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. Although DCS is formally part of TDAQ, its functionality has been logically separated here in the description of the experiment control. 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 overall system has already been discussed in Section 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 high degree of coordination between these two systems and with the LHC machine. The interaction with the LHC machine will be handled by the DCS as illustrated in Figure 5-3 and presented in detail in Chapter 11. The TDAQ system is in overall control of data-taking operations.

The general control of the experiment requires a flexible partitioning concept as it is described in Section 3.5, 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 implement the partitioning concept of ATLAS. The finest granularity of the TDAQ system is given by the segmentation of the sub-detectors into TTC zones [12-2] [12-3]. For this reason, the different sections of the sub-detectors will be logically represented in the back-end software of the DCS, presented in Section 11.4, by means of the control units, which will be operated as a Finite State Machine (FSM). A logical control unit models the behaviour of a sub-detector or a system, e.g. the Pixel sub-detector or the high voltage system. 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 or device units. The device units are responsible for the direct monitoring and control of the equipment, i.e. they represent hardware devices like a high voltage crate or a temperature sensor. According to this model, the DCS of the Tilecal, for example, may be described by a tree of logical control units: the root logical unit controls the sub-detector itself and four independent child units, which control the four sub-detector sections (two extended barrels, EB+ and EB-, and the central barrel sections, B+ and B-). The sub-detector control units are in charge of supervising the various device units. 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.

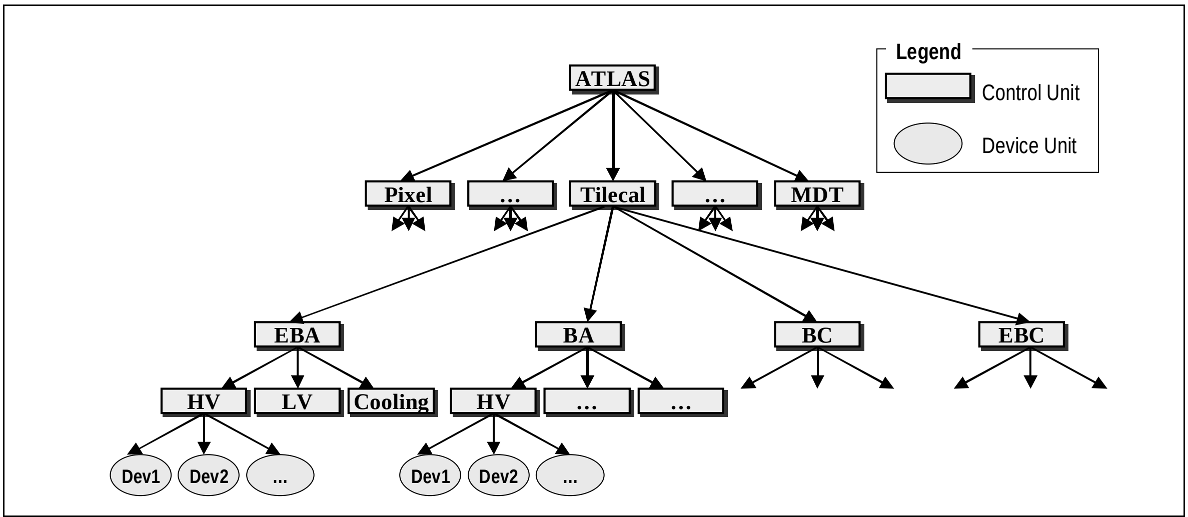

Figure 12-1 DCS Logical Architecture

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 directly from the DCS control, whereas during physics data taking and calibration procedures, commands will be sent from the TDAQ control. Therefore, a priority scheme, which avoids the possibility of issuing conflicting commands must be provided. This mechanism will be developed according to the recommendations of the JCOP Architecture Working Group [12-4].

12.3 TDAQ Control

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 detector into detectors and sub-detectors leads to a hierarchical organisation of the control system. The basis of the TDAQ control is provided by the 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 comprised of a large number of controllers which are distributed in a hierarchical tree following the functional composition of the TDAQ system.

This hierarchy 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 [12-5]. The leaf controllers on the lowest level, the so-called *local controllers*, are responsible for the control of physical hardware such as ROSs. All elements of the TDAQ system (e.g. ROSs, HLT farms etc.) employ the same type of controllers, and follow a similar structure, as discussed in Section 12.3.1 and Section 12.3.2.

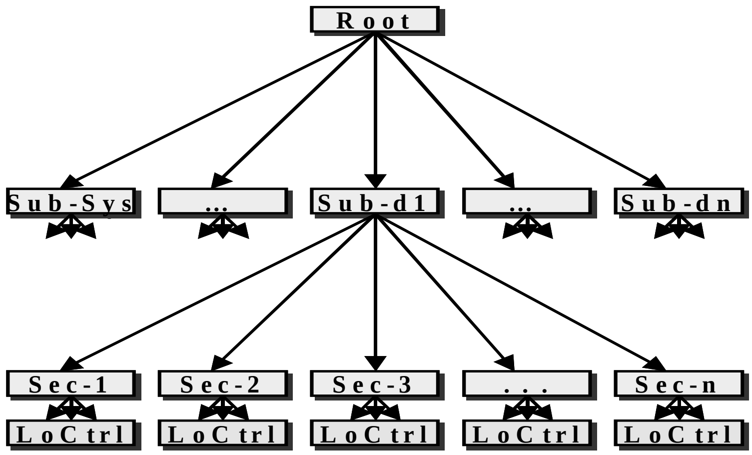

Figure 12-2 Online Software Control Hierarchy in ATLAS

A controller in the TDAQ system is characterized by its state (e.g. *Running*, *Paused*, *Stopped* etc.) given by the TDAQ state model described in Section 12.4.3. In any place of the hierarchy, a change of state may be initiated and synchronized from a given level of controller and sent down to the next lowest level. Once the requested transition has been performed, status information is returned to the requesting controller. Possible error conditions are also reported back.

A controller framework provides the facilities to handle 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 may vary widely from initializing hardware readout to EF farm control. The information on the inter-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 (controllers under its control) and for signalling its overall state to its parent. Of particular importance is the synchronisation necessary to start data-taking operations. This is performed by consecutive coordinated transitions through a number of intermediate states until data-taking is finally started as described below in Section 12.5.1. The shift operator drives the operations by issuing commands to the root controller. The inverse series of 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 monitoring and handling any error conditions. In the case of a malfunction of e.g., a detector, the controller can start corrective actions and / or signal the malfunction by issuing error messages. Severe malfunctions which are beyond the capabilities of a controller can be signalled by a state change to its parent (e.g. moving from a *Run-*

ning state to a *Paused* state). It is then the role of the parent controller to take further actions depending on the circumstances of the malfunction.

The design of the control, supervision and error handling functions is based on the use of a common expert system. Specific controllers will use dedicated rules (information which tells the controllers what to do in one of many defined situations) to perform their individual functions in addition to common rules which handle more general situations.

12.3.1 Control of the DataFlow

The DataFlow control handles all applications and hardware modules responsible for moving the event data from the detector front-end electronics and LVL1 trigger to the HLT system. It includes the control of the RoIB, the ROS, the Event Building and the SFOs. It also includes the overall control of the detector ROD crates.

There are two types of local DataFlow controllers foreseen, both making use of the Online software infrastructure in the same way, as shown in Figure 12-2. The ROS controller is tailored for the control of ROS software applications and hardware devices which do not themselves access the online software facilities. The Data Collection (DC) controller handles all other DataFlow applications and is optimized for the control of collection of processors. A version of the latter is also used for the control and supervision of the HLT and is described in Section 12.3.2. The other type, the ROD crate controller, takes over the communication handling between the other Online software services, Information Sharing and Configuration Databases, and the hardware which it controls.

Both controllers can be deployed at different levels of the control hierarchy. As an example, a DC controller can be used as top controller for all event building applications, as well as a controller for a group of them. The DataFlow controllers request information on the elements they supervise from the configuration database. Their function is to initialize, configure, start, pause and stop the hardware and software data taking elements, to monitor the correct functioning of the system, to gather operational statistics information and to perform local error handling for those kinds of errors which cannot be handled locally, but do not need to be propagated further to higher control levels.

12.3.2 HLT farm supervision

The emphasis for HLT farm supervision is the synchronisation and control of the HLT farms (LVL2 and EF) with that of the other TDAQ systems. An HLT farm is comprised of a set of sub-farms, each under supervision of a specific controller, as shown in Figure 12-3. These controllers have well defined tasks in the control for the underlying processes running on the sub-farms.

A key overall design principle of the HLT 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 a joint control and supervision system for the two sub-systems has been designed and developed.

The configuration database contains data describing the software processes and hardware (processing nodes) of which the HLT is comprised as well as a wide variety of set-up and configuration data for pre-defined data-taking scenarios. The HLT supervision and control system uses these data 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 is the sub-farm. This allows sub-farms to be dynamically included/excluded from partitions during data-taking without stopping the run. Dedicated local run controllers supervise and control each HLT sub-farm, and report to an overall farm controller. This in turn condenses the information coming from each sub-farm and reports the global state of the farm (LVL2 farm or EF farm) to the run control. The controller also provides process management and monitoring facilities within the sub-farm. An expert system will guide the actions of the sub-farm controllers (e.g. spontaneously restarting crashed processes) such that it can maintain the sub-farm in the most efficient data-taking condition.

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

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 can then be displayed to the operator or made available to other control elements (e.g. the malfunction of more than a pre-defined number of HLT sub-farms may precipitate the stopping of data-taking).

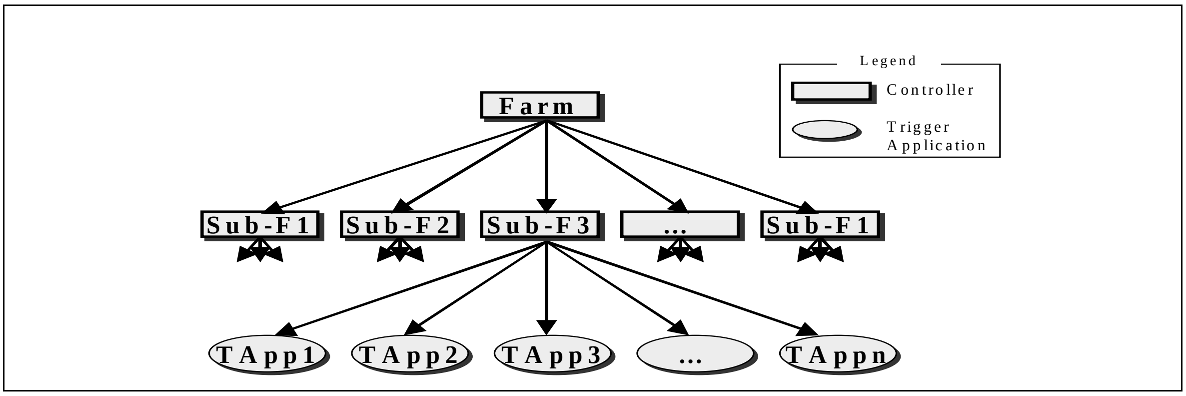

Figure 12-3 HLT farm control

12.4 Control coordination

The control of the experiment depends on the interplay between three systems: the LHC machine, the detector control and the TDAQ control. For each of them, their instantaneous status is expressed in terms of distinct logical states (e.g. *Configured, Running, Stopped* etc.). The states of a given system are highly correlated to those of other systems and only certain combinations of states are allowed at any given time. Some examples are given here:

- The LVL1 trigger will not be started unless there are circulating beams in the LHC;
- The LVL2 trigger will not been started if some parts of its configuration process failed;
- Physics data taking will not start if the detector high-voltage has failed.

The definition of these states, of their interplay, and of what actions are to be taken in what circumstances is a highly complex problem which will need the support of an advanced expert system.

12.4.1 Operation of the LHC machine

The phases of the LHC define a multitude of states required for the internal functioning of the machine. A sub-set 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 may indicate that it is safe to bring the detector to the operational state for physics data-taking, whereas abnormally high background conditions may dictate that the detectors are turned off for safety reasons.

The nominal main phases to consider for the LHC are the following:

- *Filling* — the beam is being transferred from the SPS into the LHC;
- *Ramping* — the beam is being accelerated up to its nominal energy;
- *Squeezing* — the beam is being prepared for physics;
- *Collide* — physics with stable beam;
- *Dump*, *Ramp-down* and *Recover* — restart states.

It will also be necessary to know if no beam is expected for a given time 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 controlled by means of FSMs. 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 and, together with the use of an expert system, will be vital in the coherent operation of ATLAS.

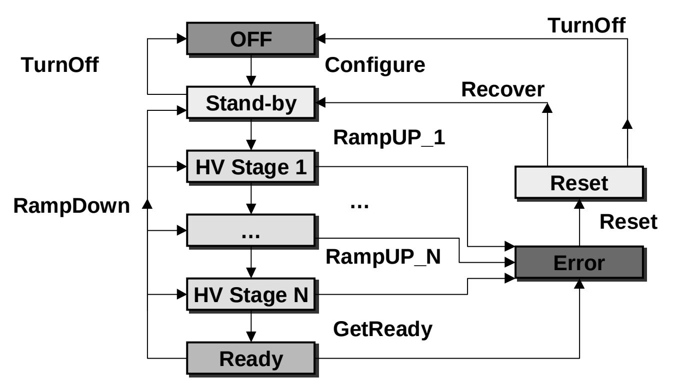

Figure 12-4 Detector states and transitions

Figure 12-4 shows an example of the internal HV states for a given sub-detector.

The sub-detector states are mainly determined by the status of the HV system. However, the status of the other systems like the cooling or low voltage systems, as well as of the external systems will also be considered to build up the state of the sub-detectors.

The starting situation for a sub-detector is the *Off* state in which the high voltage system and, in some cases, the low voltage system are switched off although other services like the cooling may be permanently operational. On request from run control, this sub-detector may transit to the *Stand-by* state after the successful configuration of the front-end equipment. In this state the low voltage system may be turned on and reduced HV values are applied to the sub-detectors. The transition to the *Ready* state will be performed through various intermediate states, which are mainly determined by the operational characteristics of the HV system of the sub-detector. The number of intermediate states is different depending on the sub-detector and is defined in the configuration database. In the *Ready* state the sub-detector equipment is ready for data taking. Many of these state changes, though requested by the run control, will actually pass for execution to the DCS (e.g. the run control will ask a detector's HV to be turned on by DCS). The DCS may also 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 sub-detector will go to the *Error* state, where dedicated pre-defined 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 may trigger a transition to the *Error* state.

12.4.3 Operation of the TDAQ states

The three main TDAQ states of *Initial*, *Configured* and *Running* (see Figure 12-5) have been introduced in Section 3.2. These states are now further sub-divided as explained.

Starting from booted but idle machines, the software infrastructure is initialized with the go command to arrive at the *Initial* state. Two state transitions are traversed between *Initial* and *Running*. 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 continued causing the LVL1 busy to be set and removed, respectively. The inverse transitions bring the TDAQ system back to the *Initial* state and it can be ended with the terminate command to arrive to the original state which has no TDAQ software infrastructure.

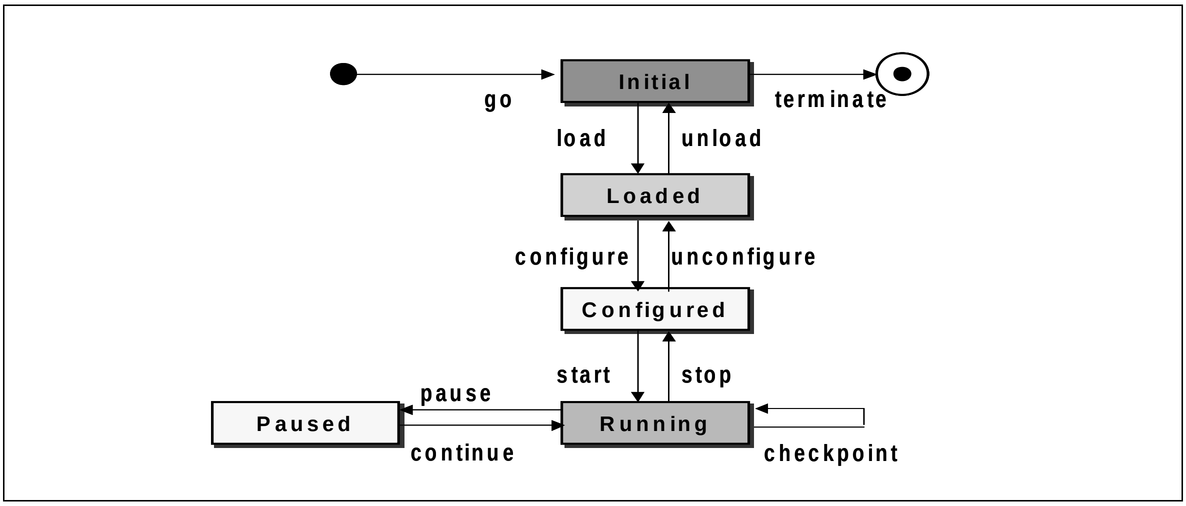

Figure 12-5 TDAQ states

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 events following the checkpoint to be 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 several intermediate states which are used for synchronisation during certain transitions that are internal to the component.

12.4.4 Inter- relation of states

As has been discussed above, coherent synchronization between system states is vital in order to ensure the quality of the data and the safe operation of the detector. Communication with the LHC is handled by DCS as described in Chapter 11. DCS transfers both the LHC states and some of its operational parameters to the TDAQ control. On the other hand, parameters measured within 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 12-6 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 example for the ramping up of the high voltages of some sub-detectors, like the Pixel or SCT tracker. 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, such as the front-end electronics.

For physics data taking the LHC will be required to signal that stable beam conditions exist. Only when the TDAQ system is in *Configured* state, and the DCS is ready, the operator may 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 to enable the optimization of the beam.

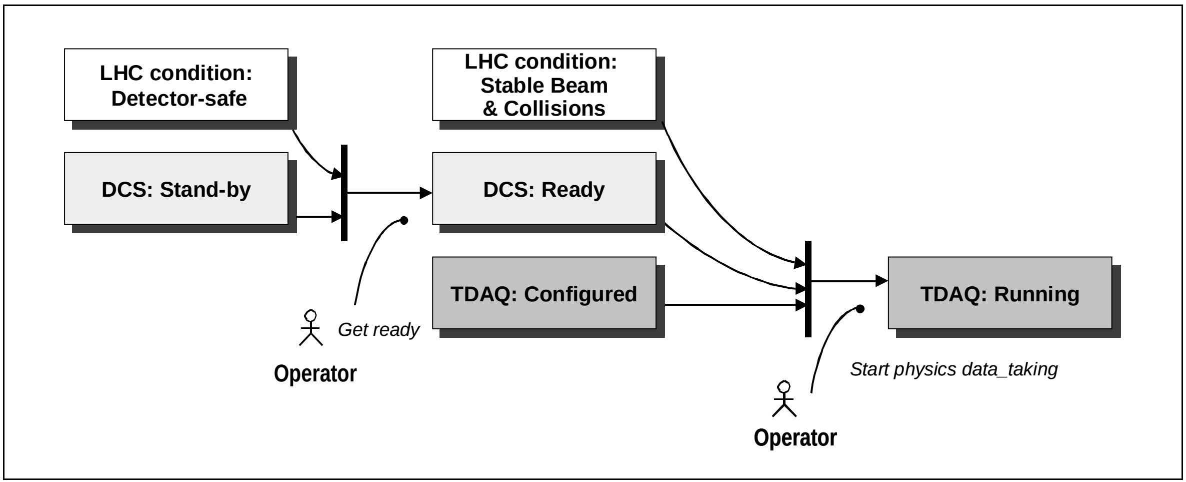

Figure 12-6 Basic example connection between the TDAQ states, the detector states and the LHC conditions Note that the cross-check of LHC conditions is performed by the DCS (not shown here)

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. However, changes of the LHC parameters observed in the detector, like the luminosity, may require actions on the DAQ system, like the change of the trigger menus.

The TDAQ system can generally be brought from the 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 control system is still needed in order to prevent dangerous operations to be engaged which could possibly damage the detector. For most TDAQ internal system tests no co-ordination with the state of the other systems needs to take place.

12.5 Control Scenarios

In this section, typical scenarios for the experiment control are presented. The first scenario describes the actions required to drive the system from its initial dormant state to the *Running* state and back again. The control of the various types of runs such as physics and calibration runs, introduced in Section 3.3, is then discussed. The procedures described rely on the ATLAS partitioning concept which is explained in Section 3.4.

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 it is assumed that the DCS is in *Stand-by* state and ready to connect to the rest of the system. During the operational phases, systems and sub-systems perform a number of operations independent from each other. 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 transitions. The aim is to prepare the TDAQ system to be ready for the start command such that the running state can be entered with a minimum time-overhead.

12.5.1.1 Initialisation

When initiating a data-taking session, the operations of the TDAQ system start from booted but idle CPUs. The operator selects a partition which is described in the configuration database or defines a new one. The process management and information distribution infrastructure, consisting of a number of servers in the distributed system (i.e. process managers, Information Services, Error Message system), is started and initialized. The hardware and software elements of this TDAQ infrastructure are then verified to ensure that they are functioning correctly (for example the servers are tested, the process managers are contacted and therefore the status of their host PCs and the network connection is verified). The sequence and synchronisation of the start-up operations follow the dependencies which are described in the configuration database. The Online Software-DCS communication software is started and the communication between the systems is initialized.

Once this TDAQ infrastructure is in place, the controllers and the application processes, which are part of the configuration, are started. The process management is de-centralized and therefore this can occur asynchronously. The data describing the selected partition is transmitted to DCS. The control communication with DCS is established. Having successfully finished this transition the TDAQ system is in the *Initial* state.

12.5.1.2 Preparation

Once in the initial state, the operator can begin to configure the system. This involves synchronizing the loading and configuring of all software applications and hardware elements which participate in the data-taking process.

The *Loading* transition, is used to initialize all the processing elements in the system including, for example, the loading of the configuration data into the various controllers. Subsequently, during the *Configuring* transition, the various elements of the system are configured according to the information obtained in the previous step, for example, the establishment of communications channels between TDAQ elements such as the L2PUs and the L2SV, or the setting of hardware or software parameters.

The preparation of the sub-detector equipment for data taking involves the issuing of commands to DCS and the subsequent execution of defined control and initialization 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 stand-alone applications. The actions are detector-specific and are defined in the configuration database. They are loaded into DCS at initialisation time. The different procedures to be performed on the equipment are previously validated and are cross-checked with the states of the external systems and of the common infrastructure to guarantee the integrity of the equipment, e.g. the LHC must have stable beam conditions and acceptable backgrounds before certain sub-detectors' HV may be turned on. In some cases, the execution of these procedures may take several minutes, depending on the actions to be taken. These actions on the detector equipment take place in parallel to the loading and configuring of the TDAQ elements.

The operations described so far may be time-consuming and should, if possible, be performed well before the run start is required, e.g. while waiting for stable run conditions to occur. The readiness of the sub-detectors for data taking is reported to the run control by the DCS which will then be in the *Ready* state. Generally speaking, the DCS must be in the *Ready* state for the operator to be able to start a run. However, it is possible to start a run regardless of the state of some or all of the sub-detectors if required (e.g. in commissioning phases when not all the detectors are available). In some cases the TDAQ control may decide to take data without certain sub-detectors or sections of a given subdetector. In this situation, the associated sub-detector run controllers are excluded from the DAQ partition used for data-taking and consequently, no commands will be issued on the subdetector DCS, i.e. the sub-detector will remain in the default state (*Stand-by*) while the other subdetectors participate in the run.

The readiness of the TDAQ system to take data is defined as the moment when the top-level run controller is in the *Configured* state. This implies that all the data acquisition and trigger systems have been booted, initialized and configured correctly and are ready to receive event data.

12.5.1.3 Data-taking

When the run is started by the operator, the LVL1 trigger is enabled and event data-taking operations commence. If necessary, a run can be paused and resumed in an orderly manner with minimum time overhead simply by inhibiting or enabling, respectively, the LVL1 trigger. In specific circumstances, such as the modification of one of a limited set of pre-defined run parameters (e.g. the beam position), the checkpoint transition, as described in Section 3.3.7, can precipitate a change in run number without having to cycle through the full run-stop and run-start procedures. The advantage of this is to minimize the time required to take account of the possible change in value of one or more of a fixed set of parameters.

Partitions comprising one or more TTC zones can be removed from the main data-taking partition, for example in case of problems with a given detector element. To do this, first the LVL1 trigger is inhibited. The detector element in question, together with its associated TDAQ resources (e.g. the ROBs associated to its RODs) are then marked as no longer forming part of the main data-taking partition. Finally the LVL1 trigger is re-enabled and data-taking can continue. The removed detector element can then be configured for stand-alone operations to allow the faulty element to be identified, tested and repaired. Once the element is functional again, it can be attached again to the main partition.

Depending on the system elements which are involved, the above actions may be able to be completed within a checkpoint transition, or they may necessitate the complete cycling through the run-stop and run-start procedures in order to re-configure the system correctly. Component failures during state transitions or while the system is in a given state which cannot be handled locally are reported through the local controllers to the top level run control. This mechanism is introduced in Chapter 6 and described in Section 10.5.3.

12.5.1.4 Stopping

When the operator stops the run, the LVL1 trigger is inhibited and data-taking is 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. Systems and sub-systems may perform different kind of actions during the *stopping* transition as compared to the *pause* transition.

12.5.1.5 Shut-down

On receipt of the shut-down command all applications and their controllers are stopped. Finally the TDAQ software infrastructure is closed down in an orderly manner, leaving the system available for a new data-taking session. 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 by ramping down and switching off the low and high voltage applied to the sub-detectors.

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 sub-detectors 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 12-7 shows an example of the experiment control including the TDAQ control and the back-end system of the DCS. The TDAQ root controller holds the highest level of control. It connects to the detector controllers, the farm controllers for EF and L2 and the DF controllers as described in Section 12.3. Each sub-detector controller supervises the controllers of its 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 TDAQ 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. However, in the case of DDC_CT connection failures the control of the sub-detectors is taken over by the DCS root control unit (called ATLAS in Figure 12-7) to ensure the safe operation of the detector equipment.

During physics data taking, only a pre-defined set of high-level commands from the TDAQ system to 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 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 described in Section 10.5. Severe problems, e.g. 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 may continue if the readout of the detector element in question is not considered vital for the given run.

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.

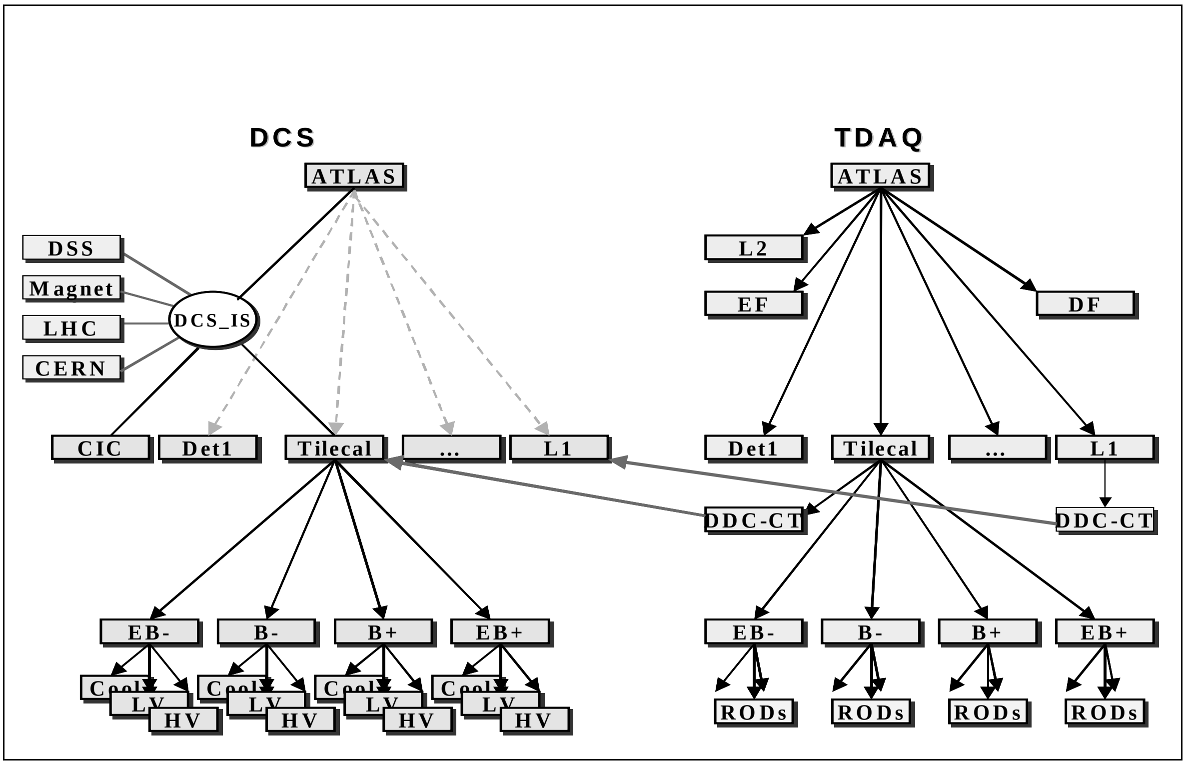

Figure 12-7 Experiment Control indicating the control command flow between TDAQ and DCS, which takes place at each sub-detector level, as illustrated here for the TileCal sub-detector.

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 different elements of a given sub-detector. From the point of view of controls, three different types of calibration can be identified:

- Procedures where only the control provided by the Online software system is required, such as the determination of the pedestals for zero suppression.
- Calibration runs entirely handled by the DCS such as 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 control from both systems. This is the case, for instance, of the calibration of the Tile Hadron Calorimeter with the radioactive caesium source, where the modules of the detector are scanned with the 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.

In the latter case, the control needs are similar to the functionality required for physics runs as presented in the previous section. Figure 12-8 shows the example of 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. It is important to note that although some calibration 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.

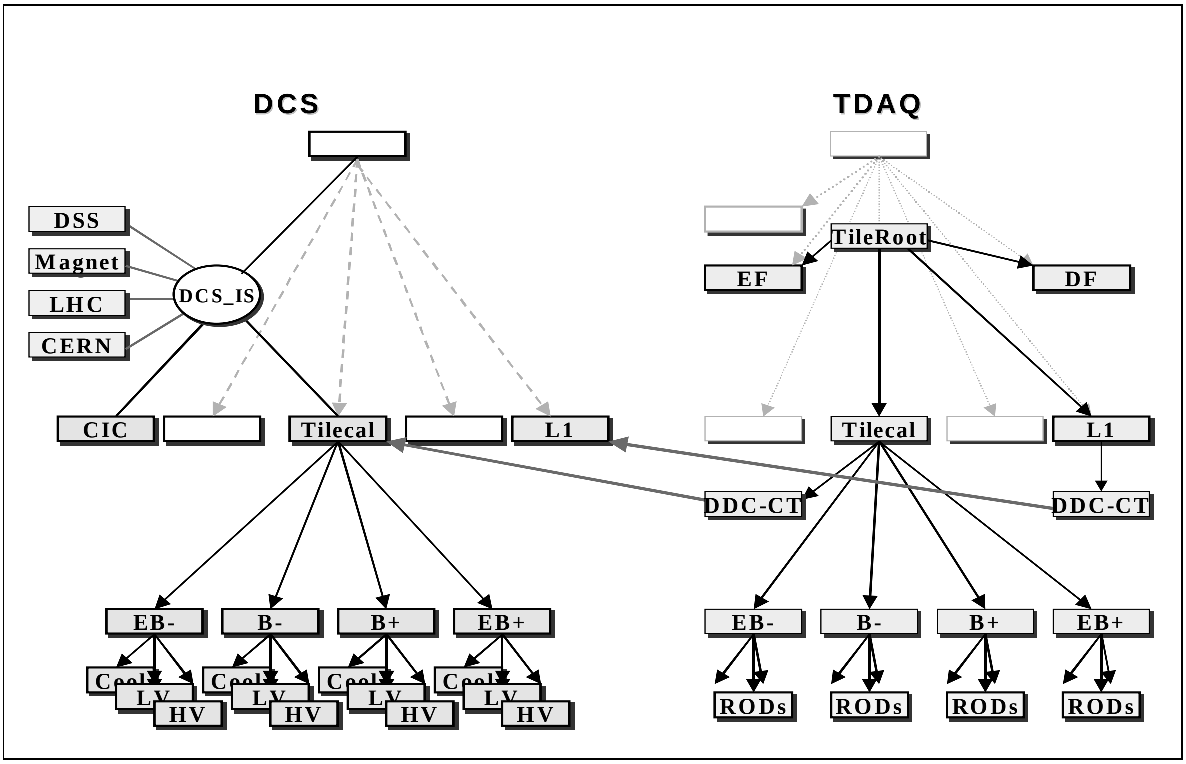

Figure 12-8 Detector Stand-alone control mode, indicating the control command flow between TDAQ and DCS at the sub-detector level, as illustrated here for the TileCal sub-detector

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 completely operational. If longer interruptions are foreseen, such as periods of machine development, the HV applied to the sub-detectors is reduced and the detector states are set to *Stand-by* or *Off*.

During shut-down periods and long term intervals without data-taking, the TDAQ system is not necessarily operational although it may be available on demand for calibration and debugging purposes. In these periods 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. The ATLAS magnets 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. For these reasons, the DCS is required to be fully operational allowing, if necessary, the sub-detectors to operate without the support of the rest of the TDAQ system. Similarly, the interfaces between the DCS and the fire alarm systems, the access security system and the DSS must be continuously operational.

12.6 References

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

12-2 *ATLAS High-Level Triggers, DAQ and DCS Technical Proposal*, CERN/LHCC/2000-17 (2000)

12-3 *ATLAS First-Level Trigger Technical Design Report*, CERN/LHCC/98-14 (1998)

12-4 JCOP Architecture Working Group, *Framework Design Proposal*, CERN-JCOP-2000-06, 2000,
http://cern.ch/itco/Projects-Services/JCOP/SubProjects/Architecture/pdf/HC7.0.pdf

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

12-6 LHC Operations project,
http://lhcop.web.cern.ch/lhcop/

12-7 *Runs and States - Global issues working group document*,
http://cern.ch/atlas-project-tdaq-giwg/Documents/Documents.htm