12 Experiment Control

12.1 Introduction

The overall control of the ATLAS experiment includes the monitoring and control of the operational parameters of the detector and of the experiment infrastructure, as well as the supervision of all processes involved in the event readout. This functionality is provided by two independent although complementary and interacting systems: the TDAQ control and the Detector Control System. The TDAQ Control is in charge of controlling the hardware and software elements in TDAQ needed for data taking. The DCS handles the control of the detector equipment and related infrastructure. The architecture of the Experiment Control has already been discussed in Chapter 5-3. The DCS is based on a SCADA system PVSS-II [12-1], whereas the TDAQ control is based on the TDAQ Online Software described in Chapter 10. These systems perform different tasks and have different requirements. Whilst the TDAQ control is only required when taking data, the DCS has to operate continuously to ensure the safe operation of the detector. The operation of the detector requires a strong coordination of these two systems with the LHC machine. The interaction with the LHC machine will be handled by the DCS as illustrated in Figure 5.5 and presented in detail in Chapter 11. The TDAQ system has the overall mastership for the control of the data-taking operations.

The general control of the experiment requires a flexible partitioning concept as it is described in Chapter 3.3, which allows for the operation of the sub-detectors in stand-alone mode, as required for calibration or debugging, as well as for the integrated operation for concurrent data taking. The overall control strategy and the control operations of the various systems are described in this chapter. Furthermore, the required coordination of the various systems involved in the scenarios for physics data-taking and calibration modes, is discussed.

12.2 Detector control

The DCS system provides the flexibility to map the partitioning concept of Atlas. The finest granularity of the TDAQ system is given by the segmentation of the sub-detectors in TTC zones. For these reasons, the different sections of the sub-detectors will be logically represented in the back-end software of the DCS by means of the so-called control units, which will be operated as a Finite State Machine (FSM). According to this model, the DCS of the Tilecal, for example, may be organized in four independent control units, which can control the four sub-detector sections. Each control unit is characterized by its state. The control units are hierarchically organized in a tree-like structure to reproduce the organization of the experiment in sub-detectors, sub-systems, etc. as illustrated in Figure 12-1. The units may control a sub-tree consisting of other control units or device units, which are responsible for the direct monitoring and control of the equipment. Each control unit has the capability to exchange information or pass commands to other control units in the hierarchy. The flow of commands and information will only be vertical. Commands will flow downwards, whereas status and alarms will be transferred upwards in the hierarchy.

The control units will support different partitioning modes. Any control unit and therefore, the related sub-tree, may be excluded from the hierarchy and be operated in stand-alone mode for testing, calibrations or debugging of part of the system. In this case the detector can be operated

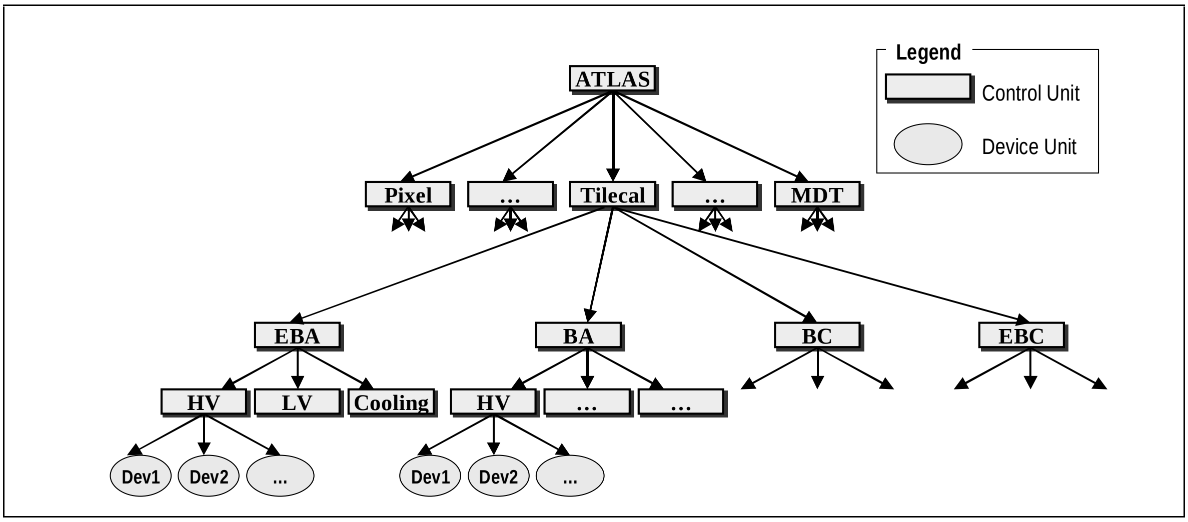

Figure 12-1 DCS Logical Architecture.

directly from the DCS graphical interface, whereas during physics data taking and calibration procedures, commands will be sent from the TDAQ control. Therefore, an ownership model, which avoids to issue conflicting commands must be provided. This mechanism will be developed according to the recommendations of the JCOP Architecture Working Group [12-2].

12.3 Online Software Control Concepts

The TDAQ system is composed of a large number of hardware and software components, which have to operate in a coordinated fashion to provide for the data-taking functionality of the overall system. The organisation of the ATLAS TDAQ system into detectors and sub-detectors leads to a hierarchical organisation of the control system. The basis of the TDAQ control is provided by the ATLAS Online Software, which is explained in detail in Section 10.5.

The basic element for the control and supervision is a controller. The TDAQ control system is built of a large number of controllers which are distributed in a hierarchical tree following the functional composition of the ATLAS TDAQ system.

This concept is illustrated in Figure 12-2. Four principle levels of control are shown. Additional levels can be added at any point in the hierarchy if needed. A top level controller named the *root controller* has the overall control over the TDAQ system. It supervises the next level of controllers in the hierarchy, the *sub-detector controllers*. It is the responsibility of the sub-detector controller to supervise the hardware and software components which belong to this sub-detector. The next control level takes the responsibility for the supervision of the sections which correspond to the TTC partitions [12-3]. The leaf controllers on the lowest level, the so-called *local controllers*, are responsible for the control of readout crates and alike. Farm supervision and ROS hardware make use of the same controllers following a similar structure, which is further discussed in Section 12.3.1 and Section 12.3.2.

A controller in the TDAQ system is characterised by its state given by the TDAQ state model described in Section 12.4.3. In any place of the hierarchy, a change of state is initiated and synchronized from the higher level controller and sent down to the next lower level. From there information is returned to the next higher level when the requested transition has been performed. Possible error conditions are also reported back to the higher level.

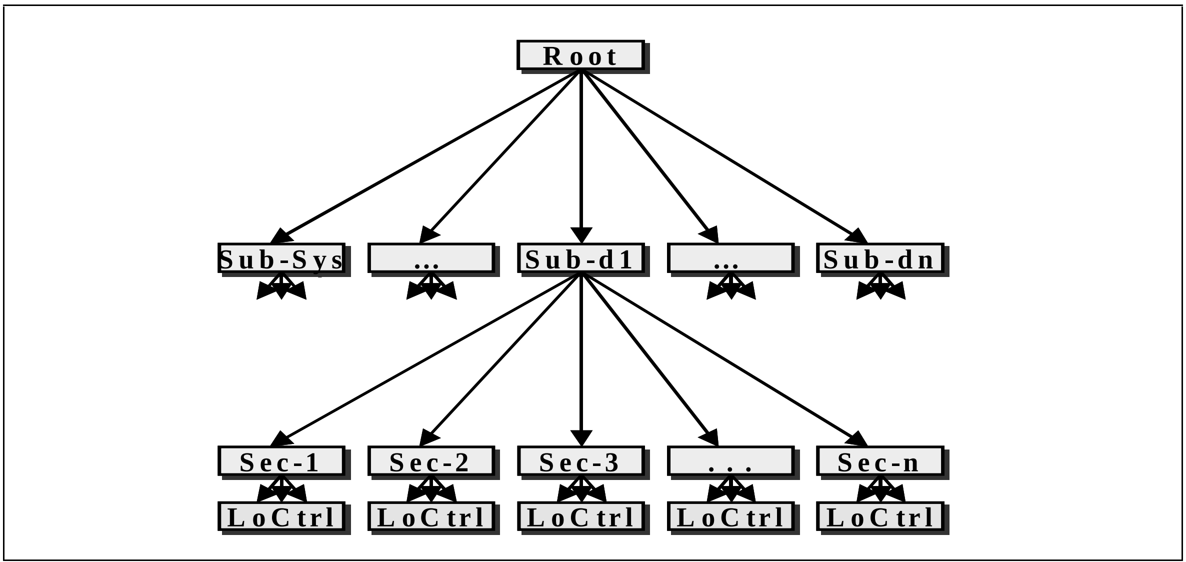

Figure 12-2 Online Software Control Hierarchy in TDAQ

A controller framework allows for the handling of the described operations in a coherent way on all the controllers in the system. It also gives the necessary flexibility to the detector expert to customize each controller for handling the individual tasks on the system under its control. These tasks take a wide range of variety from the readout of the hardware to event filter farm control. The information on the relationship of the controllers and their responsibilities for a given partition is detailed in the configuration database (Section 10.4.3).

Each controller is responsible for the initialisation and the shutdown of software and hardware components in its domain. It is also responsible for passing commands to child controllers and for signalling its overall state to its parent. Of particular importance is the synchronisation necessary to start the data-taking. This is performed by successive transitions through a number of intermediate states until data-taking is finally started as described below in Section 12.5.1. Interaction with the shift operator via the user interface drives the operations via commands to the root controller. The inverse series of synchronized transitions is traversed when data-taking is stopped.

During the operational phases, each controller is responsible for the supervision of the operation of elements under its direct control and for the observation of the operations of its children thus also providing the task of error handling. In the case of a malfunction of a detector, the controller can start corrective actions and/or signal the malfunction by sending messages. Severe malfunctions which are beyond the capabilities of a controller can be signalled by a state change to its parent. It is then the role of the parent controller to take further actions. The design of the control, supervision and error handling functionality is based on the adoption of a common expert system shell. Specific nodes will use different rules to perform their functions in addition to a common rule base which handles the generally valid aspects.

12.3.1 Control of the DataFlow

The DataFlow control encompasses the ROS/ROD control and the Data Collection control. It is comprised of the control of all applications and hardware modules responsible for moving the event data from the detector front-end electronics and LVL1 trigger to the high level triggers

(LVL2 and EF). It includes the control of the ROD crates, the RoI Builder, the ReadOut System and the Data Collection applications, such as the Event Builder.

There are two flavours of local controllers in the DataFlow foreseen, both making use of the Online software infrastructure in the same way. The ROS controller is tailored for the control of ROS software applications and hardware devices which cannot themselves access the online software facilities. The DC controller handles the different types of DC applications and is optimized for the control of computer farms. A version of the latter is also used for the control and supervision of the high level triggers and is further described in Section 12.3.2. The main difference between the two controllers is that the ROD crate controller controls a hardware device on which no standard software application is running and therefore it is the only access point to the databases as well as the only element communicating over IS / MRS to the Online system.

Both controllers can be deployed at different levels of the control hierarchy. As an example, a Data Collection controller can be used as top controller for all event building applications, as well as a controller for a group of them. In general, such a controller can be in charge of other controllers or of endpoint data taking applications.

The DataFlow controllers make use of the configuration database to extract the information on the elements they are requested to supervise. Their duty is to start, control and stop the hardware and software data taking elements, to monitor the correct functioning of the system, gather operational statistics information and perform local error handling for those kinds of errors which couldn't be handled by the data taking nodes, but do not need to be propagated further to higher control levels.

12.3.2 HLT Farm Supervision

The emphasis for HLT control is the synchronisation of the management of the computer farms with the control of the other systems of TDAQ. It is assumed that the farm for a high level trigger is divided into a set of subfarms, each under supervision of a specific controller. These controllers have well defined tasks in the control for the underlying processing tasks.

The High Level Triggers perform the final selection before sending events to permanent storage. They consist of the Second Level Trigger (LVL2) and the Event Filter (EF). The two stages of the HLT are implemented on processor farms, divided into a number of subfarms. A key design principle has been to make the boundary between LVL2 and EF as flexible as possible in order to allow the system to be adapted easily to changes in the running environment (luminosity, background conditions, etc.) Therefore communalities between the two sub-systems needed to be developed as fully as possible. Bearing this in mind, a joint control and supervision system has been designed. It is also in use for the DC described in Section 12.3.1.

The Online Software configuration database describes the HLT in terms of the software processes and hardware (processing nodes) of which it is comprised. The HLT supervision and control system uses the configuration database to determine which processes need to be started on which hardware and subsequently monitored and controlled. The smallest set of HLT elements which can be configured and controlled independently from the rest of the TDAQ system (i.e. a 'TDAQ segment') is the subfarm. This allows subfarms to be dynamically included / excluded from partitions during data-taking without stopping the run. Supervision and control for each subfarm is provided as a local run controller, which interfaces to the Online Software run control via a farm controller. The controller provides process management and monitoring facilities within the subfarm. The controller maintains the sub-farm in the best achievable state by taking appropriate actions, e.g. restarting crashed processes.

Where possible, errors are handled internally within the HLT processes. Only when they cannot be handled internally are errors sent to the supervision and control system for further consideration.

The Online Software services are used by the supervision system for monitoring purposes. For example, IS will be used to store state and statistical information which could be displayed (for example) by a dedicated panel in the Online Software graphical user interface.

12.4 Control Coordination

The control of the experiment is given by the interplay between three systems: the LHC machine, the detector control and the TDAQ control. For each of them the status of the system under control is expressed in distinct states.

12.4.1 Operation of the LHC machine

The phases of the LHC define a multitude of states [12-4] important for the internal functioning of the machine. A subset is of direct interest for the interaction with the experiment control, in particular those states and parameters which describe the condition of the beam with consequences for the operation of the detector. Phases with stable beam and low background indicate that it is safe to bring the detector to the operational state as required for physics data-taking.

The main phases to consider here are the following: *Filling* the beam from the SPS into the LHC, *ramp*, when the beam is accelerated up to its nominal energy, *squeezing* the beam, prepare for physics and *Collide*, *Physics* with stable beam, beam *Dump and Ramp-down* and *Recover*. It is also interesting to know if no beam is expected during the following hours since these periods will be used by the experiment to perform maintenance and test operations. The estimated duration of these periods is also of importance since the actions to be taken on the detector equipment will vary, e.g. HV reduction for short machine interventions or shut down in case of major problems.

12.4.2 Detector States

As it has been presented in Section 12.2, the operation of the different sub-detectors will be performed by means of FSM. The FSM approach allows for sequencing and automation of operations and it supports different types of operators and ownership, as well as the different partitioning modes of the detector. The FSM will handle the transition of the different parts of the detector through internal states.

Figure 12-3 shows the internal states for a given sub-detector. The detector states are mainly determined by the status of the HV system. However, the status of the other systems of the detector, as well as of the external systems will also be considered. The starting situation for a sub-detector is the *Off* state. This subdetector may transit to the *Stand-by* state after the successful configuration of the front-end equipment. The transition to the *Ready* state will be performed through various intermediate states, which are mainly determined by the operational character-

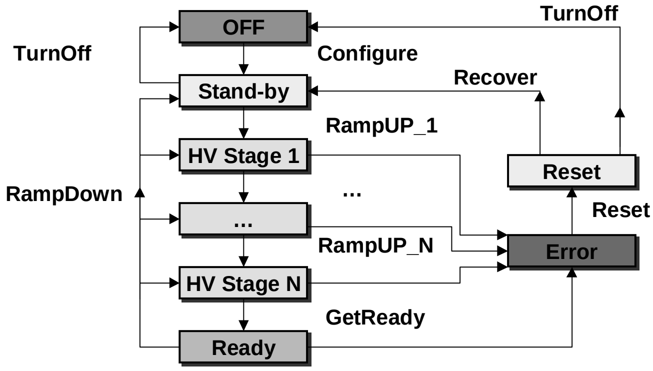

Figure 12-3 Detector states and transitions.

istics of the HV system of the sub-detector. The number of intermediate states is different depending on the sub-detector and is defined according to recipes loaded from the configuration database. In the *Ready* state the sub-detector equipment is ready for physics data taking. The DCS also permits to turn off or bring the sub-detector hardware into the *Stand-by* state in a controlled manner after a run. If an error is detected during the transition to any of these states or during data taking, the subdetector will go to the *Error* state, where dedicated recovery procedures will be applied depending on the type of failure.

The global operation of the DCS will be performed by a single FSM whose states will be built up from the states of the different sub-detectors. Any command issued at this level, which triggers a state transition, will be propagated to the sub-detectors. Similarly, any incident, which affects to the normal operation of a sub-detector, will be reported and it will trigger the state transition of the FSM to the *Error* state.

12.4.3 Operation of the TDAQ States

Three main TDAQ states from *Initial* to *Configured* and *Running* have been introduced in Section 3.1. Here the states are further sub-divided as explained in [12-5] and shown in Figure 12-4. Two state transitions are traversed between *Initial* and *Running*. Before arriving at the *Initial* state the software infrastructure is initialized. The loading of the software and configuration data is performed which brings the system to the *Loaded* state. The system configures the hardware and software involved and enters the *Configured* state. The TDAQ system is now ready to start data-taking. In the subsequent *Running* state the TDAQ system is taking data from the detector. Data-taking can be paused and the L1 busy is then set. Subsequently the run can be continued.

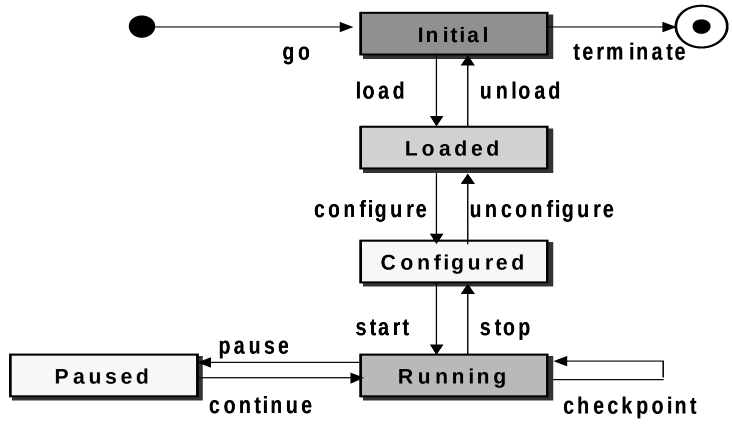

Figure 12-4 TDAQ states

The checkpoint is a transition in a running TDAQ system which is triggered by a change in conditions or by an operator. It results in the following events being tagged with a new run number and does not need the synchronisation, via run control start/stop commands, of all TDAQ elements. Some components in the TDAQ control system require TDAQ sub-states which are used for synchronisation during certain transitions.

12.4.4 Connections between States

As it has been presented in the previous sections, the LHC, the DCS, and the TDAQ system will each be operated through states. Synchronization between these systems is required in order to ensure the quality of the data and the safe operation of the detector. The communication with the LHC is handled by DCS as described in Chapter 11. It transfers both the LHC states and some of its operational parameters to the TDAQ control. On the other hand, parameters measured by the TDAQ system like luminosity, background and beam position, can be used to tune the beams and therefore, must be transferred to the LHC.

Figure 12-5 shows the overall connection for physics data-taking between the TDAQ and DCS states and the LHC conditions. The actions performed by the DCS on the sub-detector hardware are coordinated with the states of the LHC machine. This is the case for the ramping up of the high voltages of some sub-detectors, like the Pixel or SCT trackers. These sub-detectors are more vulnerable to high beam background if the high voltage is on, and hence the command to get ready can only be given when the accelerator provides beam with sufficiently low background. The sub-detector states will closely follow the operation of the LHC. However periods of particle injection or acceleration in the LHC may already be used by TDAQ to initialize and configure the different parts of the systems, like the front-end electronics. For physics data taking it must be ensured that the LHC provides stable beams and collisions and that DCS is in the Ready state. When the TDAQ system is in Configured state, the operator can give the command to start physics data-taking. During Physics data-taking bi-directional communication continues to take place to assure the correct coordination and enable the optimization of the beam.

.

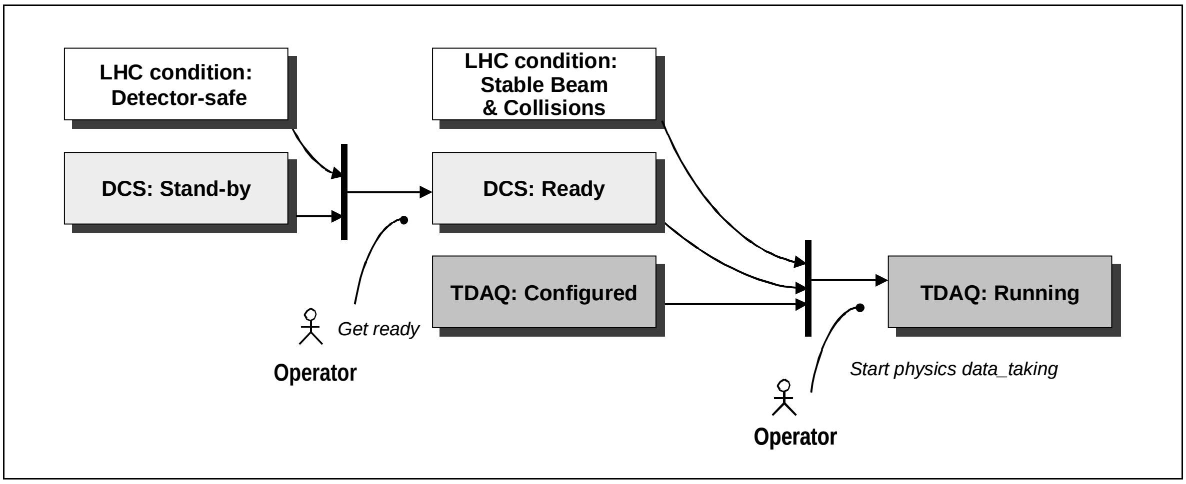

Figure 12-5 Basic example connection between the TDAQ states, the detector states and the LHC conditions

The TDAQ control is only partially coupled to the LHC and sub-detector states. State transitions of the TDAQ system are not determined by the state of the LHC. The TDAQ system can generally be brought from Initial to the Configured state while the LHC is ramping, squeezing and preparing for physics, or while the DCS prepares the detector for data taking. A new run can be triggered at any time regardless of the state of the LHC or of the detector. Data read-out may already start under poor beam conditions using only certain sub-detectors, like the calorimeters, while the high voltage of other detectors will still be set to Stand-by. As soon as the safe operation of the remaining sub-detectors is possible, the DCS will prepare them for data taking and will communicate their availability to the TDAQ system. At this moment, the physics data taking may start.

Although some calibration procedures, for example with cosmic rays or with a radioactive source, will be performed without beam, the communication and coordination with the LHC is still needed in order to avoid wrong operations and hence damage to the detector. For most TDAQ internal system tests no co-ordination with other states need to take place.

12.5 Control Scenarios

In the following section typical scenarios on the experiment control are presented. The first scenario describes the actions when driving the systems to the *Running* state and back to the original situation. Then the control of the various types of runs like physics and calibrations runs, introduced in Section 3.2.4 is discussed. The control functionalities required during commissioning runs are similar to both physics and calibration runs and therefore no separate control scenario is devoted to it. The procedures described rely on the Atlas partitioning concept which is explained in Section 3.3.

12.5.1 Operational Data-taking Phases

The TDAQ states as described in Section 12.4 are traversed when the TDAQ system is run through initialisation, preparation, data-taking and shutdown phases. As the DCS is required to always be operational, in this scenario is assumed that the DCS is in *Stand-by* state and ready to connect to TDAQ. The TDAQ states provide synchronisation points between the systems and sub-systems involved. During the state transitions, actions specific to the sub-system like initializing software and hardware elements, loading software and parameters to hardware modules and configuring them, or starting processing tasks, are performed. Time-consuming operations are preferably performed during early state transition.

12.5.1.1 Initialisation

The preparation for data taking requires the initialization and configuration of all TDAQ hardware and software elements needed for the event readout, as well as a close coordination with the DCS, which acts on the sub-detector equipment.

When initiating a data-taking session, the operations of the TDAQ system start from booted but idle machines. The TDAQ operator selects a partition which is described in the configuration database. The infrastructure, consisting of a number of servers in the distributed system (i.e. the Information Services), is started and initialized. The correct functioning of the hardware and software elements of the TDAQ infrastructure is then verified. Sequence and synchronisation of these start-up operations follow the dependencies described in the configuration database. The TDAQ-DCS communication software is started and the communication between both systems is established.

Once the TDAQ infrastructure is in place, the controllers and the application processes, which are part of the configuration, are booted. The TDAQ process management is de-centralized and can therefore occur in parallel. The TDAQ system passes the information of the chosen partition to DCS. The TDAQ controllers responsible for the command exchange with the DCS, connect to the individual sub-detectors. Having successfully finished this transition the TDAQ system is in the *Initial* state.

12.5.1.2 Preparation

Once all processes have been booted successfully the operator can cycle the system through the states. These states are used to synchronize the loading and configuring of software applications and hardware equipment which take part in the data-taking process.

During the *Loading* transition, the initialisation of all the processing elements in the system including for example, the loading of the software and configuration data, is performed. During the following transition, called *Configuring*, the configuration of a loaded system, for example the realization of connections between TDAQ elements or the setting of parameters, is performed.

The preparation of the sub-detector equipment for data taking comprises the issuing of commands from the TDAQ system to DCS with the corresponding execution of several control procedures. These commands can be associated to state transitions of the TDAQ controllers, or be asynchronous commands issued directly by the TDAQ operator or by applications. The actions are defined according to recipes previously loaded in DCS from the configuration database and are sub-detector specific. The different procedures to be performed on the equipment are previously validated and cross-checked with the states of the external systems and of the common infrastructure to guarantee the integrity of the equipment, e.g. stable beams and acceptable backgrounds must be ensured by the LHC machine. In some cases, their execution can take up to several minutes depending on the characteristics of the sub-detector. These actions on the detector equipment take place in parallel to the loading and configuring of the elements which are directly under TDAQ control. At each of these stages, further synchronization with the DCS may be provided by issuing commands.

The operations described up to here may be time-consuming and should therefore be performed a significant time before the run start is required, for example when waiting for stable run conditions. The availability of the sub-detector for data taking is reported to the TDAQ system via the DDC. Generally the DCS has to be in the *Ready* state when the TDAQ operator starts a run. However, the possibility to start a new run regardless of the state of the DCS is also provided.

When the operations are completed, the TDAQ system is in the *Configured* state and ready to receive the command for data-taking.

12.5.1.3 Data-taking

When the run is started by the TDAQ operator, the L1 busy is removed and event data-taking operations are activated. If necessary, a run can be paused and resumed in an orderly manner with minimum time overhead. On the occurrence of special conditions the checkpoint transition, as described in Section 3.2.6, can lead to a change in run number implying also here only a minimum time-overhead.

Partitions with one or more TTC zones can be split off the main data-taking partition, for example in case of problems with the respective detector part. A checkpoint transition is initiated which sets automatically the L1 busy. The information of the unavailability of the respective TDAQ resource or segment is passed on to higher level elements in the data-flow chain. The L1 busy is removed and data-taking can continue without the removed partition. The removed sub-detector can be configured for stand-alone mode to allow for testing and repairing of the faulty element. Once the removed partition is functional again, it can be joined to the main partition by once more making use of the checkpoint transition.

Depending on the TDAQ system elements which are involved, these actions may require the re-configuring of hardware or software modules and, in this case, it may be necessary to stop and re-start the run. However, it is possible to remove and join sub-farms without affecting the data-taking and without stopping the run.

Component failures which cannot be handled locally are reported through the controllers to the system via the control communication mechanism. These mechanisms are described in Chapter 6, "Fault tolerance and error handling" and in Section 10.5.3, "Control Architecture".

12.5.1.4 Stopping

When the operator stops the run, the L1 busy is set and all data-taking activities are stopped. The control and application processes involved remain active. No changes on the DCS side are foreseen, the sub-detectors remains in the *Ready* state and TDAQ in *Configured* state.

12.5.1.5 Shut-down

On receipt of the shut-down command clean-up operations in software and hardware are performed in the TDAQ system. The previously started applications and then the controllers are stopped. Finally the TDAQ infrastructure is removed in an orderly manner to leave the system in a state in which a new and independent data-taking session can be started. If no further data-taking is foreseen the Ramp Down or Turn Off commands are given to DCS in order to bring the detector to a safe state.

12.5.2 Control of a Physics Run

For physics data-taking the hierarchy of TDAQ controllers including all sub-detectors is arranged in a single common partition. The information on the type of run is transferred to the DCS system to prepare the different subdetectors for physics data-taking as described in the previous section. The successful execution of the appropriate DCS procedures to bring the sub-detectors to the *Ready* state, is then reported to the TDAQ system.

Figure 12-6 shows an example of the experiment control including the TDAQ control and the back-end system of the DCS. The TDAQ Atlas root controller holds the highest level of control. It connects to the detector controllers, the farm controllers for EF and L2 and the DC controller as described in Section 12.3. Each sub-detector controller supervises the controllers of the sections and also the controller which provides the connection to DCS for each sub-detector. The RODs are supervised by their respective sub-detector section controller. In the following, it is assumed that the DCS is in the *Ready* state and the DAQ control is in the *Running* state.

The control of a physics run is driven by the TDAQ system, which acts as the master of the experiment control by issuing commands to the DCS by means of specialized controllers called DDC_CT. There is one DDC_CT controller per sub-detector. Those controllers send commands directly to the sub-detector control units on the DCS side. This communication model implies that the TDAQ system interacts directly with the DCS of the various sub-detectors.

During physics data taking, only a pre-defined set of high-level commands from the TDAQ system on the DCS, like the triggering of the sub-detector state transitions at the start or end of the run is allowed. The command is logged and feedback on its execution is reported to the TDAQ system. The TDAQ Online software control system handles failures or time-outs from the DDC_CT in the same way as from other controllers in the system.

Global error handling and recovery is provided by the Online system control. Severe problems, for example in the HV system of a certain sub-detector are reported to the TDAQ system. Depending on the problem and on the faulty system element, the TDAQ control may decide to exclude this sub-detector from data-taking and continue the run with the remaining detectors as described in the previous section. The run continues if the readout of the detector part in question is not vital for data-taking for the type of physics chosen at the time as described in Section 12.5.1

HLT sub-farms can be removed or added to the global farm control without disturbance of data-taking activity. Breakdown and replacements of individual sub-farm nodes are handled transparently and each of such operations are logged.

Online calibration of sub-detectors may be performed by injecting calibration events, being marked as such, during a physics run without disturbing the normal data-taking activity.

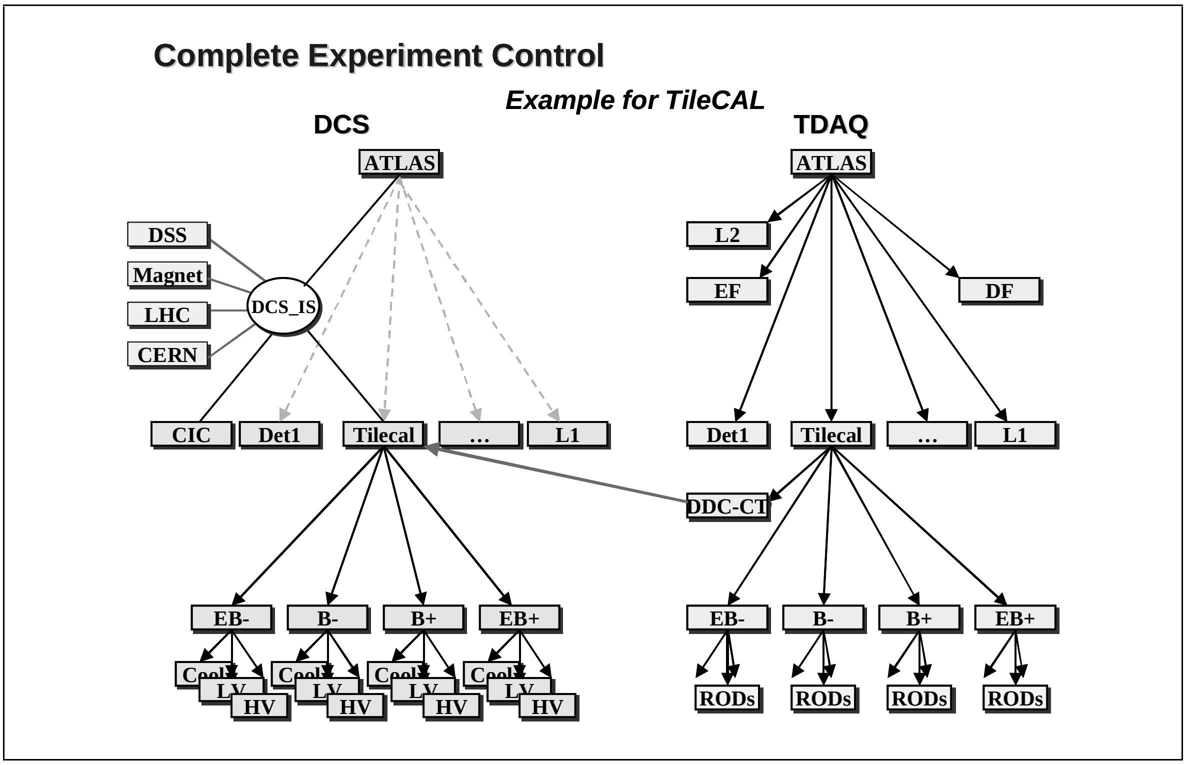

Figure 12-6 Complete Experiment Control mode.

12.5.3 Calibration Run

The calibration procedures envisaged in ATLAS profit from the flexible partitioning concept allowing for autonomous and parallel operations of the sub-detectors or even of the different sections of the sub-detectors. From the point of view of controls, three different types of calibrations can be distinguished:

- Procedures where only the control provided by the Online software system is required, like the determination of the pedestals for zero suppression.
- Calibration runs entirely handled by the DCS like the calibration of the cooling system, where the flow of the liquid is adjusted as a function of the temperature of the detector.
- Calibration procedures requiring the control provided by both systems. This is the case, for instance, of the calibration of the Tile Hadron Calorimeter with the CS source, where the modules of the detector are scanned with a radioactive source under control of the DCS. The signal produced is read by the DAQ system and the information is used to adjust the HV applied to the PMTs of the readout system.

The control needs in procedures where both systems are required, are similar to the functionality needed in the case of a physics runs presented in the previous section. Figure 12-7 shows the interplay between the TDAQ control and the DCS for calibration of the Tilecal detector. As for physics data taking, these calibration procedures are driven by the TDAQ control and commands follow the same path. The main difference with respect to physics data taking is the arrangement of the partitions. In the example presented in the figure, a TDAQ partition is defined for the stand-alone operation of the Tilecal sub-detector. The experiment control system also supports the operation of several partitions in parallel allowing for the calibration of various sub-detectors simultaneously. It is important to note that although some calibrations procedures are executed without beam or even without the control provided by the DCS, the communication with the LHC machine and other external systems must always be guaranteed since this information is of crucial importance for the integrity of the sub-detectors and for the preparation of the next run.

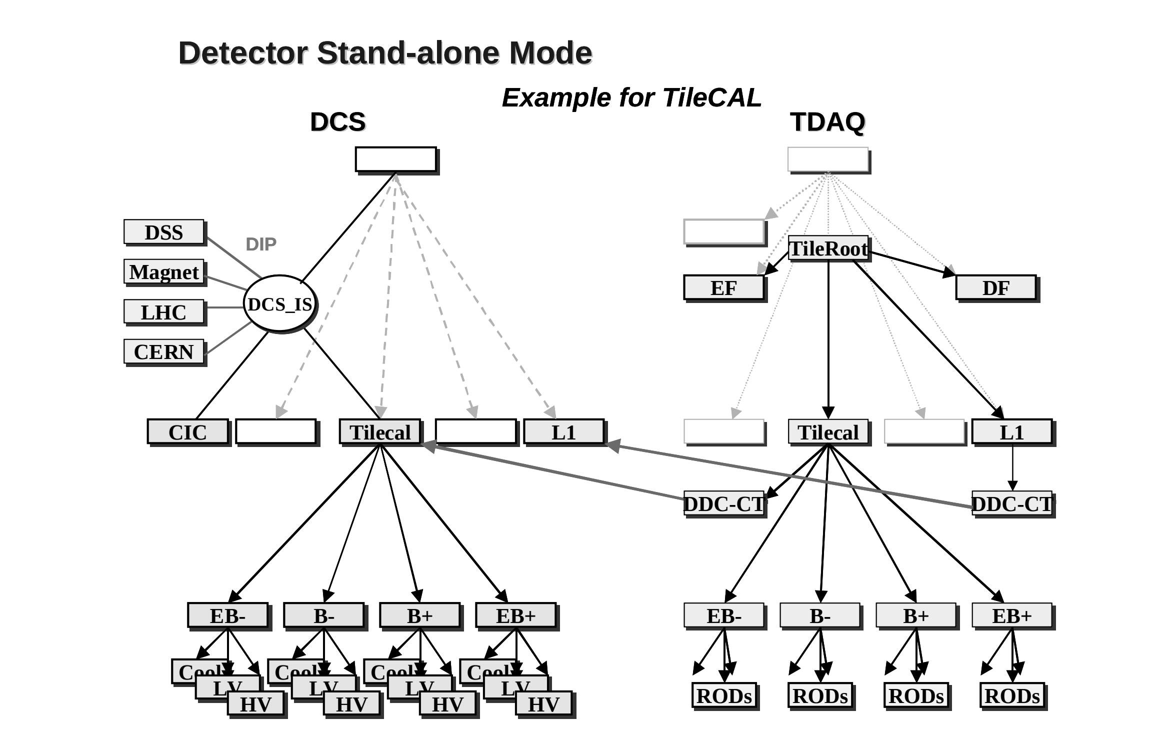

Figure 12-7 Detector Stand-alone mode.

12.5.4 Operation outside a Run

The states of the DCS and TDAQ outside a run are determined by the duration of the periods without data taking. As described in Chapter 3, during transitions between runs and short interruptions, the TDAQ system can be set to one of its intermediate states or be unavailable while the DCS remains ready for data taking. If longer interruptions are foreseen, like periods of machine development, the HV applied to the sub-detectors is reduced and the detector state are set to *Stand-by of Off*.

During shut-down periods and long term intervals without data-taking, the TDAQ system is not necessarily operational although its functionality may be available on demand for calibration and debugging purposes. However, the full functionality of the DCS is required in order to supervise the operation of the detectors and of common services. In this scenario, the DCS still allows for the stand-alone and integrated operation of the sub-detectors without TDAQ. In the former case, the operation are performed from the sub-detectors control stations, having full control of the sub-detector, whereas in the latter case, the overall control is performed from the global operation station. A number of sub-detector services like the LAr cryogenics or ID cooling stay operational. The monitoring and control of the humidity and temperature of the electronics racks, the supervision of the un-interruptible power supply system and of other detector specific equipment is also performed. Permanent access to the conditions and the configuration databases is available for DCS.

The ATLAS magnet are permanently switched on and therefore the interface with the DCS must be continuously available. The radiation levels monitored by the LHC machine must be accessible by the DCS at all times. Similarly, the interface to the fire brigade and to the access security system, as well as to the DSS must be continuously operational.

12.6 References

12-1 PVSS-II, http://www.pvss.com

12-2 JCOP Architecture Working Group, "Framework Design Proposal", CERN-JCOP-2000-06, 2000, http://itco.web.cern.ch/itco/Projects-Services/JCOP/SubProjects/Architecture/pdf/HC7.0.pdf

12-3 *Partitioning - Global issues working group document*, http://atlas-project-tdaq-giwg.web.cern.ch/atlas-project-tdaq-giwg/Documents/Documents.htm

12-4 LHC Operations project, http://lhcop.web.cern.ch/lhcop/

12-5 *Runs and States - Global issues working group document*, http://atlas-project-tdaq-giwg.web.cern.ch/atlas-project-tdaq-giwg/Documents/Documents.htm