

# MASSIVE DATA PROCESSING FOR THE ATLAS COMBINED TEST BEAM

Mireia Dosil, Port d'Informació Científica\*, Barcelona, Spain  
Frederik Orellana, Institute of Nuclear and Particle Physics†, Geneva, Switzerland  
on behalf of the ATLAS Combined Test Beam Group

## Abstract

In 2004, a full slice of the ATLAS detector was tested for 6 months in the H8 experimental area of the CERN SPS, in the so-called Combined Test Beam, with beams of muons, pions, electrons and photons in the range 1 to 350 GeV. Approximately 90 million events were collected, corresponding to a data volume of 4.5 TB. The importance of this exercise was two-fold: for the first time the whole ATLAS software suite was used on fully combined real data. Besides, a novel production infrastructure was employed for the reconstruction of the real data as well as for a massive production of simulated events.

This paper is a report on distributed Combined Test Beam Monte Carlo production with and without grid tools. In 2004, Monte Carlo production was started on the CERN LSF batch system and later it was, for the first time, performed on the LCG, with the simulation of about 4 Million events. In 2005, a more light-weight and user friendly system was tested on NorduGrid for the quick production of 210'000 photon events and 680'000 pion, muon and electron events. New Monte Carlo productions are already planned for the year 2006, using the new ATLAS production system.

## INTRODUCTION

The experimental setup for the 2004 ATLAS Combined Test Beam (CTB) was composed of elements from all ATLAS subdetectors [1]: Inner Detector (Pixel, Semiconductor Tracker and Transition Radiation Tracker), Calorimeters (Electromagnetic Liquid Argon Calorimeter and Hadronic Tile Calorimeter) and Muon System (Monitored Drift Tubes, Cathode Strip Chambers, Resistive Plate Chambers and Thin Gap Chambers) as shown in Fig. 1.

The detector was tested for six months with beams of pions, muons, electrons and photons in the energy range of 1 to 350 GeV. About ninety million events with an average size of 50 KB each, were collected and stored on the CERN tape system CASTOR [3], taking up  $\sim 4.5$  TB of space in total.

The 2004 CTB was an ambitious pre-commissioning exercise, since, for the first time, the complete software suite developed for the full ATLAS detector was used for real data. Moreover, Monte Carlo simulations were per-




Figure 1: ATLAS Combined Test Beam layout as visualized by the Geant4 simulation package [2].

formed with the specifications of the already existing collected real data, so that the physics working groups could perform accurate comparison studies.

## DATA PRODUCTION AND PROCESSING TOOLS

In order to manage simulation, digitization and reconstruction of such volumes of data, the production systems already used for the series of ATLAS Data Challenges (DC) [4] were employed. A number of adaptations and configurations were necessary in order for these systems to be usable for CTB data.

The metadata associated with each produced dataset, like beam energy, beam polarity, etc. was stored in the ATLAS metadata catalog (AMI) [5], which is used by physicists to find the run numbers corresponding to a given physics channel.

Ideally, when running production, working on the dataset level is desirable: it should be possible to start running a large-scale processing of data by specifying just the input dataset name, the output dataset name and a transformation script. At a later stage, after all jobs have finished, it should then be possible to publish information about the produced data to a central catalogue (AMI). By the efforts described in the next section we explored how much of this functionality is available with the systems described above and how suited they are for activities like CTB production.

Apart from the metadata and runtime databases, running a production job typically requires access to four databases at CERN: The primary numbers for the detector description are stored in a MySQL database called NOVA. Information about the conditions of the various detectors (e.g. sub-detector calibration) for a given time frame are stored as c++ objects streamed into POOL ROOT files. These files are registered in an Oracle database, the so-called conditions

\* Universitat Autònoma de Barcelona. Edifici D, E-08193 Barcelona (Spain)

† Université de Genève. 24, Quai Ernest-Ansermet, 1211 Geneva 4, Switzerland

database. Additional information about the liquid Argon calorimeter is kept in two additional MySQL databases.

### A GUI for the LSF batch system

For production on the CERN LSF batch system, input and output data files were read and written directly from and to CASTOR. Moreover, the job scripts themselves and additional input files, like Athena [4] job options, compiled shared libraries and directories with code, were accessed directly from the file system, using AFS. A GUI was developed to manage submission and monitoring of jobs and to have validation and bookkeeping done automatically. The GUI is an extension of the Java program AtCom [6], which was used to manage the production for the ATLAS Data Challenge 1 [7]. Effort was made to make the tool self contained, so users should not need a collection of other tools, like other MySQL interfaces, command-line tools for querying job status, etc. Also, effort was made to end up with one tool, applicable to both the reconstruction of real data as well as the simulation of Monte Carlo data, despite the different database schemas used. The result was a tool which was used by  $\sim 10$  users for processing millions of events (see below).

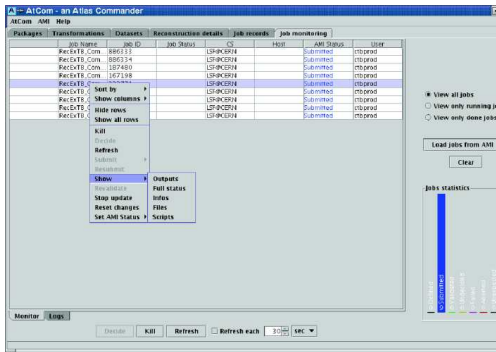


Figure 2: Monitoring panel of the GUI for running CTB production job.

### CTB simulation jobs on the ATLAS LCG production system

In spring 2005, the first simulation CTB jobs were successfully run on the LCG using the ATLAS DC 2 production system [8]. The ATLAS DC 2 production system was made of four components: a central production database (prodDB), a data management system, a supervisor daemon and an executor daemon. The system distinguishes two levels of abstraction: input datasets are transformed to output datasets by applying a *task transformation*. A dataset is a group of logical files containing events with the same physics specifications. Input logical files are transformed into output logical files by means of a *job transformation*.

CTB simulation jobs were filled in prodDB with a Python script. The supervisor took free jobs from the

database and sent them out to the LCG executor it was connected with by exchanging XML messages. The executor then submitted the jobs to the LCG. After jobs finished, a cron job ran and filled the AMI metadata database from the CTB job entries in prodDB.

### A GUI for NorduGrid

The relative ease with which new users were able to perform production tasks on the LSF batch system, using the mentioned GUI, and the amount of resources available on the grid systems in use by ATLAS, encouraged us to extend the GUI to a light-weight system for running production jobs on a Grid system. This was done using the plugin architecture of the original AtCom. We chose to start with NorduGrid, primarily because it too is light-weight and installable on several clusters of which we had control, avoiding reliance on too many externals. Several new features and work-arounds were also implemented:

- File copying moved from the job shell script to the Grid job description file.
- Reduction of the number of input parameters.
- Automatic job splitting according to number of events per output file.
- New panels for editing and deleting transformation, dataset and logical file records.

Datasets are defined by directly filling the fields of the database record. For each dataset, jobs can be defined. This is done by filling a form with typically 15 values (depending on the schema in use), allowing the use of the special variable \$i which is iterated over, and \$1, \$2, ..., to access the fields of the dataset record. The interaction with the Grid system is done via an intermediate remote shell started on a server with the NorduGrid client software installed and file system access to the production scripts (e.g. lplus.cern.ch). After initial tests, we observed the following:

- The failure rate is very low (a few percent).
- Despite the shell interaction with the Grid system, the GUI appears responsive and pleasant to work with.
- The job submission time is of the order of 1.2 seconds per job when on a fast connection. The submission time increases on a slow connection. All in all, working with more than a few thousand jobs is not manageable.
- Typing in dataset definitions is still too involved. Given the run number, most remaining parameters, like beam energy, etc. should be taken automatically from other databases.
- Filling in the job record form should not be necessary. In most cases, reasonable default values suffice.

We addressed the last two points by providing functionality for creating new datasets on the basis of existing ones, either by cloning or by using the existing ones as input datasets for the new ones. Also, the specification of the location of input and output files was moved from the logical file record to the dataset record, thus simplifying the job definition.

## DATA PRODUCTION AND PROCESSING

In this section we report on the actual data production and processing carried out with the various tools described in the previous section.

### *Real data reconstruction on LSF*

The reconstruction of CTB data was performed by a package (RecExTB) that integrated all the subdetectors' reconstruction algorithms through different steps. The program gave as output Event Summary Data (ESD) and a ROOT Ntuple needed for physics analysis.

For the massive reconstruction of real data, the CERN LSF batch system was used through the aforementioned GUI. Jobs were intentionally not submitted to resources outside of CERN since, at that time, no replication of the conditions database was foreseen and the central database server had access restrictions prohibiting worldwide access.

A subset of 400 good runs, corresponding to approximately 25 Million events, was selected by the CTB community to be used for physics studies, and reprocessing of this data was performed with all major ATLAS production releases (8.8.0, 9.1.2, 10.0.2, 10.3.0, 10.4.0). Large-scale real data reconstruction was carried out twice in 2005 with two different software releases. Half of the runs were processed with the fully combined set-up and half of them with combined calorimetry only. The average CPU time per event was of the order of 1.5 KSI2K second. All the runs were split into logical files of maximum 10'000 events each and for each job two types of output files were produced and stored on CASTOR: ROOT files with the combined ntuples (20 KB/event) and POOL files with ESD (40 KB/event). For later user analysis, all this information was kept in the AMI metadata database. These runs were rather successful, with an overall failure rate of about 2%. New reprocessing of the data is already foreseen in 2006 with the next major ATLAS software releases.

### *Monte Carlo data production on the LCG*

In view of the June 2005 ATLAS Physics Workshop held in Rome, a big Monte Carlo data production was foreseen using the ATLAS LCG production system.

The simulation of the CTB setup was performed within a simulation framework based on the Geant4 toolkit [9]. The package "CTB\_G4Sim" performed the simulation and digitization. The output of the digitization was reconstructed

with the same RecExTB package that was used for real data reconstruction.

The Monte Carlo production for the 2004 CTB was divided in two parts: a preproduction phase and a production one. During preproduction, more than 1.7 Million Monte Carlo events were simulated, digitized and reconstructed using the developed GUI for running on the LSF batch system and AMI for bookkeeping. The goal of this phase was two-fold: to test the different software pieces in order to get a stable version and on the other hand, to compare the first reconstructed Monte Carlo events with the ones from real data reconstruction. The mean processing time per event ranged between  $\sim 2$  KSI2K seconds for muon simulation and  $\sim 50$  KSI2K seconds for photon simulation.

In May 2005 the production phase started: a total sample of 392 runs corresponding to almost 4 million events were simulated using the LCG ATLAS production system. This simulation exercise was important because it was the first time ATLAS events were simulated with the same conditions as those of the good "real" runs selected by the CTB physics community. Outputs of simulation were stored on the CERN storage element. 57% of the jobs finished successfully with only one attempt as can be seen in Fig. 3.




Figure 3: Distribution of the attempt number that a CTB simulation job reached on the LCG.

Digitization and reconstruction jobs were performed using the GUI for the LSF batch system, since the CTB conditions database was not replicated worldwide as stated before.

### *Monte Carlo data production on NorduGrid*

Monte Carlo production on NorduGrid was performed with the simulation framework described above. For job submission, the GUI for NorduGrid was used (see above). The clusters involved were those forming the Swiss subset of NorduGrid [10].

Initially, 21 datasets with 210'000 photon events were simulated and digitized with release 10.4.0 of the ATLAS software.

- Output data files were copied directly from the clusters to CASTOR at CERN.
- Log files were copied to a directory in AFS at CERN.

- Bookkeeping information on the produced data was kept on the AMI MySQL server in Grenoble.
- Runtime information on the running jobs was also kept on the AMI MySQL server in Grenoble.
- Conditions and geometry data was read from central databases at CERN.
- The average simulation time per event was 25 KSI2K seconds.
- The average event size was 57 kB for the simulated data and 46 kB for the digitized data.

To summarize, this small production employed technology proven to work in previous productions, but on a set of clusters outside of CERN. Apart from problems on the clusters (network configuration, batch system configuration, software installation etc.) and problems with the ATLAS software (getting Athena jobOptions right, compiling and including the right additional code), the issues encountered were:

- Accessing the conditions and geometry databases at CERN was hampered due to their limit of 500 simultaneous connections.
- Copying output simulated data to CASTOR at CERN and reading it back via gridftp introduced unnecessary complications
- Having bookkeeping and runtime information on a server in Grenoble was an unnecessary complication.

To address these, it was decided to:

- Switch to using static replicas of the 4 central databases at CERN: these replicas were all hosted on a MySQL server running on a machine at the university of Geneva.
- Copy and read data files to and from a disk-based gridftp server instead of the CASTOR gridftp server. The replication of output files to CASTOR was then done manually in a separate step.
- Have bookkeeping and runtime information on a local mysql server and replicate the bookkeeping information to the AMI server in Grenoble after the end of production.

This proved to make the infrastructure more stable and for the next production, failure rates became very low (a few percent). The failures were always due to problems connecting to overloaded gridftp servers. The production consisted of 68 reference datasets, selected by the CTB community, containing 680'000 pion, muon and electron events. In the first run, these events were simulated and digitized with release 11.0.2 of the ATLAS software. The infrastructure put in place should make it an easy task to reprocess these events with coming releases of the ATLAS software.

## CONCLUSIONS AND OUTLOOK

Various tools and systems have been explored for producing and processing CTB data. Compared to the ATLAS Data Challenges and the ongoing centralized ATLAS production, this is a small effort, serving the needs of specific physics communities.

From the computing perspective, the result has been both to provide feedback to the ATLAS software and production teams and to put in place a light-weight system for small to medium scale production.

New CTB Monte Carlo productions are already scheduled for 2006, both using the new ATLAS distributed production system on the three grid flavours and, for producing smaller number of events, the graphical tool described above. Such productions will be of capital importance for finishing the real data - Monte Carlo comparison studies before the start of ATLAS data taking in 2007.

## ACKNOWLEDGEMENTS

We would like to thank the physicists of the ATLAS CTB group for a fruitful collaboration. Thanks are also due to the ATLAS production group who has provided several of the tools we use, and to the ATLAS database group and the LCG/ARDA group for a stimulating collaboration.

## REFERENCES

- [1] ATLAS Collaboration, ATLAS Detector and Physics Performance Technical Design Report, Technical Report CERN/LHCC 99-14, CERN, Geneva, 1999.
- [2] M. V. Gallas et al, "ATLAS detector simulation: status and outlook", IEEE Nuclear Science Symposium, Puerto Rico, 2005.
- [3] <http://castor.web.cern.ch/castor/>
- [4] ATLAS Collaboration, ATLAS Computing Technical Design Report, Technical Report CERN/LHCC 2005-022, CERN, Geneva, 2005.
- [5] S. Albrand, J. Fulachier, "ATLAS metadata interfaces (AMI) and ATLAS metadata catalogs", Computing in High Energy and Nuclear Physics (CHEP), Interlaken (Switzerland), 2004.
- [6] V. Berten, L. Goossens, C. Tan "ATLAS Commander: an ATLAS production tool", hep-ex/0305089, 2003.
- [7] G. Poulard, "ATLAS Data Challenge 1", Computing in High Energy and Nuclear Physics (CHEP), La Jolla, San Diego (California), 2003.
- [8] L. Goossens, K. De, "ATLAS production system in ATLAS Data Challenge 2", Computing in High Energy and Nuclear Physics (CHEP), Interlaken (Switzerland), 2004.
- [9] Geant 4 collaboration, "GEANT 4: A Simulation Toolkit", Nuclear Instruments and Methods in Physics Research, NIM A 506 (2003), 250-303.
- [10] S. Gadowski, C. Haeberli and F. Orellana, "Prototype of the Swiss ATLAS Computing Infrastructure", Computing in High Energy and Nuclear Physics (CHEP), Mumbai (India), 2006.