

Compendium of e-Infrastructure requirements for the digital ERA

May 2015

Contents

Foreword 3

Acknowledgements 4

CTA Data Management 5

DARIAH: Digital Research Infrastructure for the Arts and Humanities 12

European Multidisciplinary Seafloor and water-column Observatory (EMSO) 18

Fusion on the grid: transport calculations 30

Grid-empowered molecular simulations: from first principles to new materials 37

Hydrometeorological research at the computational frontier 47

The Life Science grid community 61

The e-Infrastructure for EISCAT_3D 70

Transparently scale next-generation high throughput biobank analysis between institutes, computational clusters and grids 81

WeST: a Worldwide e-Infrastructure for Structural Biology 86

Computing for the LHC 93

Summary – conclusions and recommendations 102

Foreword

Dear colleagues

On behalf of the EGI Council, I'm delighted to announce the release of this compendium of European Digital ERA Requirements, prepared in collaboration with representatives of existing and prospective EGI user communities.

The original concept behind this publication was to prepare a collection of scientific use cases, including their requirements to inform subsequent engagement and outreach activities. We also wanted to create a document that could be used to demonstrate the potential of EGI as enabler of a system of platforms for the virtual research communities, research infrastructures, the long tail of science and business, thus accelerating the implementation of the digital ERA.

In this publication, we chose to give the word to the research communities to enable them to explain, in their own words, how their science can benefit from co-development of technology, co-delivery of services, access to expertise and federation of resources in Europe and beyond.

The diversity and complexity of the scientific case studies here presented clearly demonstrate how federated and shared e-Infrastructure have become for science.

EGI is committed to support this trend by providing user-centric services and a user-centric approach to development that will allow these communities to benefit from co-development of technology, co-delivery of services, access to expertise and federation of resources in Europe and beyond.

A word of thanks to Ludek Matyska – former EGI Council Chair – who worked for the success of this initiative during his mandate.

Matthew Dovey
Chair of the EGI Council

Acknowledgements

The Editorial Team would like to thank all the contributors and authors of the European Digital ERA Requirements compendium for their efforts and dedication, and are confident that we can count on their continued cooperation.

We believe that this document will encourage other international research communities to showcase their needs for digital science and hope to include them in future editions of the compendium.

The Editorial Team

Michał Turala

Isabel Campos

CTA Data Management

G. Lamanna✉, C. Vuerli, L. Arrabito, N. Neyroud

The Cherenkov Telescope Array (CTA) Consortium

11 April 2013

✉ Giovanni.Lamanna@lapp.in2p3.fr

Overview

The proposed Cherenkov Telescope Array (CTA) is a large array of Cherenkov telescopes of different sizes and deployed on an unprecedented scale. It will allow significant extension of our current knowledge in high-energy astrophysics.

The spectacular astrophysics results from the current Cherenkov instruments in the field of very high-energy γ (gamma)-ray astronomy (VHE, energies >100 GeV), have motivated both the astrophysics and particle physics communities in designing (Actis et al., 2011) and preparing (CTA Consortium, 2013) a next-generation, more sensitive and more flexible facility, able to serve a larger community of users.

The CTA¹ was first publicly presented to an ESFRI panel in autumn 2005 and it has been considered since 2008 in the list of main future research infrastructures within the roadmaps of international agencies and forums, namely: ESFRI, ASPERA, ASTRONET, US-Decadal Survey, OECD. Currently in its pre-construction phase (also funded since 2010 through a dedicated EC-FP7 Preparatory Phase project), the CTA international consortium counts the participation of more than 1,000 members working in 27 countries in Europe, America, Asia and Africa.

The CTA will carry out observations to study the radiation at γ -rays energies. Radiation at such energies differs fundamentally from what is detected at lower energies since γ -rays cannot be generated by thermal emission from hot celestial objects. The energy of thermal radiation reflects the temperature of the emitting body, and apart from the Big Bang there has been nothing hot enough to emit such γ -rays in the known Universe.

γ -rays can be generated when highly relativistic particles, accelerated for example in the gigantic shock waves of stellar explosions, collide with ambient gas, or interact with photons and magnetic fields. The flux and energy spectrum of the γ -rays reflects those of high-energy particles. They can therefore be used to trace these cosmic rays and electrons in distant regions of our own Galaxy or even in other galaxies and directly probe and image the cosmic accelerators responsible for these particle populations. High-energy γ -rays can also be produced by decays of heavy particles such as hypothetical dark matter particles or cosmic

¹ <http://www.cta-observatory.org/>

strings, both of which might be relics of the Big Bang. For this reason γ -rays provide a window on the discovery of the nature and constituents of dark matter.

The CTA will offer worldwide unique opportunities to users with varied scientific interests and support a growing number of young scientists working in the evolving field of gamma-ray astronomy.

The CTA will, for the first time in this field, provide open access via targeted observation proposals and generate large amounts of public data, accessible using Virtual Observatory tools. The CTA aims to become a cornerstone in a networked multi-wavelength, multi-messenger exploration of the high-energy non-thermal universe. During the on going preparatory phase of the project, CTA Monte Carlo (MC) simulations campaigns are distributed on the grid via the EGI CTA Virtual Organisation.

One of the e-Infrastructure challenges under consideration will consist in defining an adequate computing model for CTA, either centralised or distributed, and in exploiting the big data technology with an increased level of complexity added by remote 'wild' experiment sites.

Directly related to the EGI CTA Virtual Team currently in progress, the second challenge will be to enable data dissemination to the scientific community through Science Gateway and Single Sign-On solutions.

Scientific case

The aims of the CTA can be roughly grouped into three main themes, the key science drivers:

- 1) understanding the origin of cosmic rays and their role in the Universe;
- 2) understanding the nature and variety of particle acceleration around black holes;
- 3) searching for the ultimate nature of matter and physics beyond the Standard Model.

Theme 1 comprises the study of the physics of galactic particle accelerators, such as pulsars and pulsar wind nebulae, supernova remnants, and γ -ray binaries. It deals with the impact of the accelerated particles on their environment (via the emission from particle interactions with the interstellar medium and radiation fields), and the cumulative effects seen at various scales, from massive star forming regions to starburst galaxies.

Theme 2 concerns particle acceleration near super-massive black holes. Objects of interest include blazars, radio galaxies and other classes of Active Galactic Nuclei that can potentially be studied in high-energy γ -rays. The fact that CTA will be able to detect a large number of these objects enables the population studies that will be a major step forward in this area. Extragalactic background light (EBL), Galaxy clusters and Gamma Ray Burst (GRB) studies are also connected to this field.

Finally, **theme 3** covers what can be called 'new physics', with searches for dark matter through possible annihilation signatures, tests of Lorentz invariance, and any other observational signatures that may challenge our current understanding of fundamental physics.

The CTA will be able to generate significant advances in all these areas.

State of the art

The latest generation of ground-based γ -ray instruments (H.E.S.S., MAGIC, VERITAS, Cangaroo III and MILAGRO) allow the imaging, photometry and spectroscopy of sources of high-energy radiation and have ensured that VHE (Very High-Energy) γ -ray studies have grown to become a genuine branch of astronomy. The number of known sources of VHE γ -rays exceeds 100, and source types include supernovae, pulsar wind nebulae, binary systems, stellar winds, various types of active galaxies and unidentified sources without obvious counterpart. H.E.S.S. has conducted a highly successful survey of the Milky Way covering about 600 square degrees, which resulted in the detection of tens of new sources. However, a survey of the full visible sky would require at least a decade of observations, which is not feasible.

Due to their small flux, instruments for detection of high-energy γ -rays (above some 10 GeV) require a large effective detection area, eliminating space-based instruments which directly detect the incident γ -rays. Ground-based instruments allow much larger detection areas. They measure the particle cascade induced when a γ -ray is absorbed in the atmosphere, either by using arrays of particle detectors to record the cascade particles which reach the ground (or mountain altitudes), or by using Cherenkov telescopes to image the Cherenkov light emitted by secondary electrons and positrons in the cascade.

Compared to Cherenkov telescopes, air shower arrays (such as MILAGRO, AS-gamma or ARGO) have the advantage of a large duty cycle – they can observe during the daytime – and of a large solid angle coverage. However, their current sensitivity is such that they can only detect sources with a flux around the level of the flux from the Crab Nebula, the strongest known steady source of VHE γ -rays. Results from air shower arrays demonstrate that there are relatively few sources emitting at this level. The recent rapid evolution of VHE γ -ray astronomy was therefore primarily driven by Cherenkov instruments, which reach sensitivities of 1% of the Crab flux for typical observing times of 25 hours, and which provide significantly better angular resolution. While there are proposals for better air shower arrays with improved sensitivity (e.g. the HAWC project), which will certainly offer valuable complementary information, such approaches will not be able to compete in sensitivity with next-generation Cherenkov telescopes.

The major current and historic Cherenkov instruments consist of up to five Cherenkov telescopes (H.E.S.S.). They reach sensitivities of about 1% of the flux of the Crab Nebula at energies in the 100 GeV to 1 TeV range. Sensitivity degrades towards lower energies, due to threshold effects, and towards higher energies, due to the limited detection area. A typical angular resolution is 0.1° or slightly better for single γ -rays. Sufficiently intense sources can be located with a precision of 10-20".

All these instruments are operated by the research groups who built them, with very limited access for external observers and no provision for open data access. Such a mode is appropriate for current instruments, which detect a relatively limited number of sources, and where the analysis and interpretation can be handled by the manpower and experience accumulated in these collaborations. However, a different approach is called for in next-generation instruments, with their expected ten-fold increase in the number of detectable objects.

Going beyond the state of the art

The CTA, which consists of two arrays of Cherenkov telescopes, will advance the state of the art in astronomy at the highest energies of the electromagnetic spectrum in a number of decisive areas, all of which are unprecedented in this field:

- **European and international integration:** CTA will for the first time bring together and combine the experience of virtually all groups world-wide working with atmospheric Cherenkov telescopes.
- **Performance of the instrument:** CTA aims to provide full-sky view, from a southern and a northern site, with unprecedented sensitivity, spectral coverage, angular and timing resolution, combined with a high degree of flexibility of operation.
- **Operation as an open observatory:** CTA will, for the first time in this field, be operated as a true observatory, open to the entire astrophysics (and particle physics) community, and providing support for easy access and analysis of data which will be made publicly available and accessible through Virtual Observatory tools.
- **Technical implementation, operation, and data access:** The goals of CTA imply significant advances in terms of efficiency of construction and installation, in terms of the reliability of the telescopes, and in terms of data preparation and dissemination.

Science performance goals for CTA include in particular:

- **Sensitivity:** CTA will be about 10 times more sensitive than any existing instrument. It will therefore for the first time allow detection and in-depth study of large samples of known source types, will explore a wide range of classes of suspected γ -ray emitters beyond the sensitivity of current instruments, and will be sensitive to new phenomena.
- **Energy range:** Wide-band coverage of the electromagnetic spectrum is crucial for understanding the physical processes in sources of high-energy radiation. CTA is aiming to cover, with a single facility, three to four orders of magnitude in energy range. Combined with the Fermi γ -ray observatory in orbit, an unprecedented seamless coverage of more than seven orders of magnitude in energy can be achieved.
- **Angular resolution:** Current instruments are able to resolve extended sources, but they cannot probe the fine structures visible in other wavebands. Selecting a subset of γ -ray induced cascades detected simultaneously by many of its telescopes, CTA can reach angular resolutions in the arc-minute range, a factor of five better than the typical values for current instruments.
- **Temporal resolution:** With its large detection area, CTA will resolve flaring and time-variable emission on sub-minute time scales, which are currently not accessible. In γ -ray emission from active galaxies, variability time scales probe the size of the emitting region.
- **Flexibility:** Consisting of a large number of individual telescopes, CTA can be operated in a wide range of configurations, allowing on the one hand the in-depth study of individual objects with unprecedented sensitivity, and on the other hand the simultaneous monitoring of tens of potentially flaring objects, and any combination in between.
- **Survey capability:** A consequence of this flexibility is the dramatically enhanced survey capability of CTA. Groups of telescopes can point at adjacent fields in the sky, with their fields of view overlapping, providing an increase of sky area surveyed per unit time by an order of magnitude, and for the first time enabling a full-sky survey at high sensitivity.

- **Number of sources:** Extrapolating from the intensity distribution of known sources, CTA is expected to enlarge the catalogue of objects detected from currently several tens of objects to about 1,000 objects.
- **Global coverage and integration:** CTA aims to provide full sky coverage from multiple observatory sites, using transparent access and identical tools to extract and analyse data. The feasibility of the performance goals is borne out by detailed simulations of arrays of telescopes, using currently available technology. The implementation of CTA does require significant advances in the engineering, construction and operation of the array, and the data access.

e-Infrastructure

Current e-Infrastructure activity

The CTA computing needs are large and expected to grow during the construction and commissioning phase (2016-2020) and the lifetime of the observatory (20 years or more), before they stabilise in the 10 following years.

The CTA Computing Model must be coherent with the aim of an Observatory: *to handle a large amount of data generated by the telescopes in remote sites and to provide modern and efficient data access at any level in time, efficiently and at a worldwide scale.*

The Computing Model should result from a series of technical solutions suitable for the data pipeline and data reduction constraints as well as for the major user requirements for CTA data management.

Two main criteria need also to be conjugated in the Computing Model:

- 1) Effective consideration of the existent ICT infrastructures and first class computer centres which already implement the computing model of major astroparticle physics experiments;
- 2) Envisaging the development and application of new solutions coherent with the fast evolution of new emerging requirements in e-Science, with the aim to provide any single user with an efficient and modern scientific analysis environment during the next 20 years.

The above criteria then have to drive the Computing Model implementation in a way that it could be than also easy to re-adapt the model to different infrastructures beyond the observatory life cycle.

The CTA observatory is expected to produce a main data stream for permanent storage of the order of 1-10 GB/s for about 1,300 hours of observation per year, thus producing a total data volume in the range of 2-25 PB per year. Data processing will require a fair amount of CPU time, with about 800 CPU days needed to calibrate and reconstruct 1 hour of raw data. The high data rate of CTA together with the large computing power requirements for Monte Carlo simulations need dedicated important computer resources. Another serious functional constraint to the Computing Model is a required Archive System, i.e. a combination of the hardware and associated services for data I/O (in the usual astronomical meaning), for permanent storage of data products that must provide the scientific users with an efficient and organised access to data.

Whereas all data processing in current astroparticle computing models takes place in a single computing centre, a distributed approach could be more adequate for the scale of CTA. The high computing needs for Monte Carlo simulations during the Preparatory Phase of CTA, are fulfilled by exploiting the EGI precisely for large simulation productions (more than 10^{11} proton, gamma, and electron induced showers are required for MC studies). A feasibility study of applications of grid solutions for CTA started within a dedicated CTA Computing Grid (CTACG) project. The CTA Virtual Organisation (CTA VO) was created in 2008 with the support of the IN2P3-LAPP computing centre and in cooperation with the CC-IN2P3 computing centre. Today the CTA VO is supported by 18 grid sites spread in seven countries, with resources of the order of some thousands of available logical CPUs and more than 600 TB of storage.

Typical MC productions consist of about 150,000 jobs producing about 600 TB of data. In order to handle such massive productions, in 2011 the consortium started the evaluation of the DIRAC (Distributed Infrastructure with Remote Agent Control) system (Arrabito et al., 2012), which is a general framework to manage the distributed activities of a user community. DIRAC has the advantage to integrate heterogeneous resources according to the evolution of the CTA Computing Model. A first deployment of a DIRAC instance for evaluation purposes was followed by the development of a dedicated CTADIRAC extension to handle the CTA specific workflows. This setup has been successfully exploited during the CTA MC campaign in 2013 enabling, in particular, to achieve a regime of 4,000-5,000 concurrent jobs with a peak of 8,000, producing about 600 TB. Further developments consist of an interface to the DIRAC File Catalogue (DFC) in order to save the file replica information as well as the Meta Data information specific to the CTA MC pipeline. Users are able to query the DFC through a dedicated web interface. A user activity has also started in 2012 to analyse MC data or to produce specific simulations. This activity represents a cumulative number of about 200,000 jobs run by a tenth of users.

Future e-Infrastructure challenges

Today the CTA consortium is in the specification phase, where the main goal is to estimate global data volumes and processing needs depending on the data acquisition rate and science requirements. The main challenge of the CTA Computing Model will be the expected volumes of data, starting at 2-25 PB per year, to be processed and managed for 20-30 years. It is not only a question of big data storage, but also data transfer from remote 'wild' telescopes sites to the science data centres.

The next phase will be an evaluation of different scenarios and associated IT architecture solutions, taking into account strengths and weaknesses related to their performance, cost, robustness, scalability, manageability and planned durability. The main architectures to be evaluated are centralised, distributed or any combination of them, with cloud and grid resources being one of these specific solutions. The usages will be data pipeline production tasks and on-demand simulation, reconstruction or analysis tasks.

In parallel with the need to handle data dissemination to the scientific community, another of the CTA challenges is to provide users with a common User Interface: a CTA Science Gateway with complex management of authentication and authorisation mechanisms.

By definition, a Science Gateway is a community-development set of tools, applications, and data that is integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customised to meet the needs of a specific community. The CTA Science Gateway will provide access to resources and services from a distributed computing infrastructure. The resources could be grid computing and storage resources, public and private cloud services, local personal computer resource, user-specific laboratory/institute storage and computing resources, together with CTA observatory storage and computing resources.

For the CTA community, the personalisation of human-computer interface and management of complex access rights is one of the main requirements of this product and explains why the on going EGI-CTA Virtual Team dedicated action targets two sub-products: Science Gateway and Single Sign On (SSO) solutions.

SSO is a property of access control of multiple related, but independent software systems. With this property, a user logs in once and gains access to all systems without being prompted to log in again at each of them. Conversely, Single Sign Off is terminates access to multiple software systems with a single action. As different applications and resources support different authentication mechanisms, the SSO system has to internally translate to and store different credentials compared to what is used for initial authentication.

Future plans

The future milestones of CTA are the end of the preparatory phase by 2014 and choice of sites at the beginning of next year. The planned four years array construction phase should start in 2016.

For ICT infrastructure, a dedicated CTA Data Management project is working, among other topics, on the proposal of the computing model implementation plan for CTA taking into account the current assumptions on data volumes and processing needs, and including a comparative study about available computing models: centralised, distributed grid, cloud, combined resources, and so on. A consultation with GÉANT is in place and will help to evaluate the ability to put together a proper network bandwidth from the array sites to the data centre.

For the Science Gateway point of view, the current EGI CTA Virtual Team dedicated to Science Gateway and Single Sign On will produce a document with suggestions and recommendations about construction of Science Gateway and Authentication system before the summer. This will be the first step to be able to start from July 2013 a production of prototypes for Science Gateway and Authentication system.

References & Bibliography

- Actis, M. et al. (2011) *Experimental Astronomy* 32:193
- Arrabito, L. et al. (2012) *Journal of Physics: Conference Series* 396 032007
- CTA Consortium (2013) CTA special issue *Astroparticle Physics*: 43: 1-35

DARIAH: Digital Research Infrastructure for the Arts and Humanities

Tobias Blanke✉

Göttingen Centre for Digital Humanities, Germany

22 August 2013

✉ tobias.blanke@dariah.eu

Scientific case

The scientific challenge

Digital research methods have recently started to enter the mainstream of humanities, arts and social sciences research. Digital humanities have existed for years as a specialised field but the recent growth in the number of centres and research projects associated with digital methods in arts and humanities (A+H) and social sciences indicate that we are at a fundamental shift. There have been significant national investments. For instance, in the first decade of the 21st century the UK Arts and Humanities Research Council and its predecessor invested about half of its research funding into projects with some kind of digital output. Most of these resources are online and open access.

Europe has played a key role in developing the state-of-the-art of digital humanities. Not only can digital humanities trace its origin to an Italian collaboration with IBM, but also nowadays there exists a vast array of collaborations in the digital humanities across Europe in the form of spontaneously funded research networks and associations. What is lacking, however, is an infrastructure that would ensure that the state-of-the-art of these collaborations is preserved and integrated, and that common best practices and methodological and technological standards are followed. The Digital Research Infrastructure for Arts and Humanities (DARIAH) aims to be this infrastructure.

State of the art

Humanities and arts have used computers for a long time. Recently, there has also been more effort to work towards more consistent cyber-infrastructure and away from *ad hoc* solutions to deliver more systematic investigations. Typical advanced digital arts and humanities activities include, for example, the use of network analysis for mapping out the biographies of historical collectives, the use of text mining for semantically enriching historical collections, or computer-aided measurements in field-based research. All these activities are very seldom joined-up, as simply the knowledge and funding is missing. The biggest current field of research is to define the new digital methodologies to meet the requirements of humanities data that is particularly fuzzy and inconsistent, as it is not automatically produced, but is the result of human effort.

Compendium of e-Infrastructure requirements for the digital ERA

Going beyond the state of the art

The state-of-the-art around DARIAH is quickly evolving. There are two major trends right now. The first one is the attempt to evolve the *ad hoc* experimentations with digital methods into systematic investigations. To this end, there are many activities around digital research methods and more and more articles of the kind appearing in traditional journals of arts and humanities. The community has begun to critically reflect on the meaning of digital methods and how digital arts and humanities projects can move beyond a quite narrow community.

The first scientific grand challenge for digital arts and humanities is to develop the reflection on digital methods towards the mainstream of arts and humanities and to embed these in everyday research activities. One example would be to further develop archive-based research in history. Our work (Blanke et al., 2010) investigated how (digital) archival content can be delivered to humanities researchers more effectively, independently of the location and implementation of that content, and with special facilities provided for customising the retrieval, management and manipulation of the content. This way retrieval of information can happen in real time, for which traditional finding aids are to be complemented by more sophisticated retrieval mechanisms, including the ability to create relevance indexes on unstructured resources, as well as the ability to combine resources in new ways. This is just one example of how digital methods can be embedded in the everyday. Others include, for instance, the digitisation of annotations (potentially distributed).

The second grand challenge for the near future is linked to what was once called the data deluge and now goes under the name of big data. The arts and humanities have seen an exponential growth in digital research material, especially in the last decade, as a result of new born-digital material or large digitisation efforts in the EU and elsewhere. DARIAH will be the decisive European step towards exploiting these large digital resources for A+H, to help fulfil the predication of Dan Atkins, the former director of the US cyber-infrastructure programme, on the future of A+H research:

“Arts and humanities are poised to achieve large benefit from e-Science methods and infrastructure as the human record becomes increasingly digitised and multimedia, and [...] with the ability now to compute across enormous collections [...]” (Dan Atkins et al., RCUK e-Science Review)

In order to achieve this, we need to rework how integration has been done. Computers do not like this kind of heterogeneous and incomplete material created by humans and typically used in the humanities and arts. We would move beyond the state-of-the-art if we finally managed to achieve the integration of humanities research material on the grand scale.

e-Infrastructure

Current e-Infrastructure activity

DARIAH partners are recognised contributors to national excellence in A+H research and successful collaborators in European research projects. The challenge for DARIAH will then be to join up national/local knowledge in a sustainable, collaborative and lasting ecosystem.

DARIAH services are based on the existing collaborations of DARIAH in digital arts and humanities, which we plan to scale to a European level.

At the moment, our service interests focus on maximising the impact and collaboration across these kinds of national projects, but over time we see a DARIAH data and service market developing that will be enriched by tools and services from the digital arts and humanities community. DARIAH will guarantee this service market place and the services necessary to join up national services such as Persistent Identifiers (PIDs) and a federated search and authentication environment.

DARIAH attempts to build services around communities, which can then be exchanged between communities in a virtual social marketplace that connects community workspaces with trusted DARIAH repositories of research data and services.




Fig.1 - DARIAH-architecture

Figure 1 describes the core architecture for our organisation of all DARIAH services. The core layer includes light-weight services that serve to sustain the DARIAH infrastructure and establish coherent operation across the open DARIAH environment. This core layer will in the medium- to long-term include a wide range of technical services. Immediately, we plan to expand the existing Persistent Identifier systems at DARIAH partners and develop an integrated community-based Authentication and Authorisation Infrastructure (AAI) for Single Sign On, so all DARIAH partners can benefit. The corresponding infrastructures are already in place and will reuse existing systems. These two services are essential for enabling interoperability across the heterogeneous data sources and decentralised services in the DARIAH ecosystem. The DARIAH PID service links various system components with relevant policies.

The intermediate layer in Figure 1, the infrastructure service environment, has services that will be supported by DARIAH but not guaranteed. The national projects will collaborate with outside initiatives and researchers to build them. For example, the DARIAH-DE supported Authority Mediation Service (AMS) deploys a network of reference data services, including library authority lists (e.g. Virtual International Authority File) as well as various dictionaries, Compendium of e-Infrastructure requirements for the digital ERA

thesauri and gazetteers. As building these resources falls under digital scholarship, many DARIAH research partners are directly involved in setting them up.

The user-facing framework (UFF) exemplifies another core principle of DARIAH. For the UFF, we document how to interact with the guaranteed core services but we also accommodate a collection of end-user tools contributed by research projects or third parties. Beyond mere documentation, tools and services ideally comply with the DARIAH service framework to foster interoperability with other DARIAH components.

Our overall architecture and compliance framework is detailed in (Blanke et al., 2011) and online². Here, also more details about existing technical work as well as plans for the marketplace and the registries can be found.

Future e-Infrastructure challenges

In the future we need to work on a more decentralised approach that will in particular develop the social components of the e-Infrastructure. As an integration project, DARIAH would like to particularly advance the local distributed knowledge at least in the immediate future. For our researchers, it would therefore be useful to have transnational access to virtual machines, data management services, persistent storage and instruments to investigate objects. This is next to the usual candidates of providing stable PIDs for resources and distributed authentication and authorisation. We already have these capacities in place in the various partner countries, but sharing these has proven to be challenging. These are the essential needs. Desirable for distinct research activities such as the analysis of manuscripts is the easy access to high-performance computation infrastructure for the occasional burst in processing needs. Furthermore, a transparent data infrastructure that allows for the combination of many small-scale but highly interrelated resources and is at the same time persistent across countries would be a great advantage. We do not have problems where one size fits it all, especially when it comes to data. A polyglot persistent infrastructure that provides seamless access from localised web data storage all the way to long-term large digital archives would thus be a great plus. This will be especially important once we work more and more with born digital material. We have begun to experiment with these kinds of alternative data infrastructures (Blanke and Kristel, forthcoming) but are a long way off from exploiting them fully.

The second future e-Infrastructure challenge is the social one, where we find new ways of collaborating across disciplines where data is part of the research process. In particular, the access and understanding to digital methods needs to be further improved and embedded in local activities. We know that a technical/scientific infrastructure can only work if it is at the same time a social infrastructure, which includes a wide range of knowledge services. We are struggling in particular with linking research activities with digital methods. Furthermore, as DARIAH is designed for the exchange of knowledge and services in a dedicated virtual social marketplace, we need open APIs to expose reusable services, as well as composition and aggregation facilities to work with these services. Finally, and maybe most importantly, we need

2

http://www.dariah.eu/index.php?option=com_docman&task=doc_details&gid=477&Itemid=200
Compendium of e-Infrastructure requirements for the digital ERA

to promote applications based on these services for several use cases and digital methods that showcase excellent research.

Future plans

There is currently a major review on-going within the DARIAH community on our future relationship with generic middleware work in the various national programmes. The review is carried out within the context of the German DARIAH-DE project and its second round of funding with the aim to set up a generic e-Infrastructure Unit based on simple infrastructure needs. DARIAH-DE can in particular rely on the work done in the associated TextGrid consortium (Hedges et al., 2013). This reflects the shift from an originally middleware-focussed (grid-based) e-Infrastructure (connecting to Globus middleware via GAT) towards a more computing centre-specific distributed storage solution. This approach will be replicated across the DARIAH communities. The set-up of grid nodes was too complicated for some of the non-computing centre infrastructure partners and partly beyond their needs. The PKI based security has too high an overhead for the end-user (user certificates handling and reissue problems).

With these lessons learned, DARIAH-DE favours simpler replication mechanisms (via IRODS) for distributed storage and user-friendlier Authentication and Authorisation mechanisms based on SAML. Within the D-Grid Initiative project IVOM, the German DAASI partner built a solution for TextGrid that integrated SAML based AAI with PKI based grid infrastructure.

There is secondly the general time plan for DARIAH ESFRI collaboration. We are about to become an ERIC with at least 9 initial members.

References & Bibliography

- Anderson, S., T. Blanke, S. Dunn (2010) Methodological commons: arts and humanities e-Science fundamentals. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1925): 3779-3796
- Anderson, S., T. Blanke (2012) Taking the Long View: From e-Science Humanities to Humanities Digital Ecosystems. *Historical Social Research-Historische Sozialforschung*, 37(3): 147-164
- Aschenbrenner, A., C. Fritze, W. Pempe (2013). Data-Driven Research in the Humanities — the DARIAH Research Infrastructure. In: *The Data Bonanza: Improving Knowledge Discovery in Science, Engineering, and Business*. Hoboken, NJ: Wiley.
- Benardou, A., P. Constantopoulos, C. Dallas, D. Gavrilis (2010) Understanding the information requirements of arts and humanities scholarship. *International Journal of Digital Curation*, 5(1): 18-33
- Blanke, T., L. Candela, M. Hedges, M. Priddy, F. Simeoni (2010) Deploying general-purpose virtual research environments for humanities research. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1925): 3813-3828
- Blanke, T., M. Bryant, M. Hedges, A. Aschenbrenner, M. Priddy (2011) Preparing DARIAH. In: *IEEE 7th International Conference on E-Science* pp:158-165. IEEE.

- Blanke, T. and C. Kristel (Forthcoming) Integrating Holocaust Research. *International Journal of Arts and Humanities Computing*.
- Constantopoulos, P., C. Dallas, P. Doorn, D. Gavrilis, A. Gros, G. Stylianou (2008) Preparing DARIAH. In: *Proceedings of the International Conference on Virtual Systems and MultiMedia, VSMM08*
- Hedges, M., H. Neuroth, K.M. Smith, T. Blanke, L. Romary, M. Küster, M. Illingworth (2013) TextGrid, TEXTvire, and DARIAH: Sustainability of Infrastructures for Textual Scholarship. *Journal of the Text Encoding Initiative* (5).
- Henrich, A., T. Gradl (2013) DARIAH (-DE): Digital Research Infrastructure for the Arts and Humanities-Concepts and Perspectives. *International Journal of Humanities and Arts Computing*, 7(supplement): 47-58.
- Romary, L. (2012). Data management in the humanities. *ERCIM News*.
- Tonne, D., R. Stotzka, T. Jejkal, V. Hartmann, H. Pasic, A. Rapp, B.S. Bribian (2012) A Federated Data Zone for the Arts and Humanities. In: *20th Euromicro International Conference on Parallel, Distributed and Network-Based Processing* pp:198-205. IEEE.
- Tonne, D., J. Rybicki, S.E. Funk, P. Gietz (2013) Access to the DARIAH Bit Preservation Service for Humanities Research Data. In: *21st Euromicro International Conference on Parallel, Distributed and Network-Based Processing* pp:9-15. IEEE.

European Multidisciplinary Seafloor and water-column Observatory (EMSO)

Giuditta Marinaro✉, Gabriele Giovanetti✉, Daniele Cesini, Paolo Favali, Laura Beranzoli, Paola Materia

on behalf of the EMSO Consortium

21 August 2013

✉ giuditta.marinaro@ingv.it, gabriele.giovanetti@ingv.it

Overview

The European Multidisciplinary Seafloor and water-column Observatory³ (EMSO) is a large-scale European Research Infrastructure (RI) and part of the ESFRI roadmap. The EMSO is composed of fixed-point, seafloor and water-column observatories. Each observatory's objective is the real-time and near-real-time, long-term monitoring of environmental processes related to the interaction between the geosphere, biosphere, and hydrosphere.

The EMSO nodes are geographically distributed throughout key-sites in European waters, spanning the Arctic, the Atlantic, the Mediterranean and Black Seas (Fig.1). The EMSO Preparatory Phase (EMSO-PP) project endorsed the European Research Infrastructure Consortium (ERIC) to create a legal organisation to operate ocean observatory infrastructures in Europe.

EMSO concluded its Preparatory Phase in September 2012 and is setting up the EMSO-ERIC to establish and manage the research infrastructure. A Memorandum of Understanding (MoU) among corresponding funding agencies of the EMSO partners has ten signatories (France, Germany, Greece, Italy, Ireland, Netherlands, Portugal, Romania, Spain, United Kingdom). The MoU remains valid for three years or until the establishment of the EMSO-ERIC.

³ <http://www.emso-eu.org>



Fig.1 - The EMSO nodes

Scientific case

Changes relating to resource availability, climate change, habitat destruction, and geo-hazards have increased society's needs for an improved understanding of the factors behind these changes and their impact. Representing the largest habitat on the Planet, the open oceans and deep seafloor play a crucial role on climactic phenomena and many of the processes occurring in the ocean affect society directly. In light of this, consecutive EU Frameworks have invested in the establishment of a European observatory network that provides *in situ* measurements of key ocean variables on a continuous and long-term basis. These measurements are essential to tackle questions at the scales necessary to understand climate change and its impacts and to improve early warnings for geo-hazards.

EMSO-ERIC will collect high-resolution data from sub-seafloor, seafloor, water column and surface, and transmit the data in real- or near-real-time to shore via satellites or cable connection and assist researchers in answering vital questions in marine and earth sciences. By establishing a set of common standards and best practices, the network will play a key role in integrating equipment and procedures across European seas.

The EMSO nodes are distributed from the Arctic Ocean to the Mid-Atlantic Ridge, the Mediterranean to the Black Sea (Fig.1). The proposed sites span the major biogeochemical provinces identified in European waters and reflect a wide range of habitats, including abyssal plains, open slopes, seamounts, canyons, ridges, faults, fluid seeps, hydrothermal vents, gas hydrates, mud volcanoes, deep-sea corals, carbonate mounds, and geo-hazard zones. EMSO will either represent a primary source of data or they will be used in a supporting manner. The nodes include cabled and stand-alone sites with moorings and benthic instruments, while communicating in real time or in delayed mode, and being serviced through annual maintenance cruises. The EMSO infrastructure also includes at the moment four submarine testing laboratories in shallow waters: Koljö Fjord in Sweden, SmartBay in Ireland, Molènes in

France and OBSEA in Spain. These laboratories are cabled and easily accessible at any season and at low cost for submarine equipment sensor tests deployed at seafloor.

Table 1 describes the actual situation in some of the EMSO sites.

Tab. 1. Description of the activities at some EMSO nodes

<i>Sites</i>	<i>Current Infrastructures and on-going research</i>
Arctic	<p><i>HAUSGARTEN</i></p> <ul style="list-style-type: none"> • Deep-sea observatory operated by AWI since 1999 • 15 permanent sampling sites • Moorings and long-term lander systems since 2000 <p><u>Research</u>: Studying the possible link between climate change-induced gas hydrate decomposition and methane release. One hydrographic section has been monitored three times over the last three years, including CH4 analyses. Atmospheric measurements are carried out with CRDS systems and air sampling. Hydroacoustic surveys with single and multibeam systems have been repeated annually.</p>
Celtic/Porcupine	<p><i>Porcupine Abyssal Plain (PAP)</i></p> <ul style="list-style-type: none"> • Active time-series station since 1980s • Some data transmitted in near real-time by satellite <p><u>Research</u>: biogeochemical and physical measurements from upper 1000 m; seafloor video and photography; biogeochemical flux measurement systems.</p>
Azores Islands	<p><i>Lucky Strike hydrothermal vent field</i></p> <ul style="list-style-type: none"> • Site active for more than a decade <p><u>Research</u>: geophysical movements of Earth (seismicity and vertical deformation); water, heat and mineral flow through vent system; behaviour of physical and chemical elements in vent fluid; variations in biogeochemistry and the ecological hotspots in vicinity of vents.</p>
Ligurian Sea	<p><i>Ligurian Sea</i></p> <ul style="list-style-type: none"> • DYFAMED - DYnamics of Atmospheric Fluxes in the MEDiterranean Sea) • Var canyon monitoring • Nice slope monitoring of geo-hazard • ANTARES (Astronomy with a Neutrino Telescope and Abyss environmental RESearch) – Earth-Sea science extension of astrophysics underwater telescope characterised by coastal upwelling, particle plumes, nutrient benthic exchange, bottom boundary layer processes, seismic monitoring. <p><u>Research</u>: sub-sea geophysics; slope stability; biogeochemical fluxes and marine ecology</p>
Western Ionian Sea	<p><i>NEMO-SN1</i></p> <ul style="list-style-type: none"> • Cabled to laboratory in harbour of Catania by electro-optical cable • Operating in real-time since 2005

	<ul style="list-style-type: none"> • Integrated with land-based networks by transmitting real-time data to National Seismological Service Centre in Rome • Test site for realisation of the underwater neutrino telescope <p><u>Research:</u> Geo-hazards (earthquakes, tsunamis, volcanoes, gravitative instabilities), climate change, bioacoustics and ambient noise.</p>
Hellenic Arc	<p><i>Series of four networks in the Hellenic study area</i></p> <ul style="list-style-type: none"> • Cabled system NESTOR • Stand-alone Poseidon Pylos • Poseidon E1-M3A • Proposed drilled observatory BUTT-1 <p><u>Research:</u> Geo-hazards (earthquakes, tsunamis, gravitative instabilities), climate change, bioacoustics and ambient noise, biogeochemical fluxes, benthic-pelagic interactions; benthic respiration; biogeochemical fluxes; photography-based ecology; seabed methane fluxes; oil and gas industry activities.</p>
Canary Islands	<p><i>PLOCAN</i></p> <ul style="list-style-type: none"> • ESTOC/PLOCAN - European station for time-series on the ocean • Operating since 1994 <p><u>Research:</u> ocean physics and biogeochemistry; decadal record of ocean acidification.</p>

The scientific challenge

The processes that occur in the ocean have a direct impact on society. Therefore, it is crucial to improve our understanding of how these processes operate and interact. To encompass the breadth of these diverse processes, sustained and integrated observations are required that appreciate the interconnectivity of atmospheric, surface ocean, biological pump, deep-sea, and solid-Earth dynamics (Fig.2) and that address these issues:

- Natural and anthropogenic change
- Interactions between ecosystem services, biodiversity, biogeochemistry, physics, and climate
- Impacts of exploration and extraction of energy, minerals, and living resources
- Geo-hazard early warning capability for earthquakes, tsunamis, gas-hydrate release, and slope instability and failure
- Connecting scientific outcomes to stakeholders and policy makers




Fig.2. Major processes in the marine environment (redrawn from Ruhl et al., 2011)

Long-term, continuous data sets from a variety of fields are necessary to build a comprehensive image of earth-ocean systems. These include:

Geoscience

Geoscience covers a range of processes, including gas-hydrate stability, submarine landslides and fluid flow along the seabed, seismic activity and geo-hazard early warning. Seismic activity and seafloor slippages, in particular, can have direct effects on human activities, such as causing damage to offshore industry infrastructure and catastrophic impact on the population through earthquakes and tsunamis. In order to perform robust forecasting, measurements need to be made out continuously over sufficiently long periods of time to be able to differentiate between episodic events and trends, or variations over shorter periods of time.

Physical Oceanography

One of the most urgent problems society faces is the effect of global warming on the marine environment (IPCC, 2007a and 2007b). For example, a rise in sea temperatures will lead to sea ice melting (IPCC, 2007a) and increasing stratification (Levin et al., 2001; Marzeion et al., 2010), which can render large areas of the ocean anoxic and uninhabitable. Therefore, detailed knowledge about ocean transport, and wind-driven and deep-ocean circulation is vital to be able to assess the how changes in the ocean affect global climate systems.

Biogeochemistry

One of the effects of increasing atmospheric carbon dioxide levels (which also leads to global warming) is an increased uptake of carbon dioxide by oceans. Up to one third of anthropogenic carbon dioxide produced today is absorbed by oceans through two processes: solubility and biological pumps. This leads to the lowering of the pH of seawater, which in turn results in the

acidification of ocean water (Feely et al., 2004; Orr et al., 2005; Raven et al., 2005; Fabry et al., 2008; Feely et al., 2008). An increasingly acidic ocean impacts the ability of marine organisms to calcify, such as calcifying primary producers, like molluscs, and corals (Orr et al., 2005; Hoegh-Guldberg et al., 2007; Tyrrel, 2008). At the other end of the spectrum, there is a limit to how much carbon dioxide the ocean can absorb. Once this threshold is reached the declining uptake of anthropogenic carbon dioxide could increase the proportion that accumulates in the atmosphere –accelerating the global warming effect.

Marine Ecology

An increased understanding of how ecosystems function is crucial to evaluating the sensitivity of marine ecosystems to anthropogenic changes. This represents one of the primary challenges in marine science over the next decades (Sutherland et al., 2006). Marine ecosystem functions maintain key services such as primary production, climate regulation, carbon sequestration and storage, and live resources, including fisheries. Compared to terrestrial systems, changes in the ecology of the oceans as a result of global warming and carbon dioxide accumulation could have numerous repercussions (Richardson, 2008).

So far, only a limited number of data sets are available allowing for the observation of climatically-driven changes in marine ecosystems by discerning between inter-annual and inter-decadal variations and secular change (Rosenzweig et al., 2008; Glover et al., 2010). The pace and scale of anthropogenic changes occurring in the oceans, such as overfishing and pollution, and the impact of these changes on marine biodiversity and ecosystems, are cause for serious concern. Ocean observatory research efforts designed to help scientists better understand marine biodiversity will provide the knowledge necessary to deploy adaptive management processes by linking variations in biodiversity, its function, and the ecological and environmental forces that drive change in a comprehensive way.

State of the art

The EMSO data infrastructure was designed to use the existing distributed network of data infrastructures in Europe and use the INSPIRE and GEOSS data sharing principles. A number of standards have been set forth that will allow for state-of-the-art transmission and archiving of data, with metadata recording and interoperability that allow for more straightforward data use and transmission. These standards include the Open Geospatial Consortium (OGC) Sensor Web Enablement (SWE) suite of standards, namely the OGC standards SensorML, Sensor Registry, Catalogue Service for Web (CS-W), Sensor Observation Service (SOS) and Observations and Measurements (O&M). OGC SensorML is an eXtensible Markup Language (XML) for describing sensor systems and processes. Following on progress from EuroSITES and others, a SensorML profile is being created that can be stored in a so-called Sensor Registry that will act as a catalogue of each EMSO sensor. This dynamic framework can accommodate the diverse array of data and formats used in EMSO, including the addition of delayed mode data.

EMSO data collected in experiments at 11 regional sites is locally stored and organised in catalogues or relational database and run by the institutions involved. Some of the EMSO observatories' data from distributed sites are harvested and archived over the long-term in three data archives: Ifremer (EUROSITES), UniHB (PANGAEA) and INGV (MOIST). A central

archive hosting web-service access to all the databases is planned for the near future. We consider a group of functions provided by the three data archives that support data quality control and preservation as a data curation sub-system.

EMSO provides advanced technology in data publication and citation through the PANGAEA system and offers capabilities for data access, standardisation/harmonisation and visualisation via MOIST data infrastructure. Presently, three regional data sites are integrated in MOIST, and one regional site is integrated in PANGAEA which offers data from several related or preparatory studies for other EMSO sites. In addition, Ifremer offers access to data from all EUROSITES sites which are shared with EMSO.

Going beyond the state of the art

The ocean sciences are currently undergoing a transition thanks to new technical capabilities for better spatiotemporal resolution of observations. The goal is to have dedicated infrastructures (reference stations) deployed at key locations to improve the understanding of processes that are closely linked, such as biogeochemical processes. This would also be an important step to improve ocean models and improve forecasting capabilities. On a small scale this has been started with the work on the Koljoe fjord observatory, a veritable 'ocean in a nutshell'. The initial results of this endeavour have shown that any future infrastructures in accordance with agenda described above must provide data with an extreme high-level of reliability: meaning data loss of less than 1%. Otherwise the entire time series of scientific data would be rendered useless. Within the EMSO this has led to the idea that new data collection, processing, and archiving concepts have to be developed to achieve this goal and justify the anticipated significant investments.

In a number of workshops and discussions it has become clear that these issues are the focus of European telecommunication companies and that ideas have already been formulated to address these issues. However, it is clear that EMSO will need to seek future cooperation with stakeholders in this field. Ideally, new developments that will generate commercial value will come from this cooperation. An open-access policy for collected data will increase the involvement of a wider pan-European community, including a new generation of researchers and from those areas where access to high-quality data is rare or absent.

One of the core services of the EMSO-ERIC is defined as the capacity to deliver basic, established, standardised data products and services to international agencies responsible for Earth monitoring. EMSO will therefore not only provide scientific research data, but it will also supply deep-ocean data to the Global Monitoring of Environment and Security (GMES) and the Global Earth Observation System of Systems (GEOSS) programmes in order to integrate and complement marine services of the GMES satellite, sea surface and subsurface observing systems. EMSO will supply the required synchronous measurements from fixed locations in four of the five areas defined by the marine services as the Global Ocean, the North East Atlantic, the Arctic, and the Mediterranean Sea. EMSO is also a partner to ICG/NEAMTWS (Intergovernmental Coordination Group for the Tsunami Early Warning and Mitigation System in the north-eastern Atlantic, the Mediterranean and connected seas) to monitor earthquake-related activities in marine seismogenic areas. In addition, it will help member states facilitate contributions deemed important to realise the MSFD vision.

Another important aspect relates to data management. An open access policy requires all data collected by EMSO be freely accessible. The immediate delivery of data from the bottom of the ocean directly to researchers' desktops across the globe turns observatory networks into 'gateways to the oceans'—making them accessible to scientists, educators, and the public alike. Consequently, it will increase our understanding of global change, which will lead to socio-economic benefits, such as improved climate forecasts, natural resource management, and human impact mitigation. A European research network of ocean observatories will transform the ability of science to inform European government policy makers and business strategists with greater certainty and efficiency than in the past.

The impact of deep-sea real-time data cannot be overstated in its contribution to mitigating the effects of geo-hazards. That is why better knowledge of deep-sea phenomena can lead to the development of new procedures and algorithms for early warning systems. Particular efforts are taking place to implement reliable real-time data transfer technologies that can be used not only by scientific end-users but also for civilian safety and protection.

e-Infrastructure

EMSO plans to exploit the power of the European Grid Infrastructure (EGI) to create a data infrastructure that will serve the vast communities of scientists studying, for example, marine mammals' underwater noise, oceanography, geophysics, astroparticle physics or ecology. This includes the special objective to also provide open access and shared tools for collaborative studies with state-of-the-art analysis algorithms. The distributed computing paradigm of the EGI e-infrastructure will be used to provide large CPU and storage capacity.

EMSO computing models and data distribution infrastructures should be similar to those adopted by the LHC community with a layered (tiered) structure. Tier0s will be created as close as possible to the experimental sites and will host raw data, Tier1s will host replicas of the data needed for the analysis and will run analysis jobs. In the future, the need for a Tier2 layer will be evaluated, but we will start with just two layers of Tier0 and Tier1. Once the infrastructure will be set up at the European level, plans call for the creation of a 'broker' service to receive user requests for data and analysis and dispatch them to appropriate grid resources.

Current e-Infrastructure activity

An EMSO and grid pilot was agreed upon between the Italian NGI and the EMSO management. This activity aims at deploying a grid-based solution for the data repository and for the basic tools needed to manage data acquired at the Western Ionian Sea node near Catania. Data will be generated by two offshore experimental sites and will reach harbours through fibre cables located offshore Catania (Fig.3).




Fig.3 - The EMSO node offshore Catania

The major data will come from hydrophones at a rate of about 400 GB/day (Table 2): 4 hydrophones with a sampling rate of 96kHz are located in the northern branch site providing about 3 Mb/s each (a bit more given to the overhead of the transmission protocol), four hydrophones with doubled sampling rate (192kHz) are located in the southern branch site with the rate doubled accordingly. So 4x3 Mb/s from the north plus 4x6 MB/s from the south totals about 36 Mb/s, which is equivalent to about 11 TB/month.

This data will be analysed in real-time by a first level of trigger, but EMSO would like to save all the raw data. Data will first be saved on a server located in the harbour and then transferred using rsync to LNS (INFN South Laboratories) in Catania where a grid site is already installed. The storage in the harbour is limited in space and organised as a round robin buffer.

Table 2 - Data rate produced daily/monthly by NEMO-SN1

<i>Instrument</i>	<i>Byte/day</i>	<i>Byte/month</i>
oceanographic sensors	13 M	390 M
magnetometers	2.5 M	75 M
seismic sensors	933 M	28 M
bio-acoustic hydrophones	372 G	11 T
station monitoring	180 M	5.4 G

Grid storage has been installed at the LNS-INFN Catania site, mounting the raw data from the local storage, common grid data management services and clients can be used to access the data.

Future e-Infrastructure challenges

- Replicate the pilot activity in other EMSO sites
- Create T1 sites
- Create high level interfaces to access and process the data. Such interfaces should be able to:
 - Perform user authentication and authorisation based on identity federation
 - Locate the requested files
 - Perform replicas and data transfers within Grid services but also towards other types of storage already in use in the community (i.e. MoistDB) and possibly interoperate with catalogues of different infrastructures
 - Catalogue interoperability
 - Allow analysis jobs (in a second phase)

Future plans

The implementation of EMSO is a complex undertaking that calls for a concerted trans-European effort. In addition to operators of existing infrastructures and involved scientific institutions, other stakeholders in the field must be encouraged to become engaged as well.

EMSO had begun building connections to supercomputing centres and to the private sector to get an overview of the state-of-the-art e-infrastructure architectures in Europe. In this phase it is critical to set-up a close communications between scientific users and the organisations that will be in charge of implementing the e-infrastructure. Only through close communication can sustainability be achieved.

The overall duration of the first phase of the EMSO implementation will be five years, beginning in 2012 (Picture 4), with a review point scheduled at year 3. During this phase activities will be focused on:

- Establishing the central EMSO-ERIC managing organisation;
- Construction and/or upgrading of EMSO's seven sites (Arctic, Porcupine, Azores, PLOCAN, Eastern Sicily, Hellenic Arc, Ligurian Sea) will commence implementation from year 1;
- Setting up a distributed data management system that will guarantee open access to data from EMSO observatories to scientists and other stakeholders.

This investment led to the construction of facilities and hardware that will be available for the first five years of the EMSO implementation phase. These installations and facilities will be provided as an in-kind contribution by EMSO partner funding agencies toward the implementation of the Research Infrastructure.




Fig.4 - EMSO-ERIC timeline

References & Bibliography

- Aymerich, I.F., J. Piera, J.J. Danobeitia (2011) Cable Planning for the European Multidisciplinary Seafloor Observatory. *Sea Technology*, 52(6):110-112.
- Beranzoli, L., P. Favali, G. Smriglio (Eds.) (2002) "Science-Technology Synergy for Research in the Marine Environment: Challenges for the XXI Century, Developments in Marine Technology." Elsevier, Amsterdam
- Embricaco, D., G. Marinaro, F. Frugoni, S. Monna, G. Etiopie, L. Gasperini, A. Polonia, F. Del Bianco, ... (submitted) Combined gas seepage and seismological monitoring along the North-Anatolian Fault in the Sea of Marmara (Turkey) by benthic observatory. *Geophys. Journ. Int*
- Favali, P., L. Beranzoli (2009) EMSO: European Multidisciplinary Seafloor Observatory, *Nucl. Instrum. Meth. A.*, 602:21-27
- Favali, P., L. Beranzoli, F. Italiano, E. Migneco, M. Musumeci, R. Papaleo, R (2011) NEMO-SN1 observatory developments in view of the European Research Infrastructures EMSO and KM3NET. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 626-627:553-556
- Fabry, V.J., B.A. Seibel, R.A. Feely, J.C. Orr (2008) Impacts of ocean acidification on marine fauna and ecosystem processes. *ICES Journal of Marine Science*, 65:414-432
- Feely, R.A., C.L. Sabine, K. Lee, W. Berelson, J. Kleypas, V.J. Fabry, F.J. Millero (2004) Impact of anthropogenic CO₂ on the CaCO₃ system in the oceans. *Science* 305:362-366
- Glover, A.G., A.J. Gooday, D.M. Bailey, D.S.M. Billett, P. Chevaldonnè, A. Colaço, J. Copley, ... (2010) Temporal change in deep-sea benthic ecosystems: a review of the evidence from recent time-series studies. *Advances in Marine Biology*, 58.
- Levin, L., R.J. Etter, M.A. Rex, A.J. Gooday, C.R. Smith, J. Pineda, C.T. Stuart, R.R. Hessler, D. Pawson (2001) Environmental influences on regional deep-sea species diversity. *Annual Review of Ecology and Systematics*, 32:51-93
- Marzeion, B., A. Levermann, J. Mignot (2010) Sensitivity of North Atlantic subpolar gyre and overturning to stratification-dependent mixing: response to global warming. *Climate Dynamics*, 34:661-668
- Orr, J.C., V.J. Fabry, O. Aumont, L. Bopp, S.C. Doney, R.A. Feely, A. Gnanadesikan, N. Gruber, ... (2005) "Ocean Acidification due to Increasing Atmospheric Carbon Dioxide". Royal Society, London. (pp:60)

Ruhl, H.A., M. André, L. Beranzoli, M.N. Çağatay, A. Colaço, M. Cannat, J.J. Dañobeitia, P. Favali, ... (2011) Societal need for improved understanding of climate change, anthropogenic impacts, and geo-hazard warning drive development of ocean observatories in European Seas. *Progress in Oceanography*, 91:1-33

Fusion on the grid: transport calculations

Francisco Castejón ✉

Laboratorio Nacional de Fusión, Asociación Euratom/CIEMAT para Fusión, Spain

22 January 2013

✉ Francisco.castejon@ciemat.es

Overview

The achievement of fusion energy needs to control the plasma phenomena, and plasmas are complex systems that obey non-linear equations whose solution requires strong computing capacity. ITER⁴ plasmas will be ready for analysis and understanding from 2020 on. Figure 1 shows the ITER tokamak. There are parallel applications that need the use of HPCs and distributed calculations that run properly on the grid (Castejón and Gómez-Iglesias, 2011). It is also possible to find workflows that mix the two types of applications making necessary the use of a grid of parallel nodes.

We propose here two examples of applications: the first one related to the study of ion kinetic transport in three-dimensional devices and the second one to the estimate of diffusive transport with different sources of particles and power. The TJ-II device (Alejaldre et al., 1990), in operation at CIEMAT in Madrid since 1997, will be also used for our calculations.

The first use case is based on the code ISDEP (Integrator of Stochastic Differential Equations in Plasmas), which is a Monte Carlo code that calculates the ion collisional transport in magnetic confinement devices by the estimate of a large number of single ion trajectories (Velasco et al., 2012).

The second use case implies the use of the transport code ASTRA in combination with applications that simulate sources and sinks of particles and power. The codes that integrate the workflow can run on HPC or grid.

⁴ The International Thermonuclear Experimental Reactor (<http://www.iter.org/>)




Fig.1 - ITER device and plasmas.

Scientific case

The scientific challenge: improving plasma confinement

Commercial fusion requires the improvement of plasma confinement or, equivalently, reducing the transport of energy and particles. The gradients of the main plasma parameters cause transport and some degradation of the confinement. Several transport regimes can be considered depending on the nature of the phenomena involved in the generation of the fluxes of energy and particles driven by the plasma gradients.

When the 3D structure of the device plus the electric field and collisions are considered as the only mechanisms we speak about Neoclassical (NC) Transport. This includes a series of approximations that can be violated in the low collisionality regimes, i. e. in the case of hot low density plasmas, and in the cases of intense radial electric field. In these regimes it is even possible that transport is not diffusive, i. e., the fluxes are not proportional to the gradients. If the NC ordering is satisfied, classical codes like the Drift Kinetic Equation Solver, or DKES (Hirshman et al., 1986), which has been recently ported to the grid (Rubio-Montero et al., 2010) can be used.

On the opposite case, when NC ordering is violated, it is mandatory to use a global Monte Carlo 5D code (3D in real space and 2D in velocity space). In our case we use ISDEP to estimate the main properties of transport in the low collisionality regime in a device as complex as TJ-II (Velasco et al., 2009).

Figure 2 shows the TJ-II device and plasma. In this case, the Hurst parameter (H) shows that transport is not diffusive (Castejón et al., 2002) ($H=0.5$ means diffusive transport, $H<0.5$ means sub-diffusive and $H>0.5$ means ballistic transport), as can be seen in Figure 3. Another important calculation that can violate the NC ordering is the one of fast ion transport. The good confinement of fast ions is mandatory in a fusion reactor, so its study is one of the key points for plasma theory.




Fig.2 - TJ-II device and plasmas.

When one does not have an *ab initio* theory of transport, it is necessary to perform transport analysis: once the sources of energy and particles are known, it is possible to obtain the fluxes and to verify its nature as well as to compare them with the experimental data. Proctr code (Vargas et al., 2007) and ASTRA (Castejón et al., 2007) have been used to perform such transport analysis on TJ-II plasmas.




Fig.3 - Profile of Hurst exponent in plasmas of TJ-II device. A ballistic behaviour is observed for the outer positions.

State of the art

The kinetic collisional transport of thermal ions when NC ordering is violated has been already calculated for the TJ-II stellarator (Castejón et al., 2009) and for the ITER tokamak (Bustos et al., 2010) by running ISDEP on the grid. The properties of this transport that violates NC ordering have been thoroughly characterised in these two works. The confinement time, the escaping points and the global fluxes are calculated. One important consequence of these calculations, with strong impact on the community is the toroidal and poloidal asymmetries of the fluxes (Bustos et al., 2011).

Fast ion transport has been estimated for TJ-II and LHD stellarators in Japan (Bustos et al., 2013). These calculations were compared with the experimental data in Japan and allowed a collaboration between Fusion NIFS group in Japan and CIEMAT.




Fig.4 - Velocity distribution function of fast ions in LHD device at four radial positions.




Fig.5 - Comparison of the measured (red points) spectrum of fast ions with the calculated one (black line) in TJ-II. The persistence of ions in the plasma is also plotted in the small chart.

Beyond the transport properties, the fast ion distribution function has been estimated for these two stellarators, which is an important achievement that allows one to estimate important plasma parameters like the slowing down time and the plasma rotation. Figure 4 shows the velocity distribution function of fast ions in LHD stellarator for several plasma positions (the parameter rho is defined as the normalised radius $\rho=r/a$, where a is the outermost minor radius of the plasma). The calculations of this distribution function implied to follow about 106 trajectories with a CPU time exhausted of 108 seconds. The comparison with experiments shows a good agreement between the calculations and the measurements (Fig.5) (Cappa et al., 2011).

Regarding the transport analysis, for the first time, calculations were performed with a varying power source, since this depends on plasma conditions that were allowed to vary in time (Chelouche et al., 2009). These calculations were performed by establishing a complex Compendium of e-Infrastructure requirements for the digital ERA

heterogeneous workflow between the ASTRA code and the TRUBA ray tracing code. The first one was running on an HPC, while the second one was running on the grid. Figure 6 show an example of ray tracing, while Figure 7 shows the simulated plasma evolution. Ray tracing calculations have given rise to other collaborations with international groups devoted to different fields, for instance, belonging to the astrophysics community (Rodríguez-Pascual et al., 2010).




Fig.6 - Rays of microwaves in the TJ-II plasma.




Fig.7- Time evolution of plasma parameters in TJ-II. Continuous lines are calculated, while dotted lines are experimental results.

Going beyond the state of the art

To improve the results of ion kinetic transport we would need to perform the following improvements and developments:

1. Calculation of fast ions on ITER device. Up to date, only thermal ion transport has been estimated. A similar exercise to the one that is performed for stellarators is necessary for such a large tokamak.

2. Establishing an automatic workflow between the Monte Carlo code FAFNER devoted to Neutral Beam Heating and ISDEP, which follows fast ions in the complex 3D geometry of the device. FAFNER2, which has been ported to the grid (Castejón and Egulior, 2003) provides the birth points of fast ions, i. e., input for ISDEP.
3. It is necessary to include the interaction between ions and waves in the code, in order to estimate the effect of heating and instabilities on ion transport.

The first two points require some computing science work, while the third one implies a deep insight on the physics of the phenomena, including the casting of the wave-particle interaction equations, in the way of Langevin, which was introduced by Castejón et al. (2008).

Transport analysis requires establishing a complex heterogeneous workflow between applications that run on the grid and other that run on HPCs. The de facto standard workflow engine in the Fusion community is Kepler, so we propose to join several codes to ASTRA using Kepler. The main codes are:

1. The ray tracing code TRUBA (Castejón et al., 2008) that runs on the grid. TRUBA provides the electron heating by microwaves.
2. The MC code FAFNER2 that provides the heating by NBI that also runs on the grid.
3. EIRENE is a large code that provides the particle source and runs on HPC (Guasp, 2007).
4. DKES is a code that provides NC transport coefficients and can run both on HPC and on the grid.
5. EUTERPE or a similar code that provides turbulent transport coefficients and run on HPC (Sánchez et al., 2010).

Taking into account all these codes, it is possible to build different workflows that can be used to understand different phenomena relevant for plasma confinement, for example, influence of fuelling, heating, collisionality, turbulence.

e-Infrastructure

Challenges and computer requirements

We need high-capability computing infrastructures, storage capacity and fast networks to connect the infrastructures in the case of establishing complex workflows between applications.

The grid calculations can be performed on the Fusion Virtual Organisation, while an HPC should be provided for the parallel applications.

Tens of TB of memory are required, especially for the turbulence and EIRENE calculations.

Timeline

The calculations for both use cases are already underway. We account with 3 persons. The first use case will take about 3 ppy (persons per year) and the second 4 ppy, so we would need the effort of our three persons during 2.5 years. Taking into account contingencies and working in parallel in the two cses, the target can be achieved in a 3-year horizon.

References & Bibliography

- Alejaldre, C., J. Alonso, J. Botija, F. Castejón, J. R. Cepero, J. Guasp, A. L. Fraguas, L. García, V. Krivenski, ... (1990) TJ-II: A Flexible Helical Stellarator. *Fusion Technology*, 17:131
- Bustos, A., F. Castejón, L.A. Fernández, J. García, V. Martín-Mayor, A. Tarancón, J.L. Velasco (2010) 3D features of ion collisional Transport in ITER. *Nuclear Fusion*, 50:125007
- Bustos, A., F. Castejón, M. Osakabe, L. A. Fernández, V. Martín-Mayor, J. Guasp (2011) Kinetic simulations of Fast ions in stellarators. *Nuclear Fusion*, 51:083040
- Bustos, A., J.M. Fontdecaba, F. Castejón, V. Martín-Mayor, M. Tereshchenko, L.A. Fernández (2013) Studies of the fast ion energy spectra in the TJ-II. *Physics of Plasmas* 20:022507
- Cappa, Á., D. López-Bruna, F. Castejón, M. Ochando, F. Medina, J.L. Vázquez-Poletti, ... (2011) Calculated evolution of the Electron Bernstein Wave heating deposition profile under NBI conditions in TJ-II plasmas. *Contribution to Plasma Physics*, 51:83
- Castejón, F., V. Tribaldos, I. García-Cortés, E. de la Luna, J. Herranz, I. Pastor, T. Estrada, TJ-II Team (2002) Enhanced heat confinement regime in the flexible heliac TJ-II. *Nuclear Fusion*, 42:271-280
- Castejón, F., S. Eguilior (2003) Langevin equations for quasi-linear wave-particle interaction. *Plasma Physics and Controlled Fusion* 45:159
- Castejón, F., L.A. Fernández, J. Guasp, V. Martín-Mayor, A. Tarancón, J.L. Velasco (2007) Ion kinetic transport in the presence of collisions and electric field in TJ-II ECRH plasmas. *Plasma Physics and Controlled Fusion* 49:753-776
- Castejón, F., Á. Cappa, M. Tereshchenko, Á. Fernández (2008) Computation of EBW heating in the TJ-II Stellarator. *Nuclear Fusion* 48:075011
- Castejón, F., J.L. Velasco, A. López-Fraguas, A. Tarancón, J. Guasp, F. Tabarés, M.A. Pedrosa, ... (2009) Flux-expansion divertor studies in TJ-II. *Nuclear Fusion* 49:085019
- Castejón, F., A. Gómez-Iglesias (2011) Advances in Grid Computing (Chapter) In: *Grid Computing for Fusion Research*, pp:215-228. Ed.: Intech. Rijeka. ISBN 978-953-307-301-9.
- Chelouche, D., R. Rabadán, S.S. Pavlov, F. Castejón (2009) Spectral signatures of photon-particle oscillations in celestial objects. *The Astrophysical Journal (ApJ) Supplement Series* 180:1-2
- Guasp, J. (2007) EIRENE code calculations for the CX spectra in the TJ-II Stellarator. CIEMAT report
- Hirshman, S. P. , K. C. Shaing, W. I. Van Rij, C. O. Beasley, E. C. Crume (1986) *Phys. Fluids*, 29:295
- Rodríguez-Pascual, M., J. Guasp, F. Castejón Magaña, A.J. Rubio-Montero, I.M. Llorente, R. Mayo García (2010) Improvements on the Fusion Code FAFNER2. *IEEE Transactions of Plasma Science*, 38:2102
- Rubio-Montero, A.J., F. Castejón, M. Rodríguez-Pascual, E. Montes Prado, R. Mayo (2010) Drift Kinetic Equation Solver for Grid (DKEsG). *IEEE T. Plasma Sci.*, 38:2093-2101
- Sánchez, E., R. Kleiber, R. Hatzky, A. Soba, X. Sáez, F. Castejón, J.M. Cela (2010) Linear and non-linear simulations using the EUTERPE gyrokinetic code. *IEEE Transactions of Plasma Science*, 38:2119HED 1
- Vargas, V.I., D. López-Bruna, J. Herranz, F. Castejón, and the TJ-II Team (2007) Experimental electron heat diffusion in ECH plasmas of the TJ-II stellarator. *Nuclear Fusion* 47:1367-1375
- Velasco, J.L., F. Castejón, A. Tarancón. (2009) Finite orbit width effects in ion collisional transport in TJ-II. *Physics of Plasmas*, 16:052303
- Velasco, J.L., A. Bustos, F. Castejón, L. A. Fernández, V. Martín-Mayor, A. Tarancón (2012) ISDEP: Integrator of Stochastic Differential Equations for Plasmas. *Computer Physics Communications*

Grid-empowered molecular simulations: from first principles to new materials

Antonio Laganà¹✉, Carlo Manuali²✉, Alessandro Costantini³✉

^{1,2} Department of Chemistry, Computational Dynamics and Kinetics Group, University of Perugia, Italy

³ INFN-Perugia and Italian Grid Infrastructure, Italy

27 August 2013

✉ ¹lagana05@gmail.com, ²carlo.manuali@gmail.com,
³alessandro.costantini@pg.infn.it

Overview

The molecular and materials sciences research experts of the Chemistry, Molecular & Materials Science and Technology (CMMST) community are using distributed computing to tackle multi-scale simulation challenges in a wide range of innovative fields, including chemical engineering, biochemistry, chemometrics, 'omic'-sciences, forensic chemistry, medicinal chemistry, food chemistry, energy production and storage, new materials, space technologies and more.

Joint endeavours with international Research Laboratories encompass programmes devoted to:

- the *ab initio* calculation of the electronic structure and assemblage of *ab initio*-based potential energy surfaces;
- the integration of quantum and/or classical equations of motion;
- the definition of additional treatments necessary to calculate the value of measurable quantities.

These programmes have been developed and/or gathered together and implemented as a coordinated realistic simulator that serve as innovative molecular science use cases after their porting on the computing grid.

Present plans for future utilisation of such use cases include the joint design, development and deployment of user-friendly interfaces to:

- a. consolidate cooperative usage of simulator components;
- b. improve the coordinated usage of high-performance and high-throughput codes;
- c. develop a quality-based collaborative credit economy;
- d. extend the use of the simulator to research-based distributed learning.

Scientific case

In recent years, molecular simulations have played an increasingly important role in the investigation of the structure, dynamics, surface properties and thermodynamics of inorganic, biological and polymeric systems. This is because they can provide a high level of detail in describing processes at sub-nanometre scale, thanks to the use of current ICT. Molecular simulations, in fact, can produce results that are not measurable using the current level of experimental technology and devices, such as surface interactions on an atomic level, flow at the exit of a micro electric thruster, or the molecular flow inside a star. In particular our attention has focused on advanced simulations such as planetary re-entry aerothermodynamics, clean combustion processes, engineering of energy production and storage and biological modulation of trans-membrane ionic fluxes.

To advance our work, a Molecular Simulator (Costantini, et al., 2010) assembled with in-house and third-party programs was implemented on the computing grid. In applying a service-oriented approach, the simulator is used as a cooperative computational engine for real-world applications with a high level of complexity aimed at supporting research, innovation and development in fields that have a strong economic and social impact. This cooperative endeavour is based on the combined expertise and efforts of over 20 European molecular science academic groups, as well as several non-European groups, and the adoption of high level ICT tools, such as the ones of the European grid, with appropriate adaptations.

The scientific challenge

The progress in theoretical and computational molecular science, together with the consequent detailed knowledge of microscopic structures and processes achieved in recent years, has enabled the assemblage of multi-scale simulators starting from first principles. We have gathered the molecular components of such simulators in an articulated use case.

The first step of the simulator consists of software components gathered into a module called INTERACTION. This module consists of several highly successful *ab initio* quantum chemistry packages which determine the electronic structure of molecular systems, within the Born-Oppenheimer scheme that separates the problem of the electrons from that of the nuclei-related properties. In the simulator, calculations of each molecular electronic structure can be performed either individually for each integration step of the nuclei dynamics equations (on-the-fly technique) of the DYNAMICS module, or globally fitted in the FITTING module by using some established routines to a global functional form, called potential energy surface (PES), that is then imported in the DYNAMICS module. Different methods and programmes are considered for this purpose. For small molecular systems, or parts of them, exact quantum calculations, that are either time-dependent or time-independent, can be run on the grid to calculate exact probabilities. On the contrary for larger systems, the quantum nature of atoms and molecules is partially neglected and, accordingly, empirical interactions (though often calibrated by a comparison with *ab initio* calculations) are adopted together with classical dynamics treatments. These are eventually complemented, when possible, by semi-classical treatments such as SC-IVR. The probabilities obtained by running related packages are input into the final OBSERVABLES module for appropriate statistical treatment leading to the evaluation of measurable quantities.

Compendium of e-Infrastructure requirements for the digital ERA

State of the art

Thanks to the support given by the COMPCHEM Virtual Organisation (VO), the CMMST community was able to implement some components of the simulator on the grid. They have been used to carry out massive computational campaigns to simulate the molecular components of some innovative technologies, such as the gas hydrated formation for energy storage, the flux of ions through trans-membrane cellular micropores, the heat transfer around spacecraft in planetary re-entry or the virtual measurement of scattered products in molecular beams apparatuses as illustrated below.

Energy storage devices: gas hydrates formation

Gas hydrates (e.g. Albertí et al., 2012) are ice-like solid inclusion compounds that result from the trapping of gas molecules within a lattice-like cage of water molecules. Many gases have molecular sizes suited to form hydrates like methane (CH_4) and carbon dioxide (CO_2). Important benefits can be derived from the exploitation of this property. For instance, at standard conditions of pressure and temperature, one volume of methane hydrate can store up to 164 volumes of gaseous methane, allowing its safer and less expensive storage and transportation. In recent years, by using the simulator components, it has been possible to assemble a new *ab initio*-based force field and run massive Molecular Dynamics calculations showing, as illustrated in Figure 1, how the SDS surfactant is able to cage CH_4 molecules in water.




Fig.1 - Screen shots (from left to right) of the folding process of the SDS- CH_4 caging CH_4 in water.

Membrane micropores: ion channel simulation

Ion channels (e.g. Skouteris and Laganà, 2006) are crucial elements of the living cell activity. Placed on the membrane of the cell, they allow particular ions to pass across from one side of the membrane to the other. They can open and close access to exert control over the ion flux and govern functions of key biological and medical importance. In this case, the components of the simulator were used to build an *ab initio*-based PES of an open-end carbon nanotube that models the micropore, and run massive reduced dimensionality quantum and full dimensional classical Molecular Dynamics calculations to compare the flux of different ions (Fig.2).




Fig.2 - Solvated ion flux through an open end carbon nanotube.

Aerothermodynamics effects in spacecraft re-entry: heat transfer evaluation from molecular dynamics

Gas and gas-surface processes (e.g. Bartolomei et al., 2012) occurring within the hypersonic flow layer surrounding spacecraft during planetary entry play a key role in energy distribution and heat transfer on the shields and flaps. Thanks to the use of simulator components, a large database of the detailed, non-equilibrium rate coefficients of the $\text{CO}_2 - \text{CO}_2$ molecular system was created as a result of extended computational campaigns aimed at determining the value of the detailed cross sections and rate coefficients. In addition, this study has enabled the determination of the dependence of the efficiency of the considered processes on the initial partitioning of energy among the various degrees of freedom.

Virtual crossed beam experiments: estimating the efficiency of elementary reactive processes intervening in combustion

The *ab initio*, fitting, classical dynamics and statistics components of the simulator have been used to simulate crossed beam measurements and reproduce the experimental signal (Laganà et al., 2012). The investigated system is $\text{OH} + \text{CO}$ whose reactive cross section has been measured using a crossed beam apparatus. The direct comparison of theoretically predicted laboratory angular distributions with measured product beam intensity raw data avoids possible uncertainties associated with inversion procedures.

Going beyond the state of the art

Our objective is to go beyond the stage in which only individual components have been both implemented on the grid and local resources (Table 1) and transform the simulator into a highly automated computational engine. To overcome those highly unsatisfactory situations in which neither HPC nor HTC are completely fit to meet the requests of complex CMMST applications, effort has to be spent in order to effectively coordinate the use of HPC and HTC e-Infrastructures (Manuali et al., 2012). The guidelines for a prototype of this have been outlined in the prototype hydrogen system to be used as a fundamental computational molecular engine for multi-scale simulations. Further tools and user-friendly interfaces should be designed, developed and deployed in order to:

- a. Scale up the complexity and the dimensionality of the molecular systems treatable by the simulator by enhancing the interoperability between high performance and high throughput platforms in order to redirect properly the different sections of the

calculations needing HPC platforms. This implies a restructuring of the codes to enable them to run on modern parallel architectures.

- b. Improve the suitability and coordination of the individual components of the simulator by enhancing automatisms and further developing and implementing the requirements of (*de facto*) standards for data formats as part of the maintenance work carried out by the COMPCHEM research groups.
- c. Further elaborate quality parameters characterizing the work of the COMPCHEM members in order to use them to evaluate the actual contribution to the collaborative work of the VO and establish a credit system fostering the growth of a collaborative community grid economy
- d. Investigate the application of the simulator to molecular science knowledge in order to extend its use in research-based distributed learning.

Table 1 - HPC and HTC requirements of some components of the simulator

Application	HPC requirements			HTC requirements		
	Cores /job	Memory (GB/core)	Total allocated time (h)	Cores /job	Memory (GB/core)	Total allocated time (h)
ABC	32	2	25k	1	2	50k
MCTDH	1	5	50k	5	1	40k
GAMESS-US	32	2	300k			
DL_POLY	64	2	100k	8	1	500k
SC-IVR	16	2	50k	1	2	50k

Such progress will enable a giant leap forward for both the quantum applications of the simulator to heavier few atom systems, including Born-Oppenheimer treatments of processes involving nuclei, and of the classical mechanical investigation of polyatomic molecules and large ensembles of molecules, including biological and omics sciences. This will also lead to important changes in the way computing resources are made available to communities and in the model adopted for cooperative research.

e-Infrastructure

The CMMST (Chemistry, Molecular and Materials Sciences and Technologies) community is currently supported by the COMPCHEM VO which relies on about 30,000 CPUs and 1.5 PB of disk storage thanks to the support of more than 25 European research institutes and universities, primarily in France, Greece, Italy, Poland and Spain (see Figure 4 and Table 2 for details). Most of the resources are shared with other VOs. This means that jobs sent to the grid infrastructure by COMPCHEM users compete with jobs from other VOs.




Fig.4 - Pie chart showing percentage of Normalised CPU time (kSI2K) consumed by COMPChem on the grid platforms of the various NGIs between May 2012 and April 2013

Table 2 - Normalised CPU time, in 1kSI2K hours, listed by date and NGI

DATE	NGI_FRANCE	NGI_GRNET	NGI_IBERGRID	NGI_IT	NGI_PL	NGI_UK	Total
May-12	166,983	167,104	46,314	241,323	19,079	1,014	641,817
Jun-12	137,670	185,567	55,952	434,669	29,243	742	843,843
Jul-12	255,055	213,168	66,182	618,265	39,178	16,931	1,208,779
Aug-12	263,501	210,714	85,777	605,767	31,060	274	1,197,093
Sep-12	67,487	59,952	28,127	84,172	8,129	100	247,967
Oct-12	87,258	163,320	976	386,956	24,010	0	662,520
Nov-12	372,586	193,233	57,720	619,342	32,957	0	1,275,838
Dec-12	307,320	148,718	250,279	932,755	47,777	0	1,686,849
Jan-13	427,404	316,119	225,897	1,011,541	47,387	0	2,028,348
Feb-13	302,648	170,382	65,623	856,630	25,627	0	1,420,910
Mar-13	837,012	201,562	124,378	1,681,296	56,199	0	2,900,447
Apr-13	177,954	53,423	27,741	727,183	16,635	0	1,002,936
Total	3,402,878	2,083,262	1,034,966	8,199,899	377,281	19,061	15,117,347
%	22.51%	13.78%	6.85%	54.24%	2.50%	0.13%	

Current e-Infrastructure Activity

To exemplify the state of the art, the CPU time consumption of COMPChem during the 12-month period from May 2012 to mid-April 2013 is plotted in Figure 5.

As apparent from Table 2 and Figure 5 there has been a factor-of-two increase from the first to the second half of the 2012-2013 period. This trend confirms the constant computing time usage increase observed in recent years, which is expected to continue in the future. However, from the planned combined use of GriF (to optimise the selection of the computed sites) and tools bridging HTC and HPC, a factor-of-ten increase over the previous year is expected. The quantities illustrated in Figures 4 and 5 and Table 2 are only a partial indication of the quality of a scientific community's activities.

More updated figures confirming this trend can be obtained from the EGI Accounting Portal hosted by Centro de Supercomputación de Galicia (CESGA). Even higher figures are obtained when adding the large share of computer time obtained by the members of the involved VOs through research projects on the supercomputers of the European large scale facilities involved in PRACE.




Fig.5 - Last year's monthly COMPCHEM Cumulative Normalised CPU time given in kSI2K

Thanks to the activities of its members, in fact, several CMMST applications have been ported in the grid environment where dedicated user support activities are provided and supported by some grid-enabled tools and services currently in use.

In particular, a selection of the most used grid-enabled tools and services in the CMMST community is presented below:

AppDB

The EGI Applications Database (AppDB⁵) is a central service that stores and provides information about tailor-made software tools for scientists and developers to use, the programmers and the scientists who developed them, and the publications derived from the registered software items to the public.

All software filed in the AppDB is ready to be used on the European Grid Infrastructure.

User interfaces and frameworks

- GRIF: Grid Framework enabling efficient and user-friendly massive scientific calculations
- GCRES: Quality of Users (QoU), Quality of Services (QoS) evaluation Framework
- GGAMESS: Front-end script for submitting multiple GAMESS-US jobs
- G-LOREP: Search tool of LOs for large communities operating on complex distributed repositories

⁵ <http://appdb.egi.eu/>

Workflows

Tools developed to govern complex ensembles of data, models and programs of an increasing number of applications offering a unified user friendly way of composing related tools. Workflows engines used by some members of the community include KEPLER, P-GRADE and GC3Pie.

Data management

- FTS (File Transfer Service) is a lightweight but fully functional set of services supporting data management
- LFC (LCG File Catalog)

Libraries

- MPI (Message Passing Interface), a library of routines providing concurrent execution of parallel programs.

Future e-Infrastructure challenges

The future e-Infrastructure challenges we face in order to keep pace with the GEMS roadmap are:

1. Allow user-friendly access to involved platforms. Access should be mimetic as possible with the user environment.
2. Seamlessly integrate the use of different platforms, such as HTC and HPC, via tools that select the proper computing resources for various tasks for the same or divergent jobs.
3. Enable appropriate coordination tools combining in a single workflow (or workflow of workflows) different packages and alternative tasks.
4. Develop new sensors for quality evaluation able to consider the user behaviours (e.g. use of the machine, tools development, research field) and of the available tools and services, such as ready-to-use production applications, interoperable tools or libraries.
5. Establish a suitably-sized community to face these challenges in a coordinated federation of Virtual Organisations (VOs) with specific interests in GEMS modules, or sections of modules.

To advance these efforts and make a significant leap forward to higher performing technologies and cooperative sustainability schemes, COMPChem has established an EGI Virtual Team⁶ and plans to further develop GriF, a framework aimed at selecting specific grid resources, that will be modified towards a deeper integration with evolved middleware products. In order to enhance the cooperative spirit of the CMMST community, GriF will be modified to capture data produced by grid sensors on the quality of both the users serviced and the services provided. All this will be complemented by GCreS, a tool turning the QoU and QoS measurements based on the data produced by GriF into a collaborative credit economy, in order to support cooperation among the community members.

⁶ https://wiki.egi.eu/wiki/Towards_a_CMMST_VRC

Finally, we aim to extend GEMS in the long-term to be used for grid distributed repositories of teaching and learning objects in research-based education. For this purpose the EGI VT has involved two large groups of chemists: the Computational Chemistry division of the European Chemistry and Molecular Science (EUCHEMS) Society and the members of the European Chemistry Thematic Network (ECTN) Association.

Future plans

The overall strategy of the present project consists of implementing the above measures through the concurrent contribution of all the members of the CMMST community participating in the design of the related virtual community. The specific aim is to extend GEMS, the grid-empowered molecular simulator, to act as a full *ab initio* generalised computational engine providing a rigorous evaluation of the contribution of the molecular structure and processes to the fate of multi-scale applications of innovative technologies. Accordingly, the following tasks must be completed:

- TASK 1: Complete the set of programs and packages necessary to implement a full quantum version of the simulator.
- TASK 2: Identify the HPC platforms suitable to support the project by opening use of these resources to the grid and the finding middleware solution best suited to operate both on the EGI grid and HPC platforms.
- TASK 3: Integrate HPC/HTC workflows by adapting the GriF server for communicating with the new middleware layer.
- TASK 4: Transform the chosen prototype application into an HPC/HTC workflow and adapt the GriF client.
- TASK 5: Modify the structure of the GriF database to accommodate further quality information and develop a credit economy based on it.
- TASK 6: Establish connectivity with the Virtual Education Community of ECTN association and EUCHEMS for dissemination.

Since the interoperation between HPC and HTC resources is still less than ideal, we rely on the collaboration among major European resource providers to fill this gap. In this regard, a research collaboration has been established between the CMMST Community and some resource centres of the Open Science Grid (OSG) and contacts with the Extreme Science and Engineering Discovery Environment (XSEDE) project are being established. In addition, the suitability of new technologies, such as cloud computing, for their use and adaptation to the CMMST community is currently being evaluated.

References & Bibliography

- Albertí, M., A. Costantini, A. Laganà, F. Pirani (2012) Are micelles needed to form methane hydrates in sodium dodecyl sulfate solutions? *J. Phys. Chem. B*, 116 (14): 4220–4227
- Bartolomei, M., F. Pirani, A. Laganà, A. Lombardi (2012) A full dimensional grid empowered simulation of the CO₂ + CO₂ processes, *J. Comp. Chem*, 33: 1806–1819
- Costantini, A., O. Gervasi, C. Manuali, N. Faginas Lago, S. Rampino, A. Laganà (2010) COMPCHEM: progress towards GEMS a Grid Empowered Molecular Simulator and beyond, *Journal of Grid Computing*, 8(4): 571-586

- Laganà, A. E. Garcia, A. Paladini, P. Casavecchia, N. Balucani (2012) The last mile of molecular reaction dynamics virtual experiments: the case of the OH (N=1-10) + CO (j=0-3) -> H + CO₂ reaction, *Faraday Discussion of Chem. Soc.* 157: 415-436
- Manuali, C., A. Costantini, A. Laganà, M. Cecchi, A. Ghiselli, M. Carpenè, E. Rossi (2012) Efficient Workload Distribution bridging HTC and HPC in Scientific Computing, *Lecture Notes Computer Science* 7333: 345-357
- Skouteris, D., A. Laganà (2006) Study of the passage of an H⁺ ion along a carbon nanotube using quantum wavepacket dynamics, *Lecture Notes in Computer Science* 3980: 757-762

Hydrometeorological research at the computational frontier

A. Parodi ✉, D. Kranzlmüller, F. Siccardi and A. Clematis

DRIHM collaboration

18 June 2013

✉ antonio.parodi@cimafoundation.org

Overview

The World Conference on Disaster Reduction defined among its thematic priorities the worldwide improvement of cooperation in hydrometeorology research (HMR) activities for the prevention and mitigation of risk associated with severe hydrometeorological events. This statement was confirmed at the 2009 joint press conference of the Center for Research on Epidemiology of Disasters (CRED) with the United Nations International Strategy for Disaster Reduction (UNISDR) Secretariat, where it was noted that flood and storm events are among the natural disasters with major impact on human life. Damage caused by extreme precipitation events strongly burdens the budgets of industry, national governments, and international organisations, and affected populations often face economic ruin.

How to reduce the impact of natural disasters linked to extreme weather events, and how to adapt their defenses against the events, more and more extreme due to the fast climate change in progress, are two issues clearly perceived as strategic.

In fact they are two aspects of the same question: are we able to tackle the inherent uncertainty of the extremes of the climate?

EU, US and worldwide authorities try to answer the demand of civil protection organisations dealing with floods, flash floods and debris flows in urban areas by enlarging the observational systems and by supporting modelling tools.

But do we have the right and adequate tools at least to predict when and where sudden disastrous events may occur? If so we will be well placed to deal with broader challenges. Certainly, more spatially extensive and temporally demanding, but conceptually similar and of the same difficulty.

At the heart of this challenge lies the ability to have easy access to hydrometeorological data and to share predictive models. These will facilitate the collaboration between meteorologists, hydrologists and earth scientists, and accelerate scientific advances in HMR, which can be used to protect civilians and the environment, also in consideration of the incoming effects of the climate change.

This scenario calls for an e-Infrastructure based approach, exploiting both grid and HPC resources, to facilitate this collaboration and provide end-to-end HMR services to combine different meteorological, hydrological and impacts models in complex workflows (Schiffers et al., 2011). This will lead the definition of a common long-term strategy, foster the development of new HMR models and observational archives for the study of severe hydrometeorological events, promote the execution and analysis of high-end simulations and analysis, and support the dissemination of predictive models as decision analysis tools for disaster management.

Scientific case

This scientific case is built around a set of three experiment suites, which are able, on one side, to address the challenges of HMR in forecasting severe hydrometeorological events and, on the other side, to provide laboratories with the integration of new ICT infrastructure and development of new working practices.

These three different experiment suites, as depicted in Figure 1, are part of the so-called hydrometeorological forecasting chain, whose end result is a prediction of a hydrological quantity such as river run-off and water level, with a large variety of models and data sources feed into the prediction.




Fig. 1 - HMR Experiment Suites.

The layers, or experimental suites, of a forecasting model are:

The rainfall layer

The rainfall layer combines an ensemble of high-resolution Numerical Weather Prediction multi-models with stochastic downscaling algorithms⁷ to produce quantitative rainfall predictions for severe meteorological events. Its accomplishment requires the integration of a suite of hydrometeorological data and models and will be highly demanding of HPC resources.

The structure of this layer is explained as follows:

- Numerical prediction of intense precipitation events, over complex topography regions, requires the use of high-resolution forecast models, and it is affected by uncertainty because of the inherently limited predictability of the rapidly developing weather phenomena that must be forecast and the sensitivity of the forecasts to initial conditions and the details of the forecast model;
- To deal with this uncertainty, there is active development of ensemble forecasting systems that can be used to make probabilistic precipitation forecasts;
- Ensemble forecasting systems: multiple numerical predictions are conducted using slightly different initial conditions that are all plausible given the past and current set of observations, or measurements. Sometimes the ensemble of forecasts may use different forecast models for different members corresponding to a multi-model ensemble approach;
- The multiple simulations are conducted to account for the two sources of uncertainty in weather forecast models: (1) the errors introduced by chaos or sensitive dependence on the initial conditions; and (2) errors introduced because of imperfections in the model, such as the finite grid spacing.

The discharge layer

The discharge layer, instead, concerns the fusion/combination of rainfall predictions (from the rainfall layer) with corresponding observations, which are input into multiple hydrological models to enable of the production of river discharge predictions. Its accomplishment is highly demanding from conceptual and data management standpoint.

The structure of this layer is explained as follows:

- Prediction of river water levels and floodplain inundation following an extreme precipitation event requires integration of precipitation forecasts with a terrain-based hydrologic and hydraulic model. The reliability and credibility of flood forecasting and warning systems cannot be guaranteed without properly incorporating the sources of uncertainty into the end-to-end forecasting and warning systems;
- The various sources of uncertainty associated with the model outputs can be classified as: model uncertainty, input uncertainty (due to imperfect forecasts of future precipitation, evaporation, slope runoff, etc.), and parameter uncertainty (due to imperfect assessments of model parameters);

⁷ Ranked fourth among HMR hot topics according to a 2011 survey conducted by the EU-funded DRIHMS project (www.drihms.eu).

- This problem becomes critical in small mountain catchments (or the headwaters of bigger river systems) and urban areas, where flood prediction requires the forecast of the precipitation field down to scales of a few square kilometers and tens of minutes;

The water level, flow and impact layer

Finally, the water level, flow and impact layer addresses the execution of hydraulic model compositions in different modes to assess the water levels, flow and impact created by the flood events and to compare them against observations through proper modelling verification metrics. Its accomplishment will be highly demanding on HPC resources.

The structure of this layer is explained as follows:

- Completing the flood forecasting chain that is initiated in the rainfall layer and continued in the discharge layer with the formulation of flooding risk scenarios that are relevant from early-warning and Civil Protection perspectives;
- Undertaking flood risk analysis for the purposes of long-term strategic planning, through investigation of the impacts of climate change on each of the case study sites.

The Scientific Challenge

Severe hydro-meteorological events are increasing in frequency and magnitude. The societal and economic implications of these events (Parodi et al., 2012), including loss of life and property damage, are the prime motivation for this scientific case to contribute towards the following aspects:

- Supporting Civil Protection decision makers with reliable information about where extreme meteorological events are most likely to occur adequately in advance;
- Protecting people and infrastructure from the direct (and propagated) impacts of severe weather;
- Targeting operational rescue activity at areas known to be at high risk.

In this context, short- and medium-term forecasting and management of severe hydro-meteorological events is highly topical and represents an important contribution to the procedures implemented by civil protection authorities. The peculiar nature of severe hydro-meteorological processes occurring in most small- and medium-size catchments in complex topography areas, such as the Mediterranean region, make them difficult to predict.

The predictive ability of these severe hydro-meteorological events can be improved with:

- Coherent and comparable observations at multiple locations;
- Denser observations in time and space;
- Better access to data (also from citizen-scientists, Bedrina et al., 2012);
- Combination of different modelling tools and post-processing tools;
- A 'joined-up science' multidisciplinary perspective.

In recent years, the quantity and complexity of the modelling tools and datasets have increased dramatically through the availability and quality of remote sensing observational data from satellites and ground-based radars, the growing use of ensemble forecasting methods, and the

increasing recognition of the need to understand the entire flood forecasting chain, from observations through to civil protection response.

However, progress in tackling this HMR challenge is slowed by the difficulty in accessing data that is scattered in different archives, different countries, and different formats. Additionally, collaboration is inhibited by the need for complex weather and hydrological models with significant high-performance computing requirements to be realistically beneficial for local, regional and international civil protection authorities.

State of the art

Hydrometeorological science has made strong progress over the last decade as new modelling tools, post-processing methodologies and observational data become increasingly available. At the same time, the importance of accurate hydrometeorological predictions under current and future climates has increased. For example, within Europe, the Mediterranean region has been identified by the Intergovernmental Panel on Climate Change (IPCC) as a potential 'climate hotspot', and a series of international programs, including HyMeX, MEDEX and MAP D-PHASE, are attempting to address the current and future hydrometeorology of the region.

However, a significant barrier to progress is the fact that the data and modelling resources (meteorological models, hydrological models and post-processing tools) required for this effort are scattered among many research institutions, national agencies including weather services and civil protection agencies.

In principle, a huge amount of hydrometeorological data and modelling tools are available, but strong scientific and technical challenges are associated with the enormous volumes, incompatible formats, and diversity of tools.

Currently, the majority of HRM chains are clumsily stitched together. It is often the case that a given subset of models belonging to the three experiment suites of Fig. 1 only fit because somebody worked for many years to 'glue' them together. In this context, adding another model and/or data set, for example, is tedious and hampers progress.

Going beyond the state of the art

The goal of hydrometeorological research is to understand the physical processes underlying hydrometeorology probabilistic forecasting chain, and to develop new and improved tools for the various layers. The daily work of HMR scientists falls generally into two categories: preparatory activities and science activities.

Preparatory activities include: find data; retrieve high-volume data; learn formats and develop readers; extract parameters, identify quality and other flags/constraints; perform filtering/mask, develop analysis and visualisation tools; select the in-house hydrometeorological models (if available); develop *ad hoc* post-processing tools.

While these tasks are necessary, they are often time-consuming, tedious, and generally regarded as diversions from the scientists' primary work.

Science activities include: exploration; initial analysis; select the best data for the final analysis; derive conclusions, and finally write the report/paper. These activities contribute in a direct way to the scientific objectives.

An ideal environment for research would take advantage of technology to minimise the first category of activities, and maximise the second.

Minimising preparatory activities involves addressing the three challenges represented by data, models and workflows. Large data sets, whether from remote sensing instruments such as radar networks, or from ensembles of numerical weather forecasts, need to be easily available. Availability implies easy to locate, easy to obtain the necessary permissions for use, and appropriate tools to read the different formats and meta-data associated with different data types.

Models, both meteorological and hydrological should also be available with a minimum of effort. As with data sets, substantial efforts are required to transport and use models at other institutions, and a better solution may be to provide transparent access to these tools at the facilities where they are maintained. Since some of the models are computationally intensive, this also requires access to high performance computing to be managed alongside the tools.

Workflow becomes a problem with the increasing need to treat the hydrometeorological forecast chain as a whole. Infrastructure is needed to allow output from alternative in tools each layer to be input into the next layer, and compare different configurations, without the need to continually develop new interface and analysis tools.

Maximising science activities can be more than simply freeing up time. If an e-infrastructure for handling large distributed data sets, a plethora of models and flexible workflows is available, new possibilities for collaboration open up. Virtual communities involving HMR and related earth science disciplines can promote interdisciplinary research and open up the research world to a broader cross-section of society.

Thus, an ideal environment for HMR is expected to:

- Allow and improve access to HMR models, tools and data without the constraints of distance, access, usability and scientific barriers;
- Change and adopt the HMR discovery process in terms of computing, simulating, data processing and visualisation;
- Enable a growing number of HMR (and related earth science) users to engage in open, cross-border and cross-discipline collaboration;
- Promote HMR cost-efficiency and related resource utilisation;
- Attract non-specialist users such as members of the public services (civil protection agencies) who will be potential beneficiaries of faster and better HMR results, and citizen-scientists who are interested in learning more about hydro-meteorology.

e-Infrastructure

The forthcoming e-Science environment for HMR will use different computing and data archiving paradigms, such as grid or High Performance Computing (HPC).

By integrating HMR resources, specialists will enter the e-Science environments more easily and at the same time stimulate use by non-specialists, such as 'citizen scientists' interested in HMR and related Earth science disciplines. Figure 2 depicts a possible view of this scenario.




Fig. 2 - Possible vision of HMR e-Science environment.

Figure 3 more detailed view, where the portal provides integrated solutions to manage and exploit the e-Infrastructure at different levels: users, resources, workflows, data, and services.




Fig. 3 - A more detailed view of the HMR e-Science environment.

It is necessary, however, to move from a conceptual view to an architectural view that describe the technical building blocks and how they interact over a grid infrastructure.

The HMR architecture is organised by the well-known layer pattern for implementing the inevitable separation of concerns. The Resource Layer (Fig. 4) contains all physical resources (computing elements, storage elements, instrument elements, sensor networks, communication networks) which are provided by the various partners in the HMR e-Science environment.




Fig. 4 - Resource and middleware layer.

Typically, these resources are provided over European e-infrastructures like PRACE, EGI or the national NGIs, but operated locally by the providers. For accessing them a grid middleware system like the Globus Toolkit, gLite, or UNICORE is required. Some of the resources may also be available in a cloud (public or private). The middleware and the access services are part of the Infrastructure and Middleware Layer.

In order to cope with the inherent heterogeneity of grid resources and grid middleware, an Interoperability Layer assists in transparently accessing the resources via the underlying layers (Fig. 5). The interoperability layer is based on standards like those defined in the Open Grid Compendium of e-Infrastructure requirements for the digital ERA

Forum (OGF⁸), or in the Organisation for the Advancement of Structured Information Standards (OASIS⁹) like the Security Assertion Markup Language (SAML).




Fig. 5 - Interoperability layer.

The HMR related services form the Service Layer, itself separated in the Basic Service Layer and the Compound Service Layer (Fig. 6). The Service Layer typically interacts with the Interoperability Layer to offer transparent access to HMR resources. It may also interact with the Middleware Layer directly.




Fig. 6 - Service layer.

The Basic Service Layer provides all services which are fundamental for the HMR community. Examples are data conversion, model access or task composition. The Compound Service Layer provides more complex services assembled from basic ones. Typical examples in the HMR context are the creation and the management of workflows, the management of model sets, or the visualisation of simulation results.

The Application Layer contains the specific HMR applications for simulations, experiments and e-Learning (Fig. 7). The Application Layer uses (basic or compound) services from the Service Layer. It may, however, also interact directly with the grid infrastructure.

⁸ <http://www.ogf.org/>

⁹ <http://www.oasis-open.org/>



Fig. 7 - Application layer.

The Application Layer is accessible via the Access Layer, which provides the capabilities to authorise access to the HMR applications via various portals or clients and through an authentication interface (Fig. 8).




Fig. 8 - Access layer.

Figure 9 depicts the complete layered architecture including the usage-relations between the layers.




Fig. 9 - The complete architecture.

Current e-Infrastructure activity

Recent efforts in Europe to develop a general European platform for e-Science, implemented with key projects such as EGEE (Enabling Grids for E-science) and EGI-InSPIRE on a European level, SEE-GRID-SCI (SEE-GRID e-Infrastructure for regional eScience) and C3-Grid (Collaborative Climate Community Data and Processing Grid) on regional or national levels, provide an ideal basis for the sharing of complex hydrometeorological data sets and tools.

In the USA, CUAHSI (Consortium of Universities for the Advancement of Hydrologic Science) has developed a services-oriented architecture for time-series data collected at fixed points; that is, the most standard form of hydrologic data. This architecture, known as CUAHSI Hydrologic Information System (CUAHSI HIS), allows for the discovery and compilation of data from multiple sources, such as multiple agencies (including the United States Geological Survey, the National Climate Data Center and the United States Environmental Protection Agency) and university researchers (more than ten universities in the USA have instituted CUAHSI HIS). A

central catalogue now indexes 4.3 billion data points. Other (related) initiatives focus more on the shared development of common tools, models and modelling frameworks.

Again in the USA, Giovanni is a web-based application developed by NASA's Goddard Earth Sciences Data and Information Services Center, that provides a simple and intuitive way to visualise, analyse, and access vast amounts of Earth science remote sensing data without having to download the data. Giovanni is an acronym for the GES-DISC Interactive Online Visualisation ANd aNalysis Infrastructure. From the researcher's point of view, Giovanni is comprised of a number of interfaces, each tailored to meet the needs of specific fields of earth science research. Each interface, known as an instance, provides functions and parameters applicable to that specific area of Earth science.

Potentially all these initiatives provide a useful basis, which may also be applied for the sharing of complex hydrometeorological data sets and tools, and may thus enable a level of international research collaboration that has been impossible in the past in the HMR community.

The substantial investment in e-Science has highlighted further gains to be made in hydrometeorological research. It is clear that the computing power enjoyed by meteorologists could also be applied to hydrologic and hydraulic studies.

At the time of writing this document, the FP7 DRIHM project community is involved in an effort to set up a Virtual Organisation on the EGI resources. The main goals of this effort are summarised in a dedicated EGI – DRIHM collaboration page¹⁰. The stepwise refinements of goals towards a feasible effort are summarised by EGI – DRIHM Roadmap¹¹ that starts from DRIHM requirement to target a timely agenda that should lead to the availability of a set of basic services that will permit to deploy DRIHM domain application services, i.e. HMR models, on the EGI resources in production mode.

Future e-Infrastructure challenges

Despite the fact that EGI resources would allow the execution of meaningful experiments, a full DRIHM production environment should go beyond EGI current capabilities and combine High Throughput Computing resources with medium- and high-end High Performance Computing resources. At the same time, data movement will be a critical aspect of any DRIHM workflow execution. The overall DRIHM vision requires the ability to run workflows that combine models belonging to different domains as exemplified in Figure 10.

¹⁰ <https://wiki.egi.eu/wiki/EGI-DRIHM:Collaboration>

¹¹ <https://wiki.egi.eu/wiki/EGI-DRIHM:Roadmap>




Fig. 10 - DRIHM models classified in related layer. A workflow may be composed of one model for each level.

This will require the availability of a DRIHM Distributed Infrastructure (DDI) that will go beyond the grid paradigm as exemplified in Figure 11. It is expected that most of the basic and middleware services are available from future European Middleware Initiative software repositories. We expect a convergence path between the current Monte Bianco release¹² and the services available from projects related to this scientific case (e.g., DRIHM).




Fig. 11 - A schematic view of DDI.

Future plans

Future plans will include the study of interoperability issues in order to combine the use of grid and non grid resources. Considering the available resources and the DRIHM time frame interoperability will be mostly addressed at the portal level. However development beyond

¹² <http://www.eu-emi.eu/products>

DRIHM should include a more general solution based on middleware interoperability. Steps towards the topic of urgent computing for hydro-meteorological operational activities should be also undertaken (Leong et al., 2013).

References & Bibliography

Bedrina, T., A. Parodi, A. Quarati, A. Clematis (2012) ICT approaches to integrating institutional and non-institutional data services for better understanding of hydro-meteorological phenomena. *Nat. Hazards Earth Syst. Sci*, 12, 1961-1968.

Leong, S. H., A. Frank, D. Kranzlmüller (2013) Leveraging e-Infrastructures for Urgent Computing. *Procedia Computer Science*, 18: 2177-2186.

Parodi, A., G. Boni, L. Ferraris, F. Siccardi, P. Pagliara, E. Trovatore, ... D. Kranzlmüller (2012) The "perfect storm": From across the Atlantic to the hills of Genoa. *Eos, Transactions American Geophysical Union*, 93(24), 225.

Schiffers, M., D. Kranzlmüller, A. Clematis, D. D'Agostino, A. Galizia, A. Quarati, ... A. Tafferner (2011) Towards a Grid Infrastructure for Hydro-Meteorological Research. *Computer Science*, 12: 45-62

The Life Science grid community

Tristan Glatard¹, Silvia Delgado Olabarriaga², Irene Nooren³,
Jan Bot³, Coen Schrijvers³

¹ Université de Lyon, CREATIS, France (biomed)

² Academic Medical Center, University of Amsterdam, NL (vlemed)

³ SURFsara, NL (Isgrid)

DATE (13 August 2013)

✉ life-science-grid-community@googlegroups.com

Overview

Life sciences comprise the fields of science that involve the study of living organisms, such as microorganisms, plants, animals, and human beings. Biomedical research is a subfield of life sciences with the aim of better understanding the mechanisms of disease, how they manifest themselves in detectable ways, and how they can be influenced to treat the patient. The final goal of biomedical research is to improve healthcare with better diagnostics, prognosis, and treatment by means of interventions with drugs, therapy of various types (e.g. radiotherapy), surgery, or changes in life style. Moreover, better understanding of diseases can help prevention and general improvement of health and well being in society.

Modern life science and biomedical research is based on a large variety of data that can be acquired in a non- or minimally-invasive manner. Research in these fields is driven by the understanding of (human) biology by analysing blood, urine, tissues or cultures using technology that is very dynamic and requires on going development of new analysis techniques and infrastructure. Matching a decrease of costs to acquire data, for example images or DNA sequences, typical modern biomedical research experiments are increasingly based on large amounts of data, requiring advanced IT infrastructures for data analysis, interpretation and modelling. Large e-Infrastructures such as EGI are currently exploited to perform the analysis of large datasets, or to run analysis that would take too long to compute on 'regular' infrastructures. Various types of user interfaces are available to facilitate access to these e-infrastructures. For example, science gateways enable the researchers themselves to run predefined analysis methods, and workflow management systems are used by researchers with programming skills.

The Life Science Grid Community (LSGC) is a federation of Virtual Organisations (VOs) operating in the field of life sciences and biomedical research. The three EGI VOs involved in LSGC (biomed, vlemed and Isgrid) offer infrastructure and support for life science and biomedical research based on technologies that require computationally intensive steps for data analysis or modelling.

The **biomed VO** was formed in 2004 during the EGEE projects as a catch-all VO for life sciences. It is one of the pioneering VOs to exploit such a large distributed and open infrastructure. Any scientist affiliated to a University or academic research group having life science activities can join the VO. Companies may also join biomed, provided that they use the resources for non-commercial activities.

The **vlemed VO** was formed during the VL-e project (Virtual Laboratory for e-Science, 2004-2008, NL). It initially included members from various Dutch organisations, but since 2008 it mostly includes members from the Academic Medical Centre of the University of Amsterdam and its direct collaborators.

The **Isgrid VO** was formed in 2008 in The Netherlands to enable the use of the infrastructure by life scientists. The e-BioGrid project¹³ (2010-2012), financially supported by BiG Grid (the Netherlands), provided life science specific support to help users to exploit the national infrastructure. This led to a growth of users with expertise on using pilot job frameworks and grid middleware, as well as using workflows on the Life Science Grid. Specific national life science support is continued from January 2013 within SURFsara, as part of the SURF national e-infrastructure organisation.

Scientific case

The scientific challenge

The promise of extracting objective information to characterise diseases (e.g. biomarkers) even before they become symptomatic, and using this information to produce patient-specific treatments (drug development), makes modern life science and biomedical research a scientific field of uttermost relevance for society. With aging of the population worldwide, and the increasing pressure on healthcare costs, it is urgent to use such information well to enable efficient treatment at low cost. This research requires the analysis, synthesis, integration and interpretation of large, and ever increasing, amounts of data. For example, one of the ultimate goals of human genetics is to identify the genes associated with a particular disease, to investigate how they affect the complete chain of biological processes that cause the disease, and to develop targeted drugs or prevention actions to alter the course of these processes. The challenge to turn this vision into action requires years of research, building on data collection, analysis, in-vitro and in-vivo experiments and collaborations.

Biomedical research is therefore both data-intensive (vast amounts of data, heterogeneous, distributed) and compute-intensive (sophisticated analysis and simulation methods). The VOs in the Life Science Grid Community have been offering a scalable infrastructure for scientists to target these challenges.

Important infrastructure requirements include storage, controlled access to, and transfer of, large datasets, as well as compute scalability, parallelisation and flexibility. The requirements for computing time differ considerably among users, from 300 up to 3 million CPU hours. Computational demands within the technology areas are expected to grow quickly.

¹³ eBioGrid project: e-infrastructure for life sciences, website: www.e-biogrid.nl

User experiences with infrastructure depend on the expertise of the user. A biologist may prefer an easy-to-use web front-end instead of a workflow engine, whereas a bioinformatician may prefer the latter. For workflow users, data provenance is an important topic. Other requirements include easy authentication and authorisation for sharing the infrastructure and (secure) data. Because biomedical research handles privacy-sensitive data, security and trust are very important requirements as well.

One of the use cases on the Isgrid VO involves pattern recognition and machine learning (Ghermann et al. 2013, de Ridder et al. 2013). The work investigates the relations between virus insertions in DNA and gene expression to better understand the influences of these mutations on cancer. The mouse models being used here will lead to better understanding of the development of cancer in humans. The researchers typically run a number of algorithms across a variety of datasets, requiring many millions of jobs which all have a limited input set and very small outputs (in the KB range). These outputs are saved to the pilot job system itself (in this case PiCaS), as this data cannot efficiently be stored on a storage resource manager. Run times for these applications are in the order of a few hundred thousand core hours per experiment.

Another Isgrid VO example is research that focuses on genetic profiling of the Dutch indigenous population. The Dutch Biobank collaboration BBMRI-NL has initiated the extensive Rainbow Project 'Genome of the Netherlands' (GoNL) because it offers unique opportunities for science and for the development of new treatments and diagnostic techniques. The current grid applications mostly focus on data imputation: estimating genetic variants based on data collected in genomic-wide association studies combined with whole genome next generation sequencing (NGS) datasets. These applications are very data and computationally intense. They use a pilot job framework developed within the group itself called RITE.

Other applications that we encounter often are the analyses of RNAseq data, image analysis and regular sequence comparisons of cancer and cardiovascular disease patients, in order to enhance treatment.

State of the art

A large amount of data acquisition devices is already available in hospitals and research organisations. Figure 1, for example, shows the growth of the number of Next Generation Sequencing (NGS) scanners installed worldwide. Another remarkable example is the area of neuroscience, which will also produce and consume a huge amount of data about the brain in the coming years as result of the two large projects that were recently started in the US and Europe: the BRAIN initiative¹⁴ and the Human Brain Project¹⁵.

¹⁴ <http://www.whitehouse.gov/the-press-office/2013/04/02/fact-sheet-brain-initiative>

¹⁵ <http://www.humanbrainproject.eu>




Fig. 1 - Next Generation Genomics: world map of high-throughput sequencers (<http://omicsmaps.com>).

In contrast to this prospect, researchers typically perform data analysis on laptops, local servers or small clusters. This limits the scale of the computations they can perform. Storage is already a serious bottleneck for the few data servers in place. Large experiments can be performed on public e-infrastructures, such as EGI, but significant effort is still required to organise the computation, data transport and execution on a distributed infrastructure. Moreover, due to privacy regulations special measures are required to use public infrastructures for this type of data. Additionally, exploiting a truly distributed infrastructure that spans the administrative domain of organisations and countries is still technically very difficult, in particular due to the opaqueness of (most) components, which makes troubleshooting a heroic task. Specialised infrastructures with adequate interfaces that expose the necessary information for their inclusion into a programmable environment are needed for more effective exploitation of advanced IT services required to address the scientific challenges in biomedical research more effectively.

biomed is a catch-all VO for life sciences, serving a broad range of scientific challenges. The VO has been historically structured in bioinformatics, drug discovery and medical imaging activities, which still represent the main topics addressed in biomed. Recent examples of biomed-supported research include an initiative to tackle Avian flu (Breton et al., 2009), a study of reproductive strategies, demography and mutational meltdown (Awad et al., 2012), and medical imaging simulation (Glatard et al., 2013).

vlemed is a small VO that uses exclusively Dutch resources. It targets biomedical researchers of one academic hospital, being used extensively for research in neuroimaging (van Wingen et al., 2012) and DNA sequencing (de Vries et al., 2012, Van Houdt et al., 2012).

The members of the **Isgrid VO** are characterised by their self-sufficiency: most of the users compose and submit their jobs manually through the grid middleware using some kind of pilot job framework (e.g. ToPoS, RITE and PiCaS) to get their work done. Some of these users act as enablers, providing scripts and support to other users that would not be able to use the grid without their help. The applications that these users run on the grid are very diverse and correspond to the research interest of the group in which these people work. Examples include pattern recognition and machine learning (Ghermann et al., 2013; de Ridder et al., 2013) and gene imputation. Within the Isgrid VO there is also an increasing number of users in the green and animal life science area, requiring data and compute infrastructure for genome analysis of animals and plants. Currently training is offered to this community to enable their big data research and analyses.

Going beyond the state of the art

Data repositories with adequate data preservation and lifecycle management procedures will be put in place. In addition, new methods will be developed to interpret the vast amount of data and information that is becoming available to understand disease, for example, to extract patterns from the data, generate and test new hypotheses using computational simulation, and integration of knowledge across scales such as space, time, population groups, cells and organisms.

e-Infrastructure

Current e-Infrastructure activity

Figure 2 shows the distribution of consumed CPU time among life science VOs from January 2010 to April 2013. The LSGC VOs – biomed, Isgrid and vlemed – account for almost 85% of the CPU consumption in the field of life sciences.

The biomed VO currently has 300 registered users, of which approximately 150 can be considered active. One of these users is a robot certificate used by a portal with more than 400 registered users. Between January 2010 and April 2013, biomed has consumed 31 million CPU hours, which is roughly equivalent to a 1070-node cluster used at 100% capacity. Most of the resources are consumed by a few users: for instance, in 2012, the five most active users were responsible for 70% of the CPU time consumed by biomed. Tools and services used to access VO resources are quite heterogeneous. Storage is mainly used through Storage Elements (mainly DPM) and the LCG File Catalogue (LFC). Since the Summer of 2012, DIRAC is the recommended solution to access computing resources. A DIRAC service is provided by the French NGI to all biomed users. A few users still use direct submission to glite WMS, mainly due to the lack of a Java API for DIRAC. Figure 3 shows the use of DIRAC in the VO between June 2012 and February 2013. Technical teams have been organised since 2010 to monitor resources and address technical issues met by users in the biomed VO (Michel et al., 2012). A total of 273 GGUS tickets were submitted by these teams during the last year.

Vlmed's membership is tightly controlled due to security policy required for handling privacy sensitive data. The resources are provided by the Dutch NGI (NGI_NL); additionally, small local clusters at the hospital are available and connected to the vlmed services platform. The VO develops and operates a science gateway (Shahand et al., 2012) based on services for workflow management based on MOTEUR and WS-PGrade, data transport between local and grid storage, and monitoring. The gateway is currently used by researchers with two profiles: biomedical researchers use the preconfigured applications from customised interfaces of the science gateway (currently 60 registered users), and some more advanced users adopt the workflow development interface (around 5 active users).

The Isgrid VO has approximately 140 users, of which half are active users. The VO Isgrid is directly (but not exclusively) related to a dedicated infrastructure in the Netherlands, called the Life Science Grid (LSG).




Fig. 2 - Repartition of the CPU time consumed by life science VOs between January 2010 and April 2013 (total: 51.9 million hours). Extracted from <http://accounting.egi.eu>




Fig. 3 - Usage of DIRAC in the biomed VO (June 2012 - February 2013)

The LSG consists of 12 relatively small gLite-clusters distributed in the Netherlands, which are placed inside the academic hospitals and other life sciences research centres in the country. The LSG is managed remotely by SURFsara, and it currently provides 1,280 cores and 0.5 PB of storage (as of August 2013). Apart from the academic network, all LSG clusters are also connected by a dedicated pool of dynamic light-paths that may be activated on demand.

The dedicated LSG hardware is available also for vlemed and the Dutch bbmri.nl VO. Besides the dedicated LSG hardware, the Isgrid VO has also access to all other resources within NGI_NL, for large-scale production runs and to accommodate the heavily-peaked usage patterns that are typical in this VO. The existence of dedicated resources for a specific community seems contradictory to the general idea of efficiency and resource sharing that is offered by grid computing, the LSG, however, is based on a well-considered concept. The locality of clusters within the academic centres and the private connectivity provided by the dynamic light-paths offers a good solution for certain legal or privacy concerns. In addition, the LSG accommodates alternative hardware configurations driven by specific requirements of life sciences: worker nodes consist of 64-core machines with 256 GB memory and 6 TB scratch space, on which a single grid job can in principle claim a large amount of resources (or the complete worker node if needed). Such hardware configurations are rare in standard grid environments, but indispensable for many applications in life sciences. Last but not least, the LSG alleviates the steep learning curve that most new life scientists on the grid face by allowing a more liberal policy with regard to local cluster usage, facilitating debugging, small workloads, and local tailoring of software stacks. As such, LSG creates an environment where the giant step from a (often windows-based) desktop environment onto the grid can be split into smaller steps (from Linux application, via small cluster-based jobs, to large grid-based production runs). This approach seems to have worked quite well. During the past years almost 80% of all CPU hours spent by the Isgrid VO are run outside the dedicated resources of the LSG, which shows that a significant part of the Isgrid workload consists of large runs that scale beyond the size of the LSG.

Future e-Infrastructure challenges

Although the infrastructure is significantly used, a few issues still limit its diffusion to a broader audience. From an application point of view, the main limitations are:

- Short runs (i.e. less than 20 minutes of cumulative CPU) are highly impacted by various overheads.
- Long runs (i.e. more than a year of cumulative CPU) need much administrative intervention to complete.
- Missing high-level user interfaces that hide the complexity of the infrastructure while still offering enough information for debugging.
- Interoperability among different infrastructures and middleware.
- In the biomed VO jobs cannot use a lot of RAM (i.e. more than 2 GB), and requirements cannot be set on RAM. The same holds for the largest clusters available for vlemed.
- In the biomed VO jobs cannot use a lot of disk (i.e. more than 2 GB) and requirements cannot be set on disk space.

- In the biomed VO, manipulating a lot of small files (more than 200) or a few big files (more than 2 GB) is not reliable and/or is highly impacted by transfer overheads.

From a VO management point of view, the main challenges are:

- VO users management: handling of user registration lifecycle, tracking of portal users using a robot certificate.
- VO operations management: consolidate GOCDB and BDII status, advertise critical resource downtimes, monitor resources availability.
- VO Data Management: cleaning procedures, management of SE free space and SE decommissioning, consistency between SEs and file catalogues.

A VO management tool called VAPOR¹⁶ is being developed to address these challenges. It will complement related services such as VOMS and the EGI operations portal.

Future plans

The biomed VO embodies the early-2000s vision of grid as an infrastructure to transparently deliver computing power and storage for scientific applications worldwide. Due to the important technical effort required for application deployment, VO coordination has focused exclusively on the delivery of technical services, to the detriment of scientific collaboration among VO members: the VO is mostly used for independent scientific projects. Fostering scientific collaboration remains a long-term objective of the VO, which will only be possible when the technical effort required to access and operate the infrastructure is reduced. biomed is now a sustainable organisation based on the volunteer contribution of its main users and resource providers. Currently, it does not critically depend on particular sources of funding or project. From a technical point of view, the main upcoming milestones are the complete migration of VO users to DIRAC for workload and data management, and the release of the VAPOR VO management tool.

The support for the vbiomed VO services and science gateway today mostly comes from FP7 projects that come to an end in October 2014. These services are operated on a voluntary basis by one small group of persons from the e-bioscience research line of the AMC. Internal discussions have been started at the AMC to investigate how to further maintain and support these services for this community in a long term.

Life science e-infrastructure support including that of Life Science Grid is currently funded through SURF. As of January 2013, SARA became part of the SURF family, which provides services to improve the quality of higher education and research. SURFsara is now responsible for the e-Infrastructure, since the BiG Grid project has ended on 31 December 2012.

¹⁶ <http://lsgc.org/en/LSGC:lsgc-vapor>

References & Bibliography

- Awad, D.B., S. Gallina, C. Bonamy (2012) Reproductive strategies, demography and mutational meltdown. *Journées scientifiques mésocentres et France Grilles 2012*.
- Breton, V., A.L. da Costa, P. de Vlieger, Y.M. Kim, L. Maigne, L., R. Reuillon, ... (2009) Innovative in silico approaches to address avian flu using grid technology. *Infectious Disorders-Drug Targets*, 9(3):358-365.
- Glatard, T., C. Lartzien, B. Gibaud, R. Ferreira da Silva, G. Forestier, F. Cervenansky, M. Alessandrini, ... (2013) A Virtual Imaging Platform for multi-modality medical image simulation, *IEEE Transactions on Medical Imaging*, 32(1):110-118
- Gehrmann, T., M. Loog, M.J.T. Reinders, D. de Ridder (2013) Conditional random fields for protein function prediction, *Pattern Recognition in Bioinformatics* (in press)
- Van Houdt, J.K., B.A. Nowakowska, S.B. Sousa, B.D. van Schaik, E. Seuntjens, N. Avonce, A. Sifrim, ... (2012) Heterozygous missense mutations in SMARCA2 cause Nicolaides-Baraitser syndrome. *Nature Genetics*, 44(4), 445-9.
- Michel, F., J. Montagnat, T. Glatard (2012) Technical support for Life Sciences communities on a production grid infrastructure, *HealthGrid 2012*, Amsterdam, The Netherlands, IOS Press.
- Shahand S., M. Santcroos, A.H.C. van Kampen, S.D. Olabarriaga (2012) A Grid-Enabled Gateway for Biomedical Data Analysis. *J Grid Computing*, 10 (4):725-742
- de Vries M, B.B. Oude Munnink, M. Deijns, M. Canuti, S.M. Koekkoek, R. Molenkamp, M. Bakker, ... (2012) Performance of VIDISCA-454 in Feces-Suspensions and Serum. *Viruses*, 4(8):1328-34.
- van Wingen, G.A, E. Geuze, M.W.A. Caan, T. Kozicz, S.D. Olabarriaga, D. Denys, E. Vermetten, G. Fernández (2012) Persistent and reversible consequences of combat stress on the mesofrontal circuit and cognition. *Proceedings of the National Academy of Sciences of the USA*, 109(38): 15508-15513

The e-Infrastructure for EISCAT_3D

Ingrid Mann^{1,2} ✉, Ingemar Häggström¹, Anders Tjulin¹, Mats Nylén^{2,3}

¹EISCAT Scientific Association, Kiruna, Sweden

²Physics Department, Umeå University, Umeå, Sweden

³Swedish National Infrastructure for Computing (SNIC)

August 2013

✉ ingrid.mann@eiscat.se

Overview

EISCAT_3D¹⁷ will be a world-leading international Research Infrastructure, using the incoherent scatter technique to study how the Earth's atmosphere is coupled to space. The EISCAT_3D multi-static radars will be a tool to carry out plasma physics experiments in the natural environment, a novel atmospheric monitoring instrument for climate and space weather studies, and an essential element in multi-instrument campaigns to study the polar ionosphere and magnetosphere. Like the existing EISCAT radars, the new instrument will be a key element for covering the aurora zone within the global network of geospace observatories, but with increased capabilities that will bring EISCAT to the cutting edge of incoherent scatter observations.

EISCAT_3D will consist of a core site with a transmitting and receiving radar arrays and four sites with receiving antenna arrays at about 100 km from the core. The system will operate in a band approximately 30 MHz wide around the centre frequency, which in the baseline design is assumed to be 233 MHz. Each antenna array will consist of a large number of single antennas whose individual signals are digitally controlled and combined to permit complex antenna pattern control. High-speed electronics changes the time delay between the antenna elements to form multiple beams and beam direction can be varied instantaneously.

EISCAT_3D covers the gap between standard atmospheric observations and space measurements, with its typical observation modes ranging up to roughly 1,200 km altitude. The location in Northern mainland Scandinavia is unique for research into the polar atmosphere

¹⁷ Plans for the new EISCAT_3D (<http://www.eiscat3d.se/>) system emerged from within the existing EISCAT Scientific Association and its user community and led to the E3D Design Study for the new instrument that was funded by EU through Framework Programme 6 (2005–2009). EU finances the present E3D Preparatory Phase Project through Framework Programme 7 from October 2010 until September 2014. EISCAT Scientific Association is the project coordinator. Other partners in the project are: in Finland University of Oulu, in Norway University of Tromsø, in the United Kingdom the Science & Technology Facilities Council, in Belgium National Instruments Corporation, and in Sweden Luleå University of Technology, Swedish Institute of Space Physics, the Research Council (Vetenskapsrådet) and the Swedish National Infrastructure for Computing (SNIC).

and for observing phenomena that are linked to the aurora. The region offers an infrastructure of research and university institutions that is unique within the Arctic. It hosts two rocket ranges and a versatile network of other instruments for active and passive observations like other types of radars and radio instrumentations, lidars, magnetometers, and optical instruments. The capabilities of the new system will make EISCAT_3D attractive for an expanding user community.

The e-Infrastructure challenges for EISCAT_3D include connecting to different databases, developing search engines, visualisation systems, and new presentation tools to reach out to the community and to account for the diversity in the research applications. High performance computing capabilities for digital signal processing are needed on each of the sites and they need network connections. EISCAT_3D will be a distributed research infrastructure with sites in the most northern region of Europe, which suggests making use of the locally existing e-Infrastructure centres nearby for storage.

EISCAT_3D will be operated by the existing EISCAT Scientific Association¹⁸ with current members China, Finland, Japan, Norway, Sweden and the United Kingdom, so that at the same time the EISCAT_3D archive needs to offer data to international users.

Scientific case

The scientific challenge

EISCAT_3D makes use of the incoherent scatter process: The charged components of the atmosphere are probed by transmitting a high-power radio wave and by receiving from several locations and analysing, with high accuracy, the back-scattered waves. The spectra of the received radio waves give characteristics of the atmospheric electron population, which in turn also reflect the characteristics of the ion and neutral populations. Measurements from multiple antenna sites in addition provide vector information. This technique being used for several decades, a well-developed theory exists, to derive physical parameters from the measured signals (Stubbe and Hagfors, 1997). The measurements are used for studying the magnetosphere and near-Earth space, the ionosphere and phenomena in the mesosphere and sometimes the lower atmosphere. The same radar instruments are also used to observe variations in refractive index of the neutral atmosphere that generate radar reflection in the troposphere and other objects (hard targets, coherent scatter signals) such as: meteor trails and head echoes, small solar system objects and the solar wind. Another example is the distinct radar echoes that are caused by irregularities in the electron distribution and linked to the presence of charged ice particles, the Polar Mesospheric Summer Echoes, or PMSE (a similar phenomenon observed in winter, the PMWE, is probably caused by somewhat different processes). Measurements with the existing EISCAT radars have enabled discoveries related to, for instance, the outflow of light ions from the ionosphere to the magnetosphere, the physics

¹⁸ EISCAT Scientific Association: is a Swedish non-profit organisation. EISCAT is currently funded and operated by research councils of Norway, Sweden, and Finland and by funding bodies in Japan, China and the United Kingdom (EISCAT Associates).

of charged dust/ice layers in the mesosphere, and the plasma processes induced in the ionosphere with the EISCAT Heater¹⁹ (e.g. Rietveld et al., 2003).

Many open questions about the coupling between the Earth atmosphere and space are related to fundamental processes in plasma physics, on small spatial scales, that are much smaller than the radar beam width. Other processes require the observation of large volumes in space. The aurora, on the other hand varies significantly with time and the smallest time constant of aurora variations is yet not determined. The capability of continuous observations is an advantage for research on ionosphere-atmosphere coupling, on auroral physics, and on planetary waves, tides and winds. Monitoring data would provide statistics for studies of, for example, PMSE and space weather phenomena. For the latter, flexibility to change observation modes would be a great advantage, so that the system could respond to sudden events like enhanced fluxes of particles with high energy reaching the Earth. These different requirements are beyond the capabilities of the existing instruments and are best addressed with the planned multi-site phased array radar EISCAT_3D.

State of the art

Most of the 15 radar systems used worldwide for incoherent scatter measurements (Arecibo, EISCAT UHF, EISCAT ESR, EISCAT VHF, Irkutsk, Jicamarca, Kharkiv, Millstone Hill, Sondrestrom and ALTAIR) have operated for more than 25 years. Except for the Jicamarca system, these radars use antenna dishes with single beams and either limited or slow steering capabilities. Several of the newer radars (MU, EAR and PANSY) are primarily for middle atmosphere research and have limited capabilities for incoherent scatter measurements. The US incoherent scatter radars PFISR and RISR are modern phased-array systems (Fig.1); single-site instruments which have neither multi-beam nor aperture synthesis imaging capabilities.

Going beyond the state of the art

The new multi-static phased-array facility EISCAT_3D will have order-of-magnitude improvements in temporal and spatial resolution, with 3-D imaging capability from the upper atmosphere to well beyond the topside ionosphere. Standard ionospheric data will be generated for an altitude range from 70 to 1,200 km. Because it can operate continuously, it will offer a novel type of monitoring data to modellers for instance for climate models. It offers the perspective of connecting its data to conventional weather model databases (of up to 60 km). It has the capability to instantaneously change its mode of operation, i.e. in case of sudden pre-defined events. Table 1 shows improvements in measuring ionospheric parameters compared to the present mainland radar.

¹⁹ The Heating facility situated next to the present radars, consists of small antenna arrays for transmission of radio waves at 3.8 to 8.0 MHz (<http://www.eiscat.uit.no/heater.html>). It will also be operated with EISCAT_3D.




Fig.1 - This sketch compares bi-static phased array (right) and conventional dish radar (left). (Prepared within the FP7 Preparatory Phase by B. Gustavsson, EISCAT Scientific Association).

Table 1 - Improvement of measurement capabilities with EISCAT_3D

Measurement Type	Integration time ^{a)} t/s		Improvement factor
	Current VHF	EISCAT_3D	
Isotropic parameters ^{b)} 110 km altitude	4.0	0.02	200
Isotropic parameters ^{b)} 300 km altitude	15.0	0.2	85
Vector velocities ^{c)} 110 km altitude	300.0	1	300
Vector velocities ^{c)} 300 km altitude	1700.0	20	80

Notes:

^{a)} The given integration times are estimates made with a numerical package prepared by Virtanen and Orispää within the EISCAT_3D Preparatory Phase (based on Vallinkoski, 1988).

^{b)} The isotropic parameters are electron density, electron and ion temperature.

^{c)} Vector velocities are ion drift velocities. (Prepared within FP7 Preparatory Phase by B. Gustavsson, EISCAT Scientific Association.)

New capabilities resulting from flexibility and the capability of large scale and small scale imaging observations include:

- multi-beam volume imaging
- aperture synthesis imaging
- fine-scale 3D imaging (sub-beam-width)
- simultaneous multiple beams / interlaced beams (sub-pulses in different directions)
- high-resolution polarisation-, phase-, and amplitude-coding instantaneous, adaptive control of beam positions

With the new system it is possible to make use of the emerging new technique of radar interferometry as well as of statistical inversion methods for powerful computer analysis in order to reach a resolution that is smaller than the beam width. The system will be capable of providing uninterrupted measurements. A number of science questions will benefit from continuous operations. They would also help to better support satellite measurements, rocket campaigns and diagnostics with other instruments. The unique feature of volumetric imaging is that it allows broad regions of the ionosphere and upper atmosphere to be mapped on a quasi-simultaneous basis. It also resolves the ambiguity in space and time that is inherent in the current EISCAT observations. This resolution is particularly important during auroral processes where the ionosphere is rapidly changing over the observing radar instrument. The same method of rapid beam steering can also be used to trace down small-scale structures at high resolution. The numerous advances that are expected from the new system are summarised in the EISCAT_3D science case that is prepared within the FP7 Preparatory Phase (Aikio and McCrea, 2012). The user communities in the EISCAT member countries have each developed their own science strategy in-line with their national research area.

The e-Infrastructure

The new system bears a number of e-Infrastructure challenges: it requires reliable computer networks between and within sites, high performance computing capabilities on site, visualisation systems and new presentation tools, and connection to different databases. The performance of the EISCAT_3D instruments is achieved by digitising the antenna signals in an early stage and modern software-based techniques make a very flexible approach to the signal processing possible, leading to novel types of data products. Utilising these techniques to the fullest will require access to large amount of storage, computing and networking capacity, so this means that a compromise between flexibility and cost must be found. For the on site data processing and preparation of data products the new system can start from the existing infrastructure of the present EISCAT.

Current e-Infrastructure activity

At present EISCAT uses a set of different sophisticated radar codes and incoherent scatter analysis tools, many of them developed in-house and in close contact with the user community (e.g. Lehtinen and Häggström, 1987). EISCAT runs a data archive and is connected to an international database. The association actively participates in projects related to improving data dissemination.

The total volume of the present EISCAT data archive for the 1981 to 2013 observations is 40 TB. The EISCAT data products can be accessed directly through the EISCAT web page or also through the Madrigal system²⁰, an upper atmospheric science database that is used by groups throughout the world. Madrigal is a robust, World Wide Web based system capable of managing and serving archival and real-time data, in a variety of formats, from a wide range of upper atmospheric science instruments and models. Data at each Madrigal site is locally

²⁰ <http://www.openmadrigal.org/>

controlled and can be updated at any time, but shared metadata between Madrigal sites allow searching of all sites at once from any site.

EISCAT at present has three major participations in EC-funded projects related to data products in space weather and environmental research. These target common operations of environmental research (ENVRI²¹), access to observations and modelling and prediction of the near-Earth space environment (ESPAS²²), and the harmonisation of data products of environmental research infrastructures in the EU with their counterparts in the US (COOPEUS²³). As a task of the COOPEUS project, EISCAT and their counterparts from the US incoherent scatter community are in the process of preparing a joint roadmap for data formats. Another project related to space weather studies is in the negotiation phase and the Association actively pursues further project participations.

Future e-Infrastructure challenges

Data flow

The overall data flow of the EISCAT_3D system is shown in Figure 2. The standard incoherent scatter data products are generated directly on site. Parallel to this data flow, a ring buffer (First in/First out, or FIFO, data buffer) will store the full data for roughly two days of continuous observations, up to 10 PB, from which other data products are derived. (Reduced storage in the initial stage of the project should permit storage of 1,000 seconds of full data volume for specific events. This time span is sufficient for following e.g. a rocket campaign, satellite passage or an auroral breakup, most other events are shorter, typically of order 60s). The standard data products are transferred to the central data archive with scalable petabyte storage and sufficient computing capacity for standard operations, image processing and search engines. Multi-static parameters are generated at the operation centre and from their transferred to the data centre. The data centre also monitors the production of the standard data products from the different sites and generates further data products. Table 3 illustrates how the different steps of the data handling are distributed between the sites, the operation centre and the data centre. A detailed plan will be made taking into account the actual locations of the sites, the operation centre and the data centre, network capacities and storage.

On site data handling

One of the main challenges, particularly in phased array radars, is on site data handling (Fig.2) because the antenna signals are digitised at an early stage at each antenna data is generated with a 120 MB/s rate. The data is processed and reduced on site through different steps generating in sequence: beam-formed data, correlated products, and fitted plasma parameter data (upper right in Figure 2). Some of the work steps are the same as in the present system. Data from this chain is also stored in a 10 PB ring buffer, from which more detailed products are derived, like observations of hard targets / coherent scatter, including meteors, PMSE,

²¹ <http://envri.eu/>

²² <http://www.espas-fp7.eu/>

²³ <http://www.coopeus.eu/>

PMWE or space debris. Normally the ring-buffer data stays within EISCAT and some small amounts are permanently stored in order to follow the performance of the system and for further development work. The ring buffer data is also available to users to generate their own defined data products. Users who wish to use the ring buffer data should apply specifically for that. The on site computing requires equipment of various types, with different life-times to be further specified during the preparation phase and in more detail during commissioning of the first built core array. During the lifetime of EISCAT_3D the actual available capacities for storage, computing and networking are likely to increase and offer room for improvement and for reducing operation costs.




Fig.2 - The expected maximum data flow for EISCAT_3D. Figure prepared within the FP7 Preparatory Phase by A. Tjulin, EISCAT Scientific Association.

Network

The estimated data rates in local networks at the active site run from 1 GB/s to 10 GB/s. Similar capacity is needed to connect the sites through dedicated high-speed network links. Downloading the full data is not time critical, but operations require real-time information about certain pre-defined events to be sent from the sites to the operation centre and a real-time link from the operation centre to the sites to set the mode of radar operation on with immediate action.

Operation centre

The radar observations will be run from an operation centre staffed with radar operations and software experts. The operation centre is located away from the sites, but staff should ideally be able to reach at least the core site for maintenance work to be carried out within one workday. The major fraction of workforce during the implementation of EISCAT_3D will be going into software development for radar operation, signal processing and software to prepare the different data products and this work will be carried out in the operation centre. The operation centre uses the on site computing capacities at the radar sites as well as the data centre. Here also for the fully operational system new radar operation modes will be developed and the standard data products that are generated at the sites are monitored and the multi-static data products are generated.

The data centre

The fully operational 5-site system will generate 40 PB/year in 2022. For the initial stage we expect several PB per year for storage. The EISCAT_3D data centre is expected to be located in one or in a few of the existing high performance computing centres in the North of Europe. Establishing external mirror sites might be relevant, since this is an international research infrastructure. The data centre will provide short-term storage for analysis and long term archiving of the data that arrives from the operation centre, with the associated metadata.

Users will access the data through a data portal and can use provided computing facilities for standard analysis routines. The archive will contain well-defined data products that are negotiated from time to time within the association. Users who plan to generate special data products with the help of high performance computing need to apply for a scientific computing project with their national computer centre. Users who plan to use selected large-volume raw data are expected provide their own storage funded through individual research grants.

Table 3 - Steps within the data flow

Operations on 5 Sites	Operation Centre	Data Centre
Beam-forming	Operation	Storage
Correlated products	Generate multistatic products	Search engines
Fitting of plasma parameters	Produce metadata	User-defined analysis
Generate specialised products	Data validation	Visualisation

Figure 3 provides a top-level view of the user interacting with EISCAT_3D through the portal. Blue downward arrows indicate the data flow. The location of the data centre is not decided yet. There are in the northern Scandinavia pre-existing e-Infrastructure centres that are well suited as hosts for the EISCAT_3D data centre. In the case of EISCAT_3D, an appealing possibility is to implement the EISCAT_3D data centre with a distributed solution. The appeal for this solution stems from the geographical distribution of EISCAT_3D in combination with the pre-existing e-Infrastructure in northern Scandinavia. Since, as discussed above, the archive needs to be stored at two independent sites, a distribution on two of the pre-existing e-Compendium of e-Infrastructure requirements for the digital ERA

Infrastructure centres will add very little extra costs (might in fact provide a modest saving) and provide the benefit of complete redundancy (with some extra costs in software complexity). Figure 3b illustrates this for the case with five antenna sites and two participating e-Infrastructure centres. In this case the archived data is redundantly stored on the two participating centres. The solid lines indicate dedicated network links. The dashed green box encloses the distributed data-centre. In this case the portal floats around - providing a single portal for the entire data centre.




Fig.3 - The data central (a) and distributed solution for the EISCAT_3D data centre. The figure is prepared within the FP7 Preparatory Phase by M. Nylén, Physics Department, Umeå University, Sweden and Swedish National Infrastructure for Computing (SNIC).

Data products and databases

The data centre offers access to standard data products, for example ionospheric measurements, middle-atmospheric measurements, PMSE and PMWE observations, and meteor observations. There will be many other data products: some user defined and user produced and some provided directly by EISCAT. From that, it is expected EISCAT will contribute to databases for space weather, satellite projects, meteors and small solar system objects, and climate modelling and meteorology. The existing system of storing incoherent scatter data is currently investigated with the goal of offering the same data products from different instruments (Tjulin et al., 2013).

Initial computing and storage requirements

During the lifetime of EISCAT_3D the actual available capacities for storage, computing and networking are likely to increase and offer room for improvement and for reducing operation costs. According to the recent planning, the computing and storage requirements for the first 5 years of EISCAT_3D, subject to further evaluation, are:

- several PB/year of archive, that needs to be mirrored at an alternate site for data-security and 50 Tflops/s user-defined computing at the Data Centre;
- about 20 PB of shorter term storage and 500 Teraflops/s of computing capacity for post-processing at the Operation Centre.
- about 100 TB RAM storage and 1 PB disk storage, about 600 Tflops/s implemented in FPGA and about 80 Tflops/s other computing for real-time processing on each site.

Future plans

The EISCAT Scientific Association is an existing organisation that operates incoherent scatter radars in the Scandinavian Arctic. EISCAT_3D will replace the radar system that the Association operates on the Scandinavian mainland. The Association is funded and operated by research councils of Norway, Sweden, Finland, Japan, China and the United Kingdom. The European Strategy Forum on Research Infrastructures (ESFRI) selected EISCAT_3D for inclusion in the Roadmap 2008 for Large-Scale European Research Infrastructures for the next 20—30 years. The European Union has funded a Design Study (Wannberg et al., 2010) and is currently funding a Preparatory Phase project for EISCAT_3D, which will finish in 2014. Users in the six EISCAT member countries are preparing for funding applications to finance EISCAT_3D and proposals in Norway and Sweden are already submitted.

The work plan below is based on an estimated funding profile and needs to be adjusted when funding is granted. Provided that adequate funding is available, the implementation of EISCAT_3D will start in January 2014 and continue until December 2021 with a work plan that consists of overlapping Preparation, Construction and Commission phases. This time line is smoothly connected to that of the EISCAT_3D Preparatory Phase Project, so that a preliminary design plan for EISCAT_3D can be prepared and evaluated by November 2014.

The Preparation Phase will address the challenging requirements posed by the Arctic environment, evaluate cost and performance of mass-produced parts, test key deployment processes and start ground preparation. Construction of the system is scheduled to start in the spring of 2016. A first functional sub-unit of the core site with transmit and receive capability (T/R array) will be installed and key capabilities will be tested and validated during the summer of 2016. The antenna arrays can then be deployed in a modular way adjusted also to the funding profile.

The Commissioning Phase will start during the erection of the first working array and comprise of detailed radar coding, development of measurement programs, work on data formats and user access to ensure that the end user can exploit the full capabilities of the radar system. The entire implementation phase will be carried out in close contact to the user community who will agree on the first science observations to be carried out during the different stages of the system and will negotiate the data products.

The major fraction of the workload during the 8-year implementation phase will be going into software development for radar operation, signal processing and software to prepare the different data products. Up to 17 specialists will be working full-time on these tasks during implementation. During regular operation at the end of the implementation phase, around ten out of 30 EISCAT staff will directly work on data production and processing including the preparation for final storage of the standard data products.

Activities with respect to data handling and distribution for EISCAT_3D are integral parts of the on going FP7 Preparatory Phase, with the Swedish National Infrastructure for Computing (SNIC) as the responsible participant. SNIC serves as a contact to the e-Infrastructure stakeholders in the Nordic countries who will monitor the progress of the project: NeIC, CSC, UNINETT Sigma, DeIC, SUNET and NORDUnet.

References

- Aikio, A., I. McCrea (2012) The EISCAT_3D Science case, prepared within the EISCAT_3D Preparatory Phase Project <https://www.eiscat3d.se/project/fp7/science-case>
- Lehtinen, M.S., I. Häggström (1987) A new modulation principle for incoherent scatter measurements, *Radio Science* 22: 625-634.
- Rietveld, M.T., M.J. Kosch, N.F. Blagoveshchenskaya, V.A. Kornienko, T.B. Leyser, T.K. Yeoman (2003) Ionospheric electron heating, optical emissions, and striations induced by powerful HF radio waves at high latitudes: Aspect angle dependence, *Journ. Geophys. Res*, 108:1141.
- Stubbe, P., T. Hagfors (1997) The Earth's Ionosphere: a Wall-Less Plasma Laboratory, *Surveys in Geophysics* 18: 57-127.
- Tjulin, A., P. Erickson, I. Häggström (2013) Gap Analysis Report on Incoherent Scatter Data. COOPEUS Deliverable 2.2 (<http://www.coopeus.eu/documents/>)
- Vallinkoski, M. (1988) *Journal of Atmospheric and Terrestrial Physics*, 50: 839-851.
- Wannberg, U.G., H. Andersson, R. Behlke, V. Belyey, P. Bergqvist, J. Borg, A. Brekke, J. Delsing, ... (2010) EISCAT_3D - A next-generation European radar system for upper atmosphere and geospace research, *Radio Science Bulletin*, 332: 75-88.

Transparently scale next-generation high throughput biobank analysis between institutes, computational clusters and grids

M.A. Swertz ✉, Heorhiy Byelas, Pieter Neerincx, Freerk van Dijk, Martijn Dijkstra, David van Enckevort

For the BBMRI-NL (VO) biobank community

✉ m.a.swertz@gmail.com

Overview

Life sciences have moved rapidly into big data thanks to new parallel methods for gene expression, genome-wide association, proteomics and whole genome DNA sequencing. The scale of these methods is growing faster than predicted by Moore's law. This has introduced new challenges and needs for methods for specifying computation protocols for example Next-Generation Sequencing (NGS) and genome-wide association study (GWAS) imputation analyses and running these on a large scale is a complicated task, due to the many steps involved, long runtimes, heterogeneous computational resources and large files. The process becomes error-prone when dealing with hundreds of samples, such as in genomic analysis facilities, if it is performed without an integrated workflow framework and data management system.

From recent projects we learnt that bioinformaticians do not want to invest much time in learning advanced grid or cluster scheduling tools, preferring to concentrate on their analyses, be closer to old-fashion shell scripts that they can fully control and have some automatic mechanisms taking care of all submission and monitoring details. Therefore we developed with BigGrid and eBioGrid²⁴ a lightweight workflow declaration and execution system to address these needs, built on top of the MOLGENIS framework for data tracking. Now, have scaled running NGS and imputation analyses from computational clusters to grids, in particular, in the nation-wide 'Genome of the Netherlands' project (GoNL, 700TB of data and about 400,000 computing hours).

Scientific case

The scientific challenge

High-throughput analysis methods have created exciting new possibilities for unraveling genotype-to-phenotype relationships. However, these experiments are heavily dependent on large computational analysis for their success. For instance, next generation sequencing

²⁴ eBioGrid project: e-infrastructure for life sciences: www.e-biogrid.nl, part of BiG Grid: www.biggrid.nl

analyses typically involve about 30 computational steps such as alignment, duplicate marking, single-nucleotide polymorphism (SNP) calling, annotation and many re-indexing and quality control steps (DePristo and Daly, 2011).

Similarly, GWAS data typically requires batching for imputation. Genomic analysis facilities typically face running many different versions of such computational pipelines on hundreds or thousands of samples. This quickly becomes a nightmare of data file logistics (raw, intermediate, result, quality controls and log data) and a computational scheduling problem (large, small, short, long jobs that may have error states and different computational back-ends).

Furthermore, different cluster and grid middleware do not provide all necessary operations to execute bio-workflows. We address these challenges in a practical solution with the software system that combines computational and data management for routine running of large bioinformatics analyses.

State of the art

In our current state, we run different workflows in a unified way and make workflow adaptation to different back-ends, such as clusters and grids environments more standard and easier. This has enabled us to develop local but then run analysis on our grid-vo BBMRI-NL.

For example, the GoNL project²⁵, we established a map of genetic variation in the Dutch population by whole-genome sequencing 769 individuals to 12X coverage (Fig.1).

To process the raw sequence data, 2,314 lanes totalling up to 50TB, we developed a pipeline which aligns each lane to the human reference genome build hg19 using BWA. Using the Genome Analysis Toolkit (GATK) (DePristo and Daly, 2011) duplicate reads are removed, re-alignment performed around known insertions and deletions from the 1,000 Genomes Project pilot, and base quality scores recalibrated. Lanes are then merged per family and re-alignment is performed around insertions and deletions found in the sequence data and in the 1,000 Genomes Phase 1 data (Li and Durbin, 2010). SNP discovery and genotyping is done using the GATK Unified Genotyper across all individuals simultaneously. The initial calls are filtered using GATK Variant Quality Score Recalibration (VQSR), followed by Phase By Transmission to compute the most likely genotype for each individual. Resulting in approximately 50,000 analysis jobs, 250,000 CPU hours and 550TB output data.

²⁵ GoNL, Genome of the Netherlands (<http://www.nlgenome.nl>)




Fig.1 - Setup of GoNL project. 250 father, mother, child trios from 5 biobanks were sequenced at 12X coverage. To increase sensitivity 8 dizygotic and 11 monozygotic twin pairs were added to the analysis.

To address the need of improvement of genotype data we created an imputation pipeline which performs QC on studydata and phases the genotypes using SHAPEIT2 (Delaneau and Zagury, 2013). Finally the imputation over genome chunks of 5MB is performed using IMPUTE2 (Howie and Donnelly, 2009). Using GoNL release 4 (499 unrelated individuals) and 1,000 Genomes phase 1 integrated version 3 (1,092 individuals) respectively as reference panel. In total around 40,000 samples, leading to 150,000 CPU hours and 100TB output, were imputed.

We use MOLGENIS compute to implement these pipelines, keep track and distribute all analysis jobs in parallel on our PBS compute cluster and the national life science grid.

Going beyond the state of the art

The Dutch University Medical Centers (UMCs) united in BBMRI-NL national biobank consortium and are embarking on many more high throughput analyses like Genome of the Netherlands, focusing on transcriptome, metabolome, proteome or methylome (see²⁶). In addition, the UMCs run other big projects, such as CTMM or TIFN, that also produce similar large data.

While most institutes have local computer clusters, for the bigger and collaborative projects they obviously need to scale their work to the (inter)national compute grids. In this use case we aim to develop seamless protocols where researchers can transparently move data and analysis pipelines from their local clusters up to the grid. We expect being able to built on best practices developed with BigGrid and eBioGrid. To this end a BBMRI-NL working group for 'big data ICT' has been established in collaboration with SURFsara and academic institutions such as UMCG, LUMC, UMCU, VU, EMC, RUMCN, and we requested light path networks to provide bandwidth to enable these integrations to work. Possible ingredients may include expanding of

²⁶ <http://bbmri.nl/nl-nl/activiteiten>

grid virtual organisation (VO) to the local UMC clusters and introduction of cloud methods there where the software diversity is beyond what the grid can accommodate.

e-Infrastructure

Current e-Infrastructure activity

In our current e-gateway solution (Byelas and Swertz, 2013), we (1) re-use best practices from the Molgenis database generator (Swertz and Jansen, 2010) to generate experiment data and workflow models for our users and (2) fulfill a direct connection to several cluster/grid infrastructures for bioinformatics.




Fig.2 - Grid/cluster infrastructure setup

On a daily basis, users from BBMRI-NL run up to 2,000 analysis jobs on the GCC cluster alone. In addition we run many jobs on the national grid.

Future e-Infrastructure challenges

We see many challenges:

- Since we use open-source software, which sometimes has quite some dependencies, we would like to deploy software easier. Unifying backends/OS, maybe by some sort of visualisation?
- This would also enable us to easily use the grid as failover site, might we have hardware issues locally.
- We still need to request separate permission to be able to login to one of the front nodes (example: ui.grid.sara.nl). Something like a transparent uniform authentication (example: surfconnext) would make life much easier.
- We have lots of big data, large storage enables us to store it all.
- The big data needs to be copied (high I/O) to several other storage and compute clusters, a good network, for example larger bandwidth (more lightpaths) would help. We now have TarGet²⁷ storage which can save much data, we still have a 1GB uplink to the outside world though.

²⁷ The Target project, ICT infrastructure for research: the project web-page: <http://www.rug.nl/target>

- Good security is a must. We work with DNA data obtained from individuals, this requires a different level of security compared to astronomers who store pictures. The grid certificates can still be quite easily used by others, maybe do something about this?

Future plans

The BBMRI-NL working group had its first meeting in Spring 2013. The meeting was also attended by people external to BBMRI, so we decided to broaden the scope to University Medical Centers and associated universities, all involved in big data research in human samples.

Timeline

Q2 2013 – Kick off workshop national big data IT working group in human bound research, organised by BBMRI-NL.

Q3 2013 - Pilot projects between UMCG, LUMC, EMC, UMCU and SURFsara as those institutes already have clusters in place. Pilot topics include: porting pipelines across the local clusters and federated grids; exploring use of virtual private cloud technology; exploration of identity federation and protocols for data sharing; (if we get awarded) use of light path technology. These projects will be used to enable BBMRI-NL big data analyses.

Q4 2013 – Presentation of pilot projects 2nd national workshop of provide report to the BBMRI-NL and NFU; contribute to BBMRI-NL project request for national infrastructure roadmap.

2014 – Implementations of selected pilot features into production into UMCs user communities.

References & Bibliography

- DePristo M., M. Daly (2011) A framework for variation discovery and genotyping using next-generation dna sequencing data. *Nature Genetics*, 43(5): 491–498.
- Byelas, H.V., M.A. Swertz (2013) Scaling bio-analysis from computational clusters to grids. In: *Proceedings of the IWSG-2013 conference*.
- Swertz, M., R. Jansen (2010). The molgenis toolkit: rapid prototyping of biosoftware at the push of a button. *BMC Bioinformatics*, 11:12.
- Li H., R. Durbin (2010) Fast and accurate long-read alignment with Burrows-Wheeler Transform. *Bioinformatics*.
- Delaneau O., J.-F. Zagury (2013) Improved whole-chromosome phasing for disease and population genetic studies. *Nature methods*, 10: 5–6
- Howie, B.N., P. Donnelly (2009) A flexible and accurate genotype imputation method for the next generation of genome-wide association studies. *PLoS Genetics*, 5.

WeST: a Worldwide e-Infrastructure for Structural Biology

A.M.J.J. Bonvin✉, M. Verlato, H. Schwalbe and A. Rosato

On behalf of the WeNMR collaboration

August 2013

✉ a.m.j.j.bonvin@uu.nl

Overview

Structural biology is concerned with the determination of the three-dimensional structures of bio-macromolecules and their complexes, contributing to society by supporting a range of applications including drug design, crop improvement, and engineering of enzymes of industrial significance. Structural biology, encompassing both experimental and computational approaches and bioinformatics, is facing challenges that require the integration of data from various methods, gaining access to sufficient computational resources, developing off-the-shelves solutions for non-expert and dealing with large data sets. There is thus a need for e-Science solutions to provide not only high profile projects but also the long tail of research with e-Infrastructure solutions for open data sharing, virtual research environments and computational resources.

Along those lines, the three-year FP7 project WeNMR²⁸ (Wassenaar et al., 2012) has set up to optimise and extend the use of the Nuclear Magnetic Resonance (NMR)²⁹ and Small-Angle X-ray Scattering (SAXS)³⁰ techniques to determine the structure and properties of proteins and other medically important molecules. The project has built a platform integrating services and streamlining the computational approaches necessary for data analysis and structural modelling, making efficient use of the European Grid Infrastructure (EGI).

WeNMR is the largest life science Virtual Research Community (VRC)³¹ supported by the EGI. It brings together a diverse group of stakeholders and has also tight connections with various European Research Infrastructure projects (e.g. BioNMR) and the ESFRI Instruct project. In addition to research communities, WeNMR also reaches out to the e-Infrastructure resource

²⁸ <http://www.wenmr.eu>; contract number RI-261571

²⁹ BioNMR animation about structural biology and NMR: www.youtube.com/watch?v=IQVvtynnjMM

³⁰ WeNMR animation explaining SAXS: <http://www.youtube.com/watch?v=2HS45Odxs8>

³¹ WeNMR animation presenting the WeNMR VRC: <http://www.youtube.com/watch?v=JUoGdfIVyNI>

providers at regional, national (NGIs) and European (EGI) levels and even worldwide (e.g. via connections to Open Science Grid and XSEDE in the US) levels and to researchers in industry.

In the future, we are keen to expand this community, both geographically and scientifically by widening our activities to encompass a large fraction of the worldwide structural biology community beyond just NMR and SAXS (i.e. by transitioning from WeNMR to WeST).

Scientific case

The scientific challenge

Proteins and their intricate network of interactions are the mainstay of any cellular process. Dissecting their interaction networks at atomic detail is therefore invaluable, as this will pave the route to a mechanistic understanding of biological function and provide the first essential step toward the development of new drugs. Atomic detail (high-resolution) information about structure and dynamics of biomolecules and their complexes is typically acquired by classical experimental methods such as x-ray crystallography and NMR spectroscopy. Compared to other structural biology methods, these provide the highest spatial resolution. They are, however, faced with many challenges, especially when the macromolecular systems under study become very large, comprise flexible or unstructured regions, exist in very tiny amounts, are membrane-associated, or when their constituents interact only transiently. In order to face these challenges, structural biologists often resort to using different types of biochemical and biophysical experiments. Often, however, the collected data are rather sparse and/or of limited information content. As a consequence, researchers have to rely increasingly on multiple different experimental and computational techniques, share data of various natures and collaborate in order to solve a specific problem, which is at the core of integrative structural biology.

In order to facilitate both the integration of techniques and the collaborative efforts which will be required to tackle the health challenges, there is a strong need for off-the-shelves e-Science solutions to provide not only high profile projects but also the long tail of research with e-Infrastructure solutions for open data sharing, user-friendly access to complex software solutions and computational resources, all of which gathered into a virtual research environment, in order to boost the research output.




Fig. 1 - e-Science solutions for integrative structural biology.

State of the art

The WeNMR project has set the first steps toward providing e-Science solutions for integrative structural biology by bringing together the NMR and SAXS research communities. To facilitate the use of NMR spectroscopy and SAXS in life sciences, the WeNMR consortium has established standard computational workflows and services through easy-to-use web interfaces. Thus far, a number of programs often used in structural biology have been made available through application portals (29 to date) that make efficient use of the European Grid Infrastructure (EGI) (see below). With over 570 registered users and a steady growth rate, WeNMR is currently the largest Virtual Organisation in life sciences, gathering users from 42 countries worldwide (>35% of users from outside Europe). The computational tools have been used so far mainly for NMR-based structural biology with SAXS portals recently been put into production. Since the beginning of the project, more than 65 peer-reviewed articles have been published by consortium members and external users in high-ranking journals for the field. It is mainly the user-friendly access to software solutions and the computational resources of the grid that attract users, together with the excellent support and training offered by the scientists of the project.

Going beyond the state of the art

With the advent of integrative structural biology, it is clear that no single researcher will possess all the necessary expertise and resources to tackle the scientific challenges ahead of us in the study of the large and complex biomolecular machines playing major roles for health (e.g. drug design), food (e.g. crop improvement) and engineering of enzymes of industrial significance. One of the challenge for e-[Life]Science will therefore be to provide an integrated virtual research environment to serve a growing number of communities that extend well

beyond the current NMR and SAXS communities served by WeNMR, while at the same time ensuring a high-quality of the already existing services. We will have to offer attractive solutions for data sharing, for supporting collaborations across disciplines and offering efficient, transparent and user-friendly software solutions to end-users, building on the most suited computing solutions for each problem (being it distributed grid, cloud or HPC solutions). In terms of e-Infrastructure needs, we foresee, next to an increased CPU demand, a strong need for open data sharing and storage solutions.

The integrated virtual research environment should particularly target the long tail of research, which covers a very large community of scientists, and, as such, should have the largest impact on the European science landscape. We are therefore keen to expand our existing community, both geographically and scientifically by widening our activities to encompass a large fraction of the worldwide structural biology community beyond just NMR and SAXS (i.e. transitioning from WeNMR to WeST).

e-Infrastructure

Current e-Infrastructure activity

The enmr.eu VO of WeNMR is registered with the European Grid Infrastructure Operations Portal as a global VO under the Life Sciences discipline. It is currently supported by grid sites from several countries (28 sites) including within Europe the National Grid Initiatives (NGIs) of Belgium [1,898 CPUs], France [241 CPUs], Germany [29,974 CPU], Italy [9,732 CPUs], the Netherlands [12,804 CPUs], Poland [18,240 CPUs], Portugal [1,248 CPUs], UK [16,287 CPU], and outside Europe, Brazil [312 CPUs], China [1,074 CPUs], Malaysia [344 CPUs], South Africa [208 CPUs], Taiwan [240 CPUs] and the US [23,427], for a total of over 113,000 CPU cores corresponding to a total computing power of 222 MSi2K, and 249.5 TB of disk space available for grid storage, all of this shared with others regional and global VOs. The dedicated computing and storage resources of WeNMR (located in Utrecht, Florence and Frankfurt) correspond to 320 CPUs for a total computing power of 781 kSi2K and 3.92 TB of disk space available for grid storage.

The enmr.eu VO has grown constantly since its establishment and became in 2012 the largest VO in the life sciences domain³², with over 580 members, 35% of which from countries outside Europe. The user distribution per geographic area is shown in Figure 2.

About 540 kSi2K.years of normalised CPU time (4.7 million kSi2K.hours) were consumed in total during last year, corresponding to about 1.8 million jobs. The enmr.eu VO accounts for about 10% of the total life sciences community grid CPU usage, being the second most used VO in life sciences, but the largest in terms of number of users. Figure 3 shows the number of jobs run and the corresponding normalised CPU time, respectively, for each geographic area hosting resources supporting WeNMR in the last year. About 18% of the jobs are executed by the dedicated resources owned by the three NMR labs (hosting 320 CPU-cores), while the remaining are executed on resources offering opportunistic access.

³² operations-portal.egi.eu




Fig.2 - Geographic grid user distribution




Fig.3 - Geographic distribution of grid jobs (left) and CPU time (right).

Future e-Infrastructure challenges

As pointed out in the previous section, 35% of the WeNMR grid users are from non-European countries, with more than 15% from US and 16% from Asia-Pacific area. One of the challenges is to extend the WeNMR e-Infrastructure to include resource providers in those regions willing to support their scientific community who already make use of the WeNMR services for their daily research activity. Ideally, the geographic distribution of resources should follow somewhat the user distribution. This does not mean *per se* that the CPU usage distribution should match the user distribution. To keep offering high quality services and attracting new users, it is important to ensure a high efficiency of jobs and thus target those sites with the most available resources at a given time, irrespective of the user distribution. Sharing access to global resources will have the advantage of load balancing across a larger pool of computational resources and of sharing and promoting specialised local knowledge and technologies.

We have already connections and support from South Africa, Latin America and the Asia-Pacific area. More sites are expected to join thanks to the recent new partnership with the Academia Sinica of Taiwan, which is coordinating the grid activities in the Asia-Pacific area. Further, considering the large number of users from North America, another challenge will be to integrate resources from the OSG (Open Science Grid) and XSEDE (Extreme Science and Engineering Discovery Environment). We are collaborating since 2010 with the SBGrid Consortium in the US³³, a global consortium of 220 structural biology groups and an OSG Virtual Organisation. A joint task force set up in 2011 has led to achieve a positive proof of concept of interoperability between WeNMR/EGI and SBGrid/OSG grid infrastructures by enabling transparent WeNMR job submission, management, monitoring and accounting. Moving from the proof of concept to a production mode is now the true challenge. A preliminary executive plan has been submitted to EGI and OSG/XSEDE management in March 2013. This will require significant work to adapt the job management scripts lying behind the WeNMR portals, and/or to implement a SAGA layer to decouple the application level from the underlying middleware stack, all of this in a manner fully transparent to the end user. This means that reciprocal submission systems need to be developed and implemented together with more flexible job/CPU brokering and accounting.

Next to worldwide expansion of resources, it will also be important to target the most suited computational resources (grid, cloud or HPC) for a given application. Along those lines we have started investigating the possibility of integrating EMI/UNICORE based resources that are part of PRACE (Partnership for Advanced Computing in Europe), the HPC European e-Infrastructure. This would greatly benefit some of our highly parallel applications (e.g. Gromacs). In terms of CPU requirements, we foresee over the coming 5 to 10 years a linear increase in computing time (current yearly usage is about 550 kSI2K CPU years). This linear rather than exponential increase is justified by the fact that, while new users join our community, old users might move to other positions or projects in their career and stop using our resources.

Finally, next to computing, data should not be forgotten. The data challenge will be to integrate our portals with open data repositories to enable users to directly deposit and share their data with their peers. Ideally, the virtual research environment hosting the various services should directly allow for upload of results to repositories (e.g. via EUDAT³⁴) or the recently launched Zenodo³⁵ repository, a collaboration of OpenAire+ and CERN. Data requirements are thus expected to increase significantly over the next 5 to 10 years, more for storage and open access purposes than for computing requirements since most of our applications are CPU- rather than data-intensive.

Future plans

The WeNMR EU e-Infrastructure project will officially end on November 1st 2013. The partners have, however, committed to make efforts to ensure continuation of the project in the future,

³³ <http://www.sbgrid.org>

³⁴ <http://www.eudat.eu>

³⁵ <http://www.zenodo.org>

in a transition period towards the Horizon 2020 programme. We are actively involved in concertation events toward the Horizon 2020 programme and will hold a networking session for the ICT2013 meeting on “ICT requirements and solutions for Structural biology and Life Sciences”. We hope to set the ground for a future consortium to meet the challenges described in this document.

To ensure the continuation and sustainability of the successful WeNMR VRC, the core partner sites (Frankfurt, Florence, Utrecht, Padua) endeavour to ensure that the WeNMR services will keep running. This means:

- **Maintaining the dedicated hardware**, which will not be a real issue because NMR and structural biology have always been associated with computing and running large facilities and laboratories will always require local computation resources.
- **Maintaining our middleware infrastructure**, which might be more problematic since this goal requires expertise and manpower. For this we anticipate that we might rely more on the respective National Grid Infrastructures of the partners involved in WeNMR for most of the grid services. We will however strive to keep all our grid services active after the end of the project.
- **Keeping the software up to date**: the software developed at the partner sites will be updated and further developed through national or European research funding obtained by these partners. New services might be added on the request of external developers and such portals will be offered for the time they are being used or supported by the developer.
- **Keeping the WeNMR VRC website operating for the user community**. This can be done provided that users, developers and WeNMR partners keep using it by adding contents and answering requests. This is a standard task of software developers in general. The VRC site and its help centre should facilitate this task. Furthermore, with a ‘single sign on mechanism’, we might allow users without certificates to use our services in the future. This should lower the hesitation barrier for potential new users to using our services and keep our community growing to meet the global challenge
- **Offering training and workshops**: this will be done in joint efforts with other projects and organisations, like EMBO or national scientific communities. We intend to ensure the user community keeps growing as usage is the key to sustainability. Will keep promoting the services offered by WeNMR and e-Infrastructure in general at international meetings.

Finally, this transition period toward 2020 will also be the perfect time to interact with- and collect requirements from structural biology communities outside NMR and SAXS to move toward a global, integrated structural biology VRC.

References & Bibliography

Wassenaar, T.A., M. van Dijk, N. Loureiro-Ferreira, G. van der Schot, S.J. de Vries, ... (2012) WeNMR: Structural Biology on the Grid, *J. Grid. Comp.* 10: 743-767.

Computing for the LHC

Jesús Marco✉

IFCA, CSIC

18 November 2013

✉ marco@ifca.unican.es

Overview

The Large Hadron Collider (LHC) at CERN is one of the best examples of international research project where e-Infrastructures, and in particular distributed computing infrastructures such as EGI, have proved to be critical to its success.

The discovery of the Higgs boson in 2012 by the ATLAS and CMS collaborations at LHC, a major outstanding milestone in science, was possible thanks to the excellent performance and large capabilities of the Worldwide LHC Computing Grid (WLCG) project, relying on the EGI infrastructure and services.

By February 2013 LHC completed a very successful data taking period, and entered what is known as Long Shutdown 1 (LS1) period, which will be devoted to the upgrade of the accelerator to reach a proton-proton collision energy of 13 TeV by 2015.

Since then, the four LHC experiments, ALICE, ATLAS, CMS and LHCb, have relied in the WLCG and in EGI to continue the processing of the data collected, and progress on the physics analysis that are resulting in an impressive collection of new results, as reflected by more than 500 publications in 2013, covering a wide range of topics.

When the upgraded LHC restarts operations in Run 2 by 2015, the computing needs of the collaborations will substantially increase, due to the higher complexity of the collisions (with an energy almost doubled), collision frequency and pile-up conditions. Processing this large stream of new data will require an upgrade in the capacities of the e-Infrastructure, more than doubling current computing power and data storage resources. The possibility of exploring new physics beyond the Standard Model, for example complex multi-parametric models as those based on supersymmetry, will require even larger resources and services.

There is a clear challenge for the next years in the evolution of the current WLCG related infrastructure, organised in a tiered structure, into a dynamic, flexible and distributed framework, likely including federated cloud resources, where EGI should play a key role.

Scientific Case

The LHC³⁶ is the world's largest and most powerful particle accelerator. It first started up on 10 September 2008, and remains the latest addition to CERN's accelerator complex. The LHC consists of a 27-kilometre ring of superconducting magnets with a number of accelerating structures to boost the energy of the particles along the way. Inside the accelerator, two high-energy particle beams travel at close to the speed of light before they are made to collide. All the controls for the accelerator, its services and technical infrastructure are housed under one roof at the CERN Control Centre. From here, the beams inside the LHC are made to collide at four locations around the accelerator ring, corresponding to the positions of four particle detectors – ATLAS, CMS, ALICE and LHCb.



Perspective of the LHC Tunnel: protons circulate in a beam pipe surrounded by superconducting magnets (in blue) used to bend their trajectory to a circumference.


These experiments use detectors to analyse the myriad of particles produced by collisions in the accelerator and are run by collaborations of scientists from institutes all over the world. Each experiment is distinct, and characterised by its detectors. The two biggest experiments, ATLAS and CMS, use general-purpose detectors to investigate the largest range of physics possible. Having two independently designed detectors is vital for cross-confirmation of any new discoveries made. ALICE and LHCb have detectors specialised for focussing on specific phenomena. These four detectors sit underground in huge caverns on the LHC ring.

ALICE (A Large Ion Collider Experiment) is a heavy-ion, 10,000-tonne ALICE detector – 26 m long, 16 m high, and 16 m wide – sitting in a vast cavern 56 m below ground close to the village of St Genis-Pouilly in France. The detector was designed to study the physics of strongly

³⁶ All information on LHC and experiments compiled on November 2013 from <http://home.web.cern.ch/about/accelerators/large-hadron-collider> and from subsections in <http://home.web.cern.ch/about/experiments>. All images and text under CERN copyright.

interacting matter at extreme energy densities, where a phase of matter called quark-gluon plasma forms³⁷.

The existence of such a phase and its properties are key issues in the theory of quantum chromodynamics (QCD), for understanding the phenomenon of confinement, and for a physics problem called chiral-symmetry restoration. The ALICE collaboration, bringing together more than 1,000 scientists from over 100 physics institutes in 30 countries, studies the quark-gluon plasma as it expands and cools, observing how it progressively gives rise to the particles that constitute the matter of our universe today.




Transversal view of ALICE (left) and ATLAS (right) detectors

ATLAS is one of two general-purpose detectors, built to investigate a wide range of physics, from the search for the Higgs boson to extra dimensions and particles that could make up dark matter. Beams of particles from the LHC collide at the centre of the ATLAS detector making collision debris in the form of new particles, which fly out from the collision point in all directions. Six different detecting subsystems arranged in layers around the collision point record the paths, momentum, and energy of the particles, allowing them to be individually identified. A huge magnet system bends the paths of charged particles so that their momenta can be measured. The interactions in the ATLAS detectors create an enormous flow of data. To digest the data, ATLAS uses an advanced 'trigger' system to tell the detector which events to record and which to ignore. Complex data-acquisition and computing systems are then used to analyse the collision events recorded. The 7000-tonne ATLAS detector - 46 m long, 25 m high and 25 m wide - is the largest volume particle detector ever constructed. It sits in a cavern 100 m below ground near the main CERN site, close to the village of Meyrin in Switzerland. More than 3,000 scientists from 174 institutes in 38 countries work on the ATLAS experiment.

³⁷ All ordinary matter in today's universe is made up of atoms. Each atom contains a nucleus composed of protons and neutrons (except hydrogen, which has no neutrons), surrounded by a cloud of electrons. Protons and neutrons are in turn made of quarks bound together by other particles called gluons. No quark has ever been observed in isolation: the quarks, as well as the gluons, seem to be bound permanently together and confined inside composite particles, such as protons and neutrons. This is known as confinement. Collisions in the LHC generate temperatures more than 100,000 times hotter than the centre of the Sun. For part of each year the LHC provides collisions between lead ions, recreating in the laboratory conditions similar to those just after the big bang. Under these extreme conditions, protons and neutrons 'melt', freeing the quarks from their bonds with the gluons. This is quark-gluon plasma.

The **Compact Muon Solenoid (CMS)** is a general-purpose designed to investigate a wide range of physics, including the search for the Higgs boson, extra dimensions, and particles that could make up dark matter. Although it has the same scientific goals as the ATLAS experiment, it uses different technical solutions and a different magnet-system design. The CMS detector is built around a huge solenoid magnet. This takes the form of a cylindrical coil of superconducting cable that generates a field of 4 Tesla, about 100,000 times the magnetic field of the Earth. The field is confined by a steel 'yoke' that forms the bulk of the detector's 12,500-tonne weight. An unusual feature of the CMS detector is that instead of being built *in situ* such as the other LHC giant detectors, it was constructed in 15 sections at ground level before being lowered into an underground cavern near Cessy in France and reassembled. The complete detector is 21 metres long, 15 metres wide and 15 metres high. The CMS experiment is one of the largest international scientific collaborations in history, involving 4,300 particle physicists, engineers, technicians, students and support staff from 179 universities and institutes in 41 countries.



Left: Transversal view of the CMS detector. Right: A view of the photomultiplier tubes used in the LHCb detector.

The **Large Hadron Collider beauty** (LHCb) experiment specialises in investigating the slight differences between matter and antimatter by studying a type of particle called the 'beauty quark', or 'b quark'. Instead of surrounding the entire collision point with an enclosed detector as do ATLAS and CMS, the LHCb experiment uses a series of sub-detectors to detect mainly forward particles - those thrown forwards by the collision in one direction. The first sub-detector is mounted close to the collision point, with the others following one behind the other over a length of 20 metres. An abundance of different types of quark are created by the LHC before they decay quickly into other forms. To catch the b quarks, LHCb has developed sophisticated movable tracking detectors close to the path of the beams circling in the LHC. The 5,600 tonne LHCb detector is made up of a forward spectrometer and planar detectors. It is 21 metres long, 10 metres high and 13 metres wide, and sits 100 metres below ground near village of Ferney-Voltaire, France. About 700 scientists from 66 different institutes and universities make up the LHCb collaboration

The scientific challenge

Since the start of LHC, an impressive progress has been made in the pursuit of its core mission, which is to elucidate the laws of nature at the most fundamental level. A giant leap, the discovery of the Higgs boson, has been accompanied by many experimental results confirming the Standard Model beyond the previously explored energy scales.

Compendium of e-Infrastructure requirements for the digital ERA

These results raise further questions on the origin of elementary particle masses and on the role of the Higgs boson in the more fundamental theory underlying the Standard Model, which may involve additional particles to be discovered around the TeV scale. The discovery of the Higgs boson is the start of a major programme of work³⁸ to measure the particle's properties with the highest possible precision for testing the validity of the Standard Model and to search for further new physics at the energy frontier. The LHC is in a unique position to pursue this programme. Europe's top priority should be the exploitation of the full potential of the LHC, including the high-luminosity upgrade of the machine and detectors with a view to collecting ten times more data than in the initial design, by around 2030. This upgrade programme will also provide further exciting opportunities for the study of flavour physics and the quark-gluon plasma.

State of the art

By 2013 the four LHC experiments are completing the processing of Run 1 data, and the analysis driving to publication of results in topics including Standard Model tests, top quark physics, b physics, Higgs analysis, search for supersymmetric particles, exotica, forward physics and heavy ion collisions, among others.

The resources available are devoted to the processing of the collected data, from RAW format to ESD/AOD/DST format and skimming to final analysis n-tuples, to simulation of different signal and backgrounds, and to the estimation of limits or discovery significance. These tasks require more than 50 million computational jobs per month, running on more than 250,000 processors, and a total storage capacity exceeding currently 180 PB on disk and 150 PB on tape.

After completion of the processing and analysis of Run 1 data, the collaborations are starting the detailed simulation for Run 2 that must be ready by 2015, in time for the first analysis with data collected at 13 TeV, to be presented in 2015 summer conferences. Given the potential discovery associated to the large increase in centre of mass energy, from 8 to 13 TeV, and the also large increase in backgrounds and signal event complexity, this is a clear challenge for LHC computing.

Going beyond the state of the art

Very relevant scientific challenges will be studied in the next years at the LHC:

Properties of the Higgs boson. Since the discovery of the 126 GeV Higgs-like particle in Summer 2012, the LHC experiments have focused on the measurements of its production rates and decay properties. Both ATLAS and CMS have released results strongly suggesting that the new particle is a Higgs boson and its properties are consistent within the experimental and theoretical uncertainties with the expectations of the Standard Model Higgs boson. After the two-year shutdown, the LHC is scheduled to operate again in 2015 at 13 TeV and it is expected to deliver 300 fb⁻¹ to each experiment by 2022. With the planned high-luminosity (HL-LHC) upgrade, an integrated luminosity of 3,000 fb⁻¹ could be foreseen by 2030. The increased

³⁸ Updated European Strategy for Particle Physics (2013):
http://council.web.cern.ch/council/en/EuropeanStrategy/ESBrochure-Strategy_Report2013.pdf

luminosity will significantly increase the measurement precision of the Higgs boson properties. LHC at 13 TeV with 300 fb⁻¹ of data is essential to firmly establish the major production mechanisms of a Higgs boson and its main bosonic and fermionic decay modes. This will also lead to about 100 MeV precision on the Higgs boson mass and the measurement of the boson spin. HL-LHC also provides unique capabilities to measure rare statistically limited SM decay modes, and increase the precision on the coupling properties with a high discovery potential for heavy Higgs bosons as predicted by many models beyond the Standard Model³⁹.

Searches for supersymmetry (SUSY) encompass a wide range of strategies aimed at different particles of the spectrum. The most common searches assume that supersymmetric partners of the Standard Model particle carry a conserved quantum number, called R-parity. If the lightest supersymmetric particle is neutral, it will typically be weakly interacting and will not be observed directly in a collider detector. Events are then characterised as containing several hadronic jets, associated with decay to the lightest particle plus missing transverse energy. No significant excess of such events has yet been observed. The results of the searches are then parametrised by limits on the gluino mass and on a squark mass, assumed common to all squark flavours. Under these assumptions current LHC results exclude masses of events up to gluino masses of 1.0 TeV and, independently, up to squark masses of 1.3 TeV. For the future stages of the LHC, we expect to be able to discover such events up to gluino masses of 1.9 TeV and squark masses of 2.3 TeV with 300 fb⁻¹, and up to 2.7 TeV with 3000 fb⁻¹. It is possible that the first signal of SUSY would not be given by the generic search just described, but would require a more specialised analysis. Special search techniques are needed in models in which mass gaps in the SUSY spectrum are relatively small so that hard jets are not emitted in particle decays, and models in which only the partners of top quarks, or perhaps only uncolored supersymmetric particles, are produced at accessible energies⁴⁰.

Model-independent searches beyond the Standard Model with rare decays and CP violation: the new particles predicted in SUSY and other extensions of the Standard Model can contribute in loop processes to the decays of beauty and charm hadrons. Precision measurements of rare processes can therefore probe very high energy scales, far beyond the reach of the LHC. One particularly interesting possibility is that the new particles may cause new sources of matter-antimatter asymmetry (so-called CP violation) and could help to explain the dominance of matter in the Universe. This approach to searching beyond the Standard Model, which is pursued extensively by the LHCb experiment, requires the collection and processing of extremely large datasets and therefore puts stringent constraints on the computing model.

³⁹ <http://www-public.slac.stanford.edu/snowmass2013/SnowmassWorkingGroupReports.html>

⁴⁰ <http://www-public.slac.stanford.edu/snowmass2013/SnowmassWorkingGroupReports.html>

e-Infrastructure

Current e-Infrastructure activity

The Worldwide LHC Computing Grid, or WLCG⁴¹, is a global collaboration linking grid infrastructures and computer centres worldwide to distribute, store and analyse the data generated by the LHC at CERN.

The WLCG is organised in three layers, or tiers, made up of computer centres which contribute to different aspects of the WLCG.

Tier-0 is the CERN Computing Centre. All data from the LHC passes through this central hub but it provides less than 20% of the total compute capacity. It is connected to other major tiers and grid services using dedicated 10 Gb/s optical-wide area links.

Tier-1 consists of eleven sites located in Canada, France, Germany, Italy, the Netherlands, the Nordic countries, Spain, Taipei, the UK, and two sites in the USA. These sites provide distribution networks, processing of raw data, data analysis, and storage facilities.

The remaining 140 **Tier-2** sites are distributed across most of the globe. Together, these sites provide approximately half of the capacity needed to process the LHC data.

Massive, multi-petabyte, storage systems and computing clusters with thousands of nodes connected by high speed networks are the building blocks of the WLCG centres. The WLCG centres use specialised tools to manage the immense disk and magnetic tape mass storage systems needed for LHC data, and allow applications to access the data for simulation and analysis, independent of the storage medium (tape or disk) that the data resides on. WLCG also uses dedicated software components, called middleware, to link up the varying hardware resources across the grid in a compatible way. This allows scientists and other users to access these resources in a uniform and secure way from anywhere in the world, turning the diverse and locally managed computing centres into a single massive virtual resource.

WLCG's European resources are integrated through the EGI infrastructure, and use the Unified Middleware Distribution (UMD). The US contribution to WLCG is provided through the Open Science Grid (OSG).

Exchanging data between WLCG centres is managed by the Grid File Transfer Service that was developed by the EGEE project. This is the proven method for securely and reliably transferring large volumes of data across distributed computing grids. It has been tailored to support the special needs of grid computing, including authentication and confidentiality features, reliability and fault tolerance, and third party and partial file transfer.

Future e-Infrastructure challenges


After the success of Run 1, ending in February 2013, the LHC is in its first long shutdown, LS1, extending through 2013 and 2014. The accelerator is scheduled to restart in April 2015 for Run 2, at a centre of mass energy of 13 TeV, where cross sections for interesting physics processes

⁴¹ <http://wlcg.web.cern.ch/>

are factors of two or more above those at 8 TeV reached in the last part of Run 1. The corresponding requests of computing resources for 2014 and a discussion of requests for 2015 and beyond were recently reported to the LHC Resources Review Board⁴², and are summarised below.

The anticipated stable LHC beam times for Run 2 will vary from 3 to 7 times 106 seconds per year, from 2015 to 2017. In 2015 the machine must be commissioned, while in 2017 it is assumed that LHC will be available for longer. The expected pileup (number of collisions in each beam-crossing) will increase from 25 (2015) to 40 (2017) for ATLAS and CMS for proton-proton running. This is a crucial parameter since event sizes and reconstruction times increase as pileup increases, with concomitant increases in simulation times and sizes. The resources requests assume 25 ns running for Run 2 that reduces pileup by allowing fewer particles in each bunch for the same luminosity. The LHC luminosity is expected to be $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ in 2015 and will increase by 50% for 2016 and 2017. The increase in centre of mass energy for Run 2 leads to greater track multiplicity in each interaction, increasing event sizes and subsequent processing times.

A recent update of the computing models⁴³ has reviewed the expected evolution of resource requirements over the next years of LHC Run 2. The figures below show the evolution expected from 2014 to 2017. Both CPU power and disk storage capacity are expected to double by 2017.



Expected evolution of requirements for CPU and disk for LHC Run 2.

Future plans

Along this last months a major review of the experiment and WLCG computing models has been undertaken, analysing the period of LHC Run 2, extending to 2017. The key points were summarised in a recent report to the Resources Review Board⁴⁴ and are reproduced in what follows:

“Regarding the Computing Model, all four experiments will take a more pragmatic view of which tasks should be run at the different Tiers compared to the original model, making use of the capabilities of the sites (capacities, connectivity, etc.) rather than a strict allocation of functions (...)

⁴² CERN-RRB-2013-094

⁴³ I. Bird et al. (Eds), Update of the Computing Models of the WLCG and the LHC Experiments (to be released in December 2013)

⁴⁴ Bird, I., S. Foffano, CERN-RRB-2013-089

for all experiments, there is a general assumption now that data can be accessed between any peers, rather in the original hierarchical model. This evolution had already begun in Run 1. For data management, all four experiments now use so-called data federations based on xrootd, which facilitates an optimised access of data from jobs, even if the data is initially remote from the site at which the job runs. Significant work has been invested in more intelligent data placement and caching, to optimise the number of files that need to be pre-placed or dynamically cached when (or just before) they are required. Data popularity services are introduced to better determine which data sets should be widely available and which data can be cleaned from disk caches (...)

Significant efforts are also being invested by the experiments and the WLCG applications area in order to improve the overall efficiency of use of modern multi-core CPU, and in future to be ready to make better use of parallelism and vector abilities of newer processors (...)

Regarding the Distributed Computing environment, there are several factors that are driving the simplification of the distributed computing environment. These are: the need to minimise the operational (staff) cost of running the grid sites; the need to be able to simply make use of opportunistic resources (clusters, clouds, HPC resources) with as little as possible set up and configuration; the need to reduce as far as possible the cost of maintaining complex grid middleware.

These factors are all leading towards a simplification of the grid middleware layer, which is realistic as the computing systems have evolved and complexity has moved towards the application layer. The ubiquitous use of pilot job frameworks have also helped simplify the system. At the same time new technologies such as open source cloud management software provides a natural way of implementing a simpler job management layer. An evolution towards a more cloud-like model for job management over the period of Run 2 is anticipated (...)"

Summary – conclusions and recommendations

This publication provides an overview of 11 scientific use cases from some of the largest, European, multinational research collaborations. Some of these are already supported by the EGI solutions.

To date EGI has been supporting 302 projects, of which 220 are currently active either nationally or internationally. EGI is enabling approximately 38,000 users of which 23,000 with personal certificate, including those from six Research Infrastructures on the ESFRI roadmap: BBMRI, CTA, EISCAT-3D, ELI-NP, LIFEWATCH and KM3NeT.

EGI will continue evolving the EGI service portfolio for the Digital European Research Area through user co-development and under the drive of user needs. EGI's human networks – such as the NGI International Liaisons and Research Champions – and the Distributed Competence Centre are key engagement channels to scientific communities, to capture and translate new requirements that drive the EGI technical roadmap. Virtual Teams, Competence Centres and dedicated projects mobilise experts and technical assets from across the EGI Collaboration to specify, develop and integrate new solutions into EGI's production infrastructure for certain scientific communities.

This compendium helps EGI and other e-Infrastructures of European relevance to identify the common needs across multiple communities and formalise recommendations to drive their strategies. To that effect, EGI will:

evolve its services to provide easier interfaces and policies for access to the services, including support of federated identify for easier authentication and authorization.

provide enabling services to Researchers by adopting a defined service portfolio and a user - centric approach to its development that will expand EGI's current service offering from EGI.eu and affiliated NGIs to retain its current research communities and attract new research communities.

provide flexible Virtual Research Environments that accelerate the ability of researchers to undertake excellent science by leveraging the expertise and connections of the NGIs to introducing technical innovations into production across Europe.

mobilize the knowhow and expertise from within the Community to establish and operate a cross-border federation of digital capabilities, knowledge and expertise for reuse by Research Infrastructures and international research collaborations.

commit the development and adoption of open standards, and will promote the interoperability of technical interfaces, access policies and support channels among e-Infrastructures of European relevance.

Disclaimer

The articles featured in this publication express the views of their respective authors within the context of their research collaboration.

Copyright

Copyright © 2015 EGI.eu