EXPLOITING THE GRID TO SIMULATE AND DESIGN THE LHCb EXPERIMENT

K Harrison[1], N Brook[2], G Patrick[3], E van Herwijnen[4], on behalf of the LHCb Grid group[a] and GridPP.
[1]University of Cambridge, [2]University of Bristol, [3]Rutherford Appleton Lab, [4]CERN.

Key words to describe the work: Grid deployment; Distributed computing; High-energy physics applications; Simulation; Bookkeeping; Monitoring of data quality

Key Objectives: Use large-scale distributed computing facilities for simulating and reconstructing high-energy particle collisions; Optimise the design of a detector for a high-energy physics experiment (LHCb)

Motivation for the work (problems addressed):

1) Design studies for the LHCb experiment require many millions of high-energy particle collisions to be simulated and reconstructed, in a distributed environment.
2) The large quantities of data being processed (in the Petabyte range) must be made available to a globally dispersed collaboration of several hundred physicists.

LHCb is a particle-physics experiment that will study subtle differences between matter and antimatter. The design and construction of the experiment is being undertaken by some 500 scientists, from 40 institutes in 14 countries around the world. The experiment will be located 100 m underground at the Large Hadron Collider (27 km circumference) being built at CERN in Geneva. The decays of more than 10^9 short-lived particles and antiparticles, known as B mesons, will be studied at LHCb each year. To optimise the detector design, and to understand the physics, many millions of particle interactions must be simulated. The first data are anticipated in early 2007.

Grid technology is being deployed so that we can use globally distributed computing resources to satisfy our requirements. A prototype system, based on our existing software, is already operational. This deals with job submission and execution, data transfer, bookkeeping, and the monitoring of data quality. DataGrid middleware is being utilised as it becomes available. In this way, LHCb is able both to produce the simulated datasets for the detector studies, and to feed back experience and ideas into the design of the Grid. This work requires fundamental research and development in the issues of:

- job optimisation - matching application demands to site attributes
- management of data replicas to facilitate optimisation
- security of systems and data in a large international environment involving thousands of users
- monitoring and error recovery for large, complex jobs involving parallelism
- storage and management of enormous datasets, running to several Petabytes
- flexible user-interface, giving users (physicists, system managers, virtual-organisation managers, etc) easy access to Grid services in a transparent manner.

Computing centres on the Grid are being integrated into the LHCb production system as they come on line. Currently these include: CERN, IN2P3 (Lyon) in France, CNAF (Bologna) in Italy, NIKHEF (Amsterdam) in the Netherlands, and the EU DataGrid Testbed. In the UK, resources are being made available for LHCb computing at the Universities of Bristol, Cambridge, Edinburgh, Glasgow, Liverpool, London (Imperial College) and Oxford, and at the Rutherford Appleton Laboratory.

The poster will describe how LHCb are using large-scale distributed computing facilities to simulate many millions of high-energy particle physics collisions in order to design their detector. The issues that will be addressed include optimising the submission of many thousands of jobs to sites distributed throughout Europe, and coping with data transfer rates that can reach Terabytes per day. The associated bookkeeping to permit efficient storage and retrieval of data, and the essential monitoring, in

a distributed environment, of job execution and data quality will also be discussed.

The Monte-Carlo production environment is heterogeneous with respect to the underlying batch system, but is currently used mainly with the Linux operating system. It is, however, also possible to run jobs under the Windows operating system, using a batch system based on CONDOR.

Job definition and submission is through a web page. A central web server, installed on a Windows platform, can submit to individual farms using AFS, or to the EU DataGrid Testbed using Grid middleware. A java servlet generates the required job scripts, as well as specifying the random-number seeds and job options. The job options, such as the version of the simulation or the type of events to simulate, can be controlled via the web page.

The submission of jobs to distributed sites and the monitoring of the status and progress of jobs in the distributed environment are performed using the PVSS SCADA system. This allows the configuration of each individual job to be checked, and allows jobs to be terminated remotely. The same system has been adopted by the LHC experiments for detector monitoring and control during data taking, and is available for both Linux and Windows platforms.

Data transfer is performed using the bbftp system, which allows parallel-stream transfers. The datasets produced during each simulation job are of the order of several Gigabytes in size. The data transfer is directly to the hierarchical mass-storage system at CERN, CASTOR. During the current data challenges, data will be stored locally at the larger computer centres, which will also accept data transferred from other external sites.

Bookkeeping is performed using java classes that interact with a central ORACLE database at CERN via a servlet. The central database will also hold information for simulated datasets stored externally to CERN.

[a] UK contacts are: N Brook (University of Bristol), K Harrison (University of Cambridge), A Khan (University of Edinburgh & University of Glasgow), G Patel (University of Liverpool), U Egede (Imperial College, London), I McArthur (University of Oxford), G Patrick (Rutherford Appleton Laboratory)