

KnowARC: Enabling Grid networks for the biomedical research community

Henning MÜLLER ^{a,1}, Mikko PITKANEN ^a, Xin ZHOU ^a, Adrien DEPEURSINGE ^a,
Jimison IAVINDRASANA ^a and Antoine GEISSBUHLER ^a

^a *Medical Informatics, University and Hospitals of Geneva, Switzerland*

Abstract. The vast availability of medical patient data in digital format creates the opportunity to use these data in medical informatics research projects. The objective is to improve future care by providing the medical staff with methods for automated data processing, including textual and visual information analysis and retrieval from medical databases. Many medical institutions do not possess a specific research computing infrastructure or the required budget for such an infrastructure to enable processing of these large amounts of data. Still, many institutions have many desktop PCs that could serve for biomedical research during the time they are little used without the need for expensive investments. The *KnowARC* project aims at building a middleware for such a simple-to-install Grid network.

This article reviews requirements for computing Grids in large hospital environments. We use the computing infrastructure in the University Hospitals of Geneva as an example, and then present the solutions that the European Union-funded *KnowARC* project plans to undertake to solve the current problems. Methods currently employed in common Grid middleware distributions are also reviewed and assessed in relation to the goals of *KnowARC*.

The computing infrastructure at the University Hospitals of Geneva is described as well as the needs and requirements for computing and storage services within this domain. A list of requirements for a Grid middleware to employ in such a challenging environment is developed. Finally, the proposed solutions and ideas of the *KnowARC* project are described in detail to present the project to a larger community. First proof of concept implementations and test results are described to illustrate how Grid networks are expected to become an important supplier of computational resources, which are required in several domains in biomedical research. A continuous process will be necessary to feed in the requirements of the biomedical domain to developers of Grid middleware to make the outcome meet the specific needs of the biomedical community.

Keywords. Grid networks, healthgrid, biomedical informatics research, information retrieval

1. Introduction

The concept of Grid computing networks is currently a hot topic in many domains. However, the fundamental principles are fairly old as already in the late 1980s the Condor [1] project proposed the idea to use idle computer time of simple desktop computers for

¹Corresponding Author: University and Hospitals of Geneva, 24 Rue Micheli-du-Crest, 1211 Geneva 14, Switzerland, henning.mueller@sim.hcuge.ch

computationally intensive tasks. The concept of Grids over the last years has been driven mainly by the high energy physics community and here in part by the CERN² (Centre Européen de la Recherche Nucléaire).

In addition to the early Condor project, several other projects were developed to provide distributed computing middleware packages such as Unicore [2] and Globus [3]. Perhaps the most widely known example about distributed computing is the SETI@home project³. This system is based on the BOINC (Berkeley Open Infrastructure for Network Computing) middleware, a fairly simple distribution mechanism for computational tasks.

Similar technology has been used for computing in biomedical domain to study protein folding⁴ and to help with research on diseases⁵. However, these projects process data that has no strict requirements for information confidentiality and thus a lightweight middleware can be used. When biomedical (patient-related) information is distributed, more sophisticated security mechanisms have to be employed to ensure information confidentiality. Thus, the needed security guarantees motivate the need for modern Grid computing solution that is more controllable than the BOINC type of Internet computing.

The physicists had requirement for a computing resource that is powerful enough to process the enormous amount of data that the next generation particle accelerator, the LHC (Large Hadron Collider) is expected to provide starting from 2007. For the requirements of the physics community the large European Union project EDG⁶ (European Data Grid) was started to advance the technology quickly and provide the community with standard interfaces. EGEE⁷ (Enabling Grids for E-science in Europe) provides a follow-up of this project to deploy a worldwide Grid with partners in 57 countries and many application fields, including a work package on biomedical Grids. Many national Grid network initiatives and projects are linked to the EGEE project to base their efforts on well-established standards.

Besides the large global projects, several projects have developed on a smaller scale and with a slightly different focus. The NorduGrid⁸ project started in 2000 [4]. The goal was to develop a lightweight middleware that would be non-intrusive for the reuse and coupling of clusters mainly in Northern European countries. One of the objectives was also to make a needs analysis and reuse existing components and standards wherever possible. Today, the result is the ARC (Advanced Resource Connector) middleware that is currently used in over 60 sites in many countries world-wide and is in continuous production use. Non-intrusiveness in this case means that the computers can be used for other software as well and do not require the system to run a single middleware, which is a requirement for many institutions. Dependency on particular versions of an operation system or compiler are also supposed to be as limited as possible.

In recent years, the term computing on demand has also developed fantasies in the computing industry. The recent emergence of services^{9,10} that enable users to pay only

²<http://www.cern.ch/>

³<http://setiathome.berkeley.edu/>

⁴<http://folding.stanford.edu/>

⁵<http://fightaidsathome.scripps.edu/>

⁶<http://www.edg.org/>

⁷<http://www.eu-egee.org/>

⁸<http://www.nordugrid.org/>

⁹<http://www.network.com>

¹⁰<http://aws.amazon.com/ec2>

for the used computing cycles serve as a prime example of industry interest. Maintaining a computing infrastructure is expensive and paying for only the needed amount is an attractive model that allows specialised providers to develop quality solutions. This model enables high levels of usage and reduces the management cost, which is particularly interesting for medium and small actors both in public and in private sectors. Some companies such as Novartis already use desktop computers of their computing infrastructure over night for computing the testing of new drug docking mechanisms [5] to find candidates for new medications. A common infrastructure for this use is the United Devices¹¹ middleware that is running on several infrastructures and platforms, including Windows, which is employed in many companies and larger organisations. Most scientific Grids on the other hand are based on Linux. Role model for many of these projects is to have a structure for storage and computing power similar to the information structure of the world wide web.

In the medical domain the healthgrid¹² association has brought together the biomedical informatics research community with health specialists to develop common solutions. So far, it has organised four conferences for this purpose. Several position papers identify a strong potential and the need for biomedical Grid networks in hospitals [6] has been published several times. However, many applications of Grids within hospitals including those of image retrieval are still in an early prototype phase [7,8,9] and not yet used in clinical practice and with routine data. A notable exception is the Mammogrid¹³ [10] that has combined routine storage of mammograms in several European countries. A more recent project is health-e-child¹⁴. Many global visions exist on how to combine large amounts of data from many levels (molecular, cell, organ, individual, population) to really make use of large-scale Grid networks [11] and use the available data up to its full potential. One of the biggest challenges with employing a Grid infrastructure in a non-grid-specialist environment is currently the installation and maintenance of such systems. Many solutions still require access to operating system resources on the used machines and extensive technical knowledge. Most scientific middlewares work only on the Linux operation systems, while failing to install on other widely used systems. In the medical field and in other large administrations there are other factors as well that limit the use of Grid networks, and these are mainly issues related to information security. Many computers contain confidential information, for example patient data or confidential administrative documents. A middleware to be employed on a large scale in an institution needs to take these security issues into account and also find political solutions by integrating the technical deciders. An enormous potential is surely present.

2. Methods

This article reviews the main goals of the KnowARC¹⁵ project from the point of view of the biomedical research community. The situation of the University and Hospitals of Geneva within the project is described in more detail including the challenges concern-

¹¹<http://www.ud.com/>

¹²<http://www.healthgrid.org/>

¹³<http://www.mammogrid.com/>

¹⁴<http://www.health-e-child.org/>

¹⁵<http://www.knowarc.eu/>

ing the infrastructure and also the main research goals. Finally, the first steps towards a hospital-wide computing infrastructure for biomedical research are explained and a short outlook is given.

3. Results

3.1. Computing infrastructure at the University and University Hospitals of Geneva

The University hospitals of Geneva (HUG) have an infrastructure and resource use that is similar to many other administrations. Large databases on servers (mostly on Sun Solaris at the HUG) exist that store the patient records and particularly the large image archive (PACS, Picture Archival and Communication System). The image archive that consists of several Terabytes is stored at the same time in two locations on two SANs (Storage Area Network) to provide a maximum of flexibility and security. With 2'200 beds and 10'000 employees, the Geneva University hospitals are one of the largest hospital groups in Europe. Over 50'000 images are currently produced per day and roughly a million events registered by the fully electronic patient record system. Besides the large servers, most computers in use are simple desktop machines that are centrally managed (for software and hardware distribution) and updated regularly. Roughly 6'000 computers exist in total and these machines are regularly renewed. Currently, a 5-year renewal cycle exists meaning that modern machines are Core2Duo CPUs with 2 GHz and 1 GB of RAM whereas the oldest computer currently in use are Pentium IV machines with 2 GHZ and 512 MB of RAM. A 100 MBit Ethernet network connects the computer on the distributed sites of the hospitals. The network is strongly secured and all outgoing traffic is basically blocked. All WWW traffic has to pass a proxy server for security control.

For research on the other side, no dedicated institutional infrastructure exists at all. Research projects with funding buy their own servers and additional desktop computers that are subsequently used within the project but only very rarely shared on a personal basis. If no funding for servers exists simple desktop computers are used for the research limiting the possibilities in domains such as data mining, text mining [12], and image processing and retrieval [13], where computational power can make a large difference with respect to possible technologies.

3.2. Constraints for using this infrastructure

The computing infrastructures in hospitals have stricter security requirements than usually faced within the University research in computer science faculties. This concerns data confidentiality as well as the use of the infrastructure. On the one hand there are research data sets that can not leave the hospital network to be computed, for example, if the data might allow the identification of the patient. Structures to anonymise the data are available and then, part of the data can possibly be computed outside of the hospital in a computing centre such as the Swiss National Super Computing Center (CSCS¹⁶). Several constraints were identified early that a to-be-developed middleware has to take into account:

¹⁶<http://www.cscs.ch/>

- The availability of the infrastructure is not always given and the use of the existing PCs can change quickly (PCs are turned off, they can be used, portable computers can be removed from the network).
- A heterogeneous infrastructure is characteristic for an organisation where computers are only renewed in regular cycles every five years, whereas most Grid networks are much easier to administer and load-balance if the infrastructure contains hardware of the same type.
- Running Clinical applications is the principal goal of the computer network group of the hospitals. The applications need to run 24 hours a day with no interruptions caused by computing research. Research is clearly only an independent little part inside the institution with little decisive power.
- Access to all data on computers is strictly regulated and a process running on a computer with clinical data should at no point have the possibility to access these clinical data.
- No maintenance personnel for the research infrastructure can be allocated at the hospitals as is practice in big computing research centres at Universities. Researchers need to maintain the computers and also adapt their applications for these computers, so simplicity is important.
- No dedicated computers for the computing tasks exist, so the real users of the PCs should not be slowed by these tasks; a user needs to be able to control the amount of computing power that he can make available for research tasks.
- Operating systems in the hospitals are dictated by the organisation. The very large majority of the hospital's computing infrastructure is based on Windows and particularly desktops are, whereas many middleware research projects are only applicable on Linux platforms.

To learn more about the possible challenges that introducing a Grid infrastructure might face in the administrations, we are currently preparing a survey among the personnel of the University hospitals, the University, and the Public Administration of Geneva. This survey is to identify further problems and also to educate the users about the possibilities and challenges that Grid computing networks can offer. Many of the future challenges to surpass are not of technical but rather of political nature.

3.3. Planned solutions to be provided by the KnowARC project

The early evaluation of challenges has also led to possible solutions for many of the issues and several of these solutions are currently being implemented in the KnowARC project:

- Use of an existing and working middleware as the basis with the need to only develop new components and adapt existing software toward these needs (ARC of NorduGrid, in our case).
- Using virtualisation to run Linux machines within the centrally administered Windows machines. This allows to use Grid software on existing windows desktops. The sandboxing through virtualisation allows also to clearly separate the instance of an operating system running on clinical data and the instance running the research task, which solves one important security issue.

- The installation process needs to become as easy as installing a simple program on windows, so only a few clicks are required. This allows everyone to quickly test the software without the hurdles of complicated installations.
- A simple task submission interface (graphical) is needed to make also the gridification of applications as easy as possible.
- Only a modifiable part of the computing power is used for the Grid, the user can decide on this. This is already enabled by the virtualisation infrastructure that we have deployed on server and desktop machines. Users can easily adjust the virtual machine through a graphical interface to donate only the specified amount of resources to Grid tasks.
- The load-balancer has to deal with the very heterogeneous architecture and has to estimate computation times.

3.4. Image retrieval on a distributed hospital infrastructure and on external Grids

Our current image retrieval system that is to be gridified is based on the GNU Image Finding tool (GIFT¹⁷, [14]). On the currently used desktop computer, the extraction of visual features from a single image takes roughly 1–2 seconds. Databases of image retrieval benchmarks such as ImageCLEFmed¹⁸ [15] contain in the order of 50'000 images, resulting in 14–20 hours calculation time. If various parameters for the features have to be tested the indexing is repeated several times. The retrieval part on the other hand is harder to parallelise but this part is currently still quick. With database sizes of around 50'000 images, response times on a desktop are in the order of 2 seconds. A particular problem is that the current infrastructure limits what we can do with the data to be able to extract features quickly. More complex features can easily take one minute per image instead of 1 second resulting in 94 hours for computing.

3.4.1. Adapting an image retrieval system for KnowARC

The feature extraction component itself is a fairly small program written in C. It does not have any special programming library dependencies. Scripts for database indexing currently take the name of a directory and then search all subdirectories for image files. For these files visual features are extracted sequentially.

A first gridification of the feature extraction takes N images at a time and sends the image files with the feature extraction source code to an ARC server. Then, the source code is compiled for each package and the features are extracted for all images in the package. After the execution the results are downloaded to a local machine for index generation. Overhead times of the job submission are fairly high and it can take 30 seconds before a block is executed. Still, as there can be many parallel packages, the overall computing and communication time is much faster than sequentially computing the visual features of all images of a large collection even without any optimisation.

The initial tests conducted on NorduGrid test sites within the virtual organisation KnowARC in Europe have given the image features 5-10 minutes after the images were sent to the Grid for processing resulting in total time for assembling the info in less than one hour that beforehand took 14 hours on a single machine. These first measurements

¹⁷<http://www.gnu.org/software/gift/>

¹⁸<http://ir.ohsu.edu/image/>

clearly illustrate that the overhead of job scheduling and data transfers are in a tolerable range. Moreover, expert opinion has proposed us to use batch sizes that result in roughly one hour of computation when submitted to European wide Grid deployment. This is in line with our initial measurements that small jobs create a fairly large overhead. While the optimal batch size will depend on the use of resources on the time of job execution, an important step has been taken in reducing the time requirements from 14 hours to just a fraction. Thus, more computing power enables us to design systems that provide more rapid response or that can use more complex feature sets, which is the final goal of this project.

3.4.2. Virtualising desktop computers

First steps have also started on virtualising standard hospital desktop computers. The windows computers now run Linux inside a virtual machine and software can be installed securely inside the sandboxed Linux operating system. Currently, we are only allowed to do this on computers of research personnel, which is sufficient for first tests (and the machines have the exact same characteristics). Six machines are currently equipped with a small distribution of Ubuntu that should make the installation of the ARC middleware relatively easy. One problem is currently that the virtual machines do not have their own IP address, which is a requirement for the middleware that we hope can be removed in the near future. Otherwise, we will be obliged to obtain additional IP addresses as well for our tests.

4. Conclusions

To apply a Grid computing infrastructure on networks in hospitals that are equally used for routine clinical practice might still be a few years away, but technology is offering solutions that could change this quickly. The possibility to virtualise even simple desktop machines allows to have a fix separation of the research and the routine parts of the computers removing at least a part of the security concerns. It also allows users of computers to close a virtual machine if this would slow down the work. Of course, much work is still ahead to make efficiency tests for virtualisation solutions to measure the amount of computing power that can really be made available.

One important step for us right now is to imagine computationally intensive solutions that are able to use a possibly unlimited computing power available to improve the current approaches in medical data analysis and medical information retrieval. Researchers in the medical imaging domains do not have the habit of doing so, and this new thinking might open up completely new possibilities.

There will most likely be two surviving business models for hospital Grids. One is the use of an existing infrastructure that allows research projects to profit from computing power partly during the day but mainly during off peak hours by using idle circles of already available CPUs. This model will mainly work for research and non-critical applications. For time-intensive critical application an external computing service provider can be used to perform extensive tasks for example during an operation or while a patient is in intensive care to optimise the outcome. A model for paying these computational requirements has to be developed.

Sharing in the scientific world needs to be extended to more domains than only knowledge with publications. Computing resources and power can equally be exchanged through the use of Grid networks and virtual organisations. Sharing of other resources such as databases for research is equally seen as an extremely important step to make research more efficient and effective [16].

Acknowledgements

This work was partially supported by the Swiss National Science Foundation (Grant 205321-109304) and the EU 6th Framework Program in the context of the KnowARC project (IST 032691).

References

- [1] Litzkov, M., Livny, M., Mutka, M.: Condor — a hunter of idle workstations. In: Proceedings of the 8th international conference on distributed computing, San Jose, California, USA (1988) 104–111
- [2] Romberg, M.: The UNICORE architecture: Seamless access to distributed resources. In: Proceedings of the 8th international symposium on high-performance distributed computing, Redondo Beach, California, USA (1999) 287–293
- [3] Foster, I., Kesselman, C.: Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications* **11** (1997) 287–293
- [4] Ellert, M., Gronager, M., Konstantinov, A., Konya, B., Lindemann, J., Livenson, I., Nielsen, J.L., Niinimäki, M., Smirnova, O., Waananen, A.: Advanced resource connector middleware for lightweight computational grids. *Future Generation computer systems* (2006 – to appear)
- [5] Peitsch, M.C.: Grid computing in drug discovery. In: Sixth IEEE International Symposium on Cluster Computing and the grid, Singapore (2006) 3–13
- [6] Müller, H., Garcia, A., Vallée, J.P., Geissbuhler, A.: Grid computing at the university hospitals of Geneva. In: Proceedings of the 1st healthgrid conference, Lyon, France (2003) 264–276
- [7] Müller, H., Michoux, N., Bandon, D., Geissbuhler, A.: A review of content-based image retrieval systems in medicine – clinical benefits and future directions. *International Journal of Medical Informatics* **73** (2004) 1–23
- [8] Montagnat, J., Breton, V., Magnin, I.E.: Partitioning medical image databases for content-based queries on a grid. *International Journal of Supercomputer Applications* **44** (2005) 154–160
- [9] Costa Oliveira, M., Cirne, W., de Azevedo Marques, P.M.: Towards applying content-based image retrieval in clinical routine. *Future Generation Computer Systems* **23** (2007) 466–474
- [10] Estrella, F., McClatchey, R., Odeh, M., Rogulin, D., Solomides, T.: Experiences of engineering grid-based medical software. *International Journal of Medical Informatics* (2007 – to appear)
- [11] Sloot, P.M.A., Tirado-Ramos, A., Altintas, I., Bibak, M., Boucher, C.: From molecules to man: Decision support in individualized e-health. *IEEE Computer* **November** (2006) 40–46
- [12] Ruch, P., Baud, R., Geissbuhler, A.: Evaluating and reducing the effect of data corruption when applying bad of words approaches to medical records. *International Journal of Medical Informatics* **67** (2002) 75–83
- [13] Müller, H., Lovis, C., Geissbuhler, A.: The medGIFT project on medical image retrieval. In: *Medical Imaging and Telemedicine*, Wujishan, China (2005)
- [14] Squire, D.M., Müller, W., Müller, H., Pun, T.: Content-based query of image databases: inspirations from text retrieval. *Pattern Recognition Letters (Selected Papers from The 11th Scandinavian Conference on Image Analysis SCIA '99)* **21** (2000) 1193–1198 B.K. Ersboll, P. Johansen, Eds.
- [15] Müller, H., Deselaers, T., Lehmann, T.M., Clough, P., Eugene, K., Hersh, W.: Overview of the imageclefmed 2006 medical retrieval and medical annotation tasks. In: *CLEF 2006 Proceedings. Lecture Notes in Computer Science* (2007 – to appear)
- [16] Vannier, M.W., Summers, R.M.: Sharing images. *Radiology* **228** (2003) 23–25