

# BABAR SPGRID - PUTTING BABAR'S SIMULATION PRODUCTION ON THE GRID

C.A.J.Brew, F.W.Wilson, G.Castelli, CCLRC – RAL, Chilton Didcot, OXON OX11 7XU, UK  
E.Luppi, D.Andreotti, INFN Sezione Di Ferrara, via Saragat 1, 44100 Ferrara, Italy

## *Abstract*

We describe enhancements to the BaBar Experiment's distributed Monte Carlo generation system to make use of European and North American GRID resources and present the results with regard to BaBar's latest cycle of Monte-Carlo production. We compare job success rates and manageability issues between GRID and non-GRID production and present an investigation into the efficiency costs of different methods of making input data, in the form of files and database information, available to the job in a distributed environment.

## INTRODUCTION

The BaBar High Energy Physics (HEP) detector [1] is based at the Stanford Linear Accelerator Center (SLAC), Stanford, US. The experiment investigates the subtle differences between matter and anti-matter by continuously colliding bunches of high-energy electrons and positrons 250 million times per second and searching for the creation of rare B-meson and anti-B-meson particles. High speed electronics and online processing rejects unwanted events resulting in a final raw event rate to tape of approximately 100Hz with each event requiring ~30kB. The raw data is sent to Padova (Italy), where the events are reconstructed. The reconstructed events are then separated (“skimmed”) into approximately 180 data streams based on physics properties using compute farms at Padova, Karlsruhe (Germany) and SLAC (US). These data streams are made available as datasets for analysis and used by 600 researchers based at 74 institutes in 10 countries. Since starting in May 1999, the experiment has recorded 3.5 billion events.

In addition to analysing the data, a major task is the simulation of the experiment. At least 3 times as many simulated events are needed as data events. Each simulated event takes ~10 seconds on a modern processor and results in 20kB of storage. Jobs of 2000 events are allocated to a distributed system based on 25 computer farms located in 5 countries with a total of ~1000 processors.

On each farm a complete BaBar software release needs to be installed locally or made available to the batch nodes via AFS [2] or NFS. A number of servers must be maintained locally: the location of the experimental and simulated data is provided by a MySQL or Oracle database; the experimental and simulated events are stored and accessed using the ROOT I/O protocol [3] either directly through NFS or, for heavily used farms, with a load-balanced, fault-tolerant file server called Xrootd [4]; an Objectivity [5] object-orientated database is required to store detector conditions, alignment and

calibration constants which are distributed to the jobs via either an Advanced Multithreaded Server (AMS) [6] or a Network File System (NFS). The system requires a great deal of local customisation with each farm maintained by a local manager. Load balancing between independent farms has to be done manually. Monitoring of jobs during production is difficult and there is no way to automatically ensure that all resources are available before commencing a production run.

Despite these drawbacks, the system has run with greater than 98% efficiency and produced 2.5 billion simulated events in 2004 alone. However, the coming years will see a tripling in the experimental data rate and a less manpower intensive system must be found to provide three times the resources. A possible approach is based on the tools and middleware provided by Grid systems such as Globus [7] or the LHC Computing Grid Project, LCG [8]. The goal is a production system with a single production manager submitting jobs to a worldwide Grid of remote sites with automatic resource allocation, job monitoring, output retrieval and cataloguing. Grid middleware, working over the Internet, provides the necessary infrastructure. Many Worker Nodes (WN) managed by Compute Elements (CE) with jobs directed by Resource Brokers (RB), metadata management by the Replica Location Service (RLS) and high volume of data kept on Storage Elements (SE).

## BABAR SIMULATION PRODUCTION ON THE GRID

We have investigated two possible solutions to this task. In the first, the minimum amount of Grid software was installed on top of legacy BaBar farms with data and databases accessed locally. Only the controlling job description was sent to the Grid using the Globus system. In the second approach, no assumption was made about local resources. The required BaBar software was pre-packaged and installed through a Grid account; data and databases were accessed remotely over the Wide Area Network (WAN) using the AMS and Xrootd protocols. The first approach offered the benefits of only minimal changes to pre-existing resources; the second offered the possibility of much reduced local customisation and support load.

### *Globus Based Submission*

In the UK, the Globus 2.4 gateway was installed on BaBar farms located at the Universities of Royal Holloway, Bristol and the Rutherford Appleton Laboratory (RAL). RAL acted as the controlling site. The Ganglia [9] monitoring program was installed locally to

provide graphical status of the batch workers. Each identical farm comprised of 40 dual CPU Linux machines, a Solaris machine used as an Objectivity server, two interactive Linux machines, one of which also acted as the OpenPBS [10] server and some NFS mounted storage. The standard production scripts were modified to submit remotely to the gateway; the jobs were passed onto the batch system and then retrieved at the end. Initially the Globus file staging functions were used to transport these files to and from the remote site, but this proved unreliable and it was replaced with GridFTP [11].

The production was automated by creating two daemons at the central site. The gateway was unable to deal with a large number of jobs so a control script was created to monitor the number of jobs running and queued at each remote site and submit more if the number fell below a certain threshold. It also interfaced to the standard BaBar tools to rebuild failed jobs and choose the next jobs for submission. The second daemon took the finished jobs, recovered the data, validated the output and updated the catalogue.

During 2004, this system was deployed at four sites across the UK. 31 million events were produced in nine months on this system with peak production rates in excess of 2.5 million events per week and more than 100 concurrent jobs at a given remote site. The failure rate in stable running was below five percent.

This system was limited to running on existing BaBar resources as the full BaBar software and servers were required. The sites were still separate farms and load balancing was not possible. No reduction in manpower was achieved as a complete local installation of the BaBar environment was required with constant updates to the local databases. A remote control programme was needed to limit the number of jobs submitted as the standard Globus gatekeeper spawned processes for every running and queued job which overwhelmed the gatekeeper and resulted in an unresponsive system. Job submission failures were also traced to poor interaction between certain firewalls and the TCP timeout behaviour.

### LCG Based Submission

The second approach used the full LCG infrastructure and prototyping was performed on the INFNGrid [12] farms directly involved in simulation production for BaBar (Fig. 1). LCG offers a number of possible advantages as the middleware was developed with HEP application use cases in mind. It provides an integrated data management system and a Resource Broker (RB) that can direct jobs to compatible resources without involvement of the submitter.

The LCG model also adds a number of challenges. The BaBar software is not available at the remote site. Therefore the complete BaBar software was reduced to a core subset that was packaged as a tar archive of about 37 MB (130 MB uncompressed). The package was then distributed on each Grid Worker Node by using procedures based on the LCG middleware to install and uninstall new software and to publish new tags. A site

manager, mapped as a special user, was able to use these procedures by simply submitting specific jobs on the Grid choosing which sites must be involved.

Three sites (Ferrara, Padova, Naples) were configured to provide Objectivity and Xrootd services that could be accessed from the Grid Worker Nodes over the WAN using AMS and the Xrootd protocol. The typical amount of conditions data read by a single simulation job was about 2% of the total and the use of the WAN reduced the CPU efficiency by about 5%, which was acceptable.

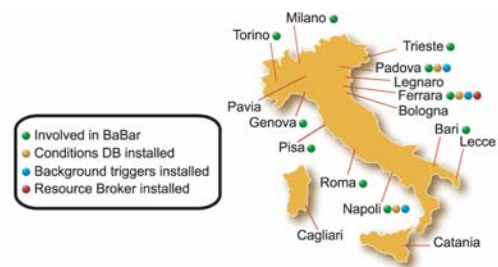


Fig. 1: Resources available for BaBar SP on the Italian Grid.

Standard BaBar software tools for simulation production were installed on the LCG User Interface (UI) and properly configured. These tools were needed to merge several collections of homogeneous events into a single collection, transfer events to SLAC, and to update the bookkeeping database at SLAC.

A JDL [13] script sets up all data needed and specifies the parameters for the simulation including type of events, run number, background triggers to use, number of events and detector conditions. Jobs were then submitted through an UI to a Resource Broker specifically installed for BaBar, located in Ferrara. The RB was able to manage Italian and European resources directly involved in BaBar. The RB performed matchmaking between resources available on each site of the Grid, published by Computer Elements (CE), and jobs requirements (memory, specific software release and type of the batch queue). Jobs were sent to CEs satisfying the requirements, and distributed to WNs, where the execution started. The simulated events were saved as ROOT I/O files and transferred from WNs to a Storage Element (SE) located in Ferrara. A conceptual schema of the flow described above is shown in Figure 2.




Fig. 2: BaBar simulation production model implemented in Italy

## Combined Approach

Between the 2004/5 (SP6) and 2005/6 (SP8) Simulation Production runs a decision to try to unify the two systems was made. The submission, recovery and monitoring daemons and scripts used for the Globus based production were adapted to work with the LCG based jobs and where feasible the existing UK BaBar resources were made available via LCG.

For the current run which started in July 2005 we are running two separate groups of sites; one in the UK with four sites and one with a single site in Italy.

## ACCESSING THE INPUT DATA

BaBar SP jobs require two types of input files in addition to the small text control files that are sent with the jobs: The Conditions Database contains the historical configuration and setup of the BaBar detector, it is stored in an Objectivity Database and accessed either over NFS or the proprietary AMS protocol. The Background Triggers are real background events recorded by the detector that are mixed with the simulated events during generation. They are stored in root files and accessed either over NFS or over Xrootd. Installing servers for this data and keeping the data up to date at each site would hardly reduce the effort required to maintain the system so a method of accessing this data remotely over the wide area network is needed. However, accessing the data over the wide area network has potential performance, security and stability issues.

### Objectivity Tests

During initial tests a single condition database installed in Ferrara was used as AMS server while jobs were submitted on Ferrara, Napoli, Trieste, Catania, Bari and Padova. A modified version of AMS, written at SLAC [14], was designed to handle more connections. It was installed in Ferrara and its performance was compared to the standard AMS. Both AMS servers used 32 threads to handle client requests. Stress tests, performed by submitting from 100 to 250 jobs of 2000 events each, showed that standard AMS supported about 30 concurrent simulation jobs before failing to respond. The number of failed submissions increased with increasing load, reaching almost 50% when 80 clients contacted the server simultaneously. Performances improved using the SLAC AMS server. The server was able to handle about 75 clients before jobs failure occurred.

The maximum number of AMS connections a client can generate could be configured in the job script. Limiting the number of connections increased the execution time for some jobs while allowing more jobs to contact the AMS in parallel. A good balance between efficiency and stability was achieved, setting the maximum number of connections to 20 for each client.

After these initial tests, 2 tests were performed accessing the AMS server at different sites. For every test, 200 jobs were submitted, each with 2000 events in order to simulate a real production flow. For both tests, the time

spent to complete jobs on each site was measured. The maximum number of concurrent clients connected during production was taken into account. The link speeds to Bari, Ferrara, Napoli and Padova were 15, 16, 32 and 160 Mb/s respectively.

Data were read using the AMS server installed in Ferrara, while production jobs have been submitted on Worker Nodes located in Ferrara, Padova and Bari. Most of the jobs run on Padova due to the high number of WNs provided. About 99% jobs successful completed (Table 1), a success rate comparable to the non-GRID production.

Site	Jobs completed	Mean elapsed time
Bari	29	10:01:04
Ferrara	29	08:02:04
Padova	141	07:28:44
<i>Total</i>	<i>199</i>	<i>07:55:47</i>

Table 1: Data read from AMS located in Ferrara

In the second test, data were read remotely, using the AMS server installed in Naples from production jobs submitted on Worker Nodes located in Ferrara, Padova and Bari (Table 2).

Site	Jobs completed	Mean elapsed time
Bari	43	09:41:43
Ferrara	12	09:00:46
Padova	130	07:38:32
<i>Total</i>	<i>185</i>	<i>08:12:30</i>

Table 2: Data read from AMS located in Naples

Due to the fast network link, jobs at Padova completed more quickly than jobs submitted on other sites. The quality of network access and data location played a basic role in turn around times and is an important parameter in choosing the optimal running conditions. The whole system was tested producing real allocations data for Monte-Carlo (more than 3,000,000 events), successfully completing the whole production cycle for about 98% of the total jobs.

A recent attempt to repeat these tests in the UK between RAL, Manchester, Oxford and Lancaster with Objectivity servers at Manchester and RAL failed due to firewall issues. It appears that the Objectivity access pattern for Simulation Production is such that some firewalls either blocking the AMS connections completely or close the connection midway through the job. This is still under investigation.

### Background Trigger Tests

Three methods of reading the Background Trigger root files were compared:

1. Directly out of a local xrootd server using the xroot protocol, this is the current standard method of reading the Background Trigger files either from a local or remote server.

2. Directly out of a local dCache[15] storage element with PNFS[16] and the lpcap preload library using the DCAP protocol.
3. Indirectly out of a dCache storage element by using the LCG data management tools to copy the root files to local disk (using the GridFTP protocol) then reading the data from there. Although we used a dCache storage element for this test the results should be comparable for any type LCG storage element.

The second and third methods tested are possible ways of storing and retrieving the Background Trigger files at the remote execution site without needing any BaBar specific services installed there.

A set of ninety-seven, 2000 event jobs were run once for each of the different access methods on the same set of hardware and the mean time to complete the job was calculated for each access method. The Background Trigger root files used by the jobs totalled 2.5GB. The results are summarized below in Table 3.

Access Method	Mean elapsed time
Xrootd	3:49:48
dCache/DCAP	3:48:36
SE/GSIFTP (inc. Transfer Time)	3:54.12
SE/GSIFTP (not inc. Transfer Time)	3:52:48

Table 3: Comparison of job time for different methods of accessing Background Triggers

It is evident from the figures that the time taken to read the data is low enough compared to the CPU time that the input method is almost irrelevant. Even when the time taken to copy the files to the local disk is taken into account the effect is negligible. Any of the tested methods is a viable method for production but it also appears that accessing the data from a remote xrootd server or even copying it from a remote storage element is feasible as long as the worker nodes have external network access.

## CONCLUSIONS

The simulation of HEP events is now feasible using Grid architecture; so far in BaBar's 2005/6 Simulation Production run more than 372 Million events have been generated on LCG based grids in the UK and Italy, (approximately 8% of the total). Reliability is beginning to approach that of local self-contained compute farms. For the 170,000 jobs run in the UK so far in the 2005/6 run we have a measured efficiency of 93% (compared to roughly 98% efficiency for "normal" production) The Grid can be interfaced to legacy software and procedures. As most non-trivial HEP programs require access to centralized or distributed resources, our prototyping shows that the major architectural problems concern the installation of these servers and their efficient access over

the WAN. Future work will concentrate on simplifying the identification of data, providing a common environment on all Grid resources and providing servers to provide regional access over the WAN to centralized databases.

## REFERENCES

- [1] B. Aubert et al., "The BaBar detector," Nucl. Instruments. & Methods, Vol. A 479 (2002), pp. 1-116
- [2] J. H. Howard, Carnegie Mellon University, "An Overview of the Andrew File System," USENIX Conference Proceedings, Winter 1988
- [3] Rene Brun and Fons Rademakers, "ROOT - An Object Oriented Data Analysis Framework," Nucl. Instruments & Methods, Vol. A 389 (1997), pp. 81-86. See also <http://root.cern.ch/>.
- [4] A. Hanushevsky, "The Next Generation Root File Server," in CHEP proceedings, Interlaken, Switzerland, September 2004.
- [5] Objectivity Inc. [online]. Available: <http://www.objectivity.com>
- [6] AMS: J. Becla, "Setting up AMS", 2002, [online]. Available: <http://www.slac.stanford.edu/BFROOT/www/Public/Computing/Databases/experts/setupAMS.shtml>
- [7] The Globus alliance, [online]. Available: <http://www.globus.org/>
- [8] CERN, "LHC Computing Grid Project", [online]. Available: <http://lcg.web.cern.ch/LCG/>
- [9] Ganglia, [online]. Available: <http://ganglia.sourceforge.net/>
- [10] OpenPBS, [online]. Available: <http://www-unix.mcs.anl.gov/openpbs/>
- [11] GridFTP, [online]. Available: <http://www-fp.globus.org/datagrid/gridftp.html>
- [12] INFNGrid: GridIt, INFN Production Grid, [online]. Available: <http://www.grid.it/>
- [13] Job Description Language: DataGrid-01-TEN-0102-0\_2, 2001 Datamat
- [14] SLAC AMS: A. Hanushevsky, "Running Multiple AMS Servers per Host", 2002, [online]. Available: [http://www.slac.stanford.edu/~abh/Objectivity/Misc/mult\\_ams.html](http://www.slac.stanford.edu/~abh/Objectivity/Misc/mult_ams.html)
- [15] dCache [online] Available <http://www.dcache.org/index.shtml>
- [16] PNFS [online] Available: <http://www-pnfs.desy.de/>