ISBN: 978-3-95450-254-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS1991-S08NC05

The Transmission of Accelerator Timing Information around CERN

C.G. Beetham, K. Kohler, C.R.C.B. Parker and J.-B. Ribes

European Organization for Nuclear Research (CERN)

CH - 1211 Geneva 23, Switzerland

Abstract

Prior to the construction of the Large Electron Positron (LEP) collider, machine timing information was transmitted around CERN's accelerators using a labyrinth of dedicated copper wires. However, at an early stage in the design of the LEP control system, it was decided to use an integrated communication system based on Time Division Multiplex (TDM) techniques. Therefore it was considered appropriate to use this facility to transmit timing information over long distances. This note describes the overall system, with emphasis placed on the connectivity requirements for the CCITT G.703 series of recommendations. In addition the methods used for error detection and correction, and also for redundancy, are described. The cost implications of using such a TDM based system are also analyzed. Finally the performance and reliability obtained by using this approach are discussed.

I. INTRODUCTION

In the planning phase of the LEP collider it was recognized that the much greater physical size of LEP compared with previous accelerators would result in a radical change in the way communications systems of all types would be implemented. The larger size implied correspondingly longer cables and more signal regenerators, and consequently much more expense. Moreover, most cables had to be routed through the LEP tunnel itself, significantly limiting the space available for future exploitation of the tunnel. It was evident that the number of cables would have to be reduced by multiplexing information for several different users onto one cable. Consequently, it was decided to install a multiplexed system throughout LEP as a general user facility.

II. ACCELERATOR TIMING SYSTEMS

A. General Features

Typically, computer control systems have a real-time response in the order of 10 to >1000 ms. Whilst this is adequate for many applications there is always a need to trigger selective hardware equipment with a finer time resolution. This is achieved by means of a timing system.

An accelerator timing system is simply a fast broadcast message transfer system. The messages (events) are normally short (typically less than 32 bits) and should have a maximum time resolution of 1 ms. The jitter of each message is determined by the transmission medium and by the sampling processes at the generator and receiver. Depending upon the type of machine, the events are either pre-programmed, according to the specific machine cycle, or triggered by external stimuli, such as emergency beam dump.

A principle difference between a timing system and a control system is that when a transmission error is detected, the control system normally retransmits the message. This is inappropriate with a timing system as, inherently, the message contains a timing reference.

Although the timing and control systems may be considered as two separate identities they are in fact strongly coupled to the machine that they are controlling. For this reason timing systems have always been designed in-house and tailored to the specific accelerator concerned. This implies producing timing equipment compatible with the chosen control system. Equipment of this type is normally not available commercially.

For a large accelerator a typical timing system consists of three parts:

- a central timing generator
- a long distance transmission network
- a local distribution system.

The timing system chosen for LEP and subsequently used on the SPS has been described elsewhere[1]. This article concentrates on the long distance transmission network.

B. Long Distance Transmission

In the case of CERN's Super Proton Synchrotron (SPS) machine, "long distance" refers to the distance between two adjacent auxiliary surface buildings located above the SPS tunnel. For the LEP machine it refers to the distance between two adjacent tunnel alcoves. For both scenarios the distance is between one and two kilometers.

For the original SPS machine the accelerator timing information consisted of a 1 ms clock and short (7 bit) trigger messages called events. The events were Manchester encoded and transmitted to each building over a video cable containing two good quality twisted pairs. One pair was used to transmit the 1ms clock whilst the other carried the events. The transmission rate was 333 kbit/s and no error detection or correction was employed. "Tap-off" amplifiers were in-

ISBN: 978-3-95450-254-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS1991-S08NC05

stalled in each building in order to transmit the timing signals around the SPS. These were eventually duplicated in order to reduce the downtime of the machine.

In addition a separate twisted pair cable was used to transmit the revolution frequency (43kHz). Identical amplifiers were used in each building; again they were duplicated. This latter system is still operational, whilst the SPS general machine timing has since been upgraded to the LEP system.

C. The Timing Distribution Standard

The local distribution of the timing information is done according to EIA specification RS-485. The data format adopted is supported by integrated circuits from National Semiconductor (the NS8342/8343 transmitter/receiver set). These ICs frame the 32 bit event into four bytes. Each frame is enveloped within a predefined start/end sequence and in addition each byte starts with a synchronizing bit and ends with a parity bit. The NS8343 receiver chip contains a seven bit error register which indicates the detection of a mid-bit transition fault, an invalid ending sequence and also a parity error. The contents of this register are used for error detection.

The entire frame is Manchester encoded. This coding scheme was developed at Manchester University in the U.K. for data recording onto magnetic media. The principle characteristic of the code is that no D.C. component is transferred to the transmitting line medium. This is achieved by ensuring that each data bit has an equal negative and positive component, as shown in Fig. 1.

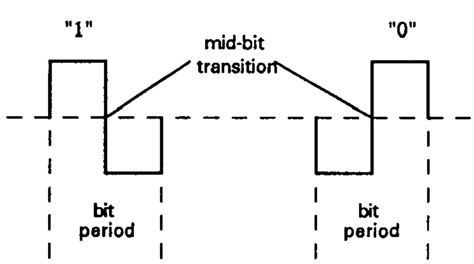

Fig. 1 Manchester Encoding

The use of Manchester coding doubles the bandwidth requirement because each bit period must contain a mid-bit transition.

III. MULTIPLEXING SYSTEMS

There are two main methods of multiplexing: frequency-division multiplexing (FDM) and time-division multiplexing (TDM). FDM uses a different frequency band for each signal; analogue filters are used to extract the required signal at an access point. These filters require tuned circuits which do not lend themselves to large-scale integration, which require careful setting up, and which are prone to drift. This, together with the general trend towards digital transmission, led to the decision to use time-division multiplexing for LEP.

In time-division multiplexing a number of bits are taken sequentially from each of the multiplexor's input ports and applied to the output port. The converse operation is applied in the other direction. This procedure is shown, simplified, in Fig. 2.

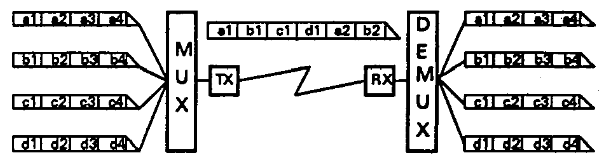

Fig. 2 Simplified TDM System

In practice TDM systems are far more complex. This is because the high-rate data streams carry extra information, particularly for synchronization purposes and in order to deal with clock rates which are only nominally synchronous.

TDM equipment is used widely around the world by PTTs and other telecommunications suppliers and users. The standards for the equipment have been set by the CCITT worldwide [4]. Since compatible equipment is available from a large number of vendors, the market is very competitive and therefore the prices are low.

As previously mentioned, it was decided to adopt TDM technology for LEP [2]. This is used to transmit the machine timing and also the revolution frequency timing information. The latter system has been fully described elsewhere [3].

The timing events are transmitted at 512kbit/s over a 2.048Mbit/s TDM link. At the 2.048Mbit/s data rate the G.703 Recommendation specifies a coding scheme known as HDB3. The TDM-Timing interface unit makes this code conversion transparent to users of the timing system.

The interface standard (CCITT Recommendation G.703) for these high speed connections is being used more and more frequently. For example, video CODECS, LAN bridges, and terminal concentrators are available with G.703 interfaces. However, one disadvantage of using a standardized approach to interfacing is that, in special situations some interface conversion is required. As will be shown later the cost of this conversion, in the LEP situation, is outweighed by the economies brought about by multiplexing.

In the case of the interface between the timing system and the TDM network it was decided to subcontract to industry the design and production.

ISBN: 978-3-95450-254-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS1991-S08NC05

IV. DESCRIPTION OF THE TDM-TIMING INTERFACE

Fig. 3 shows a block diagram of the TDM-Timing interface.

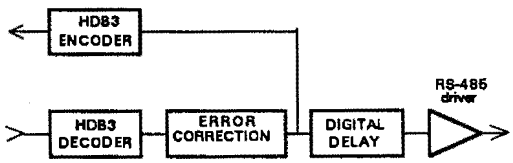

Fig. 3 TDM-Timing Interface

The interface unit was specified at CERN but designed and manufactured externally. The four principle parts of the interface are described below.

A. HDB3 Decoder/Encoder

The incoming 2.048 Mbit/s data stream is first galvanically isolated and terminated in a 75 Ohm resistive load. Thereafter, it is passed on to the HDB3 decoder. The HDB3 code is somewhat similar to Manchester coding, in that it has a low DC component and also contains clock information, but it is much more efficient in bandwidth usage. The HDB3 decoder generates a NRZ data signal and associated clock. In addition the decoder generates two alarm signals: "loss of input" which is activated when the input is no longer present, and "violation" which is generated when violations of the HDB3 coding law are detected. Both of these signals are monitored by the control system.

The HDB3 encoder simply performs the opposite function to the decoder, ie. it converts the NRZ signal to HDB3 code and passes it on to the outgoing transformer isolated driver.

B. Error Correction

It can be seen that Manchester coding inherently introduced some redundancy, as does the superimposition of each Manchester bit on the TDM links. In effect four 2.048 Mbit/s TDM bits are being used to convey one 512 kbit/s timing bit.

At the specification stage of the TDM-timing interface it was decided to exploit this redundancy and incorporate a simple error correction system. Four TDM bits are taken at a time and applied to the error correction circuitry. The four possibilities for single-bit errors when transmitting Manchester codes are indicated in Table 1.

"1"		"0"	
Correct code:	1100	Correct code:	0011
Incorrect code:	0100	Incorrect code:	1011
"	1000	"	0111
"	1110	"	0001
"	1101	"	0010

Table 1 Single-bit Errors

As can be seen from the above Table, in the event of a single bit error introduced by the TDM network there is still enough information to distinguish between a 1 and 0. In fact the scheme can also correct some double-bit errors introduced by the TDM network. As this network has a very low error rate ($< 10^{-12}$) this means that the error rate seen by the timing system is virtually zero.

C. Transmission Delay Compensation

The SPS/LEP timing system was designed so that any specified event would arrive at all of the related equipment around the machines with a one millisecond resolution encompassing a maximum jitter of two microseconds. However in physically large accelerators such as the SPS and LEP there are significant propagation and equipment delays. In order to compensate for this, delays were incorporated in the interface unit. Each delay consists of a 12 bit register incremented by the 2.048MHz. clock derived from the incoming data stream. The required value can be written and read by the control system.

D. Local Timing Output

The output of the timing interface conforms fully to the electrical characteristics defined by CCITT Recommendations V.11 and X.27 and EIA specifications RS-485. The differential signal is Manchester encoded.

V. REDUNDANCY

Redundancy, or more specifically duplication, has been installed following invaluable experience gained from the original SPS timing system installation.

The TDM links provided for the SPS and LEP timing systems have been configured in a double ring topology, one ring for each machine. The links on each segment of the ring are inherently full duplex; this feature has been exploited to provide two independently routed signals, following clockwise and anticlockwise routes around each ring. These alternate signals are used in the interface equipment for added reliability.

Besides profiting from the full duplex feature of the TDM, the interface units were also duplicated as well as the power supplies used to power the local line drivers. Switch-

ISBN: 978-3-95450-254-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS1991-S08NC05

ing to the redundant system is done automatically when necessary. In addition the selection can be made via the control system. The equipment is constructed in Euromechanics as defined in DIN 41914.

VI. COST CONSIDERATIONS

It is clear that the cost benefits brought about by the use of multiplexing will be greater as the number of channels on a link increases. In the case described here each link carries three channels: 2Mbit/s for the general machine timing, 2Mbit/s for the beam synchronous timing, and 8Mbit/s for the machine control token ring.

A careful analysis of the cost of the installation has been done, as well as an estimate of the cost involved in implementing the transmission of these three channels in the traditional way on dedicated links. It was found that these two costs were roughly equal, ie. in this case there has been, as yet, no significant cost advantage in multiplexing. However, there is considerable spare capacity to each of the 24 locations served by the network, and this allows extra channels to be put into service very quickly and cheaply, as no significant additional cabling is required. This flexibility has already proved extremely useful.

VII. CONCLUSIONS

Although in terms of bandwidth, the timing system is not a major user of the TDM system it is a vital user, ie. if the TDM is not operational, the timing system cannot work. The use of the TDM network for the transmission of timing information has proved to be more reliable than the previous transmission system which used dedicated cables and repeaters.

The error-correction features are extremely effective: since the introduction of this system there has been no lost beam time due to transmission errors. The cost of supplying the extra bandwidth for error correction is far outweighed by the advantages of improved reliability, and error correction would be an indispensable part of any future system. In fact studies of advanced error correction algorithms are presently under way with a view to using these techniques for a pilot implementation of an integrated fast control and timing system [5].

VIII. ACKNOWLEDGEMENTS

Many people have contributed to the success of the project to transmit timing information in the manner described in this paper. Particular mention must be made, however, of B. Amacker and P. Nouchi.

IX. REFERENCES

[1] C.G. Beetham, R.J. Lauckner, C.G. Saltmarsh, Overview of the SPS/LEP Fast Broadcast Message Timing System, Proc. of the 1987 Particle Accelerator Conference, Washington, D.C., Vol. 1, p. 766.

[2] C.R.C.B Parker, The Provision of Telecommunications Links for the LEP Collider, 1989 Particle Accelerator Conference, Chicago, Ill.

[3] G. Baribaud et al., The Beam Synchronous Timing System for the LEP Instrumentation, International Conference on Accelerator and Large Experimental Physics Control Systems, Vancouver, Canada 1989.

[4] CCITT "Red Book", Fascicle III.3, International Telecommunications Union, Geneva 1985.

[5] C.R.C.B. Parker, M.E. Allen, H.C. Lue, C.G. Saltmarsh, A New Data Communication Strategy for Accelerator Control Systems, International Conference on Accelerator and Large Experimental Physics Control Systems, Vancouver, Canada 1989.