Putting BaBar's Simulation Production on The Grid

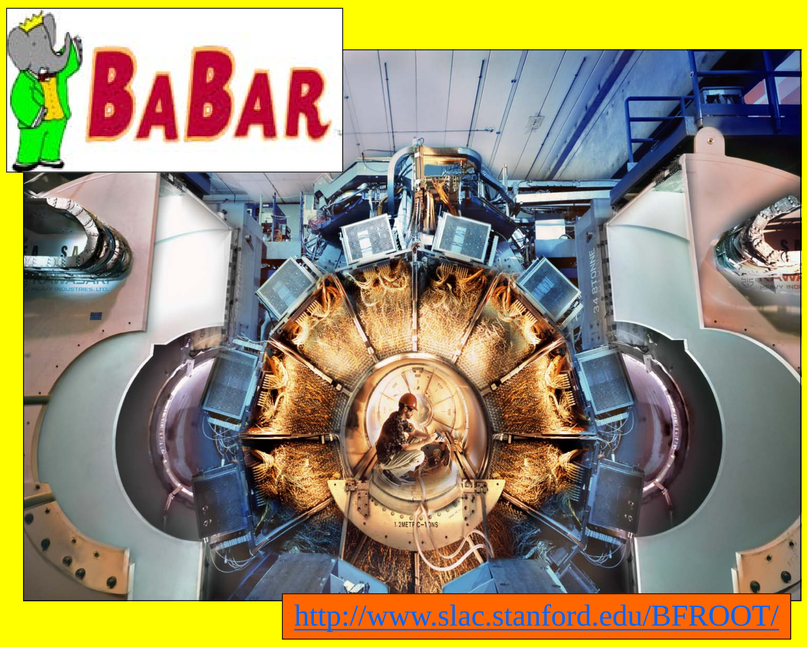

Current Setup and the Problem

- Since starting in 1999, BaBar has recorded 3.5 billion events.
- Machine improvements mean data is now taken at ever higher rates. In 2005/6 we aim to take as much data as in the previous 6 years.
- ~3:1 ratio of simulated to real events is needed.
- Each simulated event takes ~10 seconds on a modern processor and produces ~20KB of output data.
- Currently events are produced on 25 computer farms located in 5 countries with a total of ~1000 processors.
- Each site has a person who installs and maintains the software and monitors the production.
- Run with greater than 98% efficiency and produced 2.5 billion simulated events in 2004.
- 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.

Is Grid the Solution?

- The challenge is to keep up with the increased rates of data taking.
- Can use the tools and resources available to us on the LHC Computing Grid to do this.
- The goal is a production system with:
  - Single production manager defining jobs.
  - Automatic job creation and submission.
  - Worldwide Grid of remote sites.
  - Automatic resource allocation.
  - Real-time job monitoring.
  - Automatic output retrieval and cataloguing.
  - Dynamic failure detection and recovery.
- This should be implemented as a modular design capable of being interfaced to multiple Grid back-ends (e.g. LCG, OSG, ...).

Comparison of the Steps in Local and Grid SP Running

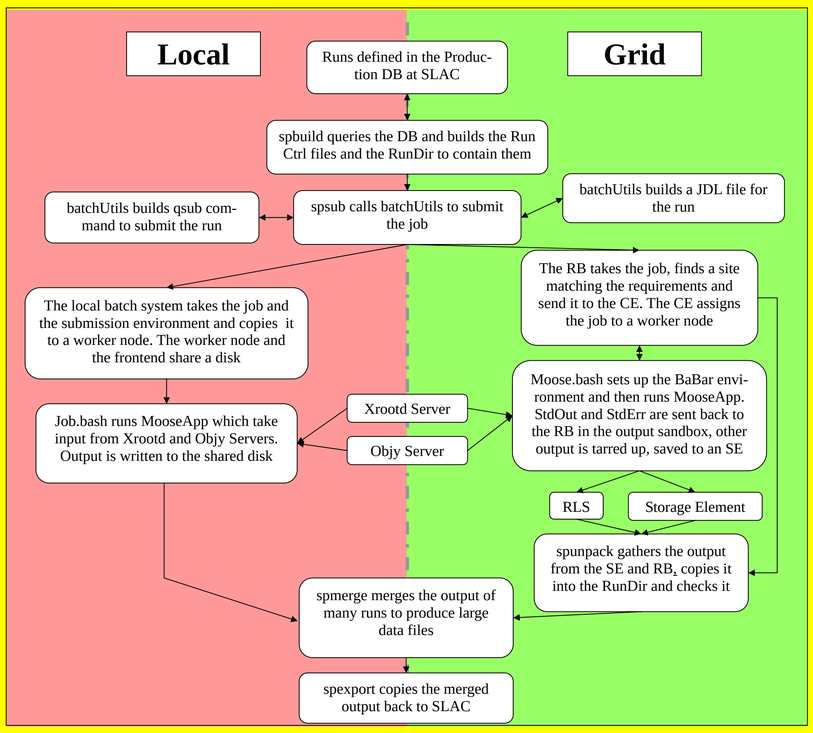

Monitoring Production

- We have developed a number of Web based Monitoring tools:

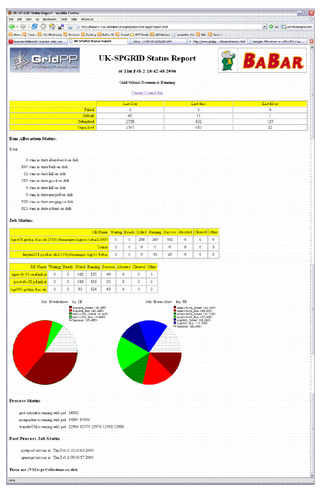

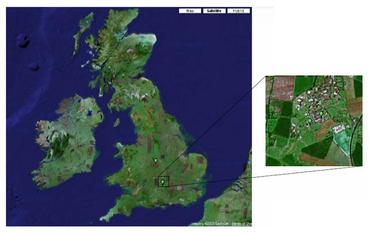

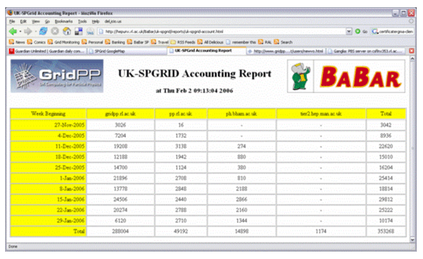

Current Status

- SPGrid is currently running two separate systems on 4 farms in the UK and on sites in Italy and is working on adding another two farms in the near future.
- Since July 2005 we have generated 370 Million events in the UK with peak rates of about 30 Million evts/week.

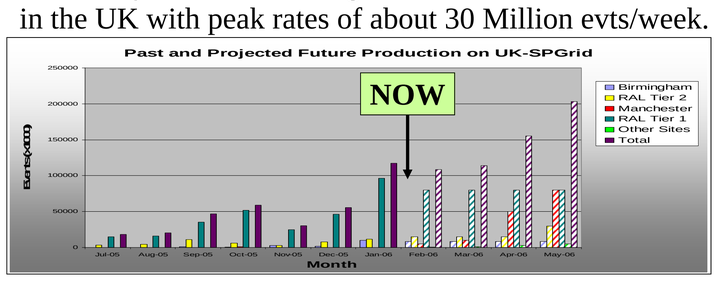

- Failed jobs are automatically detected and resubmitted.
- The overall failure rate for jobs is about 7%.
- One obstacle to increasing the number of sites for Grid production is access to the input data needed by the job.

Experimental Conditions

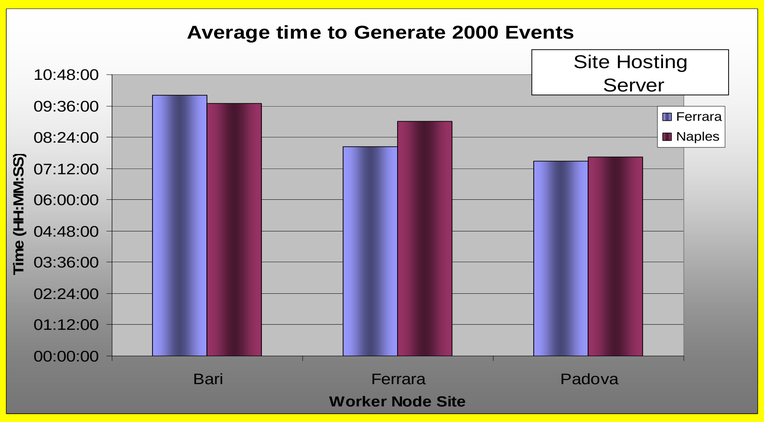

Reading the Conditions database from a remote Objectivity server has a marginal effect on the job's execution time. Unfortunately, the connection to the Objectivity server seems to be very sensitive to network glitches and firewall issues.

Accessing Input Data

- BaBar Monte-Carlo production needs access to two different types of input data:
  - Experimental Conditions and Configuration Information accessed from Objectivity Databases.
  - Background Events recorded by the detector to be mixed with the generated events. These are read from root files either via xrootd or from disk.
- Adjacent are two plots showing the effect on the length of a job of different ways of accessing this data:
  - Left: Accessing Conditions data from a local or remote server.
  - Right: Accessing the background triggers directly from xrootd and dCache and copying the root files from a local Storage Element (with and without the time to copy the files).

Background Events

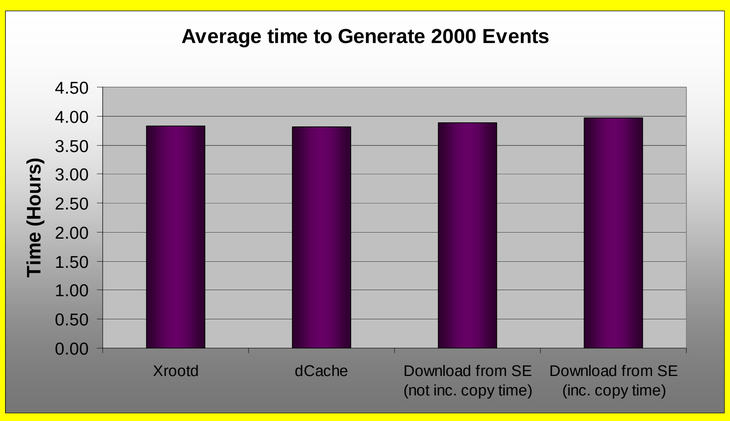

It is evident that the read rate for the background triggers is so low that is has almost no effect of the job's execution time - even including the time to copy ~3GB of data to the local disk is negligible compared to the length of the job.

Next Steps and Conclusions

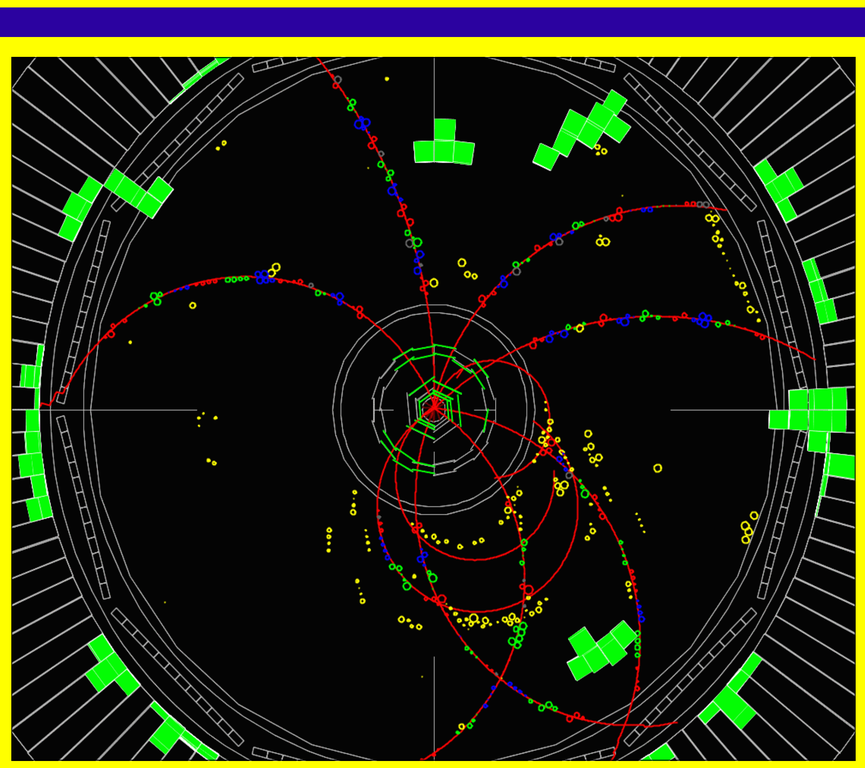

- Next Steps:
  - Expand the Production to more UK and Italian sites (including Italian Tier A resources) and look to run jobs at sites in other countries.
  - Integrate the UK and Italian resources (+ other LCG resources at BaBar sites) for a "Data Challenge".
  - Move to the Root based conditions Database and deploy servers which support multiple sites with a single protocol.
  - Develop an R-GMA based monitoring system which will allow us to track the progress of individual jobs.
- Conclusions:
  - Production of simulated data on the Grid allows BaBar to access resources it would not otherwise be able to.
  - We are already producing significant amounts of data and have doubled the production rate in the last three months and hope to do so again in the next three.

Left and right panels show an simulated $B^0\bar{B}^0$ event created on 12 October 2005 with UK-SPGrid.

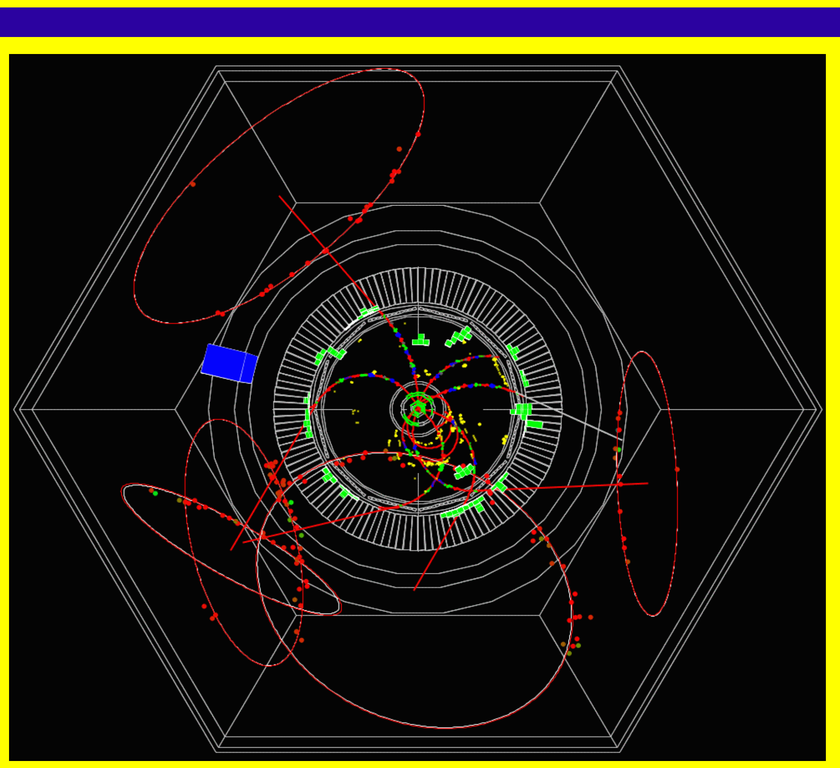

Presented by Alessandra Forti[1] on behalf of:
Chris Brew[2], Giuliano Castelli[2], Fergus Wilson[2], Daniele Andreotti[3], Elenora Luppi[3]
and the BaBar Computing Group

(1) University of Manchester (2) CCLRC - RAL (3) INFN SEZIONE DI FERRARA