Pan-European Grid eInfrastructure for LHC Experiments at CERN - SCL's Activities in EGEE

Aleksandar Belić, Antun Balaž, Aleksandar Bogojević
Scientific Computing Laboratory, Institute of Physics Belgrade
http://scl.phy.bg.ac.yu/

Abstract – *EGEE project aims to provide seamless access to computing and storage resources distributed in the pan-European Grid eInfrastructure to scientific, research and development communities worldwide. One of the prime goals of EGEE is to support computing and storage needs of LHC experiments at CERN. Scientific Computing Laboratory of the Institute of Physics Belgrade participates in this project, providing resources to ATLAS and CMS communities and ensuring that Serbian HEP researchers can access all distributed computing resources available to ATLAS and CMS. In this paper the overview of SCL's activities in EGEE is given.*

1. EGEE Project Overview

The Enabling Grids for E-sciencE [1,2] (EGEE) project brings together scientists and engineers from more than 240 institutions in 45 countries world-wide to provide a seamless access to Grid infrastructure for e-Science, available to participating scientific communities. Building upon the success of the preceding DATAGRID project, the first two-year phase of EGEE started in April 2004, and its second two-year phase, EGEE-II, in April 2006. The third phase of the project is expected to start in April 2008. EGEE-II is co-funded by the European Commission through the EU Framework Programme 6 (Activity Area FP6-INFRASTRUCTURES), which contributes one half of all eligible costs incurred. The other half is provided by participating institutes and their funding agencies.

Expanding from the two initial scientific fields, high energy physics and life sciences, EGEE now integrates applications from many other scientific fields, ranging from geology to computational chemistry. However, the majority of EGEE eInfrastructure is actually part of LCG [3] (LHC Computing Grid), and is meant for direct use of researchers participating in LHC experiments, as well for providing storage for measurement results once LHC starts with its operation. Generally, the EGEE Grid infrastructure is ideal for any scientific research especially where the time and resources needed for running the applications are considered impractical when using traditional information technology infrastructures.

The EGEE Grid consists of over 47,000 CPUs available to users round-the-clock, in addition to about 30 PB disk (30 million Gigabytes) plus tape MSS of storage, and maintains over 40,000 concurrent jobs on average. Having such resources available changes the way scientific research takes place. The end use depends on the users' needs: large storage capacity, the bandwidth that the infrastructure provides, or the sheer computing power available. A snapshot of EGEE Grid eInfrastructure is shown in Fig.1.

EGEE project manages, extends and consolidates the infrastructure, to link national, regional and thematic Grid efforts and provide interoperability with other Grids, establishing a world-wide Grid infrastructure. EGEE's infrastructure is built on top of the pan-European research network GEANT2, which allows access to a common pool of major storage, computing and networking facilities, independently of geographic location of those resources and facilities, as well as users. The resulting high capacity infrastructure greatly surpasses the capabilities of local clusters and individual resource centres, providing a unique tool for collaborative computing-intensive science (eScience).

This is all possible through the development and deployment of a special layer of software which integrates all available computing and storage resources into a unique resource. Such approach not only allows access to all resources, but also hides complexity of its implementation from end users. The usual name for this layer of software is middleware, and EGEE's flagship product is gLite open-source middleware [4], which is deployed on all Grid sites participating in the project. The details on the usage of EGEE resources can be found in the gLite User Guide [5].

User specific services from other projects and sources are integrated in EGEE's existing gLite middleware to further increase the functionality, security and ease-of-use.

One of important regional Grid projects is SEE-GRID [6], whose aim is to bridge the digital divide between South East European (SEE) region and the rest of Europe, and to provide access to Grid eInfrastructure to SEE researchers. Access to national computing and storage resources is provided by Academic and Educational Grid Initiative of Serbia (AEGIS) [7].

2. EGEE Project Activities

The activities of the EGEE project are divided into the several categories, each of which addresses the specific set of issues:

- Networking Activities (NA) provide management of the project, public outreach and communication, training and support of users, as well as policy and international cooperation;
- Service Activities (SA) provide support in deployment and management of eInfrastructure and networking support, as well as testing and certification of gLite middleware;

- Joint Research Activities (JRA) develop gLite middleware and work on quality assurance and security coordination.

The EGEE-II **Networking Activities** consist of:

- NA1: Management of the I3
- NA2: Dissemination and Outreach
- NA3: Training and Induction
- NA4: Application Identification and Support
- NA5: Policy and International Cooperation

NA1 (Management of the I3) provides high quality management which is the key to the success of the EGEE Integrated Infrastructure Initiative (I3).

NA2 (Dissemination and Outreach) in EGEE-II builds on its successful track record from EGEE, continuing to spread the fundamental messages of the project in a variety of ways, and helping to recruit new users and partners. In EGEE-II the project is moving from a development model to a role as major Grid service provider, and NA2 tunes the dissemination, outreach and communication activities to this new situation, supporting cross-activity groups and user communities, helping to foster a community feeling among both Grid users and resource providers.

NA3 (Training and Induction) aims to provide a source for acquiring the knowledge and skills needed to allow the utilisation of the EGEE infrastructure at all levels.

This means EGEE-II must:

- Educate developers of applications in order to help enrich the provision of functionalities on the EGEE infrastructure;
- Educate users and developers in the potential of Grid architectures to allow the solution of appropriate problems;
- Educate site administrators to allow them to install EGEE middleware and connect to the e-Infrastructure;
- Enhance and encourage the corporate spirit of EGEE-II and information sharing within the organisation in order to help further its goals.

NA4 (Application Identification and Support) has as its overall objective to help new communities to make use of the EGEE Grid infrastructure as part of their computing model. The most important result of the NA4 activity is the production use of the EGEE infrastructure by a large number of diverse applications. The activity drives the evolution of Grid technology through specific, challenging applications, and demonstrates that EGEE-II provides a viable computing infrastructure for several scientific research communities. These applications require, in particular, that the Grid middleware scales with a growing infrastructure, performs faster, and provides additional high-level services. EGEE-II will support these applications with dedicated manpower to ensure the maximum mutual benefit.

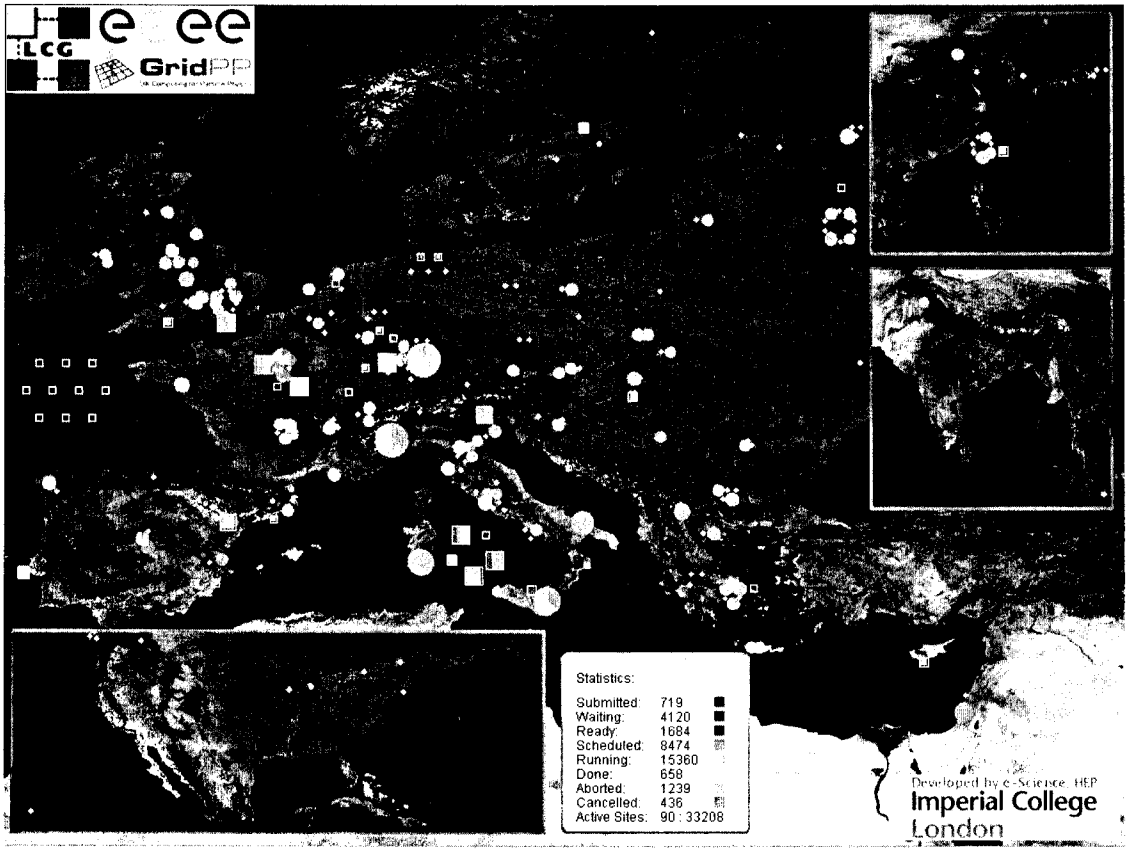

Fig. 1: Geographical distribution of the pan-European EGEE Grid eInfrastructure.

NA4 (Application Identification and Support) has as its overall objective to help new communities to make use of the EGEE Grid infrastructure as part of their computing model. The most important result of the NA4 activity is the production use of the EGEE infrastructure by a large number of diverse applications. The activity drives the evolution of Grid technology through specific, challenging applications, and demonstrates that EGEE-II provides a viable computing infrastructure for several scientific research communities. These applications require, in particular, that the Grid middleware scales with a growing infrastructure, performs faster, and provides additional high-level services. EGEE-II will support these applications with dedicated manpower to ensure the maximum mutual benefit.

NA5 (Policy and International Cooperation) activity gives support and contributes to the work of the e-Infrastructure Reflection Group (e-IRG) related work and deliverables such as white papers, roadmaps and workshops. Further, it provides roadmaps for the next generation Grid infrastructure. Currently it is envisaged that the infrastructure will be run by a European Grid Initiative (EGI) in cooperation with the National Grid Initiatives (NGIs) of each country following the research networking paradigm. NA5 also fosters cooperation between EGEE-II and the regions around Europe, such as the Baltic States and the Mediterranean, and beyond, such as US, Latin America and the Asia-Pacific region, including Japan and China.

The EGEE-II **Service Activities** consist of:

- SA1: Grid operations, support and management
- SA2: Networking support
- SA3: Integration, testing and certification

SA1 (Grid operations, support and management) activity operates, supports and manages a production quality Grid infrastructure that makes computing, storage, and informational resources accessible to users. Different user communities and Virtual Organisations (VOs) are provided with access to facilities at many resource centres across Europe and at other collaborating sites in a consistent way, according to agreed access management policies and service level agreements. The activity leverages existing national and international Grid initiatives and builds upon the work in EGEE. The key objectives of this activity are:

- The operation of Core infrastructure services;
- Grid monitoring and control;
- Middleware deployment and introduction of new resources;
- Resource and user support;
- Grid management;
- International collaboration;
- Capture and provision of requirements.

SA2 (Networking support) activity is built around the fact that EGEE-II uses the European research networks to connect the providers of computing, storage, instrumentation and applications resources with users in Virtual Organisations. SA2 provides support for resolving of any network-related issue and contributes to the development based on bandwidth-on-demand requirements.

SA3 (Integration, testing and certification) activity has the goal to manage the process of building deployable and documented middleware distributions, starting by integrating middleware packages and components from a variety of sources. It provides well defined criteria for accepting components, and runs a testing and certification activity to ensure that the middleware is as reliable, robust, scalable, and as usable as it can be. It also provides missing tools where appropriate to its staffing level and negotiates with middleware providers, both within the EGEE-II project and external to it, to provide missing functionality and services.

The EGEE-II **Joint Research Activities** consist of:

- JRA1: Middleware re-engineering
- JRA2: Quality Assurance and Security Coordination

JRA1 (Middleware re-engineering) activity works towards fulfilling the vision of EGEE-II for Grid middleware: to provide a reference open source implementation of the foundation services that are application independent and need to be deployed at all sites connected to the infrastructure. On top of this foundation, an open-ended set of application specific higher-level services that can be deployed on-demand at specific sites are either be provided directly by the project or integrated from other sources and projects. In EGEE, JRA1 produced a number of middleware releases (called "gLite") comprising security services, information and monitoring services, data services, job management services, and helper services.

JRA2 (Quality Assurance and Security Coordination) activity deals with the fact that EGEE-II has demanding quality goals which have to be quantified in advance so that they can serve as measures of success of the services provided by the project. The JRA2 activity coordinates the overall Quality Assurance aspects of the project to ensure its processes, services and deliverables are of high-quality. JRA2 also hosts overall security coordination for the project.

More details on all EGEE activities can be found on EGEE Technical Web site [8].

3. SCL's Participation in EGEE

Institute of Physics Belgrade (IPB), through its Scientific Computing Laboratory (SCL) is coordinator of the Serbian National Grid Initiative, Academic and Educational Grid Initiative of Serbia (AEGIS). SCL joined EGEE project in May 2005 as unfunded partner, just providing resources to the EGEE Grid eInfrastructure. This was a direct result of the successful participation of SCL in the SEE-GRID project. After the first phase of both projects, SCL continued its participation in SEE-GRID-2, leading Grid operations Workpackage, and joined EGEE-II project as a full partner, representing both IPB and AEGIS.

Main concerns of SCL within EGEE-II are focused on providing local Operations and Support Centre for Serbia; cooperating with EGEE South Eastern Europe Regional Operations Centre (SEE ROC) in operating the distributed Regional Operation and Support Centre, and organizing local support for training and dissemination activities. SCL is involved in four EGEE activities: NA2, NA3, NA4 and SA1.

Within **NA2**, SCL is responsible for identification of target user communities and relevant applications, and for providing dissemination and outreach activities in Serbia, including translation and distribution of EGEE press releases, and relations with the media, as well as for creation of a local EGEE-II web site. SCL's staff regularly participates in EGEE meetings. Especially intense media activity was devoted to different aspects of visits of EU DG Research Commissioner Janez Potocnik, Minister of Science Aleksandar Popović, and Minister of Finance Mlađan Dinkić to IPB and SCL in July 2006. SCL/IPB also participated in and supported the organization of Montenegro NGI inauguration and dissemination event in November 2006.

Providing of training material and organizing of training events in Serbia covering the AEGIS partners is the objective of SCL's activities within **NA3**. We contributed to the coordination, material preparation and implementation of gLite training for site administrators in Kragujevac and Novi Sad, as well as to several training events organized at IPB.

Within **NA4**, SCL provides support to HEP users and applications (ATLAS, CMS) for Serbia, identifies users and work towards introducing new user communities. In addition to the installation, configuration, and maintenance of its Grid site and all core services, SCL provided support to the local HEP community, including Belgrade ATLAS and CMS groups, and support to regional and national VOs (SEE, AEGIS). Our glite-WMSLB is used by the ESR VO as one of main WMS servers. We were also testing stability and performance glite-WMSLB. Two papers on this were presented and published in proceedings of INDEL 2006 conference, held in Banjaluka, Bosnia and Herzegovina, November 2006, and in proceedings of INFOTEH07 conference, held in Sarajevo, Bosnia and Herzegovina, March 2007. Regarding the gLite middleware development, we identified and reported several problems. As a result of our support and debugging effort, patches for some of these problems are already in certification. SCL actively participates in VO Managers' Group within NA4, and was also selected as a member of VO Deployment Working Group, created by the Technical Coordination Group of EGEE.

In **SA1** activity, SCL is responsible for user support and resource induction in Serbia. SCL also provides support to the local users and for the integration of AEGIS resource centres into the EGEE infrastructure. We are also in charge of interacting with SEE ROC and other SEE partners for the smooth operation of the distributed ROC. As a part of SEE ROC, SCL contributes to ROC management, management of resources, and ROC representation and coordination. SCL directly works on regional coordination of gLite deployment, and is also active in the regional support for operational problems and their oversight and management. We provide VO support and are very active in providing support through the regional SEE helpdesk. We also provide support in managing the installation and configuration of Grid sites within the ROC, support in monitoring deployment, and take the active role in resolving of deployment issues. We run RB, MyProxy, glite-WMS, and top-level BDII services, used by regional ATLAS, CMS and SEE users, as well as by ESR VO. Our glite-WMS was temporarily used during August 2006 for official SFTs sent from CERN to all gliteCEs. It is also used for certification of sites within the ROC. We provide regular feedback to SEE ROC regarding middleware deployment, and present our experiences on EGEE-SEE Wiki [9]. Feedback to SEE ROC is also provided regarding all operational and service management issues encountered. We addressed accounting issues within the SEE ROC, isolated several accounting-related problems and suggested solutions, including regional verification of gCE accounting configuration. We were also working on test deployment of SL4 WNs (both i386 and x86_64 architectures), have provided detailed guides for their installation and deployment, and have been the first within the SEE region to deploy them.

The third phase of EGEE project, EGEE-III, is in preparation for submission to the FP7-INFRASTRACTURES-2007-2 call for proposals. SCL will be partner in this project, continuing with the activities it performed in EGEE-II.

4. SCL's Resources Committed to EGEE

SCL participates in EGEE-II pan-European Grid eInfrastructure with its site AEGIS01-PHY-SCL, This Grid site is also part of the regional SEE-GRID-2 infrastructure, and of the national AEGIS infrastructure. Apart from computing and storage resources, it provides core services, to all these infrastructures, enabling seamless access to resources. AEGIS01-PHY-SCL, as the largest, is the Tier-0 site in AEGIS infrastructure (comprising of 5 additional smaller sites), providing all core services and manages AEGIS VO. All computing nodes and core services nodes at SCL's Grid site are interconnected by the star topology Gigabit Ethernet network through three stacked high-throughput Layer 3 switches (3COM 3870), each node being connected to the switch by two Gigabit Ethernet cables in channel bonding.

Currently 41 nodes is included into the AEGIS01-PHY-SCL Grid site, providing 132 CPUs in computing nodes amounting to 130 kSI2K computing power according to the SpecInt benchmark. Eight nodes provide core services necessary for the operation of gLite-based Grid infrastructure (BDII, RB, WMS, MyProxy, VOMS, LFC, UI). Additional 8 computing nodes (each with two dual-core CPUs, totalling to 32 CPU cores) are just being assembled to the cluster. This will add additional 57 kSI2K of computing power, amounting to overall 187 kSI2K and 164 CPUs. Substantial upgrade to SCL's high-

performance computing facilities is expected until the end of this year, funded by the Serbian National Investment Plan. Fig. 2 shows the time evolution of computing resources at SCL over the last several months, as well as the usage of AEGIS01-PHY-SCL Grid site. As can be seen, from the bottom graph, the site is overloaded with many jobs, waiting to be executed on SCL's cluster. Those jobs are submitted for execution by one of user communities supported by AEGIS01-PHY-SCL (ATLAS, CMS, AEGIS, ESR, SEE, SEEGRID).

In terms of storage resources, SCL just acquired new storage element which will provide 27 TB of disk space to supported VOs, organized physically in three disk servers. Through the Disk Pool Manager interface, this storage element will be accessible as a unified disk space, allowing seamless usage by the end users.

These latest equipment upgrades are made possible by FP6 project CX-CMCS [10], whose aim is to transform SCL into a Centre of excellence in the fields of Grid computing and simulations of complex systems.

As is already evident from the bottom graph of Fig. 2, AEGIS01-PHY-SCL is an extremely sought-after site, due to its quality and reliability. Fig. 3 gives a comparison of the number and structure of the jobs executed on Grid sites in the SEE region. We can see that AEGIS01-PHY-SCL plays a central role within the region, as well as that its workload is primarily devoted to national and regional VOs (AEGIS, SEE, SEEGRID) on the one hand, and to HEP VOs (ATLAS, CMS) on the other. It is important to note that high energy physics jobs play a central role within AEGIS VO as well, primarily through extensive path integral Monte Carlo simulations of systems of many degrees of freedom [11].

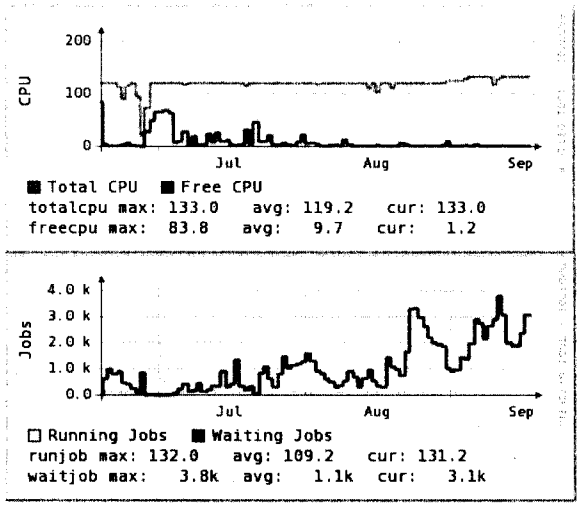

Fig. 2: SCL's resources committed to the EGEE Grid eInfrastructure: No. of CPUs (top); No. of jobs (bottom).

Detailed accounting data is presented in Table 1, where the elapsed CPU hours are given across user communities and SEE countries. It is seen that HEP usage is predominant throughout the region. The already mentioned reliability and stability of AEGIS01-PHY-SCL site can be seen directly from the fact that Serbia provides comparable total accounting time as Romania and Turkey, countries with substantially greater numbers of CPUs. Overall, the SEE region dedicates 53% of its accounting time to high energy physics usage, outside of national VOs. It is important to note that the SEE infrastructure plays a key role within the pan-European Grid effort: from the start of the EGEE-II project, SEE sites have provided around 20% of the overall accounting time.

COUNTRY	HEP	National	Regional	Other	TOTAL	Percentage
Serbia	336,625	479,295	128,508	45,866	990,294	8.39%
Bulgaria	476,328	0	174,351	179,573	830,252	7.03%
Cyprus	52,548	0	76,345	53,916	182,809	1.55%
Greece	3,177,134	0	1,168,242	2,037,093	6,382,469	54.07%
Israel	175,284	0	67,174	104,079	346,537	2.94%
Macedonia	0	0	6,535	8,389	14,924	0.13%
Romania	1,567,099	0	73,085	57,993	1,698,177	14.39%
Turkey	458,673	451,706	218,742	229,542	1,358,663	11.51%
TOTAL	6,243,691	931,001	1,912,982	2,716,451	11,804,125	100.00%
Percentage	52.89%	7.89%	16.21%	23.01%	100.00%	

Tab. 1: The accounting time in elapsed CPU hours committed to the EGEE Grid eInfrastructure by the South Eastern Europe countries in the period April 2006 – September 2007. The distribution of accounting time committed by each country is given according to categories: AEGIS (Serbian national AEGIS VO usage), HEP (high energy physics – ATLAS, CMS, ALICE, LHCB), Regional (SEE, SEEGRID, Other (all other EGEE VOs, including national VOs of other countries), The percentage column gives the share of accounting time provided by each country, while the percentage row gives the share of each category in the overall accounting time.

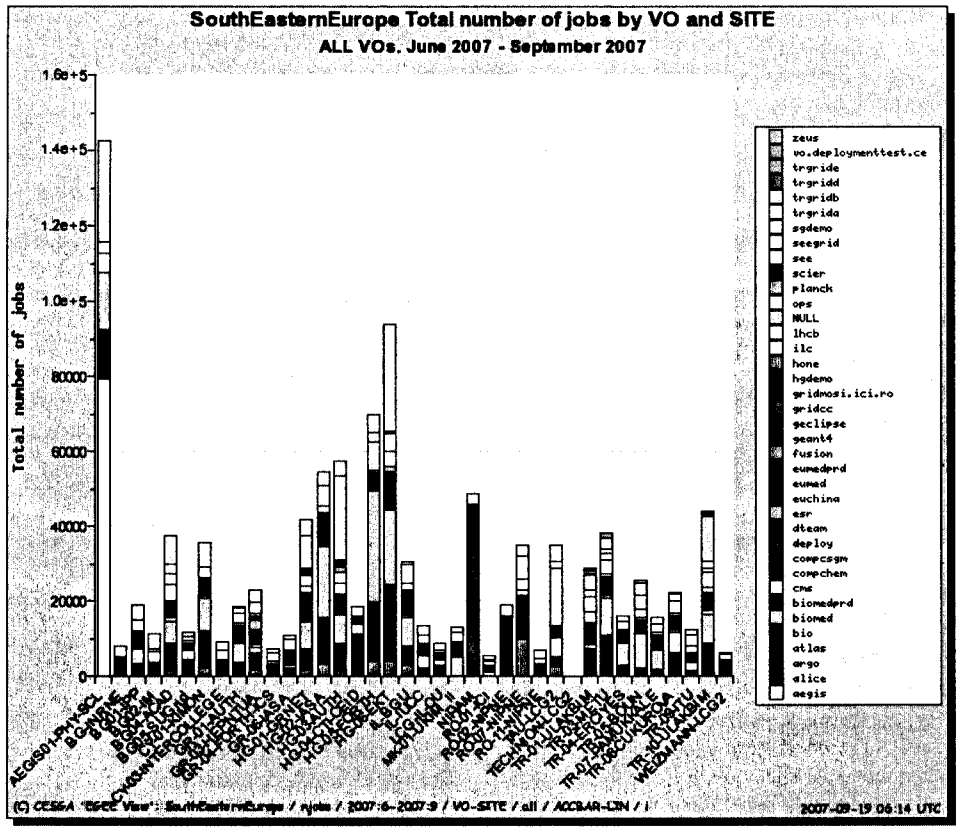

Fig .3: The total number of jobs executed on the EGEE Grid eInfrastructure in the region of the South Eastern Europe by site and by VO from June to September 2007

5. Conclusion

Serbia's effort within European Grid development and exploitation is substantial, particularly if one has in mind the fact that this whole contribution is provided by a single high-performance computing facility at SCL, as well as that Serbia is a relative late-comer to production-level Grid infrastructure. It is to be expected that Serbia's contribution to Grid infrastructure will grow in the coming years. This can be seen from the pace of upgrades of AEGIS01-PHY-SCL, and even more importantly from the robust growth of the national Grid initiative AEGIS in terms of participating institutions and deployed infrastructure. We expect that AEGIS will grow to a fully operational and self-sustainable national Grid infrastructure, to be used by researchers from Serbia. This growth potential is clearly reflected in the increasing role played by AEGIS01-PHY-SCL site within the regional SEE Grid eInfrastructure as is seen in Fig.3 showing the number of executed jobs per site during the last quarter.

HEP has been the initiator of the Grid paradigm in distributed computing. At global and European levels CERN's LHC experiments continue to be the principle users, as well as the driving force behind the quest for stable, production-level Grid infrastructures. The central role of HEP is mirrored at the level of the SEE region as a whole, and in Serbia. The Grid-ready applications used by Serbia's HEP community, as well as its size and level of development, are strong indicators that they will remain key users of Grid infrastructures in the foreseeable future.

Acknowledgements: This work was supported in part by the Ministry of Science of Serbia through the project OI141035, and by the European Commission through FP6 projects CX-CMS (INCO-026343), EGEE-II (INFSO-RI-031688) and SEE-GRID-2 (INFSO-RI-031775).

References

[1] http://www.eu-egee.eu/
[2] http://egee.phy.bg.ac.yu/
[3] http://lcg.web.cern.ch/LCG/
[4] http://glite.web.cern.ch/glite/
[5] https://edms.cern.ch/document/722398/
[6] http://www.see-grid.eu/
[7] http://aegis.phy.bg.ac.yu/
[8] http://egee-technical.web.cern.ch/egee-technical/
[9] http://wiki.egee-see.org/index.php/Main_Page
[10] http://cx-cmcs.phy.bg.ac.yu/
[11] A. Bogojević, A. Balaž, and A. Belić, Phys. Rev. Lett. **94**, 180403 (2005); Phys. Rev. **B72**, 064302 (2005); Phys. Lett. **A344**, 84 (2005); *ibid.* **A345**, 258 (2005); Phys. Rev. **E72**, 036128 (2005); Phys. Low Dim. Struct. **1**, 49 (2006); *ibid.* **1**, 52 (2006).

D. Stojiljković, A. Bogojević, and A. Balaž, Phys. Lett. **A360**, 205 (2006).

J. Grujić, A. Bogojević, and A. Balaž, Phys. Lett. **A360**, 217 (2006).