

Timing Distribution at the LHC

B.G. Taylor

CERN, 1211 Geneva 23, Switzerland
Bruce.Taylor@cern.ch

for the RD12 Collaboration

Abstract

This paper describes a unified approach to fast timing distribution at the LHC. The timing signals for each ring of the machine will be encoded and transmitted over optical links from the RF system to the PCR, where beam-synchronous messages will be added. High power laser transmitters will then broadcast the signals over singlemode optical fibres to the four LHC experiments, to the test beam areas, and to the beam instrumentation located around the LHC ring and on the SPS transfer lines. At the experiment areas, trigger information and local synchronous commands and data will be added. The regenerated signals will then be broadcast over multimode passive optical networks to several thousand destinations.

I. INTRODUCTION

Ever since the invention of the synchro-cyclotron in 1946, timing and synchronisation have played a key role in the field of particle accelerators. At the LHC, “fast” timing signals must be distributed to all the experiments and to the beam instrumentation of the machine. These signals are derived from the LHC RF generators and will be synchronous with the circulating beams, so that their frequencies will vary a little during acceleration. At 7 TeV, the bunch clock frequency will be about 40.07897 MHz while the orbit frequency will be 11.2455 kHz. They are distinct from the “slow” LHC timing signals, having a granularity of 1 ms, which will signal machine events and distribute UTC time for data tagging and post mortem applications.

The 1917 Sopwith Camel (Fig. 1) illustrates a common attribute of timing and synchronisation systems. Although the Camel was notoriously difficult to fly, it was one of the most successful fighter aircraft of its day. Much of that success can be attributed to the use of twin Vickers machine guns, synchronised to fire at 500 rounds/min between the blades of the propeller rotating in front of them. In the event of a timing error, a few hits on the wooden blades were sufficient for the plane’s own propeller to be shot away.

This problem illustrates a typical characteristic of timing systems: while they do not themselves do anything particularly exciting, if they do go wrong they can lead to quite a disaster! At the LHC, distributing correctly synchronised signals to several thousand electronics channels presents some interesting challenges and this paper describes briefly how these have been addressed. The work has been done in the framework of the RD12 TTC Project [1], which comprised members from all the LHC experiments, the Microelectronics and Beam Instrumentation Groups and two industrial partners.



Figure 1: Sync or swim – the Sopwith Camel

Table 1 gives a little glossary of some TTC terms. This is not a comprehensive list of all the system components that were developed in the project, as some items like the subminiature optical fibre connectors, encoder and VCXO/PLL modules were not given special TTC names. The TTCpr PMC module was developed by Argonne National Laboratory.

TTC	Timing, Trigger and Control
TTCbi	Beam instrumentation interface
TTCcf	Clocks fanout
TTCex	Laser encoder/transmitter
TTClc	Laser controller
TTCmi	LHC machine interface
TTCmx	Laser minitransmitter
TTCoc	Optical tree coupler
TTCos	Orbit synchroniser
TTCpr	PMC receiver
TTCrm	Receiver mezzanine
TTCrx	Receiver ASIC
TTCsr	Simple receiver
TTCtx	Laser transmitter
TTCvi	VMEbus interface
TTCvr	VMEbus receiver
TTCvx	LED transmitter

Table 1: TTC glossary

The unified approach to TTC distribution developed by RD12 provides for the broadcasting of the fast timing signals through all the transmission stages from the RF generators of the LHC machine to the outputs of the timing receiver ASICs at the experiment and beam instrumentation destinations. That general path will be followed in this review of the system.

II. LHC BUNCH STRUCTURE

Commencing with the timing of the LHC machine, it should first be noted that there has been an important change to the bunch structure described in the Yellow Book [2]. Initially it was proposed to accelerate trains of 84 bunches in the PS, 81 of which would be injected into the SPS, 3 being lost in the PS ejector. The difficulty with this configuration is that it would be dirty in the PS and SPS machines and there would be longitudinal stability problems in the PS with the 84-bunch trains.

Various solutions to these problems have been proposed and the one that has been retained as the current baseline foresees the acceleration of PS trains of 72 bunches which will be entirely injected into the SPS.




Figure 2: Revised LHC bunch structure

In order to maintain an acceptable filling factor in the LHC with the 72-bunch PS trains, the SPS batches which will be injected into the LHC will comprise groups of 3, 3 and then 4 PS trains (see Fig. 2). So whereas formerly we had quasi 12-fold symmetry in the LHC (the last PS train in the last SPS batch being suppressed to allow for the risetime of the LHC extraction kicker), we now have quasi 4-fold symmetry with corresponding implications for the TTC synchronisation algorithms [3].

As a result of this change the number of bunch crossings per orbit will be reduced from 2835 to 2808. Note that this applies only to ATLAS and CMS – since these experiments are diametrically opposite each other it is possible to phase the beams to make the LHC extractor gaps coincide at both of them. That is not possible at ALICE and LHCb, which will result in the loss of a further 188 bunch crossings per orbit at these experiments. It should also be noted that it is now expected that there may be quite a substantial initial running period with 75 ns instead of 25 ns bunch spacing. The expected rms collision length remains about 180 ps.

In the case of the LHC bunch structure for heavy ions, there may be several resynchronisations during each orbit and there could be gaps which are a non-integer number of bunch intervals in length. But neither of these factors should be a cause for alarm, for in this case the bunch spacing concerned will be 100 ns or 125 ns. The TTC system will continue to distribute a 40.079 MHz clock during this mode of operation and the bunch crossings will remain in phase with this clock.

III. CLOCK ARTEFACTS

On the other hand, at times there may be some artefacts in the distributed clocks. There could be a 1 ms hole in the SPS

RF/5 clock occurring once before each transfer from the PS, because the LHC machine will be the master of the timing and the SPS has to be synchronised such that the SPS batches are injected into the correct part of the LHC orbit. That will be done by calculating back to the PS, so the PS and SPS have to be resynchronised before each injection and during this time there may be an interruption in the clock.

The situation with colliding beams at the LHC will be more comfortable. In that case there could be a 1 ms hole in the bunch clock which will occur once, and once only, before the very first injection from the SPS into the LHC. The reason for this is that a general RF system reset will be made prior to each LHC run in order to ensure that all the dividers have the correct phase and there may be an interruption to the clock while this is applied.

During these clock holes, the TTC system will continue to distribute a 40.079 MHz clock to the experiments. But developers should be aware that there may be a momentary phase perturbation when the system resynchronises with the real clocks when they are restored after the interruptions.

IV. DISTRIBUTION ARCHITECTURE

At present the RF timing generators are located in the BA3 Faraday Cage adjacent to the Preveessin Control Room (PCR). Four clocks are available: the constant frequency 40.079 MHz LHC bunch clock, a pseudo LHC orbit signal obtained by dividing the clock by 3564, the real SPS orbit signal and the ramping SPS 40 MHz clock obtained by dividing the SPS RF by 5. The PCR transmitters can each broadcast only one orbit and one clock signal simultaneously. The selected pair are encoded and used to modulate a high power laser, the output from which is split by a 1:32 optical tree coupler and broadcast via optical fibre to different destinations around the CERN sites. At present these destinations include the test beam areas in the North and West halls and labs where beam instrumentation and TTC development work are being done. Finally they will include the LHC experiment areas.




Figure 3: OTDR PCR – B4

At the experiment areas the signals will be received by a TTC machine interface (TTCmi) in which they are decoded,

adjusted in phase and cleaned up by a VCXO/PLL before multiple copies of the two clocks are distributed to the local TTC transmitters which broadcast the timing at the experiments themselves.




Figure 4: Overall block diagram of TTC distribution

The optical part of the distribution system operates at 1310 nm since very low jitter is required and the chromatic dispersion of normal optical fibre is negligible at that wavelength. That compares with a dispersion of about 80 ps/nm.km at the short wavelength window at 830 nm. Hence at 1310 nm one can use not only laser transmitters but also LED transmitters in low-cost test setups. There are other reasons why operation at 1310 nm is attractive. Laser safety requirements are less stringent at that wavelength and there is some evidence that irradiation-induced fibre attenuation is less than at 830 nm [4]. In addition, fibre having a given multimode dispersion at 1310 nm is less expensive than for the same multimode dispersion at the short wavelength.

Initial tests were carried out with fibres laid from the PCR through the ATLAS surface building and the CERN telephone exchange to Building 4. The line of sight distance for this path is only 2.5 km but the fibres were laid around the SPS ring and their total length is 6.5 km (see Fig. 3). The tests were carried out with both singlemode and multimode fibre, as a result of which it was decided to use singlemode fibre for broadcasting the signals from the PCR to the experiment areas

while retaining multimode fibre for distributing the TTC signals at the experiments themselves [5]. This is primarily because a large number of optical tree couplers are required at the experiments and multimode couplers are cheaper than singlemode couplers by a factor of 2 - 3. The transmission distances at the experiments are only of the order of 100 m so that multimode dispersion is not a serious problem there.

In the final configuration (Fig. 4), the timing generators will be relocated in a new Faraday Cage to be constructed at SR4. It has been decided to retain the high-power transmitters at the PCR, so the signals will be sent from SR4 to the PCR by an optical link using the standard TTC system components and existing Sumitomo phase-stabilised fibres left over from LEP. As only one fibre is required for the encoded signals using the RD12 protocol, spare fibres are available. From SR4 the fibre cable follows the old railway line before cutting across country to the PCR and its total length is 9.5 km.

Phase stability measurements have been made for a period of several weeks over the 28.6 km fibre loop from Building 4 through SR1 and the PCR up to SR4 and back again. They indicate that during a typical LHC run it should not be necessary to make adjustments of more than one or two of the TTCrx timing receiver fine deskew steps of 100 ps. The winter-summer seasonal variation will be much greater than this but, as it is very slow, compensation will be easily made.

V. PCR TTC TRANSMITTERS

Fig. 5 shows three of the four high-power TTC transmitters that have been installed at the PCR. One is for the SPS and transfer lines and two are for the LHC rings, which can have slightly different radio frequencies during certain beam steering operations. The fourth transmitter is a spare and a fifth transmitter will be installed shortly as an additional reserve. The outputs from the root optical tree couplers in these transmitters are connected to the main fibre optic patch panel in the adjacent TDM room, where they can be manual patched through to many destinations around the CERN sites.

The signal encoding and formatting [6] were chosen specifically to meet the needs of the TTC system. Two TDM channels are transmitted. The A channel has very low latency and is used for broadcasting the Level-1 trigger accepts at the experiments and the LHC orbit signal from the PCR. The B channel provides for the transmission of framed and formatted commands and data. These can be in the form of short-format broadcast commands or longer format individually addressed or broadcast commands and data. The 2-channel system means that it is not necessary to interfere with the flow of time-critical signals on the A channel in order to transmit information on the B channel.

Biphase mark encoding is used as an appropriate compromise between the conflicting demands of high channel efficiency and low systematic jitter. At 160.316 MBaud, 4 symbols can be broadcast in any 25 ns bunch crossing interval. There are no restrictions on the transmission of 3 of these symbols, but if a lengthy sequence of the 4th symbol is received the TTCrx detects that there is an A/B channel phase error and can correct it automatically. Fig. 6 shows the performance of the encoders in the PCR transmitters. The

histogram indicates an rms jitter of 7 ps relative to the 40.079 MHz clock input to the transmitter while broadcasting PRBS data at the maximum rate on both channels.




Figure 5: PCR TTC transmitters

The first real tests of the system were carried out in May 2000 and October/November 2001, when LHC-structured test beams were delivered by the SPS to the North and West areas. During these periods the TTC system distributed a constant frequency 40.079 MHz clock and after each acceleration ramp the SPS was resynchronised to that clock before the beams were extracted to the test beam areas.

Thus the SPS was operated in the same way that it will be when it is used as an injector for the LHC machine [7]. On the other hand the orbit signal that was distributed was the real

SPS one, which swings about 29 Hz during acceleration from 25 to 450 GeV. Metastability is avoided by synchronisers at the TTC transmitters.




Figure 6: PCR transmitter encoder jitter

VI. TTCMI MACHINE INTERFACE

To receive these signals a number of TTC machine interface (TTCmi) crates (see Fig. 7) have been manufactured. The ATLAS TTCmi is installed at H8 in the North area, where it was used by the inner detector groups in 2000 and also shared with the TileCal team in 2001. The main CMS TTCmi is used by the tracker group at X5 in the West area, while additional systems have been used by the muon group in the GIF area and at H2 and H4. The LHCb TTCmi is installed at X7 and additional reduced TTCmi crates have been provided to the beam instrumentation group in SL Division and the ESS support group in EP. Since it is essential that the TTCmi be operational for an LHC experiment to take data, a complete spare has been made for each of them. The spare systems will be installed and run online continuously, fed via separate optical fibres from the PCR to provide an additional measure of redundancy.



Figure 7: TTCmi minicrate

The timing distribution system performed very satisfactorily during all the test beam runs. Using a TDC with

60 ps binning, the CMS test beam monitor group at X5 measured a bunch length of 2.3 ns, which compares with 2.5 ns expected by the SPS machine specialists [8].

The TTCmi minicrates provide a standardised interface between the timing signals broadcast from the PCR and the local TTC distribution systems at the experiments. A detailed user manual is available [9]. 14 TTCmi have been produced, some of them rather economically by upgrading old TTC transmitter minicrates.





Figure 8: TTCmi block diagram

A block diagram of the TTCmi is shown in Fig. 8. The power of the optical signal received from the PCR is adjusted to the optimum level of about -20 dBm by a local attenuator. The signal is received by an integrated photodiode/preamp and decoded by a TTCrx timing receiver ASIC. The orbit signal is passed through a global phase adjuster which can set the phase in 3564 steps of the 25 ns bunch crossing interval, permitting phase compensation for any location of the TTCmi around the LHC ring.

The TTCrx delivers a raw 40.079 MHz bunch clock with an rms jitter of about 80 ps, which is rather large for the primary timing reference for an LHC experiment. So the TTCmi incorporates a low-noise 160.316 MHz VCXO in a PLL with a narrow loop bandwidth which cleans up the clock and delivers a 40.079 MHz signal with an rms jitter of about 7 ps. Multiple copies of the two clock signals are output electrically. It is also possible to plug optical transmitters into the TTCmi to broadcast either the bunch clock or the encoded TTC signal if it is used in conjunction with an optional TTCvi.

By having a PCR transmitter and TTCmi co-located, the performance of the system has been measured around a 13 km fibre loop, comparing the jitter of the TTCmi clock output with the clock input to the transmitter (see Fig. 9). During the SPS runs with LHC-structured test beams, special tests were carried out over a period of several days to verify that no phase slips occurred in any of the multiple PLLs in the complete timing distribution system.



- Trace A:* Biphas mark encoded output from PCR transmitter. The central transitions represent the LHC orbit signal, occurring once per 3564 bunch clock periods.
- Trace B:* 200 ps/div window on bunch clock recovered by the TTCrx.
- Trace C:* Recovered bunch clock after cleanup by PLL in TTCmi.
- Trace D:* 200 ps/div window on Trace C (40.079 MHz bunch clock).
- Histogram E:* Jitter of recovered bunch clock after PLL cleanup. (7 ps rms).

Figure 9: TTCmi performance (13 km fibre)

A. LHCrx

The LHCrx module (see Fig. 10) is used in the TTCmi for receiving the signals from the PCR transmitters and at these transmitters for receiving the timing signals from SR4. It incorporates an independent monitor receiver for diagnostic purposes. A total of 20 LHCrx have been made and they have all been updated with the latest version of the TTCrx which has a watchdog circuit to relock the PLL automatically after an optical signal interruption.

The TTCrx internal coarse deskew circuits provide for phase adjustment over a range of 16 bunch-crossing intervals (about 400 ns). This is sufficient to compensate for the differences in particle time-of-flight and optical fibre path length between different destinations at one experiment but not for the phase differences in the orbit signal received from the PCR at different points around the LHC ring.

The orbit phase adjustment provided in the LHCrx module allows the phase to be adjusted in 3564 steps of 25 ns. The local BPTX beam monitors at each experiment will allow the phases of the circulating beams to be observed by noting the locations of the large extractor gaps. It is then a simple matter to set the LHCrx adjustment to obtain the orbit signal phase required to bring the destination TTCrx bunch counters within the adjustment range provided by their internal deskew circuits. The setting should not require to be changed after being made for a particular experiment location.



Figure 10: LHCrx




Figure 11: TTCcf modules

The LHCrx provides ECL outputs for further processing. The bunch clock output is an AC-coupled 40.079 MHz square wave while the orbit output is a DC-coupled negative-going pulse train of duration 1 μ s and period 88.924 μ s. These are the signal characteristics for which the transmitter encoder clock and TTCvi orbit input circuits were designed.

B. TTCcf

In order to choose the configuration of the electrical fanout of the bunch and orbit clocks from the TTCmi, the LHC experiment collaborations were asked to estimate the number of TTC partitions foreseen. The current ATLAS model [10] has 34 partitions and CMS foresees about 32. To allow some spares the TTCmi was designed to provide over 40 outputs each of the bunch and orbit clocks when it is equipped with 5 TTC clocks fanout (TTCcf) modules (see Fig. 11).

Each TTCcf is a dual 1:9 fanout module of which one half would normally be used for each clock. But the two halves are identical and can be used interchangeably to match the requirements of experiments that need more copies of the bunch clock than the orbit signal and *vice-versa*. For lowest jitter the signals are standard ECL, DC coupled since the orbit signal has a low repetition rate and is unbalanced. Since the

bunch clock is a square wave it may be AC coupled at sources and destinations. 36 TTCcf modules have been produced, sufficient to equip the main experiment TTCmi crates to drive up to 40 trigger partitions each.

It is expected that for trigger latency reasons the TTC transmitters for all the partitions will be grouped together in the vicinity of the central trigger processor so that the coaxial cables by which they receive the clock signals from the TTCmi will be short. For longer distances low-loss coaxial cable should be used with adapters, rather than the miniature solid dielectric type.

C. TTCmx

For even longer distance transmission, or where ground offsets are troublesome, TTCmx laser transmitters can be plugged into the TTCmi crate in place of some of the TTCcf modules.



Figure 12: TTCmx laser transmitter

The TTCmx module (see Fig. 12) provides 4 optical outputs at a level of +1 dBm, each of which can be fanned out by a 1:32 optical tree coupler to broadcast to a total of 128 destinations. An electrical output is provided to allow TTCmx modules to be daisy-chained to broadcast to larger partitions. The transmitters can be used to send either the 40.079 MHz bunch clock alone, or the composite encoded TTC signal. In the latter case a TTCrx with photodiode/preamp is required at the receiving end.

VII. TTCVI VMEBUS INTERFACE

Continuing further towards the receivers we reach the TTCvi, the VMEbus interface to the system (see Fig. 13). This is a fully programmable VME module [11]. On the A side it provides for the selection of the trigger source, which can be either the central trigger processor, or a local trigger source, or an internal random trigger generator. On the B side, the TTCvi provides for the generation of either asynchronous or synchronous commands, which can be in the short or long format (the short format without address being only for broadcasting). The synchronous commands can be generated at

any specified time relative to the LHC orbit and they can be programmed to occur just once, or as a sequence, or repetitively (i.e., occurring once per LHC orbit). The latter mode is used for broadcasting the important bunch counter reset signal that is used for adjusting the phase of the bunch counters in the TTCrx receivers.




Figure 13: TTCvi VMEbus interface

The TTCvi has proved very reliable, there having been only one failure in an early module, and a total of 83 modules were produced at CERN. Over half of the modules were a Mk II version incorporating several modifications and upgrades, including a burst mode which was required for beam instrumentation applications. A further 80 modules have been ordered from EFACEC and CERN has an option to purchase more if required.

VIII. TTC TRANSMITTERS

A range of laser transmitters is available for broadcasting the TTC signals at the LHC experiments (see Fig. 14). All the transmitters use multi-sourced InGaAsP low-threshold MQW hermetic-package Fabry-Perot lasers with temperature-compensated bias current and automatic power control by rear facet monitor. A detailed user manual for the laser transmitters is available [12].

First there is the TTCtx, a VMEbus module which has 14 optical outputs at a level of about +1 dBm. Each of these outputs can be split by a 1:32 optical tree coupler so that the module can drive up to 448 destinations. The module can be used as a single one, or it can be split into two halves with separate inputs for driving smaller partitions. On the other hand, for broadcasting to larger partitions, there is a facility to allow several TTCtx to be daisy-chained together. A crate full

of these modules can broadcast to a total of 8960 destinations, divided into between 1 and 40 trigger partitions.



Figure 14: TTCex and TTCtx laser transmitters

The TTCex has 10 outputs and so can broadcast to 320 destinations. This module can also be operated as a single unit or divided into two halves. Unlike the TTCtx, it incorporates dual biphasic mark encoders driven by a common VCXO/PLL with very low phase jitter. So one TTCex can drive two small partitions, whereas for large partitions it can be expanded by one or more TTCtx and an electrical encoder output is provided for that purpose. The output from each biphasic mark encoder has an rms jitter of less than 10 ps relative to the 40.079 MHz clock input to the module when transmitting PRBS data on both channels at the maximum rate.

The TTCmx (see Fig. 12) has already been mentioned in connection with the TTCmi. This is a 3-unit high version of the TTCtx that has 4 outputs and so can broadcast to 128 destinations. One of the main uses of these modules will be in TTC repeaters for the LHC beam instrumentation systems, which will have a somewhat different timing distribution architecture from that of the experiments. TTCmx modules can be housed in minicrates equipped with only a 5v power supply, which are much cheaper than VMEbus crates.

So far only 50 of these laser transmitters have been made, since most developers are working with small lab setups having few destinations. TTCtx are used by COMPASS while TTCex are used by the CMS Tracker, HCAL, trigger and DAQ groups, by the ALICE trigger group and by the ATLAS TileCal group at the H8 test beam. LHCb will require only TTCtx modules because the encoder function will be incorporated in their readout supervisor module [13], which replaces the TTCvi in that experiment.

For initial development work a LED transmitter (TTCvx) [14] is available which has 4 outputs at a power level suitable for driving one receiver each (see Fig. 15). Since the TTCvx PLL does not use a quartz oscillator it has higher encoder jitter than the laser transmitters but, on the other hand, this means that it can be driven with a wide range of clock frequency and not just 40.079 MHz.




Figure 15: TTCvx LED transmitter

IX. OPTICAL TREE COUPLERS

Passive optical tree couplers are used to fanout the outputs from the laser transmitter modules. Fused biconic taper (FBT) technology is preferred for the multimode couplers for broadcasting the TTC signals at the experiments. This is the classic technology which is still the least expensive and probably the most reliable [15], especially in a radiation environment, as it avoids the fibre-waveguide interface required with planar technologies. With multimode FBT technology a 1:32 fanout can be obtained in a single fusion.



Figure 16: 1:32 multimode optical tree coupler

In the case of the singlemode couplers for the PCR transmitters a 1:32 fanout is not possible in one fusion and these are made by splicing eight 1:4 couplers to two intermediate 1:4 couplers fed by one 1:2 coupler at the root. In spite of this, the excess loss is only slightly greater than for the single fusion multimode coupler and the uniformity remains good, with a spread of about 2 dB among the output ports of any given coupler.

Although bare couplers can be integrated in custom detector electronics and used for inter-module distribution within crates, additional protective packaging is required for mounting them in standard distribution racks. A number of enclosures for optical tree couplers have been developed. Fig. 16 shows a 1:32 coupler mounted in a 1U rack tray. These can be stacked above each other to achieve the desired fanout. Industrial cable breakout enclosures of this size generally provide for 24 outputs.



Figure 17: Cassette-mounted 1:32 singlemode tree coupler

Fig. 17 shows the cassette enclosure which was developed for the singlemode 1:32 couplers for the PCR transmitters. The assembly of such a configuration is quite delicate but the finished packaged coupler is very robust.

X. LATENCY

Preliminary data on the latency of the TTC system components are available [12]. The delays are measured from the A input of the transmitters through to the L1A output of the TTCrx. In the case of the TTCex transmitter, for example, the delay is 68 ns plus the propagation delay of 4.9 ns/m in the fibre. Note that the measurements are made with the internal deskews in the TTCrx set to minimum. The overall latency calculations might very well be made for the deskews set to their mid values to allow for subsequent adjustment. The internal delay in the optical tree couplers should not be forgotten – it is about 5.7 ns for the present version and 11 ns in earlier models having longer internal fibres.

Optical fibre patchcords for interconnecting TTC transmitters, couplers and receivers have been standardised in the CERN stores. They are equipped with ST/PC connectors and are currently stocked in lengths of 0.5m, 1.5m, 5m, 15m and 50m.

XI. SUBMINIATURE OPTICAL FIBRE CONNECTOR

Continuing to the receiving end of the fibre, an interesting development was carried out by one of the industrial partners in the project (Lemo SA, well known for their coaxial connector products). Lemo developed a subminiature rad-hard connector family for single optical fibres which includes a

low-mass active device mount and a bushing for connector-connector interfacing.

Although these parts are much smaller than standard telecoms components (see comparison in Fig. 18) they have a performance compatible with singlemode operation (0.25 dB insertion loss and 45 dB return loss with 9/125 μm fibre). The connectors have a zirconia ceramic ferrule although its diameter has been reduced from 2.5 mm to 1.5 mm.




Figure 18: Subminiature single optical fibre connector

Lemo took up this development on their own initiative and without any financial input from CERN. They are promoting the family of components on the open market as “the world’s smallest snap-on fiber optic connector”. The parts have passed quite severe aviation industry tests and there is a particular demand for the interconnector configuration as an alternative to splices.



Figure 19: 1B43 photodiode/preamp

XII. OPTOELECTRONIC RECEIVERS

Cost is an important issue for the TTC optoelectronic receivers since a total of well over 10,000 will be required. Early systems employed the HFBR-2316T, a synthetic packaged device which contains an InGaAs photodiode with a Si bipolar preamp having a bandwidth of 125 MHz. It costs about \$30. The current baseline choice is the TRR-1B43-000, (see Fig. 19) which costs only \$8 each in quantity 1000.

Unlike the 2316, the differential preamp of the 1B43 incorporates AGC so that this component has a large dynamic range. In spite of the AGC action, the variation of phase delay with input optical signal level is a relatively modest -25 ps/dB at -19 dBm.

In the normal configuration the photodiode/preamp is connected directly to the custom TTCrx timing receiver ASIC which has been developed by the CERN Microelectronics Group. However, some TTC systems use a configuration in which a single optoelectronic receiver drives a short ECL bus to which several TTCrx are connected. Complete modular optoelectronic receivers suitable for this application are available.

XIII. TTCRX RECEIVER ASIC

It is currently estimated that a total of about 20,000 TTCrx [16] will be required for all the LHC experiments and the machine applications. The design of the TTCrx has evolved with time. With the latest version (see Fig. 20) an external poststamp is not required between the photodiode/preamp and the TTCrx input unless these components cannot be mounted in very close proximity to each other. If it is inconvenient to locate the TTCrx near the front panel of a module, the photodiode can be placed next to the TTCrx and its optical input linked to a bulkhead bushing on the panel via a short optical fibre.




Figure 20: Rev. 3 (DMILL) TTCrx

It is also no longer necessary to employ an external configuration PROM with the TTCrx, an important requirement in a radiation environment. Each chip can now read its local ID from the subaddress and data bus following a reset.

The initial TTCrx design in 1 μm CMOS was ported to the 0.8 μm rad-hard BiCMOS DMILL technology in 1999 [17]. A yield of about 75% was obtained from the first MPW run and the packaged chips were subjected to irradiation tests [18]. It was found that the cumulative changes in jitter characteristics and clock deskewing linearity were acceptable for 8 Mrad gamma radiation from a 10 keV source and a neutron fluence of $5 \times 10^{13}/\text{cm}^2$.

The TTCrx includes a correction machine for single event upsets (SEU) which monitors the 10 critical configuration registers, each of which is protected by 4 Hamming check bits. The operation of the SEU correction machine can be monitored via the I²C bus. No errors were detected when the TTCrx was irradiated with 10^{11} protons or 1.2×10^{10} neutrons/cm² with an energy of 60 MeV. However, it was found that when the associated photodiode was irradiated with protons of energy exceeding 20 MeV, the SEU generated by direct ionization caused an unacceptable rate of loss of phase lock in the TTCrx.

To mitigate the effects of such SEU the TTCrx design was modified and an engineering run of 8 wafers was made in 2000. 1650 chips of the proven Rev 3.1 design (which can be used in a non-radiation environment) and 1550 of the modified Rev 3.2 version were received, the overall yield being 81%. Since the modified design was found to be entirely satisfactory, a new mask set for only that version was prepared for the second engineering run, which was launched in August 2001. During this run the foundry encountered a processing problem with the large PMOS devices, which considerably delayed deliveries. 3 wafers were received in February 2002, providing 1450 chips with a yield of 80%, and after modification of the process the final batch of 7 wafers was received in July 2002.

TTCrx packaging has also evolved, from the initial large PGA packages through a custom 15x15 mm 100 BGA design to the current standard-catalogue 13x13 mm 144 fpBGA solution. Some users have expressed concern about the use of this relatively exotic 1 mm pitch package but very satisfactory results have been obtained from assembly shops which are adequately equipped to handle it. Some of these services carry out 100% radiography control of each mounted ASIC.

It is currently planned to launch a TTCrx production run of 50 wafers in October 2002, which is expected to provide about 22,000 good TTCrx chips if a yield of 70% is attained. Packaged chips from this run should be available in May 2003. A new test fixture and enhanced test vector suite are being prepared for production testing using the in-house IMS IC tester.

XIV. TTCRX DEVELOPMENT BOARDS

A TTCrm mezzanine test board was initially produced for the convenience of developers evaluating the TTCrx for their applications. The board, which carried the photodiode/preamp and the supporting components required for the TTCrx, proved popular and several groups incorporated it in their final system designs. In some cases this allowed the motherboard to be of a less expensive class than that required for the TTCrx package.

A new version of the TTCrm (see Fig. 21) has been produced for the DMILL 144 fpBGA TTCrx and this can be

equipped with jumpers to specify the local address of the chip. A future version of the TTCrm could incorporate the QPLL auxiliary chip described later.



Figure 21: TTCrm mezzanine board

A number of other platforms are available to assist users with their initial development work. The TTCvr [19] is a VMEbus module developed by CMS which accepts the TTCrm mezzanine. It has an A24/D32 VME interface and output buffers and the functionality is provided by a Xilinx XC4006E.



Figure 22: TTCpr PMC receiver module

The TTCpr [20] is a PMC mezzanine (see Fig. 22) which has been developed by ANL for the ATLAS tile calorimeter DAQ. It has 4 blocks of 8K x 16b FIFO and is based on an Altera 10K30A FPGA, the configuration file for which is stored in an adjacent socketed serial PROM.

XV. QPLL AUXILIARY CHIP

For some applications, such as clocking high-speed serializers and high-resolution ADCs and TDCs, the jitter of the 40.079 MHz clock output from the TTCrx is too high. This problem was solved in the TTCmi by cleaning up the raw clock output from the chip with a VCXO/PLL having an rms jitter of about 7 ps. This early design used a bulky quartz oscillator followed by discrete multiplier stages with high-Q LC circuits to reduce subharmonic feedthrough to a low level.

It employed an ECL phase comparator/frequency discriminator and achieved low static phase error with an active loop filter based on an IC op amp with very high DC gain.




Figure 23: QPLL block diagram

A VCXO/PLL ASIC is now being developed by the Microelectronics Group to perform this function. This QPLL chip (see Fig. 23) is being designed in rad-tolerant $0.25\ \mu\text{m}$ CMOS technology and will be contained in a tiny ($4\ \text{mm} \times 4\ \text{mm}$) LPCC-24 package. Only the quartz crystal is external.

The QPLL will provide three LVDS clock outputs in two frequency multiplication modes from a CMOS or LVDS $40.079\ \text{MHz}$ clock input and is expected to have a pk-pk output jitter of less than $50\ \text{ps}$. A preliminary data sheet for the QPLL is available [21] and an MPW submission of the prototype design is planned for September 2002.

XVI. TTCBI BEAM INSTRUMENTATION INTERFACE

The LHC Beam Instrumentation Group is developing the TTCbi (see Fig. 24), another PMC mezzanine which provides the standard interface from the TTC transmitters at the PCR to the beam instrumentation that will be located around the LHC rings and along the SPS-LHC transfer lines.



Figure 24: TTCbi PMC module

Three prototypes have been extensively tested in the SPS during 2002 and a Mk II version having enhanced programmable trigger facilities is currently being designed. The specifications [22] should be frozen by the end of the year. A LynxOS driver will be available for the module.

The LHC experiments can also use these TTCbi modules to receive useful messages that will be broadcast on the B channel by the PCR transmitters about LHC “machine events” (such as start ramp, coast or dump) and information about the status of the machine (operating mode, beam type and energy, number of bunches, etc). Absolute time (with the granularity of the orbit period) will also be broadcast in 64-bit UTC format.

XVII. LASER SAFETY

The $1310\ \text{nm}$ optical signal broadcast from the PCR to the TTCmi at each experiment area is IEC 60825 [23] Class 1 (inherently safe), even prior to the level-adjustment attenuator. The high power Class 3B PCR transmitters themselves are in a “controlled access” area and before leaving this area their outputs are divided by 1:32 optical tree couplers having an insertion loss of about $19\ \text{dB}$.

The TTCex, TTCtx and TTCmx laser transmitters for local TTC signal distribution at the experiments (and LED transmitters such as the TTCvx) are also Class 1. However, the direct individual outputs from the laser transmitters should not be merged into a fibre ribbon equipped with an array connector since this could create a single extended source of greater hazard level. (Multifibre cable or ribbon fibres can be used as long as they are terminated by a single fibre breakout). The laser transmitter modules have a manual disable switch, laser emission indicator, warning label and facilities for interlock connection.

On the other hand, for the distribution fibres to the final TTCrx destinations (i.e., the outputs from the 1:32 optical tree couplers) the hazard level would remain Class 1 even if they were merged in ribbon fibres terminated with MT-12 connectors. Notwithstanding this, it is strongly recommended never to view output optical connectors or fibre ends either with the naked eye or with any kind of optical instrument.

XVIII. CONCLUSION

The unified timing distribution system reviewed in this paper provides for the transmission of the LHC fast timing signals from the RF generators of the machine through to the destinations at all the experiments, the test beam areas and the beam instrumentation around the ring and in the SPS transfer lines. It then provides facilities for the distribution of the signals over the TTC backbone at the experiments themselves. The use of a single CERN-wide system for the distribution of these timing signals is expected to result in cost savings and operational and maintenance advantages.

Work on the system is now progressing from the development to the implementation phase and the RD12 Project will terminate at the end of 2002. Detailed documentation on the system components and up-to-date status information is given on the TTC website [1]. Users are also encouraged to subscribe to the TTC mailing list [24].

XIX. ACKNOWLEDGEMENTS

RD12 would like to acknowledge the valuable contributions and support received from many other individuals and groups at CERN and related institutes. In

particular the collaboration enjoyed excellent co-operation with members of the Controls and Hadron RF Groups in SL Division and colleagues in the Electrical Engineering Group of ST Division at CERN.

XX. REFERENCES

- [1] TTC website: <http://ttc.web.cern.ch/TTC/intro.html>
- [2] "The Large Hadron Collider - Conceptual Design", CERN/AC/95-05(LHC), 20 October 1995.
- [3] Varela, J., "Timing and Synchronization in the LHC Experiments", Proc. 6th Workshop on Electronics for LHC Experiments, Cracow, Poland, 11-15 September 2000, pp. 77-87.
- [4] Leskovar, B., "Radiation Effects on Optical Data Transmission Systems", IEEE Trans. Nuclear Science, Vol. 36, 1989, pp. 543-551.
- [5] Taylor, B.G., "LHC Machine Timing Distribution for the Experiments", Proc. 6th Workshop on Electronics for LHC Experiments, Cracow, Poland, 11-15 September 2000, pp. 312-317.
- [6] Taylor, B.G., "TTC Distribution for LHC Detectors", IEEE Trans. Nuclear Science, Vol. 45, 1998, pp. 821-828.
- [7] Baudrenghien, P. et al, "SPS Beams for LHC: RF Beam Control to Minimize Rephasing in the SPS", presented at the 6th European Particle Accelerator Conference, Stockholm, 22-26 June 1998.
- [8] Beaumont, W. et al, "The CMS Tracker front-end and control electronics in an LHC-like beam test", Proc. 6th Workshop on Electronics for LHC Experiments, Cracow, Poland, 11-15 September 2000, pp. 135-140.
- [9] Taylor, B.G., "TTC machine interface (TTCmi) User Manual", RD12 working document, CERN.
- [10] Farthouat, P. et al, "PRR/Design Review of the ATLAS TTC", ATC-RD-ER-0010, 10 June 2002.
- [11] Farthouat, P. and Gällnö, P., "TTC-VMEbus Interface (TTCvi)", RD12 working document, CERN.
- [12] Taylor, B.G., "TTC laser transmitter (TTCex, TTCtx, TTCmx) User Manual", RD12 working document, CERN.
- [13] Jacobsson, R. and Jost, B., "The LHCb Timing and Fast Control system", Proc. 7th Workshop on Electronics for LHC Experiments, Stockholm, 10-14 September 2001, pp. 227-231.
- [14] Gällnö, P., "TTCvx Technical Description and Users Manual", RD12 working document, CERN.
- [15] "Qualification and Lifetime Testing of Corning™ Couplers", Application Note AN-119, Corning France, January 1991.
- [16] Christiansen, J., Marchioro, A., Moreira, P. and Toifl, T., "TTCrx Reference Manual - A Timing, Trigger and Control Distribution Receiver ASIC for LHC Detectors", RD12 working document, CERN.
- [17] Status Report on the RD-12 Project, CERN/LHCC 2000-002, 3 January 2000.
- [18] Toifl, T. et al, "Measurements of Radiation Effects on the Timing, Trigger and Control Receiver (TTCrx) ASIC", Proc. 6th Workshop on Electronics for LHC Experiments, Cracow, Poland, 11-15 September 2000, pp. 226-230.
- [19] Pollet, L. et al, "TTC Emulation Board Users Manual", TriDAS Internal Note, CERN, 6 October 1998.
- [20] Dawson, J.W. et al, "TTCpr: A PMC Receiver for TTC", Proc. 7th Workshop on Electronics for LHC Experiments, Stockholm, 10-14 September 2001, pp. 215-216.
- [21] "Quartz Crystal Phase-Locked Loop (QPLL)", Preliminary data sheet, CERN Microelectronics Group, June 2002.
- [22] Savioz, J.-J., "TTCbi - A Timing, Trigger and Control Interface for the Beam Synchronous Timing System", Draft LHC Project Document, CERN.
- [23] "Safety of laser products - Part 2: Safety of optical fibre communications systems", IEC 60825-2, 1993-09.
- [24] CERN listbox lhc-exp-ttc.