ATLAS SCT Clock and Control

R.B.Nickerson, A.R.Weidberg
Oxford University, 15 May 1998

Content

This document is intended to draw together existing information on the use of the SCT clock and control system. It describes the hardware system in outline only.

It is to be anticipated that this is a precursor to final documentation as there remain details that are not yet fixed.

Introduction

The ATLAS wide clock and control signals are distributed with the system developed by Taylor et. al.[1] , the TTC system. A fibre from this system arrives at each SCT readout crate which has modules which require clock and control. The signals are internally distributed within the crates, probably on the back-plane[2]. The subset of these signals which are transmitted to the front-end electronics are distributed using a system based on the low-speed fibre optic links (40Mbit/s). All clock and control information, including L-1 triggers, are transmitted to each module in the detector down a single such fibre link. VCSELs and a bi-phase mark encoding scheme are used[3].

Since a single serial link is used for clock, slow and fast control there are restrictions on the simultaneous use of front-end features[4]. The only restrictions which directly affect clock and control are

a) L1 Triggers cannot be closer than 3 bunch crossings (BCs) and

b) fast controls must be separated by 5BCs from each other (except for calibration pulses) and by 3BCs after a L1 trigger or 5BCs before one.

Since L1 triggers must be separated by more than 3BCs for other reasons, the only real impact on clock and control is that fast resynchronisation triggers from the TTC cannot be issued within 3BCs of a L1 trigger. All timing delays associated with serial encoding of the L1 and fast control signals are accounted for in the SCT clock and control system: The different latency for a fast reset, bunch and L1 counters and a L1 trigger are allowed for internally so that the TTC system is required only to issue fast counter resets (and other fast controls) with the same time relationship to the relevant BC as the L1 triggers have.

Time Synchronisation & Cosmic Ray Data Taking

.There are four basic issues in synchronisation during normal data taking.

•Synchronisation of local clocks at each detector element with the particle arrival times. This adjustment is over one BC, 25ns, with 1ns accuracy required[6].

•Internal synchronisation of L1 triggers across all the detector elements so that each element reads the same BC in response to a TTC L1 trigger. As the triggers are encoded in the same signal as the clock in the SCT system, this is a requirement to correctly identify the clock cycle, *ie* with a resolution of 25ns.

- Synchronisation of the L1 trigger signals delivered by the TTC with the correct BC in the detector. Since this signal also arrives encoded with the signal which will be decoded into the local clock it also requires a resolution of 1BC, *ie* 25ns.

- Monitoring of synchronisation drifts.

In addition there is the question of how to provide correct timing of the various detector elements during cosmic ray data taking prior to LHC turn on.

These issues are all addressed by the SCT clock and control system.

Overview of Hardware Scheme

Figure 1 shows the layout of the scheme. It should be noted that the distribution of functionality between modules is indicative only; the final designs are not agreed.

The SCT TTC crate delivers TTC information via one fibre to each SCT R/O crate. The crate TTC receiver module decodes this information and passes the relevant subset along the back-plane to the modules in the crate. The backplane is designed to ensure all the clocks within a crate are synchronised. This module also acts to allow stand-alone operation without the TTC system and injection of slow control information from the local computer. The module also has the ability to delay the fast signals by integral numbers of beam crossings. This is where the overall TTC adjustment resides in the SCT system.

The SCT-relevant TTC information on the back plane is received by the clock and control distribution boards (which may be part of the RODs) on which they are encoded into the biphase mark and transmitted to the front-ends. There are 4088 such fibres in the system. At the receive end, the clock is reconstructed and the control data stream separated (but there is no decoding of the data stream). The 40MHz clock and the data is distributed to a single module electrically.

Each module must receive the clock at the appropriate time. The time relative to the machine RF, τ_i, will depend on many factors such as location with respect to the interaction point and propagation delay through the electronics. The scheme of French et al[5]. will be employed to partially addresses the issue of how to cross-synchronise elements within a detector. This works by arranging neighbouring modules to operate either from its own clock, or from the clock that belongs to its neighbour. In ‘clock through mode’ the clock (divided by two) is returned along the data link so that by comparing the phase of the returned clock and the returned neighbour clock, the relative synchronisation of neighbour clocks can be determined and adjusted to be correct. By arranging neighbours in a chain, large groups of modules can be internally self-consistent in terms of synchronisation. In the case of the silicon barrel, the current plan calls for groups of 24 modules to be so connected. Vastly larger numbers would not be useful due to the build up of errors.

Cross synchronisation of the groups of measured neighbours will be given by measurements of propagation delays made during assembly of the detector. Assuming the fibre lengths can be controlled to 5 cm and that the uncertainty in the propagation delays through the VCSEL and PIN (plus the driver and receiver electronics) is less than 4 ns, all the clocks will arrive within 5 ns of the correct time (±2.5 ns) Assuming the errors are random, yields one nanosecond accuracy for the average propagation delay for each group of 24 neighbours, which is sufficient. This procedure should get the synchronisation close to correct: the fine tuning will be done using a different method.

The jitter in the arrival time of the clock relative to the transmitted clock has been measured for a prototype link. For the particular system measured the jitter was less than XXX ps full width at the base of the distribution (sigma less than xxx ps), which is more than adequate. Over a period of several hours the drift in transmission time over a 50m long link was <xxx ps. Both the jitter and longer term drift are functions of the electronics design; the systems tested were simple and no attempt had been made to control these parameters. Considerable improvement is to be expected in the final electronics. The expansion coefficient of glass is typically one part in 10^5, so that for 50m of fibre a temperature change of 10C results in a 25 ps change in transmission time. The temperature dependence of the refractive index is similarly small. These are encouraging numbers for the SCT application.

Synchronisation Procedure

a) **Synchronisation of Local clocks to Particle Crossing Times**

This procedure is automated and under the control of the local SCT computer via the local bus. It does not require the TTC system. It is performed infrequently, at the instigation of an expert operator who has determined a need to do so.

- The front-end chips are switched into clock through mode. (All at once)
- The RODs measure the phase of the clocks returned via the data links relative to the local copy of the system clock. This information is reported to the local computer.
- The front-end chips are then instructed to use the neighbour clock (all at once)
- The RODs measure the phase of the clocks returned via the data links relative to the local copy of the system clock. This information is reported to the local computer.
- The SCT computer evaluates corrections based on the known geometry of the detector, the relative phases of clock and neighbour clock, and any ‘bad channel’ data base. The corrections are loaded into the SCT clock and control distribution modules, and reported to central DAQ for use during normal data taking.
- The procedure is repeated until it converges (The delay step size is not necessarily well controlled, so iteration may be required).
- Average shifts for each group of 24 modules (barrel) are evaluated and the mean overall group delay calculated. (Used in next phase of process)

Estimated time for this procedure is a minute or two.

b) **Adjustment of local copies of L1 Trigger to all read the Same Event**

The procedure in a) results in all front-end clocks being correctly delayed relative to the clock in the distribution modules, ie with respect to the copy of the TTC clock within one SCT crate. It also provides the mean delay for each group of neighbour modules relative to the clock in the distribution modules.

- Using the data base information on propagation delays for each link, and the mean shifts calculated in a), the SCT computer evaluates an integral number of clock cycles delay to be added to each distributed clock to ensure that all the modules supplied from an SCT crate read the same BC.
- These are loaded into the crate timing distribution modules for each copy of the clock sent to individual modules in the crate. OR they are loaded into the individual modules.

Note that further internal synchronisation can be achieved only by inter-crate communication between the crate timing distribution modules, or *by calibration of the TTC fibres* so that they are all synchronised well within one BC. The latter is assumed, inter-crate communication of L1 timing is not envisioned. (The TTC fibres do not need to be the same length, but the lengths need to be in the data base along with the lengths of the fibres which go down to the front-ends).

This procedure should take a minute or two.

c) **Adjustment of TTC L1 Trigger to read the correct BC**

After procedures a) and b) above, there is only one time left to adjust, namely that of the TTC L1 trigger relative to the beam crossing it is intended to identify. This can be accomplished using the methodology of Taylor et al.[6] the essence of which is as listed below (The exact details may be found in reference 6). Note that this will require separate adjustment of each crate in the SCT system since each has a separate TTC signal. However since we assume we have calibrated the TTC fibre lengths, the adjustments should be the same to within a few ns for all the crates. To identify the correct crossing, all the data can be used simultaneously, simplifying the process.

- The overall synchronisation for each SCT crate takes advantage of bunch structure in the machine. In particular the abort gap has an extended period of empty bunches followed by full bunches. In the case of the SCT system, the easiest synchronisation procedure is probably as follows:
- The TTC system is set to deliver L1 triggers corresponding to the first full bunch after the abort gap.
- The local SCT system scans the trigger distribution latency delay from the earliest possible time in full BC steps (ie *25 ns*). At each delay position events are collected. The number required depends on the machine luminosity: at 1% of full luminosity and higher, one event is sufficient. The data is read into the local SCT computer.
- As the timing is correct at this stage, except for the BC identification, the data will show a jump in mean occupancy at the transition from empty to full BCs at the abort gap boundary. This jump will be identified by the local computer and the required local latency for each crate recorded in the appropriate data base.
- BC counter and L1 counter synchronisation is then done by receiving the appropriate command from the TTC system synchronised with the BC that commands must be executed in.

This procedure assumes the TTC timing is pre-adjusted to the correct value, which is possible for the SCT because it does not participate in the L1 trigger. This is a considerable simplification over the calorimeter or muon system

Final Adjustments

Procedures a) - d) above should result in all the modules being timed to within a couple of ns of the correct value. For cosmic running no further adjustment is possible. For real data taking, there is enough data to fine tune each clock on the basis of the efficiency plates for the detectors. This will be executed by ‘expert’ software on the basis of experience gained during the early running. The process will probably involve off-line analysis since the ideal timing will be different according to where the particle traverses the detector. In order to examine the extremes within a detector, the tracks will need to be reconstructed and the plates examined according to area of the detector, angle of track etc.

100 events times 1/(occupancy) (so 10k for 1% and 100k for 0.1%) are taken with beam in the machine with the timing at each of 25 settings separated by 0.5ns about the nominal value.

It is assumed that this can be done in a mode where only the SCT is written to tape, and it is written after the L1 trigger. The total data sample size is then about 50Gbytes and could be acquired in a few minutes.

This procedure would be automated and executed at the start of every nth data taking run, where n would be chosen on the basis of experience. Any adjustments required would be done ‘by hand’ after expert consideration.

d) **Monitoring**

Real-time monitoring is achieved using the clock feed through mode. A data base is retained which records the time of the returned clock relative to the local crate clock for each fibre. At the start of each run, the system is switched into clock feed through mode and the time of each returned clock recorded. It is anticipated that these numbers will always be within a few hundred ps of the data base value. Any fibre which fails to satisfy this condition, or shows a time dependent drift will be drawn to the attention of the shift crew by ‘expert’ software.

If there was significant deviation from the correct value, the procedure described above would need to be repeated. No conditions have been identified which would cause this to be required, so it is anticipated that re-synchronising would be a rare event.

In addition to this, off-line monitoring of strip efficiencies would be used to detect any more subtle drifts. It is to be anticipated that the efficiency would be affected fastest in certain geometrical regions of each module and these would be continuously monitored for any signs of systematic drift in efficiency.