The Detector Algorithms framework

ALICE detector

ALICE is the heavy-ion detector designed to study physics of strongly interacting matter and the quark-gluon plasma at the at the CERN Large Hadron Collider.

The detector includes high resolution tracking (silicon detectors, large time-projection chamber), particle identification, and triggering elements. It features two large magnets, a main solenoid and a dipole on the Muon arm.

It primarily targets heavy-ion lead-lead collisions, but it also has a substantial physics program with proton-proton and proton-ion collisions.

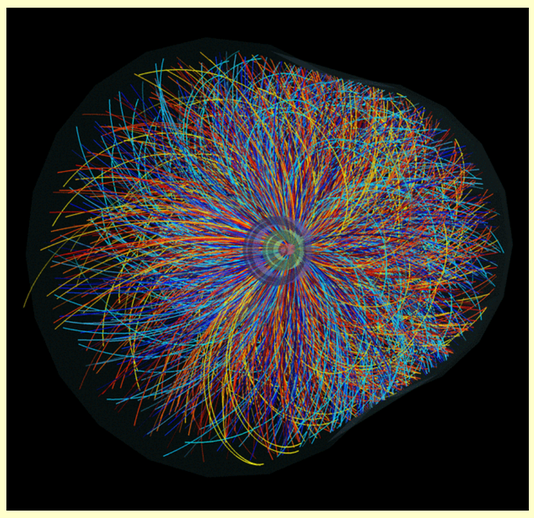

An example of lead ion collision as seen by ALICE

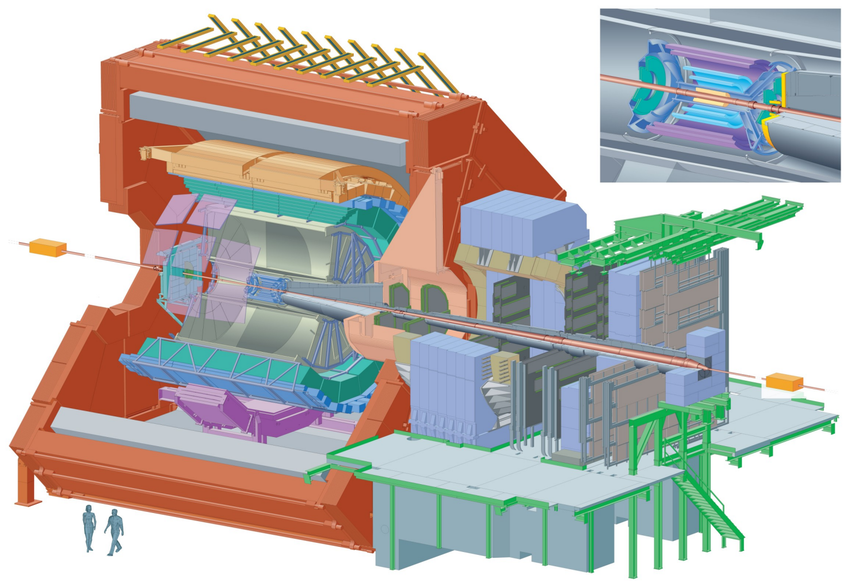

The ALICE detector

Calibration procedures

The 18 ALICE sub-detectors require specific calibration tasks to be performed regularly in order to achieve the most accurate physics measurements. The corresponding set of procedures involves events analysis in a wide range of experimental conditions.

A dedicated framework has been designed and implemented to achieve as much as possible the detector calibration directly online in the DAQ. These calibration tasks may be done either in dedicated runs, or in parallel to physics data taking.

ALICE data-acquisition system (DAQ)

The DAQ handles the data flow from the sub-detector electronics to the archiving on tape.

It is capable of sustained recording rates of more than 3GB/s.

In 2011, ALICE DAQ recorded 2.5PB of data to permanent storage.

DA architecture

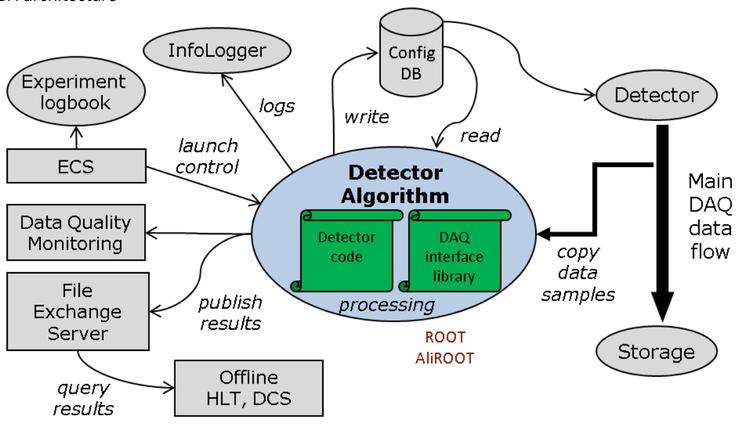

Detector algorithms (DAs)

DAs are provided by the sub-detector teams, using the global framework to develop detector-specific calibration procedures.

Each DA grabs detector data (physics or calibration events) and produces results online. These results can be reused directly online (e.g. to configure the detector electronics or give quality feedback to the Data Quality Monitoring system), or shipped offline (to be post-processed and used in event reconstruction).

A DA consists of some specific detector code (to analyze events and produce results according to a given calibration task), using a support library to interact with the external components (read configuration, grab events, log messages, export results, deal with the control system commands).

Software life cycle

Deployment in production is done after a strict release procedure, including building in a standardized environment and testing on reference data sets.

50 different DAs regularly used online.

328 DA packages upgrades between 2007 and 2012.

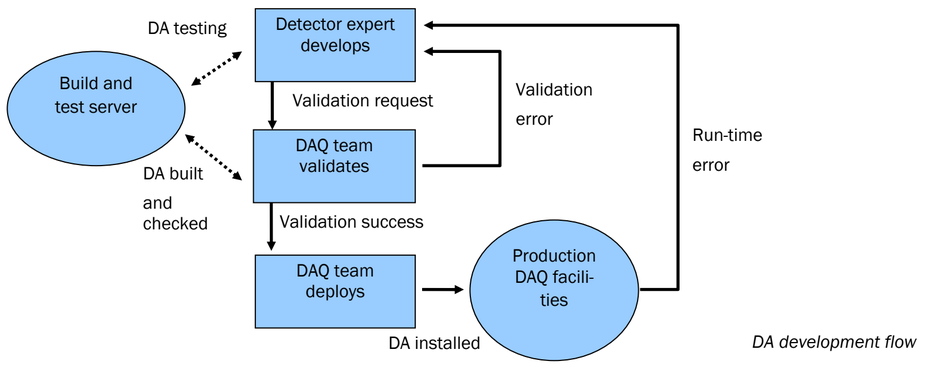

DA development flow

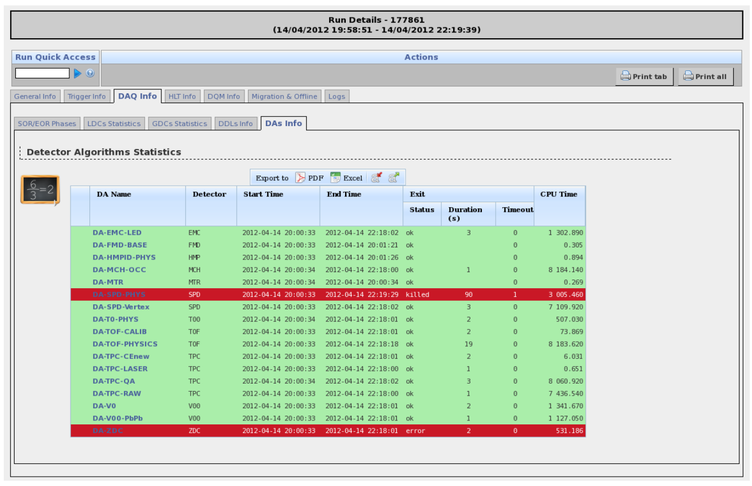

Run Details - 177861
(14/04/2012 19:58:51 - 14/04/2012 22:19:39)

Run Quick Access | Actions | Print tab | Print all

General Info | Trigger Info | DAQ Info | HLT Info | DQM Info | Migration & Offline | Logs

SOR/EOR Phases | LDCs Statistics | GDCs Statistics | DDLs Info | DAs Info

Detector Algorithms Statistics

Export to PDF Excel

DA Name	Detector	Start Time	End Time	Exit Status	Exit Duration (s)	Exit Timeout	GPU Time
DA-EMC-LED	EMC	2012-04-14 20:00:33	2012-04-14 22:18:02	ok	3	0	1 302.890
DA-FMD-BASE	FMD	2012-04-14 20:00:33	2012-04-14 20:01:21	ok		0	0.305
DA-HMPID-PHYS	HMP	2012-04-14 20:00:33	2012-04-14 20:01:26	ok		0	0.894
DA-MCH-OCC	MCH	2012-04-14 20:00:34	2012-04-14 22:18:00	ok	1	0	8 194.140
DA-MTR	MTR	2012-04-14 20:00:34	2012-04-14 20:00:34	ok		0	0.269
DA-SPD-PHYS	SPD	2012-04-14 20:00:33	2012-04-14 22:19:29	killed	90	1	3 005.460
DA-SPD-Vertex	SPD	2012-04-14 20:00:33	2012-04-14 22:18:02	ok	3	0	7 109.920
DA-T0-PHYS	T00	2012-04-14 20:00:34	2012-04-14 22:18:01	ok	2	0	507.030
DA-TOF-CALIB	TOF	2012-04-14 20:00:33	2012-04-14 22:18:01	ok	2	0	73.869
DA-TOF-PHYSICS	TOF	2012-04-14 20:00:33	2012-04-14 22:18:18	ok	19	0	8 183.620
DA-TPC-CEnew	TPC	2012-04-14 20:00:33	2012-04-14 22:18:01	ok	2	0	6.031
DA-TPC-LASER	TPC	2012-04-14 20:00:33	2012-04-14 22:18:00	ok	1	0	0.651
DA-TPC-QA	TPC	2012-04-14 20:00:34	2012-04-14 22:18:02	ok	3	0	8 060.920
DA-TPC-RAW	TPC	2012-04-14 20:00:33	2012-04-14 22:18:00	ok	1	0	7 436.540
DA-V0	V00	2012-04-14 20:00:33	2012-04-14 22:18:01	ok	2	0	1 341.670
DA-V00-PbPb	V00	2012-04-14 20:00:33	2012-04-14 22:18:01	ok	1	0	1 127.050
DA-ZDC	ZDC	2012-04-14 20:00:33	2012-04-14 22:18:01	error	2	0	531.186

Example view of the ALICE e-logbook DA statistics page for a run.

Control and bookkeeping

The "DA launcher" starts each DA and follows their execution. It reports to the DAQ runControl the DA status (running, error, exited) and executes commands (stop or abort). It enforces configured resource limitations (e.g. memory used), and if necessary kills the DA process if unresponsive.

The launcher also integrates bookkeeping features. It redirects the DA output messages to the central DAQ logging system, and collects run-time statistics (exit status, execution time, etc) for the experiment e-logbook. These records are remotely accessible by detector experts through a Web interface.

Parallelization

The DA framework provides a set of C++ classes to ease the development of multi-threaded DAs, and be able to use the full power of multi-core CPUs. It involves independent processing of events in parallel threads and aggregation of partial results.

Detector specific code is effectively implemented by filling a few virtual methods (e,g. processEvent()). The control and data flow between the runtime threads are then transparent to user.

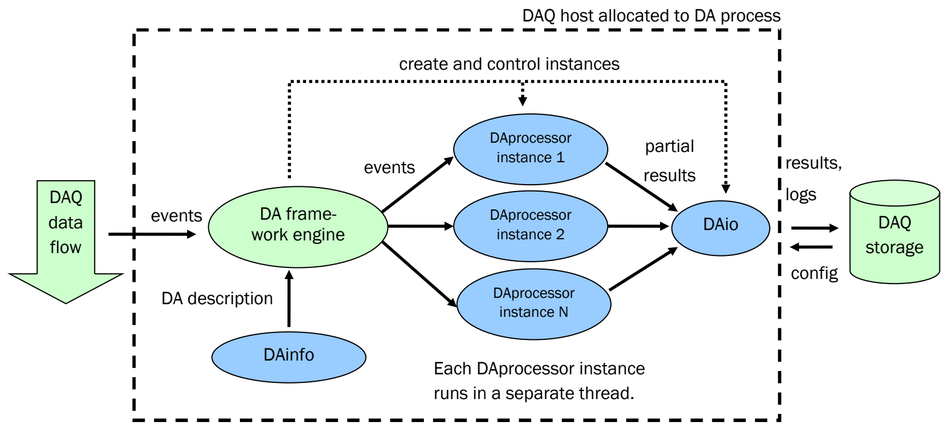

Contact: sylvain.chapeland@cern.ch

Conference on Computing in High Energy Physics, May 2012, New York, USA