Grid Technology

Research Assistant, Dr Oxana Smirnova, sings the praises of the distributed supercomputer...

It was only a few years ago when the tendency of deploying bigger and bigger supercomputers for scientific research and industry started to give way to a new paradigm, the Grid. The idea of joining the existing variety of computing resources by using a layer of dedicated tools, the Grid middleware, became particularly appealing for large multinational collaborations, eager to merge their distributed computing power into a single, well managed facility. Applications from such areas as aeronautics, space sciences, interactive visualisation and teleimmersion were the first to try the emerging Globus Toolkit™ (www.globus.org), which became a defacto standard for the Grid middleware. In 2001, the huge community of high energy physics, famous for inventing the world wide web, made a bold step, committing resources to the EU-sponsored DataGrid project [www.eu-datagrid.org], aiming at developing a solution satisfying their needs of high-throughput, data-intensive distributed computing.

For us, the high energy physicists in the Nordic countries, the Grid approach to computing appeared to be the most efficient way of joining our resources in order to support the efforts of achieving the common goal: construct, deploy and utilise the future ATLAS detector, which is being built at the European Particle Physics Laboratory CERN (www.cern.ch). Therefore, simultaneously with the EU DataGrid, the Nordic Testbed for Wide Area Computing and Data Handling (NorduGrid) project (www.nordugrid.org) was launched, aiming at building a Grid infrastructure based on the excellent NorduNet networks (www.nordu.net).

The NorduGrid participants include universities and research centres in Denmark, Sweden, Finland and Norway. The active phase of the project started in May 2001, and involves the Niels Bohr Institute (Denmark), Lund and Uppsala Universities (Sweden), Universities of Oslo and Bergen (Norway) and Helsinki Institute of Physics (Finland). National supercomputer centres in Norway (Parallab) and Sweden (NSC, HPC2N) entered the project by providing the computing resources for our research.

Our original intention was to rely on the software tools provided by Grid R&D projects worldwide, as a handful of such started to emerge. The driving set of requirements was

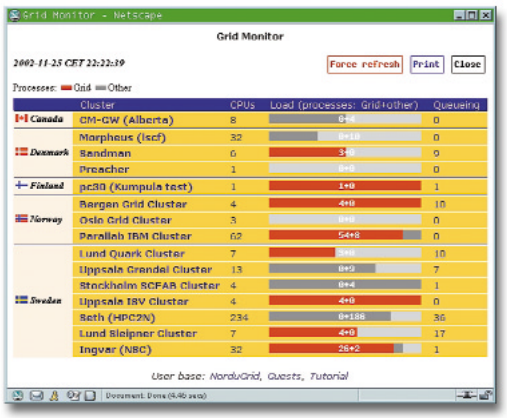

to provide the researchers with a reliable, distributed facility, to be used in their daily work. However, during the evaluation of existing and developing tools, it became clear that none satisfies these requirements. Meanwhile, the needs of the ATLAS collaboration to start computing simulation of the physics processes and the detector performance were pressing, and by no means did Nordic physicists want to be sidelined. Therefore, in February 2002, it was decided to change the priorities of the NorduGrid project from building a testbed to developing a set of own middleware tools, allowing deployment of production level Grid services in the Nordic countries.

Apart from the demanding task of quickly developing a production-scale Grid solution, the biggest challenges faced by us were the variety of available computing resources and a very distributed nature of both computing and human power. The computers range from single notebook portables to huge production clusters of hundreds of processors and supercomputers. Thankfully, most of the research facilities in the Nordic countries use Linux as the operating system of choice, but even this comes in many flavours. The project members are separated by many hundreds miles, which made internet an essential service and an integral part of the developers team. The acquired experience and developed tools have the ability to plug into the system of any computer worldwide, be it in Canada or Japan.

The architecture of the NorduGrid middleware is quite simple: each computer resource is being equipped with a set of tools, of which some are publishing the necessary status information

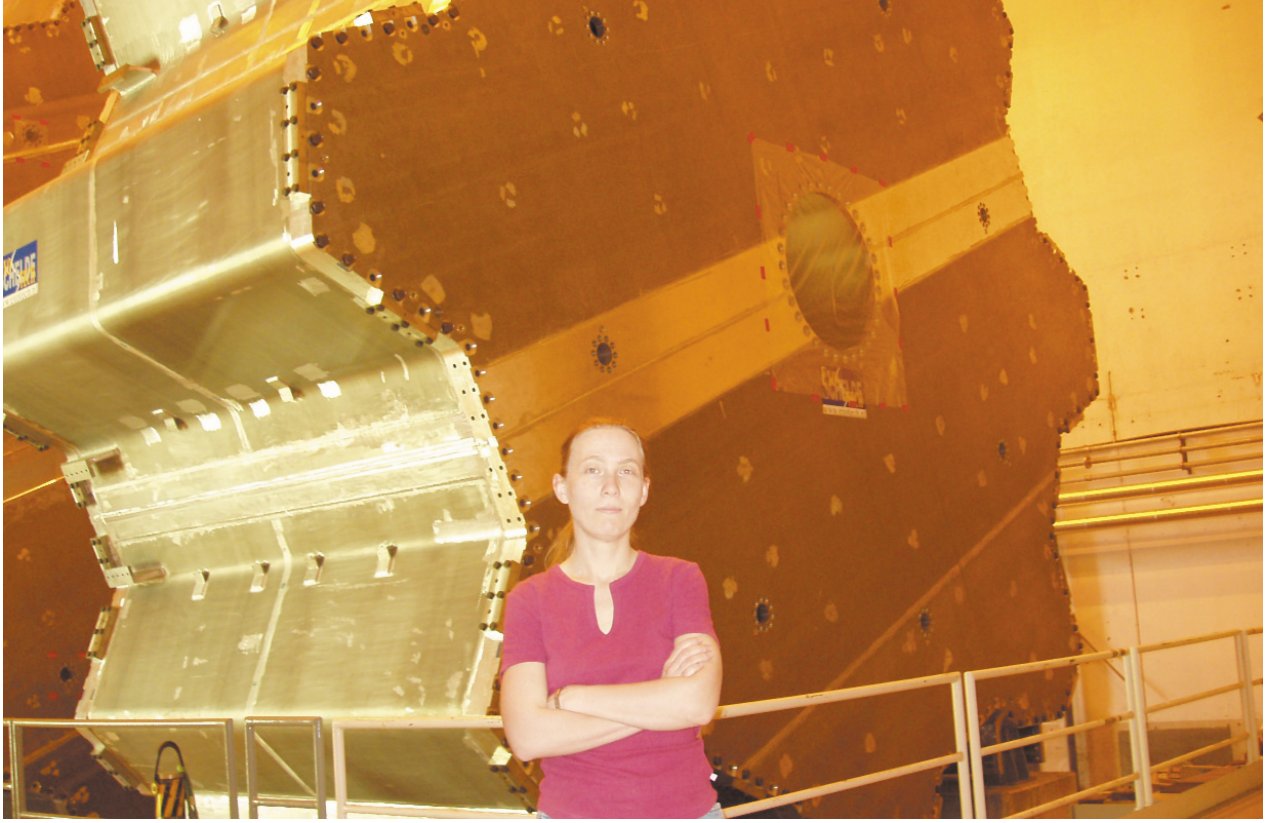

into a common distributed database (Information System), and others are in charge of receiving job descriptions from remote users and submitting them locally (Grid Manager). This is the so-called server part, while the client (User Interface) can be installed by any user at any desktop Linux PC. You may wonder, how does the client know where to submit a task? Quite easily; it polls the Information System and makes a matching of user requirements to the existing resources. If more than one resource matches, the User Interface makes a choice based on its estimation of a best resource. All the data transfer is made either by the User Interface, or by the Grid Manager, depending on the structure of the task. The middleware relies almost entirely on the Globus Toolkit™; however, most high-level services had to be re-written or created anew. An essential feature of all products developed by the NorduGrid project tools is their non-invasiveness, as our goal was to create a middleware layer, easily deployable on top of any system. Such systems can join or leave the Grid at any time, thanks to the deployed soft registration mechanism, which keeps the whole testbed dynamic and evolving.

The first live job submission on the NorduGrid – powered testbed was made during the annual meeting of the Swedish Particle Physics community on 14th April 2002. In May 2002, first tests of the ATLAS detector simulation program were made, using five test clusters in Denmark, Norway and Sweden. These tests were so encouraging that NorduGrid signed up to be the Nordic part of the so-called ATLAS Data-Challenge 1. ATLAS has planned a series of computing challenges of increasing size and complexity to validate its computing model and the software suite, and to incorporate the use of Grid middleware as soon as possible. The first one, ATLAS Data-Challenge 1, took place from July to September 2002 and consisted of large-scale physics simulation based on requests from different physics communities. 39 different institutes from around the world participated and NorduGrid was among them. The simulation performed by NorduGrid in this challenge consisted of 1,300 jobs using 176 Gb of input data and producing in excess of 800 Gb of output. All output files were, at the end of the simulations, automatically stored on a storage-server in Norway and registered into a database. All in all, the challenge was successfully met by NorduGrid.

Already by September 2002, not only the ATLAS researchers were using the growing number of the Nordic Grid resources for large-scale production, but also other Nordic researchers, most notably, in theoretical physics, were able to use the Grid in order to perform calculations they never would be able to accomplish in such a short time before.

For the time being, the NorduGrid celebrates its second anniversary by having put together almost 1,000 processors in the Nordic countries, as well as some in Switzerland, Canada and Japan. The basic Grid infrastructure and services are set up and being constantly upgraded to meet the daily needs of the users. The NorduGrid provides a unique, geographically distributed, international testbed, operating reliably 24/7 without any necessity to be manned around the clock. The middleware is distributed freely and, together with more information, is available at our website.

SWEDEN

Dr Oxana Smirnova
Research Assistant
oxana.smirnova@hep.lu.se
www.nordugrid.org