Meeting the particle physics computing challenge

Professor Tony Doyle describes the context and latest developments in Grid computing at the start of the GridPP2 project…

In 2007, following more than 10 years of preparatory work, the Large Hadron Collider (LHC) at CERN, Geneva, will start to collide protons with an energy equivalent of seven trillion volts to recreate the conditions that prevailed in the universe at the earliest moments of the 'Big Bang'. A billion interactions will be generated every second and, amongst them, perhaps one event will involve the production of a Higgs boson, responsible for the intrinsic mass of all the other fundamental particles. This will rapidly decay and its daughter particles will be detected in massive 20 metre high detectors, with up to 20 million readout channels. The data will be efficiently filtered using dedicated electronics but the data rates will still be enormous, with around 10 trillion bytes (10 PetaBytes) of data produced each year from each of the four experiments. Thousands of physicists from all around the world will be eager to analyse the first data. Tens of millions of lines of analysis code will be written and the required processing power will be more than 100,000 processors operating continuously over many years. This is the nature of the LHC Computing Challenge.

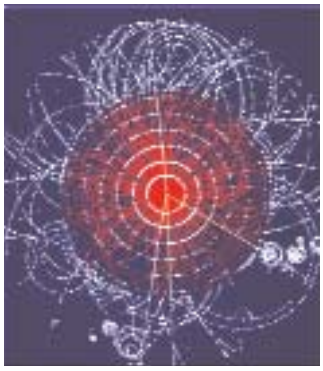

This complexity is illustrated in the figure above of a single simulated event where a Higgs particle decays into four (highlighted) particles amongst the many generated and detected in the ATLAS detector. As described in a recent article by Dr Sarah Pearce, 'An open source Grid for science', the Grid is the chosen technology, a hardware and software infrastructure that provides dependable, consistent, pervasive and inexpensive access to high end computational capabilities. The system must allow sharing of data between thousands of scientists with multiple interests; link major and minor computer centres across the globe; ensure all data is accessible anywhere, anytime; grow rapidly, yet remain reliable for more than a decade; cope with different management policies at different centres; ensure data security; and be up and running routinely by 2007. The Grid is a practical solution to meet this and many other challenges in various areas of science, as part of the UK's e-Science programme. In this article, we concentrate on the scale of the infrastructure, the tests that have been performed in recent months and the plans for the second phase of the GridPP project, going from prototype to production.

In September 2001, the UK defined its aims in this area and the three year GridPP project commenced. The aim was to develop a highly functional prototype Grid deployed across the UK, consisting of 2,000 CPUs, capable of accessing a PetaByte (PB) of data and linked to computer centres around the world. 2004 was a pivotal year, marked by extraordinary and rapid change with respect to Grid deployment, in terms of scale and throughput. Following three years of intensive developments at 20 UK institutes, the required infrastructure was built up. The UK Grid is integrated seamlessly with the international LHC Computing Grid (LCG). The scale of the prototype was significant, with 10,000 CPUs linked and able to access 5PB of data across more than 100 institutes worldwide.

'The Grid is a practical solution to meet this and many other challenges in various areas of science, as part of the UK's e-Science programme.'

Sites across Europe are linked to those in the Far East and North America. This enabled the Economist to declare the LCG the world's largest Grid in October 2004. In the UK, hundreds of computers are currently accessible at each of 16 institutes. In addition, four sites linked via the UK e-Science National Grid Service are introducing the same 'middleware'.

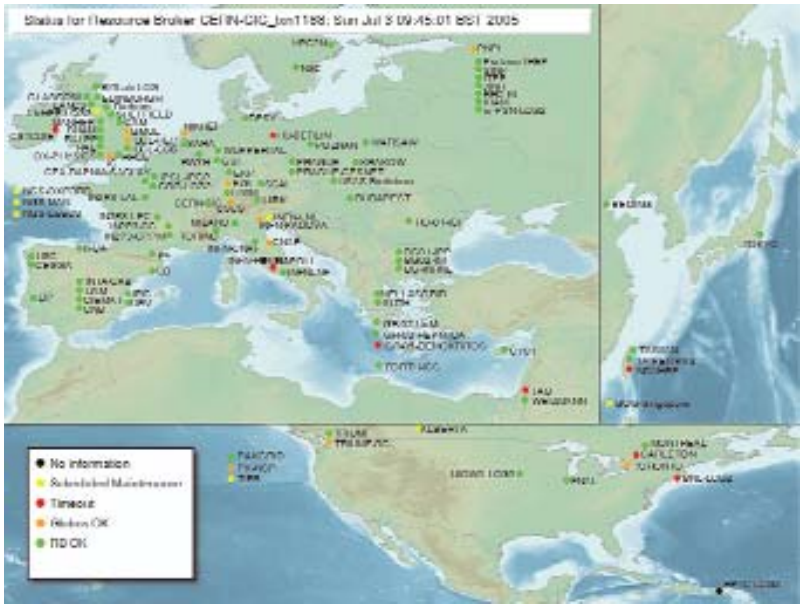

Middleware is the key to a successful Grid: the initial software stack of more than one million lines of code was developed and tested in conjunction with the European DataGrid and Enabling Grids for E-sciencE projects, building upon earlier work by the US-based Globus and Condor projects. This enables a user in Glasgow, armed with his/her virtual passport (a digital certificate), to submit an analysis job to a resource broker. As a member of a recognised experiment (or Virtual Organisation, VO), the job will run a few minutes later on any of the sites with pre-installed VO specific software and the required hardware for the job. This is a powerful generic methodology, applicable in many branches of science and industry, where access to large-scale computing resources are required, typically with demanding timescales, across the globe.

The system was intensively tested during 2004 via a set of planned stress tests by each of the four LHC experiments. Individual experiments accumulated up to 400 CPU years' worth of test data, with individual jobs running for up to a day, and a peak load of almost 6,000 simultaneous jobs achieved in August.

The GridPP project recently reached its halfway mark and started its second phase. Numerous issues have been identified that are now being addressed as part of GridPP2 planning in order to establish the required resource for particle physics computing. Further site and middleware validation tests are needed in order to improve the overall Grid efficiency. Each experiment needs to develop an individual application interface: in this regard, it is noteworthy that the system developed for the large LHC experiments has been shown to work effectively for other less resource-intensive applications. Addressing analysis group computing within the experiments, developing distributed file and database management systems, installing and validating experiment software distributions, setting up production accounting systems and creating an environment where everyone can share resources are all areas requiring further development.

The aim of GridPP2 is to deliver a 'Production Grid': a robust, reliable, resilient, secure, stable service delivered to end-user applications. The collaboration aims to develop, deploy and operate a 10,000 processor, multi-PetaByte Grid. During the last three years, the GridPP project has been at the forefront of developing and deploying pioneering Grid middleware. Tests of the prototype system have given us some confidence that the final system will be capable of providing the required resource for LHC data analysis – in this way, we plan to meet the particle physics Computing Challenge.

Professor Tony Doyle
Department of Physics and Astronomy
University of Glasgow
Glasgow G12 8QQ
Tel: 0141 330 5899
A.Doyle@physics.gla.ac.uk
www.gridpp.ac.uk

UNIVERSITY of GLASGOW