University of Glasgow | Experimental Particle Physics

GLAS-PPE/2009-03
21st May 2009

Department of Physics and Astronomy
Experimental Particle Physics Group
Kelvin Building, University of Glasgow
Glasgow, G12 8QQ, Scotland
Telephone: +44 (0)141 330 2000 Fax: +44 (0)141 330 5881

Deploying a Resilient Grid for UK Particle Physics Grid

D.I. Britton[1], A.J. Cass[2], P.E.L. Clarke[3], J. Coles[4], D.J. Colling[5], A.T. Doyle[1], N.I. Geddes[6], J.C. Gordon[6], R.W.L. Jones[7], D.P. Kelsey[6], S.L. Lloyd[8], R.P. Middleton[6], G.N. Patrick[6], R.A. Sansum[6] and S. Pearce[8]
[1]Dept. Physics & Astronomy, University of Glasgow, Glasgow G12 8QQ, UK.
[2]CERN, Geneva 23, Switzerland.
[3]National e-Science Centre, Edinburgh - EH8 9AA, UK.
[4]Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, UK.
[5]Dept. Physics, Imperial College, London SW7 2AZ.
[6]STFC Rutherford Appleton Laboratory, Didcot OX11 0QX, UK.
[7]Dept. Physics, Lancaster University, Lancaster LA1 4YB, UK.
[8]Dept. Physics, Queen Mary, University of London, London E1 4NS, UK.
Email: d.britton@physics.gla.ac.uk

Abstract

The start-up of the Large Hadron Collider at CERN in Geneva presents a huge challenge in processing and analysing the vast amounts of scientific data that will be produced. The architecture of the worldwide Grid that will handle the 15PB of particle physics data annually is based on a hierarchical tiered structure. We describe the development of the UK component (GridPP) of this Grid, from a prototype system to a full exploitation Grid for real data analysis. This includes the physical infrastructure, the deployment of middleware, operational experience and the initial exploitation by the major LHC experiments.

Deploying a Resilient Grid for UK Particle Physics

D.I. Britton[1], A.J. Cass[2], P.E.L. Clarke[3], J. Coles[4], D.J. Colling[5], A.T. Doyle[1], N.I. Geddes[6], J.C. Gordon[6], R.W.L. Jones[7], D.P. Kelsey[6], S.L. Lloyd[8], R.P. Middleton[6], G.N. Patrick[6], R.A. Sansum[6] and S. Pearce[8]
[1]Dept. Physics & Astronomy, University of Glasgow, Glasgow G12 8QQ, UK.
[2]CERN, Geneva 23, Switzerland.
[3]National e-Science Centre, Edinburgh - EH8 9AA, UK.
[4]Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, UK.
[5]Dept. Physics, Imperial College, London SW7 2AZ.
[6]STFC Rutherford Appleton Laboratory, Didcot OX11 0QX, UK.
[7]Dept. Physics, Lancaster University, Lancaster LA1 4YB, UK.
[8]Dept. Physics, Queen Mary, University of London, London E1 4NS, UK.

Phone: +44 141 330 5454 Fax: +44 141 330 5881 E-mail: d.britton@physics.gla.ac.uk

Abstract - The start-up of the Large Hadron Collider at CERN in Geneva presents a huge challenge in processing and analysing the vast amounts of scientific data that will be produced. The architecture of the worldwide Grid that will handle the 15PB of particle physics data annually is based on a hierarchical tiered structure. We describe the development of the UK component (GridPP) of this Grid, from a prototype system to a full exploitation Grid for real data analysis. This includes the physical infrastructure, the deployment of middleware, operational experience and the initial exploitation by the major LHC experiments.

I. THE LHC COMPUTING CHALLENGE

The Large Hadron Collider (LHC) [1] is set to become the world's highest energy particle accelerator. Protons, with energies of up to 7 TeV, will be collided at 40 million times a second to recreate the conditions of the earliest moments of the "Big Bang". Positioned around the 27 km superconducting collider are four major experiments - ALICE [2], ATLAS [3], CMS [4] and LHCb [5] - which will record the particle interactions of interest. These experiments contain a total of ~150 million electronic sensors and produce about 15 PB of data a year at a rate of about 700 MB/s. The processing and analysis of these data will require 80PB of storage and an initial CPU capacity of 100,000 processors operating continuously, which will need to double by 2010 [6].

Particle physicists have chosen Grid technology to meet this challenge, with the computing and storage distributed worldwide. This paper provides an overview of this worldwide LHC Computing Grid (wLCG) with an emphasis on the component developed and deployed in the UK by the GridPP collaboration.

II. THE WORLDWIDE LHC COMPUTING GRID

The data from the LHC must be shared by 5000 scientists located at ~500 institutes throughout the world. Copies of the data must be kept for the lifetime of the LHC and beyond - a minimum of 20 years. At present, the wLCG is comprised of ~250 computing centres based on a hierarchical tiered structure [7]. Data flows from the experiments to a single Tier 0 Centre at CERN, where a primary backup of the raw data is kept. After some initial processing, the data is distributed over an optical private network, LHCOPN [8], with 10 Gb/s links to eleven major Tier 1 centres around the world. Each Tier 1 is responsible for the full reconstruction, filtering and storage of the event data. These are large computer centres with 24x7 support. In each region, a series of Tier 2 centres then provide additional processing power for data analysis and Monte-Carlo simulations. Individual scientists will usually access facilities through Tier 3 computing resources consisting of local clusters or even their own desktops/laptops.

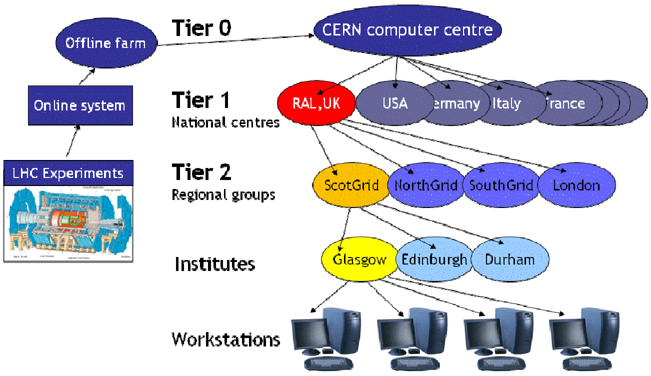

Fig.1: Hierarchical Tier structure of the wLCG.

The hierarchical nature of the wLCG is an optimisation for particle physics that reflects the data-processing requirements and is not necessarily the best design for other applications. WLCG is based on two major Grid infrastructures: The Enabling Grids for E-Science} project or EGEE [9] and the US OpenScience Grid [10]. The Scandinavian countries also contribute the NorduGrid infrastructure.

III. THE GRIDPP ARCHITECTURE

The UK component of the Grid infrastructure for particle physics has been built by the GridPP collaboration through the joint efforts of 19 universities, the Rutherford Appleton Laboratory (RAL) and CERN. The initial concept was first taken through a prototype stage in 2001-2004 [11] and this was then followed by building a production scale Grid during 2004-2008. With the start of LHC commissioning in 2008, the project has just entered its exploitation phase.

The UK Tier 1 Centre is located at Rutherford Appleton Laboratory (RAL in Fig.1). The 2009 hardware, currently being procured, will roughly double the existing cluster of over 5000 job slots and 340 disk servers providing 2.3 PB of disk storage. A Sun SL8500 tape robot provides 10,000 media slots, 18 T10K tape drives and a storage capacity of 5 PB. A hierarchical storage management system is used, based on the CASTOR 2 system [12]. Each of the LHC experiments uses separate CASTOR instances to avoid resource conflicts and to minimise problems with data flow.

Upon receiving a share of the raw data from CERN, the Tier 1 reconstructs the particle interactions (so-called "events") from the electronic information recorded in the various detector components. The extracted physics information is then used to select, filter and store events for initial analysis. Datasets are then made available to the Tier 2 centres for specific analysis tasks.

GridPP has developed four Tier 2 Centres which are federations of computing facilities, as shown in Fig.2. These groupings of institutes (London, SouthGrid, NorthGrid and ScotGrid) allow coordination at a level that is sensitive to regional issues and priorities, and enables more efficient use of technical experts. These centres focus on providing computing power for generating simulated Monte-Carlo data and on the analysis of data by individual physicists.

The distributed Tier 2 resources currently provide ~10,000 job slots and approaching 2 PB of disk storage. The 17 individual sites vary in size from large centres such as Glasgow providing 1,900 job slots (2,913 KSI2K) to small departmental clusters of only a few machines.

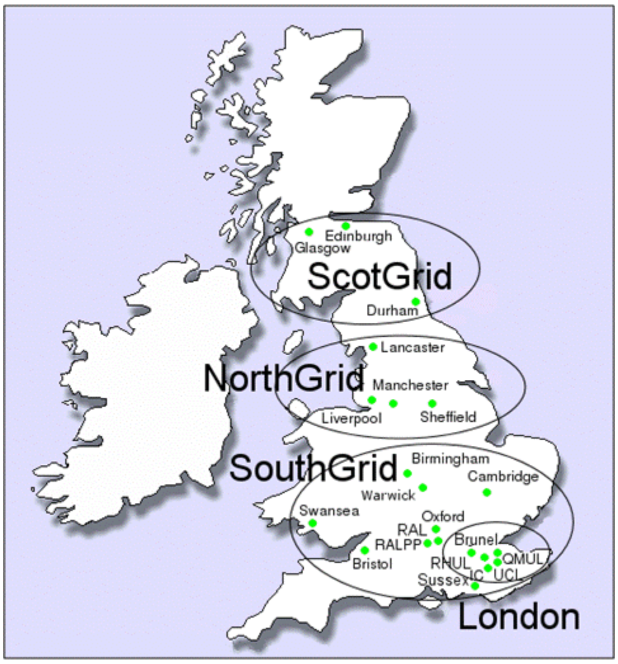

Fig.2: The UK Tier-2 Centres.

IV. GRID MIDDLEWARE

GridPP uses the the gLite middleware, currently version 3.1, from the EGEE project which has been adopted by wLCG across all of its sites. A few of the key components are described below.

1) Workload Management – WMS: A Grid job is specified using a Job Definition Language, based on the Condor ClassAd language, and this is submitted in a script together with the necessary program and input files through the Workload Management System [13]. A Resource Broker component accepts each job and matches it to a suitable site for execution. Other components transmit the job to the relevant site and manage any output. Sites accept jobs through a gatekeeper machine known as a Computing Element, which then schedules the job to run on a worker node within the cluster. Small output files are transmitted back through the WMS, while larger data files may be written to a Storage Element and catalogued.

2) Data Management: GridPP supports the small scale Disk Pool Manager (DPM) and medium scale dCache storage solutions at the Tier 2 sites, in addition to the large-scale CASTOR system at the Tier 1 centre. The problem of file access across a heterogeneous Grid with multiple mass storage systems is addressed by the Storage Resource Management (SRM) interface, a protocol that allows clients to retrieve and store files, control their lifetimes as well as reserving file-space for uploads, etc. The concepts of storage classes and space tokens were introduced recently [14] to enable experiments to place data on different combinations of storage device and to dynamically manage storage space.

Data distribution between the Tier 0, Tier 1 and Tier 2 centres is performed using the File Transport Service (FTS). This uses unidirectional channels to provide point-to-point queues between sites with file transfers handled as batch jobs. FTS provides prioritisation and retry mechanisms in the case of failure.

3) Distributed Databases and Catalogues: In addition to the interaction data (physics "events"), experiments also require a large amount of non-event data, describing the running parameters and calibration constants associated with each sub-detector. This time-varying data is essential for the reconstruction of physics events and is stored in a conditions database. Users and production programs also need the ability to locate data files (or their replicas) and this is achieved through the LCG File Catalogue (LFC). The LFC contains the mappings of logical file names to physical files, along with any replicas on the Grid.

These worldwide distributed databases have been set up by the LCG 3D (Distributed Deployment of Databases for LCG) project using Oracle Streams technology to replicate the databases to the external Tier 1 centres outside CERN [15]. At the UK Tier 1 centre, there are independent multi-node database clusters for the conditions databases of both the ATLAS experiment [16] and the LHCb experiment [17]. This ensures high availability and allows for large transaction volumes. The LFC is also based on an Oracle Enterprise relational database, but the ATLAS and LHCb catalogues reside on multi-node clusters shared with other services. These are highly configurable allowing for all services to run even if a node fails.

IV. USER ACCESS AND EXPERIMENT SOFTWARE

Particle physicists access the Grid from a User Interface (UI) which is a local computer with the user-level client tools installed. Authentication is based on digital X.509 certificates, which in the case of the UK are issued by the National Grid Service. Individual physicists belong to a Virtual Organisation (VO) representing their individual experiment and each computing site in the UK decides which VOs can use its facilities and the appropriate level of resource. A Virtual Organisation Membership Service (VOMS) provides authorisation information; specifically the roles and capabilities of a particular member of a VO. At the beginning of each session, a proxy certificate is obtained for a limited lifetime (typically 12 hours).

Experiments have developed front-ends to simplify the definition, submission and management of jobs on the Grid. The Ganga interface [18], a GridPP supported project in wide-spread use by the ATLAS and LHCb experiments, allows jobs to be run either on a local batch system or on the Grid and provides all of the facilities for job management including submission, splitting, merging and output retrieval. Ganga can be used via three methods: An interactive interface; in a script; or through a graphical interface. The adoption of Ganga enables a physicist to exploit the Grid with little technical knowledge of the underlying infrastructure. Over 1000 unique users of Ganga have been recorded in 2008 as data analysis programs have been prepared for the start of the LHC.

Similarly, experiments have written their own data management layers which sit on top of the standard Grid services. For example, the GridPP supported PhEDex (Physics Experiment Data Export) system of the CMS collaboration provides a data placement and file transfer system for the experiment [19]. This is based on transfer agents, management agents and a control database, providing a robust system to manage global data transfers. The agents communicate asynchronously and between centres the system can achieve disk-to-disk rates in excess of 500 Mbps and sustain tape-to-tape transfers over many weeks.

V. DEPLOYMENT, OPERATION AND SERVICES

In order to achieve a high quality of service across the wider Grid, the wLCG/EGEE infrastructure is divided into ten regions, each with a Regional Operations Centre (ROC) responsible for monitoring and solving operational problems at sites within its domain. GridPP sites are covered by the UK/Ireland (UKI in Fig.3) ROC.

Clusters are usually monitored by open source programs such as Nagios and Ganglia. These can display variations in cluster load, storage usage and network performance, raising alarms when thresholds are reached or components fail. Monitoring for the LCG Grid is via the Service Availability Monitoring (SAM) system. This regularly submits Grid jobs to sites and connects with a number of sensors which probe sites and publishes the results to an Oracle database. At the regional level, problems are tracked by Global Grid User Support tickets issued to the sites affected or to the middleware/operations group.

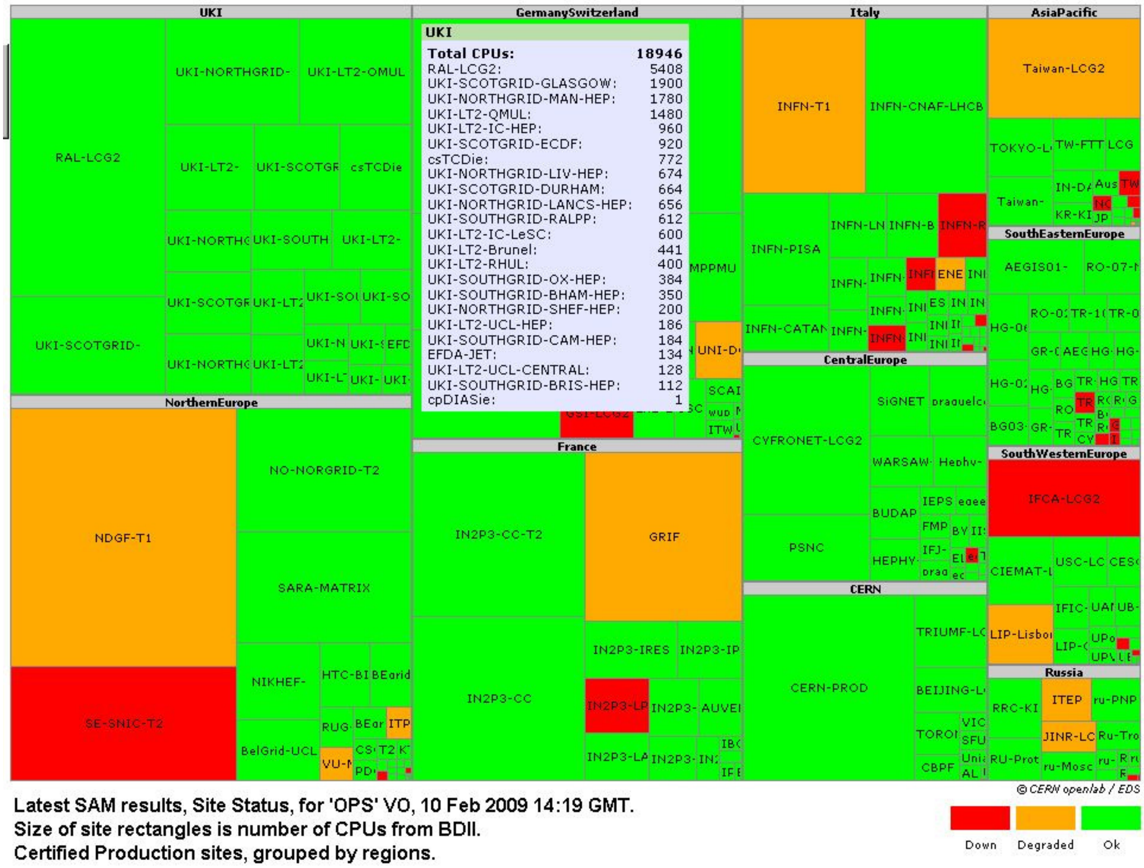

UKI	
Total CPUs:	**18946**
RAL-LCG2:	5408
UKI-SCOTGRID-GLASGOW:	1900
UKI-NORTHGRID-MAN-HEP:	1780
UKI-LT2-QMUL:	1480
UKI-LT2-IC-HEP:	960
UKI-SCOTGRID-ECDF:	920
csTCDie:	772
UKI-NORTHGRID-LIV-HEP:	674
UKI-SCOTGRID-DURHAM:	664
UKI-NORTHGRID-LANCS-HEP:	656
UKI-SOUTHGRID-RALPP:	612
UKI-LT2-IC-LeSC:	600
UKI-LT2-Brunel:	441
UKI-LT2-RHUL:	400
UKI-SOUTHGRID-OX-HEP:	384
UKI-SOUTHGRID-BHAM-HEP:	350
UKI-NORTHGRID-SHEF-HEP:	200
UKI-LT2-UCL-HEP:	186
UKI-SOUTHGRID-CAM-HEP:	184
EFDA-JET:	134
UKI-LT2-UCL-CENTRAL:	128
UKI-SOUTHGRID-BRIS-HEP:	112
cpDIASie:	1

Latest SAM results, Site Status, for 'OPS' VO, 10 Feb 2009 14:19 GMT.
Size of site rectangles is number of CPUs from BDII.
Certified Production sites, grouped by regions.

Fig.3: The wLCG SAM tests.

The LHC experiments also need to have a clear picture of their Grid activities (e.g. job successes/failures, data transfers, installed software releases) and so have developed their own SAM tests which probe the site infrastructures from their own specific standpoint. Dashboards have been developed to present this information in a concise and coherent way. In addition, GridPP has found it useful to also develop a suite of tests which regularly exercise the essential components in the experiment computing chains. For the ATLAS experiment, jobs are sent every few hours to each UK site on the Grid and perform an exemplar data analysis. The continuing challenge is to improve the robustness and reliability of the Grid for experiments.

VII. SERVICE RESILIENCE

GridPP is now at the start of the exploitation phase of a Grid that it has developed and deployed over the last 6 years. Of key importance in this phase is the reliability of the service. The Grid must be made as resilient as possible to failures and disasters over a wide scale, from simple disk failures up to major incidents like the prolonged loss of a whole site. One of the intrinsic characteristics of the Grid approach is the use of inherently unreliable and distributed hardware in a fault-tolerant infrastructure. The term "service resilience" is about making this fault-tolerance a reality.

The approach taken is to deconstruct the Grid into a set of identifiable services. Each service must then be made resilient using appropriate methods which may include redundancy; automated fail-over; manual fail-over; or temporary alternatives. Moreover, scenario planning must consider the types of problems that might happen, the correlations between them, and whether, for example, off-site back-up services need to be prepared in addition to on-site solutions. When considering large-scale events (such as loss of the Tier 1 centre or the loss of the OPN network) the response of each experiment must be clearly understood before sensible contingency plans can be formulated. In many cases, these experiment plans are only just crystallising and our planning has had to remain at a more abstract level.

During the April 2008 WLCG workshop it was noted that 80% of operational problems occur due to problems in the design and development of the software [20], which illustrates that the struggle to ensure resilience must be addressed in the software development, and not just the deployment, cycle. At the deployment level, increasing the resilience of the infrastructure can be done in a number of generic ways such as:

- *Increasing the hardware's capacity to handle faults.* This may be done by adding spare Power Supply Units preferably on alternative power routes; and by adding disks in a RAID (Redundant Array of Inexpensive/ Independent Disks) configuration. For many services the configuration used is RAID1 (mirrored disks).

- *Duplicating services or machines.* This approach is being used for many of the gLite services. Several instances are created and then addressed using a DNS round robin. Some of the services have been developed with failover in mind so that a service running on two machines in a load balanced manner can in fact run on just one if there is a failure. An approach like this is used for critical databases in the form of Oracle Real Application Servers (RACs).

- *Implementing automatic restarts.* In addition to hardware redundancy there are software safeguards that can be employed to improve reliability. The most obvious example is the use of automatic daemon restarts. There are various reasons why a daemon may stop running, but by checking regularly and automatically restarting missing daemons deeper problems are avoided.

- *Providing fast intervention.* Fast intervention allows problems to be caught early. It is greatly improved by close monitoring and good quality alarms coupled to an efficient call-out system. This is supported by the development of fast instancing of nodes/services and the provision of appropriate manpower, testing and documentation.

- *In depth investigation of the reason for failure.* Having an incident report and follow up procedure ensures that the recurrence of problems is reduced. Within GridPP we have now implemented an incident report template in order to ensure consistent and complete follow up.

Currently, the majority of the services run by GridPP have some level of resilience. The degree of resilience enabled by the middleware varies, but is improving in many cases. The FTS and SRM implementations in particular will benefit from improvements expected in future middleware releases and the increased use of virtualization has beneficial implications for many other services.

VIII. EXPERIMENT EXPLOITATION

Over several years, the LHC experiments have individually exploited the evolving Grid infrastructure through a series of "data" and "service" challenges. These have been based on large simulated data samples, typically containing many millions of events. Data have been generated and then processed by the Grid, exercising the data processing chains prepared for real data. In preparation for the switch-on of the LHC, all of the experiments simultaneously conducted a Common Computing Readiness Challenge at the start of 2008 (CCRC08), with phases in February and May. In proton-proton mode at nominal luminosity, the four LHC experiments are expected to produce a total data rate of ~1600 MB/s from the Tier 0 to the eleven external Tier 1 Centres. The UK share is estimated to be of the order 150 MB/s and, from Fig.4, it can be seen that this was achieved during both phases of CCRC08 in February and May 2008.

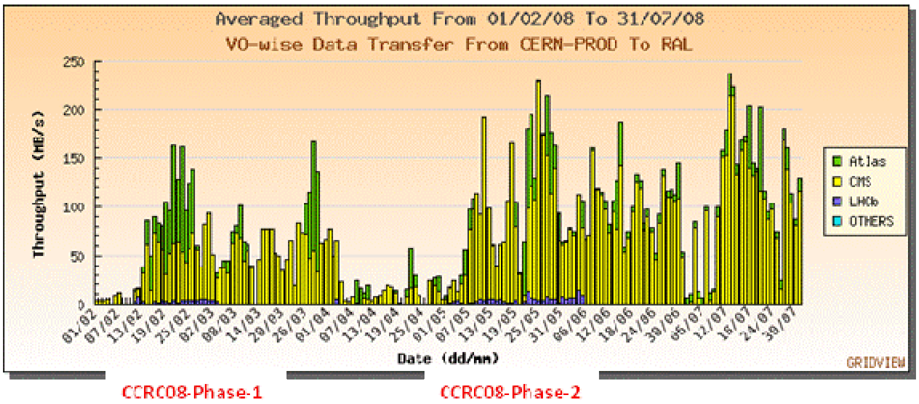

Fig.4: Averaged daily throughput from CERN to UK

Monte-Carlo techniques play a vital role in the modelling, analysis and understanding of particle physics data. By simulating a sample of events, equivalent to the size of the expected dataset, many thousands of times, the precision with which a parameter will be measured when real data arrives, can be studied. Typically, this is undertaken on the Grid by farming out many hundreds of simulation runs and then collecting the results to produce a set of statistical distributions or to show a trend. This is illustrated in Fig.5a, which shows the value of a systematic detector effect for various species of particle reconstructed in the LHCb experiment which could bias the measurement of a real asymmetry between the decays of the neutral B meson (particle) and anti B meson (antiparticle), important in understanding the prevalence of matter in the Universe. To produce the physics asymmetry in Fig.5b, a simulated "experiment" with one million signal events (comparable to six months of real data taking) was run ~500 times with different input configurations. This was repeated for several different models leading to a very large number of simultaneous jobs on the Grid which consumed over 100,000 hours of CPU time [21].

Another example [22] of the power of the Grid comes from the CMS experiment, which simulated, processed and analysed a sample of over 20 million physics events containing pairs of muons. The aim was to perform a measurement of the mass spectrum of muon pairs and isolate any resonances after processing and filtering the data. A fit to the final subsample of events after track isolation is shown in Fig.6, where the peak due to the decay of the Z-boson to a pair of muons can be seen prominently sitting on the background from quantum chromodynamics (QCD).

IX. OUTLOOK

A working Grid for particle physics has been established in the UK with the necessary resources for the early exploitation of data from the Large Hadron Collider. The LHC is poised to become the frontline facility for particle physics over the next decade and GridPP is a vital component of the worldwide infrastructure built to process and analyse the data recorded by the four main experiments.

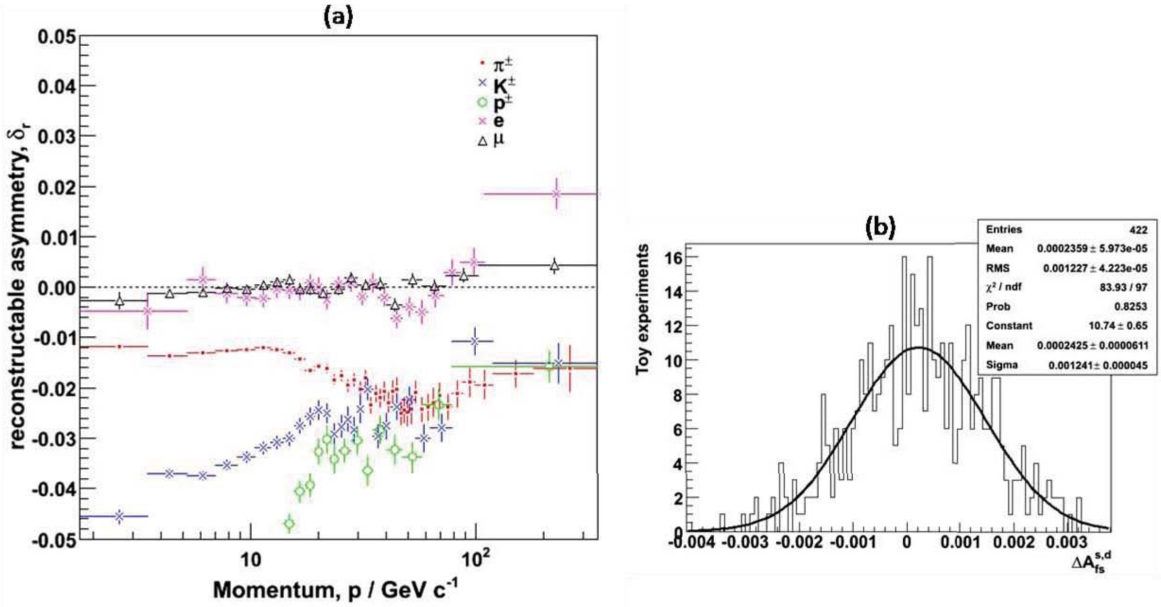

Fig.5: (a) Reconstruction asymmetry for various reconstructed particles in the LHCb detector, and (b) example contribution to the asymmetry parameter from a set of 422 Monte-Carlo simulations of one million event samples.

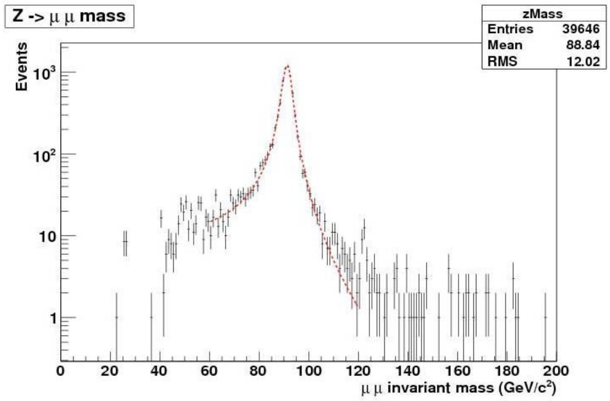

Fig.6: Simulated di-muon invariant mass spectrum from CMS.

The commissioning of the collider started with a series of injection tests in August and September 2008. On 10 September 2008 the first circulating proton beam through the entire 27 km of the LHC ring was achieved, closely followed by several hundred orbits. Particle interactions originating from the beam were observed in all four experiments. A technical fault with the accelerator developed on 19 September and proton-proton collisions are now scheduled for autumn 2009. GridPP will then be ready to record the first real LHC collision data following eight years of development.

ACKNOWLEDGEMENTS

We acknowledge financial support from the Science and Technology Facilities Council (STFC) in the UK and from the EGGE collaboration. We wish to thank the many individuals in the UK, both in the Tier 1 Centre and the Tier 2 institutes, who have helped to build GridPP. This Grid project would also have not been possible without the major contributions from our wLCG and EGEE colleagues at CERN and around the world. The CMS collaboration and Rob Lambert from the LHCb collaboration are thanked for making available the results of their simulations.

REFERENCES

[1] L. Evans and P. Bryant, (eds.), “LHC Machine”, *JINST* 3, S08001, 2008.

[2] K. Aamodt, et al. “The ALICE experiment at the CERN LHC”, JINST 3, S08002. 2008.

[3] G. Aad, et al. “The ATLAS Experiment at the CERN Large Hadron Collider”, *JINST* 3, S08003, 2008.

[4] S. Chatrchyan, et al. “The CMS Experiment at the CERN LHC”, *JINST* 3, S08004. 2008.

[5] A.Augusto Alves Jr, et al. “The LHCb Detector at the LHC”, *JINST* 3, S08005. 2008.

[6] J. Knobloch, (ed.), “LHC Computing Grid Technical Design Report”, LCG-TDR-001 and CERN-LHCC-2005-024, CERN, 2005. See http://lcg.web.cern.ch/LCG/tdr.

[7] J. Shiers, “The Worldwide LHC Computing Grid (worldwide LCG)”. *Comp. Phys. Commun.* 177, 219-233, 2007. See also http://lcg.web.cern.ch/LCG/ .

[8] D. Foster, (ed.), “LHC Tier-0 to Tier-1 High Level Network Architecture”, 2005. See https://twiki.cern.ch/twiki/pub/LHCOPN/LHCopnArchitecture/LHCnetworkingv2.dgf.doc.

[9] B. Jones, “An Overview of the EGEE Project”, In *Peer-to-Peer, Grid, and Service-Orientation in Digital Library Architectures, Lecture Notes in Computer Science*, 3664 pp 1-8, Berlin: Springer. (doi:10.1007/11549819_1), 2005. See also http://www.eu-egee.org/

[10] R. Pordes, et al., “The Open Science Grid status and architecture”, *J. Phys.: Conf. Ser.*, 119, 052028, (doi:10.1088/1742-6596/119/5/052028), 2008. See also http://www.opensciencegrid.org/

[11] P.J.W. Faulkner, et al., “GridPP: Development of the UK computing Grid for particle physics”, *J. Phys. G: Nucl. Part. Phys.*, 32, N1-N20, (doi:10.1088/0954-3899/32/1/N01), 2006. See also http://www.gridpp.ac.uk/

[12] G.L. Presti, et al., “CASTOR: A Distributed Storage Resource Facility for High Performance Data Processing at CERN”, *Proc. 24th IEEE Conf. on Mass Storage Systems and Technologies*, 275-280. (doi:10.1109/MSST.2007.4367985), 2007. See also http://castor.web.cern.ch/castor/

[13] P. Andreetto, et al., “The gLite Workload Management System”, *J. Phys.: Conf. Ser.*, 119, 062007, (doi:10.1088/1742-6596/119/6/062007), 2008.

[14] F. Donno, et al., “Storage Resource Manager Version 2.2: design, implementation, and testing experience”, *J. Phys.: Conf. Ser.*, 119, 062028, (doi:10.1088/1742-6596/119/6/062028), 2008. See also http://sdm.lbl.gov/srm-wg/doc/SRM.v2.2.html

[15] D. Duellmann, et al., “LCG 3D Project Status and Production Plans”, *Proc. of Computing in High Energy and Nuclear Physics*, Mumbai, Feb. 2006. See also http://lcg3d.cern.ch/

[16] F. Viegas, R. Hawkings and G. Dimtrov, “Relational databases for conditions data and event selection in ATLAS”, *J. Phys.: Conf. Ser.,* 119, 042032, (doi:10.1088/1742-6596/119/4/042032), 2008.

[17] M. Clemencic, “LHCb Distributed Conditions Database”, *J. Phys.: Conf. Ser.,* 119, 072010, (doi:10.1088/1742-6596/119/7/072010), 2008.

[18] A. Maier, “Ganga - a job management and optimising tool”, *J. Phys.: Conf. Ser.,* 119, 072021, (doi:10.1088/1742-6596/119/7/072021), 2008.
See also http://ganga.web.cern.ch/ganga/

[19] L. Tuura, et al., “Scaling CMS data transfer system for LHC start-up”, *J. Phys.: Conf. Ser.*, 119, 072030, (doi:10.1088/1742-6596/119/7/072030), 2008.

[20] “On Designing and Deploying Internet Scale Services”, 21st LISA conference, 2007.

[21] R.W. Lambert, “LHCb Hybrid Photon Detectors and Sensitivity to Flavour Specific Asymmetry in Neutral B-Meson Mixing”, PhD thesis, U. of Edinburgh, 2008.

[22] CMS Collaboration, “Towards a measurement of the inclusive $W \rightarrow \mu\nu$ and $Z \rightarrow \mu^+\mu^-$ cross sections in pp collisions at Sqrt(s) = 14 TeV”, Report CMS-PAS-EWK-07-002, CERN, 2008.