

## 16 Timing, trigger and control

### 16.1 Introduction

The ATLAS readout elements, such as the front-end electronics, the readout drivers (ROD) and possibly the readout buffers (ROB), need the bunch-crossing signal (BC clock) and the level-1 accept signal (L1A). As specified in Ref. [16-1], each readout element has to produce an event identifier (EVID) and a bunch-crossing identifier (BCID) with each event. It is therefore necessary to make available some synchronization signals, in order to maintain coherence between these identifiers across the experiment. These synchronization signals are the bunch counter reset (BCR) and the event counter reset (ECR). During test and calibration periods, the readout electronics also need to receive test and calibration signals.

The Timing, Trigger and Control (TTC) system allows the timing and trigger signals to be distributed to the readout electronics. The timing signals comprise the LHC clock (BC clock) and the synchronization signals (BCR, ECR). The trigger signals include the L1A, test and calibration triggers. The TTC allows the timing of these signals to be adjusted. The requirements on the TTC are defined in Ref. [16-2].

The ATLAS TTC system is based on the optical fan-out system developed within the framework of RD12 [16-3], which allows signals to be distributed from one source to up to 1024 destinations. The system can be partitioned and subdetectors can be run with the central ATLAS timing and trigger signals, or independently with their own specific timing and trigger signals. The TTC system receives the LHC 40 MHz clock (BC clock) and the ORBIT signal from the LHC, the L1A signal from the central trigger processor (CTP), and commands and data from the CTP and subdetector-specific electronics. Encoding allows this information to be transmitted on a single optical link which is fanned out to a maximum of 1024 destinations. At the receiving end, an ASIC decodes the incoming signal and makes available the BC clock, the L1A signal, the ECR and BCR signals, the EVID and BCID numbers, and the user commands and data. Provision is made to adjust the timing of all the signals. The context diagram of a partition is shown in Figure 16-1.

Some of the subdetectors (e.g. the liquid-argon calorimeters) will use the TTC system to bring signals directly to the on-detector front-end electronics; others (e.g. the inner tracker) will use the TTC system only as far as the RODs, at which point they will change the functionality and the protocol of the signal transmission. Figure 16-2 summarizes the different configurations.

The way a subdetector will use the TTC system will depend on its specific requirements. Most of the subsystems will use more than one partition to allow the concurrent running of different parts of the subdetector in different trigger modes during commissioning or calibration periods. The TTC team will work closely with the subdetector communities in order to define the optimum number of partitions for each subsystem, and to determine how the TTC should be used to set up the timing.

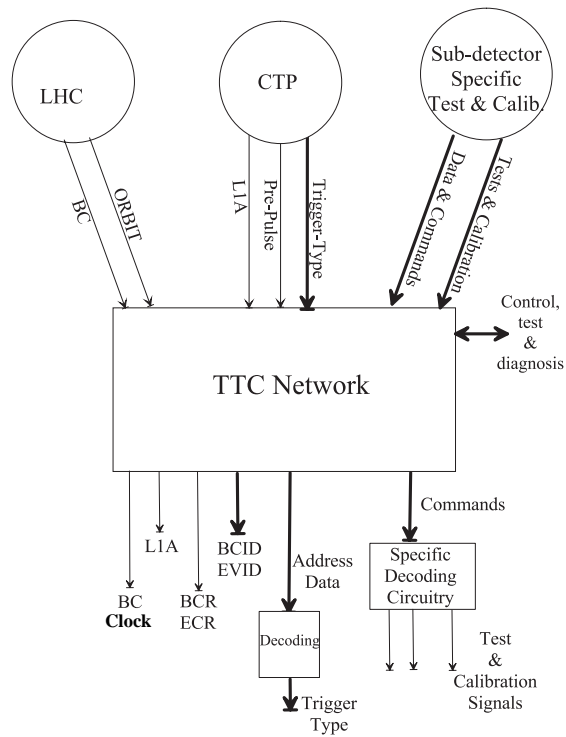


Figure 16-1 TTC context diagram.

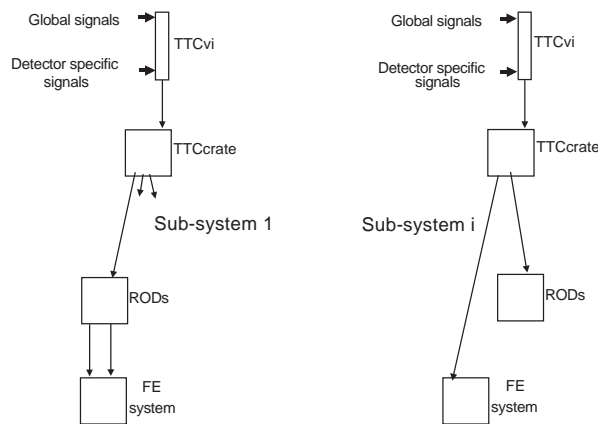


Figure 16-2 TTC configurations.




Figure 16-3 LHC bunch structure.

In ATLAS, the TTC system is used in different ways:

- In normal running, each TTC partition receives its clock from the LHC and the L1A from the CTP. The BCR is derived from the LHC ORBIT signal. After each L1A, an 8-bit trigger-type word is forwarded to the destinations as well as (optionally) a 24-bit event ID. The trigger-type word is formed in the CTP and contains information on what gave rise to an L1A, while the 24-bit event ID is formed in the TTC VME interface (TTCvi). The TTC system can also transmit specific subdetector data and commands without introducing deadtime, e.g. test pulses when there are no bunches (LHC gap), and front-end parameters (e.g. delay values).
- During commissioning and for test and calibration runs, triggers can be injected locally in each TTC partition.

The specifications of the different components of the backbone system are given in Sections 16.2–16.4. A description of the functionality of the system is given in Section 16.5.

## 16.2 Interfaces

The TTC system has interfaces to the following: the LHC machine, the ATLAS CTP, the specific test and calibration logic of the subdetectors, and the front-end and readout electronics of the subdetectors (see Figure 16-1). These interfaces are described below.

### 16.2.1 Interface to the LHC machine

The LHC provides a continuously running clock signal (BC clock), although there are empty bunches as shown in Figure 16-3. The LHC also delivers an ORBIT signal at a fixed time within the LHC cycle allowing synchronization within this cycle. The BCR signal is derived from this ORBIT signal.

Both BC and ORBIT signals are fanned out electrically to the TTC partitions.

### 16.2.2 Interface to the central trigger processor

The CTP provides to the TTC:

- the L1A signal which has to be broadcast to all of the readout elements;
- a prepulse (PPS) signal that can be issued a predefined number of clock cycles before a test trigger is generated. This signal can be used to initiate the transmission of a test pulse to the front-end electronics at the correct time with respect to the arrival time of the test L1A;
- an 8-bit trigger-type word with each L1A. This word is transmitted to the readout elements and allows different types of events (e.g. physics events, test events, calibration events) to be distinguished.

All of these signals are electrically fanned out from the CTP crate to the TTC partitions.

### 16.2.3 Interface to subdetector test and calibration electronics

The TTC system receives from the subdetector-specific electronics, timing and trigger signals which can be used during commissioning, test and calibration periods. These signals are transmitted to the front-end and readout electronics as TTC commands.

### 16.2.4 Interface to the readout electronics

The TTC system provides to the readout electronics the following signals or data:

- the BC clock;
- the L1A signal;
- the BCR and ECR signals;
- the BCID and EVID words with each L1A;
- the trigger-type word with each L1A;
- commands and data, including test and calibration pulses.

The phase of the BC clock is adjustable in steps of 100 ps within 25 ns. L1A, BCR and commands can in addition be delayed by up to 15 BC clock cycles.

## 16.3 TTC backbone components

The TTC backbone components are:

- the TTC crate and fibre network [16-4];
- the TTC VME interface (TTCvi) [16-5];
- the TTC receiver chip (TTCrx) [16-6].

The TTC crate receives electrical signals from the TTCvi and the LHC machine, and performs the encoding and the electrical-to-optical conversion. An optical tree network with optical passive fan-outs distributes the encoded optical signals to up to 1024 destination nodes. These nodes consist of an optical-to-electrical conversion device followed by a receiver ASIC (TTCrx) which decodes the incoming frame and makes available all of the timing, trigger and control signals. The TTCvi is a VME module which provides the trigger and control signals to the TTC crate in the form of two encoded signals (A-channel and B-channel). This module allows the user to select the trigger source and to generate commands at known times. A TTC

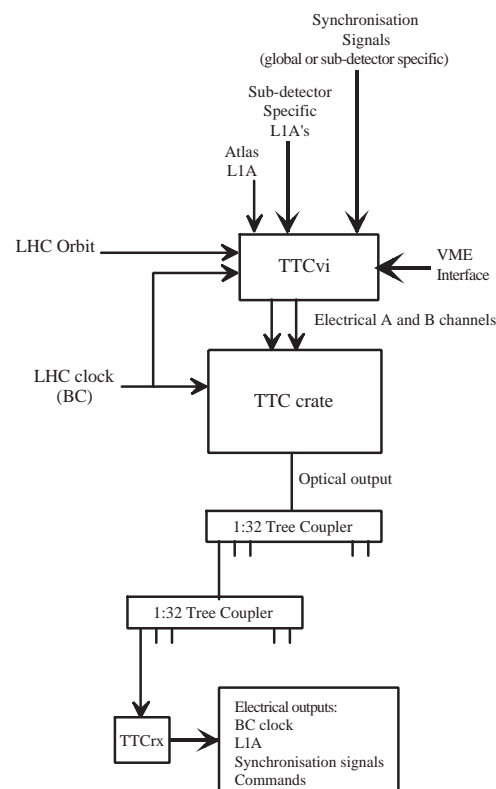


Figure 16-4 A TTC partition.

partition consists of a TTCvi, a TTC crate, an optical network and as many receivers as necessary. An example is shown in Figure 16-4.

### 16.3.1 The TTC crate

The TTC crate receives two electrical signals from the TTCvi and performs the final encoding and the electrical-to-optical conversion. The TTC crate includes a high-power laser transmitter incorporating a 1310 nm laser head module. At this wavelength the chromatic dispersion of normal optical fibre is negligible.

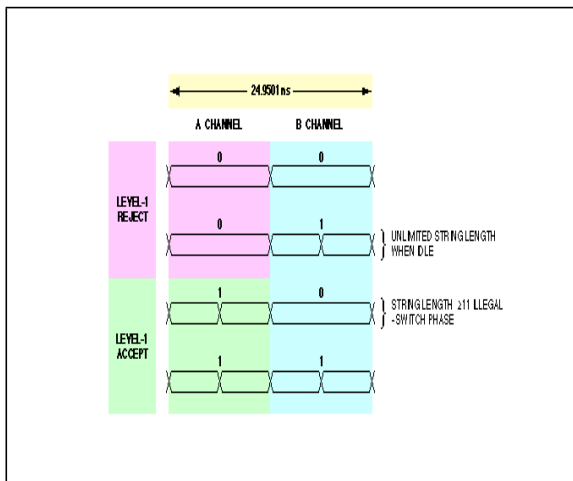


Figure 16-5 TTC two-channel encoding.

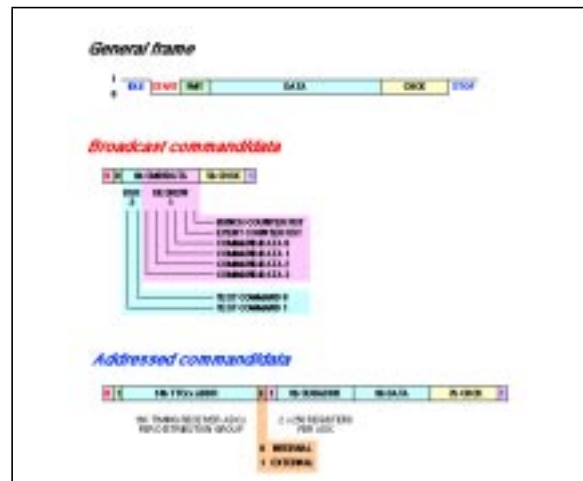


Figure 16-6 TTC encoding format.

The crate incorporates a 160.32 MBd biphasic mark encoder which time-division multiplexes the two input signals (A-channel and B-channel, Figure 16-5) using a balanced d.c.-free code. The primary phase-locked loop (PLL) plus encoder jitter when transmitting data is about 7 ps r.m.s.. A narrow-bandwidth PLL with a low-noise voltage controlled oscillator (VCXO) having low subharmonic feedthrough allows an r.m.s. output jitter of less than 10 ps to be maintained with an input clock reference jitter of several hundred picoseconds. The low-latency A-channel is dedicated to the transmission of the L1A signal. The B-channel transmits framed and formatted broadcast and individually-addressed commands and data using forward error-correction for high reliability. The format supports up to 16,000 receivers per distribution zone (partition), each of which may be associated with up to 256 internal and external subaddresses. Figure 16-6 gives the encoding format.

### 16.3.2 The TTC VME interface (TTCvi)

The complete specification of the TTCvi module is available in Ref. [16-5]. A simplified block diagram of this module is given in Figure 16-7.

The TTCvi delivers the A-channel and B-channel signals to the TTC transmitter crate. These two signals carry timing, trigger and control information.

The TTC A-channel is used only to transmit the L1A signal.




Figure 16-7 TTCvi simplified block diagram.

The TTC B-channel is used to transmit framed and formatted commands and data. These can be either:

- Short-format synchronous or asynchronous broadcast command/data cycles. If synchronous, the timing of these cycles relative to the LHC orbit is controlled precisely. Such cycles are used for the broadcasting of the BCR signal and for the transmission of other fast synchronous broadcast controls and test commands. These can be decoded at the receiving end (after the TTCrx) and used to produce test pulses, calibration pulses, etc.
- Long-format asynchronous individually-addressed or broadcast command/data cycles. The timing of these cycles with respect to the LHC orbit is indeterminate and they are not individually deskewed in the TTCrx ASICs. They are used for the transmission of parameters and non-time-critical commands. For instance, they are used to transmit the trigger type after each L1A signal.

The TTCvi is the master of a TTC partition. It is fully controlled through a VME interface by the partition user. It is at this level that the user decides to include this partition in the global ATLAS readout system or to run it in independent test or calibration runs. In this last case it uses subdetector-specific timing and trigger signals. When the partition is in the global ATLAS readout system the TTCvi is controlled by the DAQ run control.

### 16.3.2.1 Trigger inputs

Although in normal running mode the trigger input to the module is the L1A signal provided by the CTP, the TTCvi allows other trigger sources to be selected for test or calibration purposes without modifying the cabling. The main L1A sources are:

- The CTP L1A. The latency introduced by the module on the L1A coming from the CTP is minimized and is about 3 ns.
- Three additional local L1A inputs for subdetector-specific purposes.
- An internal random trigger generator. The number of L1A signals per unit of time follows a Poisson distribution with a mean rate programmable from about 1 Hz to 160 kHz.

### 16.3.2.2 ORBIT input

The ORBIT signal is a square wave of period 88.924 s which is received from the LHC machine and used for the generation of signals which are synchronized to the LHC orbit. Adjustment of the phase of the ORBIT signal permits a global control of the timing of the entire TTC system relative to the LHC bunch structure.

### 16.3.2.3 INHIBIT signals

In order to be able to send B-channel signals at known times within the LHC cycle, four independently-programmable timing signals called INHIBIT are generated within the TTCvi module. Each INHIBIT signal starts a programmable number of clock cycles after the start of the ORBIT signal and has a duration of a programmable number of clock cycles. Transmission of the associated synchronous command commences at the end of the INHIBIT signal duration.

One INHIBIT is used to initiate the transmission, during the long gap in the LHC cycle, of a broadcast command containing the BCR. The three other INHIBIT signals are available for the generation of other synchronous commands.

### 16.3.2.4 Generation of B-channel cycles

The TTCvi permits synchronous and asynchronous short- and long-format B-channel command/data cycles to be generated in a number of different ways:

#### **Short- and long-format asynchronous cycles.**

Asynchronous cycles may be initiated by writing the required data (a single byte for short-format or two 16-bit words for long-format) to specified TTCvi VME addresses. Normally, short-format cycles are used for broadcast commands or data, while long-format cycles are used for individually-addressed commands or data. However, a broadcast of 16 bits of data can be made with long-format cycles if TTCrx address zero is chosen. The timing of these cycles is not synchronized with the LHC orbit.

#### **Preloaded synchronous or asynchronous cycles.**

Four VME-addressable FIFOs are provided which may be preloaded with commands and data to be transmitted by B-channel cycles. For each of the four channels, the actual transmission of the preloaded information is initiated by a signal called B-Go. If synchronous mode is selected, the B-channel cycle is generated at the end of the corresponding INHIBIT signal. If

asynchronous mode is selected, the B-channel cycle is generated when the B-Go signal occurs. In addition, a special mode has been implemented to allow the download of a large amount of data (e.g. at the beginning of a run if the TTC system is used to download parameters). In this mode, the B-channel cycle is initiated as soon as the selected FIFO contains data.

#### Event-number and trigger-type cycle.

After each L1A is transmitted, the contents of a 24-bit event counter in the TTCvi can optionally be broadcast together with an 8-bit trigger-type word received from the CTP via a front panel connection. This broadcast is made asynchronously and takes about 4.4  $\mu$ s if the B-channel is free. Internally, the event number and the trigger type are stored in FIFOs to avoid any losses due to the random time of arrival of L1A.

### 16.3.3 TTC receiver chip (TTCrx)

The TTCrx (Figure 16-8) is a custom ASIC that receives control and synchronization information from the central TTC system (after optical-to-electrical conversion) and makes it available to the readout electronics. The TTCrx can be programmed to compensate for particle times of flight and for propagation delays associated with the detectors and their electronics.




Figure 16-8 TTCrx diagram.

One of the main functions of the TTCrx is to recover and distribute the 40.08 MHz BC clock with minimal jitter. It also makes available to the readout electronics the level-1 trigger-accept decision (L1A signal) and its associated BCID and EVID numbers.

Each TTCrx IC is identified in the distribution network by a unique 14-bit channel-ID number. The ASIC control logic identifies the A- and B-channels, deserializes the data in the B-channel and continuously monitors it to look for the presence of its channel-ID number.

The TTCrx IC has already been prototyped in a standard 1  $\mu$ m CMOS digital process. The ASIC footprint is 20 mm<sup>2</sup> and has been packaged in a 100-pin BGA package. A new version is being developed in a radiation-hard technology. A prototype should be available in spring 1999.



### 16.3.3.1 TTCrx characteristics

#### Bunch crossing

The BC clock is recovered with minimum jitter of the order of 40 ps r.m.s.. The recovered clock is then fed to a programmable fine deskew unit where two different clock phases are generated. The phases of the two clocks can be controlled independently in steps of 104 ps between 0 and 25 ns.

#### Level-1 accept, bunch-crossing identifier and event identifier

The TTCrx decodes the input frame and extracts the L1A signal with minimal latency (three BC clock cycles). With each L1A, the TTCrx delivers a 12-bit BCID number and a 24-bit EVID number on a 12-bit bus. Hence, three clock cycles are necessary for the readout electronics to get these numbers. The L1A signal can be delayed in the TTCrx by up to 15 BC clock cycles.

#### Bunch counter reset and event counter reset

The BCR and ECR are decoded in the chip and made available. The BCR signal can be delayed in the TTCrx by up to 15 BC clock cycles.

#### Broadcast commands

The other broadcast-command bits are made available as an 8-bit word. External decoding of these commands is necessary (for instance to generate a test pulse). The command word can be delayed by up to 16 BC clock cycles.

#### Set-up

There are four registers used to define the following delays in the TTCrx:

- fine deskew for BC clock 1 (0 to 25 ns with 104 ps step);
- fine deskew for BC clock 2 (0 to 25 ns with 104 ps step);
- L1A coarse deskew (0 to 15 BC clock cycles);
- command word (including BCR and ECR) coarse deskew (0 to 15 BC clock cycles).

These delays can be programmed via commands on the B-channel of the TTC network or through a local I<sup>2</sup>C interface for the new version.

#### Error detection and correction

Provision is made for error detection on the incoming frame, and internal error registers are incremented each time an error is detected. These error registers can be accessed through the I<sup>2</sup>C port.

## 16.4 Partitioning

As mentioned in Section 16.1, a TTC partition consists of a TTCvi, a TTC crate, an optical network and as many TTCrx chips as necessary (up to 1024). Any signal, data or commands transmitted on a particular partition reach only the TTCrx chips belonging to that partition. Partitioning the TTC system allows different subdetectors to work concurrently in the following instances:

- during test or calibration periods with independent timing and trigger signals;
- during the setting up of a physics run when the TTC network may be used to download parameters in the front-end and readout electronics.

Not only is it necessary to have independent partitions per subdetector, but also to have further subpartitions within each subdetector. This is for the following reasons:

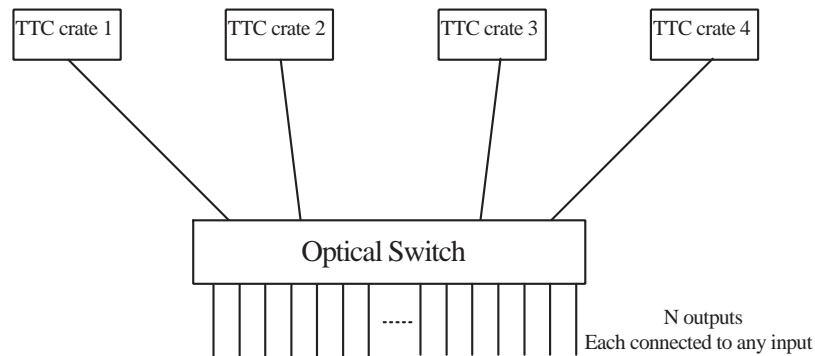
- During assembly and commissioning, a subdetector is divided into independent parts.
- One could wish to run different parts of a subdetector in different modes (e.g. test mode and calibration mode).
- The fibre length from the TTC crate to the readout electronics may vary a lot from one part of the detector to the other and hence the delay compensation foreseen in the TTCrx chip may not be sufficient.

The exact number of subpartitions needed per subdetector is still being discussed. Table 16-1 shows the numbers according to present assumptions. This model leads to a total of 35 TTC partitions.

**Table 16-1** Number of partitions per subdetector.

Subdetector	Number of partitions	Remarks
Pixels	2	
SCT	4	Barrel right & left, end-cap right & left
TRT	4	Barrel right & left, end-cap right & left
EM liquid-argon calorimeter	4	Barrel right & left, end-cap right & left
Hadronic end-cap calorimeter	2	End-cap right & left
Forward calorimeter	2	End-cap right & left
Tile calorimeter	4	Barrel right & left, extended barrel right & left
TGC	2	End-cap right & left
RPC	2	Barrel right & left
MDT	4	Barrel right & left, end-cap right & left
CSC	2	End-cap right & left
Level-1 calorimeter trigger	1	
Level-1 muon trigger	1	
Level-2/DAQ	1	

It has to be noted that the current model assumes a modest number of partitions per subdetector (of the order of four or fewer). If a subdetector wants more partitions, then the model may not be cost-effective and a specific solution will have to be found. One could, for instance, think of adding an optical switch in between the TTC crates and the distribution network, as shown in Figure 16-9.



**Figure 16-9** Possible use of an optical switch to get a high number of TTC partitions.

## 16.5 Functionality

This paragraph describes how the TTC system will be used, what functionality is available and how it is controlled. It does not touch upon the question of how the timing is set up; this is addressed in Chapter 19.

Figure 16-10 shows the different delay adjustments and their location in the TTC system.

### 16.5.1 Initialization

At the start of a run or a test or calibration period, one has to:

- define whether a particular TTC partition is to run independently or is to be part of the global ATLAS data-taking;
- select the signals to be used accordingly (clock, trigger source, etc.);
- configure necessary parameters in the TTCvi driving the TTC partition;
- set all the programmable delays in the TTCrx chips.

The first three operations are done through the VME interface of the TTCvi while the fourth one can be done either through the TTC network or through the I<sup>2</sup>C interface of the TTCrx chips. This last method is more convenient when the TTCrx chips are located at the ROD level. The TTCrx registers are loaded in the same way as the other registers of the readout chain (and can be read back and checked).

The initialization is performed within the framework of the DAQ/run-control system (and not the DCS) as, from the DAQ viewpoint, there is a direct correlation between a TTC partition and a DAQ partition.

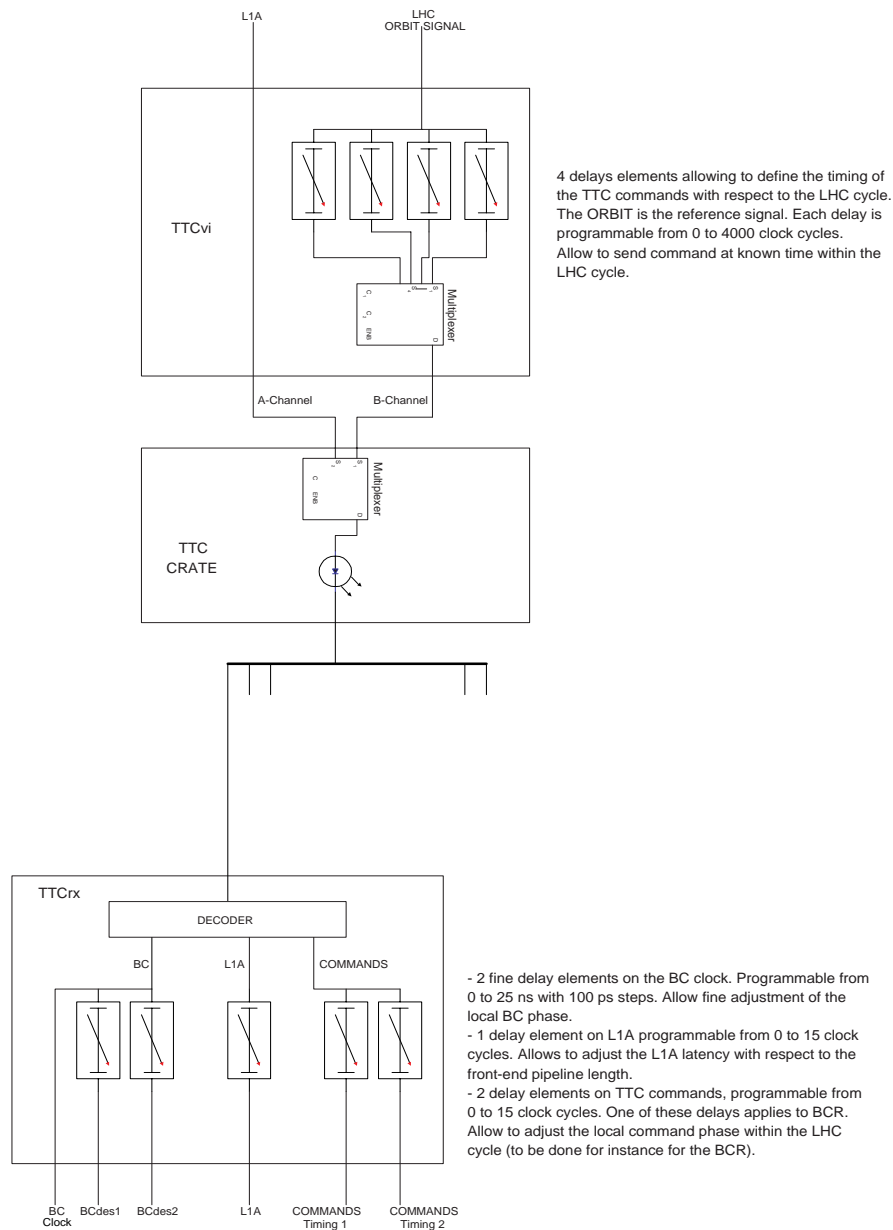


Figure 16-10 Different delay elements in the TTC system.

## 16.5.2 Synchronization

With regard to synchronization, one has to consider the following:

- synchronization with respect to the LHC cycle;
- maintaining a consistent BCID across the complete experiment;
- maintaining a consistent EVID across the complete experiment.

### 16.5.2.1 Synchronization within the LHC cycle

The LHC cycle is shown in Figure 16-3. The LHC machine provides a synchronization signal (ORBIT) which exhibits a transition edge at a known time in the cycle. This signal is used by the TTCvi to build up to four internal synchronization signals (INHIBIT). These signals allow the generation of commands (including BCR) at known times within the LHC cycle.

### 16.5.2.2 BCID synchronization

All the necessary tools to maintain a coherent BCID across the experiment are available. In the TTCvi, one of the INHIBIT signals is devoted to the BCR command. The phase of this command with respect to the LHC cycle is properly adjusted in the TTCvi. At the receiving end (TTCrx) one can locally adjust the phase of the BCR signal in the range 0–400 ns. This allows a misalignment of up to  $\pm 200$  ns between the TTCrx's of a given partition; this is deemed to be enough to cover differences in cable or fibre lengths or differences in time of flight in the detector. BCR will be issued at the end of each LHC cycle in order to reset the bunch-crossing counters in the TTCrx. BCID 0 will be allocated to the last BC of the 3.17  $\mu$ s gap at the end of the LHC cycle so that the LHC cycle starts with BCID 1. This choice may be subject to change if it gives rise to problems.

### 16.5.2.3 EVID synchronization

The ECR signal is necessary to reset the local event number counters if necessary (e.g. if a de-synchronization has been detected). The system must guarantee that there is no L1A during the time this signal is broadcast to all the TTCrx's. This can easily be achieved by introducing dead-time at the CTP level.

## 16.5.3 Test and calibration signals transmission

During test and calibration periods one has to be able to send pulses and triggers to the readout electronics. The TTCvi allows this to be done in different ways, without requiring any hardware intervention. In addition to the ATLAS L1A signal coming from the CTP, three other trigger inputs are available. The test signals can be transmitted as synchronous commands on the TTC and any one of the available B-Go signals can be used (in this case the relative phase between the test pulse and the trigger is to be adjusted externally), or sequences of commands with timing defined by the internal INHIBIT signals can be used. In this latter case the user has to decode the TTCrx command outputs to locally generate test and trigger signals. The INHIBIT signals are available as NIM signals on the TTCvi front panel and can be used to fire external logic if needed. The command output of the TTCrx can be resynchronized with one of the two deskewed clocks available (Bcdesk1 and Bcdesk2). This allows the readout electronics to have, for instance, a test signal with an adjustable timing with respect to the clock.

The prepulse (PPS) signal that can be issued by the CTP a known time before a random L1A is asserted, will be connected to one of the B-Go inputs and used to transmit a test pulse to the front-end electronics.

### 16.5.4 Trigger type transmission

It is necessary to have available, with each L1A, some information on the type of the trigger (for example, calibration, physics). This information may be used in the RODs. The TTCvi receives an 8-bit trigger-type word from the CTP and transmits it to the TTCrx. This transmission cannot be in phase with L1A, as there is not enough bandwidth available to transmit this word within five BC (the BC of the L1A followed by the four-BC dead time). The trigger type is transmitted in an asynchronous way when the TTC system is available (i.e. when there is no transmission of synchronous commands such as BCR). The trigger type is available at the output of the TTCrx 1–2  $\mu$ s after the L1A has been issued, assuming there were no L1As in the previous 2  $\mu$ s. If consecutive L1As occur within a short period of time, the trigger type transmission is queued and the trigger type of a given L1A will become available later at the output of the TTCrx. This is deemed to not be a problem as the trigger type is used only at the ROD level where the event data are subject to similar fluctuations in arrival time.

## 16.6 Latency

The TTC system introduces latency (delays) on the L1A signal at different places:

- The L1A signal has to be fanned out and transported from the CTP to the TTCvi.
- The L1A has to be transformed in the TTCvi to provide the A-channel signal to the TTC crate.
- The signal has to be transmitted (cable) from the TTCvi to the TTC crate and, in the TTC crate, it has to be encoded and converted from electrical to optical.
- The light has to travel through the optical network.
- At the receiving end, there is an optical-to-electrical conversion and the TTCrx has to decode the incoming frame and make available the L1A signal.

The time to transport the L1A signal from the CTP to the TTCvi has been estimated to be 50 ns, which includes a possible fan-out stage after the CTP and about 8 m of cables. This is based on the assumption that the TTCvi's are in the vicinity of the CTP. This scheme may not be optimal in terms of practicality as a lot of subdetector-specific electronics is connected to the TTCvi and subdetector groups may wish to have their TTC electronics close to their readout electronics.

The part of the latency due to the optical network depends on the fibre lengths and is very difficult to evaluate today. A worst-case length of 80 m has been taken. This corresponds to the maximum distance between any on-detector front-end electronics element and the furthest crate in the underground control room (USA15).

**Table 16-2** TTC latency.

Element	Latency (ns)
CTP to TTCvi	50
TTCvi	3
Cable to TTC crate	3
TTC crate	22
Fibre ( $\leq 80$ m)	400
TTCrx	75
<b>Total</b>	<b>553 (23 BC)</b>
<i>Total without cables and fibres</i>	<i>100 (4 BC)</i>

The final latency will have to be computed for each sub-system when the physical locations of the different elements are known. It should also include the latency introduced by a change of protocol on the line as is done, for instance, for the inner tracker.

The latency in the TTCvi-TTC crate has been optimized in three ways. First, there is no re-synchronization of the L1A signal with the BC clock used by the TTCvi; second, only three ECL stages are necessary for the TTCvi to produce the A-channel signal; and third, the phase of the BC clock used in the TTC crate is adjusted in such a way that its rising edge is properly set with respect to the L1A phase. The total latency of the TTC system is summarized in Table 16-2.

## 16.7 Test and monitoring

There is a need for a global timing adjustment of the experiment (Chapter 19) which has to be done regularly and at least before each run. This timing adjustment sequence is a very good test of the TTC network itself as it requires the transmission to be fully efficient and all the timing characteristics of the delivered signals to be good.

The monitoring of the system is done in the following ways:

- The information carried by the TTC is part of the event readout data (BCID number, EVID number, trigger-type word) and this information is checked at every stage of the readout chain. This allows the basic transmission on the TTC network to be monitored.
- The TTC includes an error detection mechanism and each TTCrx maintains an error register. These registers will be read out regularly.
- The timing characteristics of the delivered signals (jitter on the BC clock, delay stability, etc.) will require a fine analysis of the readout data.

The test of the TTC system is done via VME for the TTCvi and via the I<sup>2</sup>C interface for the TTCrx.

## 16.8 Experience from prototyping

Each part of the TTC system has been prototyped and tested and some experience acquired. The system works, but there are areas where performance needs improvement. The main problems discovered up to now concern the receiver chip, for which some nonlinearity in the delay elements and some sensitivity to temperature variations have been discovered. These problems have been solved either by a design modification or by the development of a special small ball-grid-array (BGA) package.

There is still a lack of experience in the use of the system in a readout environment, but this will be addressed during the years 1998 and 1999 as most of the subdetector groups are planning to use this system for beam tests.

## 16.9 Construction, assembly procedure and schedule

As mentioned in Section 16.8, the prototype of the full system has not yet been validated by extensive use in a real data-taking environment. For this reason the specifications of the different parts of the system are not yet finalized. It is expected that after two years of utilization in the test-beam environment, the final specifications will be issued and the final design will start in 2000. The production of the TTC crate and TTCvi will then be subcontracted to external firms as the total amount of modules or elements needed is high.

The TTCrx chips will be ordered centrally by CERN from the selected foundry and made available to the community.

A schedule for the TTC system production is shown in Figure 16-11.




Figure 16-11 Schedule for the TTC system production.

## 16.10 Quality assurance and review procedures

A set of procedures and rules must be followed to assure the required quality of the system.

The design and specification of each element will be subject to at least two review exercises — a Preliminary Design Review (PDR) and a Final Design Review (FDR) — with the possibility of an Interim Design Review (IDR) between them.



The PDR will be used to examine and assess the full requirements and specifications for the subsystem element, to identify any missing functionality, to ensure full compatibility with all connecting subsystems and to determine overall feasibility. This is perhaps the most important part of the review procedure, as it will determine the direction of subsequent engineering effort. Detailed written specifications will be supplied to the review group two weeks in advance of the review itself, and following the review the final agreed conclusions will be distributed to the community.

IDRs may be held at any time during the design phase, but most usefully at the completion of schematic capture when many engineering issues (timing margins, interfaces to other subsystems, detailed latency calculations, etc.) can be explored. By monitoring progress at this stage some potential problems may be detected and resolved early, thereby minimizing wasted effort.

The FDR will be held before the module design is sent for manufacture, and is intended to catch any design or engineering errors before budgetary commitment. This review will necessarily be of a more technical nature than the initial review, but there should be few problems to detect by this stage. It will be merged with the ATLAS production readiness review (PRR).

This review work has already been done for the prototyping. A User Requirement Document (URD) for the TTC system has been written and reviewed for a first time by representatives of each subsystem.

Prototypes of each element exist and will be used by the subsystems in the coming year. During that time, the final specifications will be written. The PDR will review them as well as the performance and functionality of the prototype system.

The construction of the TTC crates and of the TTCvi will follow usual rules, namely:

- Electrical tests of the PCB before assembly.
- Burn in of the boards after assembly by leaving them under power without cooling for two days.
- Test of the full boards.

Each element will receive a serial number and a database will be set up to store the information relative to each board: origin of the components which populate the board, results of tests, history of failure and location.

In order to facilitate the test and the maintenance of the boards, the JTAG test bus will be used as much as possible.

## 16.11 References

- 16-1     *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ-NO-103, June 1998.  
<http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps>
- 16-2     *Timing, Trigger and Control System (TTC) (User) Requirements Document (Draft version 0.2)*, ATLAS working document, ATL-DA-ES-0004, February 1998.
- 16-3     *RD12 Status Report*, CERN/LHCC/97-29, April 1997.

- 16-4     TTC crate information, RD12 web page.  
          <http://www.cern.ch/TTC/intro.html#Transmitters>
- 16-5     *TTC-VMEbus Interface: TTCvi, Revision 1.5*, RD12 working document, November 1997.  
          <http://www.cern.ch/TTC/TTCviSpec.pdf>
- 16-6     *TTCrx Reference Manual, version 2.2*, RD12 working document, July 1997.  
          <http://pcvlsi5.cern.ch:80/MicDig/ttc/MANUAL22.PDF>