# DEISA Extreme Computing Initiative (DECI) and Science Community Support

Alison KENNEDY [a;1]
and Hermann LEDERER [b;2]
[a] *EPCC, The University of Edinburgh, UK*
[b] *Rechenzentrum Garching der Max-Planck-Gesellschaft, Germany*

**Abstract.**
The DEISA Consortium continues its work in the DEISA2 project with EU FP7 support from 2008 to 2011. DEISA2 focuses on the provisioning and operation of infrastructure services which allow its users to work efficiently within a distributed high performance computing environment. DEISA supports both single projects (via DECI) and Virtual Communities.

## Introduction

DEISA is a consortium of leading European national supercomputing centres, founded with the purpose of "fostering and advancing computational science in Europe" in the area of High Performance computing by providing researchers with CPU, applications support and user services [1,2]. The consortium is currently funded through the DEISA2 project with EU support from 2008 to 2011 [3,4]. Researchers have two ways to apply for access to the DEISA infrastructure - via the single-project access route known as the DEISA Extreme Computing Initiative (DECI) or as a Virtual Community. This paper describes these two opportunities and the types of research communities who have availed themselves of DEISA's services.

## 1. DEISA Extreme Computing Initiative

The DEISA Extreme Computing Initiative (DECI) is a scheme through which European scientists can apply for single-project access to world-leading computational resources in the European HPC infrastructure for a period of up to 12 months per project [5].

DECI was introduced in 2005 to enable world-leading European computational scientists to obtain access to the most powerful national computing resources in Europe regardless of their country of origin or work, and to enhance DEISA's impact on European science and technology at the highest level. Through an annual call, a number of capability computing projects are selected by peer-review on the basis of innovation and sci-

[1] Corresponding Author; E-mail: A.Kennedy@epcc.ed.ac.uk
[2] E-mail: Lederer@rzg.mpg.de

entific excellence. The consortium has designed, deployed and operates a complex, heterogeneous supercomputing environment with an aggregate peak performance in excess one PetaFlop/s.

Successful projects are given access to the exceptional resources in the DEISA infrastructure (on HPC architecture selected for its suitability) and are offered applications support to enable them to use it productively. The number of proposals received and the number of CPU cycles requested by applicants has grown year on year since DECI's inception, showing that there is a continuing demand from European researchers for a single-project access scheme. DECI will continue to operate until 2011 through FP7 DEISA2 funding, as one of a range of access initiatives offered by DEISA to the European HPC user community.

DECI is of key importance in continuing to build a European HPC user community, supported in their use of top-level HPC facilities by applications experts from leading European HPC centres. More than 500 different researchers from 25 European countries have participated in the scheme to date as Investigators or related scientific users, along with collaborators from four other continents. The scheme aims to enhance Europe's international standing in science .

DECI also aims to facilitate a better understanding of the likely requirements of future users of the Tier-0 systems (European leadership-class supercomputers) by collecting real use-case information about what European computational scientists want and about the differences between usage of national and European resources and facilities.

DECI-4 call attracted 66 proposals for challenging European computational science projects, requesting over 134 million processor(-core)-hours and asking for significant application support. The call was oversubscribed by a factor of around 3, both in requests for CPU and for applications enabling effort. 49 million processor(-core)-hours were available for distribution, and these were allocated to 42 projects, which have been given access to the infrastructure between 1 January 2009 and 31 December 2009 [6].

The most recent call (DECI-5) attracted 75 proposals, requesting 220 million processor(-core) hours. The 69 million hours available have been awarded to 50 projects [7].

Figure 1 shows how the demand for CPU is increasing at a faster rate than supply. Each partner commits 5% of their CPU to DEISA (although some partners contribute additional resources) so the annual increase in CPU available via DEISA broadly reflects the average increase in CPU available at the partner sites.

Figure 2 shows how the average and the median amount requested has grown during the lifetime of DECI. From this we can see that while the median amount requested has increased broadly in line with the increase in availability of CPU, the average amount requested has seen a very sharp increase. This indicates an increase in the number of projects requesting a large amount of CPU cycles.

### *1.1. Overview of DECI-4 Projects*

The following section gives an overview of the projects running on the DEISA infrastructure during 2009 to show the range of science supported.

#### *1.1.1. Projects by Applications Area*

The 42 DECI projects selected for DECI-4 were self-categorised (by the PIs) into six broad scientific disciplines - Astronomical Sciences, Biological Sciences, Earth Sci-

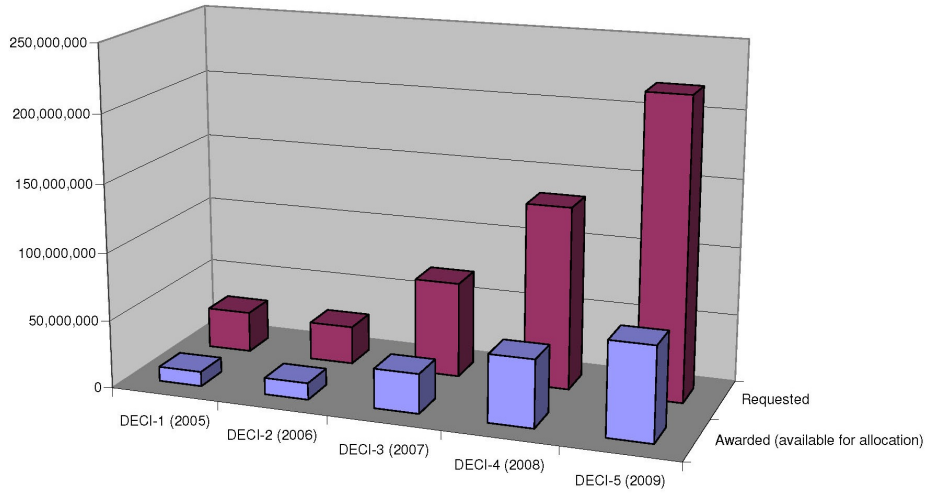


**Figure 1.** CPU in DEISA is increasing at a faster rate than supply

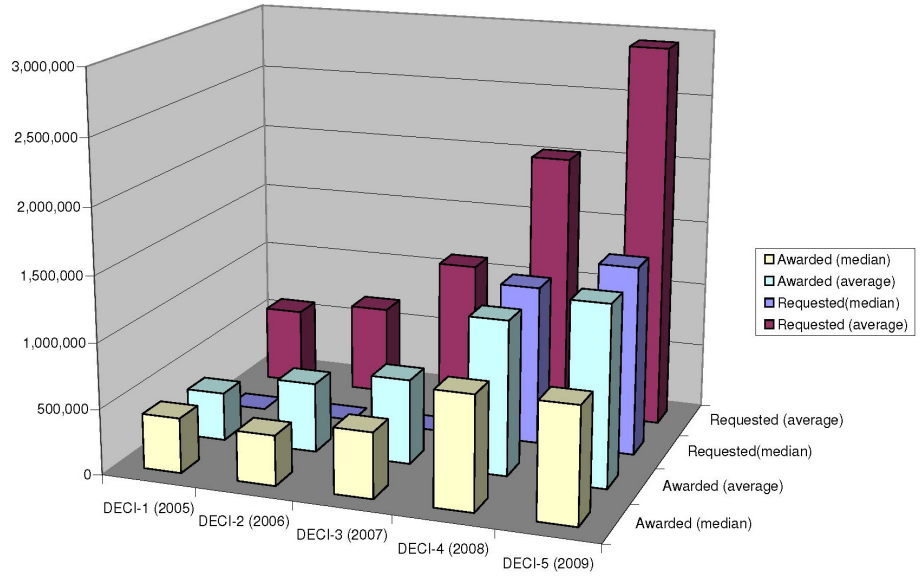


**Figure 2.** Average and median amount requested during the lifetime of DECI

ences, Engineering, Materials Science and Plasma/Particle Physics. The number of awards by discipline is shown in Figure 3. Overall, about two thirds of the proposals were accepted, although the amount of resources awarded to a number of projects was scaled back to enable as many projects as possible to be supported.

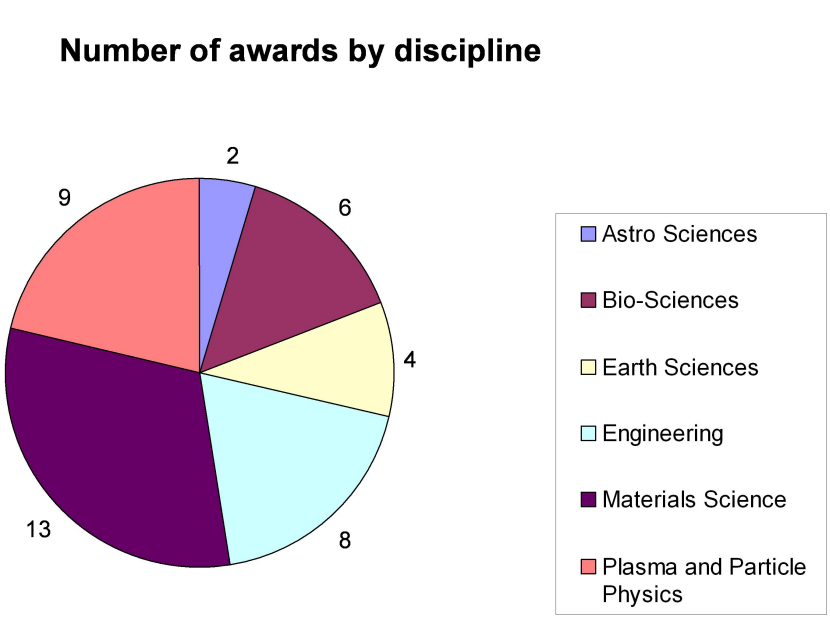


**Figure 3.** Number of CPU awards by discipline for the DECI-4 projects

### *1.1.2. Projects by Nationality of Applications*

There were 115 investigators named in the 42 accepted DECI-4 projects - an average of 2.7 investigators per project. Of these investigators, 88 (77%) came from countries with a DEISA partner site, 19 (17%) from other European countries and 8 (7%) from large counties outside of Europe (China, Japan, USA). This suggests that successful project proposals were generally based on collaborative European science. All of the 44 PIs (two projects had two PIs) were Europeans, offering further proof that the research being undertaken focuses on European research priorities, with 39 (89%) of the PIs coming from countries with DEISA sites and 5 (11%) coming from other European countries.

As many of the countries with large computational research communities are DEISA partners, it is to be expected that a large majority of the PIs and CoIs will continue to come from countries with DEISA sites. However, we are working to ensure that future calls are publicised extensively throughout Europe. The figures from the latest call (DECI-5) indicate that we are being successful in this respect, showing an increase in the number of investigators from non-DEISA European countries to 20% of the total.

In an effort to find out more about the degree of co-operation and collaboration within DECI projects which were accepted, we analysed the information on investigators in further detail. The number of investigators varied from one to eight, with the median number of investigators being three. During the five years in which DECI has been operating, the average number of investigators per project has risen steadily, year-on-year from 1.03 to 2.96.

A similar analysis was undertaken of the number of institutes involved in each project. This revealed that the number of institutes involved in each project varied from one to six, with the median number of institutes being two.

Finally, we looked at the number of countries collaborating in successful proposals. This showed that 23 (54.8%) of the projects involved scientists from more than one country, with 19% involving scientists from three or more countries.

Here is the evolution of the proportion of projects involving researchers from three or more countries:

- DECI-1: 0%
- DECI-2: 21.4%
- DECI-3: 21.7%
- DECI-4: 19.0%
- DECI-5: 26.0%

This shows that the proportion of projects involving collaborations of scientists from three or more countries in increasing every year, indicating that DECI is supporting pan-European scientific collaboration, and that the DEISA infrastructure is attractive to European researchers. Of the 364 scientific investigators who have used the DEISA infrastructure via DECI, 281 (77%) have been involved in one project to date with a further 83 (23%) involved in two or more projects. These figures indicate that DEISA is being successful in reaching out to new groups and collaborations. But they also show that scientists who have used the infrastructure recognise its value and often apply to use it subsequently for appropriate collaborative projects.

Overall, the statistics which we have collected suggest that DECI is being successful in attracting high quality collaborative proposals involving scientists from more than one European country, and in involving partners from outside Europe where this is appropriate.

### *1.1.3. CPU awarded*

As can be seen from Figure 4, the average number of standard core hours awarded to a DECI project was 1,178,955. The median award was 867,792. Astronomical science, earth science and plasma physics projects on average were awarded more resources than average per project while biological sciences, engineering and materials science were awarded fewer resources than average per project.

However, as can be seen from Figure 3, the largest amount of resources went to materials science projects (27%) closely followed by plasma and particle physics (also 27%). The remaining 48% of the resources were shared between the other four scientific areas.

### *1.1.4. Scientific codes used*

In all projects accepted so far 180 different codes have been specified for usage in DEISA.

The big majority (150 codes or 83%) consists of pure MPI codes, followed by hybrid OpenMP/MPI codes (21 codes or 12%). There has also been a very small number of pure OpenMP codes (six codes or 3%), one pthread-parallel code and two serial codes for pre- and post-processing purposes only.

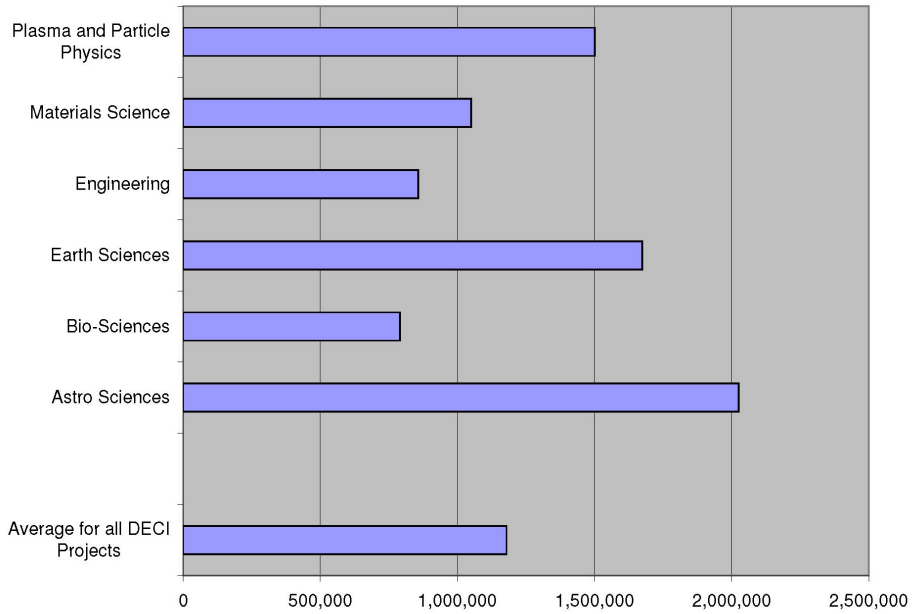


**Figure 4.** Average number of standard CPU hours by project by discipline

This large number of codes used in DEISA also can provide some insight into the codes used per science area. Since some codes are used in different science areas, the sum will be larger than the total number of codes of 180. The field is led by Plasma and Particle Physics with 43 codes (24%), followed closely by Materials Science with 41 codes (23%). The next three close groups are Engineering (27 codes or 15%), Life Sciences (26 codes or 15%), and Astronomical Sciences (25 codes or 14%). Earth Sciences have been represented with 19 different codes (11%). In addition, Informatics, not shown in Figure 3, was represented with three codes.

Top ranking codes with highest usage over several projects include: NAMD (Materials and Life Sciences, 19 times), CPMD (Materials and Life Sciences, 18 times), CP2K (Materials and Life Sciences, 15 times), GROMACS (Materials and Life Sciences, 12 times). In Astronomical Sciences, GADGET has been used most frequently (four times), in Earth Sciences ECHAM and derivatives (seven times), in Plasma and Particle Physics ELMFIRE, EUTERPE, GENE and ORB5 four times each.

In particular, the large number of different codes used by DECI projects indicates the wide variety of science being undertaken via DECI and reinforces the need for close partnership with users in applications enabling and code tuning to help them achieve their objectives. There is obviously a strong demand for the sort of complementary computing support offered by DECI.

Thus a remarkable high number of codes and their respective projects were able to benefit from the BlueGene/P systems equipped with large numbers of processor-cores for the use of which good scalability to at least 1024 or 2048 cores was mandatory.

## 2. Virtual Community Support

DEISA has no set definition of a Virtual Community but seeks to offer an alternative access mechanism for larger, loosely or closely coupled European research consortia to provide them with access to DEISA facilities for a longer period of time than the annual cycle of a DECI project can guarantee.

Any Virtual Community supported by DEISA is given an allocation of resources. Responsibility for the allocation of resources amongst the members of the community is delegated to the community itself.

This initiative by DEISA has given a strong impetus to HPC support to Virtual Communities within Europe. An open call for "Expressions of Interest for Community Support" (EoIs) was published in December 2008 [8]. By the closing date in February 2009, seven Expressions of Interest had been received from four different areas of science (Climate Research; Fusion Research; Astro Sciences; and Life Sciences) from Virtual Communities with undisputed scientific reputations as well as a scientific critical mass based across many European countries.
The Virtual Communities who submitted EoIs were as follows:

- Fusion Research: EFDA and EUFORIA
- Climate Research: ENES
- Astrophysics/Cosmology: VIRGO/COSMOCOMP and LFI-PLANCK
- Life Sciences: VPH NoE and VIROLAB

Their motivation for applying to DEISA resources and support is discussed in the following section.

### *2.1. Fusion Research*

DEISA has a tradition in supporting fusion energy research in the field of supercomputer simulations and has now started with direct community support [9].

#### *2.1.1. European Fusion Development Agreement - EFDA*

The decision to proceed with ITER - a joint international research and development project that aims to demonstrate the scientific and technical feasibility of fusion power as a viable future energy option -

has focussed EU attention on its own need to have sufficient theory and modelling capabilities to adequately support and exploit the ITER project. This implies having to develop appropriate physics-based models and having adequate computational facilities to apply these models to key ITER-scale problems.

EFTA looked to DEISA to provide High Performance Compute resources, access to state-of-the-art HPC architectures, and application enabling assistance to support physics research in areas such as turbulence, fast particle physics, nonlinear MHD and extended MHD, edge physics, scrape-off-layer (SOL) and plasma wall interactions (PWI), radio-frequency heating and current drive and physical modelling of materials.

#### *2.1.2. EUFORIA*

EUFORIA is an FP7 project with 14 countries involved which aims to enhance the modelling capabilities for ITER-sized plasmas through the adaptation, optimization and inte-

gration of a set of critical applications for edge and core transport modelling that targets different computing paradigms. Code porting and optimization was undertaken by the EUFORIA partners who then asked DEISA for access for production runs to supercomputing resources, access to state-of-the-art supercomputers, and technology support for accessing DEISA from their gateway machine in Italy.

### *2.2. Climate Research*

#### *2.2.1. European Network for Earth System Modelling - ENES*

A major challenge for the climate research community is the development of comprehensive Earth system models capable of simulating natural climate variability and human-induced climate changes. Such models need to account for detailed processes occurring in the atmosphere, the ocean and on the continents including physical, chemical and biological processes on a variety of spatial and temporal scales. They have also to capture complex nonlinear interactions between the different components of the Earth system and assess how these interactions can be perturbed as a result of human activities.

An important task is to develop an advanced software and hardware environment in Europe, under which the most advanced high resolution climate models can be developed, improved, and integrated. The European Network for Earth System Modelling (ENES), comprising about 50 public and private institutions involved in Earth system and climate research, was initiated in 2001.

ENES has asked for High Performance Compute resources, access to state-of-the-art HPC architectures, community data repository, application enabling, and technology support.

### *2.3. Astro Sciences*

#### *2.3.1. LFI-PLANCK*

Planck is a European Space Agency satellite that will be launched in 2009 to study the cosmic microwave background. LFI-PLANCK is a project within the ESA PLANCK mission to study the birth of the universe; six European countries are involved.

Planck will have a major impact in cosmology and will be valuable for astrophysics also. Planck data analysis is a complicated and computationally demanding task, where simulation work and Monte Carlo studies play a crucial role.

Planck LFI Virtual Community has requested High Performance Compute resources, a Community data repository (with associated technical support for managing large amounts of disk and data) and application enabling assistance for porting applications.

#### *2.3.2. VIRGO/COSMOCOMP*

The Virgo Consortium (founded 1994) is an international group which combines computer, software and personnel resources to carry out top-end simulations of the formation of cosmic structure. It has repeatedly pushed the envelope in this field by carrying out the largest simulations ever of major problems of interest. With respect to the following services: High Performance Compute resources, Access to state of-the-art HPC architectures, Community data repository, Application enabling, Technology support, VIRGO

has expressed potential interest in all of them, but the last three are of particular interest for the data archive facility of the Virgo Consortium.

### *2.4. LIFE Sciences*

#### *2.4.1. Virtual Physiological Human - VPH*

VPH is an FP7 Network of Excellence with 7 European countries involved.

Virtual Physiological Human lists, as its main target outcome, patient-specific computer models for personalised and predictive healthcare and ICT-based tools for modelling and simulation of human physiology and disease-related processes. The initiative consists of various projects with a need to access compute resources on an EU wide basis, in order to support their scientific objectives.

The following services have been asked by the VPH Consortium to be supported by DEISA: HPC resources, data storage resources, Grid Interfaces (Unicore 6/Globus 4) - OGSA BES interfaces, advanced reservation capabilities (HARC), computational steering (RealityGrid steering library), emergency computing tools (SPRUCE), application deployment support, and user support.

#### *2.4.2. VIROLAB*

The mission of the EU FP6 project ViroLab is to develop a virtual laboratory for infectious diseases. In future years, genetic information is expected to become increasingly significant in many areas of medicine. This expectation comes from the recent and anticipated achievements in genomics expected to be of use for the prevention, diagnosis and treatment of diseases.

To continue the HIV drug binding affinity simulation work undertaken by the project, for which DEISA community support was started in 2008, a continuation of the support with additional compute resources was requested until the end of the project in 2009.

### *2.5. Summary for Virtual Communities*

DEISA requested the EoIs to specify the requirements of their community as follows:

What types of DEISA services and resources will be of community interest for the collaboration?

1. High Performance Compute resources,
2. Access to state-of-the-art HPC architectures
3. Community data repository,
4. Application enabling,
5. Technology support

According to the answers we received, all communities expressed needs for multiple services, but no one community had immediate need of all of the services offered by DEISA. However, the flexible support model on offer allows the mix of services to be tailored to the specific requirements of a Virtual Community and to be altered over time.

## Conclusions

DEISA has established two complementary and effective ways for challenging supercomputing projects in Europe: The classic project-oriented DEISA Extreme Computing Initiative DECI and the support of Virtual Science Communities in the HPC domain. Both ways have found a high user acceptance and thus significantly contribute to the advancement of computational sciences in Europe.

## Acknowledgments

The authors thank the European Commission for support through contract RI-222919

## References

[1] Lederer, H.: DEISA - towards a persistent European HPC Infrastructure, in: eStrategies: Projects, 2009 (ISSN 1758-2369), 29-31, 2009

[2] Supercomputing gets its own superhero, see http://cordis.europa.eu/ictresults/index.cfm/section/news/tpl/article/BrowsingType/Features/ID/90477

[3] Lederer, H.: DEISA2: Supporting and developing a European high-performance computing ecosystem, Journal of Physics: Conference Series 125 (2008) 011003; doi:10.1088/1742-6596/125/1/011003

[4] For further details see http://www.deisa.eu

[5] Lederer, H.: DECI - The DEISA Extreme Computing Initiative, inSiDE, Vol 6 No 2, Autumn 2008, eds. H-G. Hegering, Th. Lippert, M. Resch; Gauss Centre for Supercomputing, 2008

[6] DEISA Announces Extreme Computing Awards, see http://www.hpcwire.com/offthewire/DEISA-Announces-Extreme-Computing-Initiative-Awards35540714.html

[7] DEISA Announces Extreme Computing Initiative Awards, see http://www.hpcwire.com/offthewire/DEISA-Announces-Extreme-Computing-Initiative-Awards-63714257.html

[8] DEISA Announces Virtual Community Support Initiative, see http://www.hpcwire.com/industry/government/DEISA-Announces-Virtual-Community-Support-Initiative-36014839.html

[9] Commission puts Europe's supercomputers on the path to sustainable energy sources, see http://europa.eu/rapid/pressReleasesAction.do?reference=IP/09/117&format=HTML&aged=0&language=EN&guiLanguage=en