

The BEgrid Tribune

About BEgrid

BEgrid is the Belgian computing grid infrastructure for research featuring about 500 CPUs and 4 Tbyte of storage. This grid initiative is coordinated by BELNET and strongly supported by the Flemish Government.

Universities and high schools participate in BEgrid by providing computing resources and manpower to manage their part of the grid.

For more information: see <http://www.begrid.be/>



Editorial

Dear Reader,

I am happy to present you the first issue of the BEgrid Newsletter. Some of you already know what BEgrid is and are aware of its activities. For those of you that do not know: BEgrid is the Belgian computing grid for research, coordinated by BELNET. See the "About BEgrid" on this page.

We want to communicate more and better about this important initiative and from now on you will find this online publication two times a year in your inbox. Its goal is to inform you about the evolution of the BEgrid infrastructure and grid related activities, especially those centred around BEgrid. We are also aware that we certainly do not reach all researchers who might be interested in grid computing and in the use of the BEgrid infrastructure but you can help us by distributing this newsletter to your colleagues, friends, boss, ...

This first issue is entirely dedicated to student work on the theme of grid computing. This work includes thesis work carried out to obtain a master degree in computer science or engineering or projects done in the last year of a "High School". Most of the work is done by institutes that already participate in BEgrid. It is highly probable that more research in grid computing, of which we are not aware, is going on in Belgium.

Feel free to send us more information about what is happening in your institute about grid computing and we will be happy to include that information in a next issue.

We hope that you will read through the whole newsletter and automatically arrive at the announced grid events. The next Belgian grid event takes place on the 12th of June at Ghent University and is realised in collaboration with BEgrid. More information will soon be available, look at <http://www.begrid.be/>.

We wish you a good reading.

All comments and suggestions can be sent to rosette.vandenbroucke@belnet.be

Rosette Vandenbroucke
BEgrid coordinator

Development of Grid applications at the Hogeschool Antwerpen

Distributed motion detection calculations on the BEgrid Development of a grid portal for PDA devices

During the academic year 2005-2006, three last year students, Kristof Doms, Maikel Michielsen en Robby De Maarsschalck, of the department Elektronica-ICT of the Hogeschool Antwerpen (HA), worked, under the supervision of Prof. Filip Van der Schueren and Prof. Jan Broeckhove (UA), on projects where a computational grid (BEgrid) was the key subject.

Their project was carried out in collaboration with the CoMP group of the Universiteit Antwerpen (UA).

As a first prerequisite, they had to fulfill a complete build up of a grid system, according to BEgrid specifications; their first installation was used to acquire the necessary knowledge to build the grid software infrastructure using a Quattor server, a system administration toolkit for automated installation, configuration and management of Unix clusters. During this setup, they spent some time in building the software repository and editing template files from scratch. When the LCG-2 middleware needed an upgrade to version LCG 2.7, a second and more automated installation was realised making use of the centralised BEgrid configuration server. They just had to setup a local Quattor server, together with a small number of site specific template files. This local Quattor server behaves first as a Client to the Central Quattor server and becomes later the site specific installation server.

As a first conclusion, they proved that the new automated installation, by using the Centralised BEgrid configuration server, makes life easier when setting up a new BEgrid infrastructure.

While setting up their own grid infrastructure was a starting point, they also worked on two applications.

A first application was the realisation of a video surveillance system by means of IP cameras. The application consisted of three separate parts, all written in C. The first part controlled the recording of a stream of jpeg files from a number of IP cameras. A carefully designed meta information structure was needed to record not only the images, but also camera and time information. This data was recorded using XML. The second program then calculated the - occasional - motion between images, parameterised by some threshold parameters; motion occurrence was again written to an XML structure. The last program was able to deliver a movie of a selected set of images, together with a timestamp, to make the motion occurrence concrete.

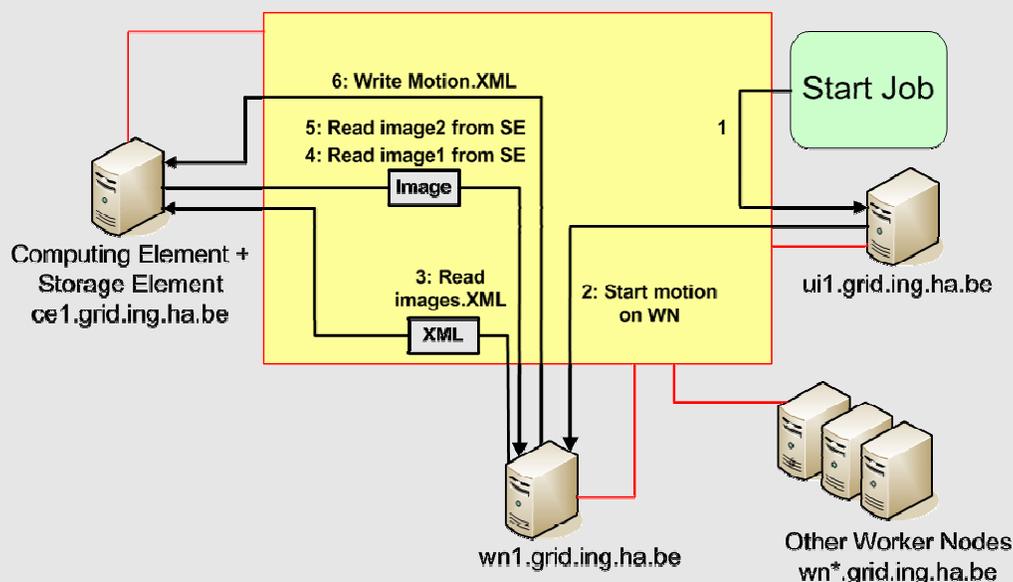


Figure 1: the flow of activities when calculating motion

The motion detection project showed that grid computing is not only limited to scientific applications, but proves that the BEgrid infrastructure can also be successfully applied to more general applications, where dispersed data and processing power is needed.

The second application was the design of a grid portal, specially made for a PDA. When the appropriate security requirements were met, the grid portal allows users to interact with the grid's resources, more specifically to monitor, to start, query and stop grid applications from the PDA. Because the emphasis was on portability, the information and layout were adjusted for small screens. To work operating system independently, Web services and XML were used to build the system.

This last project showed that with a limited knowledge of grid computing, users can easily make use of grid resources, even in a mobile environment.

A more detailed description of both projects can be found in a 6 pages paper, which is available online at <http://www.e-lab.be>; the paper can also be requested by e-mail: f.vanderschueren@ha.be.

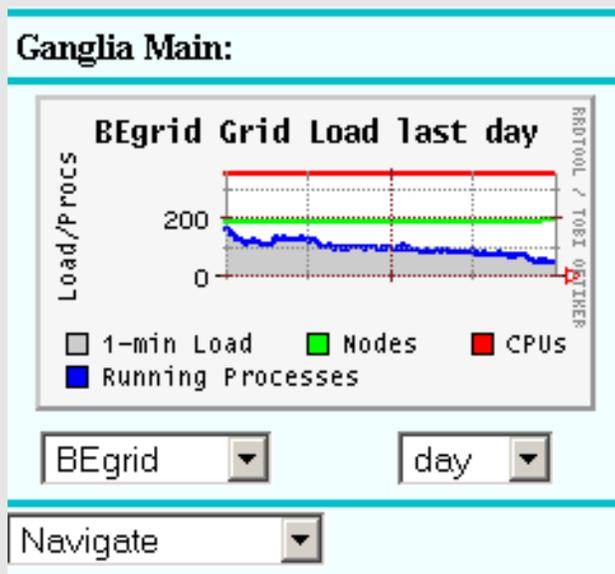


Figure 2: a small monitoring diagram, created by ganglia, and viewed on our PDA device



An application at Ghent University that made use of BEgrid

Yue Wang defended the thesis “Fast Monte Carlo based reprojector for iterative fanbeam SPECT reconstruction” to obtain a master in engineering. Promoter was Yves D'Asseler with Co-Promoter Ignace Lemahieu and supervisor Steven Staelens

Nuclear medicine imaging involves the use of radioactively labelled pharmaceuticals to diagnose and assess disease in the human body. Single photon imaging (SPECT) requires at minimum just one detector fixed in one position to obtain a two dimensional projection of a three dimensional radiopharmaceutical distribution. A typical imaging system suitable for this task is the gamma camera. A typical gamma camera consists of a collimator which limits the angle of incidence of the detected photons. The radionuclide mostly used in SPECT is Tc-99m which has a half life of 6.03h and which emits photons of 140.5 keV. SPECT application fields in order of importance are cardiac imaging, bone scanning, neurology, renal and liver function imaging.

An estimate of the three dimensional radiopharmaceutical distribution needs to be reconstructed from the aforementioned acquired two dimensional projections. State-of-the-art reconstruction makes use of the Monte Carlo technique. Monte Carlo methods are statistical simulation methods wherein a sequence of random numbers is applied. In order to do Monte Carlo calculations some a priori information about the occurring physics processes is needed; this information is expressed in the form of probability density functions (pdfs).

Monte Carlo simulations are CPU-bound computations, in the sense that most of the work corresponds to numerical calculations related to particle transport. Thus, simulations usually take up to several days to complete with state-of-the-art single-CPU computers. However, it is also generally true that Monte Carlo simulations are well suited for parallelized solutions to the computational problem

In order to reduce the overall computing time of Monte Carlo experiments, a parallel computing platform for running such simulations, in a cluster of computers was developed which significantly shortens the setup time and provides fast data output handling. The approach is platform independent in the sense that the simulations are virtually separated so that the user obtains a number of fully resolved independent job execution macros accompanied by a platform specific submit file. Moreover, the software is fully automated and requires no interaction whatsoever from the user. Therefore it is an excellent candidate for deployment on the BEgrid which enables us to fully exploit the scalability and offers the opportunity to improve the clinical relevance of our datasets.

Our group has designed iterative reconstruction software using a MonteCarlo reprojector for mainstream SPECT scanners. However, for brain imaging, fanbeam collimators are used to increase the sensitivity in order to counteract a low activity uptake. Main indicated pathologies are Alzheimer, Parkinson, ADHD, OCD and addiction. It is the purpose of this thesis to extend the existing reconstruction software to these fanbeam scanners.



Ant Colony Optimization for the Routing of Jobs in Optical Grid Networks

Sofie Demeyer from the Ghent University worked on “Ant Colony Optimization for the Routing of Jobs in Optical Grid Networks” to obtain her master in engineering.



Ant Colony Optimization is based on the natural behaviour of ants foraging for food. By leaving a trail of pheromone ants can give other ants of their colony directions to the closest food resource. We translated this concept to be used for the routing of jobs in optical grid-networks. In different sources in the network digital ants are generated. Their task is to travel to the network looking for available resources. When a resource is found, the ant will return on the same path leaving some information (“the pheromone trail”) in each of the routers on the way. Ants that find the closest resource can return more quickly to their source

than ants that find a resource further away. Because more ants can visit this closest resource in a certain amount of time the pheromone trail on this path will be the strongest. A job searching a resource in the network will make use of the information left by the ants in the routers. So the job will be attracted to the closest resource because the pheromone trail on this path is the strongest.

In this way jobs can be routed dynamically in grid networks without the use of an intermediate broker. We developed different ways to implement this ant colony optimization and results have shown that this anycast routing performs better than the usual unicast routing.



OGSA Mapping of Jini Concepts



To obtain a master in Computer Science, W. Cuypers worked on “OGSA Mapping of Jini Concepts”. His promoter was Prof. Jan Broeckhove and supervisor was Kurt Vanmechelen.

The development of computing grid architectures aims at providing an answer for the computation requirements of researchers in a robust, secure and efficient way. Very often these requirements are too demanding to be satisfied with the provision of local architectures like clusters and multiprocessor systems. The solution is to switch to distributed architectures that integrate available resources on a large scale.

The conception of such architectures is a worldwide research topic where a large number of research groups contribute. To further interoperability between different grid systems the OGF (Open Grid Forum) developed the Open Grid Services Architecture (OGSA). This standard maps a grid architecture to a service oriented model that focuses on web-service standards.

Sun Microsystems has developed the Jini framework to cope with inherent problems that occur during the development of distributed systems in a unified framework.

The development of the Jini framework started in 1997; mid 2003 a revision that had a strong focus on protocol flexibility and security aspects was made available. Jini is based on Java concepts and extends these concepts to the network level creating in this way a completely service-oriented programming model.

Some current research in grid computing aims at building a service-oriented grid architecture by using Jini.

This thesis will try to formulate an answer to the question “Can the concepts of the Jini framework be mapped on the OGSA standard?” This answer will indicate if it will be possible to make Jini based grid architectures interoperable with other grid systems.



Master theses about grid computing conducted at the ULB

At the Université Libre de Bruxelles, three groups have conducted three joint master theses during the academic year 2005/2006: the groups of Particle Physics and Distributed Systems from the Faculty of Sciences as well as the Computer and Network Engineering group of the Faculty of Applied Sciences. This work was done under the supervision of Daniel Bertrand, Othmane Bouhali, Raymond Devillers, Joël Goossens, Pascal Vanlaer, Shkelzen Rugovac, Stijn De Weirdt and Esteban Zimanyi.

Jean François Roche worked on “Grid and cluster monitoring”. A monitoring system should be able to track the different grid events in real time. The main concepts of the subject were introduced as well as related grid requirements. Two monitoring solutions were deployed and compared. Future improvements were also proposed.

Jérôme Vos concentrated on “Etude comparative d’un gestionnaire de stockage pour la ferme de calcul du service de Physique des Particules dans le cadre de l’expérience CMS au collisionneur LHC”. This work was performed in the framework of the Belgian participation to the CMS experiment at the Large Hadron Collider at CERN. The thesis focused on the test and installation of a distributed storage solution called dCache. This latter has proven to be capable of managing the storage and exchange of several hundreds of Tera Bytes of data, transparently distributed among many disk storage nodes. The work performed also included the participation to the successful performance tests conducted by the CMS worldwide grid platform.

The thesis of Guillaume Desmottes “Déploiement et configuration des intergiciels européens de grilles de calcul” has been nominated for the “AScBr award”. AScBr is the « Association des diplômés en Sciences de l’Université Libre de Bruxelles ». It addressed the installation, configuration and administration of grid middleware. The EGEE middleware gLite has been at the core of this study. Several key services have been tested and problems were solved and/or reported. Another important contribution was the effort to adapt the middleware to the centralized CERN installation and administration suite called Quattor.



BEgrid and the Royal Military Academy



To obtain the degree of Polytechnical engineer in Telecommunications Christophe Joris worked on the thesis “Construction d’un système parallèle et analyse de sa performance” under the direction of Major Mees.

The goal of this work was to do a performance analysis of different parallel architectures in order to compare those systems.

As the time available for this work did not permit to install a gridcluster with the BEgrid middleware it was decided to work with a computing cluster installed at the RMA (Royal Military Academy) and with the BEgrid infrastructure to which access could be obtained. The algorithm that was chosen to compare both architectures was the Gauss elimination method and its parallel implementation was used. The Gauss elimination method is an algorithm that works with a techni-

que of triangularization of matrices to resolve a system of linear equations. Work had to be done to measure execution times but also to check that the obtained results were correct. It was also necessary to compare the execution times of the parallel and the sequential version of the algorithm in order to study the overhead introduced by the addition of the parallelism.

The work carried out showed that the use of the BEgrid infrastructure delivered the best results. However it has to be said that the cluster of the RMA could give better results with an extension of the memory of the systems and with a better adapted operating system. It has also to be noted that at BEgrid there is a waiting time before a job goes into execution and hence BEgrid is more useful for jobs with a large execution time. Small jobs have a quicker turn-around time on the local cluster.



The integration of BEgrid in CoBRA A development from the Universiteit Antwerpen



Developments in the field of computing grids were growing these last years. The range of applications in the research world that make use of grid computing has also broadened since the evolution of the EU DataGrid (EDG) to Enabling Grid for E-sciencE (EGEE) in March 2004. In a first phase those European grid infrastructures only supported high energy physics and biomedical experiments but now astrophysics, computational chemistry, geology and others use grid computing cycles. The success of EGEE can be accounted to the large infrastructure that has been built up during the lifetime of the project and the help of the European Commission for the development of middleware, security solutions. The connection of BEgrid to EGEE at the end of 2005 and the connection of grids in the USA and Asia are examples of this success.

There are probably many projects that try to integrate EGEE/BEgrid middleware into their computing solution. At the Universiteit Antwerpen a project was started to connect CoBRA to BEgrid, in the same way as BEgrid is connected to EGEE.

CoBRA is a lightweight gridsystem that has been developed at the department of Computational Modelling and Programming of the Universiteit Antwerpen. It offers the possibility to integrate as well normal desktops as clusters in a shared resource pool. In a typical scenario desktops of an organisation can run the CoBRA client that will enable the use of idle processor cycles by applications that require computing time in the CoBRA grid.

To execute a job submitted within CoBRA, as is the case with other grid middleware, a number of steps have to be taken.

First of all the job has to be assigned to a Worker Node (WN) by the Resource Broker (RB).

The WN has to fetch the input data needed to execute a job from the Storage Element (SE) before the job can be executed. Finally the output data has to be sent to the SE.

However there is a fundamental difference in the way jobs are described in CoBRA and BEgrid. In BEgrid jobs are described in a textfile that defines the required input data, the program to be executed and the output files that have to be returned. In CoBRA that same information is included in a java file. For general jobs in CoBRA it is impossible to translate this java code to a textfile that can be used in BEgrid.

A schema has been developed to be able to address BEgrid from within CoBRA. CoBRA starts to fetch the input data and sends this information together with the CoBRA job to a BEgrid SE. Then a bootstrap job is submitted by the BEgrid RB, assigned to a WN and executed. The bootstrap job will fetch the CoBRA job at the SE and start the execution. The CoBRA job will fetch the input data at the SE, finish the job and send the output data to the SE. The last step consists in fetching the output data at the SE and sending it to the desired location.

This schema is implemented in CoBRA by using the Commodity Grid Kit (CoGKit). The CoGKIT is a middleware that enables the development of high level applications that have to communicate with grid systems. Adapting to new incarnations of BEgrid will also be very easy with the CoGKit.

BEgrid Participants

BELNET- Hogeschool Antwerpen- Katholieke Universiteit Leuven - Universiteit Antwerpen- Universiteit Gent - Université Libre de Bruxelles - Vlaams Instituut voor de Zee - Vrije Universiteit Brussel



KATHOLIEKE UNIVERSITEIT
LEUVEN



Vrije
Universiteit
Brussel

Some ongoing research

More research is going on at the Universiteit Antwerpen, the following theses centered around the topic of grid economics with promotor Prof. Jan Broeckhove and supervisors Kurt Vanmechelen and P. Hellinckx, are in preparation:

- On the Integration of Combinatorial Auctions in Grid Resource Management Systems (G. Van Eyck)
- Assessment and Java port of the Market Based Resource Allocation system Tycoon (J. Delforge)
- Establishing price levels for market based resource allocation: what about agents? (W. Depoorter)
- A comparative study of commodity models and auction based models for realizing Grid economies (R. Vandenbossche)
- Grid Monitoring for the Network Weather Service (G. Heiremans)
- Moving simulated market-based allocation into the real-world: GES and CoBRA (S. Verboven)
- Studying the state-of-the-art in Grid simulation and its application to network modeling in the GES (J. Buys)
- Parameter sweep prediction (S. De Munck)

At the ULB a thesis is being made by Eric Robette under the direction of Prof. P. Van Binst and Rosette Vandembroucke about the inclusion in BEgrid of the computing power of the student PC-rooms.

Coming Grid Events

7-11 May, OGF/EGEE User Forum, Manchester, U.K.

The second User Forum will be co-located with the OGF20 meeting in May in Manchester, UK. The OGF20 (Open Grid Forum) meeting will be held May 7th-9th and the EGEE User Forum will be held in Manchester May 9th -11th. This combined event will further strengthen the links between EGEE and the Open Grid Forum, bringing users and standards bodies together to ensure that the future of the Grid is complemented by the establishment of key standards. See <http://egee-intranet.web.cern.ch/egee-intranet/User-Forum/>

12 June, Grids for Science and Business, 't Pand, Gent

“Grids” is one of the newest and most promising ICT paradigms. For computing and/or data storage, it offers unprecedented solutions at very affordable costs. Hundreds of grid development and deployment projects are active worldwide, particularly in Europe – and also in Belgium. The technology will soon move from academic and scientific circles to industry and institutions, as large commercial companies are entering the field. This one-day International Symposium aims at bringing together a number of active parties and everyone interested in the short to medium-term evolution of grid technologies.

The target public includes industry as well as academia and public institutions, as future grid applications may apply in different fields and involve many parties.

This event is organized by the Flemish Interdisciplinary institute for BroadBand Technologies (IBBT) and Ghent University (UGent) in collaboration with the Brussels universities (ULB and VUB), with the support of the projects BEgrid, EGEE and EUROLABS.

16 October, BEgrid Seminar, Brussels

The main annual BEgrid event.

Contacts

BEgrid
BELNET
Wetenschapsstraat,4 - Rue de la Science, 4
1000 BRUSSELS
Tél: +32 2 790 33 95
e-mail: begrid@belnet.be

