to provide a coherent set of information to the users, also for engineering purposes. In addition, it may be necessary to have some subset of the acquired data processed at the telescope prior to moving it off the summit. The purposes of such processing might include: 1) Data quality assurance; 2) Data volume reduction; 3) Evaluation of the solar structures being observed.

The data volume daily generated could be reduced to 2/3 with evaluation of observations quality and to 1-5 to 1/10, depending on the instrument, with on-site data processing. In addition, a lossless data compression can be applied to the data, such as that performed by JPEG-2000 or FPACK, to compress most integer-format data by a factor of approximately two. The JPEG-2000 is a lossless wavelet-based compression algorithm already applied to SDO/AIA data and tested on data set such as the one that shall be generated with the EST. It is worth to consider that the data production by the LHC project LHC (basically 1 DVD/s, ~5 GB/s) is currently managed by using grid computing, state-of-the-art data storage facilities, and high-bandwidth networks. Therefore, independent of the capacity to reduce the EST data volume (~80 GB/s) on-site and of the improvement of information technologies by the start of EST operation, the management of data generated by the EST shall rely also on European e-Infrastructures (e.g. http://cordis.europa.eu/fp7/ict/e-infrastructure/home_en.html).

3. REFERENCE MODEL

The system requirement analysis shows that the standard concept of proposal submission/approval/execution and then data retrieval applies also for EST in order to exploit its scientific outputs at optimum level. Figure 2 shows the reference model considered for the operation of the EST infrastructure.

The Principal Investigator shall prepare its observation program by using the Observing Program Management tools. The observation program shall consist of a set of individual Observation Sequences. The Observing Program Management shall provide tools for the submission of both Phase I and Phase II proposals. The submission shall result in a script to be executed automatically. The script shall contain commands for the execution of the observation and a description of the observation; it shall be stored in the Operation Repository. Staff Astronomers and TAC shall be allowed to browse all Observations stored in the Operation Repository by using the Inspector. Staff Astronomer shall be allowed to prepare long- to mid- term plans for telescope operation based on accepted observations and other defined constraints by using the Scheduler.

At the time the accepted observation shall be carried out in queued observing mode, the Scheduler shall select the best candidate Observation that matches the actual observing conditions and the ones required in the submitted proposal; then it shall pass the Observation to the Sequencer. The Sequencer shall take as input the Observation description and depending of the required instruments, observing mode, and observation set-up, it shall coordinate the various subsystems that shall control the operation of telescope and instruments to perform the observation with the required conditions. This shall include sending of the configuration, pointing and guiding commands to the Observing Engine.

The Observing Engine shall manage to coordinate operation of all the sub-systems controlled by the Telescope Control System and by the Instrument Control System during the execution of the observation sequence. During the execution of the observation, the Observing Engine shall be provided with the outcomes of some common services e.g. by the Logging and Alarms, Monitoring, Configuring service Sub-systems.

The science data shall be generated during the execution of the observation sequence by the Detector Control Service that shall be coordinated by the Instrument Control Sub-system. Metadata describing telescope and environmental conditions during the observation execution shall be also generated; they shall be associated with the corresponding science data most likely only after completion of the observation run and stored in the Operation Repository.

Once the raw science data have been generated, they may be processed with the Data processing Sub-system for quick-look analysis and quality control, by using, e.g., most recent relevant calibration data stored in the Operation Repository. Staff Astronomers and Visiting Astronomers shall be allowed to display the result of the data processing and to apply quick-look analysis on them, by using the tools provided by the Data Processing and Inspector Sub-systems.

After completion of observation sequences and telescope operations, all the science data generated by the EST shall be associated with the corresponding metadata, and analysed with the tools provided by the Data Processing and Data Management Sub-systems to reduce the data volume on-site. Then the data shall be moved to the first-support data centre for completion of the data processing and for reduction of the volume of the data to be distributed to other processing centres with the tools provided by the Data Archive Management Sub-system. With some intermediate processing steps, the data shall be processed and formatted with VO compliant standards, then distributed by the Science Database following the policy that shall be defined for the distribution of EST data.

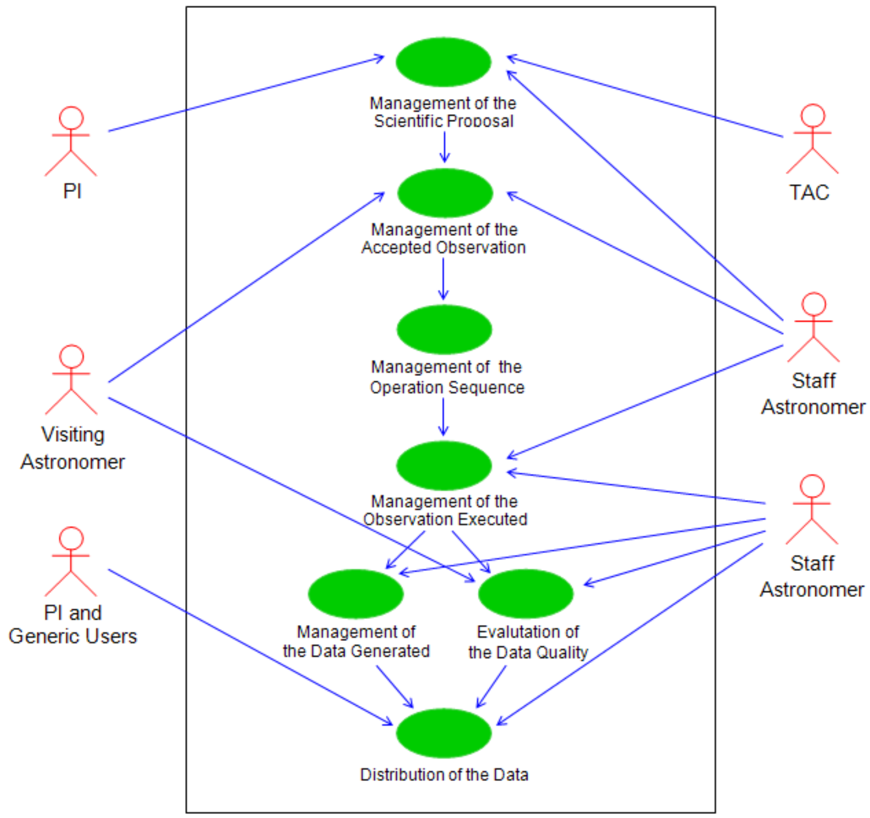

Figure 2. Reference model.

4. DATA HANDLING SYSTEM DESIGN

Figure 3 sketches out the hardware components for the Data Handling block All the data derived from the instruments shall be written to a central data store as it is acquired. This central store shall utilize a shared pool, which shall be partitioned among experiments and instruments in a dynamic manner. By using a homogeneous storage facility, the data can be efficiently managed using a single set of applications.

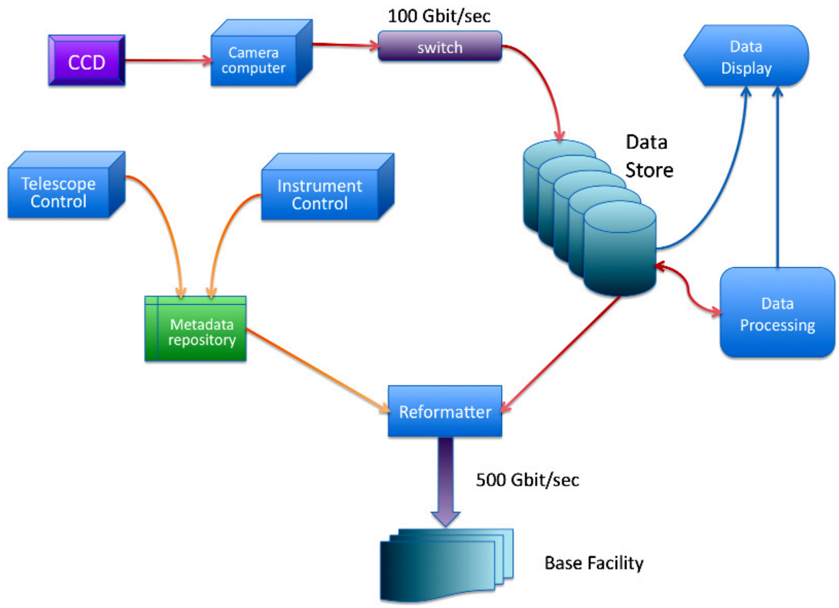

Figure 3. Physical model of the hardware components of the Data Handling block.

In the reference design, the central storage area, with a size of approximately 25 PB, is based on traditional hard-disk technology. Alternatives may be explored closer to the time of construction, depending on technological development of technologies. Assuming a hard disk size of 50 TB in 2018 and a fully redundant RAID configuration for speed and security, approximately 1000 disks shall be required for the data store. This size of data store shall require several full racks, plus space for the associated electronics and cooling. The total power consumption for the central store may be of order of 30 kW. Given the scaling of disks capacity, the proposed data store is comparable presently to the implementation of a 600 TB data store, complicated but not a serious challenge. Petabyte storage arrays are already available from multiple vendors.

As introduced above, the connection from the detector computers to the data store shall occur through a dedicated layer of the control network, to allow for a highly efficient point-to-point transmission of the camera data stream. To maintain the high data rates, enterprise class hardware shall be needed as well as fiber optic connections. No current technology can support the transfer rate (100 Gbit/sec) from the large format detectors expected for the EST, but development of 100 Gbit Ethernet promises to reach these speeds in the near future at a moderate cost (with the link latency still to be determined). Other technologies, such as PCI Express, Infiniband, Myrinet may offer similar transfer rates in the coming years. Therefore a single connection may suffice for almost any but the most extreme EST detectors (12k x 12k x 100 Hz). For such detectors or if technologies do not reach the desired speeds in the expected timeframe, multiple connections could be used for a single camera, streaming data from different taps on the chip over separate physical connections.

The processing needs at the telescope shall require processing nodes on-site. Speckle interferometry on the images acquired at the full rate of 100 frames per second would require a sizable computing cluster, comprising several 100 cores, with significant RAM and fast interconnects, that may be prohibitive to manage on-site. The implementation of a system based on GPU's promises a solution to this problem. At the moment, blind-deconvolution approaches seem not feasible on-site.

All the metadata collected from the instrumentation and telescope systems shall be stored in real time in the central repository. Some metadata may be recorded at near real-time rates for each images (acquisition time, instrument configuration, AO seeing information). Other metadata from instruments or the observatory control systems may be recorded at slower rates. Receiving data from many subsystems at a variety of rates, the repository shall be strongly transaction driven. The most reliable system for recording such information at these rates shall be a database system, most likely a well tuned relational database. Such a system shall have to support the writing of information at a high rate into the database, while at the same time supporting queries of the database contents from other telescope systems in near real-time (perhaps with a lower priority). An in-memory database might support such demanding real- time uses, while this information could transition to a traditional disk-based database to maintain a historical record of the database contents (e.g. for engineering purposes). The entire database should probably be mirrored to the first support center or long-term archives (though not in real time).

The metadata repository shall also be queried to produce the FITS or other metadata headers to attach to the raw Level-0 data in the data store in the process of converting those data to Level-1 formatted data. A timestamp and other essential image identifying information, written both in the metadata repository and as a small tag attached to each Level-0 image, would be used to associate the appropriate metadata to the raw data frame. The reformatting of the data may take place on a resource-available basis, possibly after the end of the daily observations, or on-demand for certain small datasets. This latter would facilitate the processing of limited datasets through a reduction pipeline for immediate data quality assurance purposes.

The data display system shall extract the necessary data from the data store and possibly the metadata repository. These data may need to be partially processed prior to display, so some computing power shall be incorporated in the display system. The algorithms to be applied may be fixed, with the possibility of incorporating new software as needed. In order to extract images from the data store and read them for display at rates of up to 10 frames per second and 0.25-0.50 sec latency, a high-speed connection to the data store shall be required. Proper access controls shall need to be implemented to avoid resource contention or the reduction of the data acquisition rate due to data display tasks. Multiple data displays may be operating simultaneously showing different data.

Figure 4 shows the general strategy for the data distribution beyond the telescope, down to one or several processing centers, and to the terminal VOCDB.

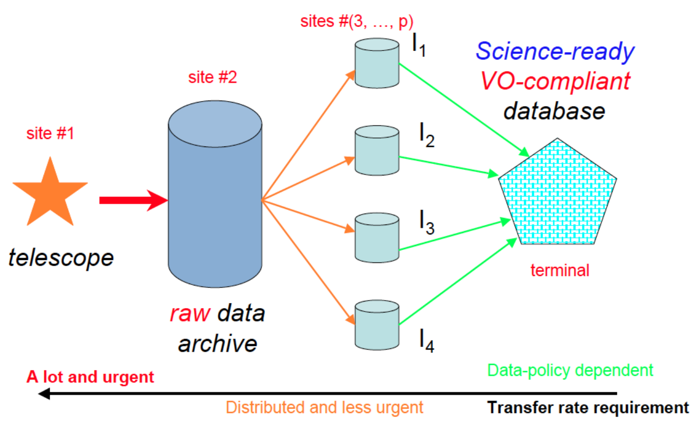

Figure 4. View of the dataflow from the telescope to the first support data center , dedicated processing centers distributed across Europe and, finally, the EST VOCDB.

The EST reference design calls for the data to be transferred away from the telescope to an intermediate, first support, data storage facility at the base of the mountain. This facility shall allow a more convenient and economical installation of a portion of the data handling facilities and better working conditions for the staff. The base center shall likely to be somewhere in the Canary Islands (e.g., IAC in Tenerife), and connected with a very fast dedicated Ethernet connection to the telescope (see e.g., the LSST case). To date, 10 Gbps (per wavelength) dark fiber links are available between OT-IAC (Tenerife) and ORM-CALP (La Palma)-IAC (Tenerife). A slower (35 Mbps) back-up link is also available between ORM-OT. Rates significantly greater than 10 Gbps per wavelength are expected at the time of EST operations. In particular, in the 1-2 years to come, 40 Gbps per wavelength connections will be available so that it can be reasonably expected that 100 Gbps per wavelength connections will be freely available by the beginning of the EST commissioning period.

The base center shall comprise a large data store, with a pool capacity at least several times that necessary at the telescope, on the order of 50 PB. It may be easiest to use a hard-disk based storage system also for this purpose, although tape-based or other removable media might be considered especially for longer term storage. Presumably the data shall only be staged here temporarily before onward transfer to other long-term archives. Given the potential slower rate of data transfer away from the Canary Islands than within continental Europe, it may be necessary to process and reduce the data volume at this stage. This may require significant processing capacity at the base facility.

After this first stage, the data would continue to flow to continental Europe, down to dedicated "instrument PI institutes" (hereaftrer referred to as {In}), acting as specific processing centers. This shall have the advantage that EST infrastructures at the Canary Islands take charge of telescope operations only, while production of science-ready data is in charge of dedicated resources. Another point is that network capabilities within mainland Europe may be still more favourable 10 years from now then the one available between the Canary Islands and the peninsula. The ongoing ACE (Africa Coast to Europe) project is deploying WDM technology submarine cables that could be utilized to flow the data from the Canary Islands to continental Europe. With WDM, cable capacity can be increased without additional submarine work. With an overall potential capacity of 5.12 Tbps, the system will support the 40 Gbps technology by 2012. The relevant stages of this new cable system are France (northwestern), Portugal (Lisbon), and Spain (Tenerife). Then, the GEANT2 network should be suitable for the flow of EST data within continental Europe between the various centres involved in the tasks of producing science-ready data and the ones storing/diffusing them.

In relation to the data flow off the mountain and processing of large data volumes, it is important to look on EU projects such as DANTE, for advanced network for research and education, EGEE, for grid computing and the processing of large datasets, PRACE, for advanced computing.

Each processing center {In}, responsible for its instrument data reduction pipe-line, and with the task of providing VO-compliant, science-ready data to the VOCDB, should be equipped with a high performance computer able to handle several 10's of TB produced daily, on average. The {In} could be replaced by a single European Solar Physics Institute where all the expertises for producing EST science-ready data can be gathered. Then, the science-ready data shall be moved or duplicated from the {In} to the mainland Europe VOCDB. This center shall be in charge of the long-term data storage and the VO-diffusion of EST data. The VOCDB shall take charge of the interoperability with the VOCDB from other facilities, e.g. ATST. This could be achieved with the EST-VOCDB either serving as a duplication center for ATST science-ready data or using virtualization tools for accessing the data stored on the ATST-VOCDB.

In the reference design, the long-term data storage utilized the current standard technology of rotating magnetic disks. These "hard disks" have undergone consistent technological development and have a large commercial base driving future development. It might be reasonable to expect only a factor of 10-25 improvement in the hard-disk areal density in the coming decade. This would result in hard disks with upper capacities of 20-50 TB (assuming the same number of platters) when EST becomes operational.

In the next years to come, rotating disk technologies will remain the best choice allowing the storage of large volumes. However, there are several new data technologies being developed for high-density, persistent data storage. These include 3D optical of Holographic (HVD) storage and Magnetoresistive (non-volatile) RAMs (MRAM) that have started to come into the market (e.g., InPhase Technologies/DSM, Freescale). In addition, SSD-based products (e.g., nimbusdata.com) could provide faster access to the data than rotating disk technologies, although capacities are generally lower than for HDD (see e.g., fusionio.com). SSD is growing fast and beside the fast access (up to 1 Gbps), this is also a technology much less demanding than HDD in terms of necessary power (impact on operations costs). This technology could be of some interest also for the handling of the data at the telescope. Following the evolution of technology in the next years to come, hybrid storage could represent a solution for the storage of data generated by the EST, with the hierarchy RAM/SSD/HDD/tape (e.g., LTO-5) depicted by the increased storage capacity and decreased access time.

The VOCDB shall run a significant amount of science-ready data available on-line. In order to allow for a comfortable time-response to the users, the data frequently accessed together should be co-located and re- arranged when moving from the {In} processing centers to the VOCDB. Indeed, the data access from users is driven by an object-scientific interest. The science-ready data from the {In} sets derived with the instruments should be re-organized in {On} storage units, depending on a few object classes, such as, e.g., active regions, quiet sun, 2nd solar spectrum, prominences and spicules. At the moment, a RDBMS like MySQL seems to fit in with the requirements and reference design sketched above. It is a freeware, widely used in science and commercial databases. iRODS is the state-of- the-art software for data storage visualization that would allow for the handling of the {On} storage units and the remote access to other VOCDB units.

5. CONCLUSIONS

We have summarized the main functional and technical requirements for the definition of the Data Handling System of the EST. We have also described the technical alternatives identified to date and outlined the system architecture that was presented at the final review meeting of the EST design study project.

ACKNOWLEDGMENTS

The conceptual design study of EST has been partially supported by the European Commission through the collaborative project Nr. 212482 "EST: the large aperture European Solar Telescope".

REFERENCES

[1] Collados, M., et al., "European Solar Telescope: project status" Proc. SPIE, 7733, 77330H (2010)

[2] Collados, M., et al., "European Solar Telescope: Progress status" Astron. Nach., 331, 615 (2010)

[3] Di Marcantonio, P., et al., "Conceptual design of the control software for the European Solar Telescope" Proc. SPIE, 8451 (2012)