SOFTWARE FOR THE CMS COSMIC CHALLENGE

G. Bruno, Université catholique de Louvain, Louvain-la-Neuve, Belgium
on behalf of the CMS Collaboration

Abstract

At the end of 2004 CMS decided to redesign its offline software framework. The new software now includes a completely revisited event data model and is fully integrated with a database infrastructure for handling calibration and alignment data. This new software will be used in the first months of 2006 for the so called Magnet Test Cosmic Challenge (MTCC). The MTCC is a slice test in which a small fraction of all the CMS detection equipment is expected to be operated in the 4 T solenoid of the experiment. Cosmic rays detected in the muon chambers will be used to trigger the readout of all detectors in the general data acquisition system. Prior to data taking, the detectors and their readout electronics must be tuned and synchronized with dedicated software procedures. Local reconstruction must be carried out online and offline in all sub-detectors for event selection and monitoring purposes. Global reconstruction, linking different sub-detectors, is expected to be performed mainly offline. One of the main goals of the offline analysis of the MTCC data is the validation of the muon hardware alignment system functionality. A number of visualization tools are also expected to be used for validation purposes and monitoring.

THE MAGNET TEST COSMIC CHALLENGE

The Compact Muon Solenoid (CMS) is a general purpose detector [1] that will be operated at the Large Hadron Collider (LHC) [2].

Before the LHC start-up, CMS has scheduled a single combined test of its detecting equipment. The superconducting solenoid and a fraction of all sub-detectors, with the exception of the pixel and the preshower sub-detectors, will be operated from April 2006 in the so called Magnet Test Cosmic Challenge (MTCC). Participation of the muon RPC sub-detector in the MTCC is still uncertain. The L1 Trigger electronics of the muon detectors will be used to detect cosmic muons and provide a signal to trigger the readout of all detectors in the global DAQ.

The muon Drift Tubes (DT) setup consists of 14 stations distributed on 4 layers, as shown in Figure 1. Each station provides measurements to build two φ-view segments and one θ-view segment. Thirty-eight Cathode Strip Chambers (CSC) chambers are going to be operated at the MTCC. Each chamber can provide one three-dimensional segment. Two Electromagnetic Calorimeter (ECAL) supermodules and 15 Hadron Calorimeter (HCAL) wedges will be present at the MTCC. The HCAL is also equipped with L1 Trigger electronics, which, however, is not expected to be used as primary L1 Trigger. The MTCC setup is completed with 133 single and double-sided Silicon Strip Tracker (SST) modules. Up to six SST detecting layers can be traversed by a straight track.

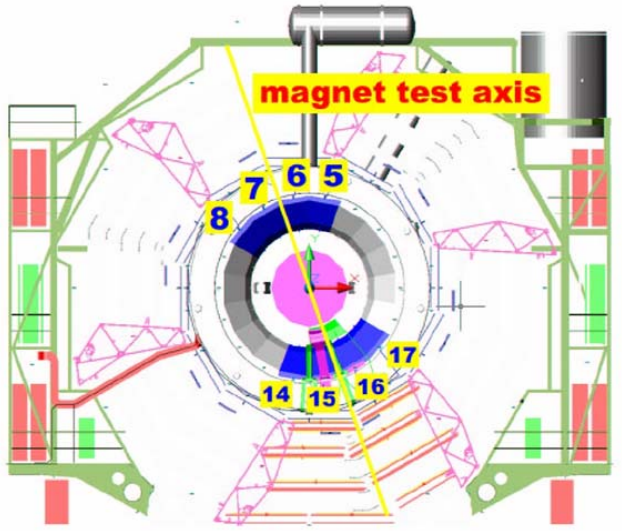

Figure 1: CMS cross section. The muon DT and the barrel calorimeter detectors that will be operated at the MTCC are visible.

The fraction of the final system represented by the setup described above and the expected data volumes are reported in Table 1.

Table 1: Fraction (with respect to final system) of each sub-detector present at the MTCC and data volumes with electronics in Zero Suppression (ZS) mode and in non ZS mode as expected at the MTCC.

Detector	Fraction of final system (%)	ZS data volume (kB/ev)	Non ZS data volume (kB/ev)
DT	5	1	
CSC	8	20	
HCAL	10	5	30
ECAL	5	<5	80
SST	1	<1	200
Total		30	330

Online event selection in the CMS experiment occurs in only two physical systems: the Level 1 Trigger and the High Level Trigger (HLT). The HLT selection is entirely software implemented and takes place in the Filter Farm (FF) of the CMS Data Acquisition system (DAQ) [3]. The FF nodes, called Filter Units (FUs), make use of the full event information to reconstruct in the most accurate way the physics objects in each event.

The DAQ system at the MTCC will comprise a two stage event builder, as in the final system, and a FF equipped with 16 Dual-Xeon 2.8 GHz processors. The maximum data on disk rate is expected to be about 50 MB/s. Therefore, given the data volumes reported in Table 1, the maximum event on disk rate will be of the order of 1 kHz in case the readout electronics is operated in Zero Suppression (ZS) mode. This value becomes 100 Hz if all detectors are read out in non-ZS mode.

A large fraction (30%) of the muon hardware alignment system [4] will also be tested at the MTCC. This laser based system measures to an accuracy of 100 μm the positions of the muon support structures with respect to each other and to the main inner tracker support structure. The hardware alignment system data is readout through a dedicated data acquisition system.

SOFTWARE ARCHITECTURE

At the end of 2004 CMS decided to redesign its offline software framework. One of the main goals of the MTCC is to test the full basic functionality of the new software, called CMSSW, in an experimental environment.

Event Data Model

The CMSSW event data model [5] requires that all event data processing modules communicate only through a single data structure called the "Event". Modules are executed as plug-ins in an order specified in the job configuration. A special input module can be used to insert in the Event the raw data received from the CMS DAQ. Thus, the offline software can run in a completely transparent fashion both in purely offline applications and in the HLT.

Reconstruction modules get from the Event the needed input data and put higher level reconstructed data back into it. Analyzer modules only have read access to the Event. They are widely used for monitoring as described later. Filter modules are used online to avoid forwarding events to the Storage Manager application, whose task is to collect accepted events from several FUs nodes and assign them to event streams identified on the basis of the HLT pattern.

Non-Event data and Database Infrastructure

The data processing modules need non-event data like cabling, calibrations and alignment information. CMS non-event data is handled with a complex Database (DB) infrastructure fully integrated with the offline software. Three DBs can be distinguished:

- OMDS: DB containing the configuration and conditions data of each sub-detector
- ORCON: DB containing the non-event data needed by the CMSSW applications running the HLT selection and monitoring tasks in the FF.
- ORCOFF: DB containing the non-event data needed by the CMSSW applications running pure offline event data analysis.

CMSSW is interfaced to ORCON for the HLT and to ORCOFF for offline operation through the Object-Relational database access in POOL (POOL-ORA) [6]. The essence of the object relational approach of POOL is that the transient object model drives the DB data model. The designers of the offline data model are unaware of the tabular representation of the data in the DB. The DB schema is automatically created from the definitions of the persistent-capable objects. The data are retrieved from the underlying relational DB and then materialised as C++ objects in memory by following the SEAL [6] dictionary information.

The data processing modules have access to all non-event data via the "EventSetup" data structure [7]. The EventSetup is made available to modules on an event by event basis. A new set of non-event data is automatically read from ORCON or ORCOFF and delivered to the module that requested it whenever the "interval of validity" of the currently used data expires. The interval of validity of a given set of non-event data is defined as the set of consecutive events to which it applies. For pure offline data analysis, the intervals of validities of non-event data are set by experts prior to the availability of a given dataset. For the HLT a single set of non-event data is used throughout a run. If the run conditions change to the point of compromising the quality of the data being taken, then the run will be stopped and a new one will be started with a new set of non-event data corresponding to the new conditions.

The OMDS has components specific to each sub-detector. This DB contains the data used to configure the readout electronics, which typically consists of pedestal values and a number of parameters needed to optimally tune the electronics through which signals are readout and processed. OMDS will also host the conditions data produced during data taking by the environmental and power supply monitoring sensors of each sub-detector.

Some of the configuration and conditions data in OMDS may be needed by the CMSSW data processing modules. To satisfy this requirement data transfer procedures from OMDS to ORCON and ORCOFF have been put in place. The technology used to transfer data from OMDS to ORCON consists of SQL scripts that produce materialized views of the OMDS tables. In this step some data transformation may turn out to be necessary in order to present the non-event data to the reconstruction modules in the most optimal format. Any transfer between ORCON and ORCOFF takes place as direct ORACLE streaming.

Calibration data computed offline, like the alignment constants, has to be stored in both ORCON and ORCOFF.

In this case the data flows from ORCOFF to ORCON again through direct ORACLE streaming.

Data Acquisition and Monitoring

Events selected by the L1 Trigger are distributed by the DAQ event builder to a number of FU nodes of the FF. It is in these nodes that the full event data is inspected for the first time. Each of the FU nodes runs a CMSSW application. The configuration of the CMSSW application comprises the execution of one or more sequences of modules. One such sequence normally consists of a number of reconstruction modules followed by a filter module. Additionally, the FU CMSSW application can be configured to compute also monitorable elements (numbers, histograms, etc.). In this case additional sequences of modules are scheduled for execution. Monitoring sequences consist of reconstruction modules followed by analyzers modules. The latter make use of the monitoring framework of CMSSW [8] to book and fill monitorable elements. The monitorable elements are transferred to special “Monitorable Consumers” applications. Collection of monitorables from multiple sources, like the FUs of the FF, is also supported by the monitoring framework. The Monitorable Consumers receive updated monitorables at a configurable frequency and their execution can be controlled interactively.

The job configuration is always automatically pre-processed by a scheduler application to guarantee that each module is executed only once, even if present in more than a sequence.

The FF is organized in sub-farms (see Figure 2). Events selected by the FUs of a sub-farm are forwarded to a special application, called Storage Manager, which decouples the FUs from the computing services responsible for permanent data storage or data delivery to remote GRID sites.

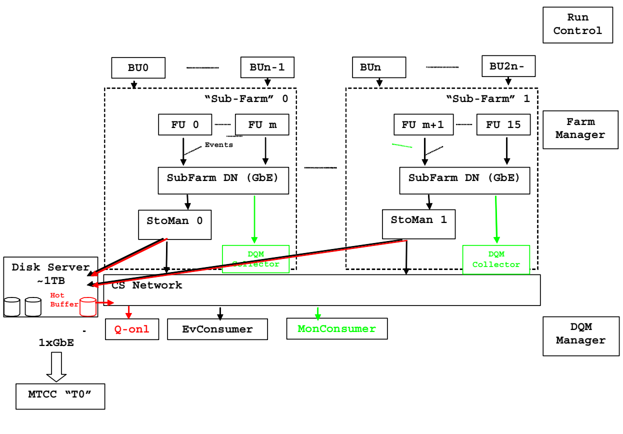

Figure 2: CMS Filter Farm architecture

The Storage Manager can also forward full events to special “Event Consumer” applications for monitoring purposes. Event Consumers are again CMSSW applications that make use of the monitoring software framework to book, fill and analyze monitorable elements. Event Consumers can also be configured to access selected events on disk. In this case they are pure offline monitoring applications.

SOFTWARE OPERATION AT THE MTCC

Event Topology and Rates

A dedicated simulation tool has been used to evaluate the event topology and rates at the MTCC. The predicted rate of cosmic muons traversing the DT chambers is 2.3 kHz. Only about 10 Hz are expected to traverse also at least one SST module. As discussed previously, the maximum event on disk rate at the MTCC will be of about 1 kHz. However, with the exception of limited data taking periods, it is expected that the average rate on disk will be at least one order of magnitude lower than this maximum value.

Online Event Selection

Given the numbers presented in the previous subsection, special attention is put on ways of guaranteeing that the tiny fraction of “SST events” (defined as the L1 triggerable events in which at least one SST module is traversed) are selected and stored. One solution would be to use dedicated L1 Trigger signals. These signals can be obtained by imposing angular requirements on the muon segments reconstructed by the muon trigger electronics. These methods are not optimal for two reasons. First, the selection efficiency and purity provided by these triggers are limited both by the instrumental precision with which the area covered by the SST can be illuminated and by the interactions undergone by the muons between the muon chambers and the SST (multiple scattering, magnetic field deflection, etc.). The second reason is related to the difficulty in operating these exclusive triggers in parallel to other more inclusive ones, which are more interesting to the Muon and Calorimeter sub-detectors.

A more powerful method consists of allowing the maximum possible L1 Trigger rate and performing the final event selection in the HLT. “SST” events are expected to be easily identified with a simple and robust ADC counts method, possibly limited to channels associated to reconstructed clusters. This method requires the SST local reconstruction, in addition to a dedicated filter module, to run online. A study aiming at optimizing the selection in terms of efficiency and purity achievable is ongoing. Additional muon and calorimetric streams could be produced in a similar way. Muon streams could also be produced on the basis of the L1 Trigger data, which is also available to the offline software in the HLT.

Non-Event data handling

Prior to physics runs, all sub-detectors must be optimally tuned and synchronized. This operation is normally carried out with custom XDAQ [9] applications. A special case is represented by the SST sub-detector that has decided to implement all its commissioning software in CMSSW. All details can be found in [10].

The last step of the commissioning procedures is normally represented by the transfer of the OMDS data needed by the offline software to ORCON and ORCOFF. The ECAL sub-detector has developed an interesting solution for computing and transferring new sets of non-event data. For instance, the computation of a new set of pedestals and its transfer to OMDS can automatically occur during data taking whenever the monitoring system detects a significant variation of the pedestals.

Monitoring and Visualisation

All sub-detectors are expected to use their CMSSW based monitoring suite at the MTCC. Data at different reconstruction levels will be monitored. Accordingly, the proper reconstruction modules will have to be executed online.

Track reconstruction is not expected to be fully validated for the MTCC. For this reason, the quality of the data taking will have to be assured with just the lower level Muon and Tracker reconstruction based monitoring.

The general visualisation tool IGUANACMS [11] is available and functional in CMSSW and will be used at the MTCC. IGUANACMS can be used as an event display where various reconstructed objects can be visualised in different views. A screenshot of an HCAL beam test muon event is shown in Figure 3. IGUANACMS has proved to be extremely useful also as a debugging tool.

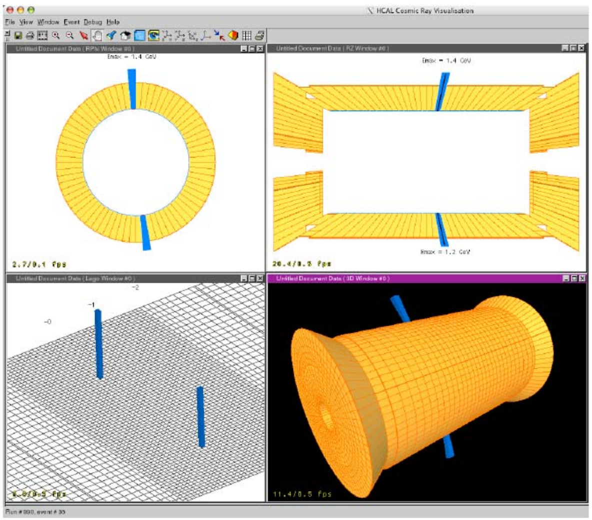

Figure 3: Screenshot of the IGUANACMS visualisation of an HCAL beam test event

Another lightweight visualisation tool that will be used by the SST community is documented in [12].

Offline data analysis

One of the main challenges of the MTCC offline data analysis is represented by the validation and cross check of the measurements provided by the muon hardware alignment system. Track reconstruction linking the muon and SST detectors will have to be carried out not only to demonstrate the improvement obtained with the alignment constants provided by the laser system, but also to compute an independent set of such constants for validation purposes. A specialized software tool has been developed to compute and store in ORCOFF and ORCON the alignment constants for the CMS geometry starting from the laser alignment system measurements. The entire EventSetup software infrastructure necessary to produce an aligned geometry out of these constants and present it to the reconstruction software is in development and is expected to be in place by the start of the MTCC.

CONCLUSIONS

After less than one year and a half of development, the new CMS offline software is ready to be tested at the MTCC. It will be used for commissioning, online event selection, monitoring, visualization and offline analysis.

REFERENCES

[1] CMS Collaboration, Technical Proposal, December 1994, CERN/LHCC 94-038.

[2] http://lhc.web.cern.ch/lhc/LHC

[3] CMS Collaboration, DAQ Technical Design Report, December 2002, CERN/LHCC 2002-26.

[4] CMS Collaboration, Muon Technical Design Report, December 1997, CERN/LHCC 97-32.

[5] C. Jones et al., "The new CMS Event Data model", to be published in the Proceedings of CHEP06, Mumbai, India, Feb. 2006.

[6] http://lcgapp.cern.ch/project/persist/

[7] C. Jones, "Access to Non-Event data for CMS", to be published in the Proceedings of CHEP06, Mumbai, India, Feb. 2006.

[8] C. Leonidopoulos, "Physics and Data Quality Monitoring at CMS", to be published in the Proceedings of CHEP06, Mumbai, India, Feb. 2006.

[9] L. Berti et al., "Using XDAQ in application scenarios of the CMS experiment", in the Proceedings of CHEP 03, La Jolla, California, Mar 2003, hep-ex/0305076. http://sourceforge.net/projects/xdaq/

[10] R. Bainbridge, "Commissioning Procedures and Software for the CMS Silicon Strip Tracker", to be published in the Proceedings of CHEP06, Mumbai, India, Feb. 2006.

[11] G. Alverson et al., "IGUANA Architecture, Framework and Toolkit for Interactive Graphics", in the Proceedings of CHEP03, La Jolla, California, Mar 2003, cs.SE/0306042.

[12] M.S. Mennea and G.Zito, "Use of Interactive SVG Maps and Web Services to Monitor CMS Tracker, Proceedings of the IEEE-NPSS Real Time Conference 2005, Stockolm, Sweden, June 2005.