Development of An On-Line Data Quality Monitor For The Relativistic Heavy-Ion Experiment ALICE

Özgür Çobanoğlu*
*University of Turin & INFN of Turin
Physics Department
Turin, Italy
Email: Ozgur.Cobanoglu@cern.ch
cobanogl@to.infn.it

Pierre Vande Vyvre†
†CERN
PH
Geneva, Switzerland
Email: Pierre.Vande.Vyvre@cern.ch

Ferhat Özok‡
‡İstanbul Üniversitesi
Fizik Bölümü
Istanbul, Turkey
Email: Ferhat.Ozok@cern.ch

for the ALICE collaboration

Abstract— **The on-line data monitoring tool developed for the coming ALICE experiment at LHC, CERN is presented. This monitoring tool which is a part of the ALICE-DAQ software framework, written entirely in C++ language, uses standard Linux tools in conjunction with the data display and analysis package ROOT, developed at CERN. It allows checking the consistency and quality of the data and correct functioning of the various sub-detectors either at run time or during off line by playing back the recorded raw data. After discussing the functionality and performance of this package, the experience gained during the test beam periods is also summarized.**

I. INTRODUCTION

The aim of high-energy heavy-ion physics is the study of strongly interacting matter at extreme energy densities (QCD thermodynamics). Statistical quantum chromo dynamics (QCD) predicts that, at sufficiently high density, there will be a transition from hadronic matter to a plasma of deconfined quarks and gluons -a transition which in the early universe took place in the inverse direction some 10^{-5} seconds after the big bang and which might still play a role today in the core of collapsing neutron stars [1].

ALICE is one of the four experimental set-ups currently built for the Large Hadron Collider (LHC) - the largest accelerator all over the world. ALICE was first proposed as a central detector in 1993 and complemented by an additional forward muon spectrometer designed in 1995. Its goals are to measure flavour content and phase-space distribution, event by event, for a large number of particles whose momenta and masses are of the order of typical energy scale involved (temperature ≃200 MeV). The experiment is designed to cope with the highest particle multiplicities anticipated for Pb-Pb reactions, that is, up to 8000 charged particles per unit of rapidity ($_{ch}/ \quad = 8000$) [1].

The high data volume (of the order of one peta byte per year) and the sophisticated trigger system requires an on-line data quality monitor in order to have immediate information about the ongoing data acquisition as well as for the detector commissioning period [2]. The ALICE MOOD (Monitor Of On-line Data and Detector Debugger) [3] developed for this purpose is a part of ALICE Data Acquisition and Test Environment (DATE) framework and is based on the ROOT [4] libraries developed at CERN.

II. IMPLEMENTATION

MOOD is written in C++ based on the ROOT framework and uses the DATE monitoring library. It is a high level user friendly package with pop-up menus which also allows low level execution of generic stream handling functions by its command line utility mooDump [3][6].

The tool will be used to monitor, at run time, sub-events and complete events. The access to on-line/off-line data is achieved by DATE [7][5] monitoring library functions. The graphical user interface (GUI) and some of the fast analysis functions [8][9] are based on the ROOT libraries.

The system parameters which will then define the monitoring policy and the visualization algorithm, that is, the character of MOOD as a monitoring client can be set via GUI or they can be loaded from file. The parameters consist of two components; the first component is used while calling the monitoring functions to communicate with the DATE system and the second one is used for defining the visualization algorithm.

III. FUNCTIONALITY

Detector development teams concentrate on the quality of the data stream created by their particular detector, from single channels up to more global views. Faulty and misconfiguration conditions must be discovered as quickly as possible to recover the quality of the data. For that purpose, MOOD can run at different levels of the DAQ architecture shown in Fig. 1. It can monitor subevents in the LDC (Local Data Concentrator), complete events in the GDC (Global Data Collector) and entire raw data files in the TDS (Transient Data Storage). The tool can perform on-line and off-line monitoring session locally and remotely.

The TPC and the ITS-SDD detectors are used in this paper as example on how the raw data are processed and the

0-7803-9183-7/05/$20.00 ©2005 IEEE.

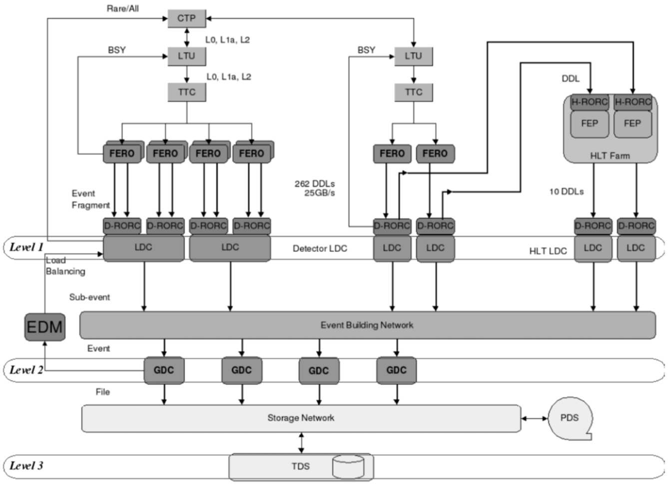

Fig. 1. The simplified scheme of ALICE DAQ system.

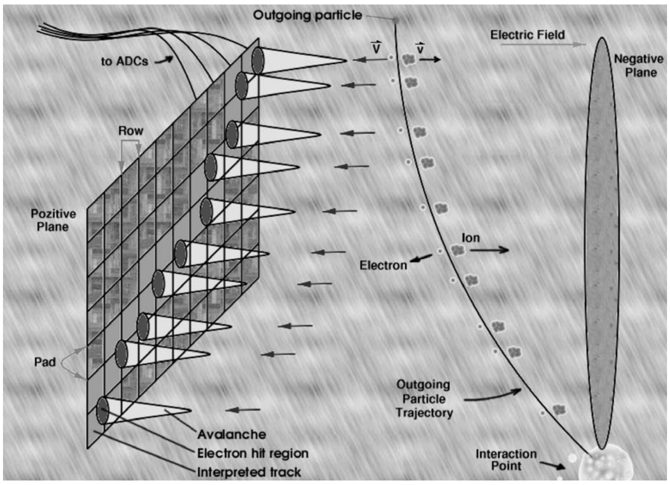

Fig. 2. Generic tracking principle.

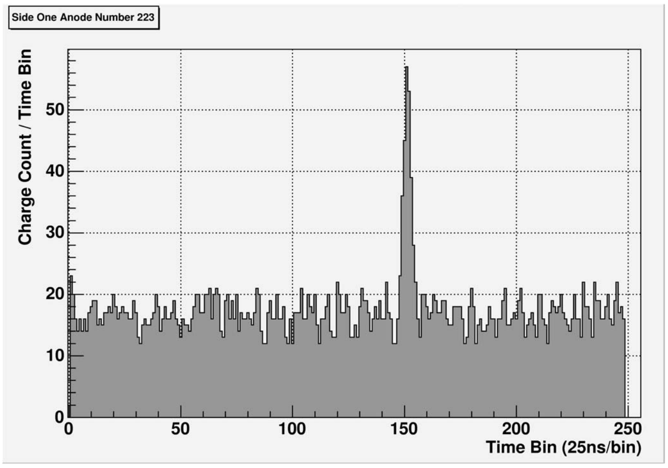

Fig. 3. An anode projection histogram of ITS-SDD.

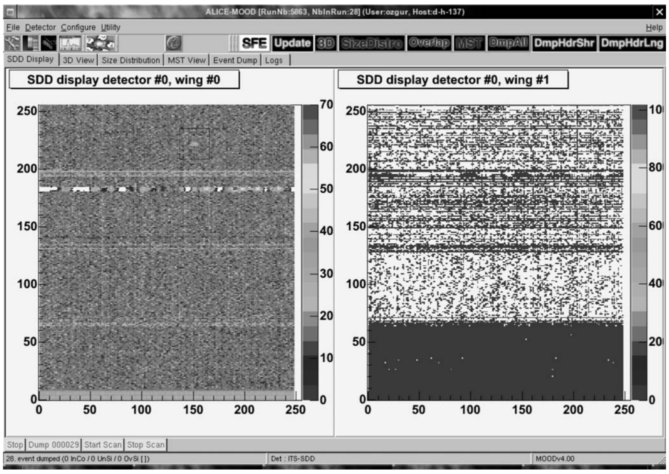
ALICE-MOOD [RunNb:5863, NbInRun:28] (User:ozgur, Host:d-h-137)

28. event dumped (0 InCo / 0 UnSi / 0 OvSi [ ])

Det : ITS-SDD

MOODv4.00

Fig. 4. Two dimensional ITS-SDD event.

physics information is visualized. Generic tracking principle of a device is shown in Fig. 2.

The ions created by the outgoing particle in the tracker can not recombine due to the existence of an electric field strong enough to separate the electron from its ion. Electrons then begin to drift along the volume till they reache the positive pad plane. The drift time of an electron is proportional to the distance to the pad plane. Ionizations are visible as peaks in the time proportional histograms which are created by the ADCs that sample the related pads. The distance of these peaks are used as one of the dimensions (e.g. z axis for TPC and x axis for ITS-SDD) while displaying the event. The cones indicate secondary ionizations (in case such an ionization occurs) due to the acceleration of the electrons created during primary ionizations and the ellipses indicate the real hits. The flat coloured region on the pad plane (this would be one dimensional anode array for ITS-SDD) is the interpreted result of the real trajectory according to the algorithm MOOD uses.

The sample number that holds the maximum ADC count (z value) is taken into account as the 'hit' location as shown in Fig. 3, although, a cluster finding analysis should be performed to be able to calculate the precise hit positions. The pad and row numbers (anode number for ITS-SDD) are treated as x-y plane coordinates (y ordinate for ITS-SDD) and with the z value, a 3D view of the event which is being processed can be rebuilt. In Fig. 4, one hit which has the dimensions of several anodes and several time bins is visible on the left top side of the first wing. The x and y axis represent the time and anode number respectively. Color code shows the collected charge.

MOOD has three types of displays on which the raw data are visualized. The first type is based on viewing every single datum coming from each of the sub-detector systems without performing any analysis algorithms (e.g. Fig. 3). As shown in Fig. 5, a time projection histogram of an individual TPC pad always shows uninterpreted raw data.

The second type of display consists of a low level reconstructed event on which cluster analysis is not applied. This type of a display in MOOD aims showing the general status of what is going on inside the detector. As shown on the upper most subframes in Fig. 6, a 3D display of a TPC event contains always interpreted data. In time projection histograms, the time

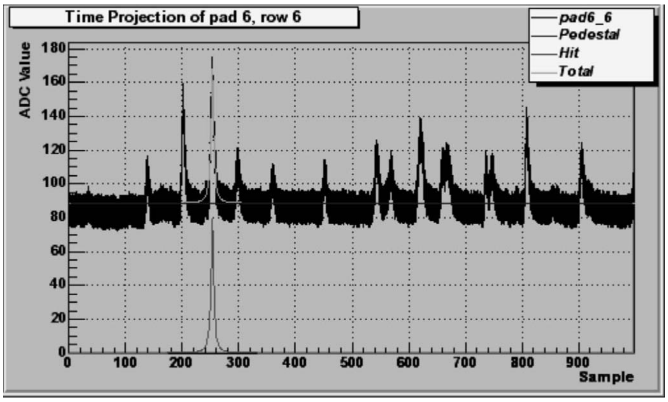

Fig. 5. A time projection histogram of an individual TPC pad.

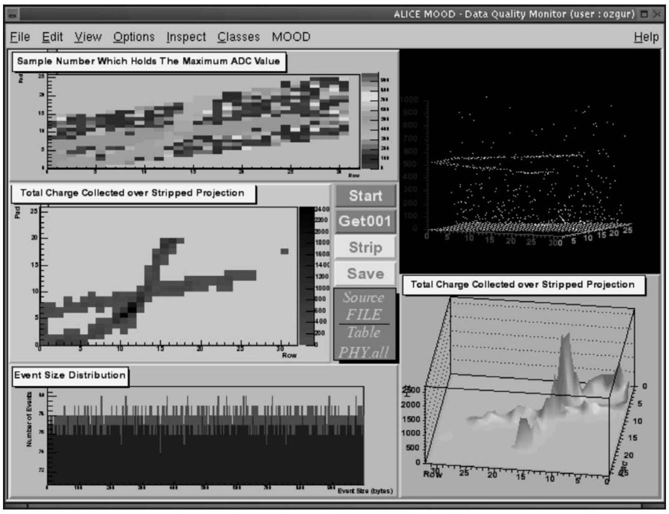

Fig. 6. Compact TPC display.

bin which holds the maximum ADC count can be taken into account as a hit and this value is used to create the 3D view. Another approach MOOD uses for fast reconstruction of the TPC tracks is subtracting the average of the histogram from the maximum ADC value for each time projection histogram. The resulting value is taken into account as the deposited energy.

These approximations are particularly suitable for test beams as the events usually contain only a few tracks. Since the cluster finding is not applied, the hits and trajectories are not precise but the calculation is fast. The thickness of the trajectories (and the diameter of the hit point on an ITS-SDD layer as shown in Fig. 4 are directly related to the track resolution of the detector device. The occupancy within TPC can exceed 50% of the total volume with Pb-Pb collusions at the LHC and for all the sub detectors of ALICE, the hit and track resolutions are two of the challenging components of the experiment.

The third type of display concerns the format and size of the data. The size distribution of ALICE events is one of the most important components that must be monitored since it is related to many of the system parameters as a whole.

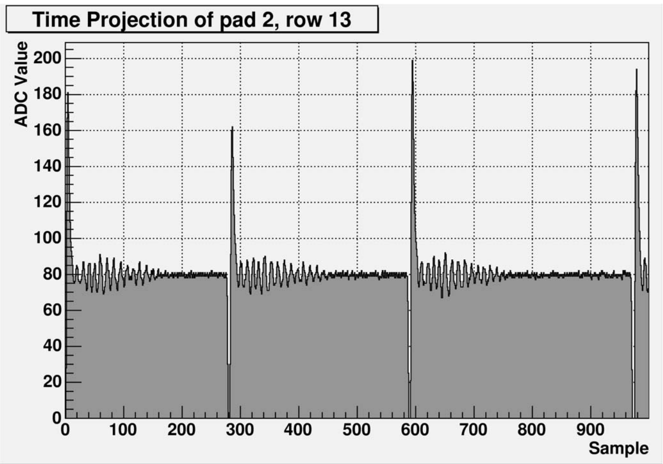

Fig. 7. A time projection histogram of a pad of the ALICE TPC.

MOOD is able to collect statistics for various types of event size distributions as shown on the left bottom subframe in Fig. 6 (e.g. for specific hosts and/or trigger types).

The tool also performs consistency checks on the raw data format both concerning the stream format of the DATE system and the internal format of the payload produced by a detector. MOOD is able to dump event headers and payloads in several different formats in a bit-by-bit basis to ease the diagnosis procedures. It is also able to create configuration files required by the FEEs of the detector systems to be used for fine tuning the detector system itself recursively.

IV. CONSIDERING THE DETAILS

A. Auxiliary Windowlets

The tool has auxiliary monitoring windows which can be used for detecting malfunctions that are not visible on the plots located on main window. Fig. 7 shows a time projection histogram of a TPC pad that would lead to a misleading colour code creation. Such a time projection histogram would always produce a colour code regardless of whether it contains real ionization locations (hits). Since as shown in Fig. 5, MOOD selects the highest peak to calculate the colour code or z value in both algorithms.

There may be some hardware problems that can not be seen on main window but on auxiliary ones. A time projection histogram, in which some of the samples of the pad's ADC outputs the same value regardless of the real situation, can also be produced as a result of hardware malfunction. When the pedestal is not approximately constant, although there is not even a peak, a misleading colour code can be produced in this case as well. Because the pedestal subtraction function assumes that pedestal is the average of corresponding histogram plus a constant that is usually less than a minimum value. MOOD calculates the pedestals at run time and does not load them from a file; this slightly slows down the process but makes the tool able to adapt itself to changes in pedestal levels due to changes in run time conditions.

B. Command Line Utility mooDump

The package has a command line utility named mooDump [6][10] for checking on-line/off-line DATE streams in a lower level. mooDump can search for a specific pattern within the DATE stream and outputs statistics related to the monitoring session. The tool can extract events that fit into a specific requirement and isolate them as binary DATE event files for further diagnosis.

mooDump runs in batch mode and does not require a GUI. Its installation and maintenance is easy as it does not depend on ROOT but only DATE monitoring library. mooDump is aware of the detectors. It is able to perform detector specific monitoring sessions, loading the external mapping files where necessary and outputting 2D or 3D histogram files in case it is required.

V. CONCLUSION

The tool is currently able to produce statistical distributions (e.g. event size distributions) and to visualize the TPC, TRD, ITS-SDD, ITS-SPD, ITS-SSD, HMPID, Muon Tracker and MST test data taken during test beam studies performed at CERN.

No cluster finding algorithms is implemented, since it requires much CPU time. Visualizing event without cluster analysis is treated better since it enables consideration of the pads that would make contribution while finding the precise hit locations in 3D during cluster analysis.

A MOOD plugger can host one detector object at a time; so, in case two detector specific monitoring sessions are to be created on the same host machine, two instances of MOOD must be invoked. The cost of adding a new detector object in terms of memory usage is as shown in Table. I.

MOOD is tested on SLC3, the default OS used at CERN for DAQ activities as well as other major distributions of Linux OS. The tool is currently supported on AMD and Intel 32-bit machines.

DATE monitoring library provides the functionality to stop the data flow when the related physical buffers are full. This feature enables MOOD and mooDump to monitor all the events within the stream in the cost of slowing down the data flow. Even though, the environment which MOOD runs within is a real time system, MOOD itself is not intended to be a real time data quality monitoring tool. As a result, monitoring speed depends on the monitoring policy (e.g. all, calibration, physics events, etc) and monitoring type (e.g. remote, local, short or long path within the network).

The final goal for MOOD is to be able to monitor all the ALICE detectors. Not only the tool should be able to monitor the data that comes from every single sample of all the sub detector systems but it should also be capable of making a low level reconstruction in order to view the super event in a convenient form.

TABLE I

MEMORY CONSUMPTION OF DETECTOR MODULES

Detector Module	Description	Memory (Mb)
MUON-TRK	Muon Tracker	64
TPC	Time Projection Chamber	34
TPCc	Compact TPC	31
ITS-SPD	Silicon Pixel Detector	25
ITS-SSD	Silicon Strip Detector	24
ITS-SDD	Silicon Drift Detector	23
HMPID	High Momentum PID	23
TRD	Transition Radiation Detector	23
DATE Low Level	No detector dependence	22
Plugger alone	Detector object carrier	20

The TPC detector system alone will consist of more than 57×10^4 individual pads [1]. This is also the number of the time projection histograms which are to be used during visualizing the TPC part of the event in 3D. Considering the number of samples within these histograms, it is obvious that some automatic malfunctioning component search algorithms could be applied.

Integration to AliROOT [11], the off-line analysis framework of ALICE, is foreseen. This will provide a code base which is much more spread over the off-line developers and users so that the maintenance and particularly the further enhancements for the new detector systems (e.g. PHOS, FMD and TOF) within MOOD is expected to be easier.

ACKNOWLEDGMENT

Ö. Çobanoğlu would like to thank Gökhan Ünel for his continuous personal and professional support and encouragement.

REFERENCES

[1] ALICE Collaboration, 1995, *ALICE Technical Proposal for A Large Ion Collider Experiment at the CERN LHC*, CERN/LHCC/95-71, LHCC/P3, ISBN 92-9083-077-8

[2] ALICE Collaboration, 2004, *ALICE Technical Design Report of Trigger, Data Acquisition, High Level Trigger and Control*, CERN-LHC-2003-062, ALICE TDR 10, ISBN92-9083-217-7

[3] Ö. Çobanoğlu, 2003, *Development of an On-Line Data Quality Monitor For The Relativistic Heavy-Ion Experiment ALICE*, Master Thesis, İstanbul Üniversitesi, Fen Fakültesi, Fizik Bölümü, İstanbul

[4] Brun R., Rademakers F. et al., 1997, *ROOT - An Object Oriented Data Analysis Framework*, Nuclear Instruments and Methods in Physics Research, A 389, 81-86

[5] http://aldwww.cern.ch, *ALICE-DAQ web site*

[6] ALICE DAQ Group, *DATE 5 User's Manuel*, Being Prepared, CERN

[7] ALICE DAQ Group, *DATE 4 User's Manuel*, Internal Note/DAQ ALICE-INT-2002-036, CERN

[8] Brun R., Rademakers F., Panacek S., Buskulic D., Adamczewski J., Hemberger M., West N. et al., 2002, *ROOT Users Guide 3.02c*, CERN

[9] http://www.root.cern.ch, *ROOT web site*

[10] http://aldwww.cern.ch/MOOD, *MOOD web site*

[11] http://aliweb.cern.ch/offline/, *AliROOT web site*