The Design and Performance of the ZEUS Central Tracking Detector z-by-Timing System

D.S. Bailey[1], N. Harnew, T. Khatri[2], M. Lancaster[3], I.C. McArthur, J.D. McFall, J. Nash[4], P.D. Shield, S. Topp-Jorgensen, F.F. Wilson[5].
University of Oxford, Department of Nuclear Physics, Keble Road, Oxford OX1 3RH, UK.

B. Foster, G.P. Heath, C.J.S. Morgado[6].
University of Bristol, H.H. Wills Laboratory of Physics, Tyndall Avenue, Bristol BS8 1TL, UK.

R.C. Carter, M.D. Jeffs, R. Milborrow, M.C. Morrissey, D.A. Phillips[7], S.P.H. Quinton, G. Westlake, D.J. White.
Rutherford Appleton Laboratory, Chilton, Didcot, Oxon OX11 0QX, UK.

J.B. Lane, G. Nixon, M. Postranecky.
University College, Department of Physics and Astronomy, Gower Street, London WC1E 6BT, UK.

Abstract

The ZEUS Central Tracking Detector utilizes a time difference measurement to provide a fast determination of the z coordinate of each hit. The z-by-timing measurement is achieved by using a Time-to-Amplitude Converter which has an intrinsic timing resolution of 36 ps, has pipelined readout, and has a multihit capability of 48 ns. In order to maintain the required sub-nanosecond timing accuracy, the technique incorporates an automated self-calibration system. The readout of the z-by-timing data utilizes a fully customized timing control system which runs synchronously with the HERA beam crossing clock, and a data acquisition system implemented on a network of Transputers. Three dimensional space-points provided by the z-by-timing system are used as input to all three levels of the ZEUS trigger and for offline track reconstruction. We measure a z resolution of 4.4 cm for multi-track events from positron-proton collisions in the ZEUS detector.

---

[1]Now at University of Bristol, H.H. Wills Laboratory of Physics, Tyndall Avenue, Bristol BS8 1TL, UK.
[2]Now at A.T. Kearney Ltd., 130 Wilton Road, London SW1V 1LQ, UK.
[3]Now at LBL, Berkeley, CA94720, USA.
[4]Now at Unipart Information Technology, Cowley, Oxford OX4 2PG, UK.
[5]Now at PPE Division, CERN, 1211 Geneva 23, Switzerland.
[6]Now at M.A.I.D. PLC., 38 Leicester Square, London WC2H 7DB, UK.
[7]Now at Data & Research Services PLC., Linford Road, Milton Keynes MK14 6LR, UK.

1 Introduction

This paper describes the readout and performance of the ZEUS Central Tracking Detector z-by-timing system during its first three years of data-taking. ZEUS is a multipurpose detector [1] which records collisions of 27.5 GeV electrons or positrons with 820 GeV protons at the Hadron Electron Ring Accelerator (HERA) at DESY. The Central Tracking Detector (CTD) [2, 3] is a cylindrical drift chamber which surrounds the beam pipe at the $e-p$ interaction region and is located directly inside a superconducting magnet which provides a 1.43 T axial field. Its purpose is to measure momenta of charged particles between polar angles of $7.5° < \theta < 164°$, provide dE/dx information to enhance particle identification and to provide tracking information to all three levels of the ZEUS trigger.

The objective of the ZEUS First Level Trigger (FLT) [1, 4] is to enhance the sample of genuine electron-proton collisions by reducing the 30 kHz background rate from proton beam-gas and beam-scraping interactions down to a maximum of 1 kHz. The CTD Track Trigger is a component of the ZEUS FLT which contributes to this reduction. Firstly it identifies those events which have a vertex consistent with originating from the $e-p$ interaction region. Secondly it can select interesting physics processes that may deposit only small amounts of energy in the ZEUS calorimeter but have a distinctive track topology (e.g. elastic ρ^0 and J/ψ production [5, 6]).

The CTD has a multi-cell stereo superlayer wire geometry, similar in design to the CDF chamber [7]. Five of nine superlayers are axial, having wires running parallel to the z axis[1]; the remaining four superlayers have $\sim 5°$ stereo angles which allow three-dimensional track reconstruction of the z coordinate. The stereo information from the CTD provides an accurate offline measurement, however it cannot be used in the ZEUS FLT because of insufficient processing time. Hence for the FLT we have chosen the technique of z-by-timing which can provide fast (in our case a digitization within 48 ns), three-dimensional space-point measurements of the CTD hit information. The method utilizes a time difference measurement in which the z coordinate is proportional to the time difference between the induced pulse arriving at each end of the wire, the constant of proportionality being nominally half the speed of light [8, 9, 10]. By comparison, measuring the z position by the method of charge division would be inadequately slow because of need to integrate charge over many tens of nanoseconds.

HERA is designed to run with 220 bunches in each of the electron and proton rings with a time interval between crossings of 96 ns. The 96 ns beam crossing interval puts special demands on triggering and readout of the ZEUS detector. To ensure that ZEUS is sensitive to every HERA beam crossing, the data through the z-by-timing system and the FLT processors must be pipelined [1, 8]. This

[1]The ZEUS coordinate system is a right handed system with the z axis pointing in the proton beam direction, the x axis pointing to the centre of HERA and the y axis vertically upwards. Hereafter we refer to the proton ($+z$) direction as the 'forward' direction and the electron ($-z$) direction as the 'rear' direction.

requirement is the overriding influence on the design of the readout architecture and trigger system.

As well as being an important component of the ZEUS FLT, the z-by-timing information is also used as input to trigger levels 2 and 3. The Second Level Trigger (SLT) [1, 11] reduces the maximum 1 kHz First Level rate down to a maximum of 100 Hz by refined processing of track and calorimeter information. Here the z-by-timing information is used as input to a track finding and vertex fitting package, implemented on a network of Transputers. At the Third Level Trigger (TLT) the space-points provided by the z-by-timing system are used in full track reconstruction to obtain an accurate vertex measurement (giving a stand-alone resolution of $\sim$40 mm). The output from the TLT is written to magnetic storage at a rate of $\sim$3-5 Hz.

Offline, not only does the z-by-timing system provide an aid to track finding at the first stage of track reconstruction but it can, and has been, used as an independent tracking system in its own right. During the first year of HERA data-taking, the z-by-timing system provided the sole tracking readout system in ZEUS. Throughout the first three years of HERA operation, the CTD and the z-by-timing system have run successfully and reliably.

This paper is outlined as follows. The CTD is described in Section 2 and an overview of its readout is given in Section 3. Section 4 describes the CTD amplification system and the implications for the linearity of the z-by-timing measurement. The z-by-timing readout cards, which incorporate time-to-amplitude conversion and digitization, are described in Section 5. Section 6 discusses the CTD calibration system, including pulse generation and calibration control. Section 7 describes timing control and data acquisition (DAQ) of the z-by-timing data. The performance during the first three years of HERA operation is detailed in Section 8. Finally, Section 9 contains a brief summary.

2 The Central Tracking Detector

The CTD is one of the inner tracking detectors of ZEUS. It is a cylindrical drift chamber with active inner and outer radii of 190 mm and 785 mm respectively, and has an active length of 2.03 m. It surrounds the beam pipe at the $e-p$ interaction region and is located inside a superconducting solenoid which provides a 1.43 T axial magnetic field. The chamber axis is coincident with the beam axis.

Fig. 1 shows an octant of the CTD endplate normal to the beam axis in the $r-\phi$ projection[1]. The total number of wires in the detector is 24192 of which 4608 are sense wires. The wires are arranged in nine 'superlayers'. The five odd numbered superlayers are axial, their wires running parallel to the beam axis (0°). The four even numbered superlayers have small stereo angles, their wires being

[1]Here r is defined as the radial distance from the z axis, ϕ is the azimuthal angle measured with respect to the x axis. The polar angle θ is defined with respect to the proton beam direction.

rotated by approximately $\pm 5°$. This combination of axial and stereo superlayers allows accurate three-dimensional reconstruction of charged particle tracks at the offline analysis stage.

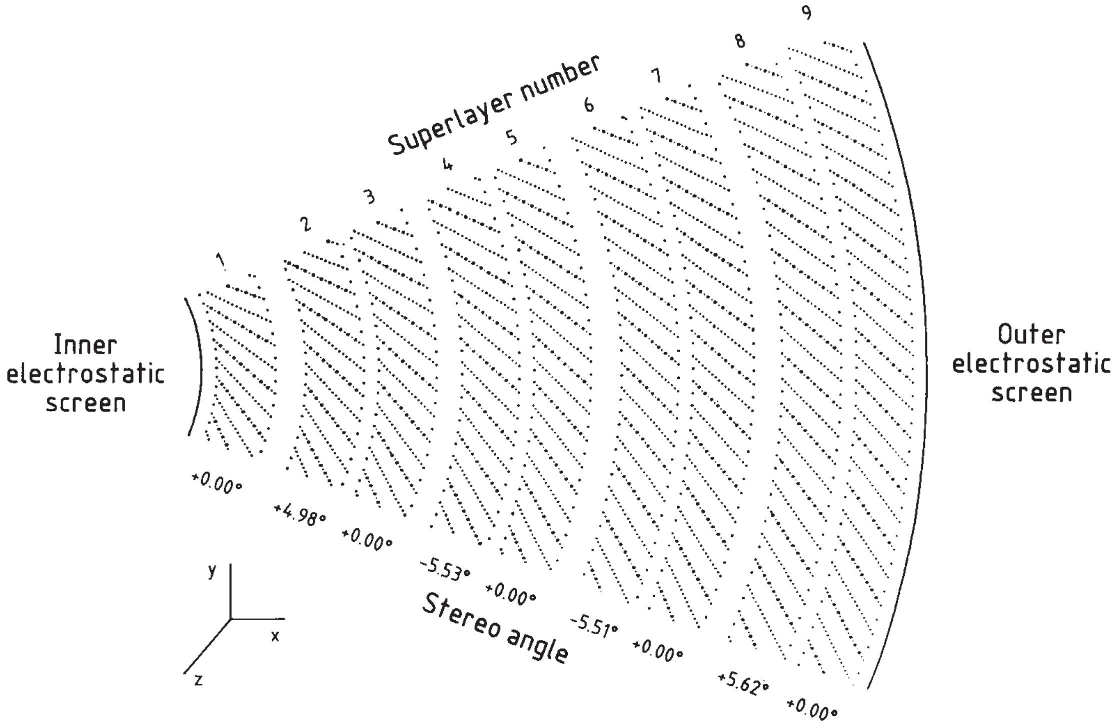

Figure 1: The layout of an octant of the CTD at the chamber endplate. Sense wires are indicated by highlighted dots.

A superlayer is divided into a number of 'cells', each containing eight sense wires. The cell boundary is defined by planes of field wires. The eight sense wires, separated by field-shaping wires held at ground potential, are arranged in a straight line at 45° to a radial line from the chamber axis. This arrangement has a number of advantages for the chamber performance [2, 7]. The separation of adjacent sense wires is approximately 8 mm and the distance from the sense wires to the cell boundary, which defines the maximum drift distance in the cell, is typically 25 mm. The sense wires are tungsten and have a diameter of 30 μm and a resistance of 80 Ω/m.

A total of 704 wires are instrumented with z-by-timing readout. All wires in superlayer 1 are equipped in order to give maximum acceptance for tracks in the forward proton direction (note that tracks from high Q^2 $e-p$ collisions are predominantly produced in the forward proton direction due to the asymmetric beam energies, whilst the background is similarly forward since it arises from proton collisions with stationary gas molecules). For less forward going tracks, the lever arm, and hence the vertex determination, is greatly improved by adding a relatively small number of hits from additional axial superlayers. Hence as a compromise between improved performance and cost, axial superlayers 3 and 5 are half instrumented with readout.

The voltages on the wires in a cell are used to define an operational drift field (E_d) and sense wire surface field (E_s). The choice of these parameters is dependent on the amplifying gas mix in the chamber which in turn is an important consideration for efficient and safe operation of the CTD. We have operated the CTD and its prototypes with a variety of gas mixes at 3 mbar above atmospheric pressure, and their constitutions are given in Table 1. All gases and operating conditions listed result in azimuthal drift with velocity $\sim$50 μm/ns, a maximum drift time of 500 ns (equivalent to five beam crossings of the HERA accelerator) and a sense wire surface field defining a gas gain of approximately 1×10^5 [12]. To ensure the condition of azimuthal drift, the tilted cell geometry requires the choice of a 45° Lorentz angle. Our preferred gas mix is 50:50 argon:ethane with a small admixture of ethanol ($\sim$1.6%). With a 1.6 kV/cm drift field this gives a 45° Lorentz angle in a 1.8 T magnetic field, the nominal field of ZEUS. Unfortunately technical difficulties have limited the field to 1.43 T which has put stringent constraints on chamber operation. In addition, to date we have chosen to operate the CTD with 'safe' argon:CO_2:ethane gas mixes during the relatively low luminosity running of HERA. This has meant we have had to run the CTD at a drift field close to 1.2 kV/cm, below which the field shaping becomes poor [2]. This has implications on chamber performance, especially on the hit efficiency at cell boundaries. We will return to this point in Section 8.

3 Overview of the CTD Readout

Fig. 2 shows a block diagram of the CTD front-end readout system. The major components are as follows:

- Preamplifiers mounted directly on either end of the CTD end-flanges amplify the chamber pulses and drive the signals down 42 m of coaxial cable to postamplifiers, which provide most of the electronic gain.

- A 104 MHz Flash-Analogue-to-Digital Converter (FADC) system (hereafter referred to as the $r-\phi$ FADC system) samples the amplified pulses from all 4608 sense wires and measures drift times [13, 14]. In this way, $r-\phi$ trajectories of charged particles in the plane perpendicular to the beam are measured to a precision of about 190 μm on each wire. The performance of this readout system will be the subject of another paper.

Nominal gas mix by volume (%)			Added Ethanol Fraction (%)	B (T)	E_d (kV/cm)	E_s (kV/cm)	v_d (μm/ns) ($\pm 1\%$)	$\theta_{Lorentz}$ (degrees) ($\pm 0.5\%$)	Year of HERA operation
Argon	CO_2	Ethane							
50	0	50	1.6	1.80	1.50	232	51.0	45.0	T.B.
85	13	2	0.95	1.80	1.35	192	49.0	45.0	T.B.
90	8	2	0.84	1.43	1.20	172	47.3	44.7	1992/3
85	8	7	0.84	1.43	1.22	182/4	50.5	43.4	1994
83	5	12	0.84	1.43	1.22	183	48.3	44.7	1995

Table 1: The nominal operating conditions for various CTD gas mixes which have been used over the course of the HERA running period (T.B. signifies test beam operation prior to HERA startup). E_d and E_s are the drift field and sense wire surface fields respectively which result in a drift velocity v_d and Lorentz angle $\theta_{Lorentz}$ in an applied magnetic field B.

- The z-by-timing readout system digitizes the time difference between the arrival of the pulses from the two ends of the chamber using a pipelined Time-to-Amplitude Converter (TAC). The intrinsic resolution of the TAC is $\sim$36 ps.

The readout philosophy of the CTD utilizes the geometrical symmetry of the chamber. We have chosen to electronically divide the CTD into 16 azimuthal sectors in ϕ, each of which is serviced independently by its own rack of electronics. All electronic readout modules are housed in custom-built 9-Unit (9U) high *Teradyne* crates [15]. The electronics for each sector are contained in three crates, stacked within a single rack. Fig. 3 demonstrates the layout. The 'Postamplifier Crate' contains postamplifier cards for that sector. From the postamplifiers the pulses are split into the 104 MHz FADCs located in the '$r - \phi$ Crate', and the z-by-timing electronic cards, located in the 'z-Crate'.

Each z-Crate contains eleven z-by-timing cards. It also contains seven or eight CTD FLT cards which process the z-by-timing data from that sector. The trigger cards dictate the need for custom crates and backplanes because of the large number of interconnections required, 360 per slot.

A calibration system serves the z-by-timing system and also provides a general diagnostic tool to check electrical continuity. The system includes the following :

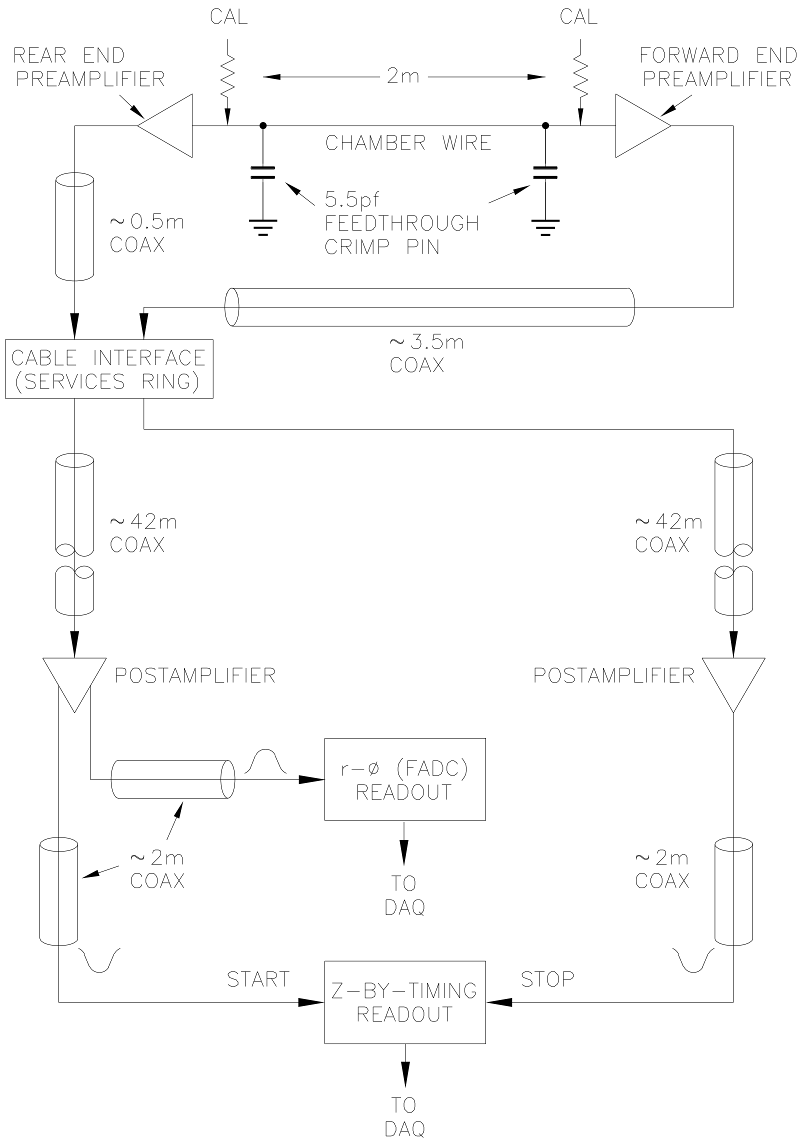

Figure 2: A block diagram of the CTD front-end readout system.

- Programmable 'Calibration Controller' modules (described in section 6.2), one located in each z-Crate. These modules supervise a calibration run sequence.
- 'Calibration Driver' modules (described in section 6.3), one located in each Postamplifier Crate. These modules generate pulses for transmission to the chamber.

The readout of the z-by-timing and trigger cards requires modules which control the timing, pipeline addressing, and local DAQ. The system includes the following :

- Readout Controller (ROC) modules [16, 17] (described in section 7.2), one located in each z-Crate. Each ROC supervises the local DAQ of the z-by-

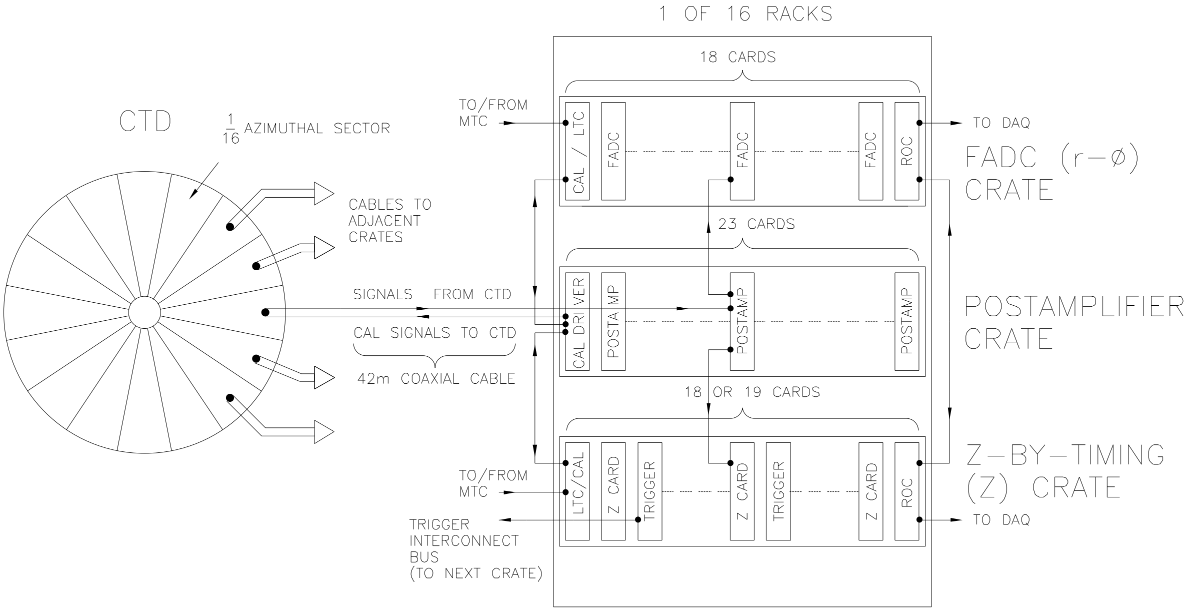

Figure 3: The crate layout for $\frac{1}{16}^{th}$ of the CTD.

timing and trigger data for that crate.

- A Master Timing Controller (MTC) module (described in section 7.3), housed in an external rack. The MTC synchronizes the CTD readout and provides the interface to the Global First Level Trigger (GFLT).

- Local Timing Controller (LTC) modules (described in section 7.4), one located in each z-Crate. Each LTC controls the timing and addressing of the pipelines and buffers of the z-by-timing and trigger cards in that crate. Each LTC and Calibration Controller share a a single 9U card.

The $r - \phi$ Crate contains its own LTC, Calibration Controller and ROC in addition to the 104 MHz FADC readout cards.

A schematic diagram of the CTD readout architecture is shown in Fig. 4. Triggers at three levels select the events and reduce the throughput rate to a maximum of 1 kHz, 100 Hz, and 5 Hz respectively at each stage. Since an FLT decision cannot be made within the HERA beam crossing time of 96 ns, it is necessary to store the data from all ZEUS components in local pipelines. Each pipeline contains a record of the component data for at least the previous 5 μs, the time to make an FLT decision. Trigger processing is divided between local component triggers and the GFLT. The CTD FLT completes its internal trigger calculation after 2 μs and passes information for a particular crossing to the GFLT. The GFLT then issues a decision on the strength of the information supplied from all components. If no GFLT 'Accept' is sent back to the individual components, the event data in the pipelines are overwritten. Events accepted by the GFLT are transferred from the local pipelines to the SLT. As with the FLT, trigger processing at the SLT is divided between local component triggers and the Global Second Level Trigger (GSLT). The CTD SLT is implemented on a network of Transputers which run complex trigger algorithms and contribute a reduction in the event rate to a maximum of 100 Hz. Events in this case are stored in local software buffers. For events accepted by the GSLT, the digitized data from all components are collected together in the Event Builder. From here events are passed to the TLT, a processor farm, where for the first time trigger processing can be performed on the complete event data from all components. The TLT runs a first-pass data analysis algorithm (which includes track finding and fitting). Events accepted by the TLT are passed to mass storage at a maximum rate of 5 Hz.

4 The CTD Amplification System

The amplifiers of the CTD have been described elsewhere [3], however since they are a crucial component of the z-by-timing measurement the relevant features are described in this section.

The preamplifiers are assembled in groups of eight channels onto printed circuit boards which form the preamplifier cards. These are mounted directly onto the chamber endplates, each card services a single cell of wires. Each of the 4608 sense wires is connected to a preamplifier channel at the rear end of the chamber to provide the $r-\phi$ drift-time measurement. To achieve the z-by-timing measurement, the 1152 axial sense wires in superlayers 1, 3 and 5 have identical preamplifiers mounted on the forward end of the chamber. A total of 704 sense wires provide hit information to the z-by-timing readout system; innermost wire 1 to outermost wire 8 in superlayer 1 (256 channels) and wires 1, 3, 5 and 7 in superlayers 3 and 5 (192 and 256 channels respectively).

The voltage gain of the preamplifiers is constrained to a relatively low value, 2.5, in order to minimize power dissipation (16 mW per channel). The first transistor stage of the preamplifier has a common-base configuration such that the emitter current can be tuned to actively terminate the sense wires in the characteristic impedance of the chamber, 360 Ω. The need to minimize pulse reflections caused by improper termination at the ends of a wire is well known, and is crucial to the technique [9, 18]. These reflections result in a non-linear response of the time difference measurement as a function of z position of the hit. As described in [3, 9], mismatch of the characteristic impedance is caused by a 7 pF capacitance at the chamber endplates introduced by the wire mountings, and the 4 pF input capacitance of the preamplifiers. The introduction of a 1μH inductor at the front-end of the preamplifier card forms a π-section of transmission line and improves chamber termination and hence linearity. A measurement of the residual non-linearity is described in section 5.1 and shown in Fig. 7(a).

The signals from the preamplifiers are transmitted to postamplifiers via coaxial

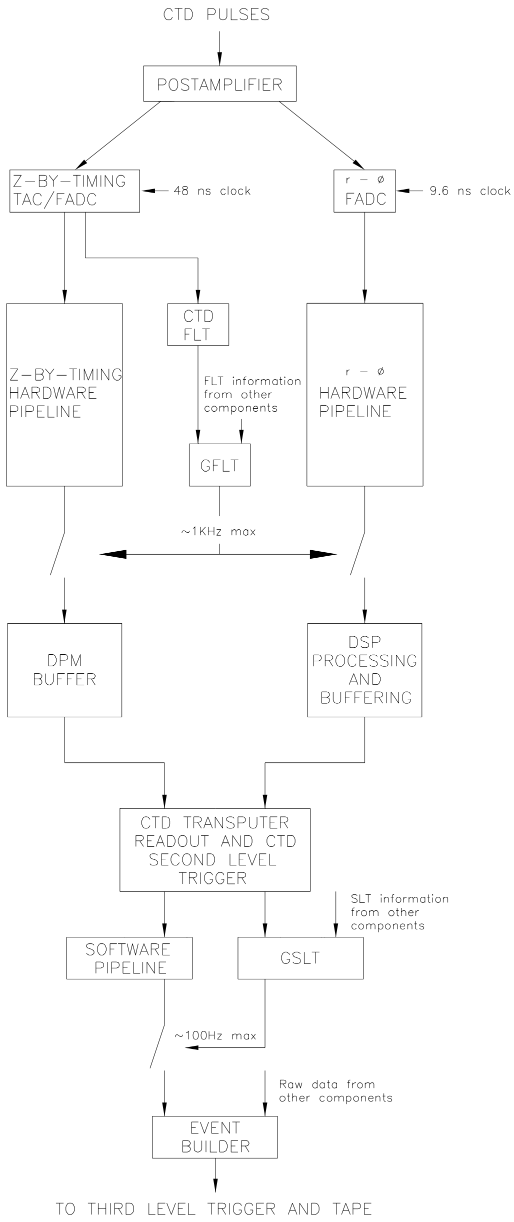

Figure 4: A schematic diagram of the CTD DAQ architecture.

cables with total lengths of 45.5 m and 42.5 m from the forward and rear ends of the chamber respectively. Two types of signal cable are used. Thin 1.6 mm outer diameter (O.D.) foam-core coaxial cable, lengths of approximately 3.5 m and 0.5 m from the forward and rear ends respectively, connect the preamplifiers to bulkhead connectors at the 'Services Ring'(see Fig. 2 and Ref. [3]). At the Services Ring, the thin coax is interfaced to 42 m of high quality foam-core 2.8 mm O.D. coaxial cable, which transmits the signals to the postamplifiers. Ten cables are used for each cell of eight wires: eight signals, one power and one calibration. The attenuation of the cable at 100 MHz is 46 db/100m for the 1.6 mm O.D. cable and 22.5 db/100m for the 2.8 mm O.D. cable. High frequency compensation circuitry on the preamplifier ensures that the preamplifier plus cable response is flat to ≈100 MHz.

The postamplifiers provide an additional voltage gain of 80 for the z-by-timing system and also split the pulses to the z-by-timing and $r - \phi$ FADC systems. The basic gain for each channel is provided by an NE592 video amplifier.

Maintenance of bandwidth throughout the system is crucial for the z-by-timing technique. The complete amplification system (preamplifier, cable and postamplifier) has a 3 dB bandwidth of 80 MHz. It is also essential to maintain low noise; the RMS noise amplitude from all sources is approximately 220 nV/$\sqrt{Hz}$ (dominated by the postamplifier). The amplification system described above results in average 10–90% risetimes of 10 ns for chamber pulses at the input to the z-by-timing readout cards.

5 The z-by-Timing Readout System

Signals from the postamplifiers are fed into the z-by-timing readout cards whose function is to digitize the time difference from each wire into a 7-bit number. The sensitive analogue circuitry of a single channel resides on an analogue 'daughter-board' which is equipped to make the timing measurement and provide its calibration. Four daughter-board channels are mounted on a single 'mother-board' which contains the digital readout circuitry. In this way a single cell of eight wires is serviced by two cards. The mother-boards are 400 mm deep and are housed in a 9U high Teradyne crate with a 20 mm pitch. Each z-Crate contains eleven z-by-timing cards, making 44 channels in total. The crate also contains seven or eight CTD FLT modules which process the z-by-timing data from that sector.

5.1 The Analogue Time Digitization

A block diagram of the front-end of a single analogue readout channel is shown in Fig. 5. The digitization process can be divided into stages. At the input of the card, the analogue 'start' and 'stop' pulses are sampled by constant fraction discriminators. The constant fraction of 0.5 and delay of 5 ns have been carefully chosen to provide accurate timing and to match the input pulse shape characteristics with minimal timing slew [19]. The discriminated pulse from the rear end of the chamber provides the start pulse to the TAC, the pulse from the forward end provides the stop. The additional 3 m of cable which runs along the length of the CTD from its forward end to the Services Ring ensures that the start pulse will always arrive before the stop for a valid hit. There is a programmable digital delay in the stop circuitry to finely tune the relative timing (see section 5.2).

The TAC generates a voltage level proportional to the time difference between the start and stop signals which is then sampled by an 8-bit FADC; a similar technique has previously been used by OPAL [20]. A timing diagram for the conversion

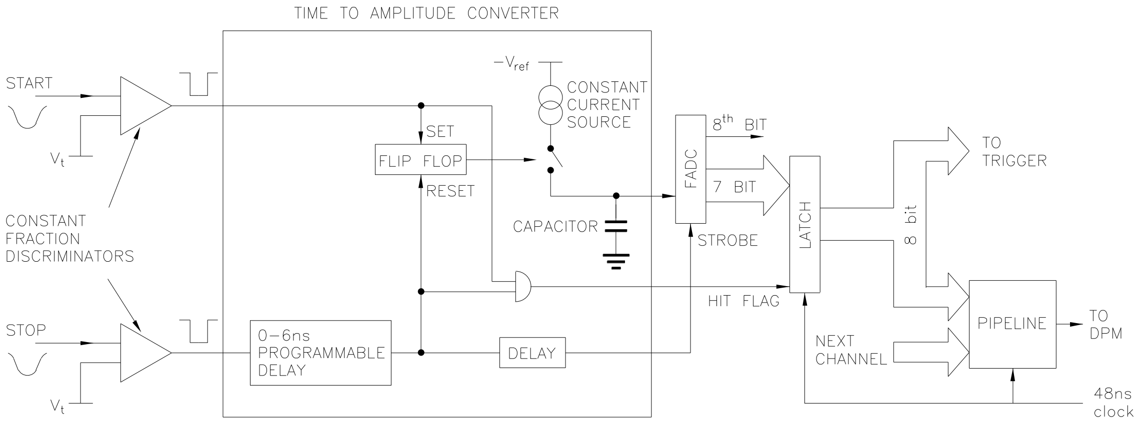

Figure 5: A schematic diagram of the front-end z-by-timing readout from Time-to-Amplitude Converter to pipeline. All circuitry, with the exception of the pipeline RAM, resides on a single analogue daughter-board.

is shown in Fig. 6(a). On arrival of a start pulse a capacitor is charged at constant current until a stop pulse is received. A strobe derived from the stop pulse (suitably delayed to allow for the maximum ramp time and for the settling of the flat top) triggers the sampling of the stored charge. This sampling is achieved by the 8-bit FADC; only the seven least significant bits are used to digitize the voltage, the eighth bit is used as an overflow. The 7-bit output of the FADC is subsequently combined with an 8^{th} bit validity flag ('hit flag') derived from a valid start-stop pair, shown in Fig. 5. The 8-bit data are then held in an output latch awaiting a synchronous copy into a RAM (pipeline) memory. A conversion will occur only if a start is followed by a stop within a specified time. If a valid stop has not been received within $\sim$40 ns after the arrival of a start, there is a 'time-out' and the ramp is reset ready for the next hit.

The full range of time differences in the chamber is $\sim$16 ns. This is derived from the full range of travel times of pulses down the 2.03 m wire length (nominally 13.3 ns), and additional contributions from small non-linear and end effects. Seven bit sampling was chosen so as to give a negligible contribution to the overall resolution arising from the TAC binning. (Note that the intrinsic timing resolution of the 7-bit TAC is $\sim$16 ns/$(128\times\sqrt{12})$=36 ps, equivalent to $\sim$5 mm of wire length. This represents a negligible contribution to the overall resolution of $\sim$4.4 cm, the measurement of which will be described in Section 8.)

The time between multihits of the z-by-timing measurement is determined by the maximum time between a valid start-stop combination ($\sim$16 ns), the sampling time ($\sim$15 ns flat-top), and the discharge and re-settling time of the capacitor ($\sim$10 ns). Added to this is the time for the ramp to reach the upper threshold level, chosen to be 5 ns, which ensures that operation is always on a constant slope. This results in a maximum ramp length of $\sim$21 ns for any valid start-stop combination. We have chosen to clock the system at 48 ns which is double the HERA beam crossing

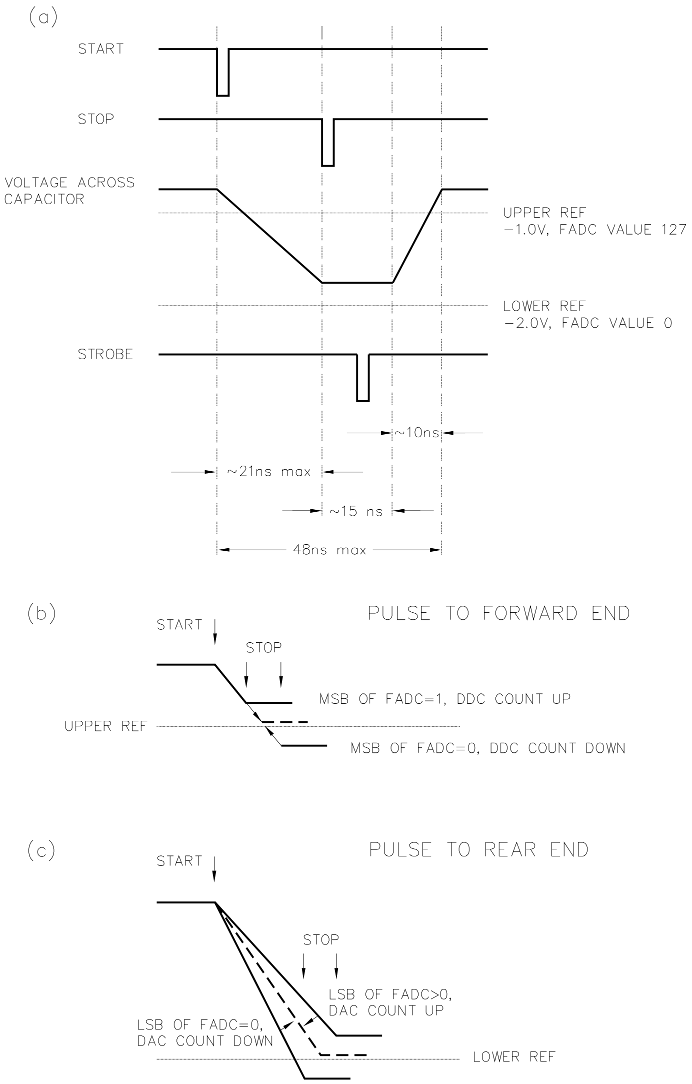

Figure 6: Timing diagram showing the voltage across the capacitor (a) in normal data-taking mode, (b) showing a calibration pulse to the forward end of the chamber and (c) showing a calibration pulse to the rear end of the chamber. 'LSB' and 'MSB' are the least and most significant FADC bits respectively.

frequency. Since the maximum conversion time of a valid hit is always less than 48 ns, we gate the system so that a second hit *cannot* occur within 48 ns of the first. Although a prompt conversion could in principle allow the system a faster recovery time, we have chosen a solution which brings uniformity to all channels for all hits. The second hit resolution of 48 ns nicely matches the equivalent multihit resolution of the $r - \phi$ FADC system, approximately 50 ns. This corresponds to a 2.5 mm drift distance.

Every 48 ns, the 8-bit data from the FADC output latch are clocked into a

8×1k bit RAM pipeline memory (details of the clocking procedure are given in section 7.4). In the usual event of null data from a given wire, zeros will be written into the pipeline; a zero hit flag means that no start-stop combination has been received. In the case of a valid hit, the location of the data in the pipeline memory relative to the beam crossing location in which the interaction occurred provides a coarse, independent measure of the $r - \phi$ drift time. The error on the drift-time measurement is approximately $48/\sqrt{12}$ ns and, in combination with the digitization of the z coordinate, gives a stand-alone three-dimensional space-point measurement. This is an important feature of the system.

The performance of the TAC digitization is demonstrated in Fig. 7(a). This shows the distribution of digitized time difference as a function of the reconstructed z coordinate of a hit for all instrumented wires, measured in the ZEUS experiment using tracks from $e - p$ collisions. The solid line represents the simplest linear form for the response, constrained by digitizings of 0 and 127 at the wire ends. As discussed in section 4, imperfect termination of the chamber sense wire results in non-linearities in the time-to-distance relationship. The deviation of the response from linearity is shown in Fig. 7(b), with a quintic fit superimposed.

The hit data from each FADC output latch are clocked to the CTD FLT [4], which determines whether a track is consistent with originating from the interaction region by considering z/r versus r patterns (where r is the radial position of the wire). The conversion from r to z/r is implemented using a preprogrammed ROM look-up table, located on the mother-board of each z-by-timing card. The data are also corrected for non-linearities into the FLT by preprogramming the look-up tables according to the quintic parameterization of Fig. 7(b).

5.2 Calibration Control

Calibration of the z-by-timing system is crucial to its successful operation. Since we require sub-nanosecond accuracy, the system is extremely sensitive to small delay changes, e.g. temperature fluctuations in cables and electronics. We have designed the system to calibrate out such changes either during data-taking or non data-taking periods. The logic necessary to perform z-by-timing calibration is implemented entirely in hardware with the only software interaction being the preparatory configuration of the front-end cards. For effective operation, a channel is calibrated when the full 7-bit dynamic range of the FADC corresponds to the full length of the chamber wire, i.e. the extreme FADC values 0 and 127 are output for tracks which deposit ionization at the wire end-points.

A schematic diagram of the components used in the calibration process is shown in Fig. 8. Timing diagrams for the calibration sequence are shown in Figs. 6(b) and (c). During calibration, pulses of a size and shape similar to real chamber pulses are sent alternately to the inputs of the appropriate preamplifiers at either end of the chamber (details of the pulse injection procedure will be given in section 6). These inputs closely approximate the physical ends of the sense wire. To maintain calibration, two parameters in the TAC are available for adjustment. The first

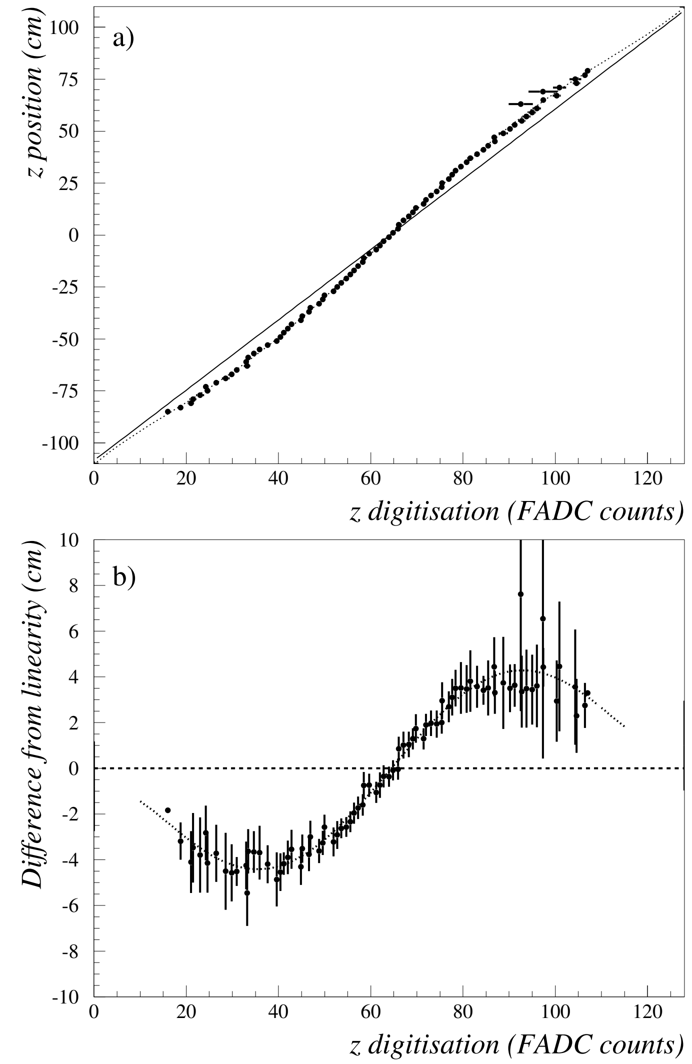

Figure 7: (a) The z-by-timing digitized time difference as a function of reconstructed z coordinate of the hit, measured using $e-p$ collision data. The solid line represents linearity, with the ends of the line constrained to the chamber end-points. (b) The deviation of the response from linearity. A fit to these data, represented by the dotted line, is described by an antisymmetric quintic function of the form $z = P_1(t-P_0) + P_2(t-P_0)^3 + P_3(t-P_0)^5$. This allows a conversion between time (in units of FADC counts) and z coordinate of the hit. The fit parameters are as follows : P_0=64.2; P_1=2.16; P_2=-2.30×10^{-4}; P_2=-2.99×10^{-8}.

is the relative arrival of the start and stop signals, controlled by an 8-bit 0-6 ns programmable Digital Delay Chip (DDC), which adds delay to the stop line. When the forward end is being pulsed, the value of the DDC is increased by a single bit if the FADC issues an overflow (i.e. the 8^{th} bit is set), or decreased by a single bit otherwise. Hence at the point of calibration the value issued by the FADC ideally oscillates between 127 and 128. The second parameter is the slope of the constant current ramp which charges the capacitor. When the rear end is being pulsed, an 8-bit Digital to Analogue Converter (DAC) is incremented if the FADC issues a value greater than 0 and this increases the slope of the current ramp. Conversely if the FADC value equals 0 the DAC is decremented and this decreases the slope. Hence calibration is ideally marked by oscillation between no FADC bits set and the first bit set. The delay and slope parameters are not independent, hence the need to pulse the ends alternately.

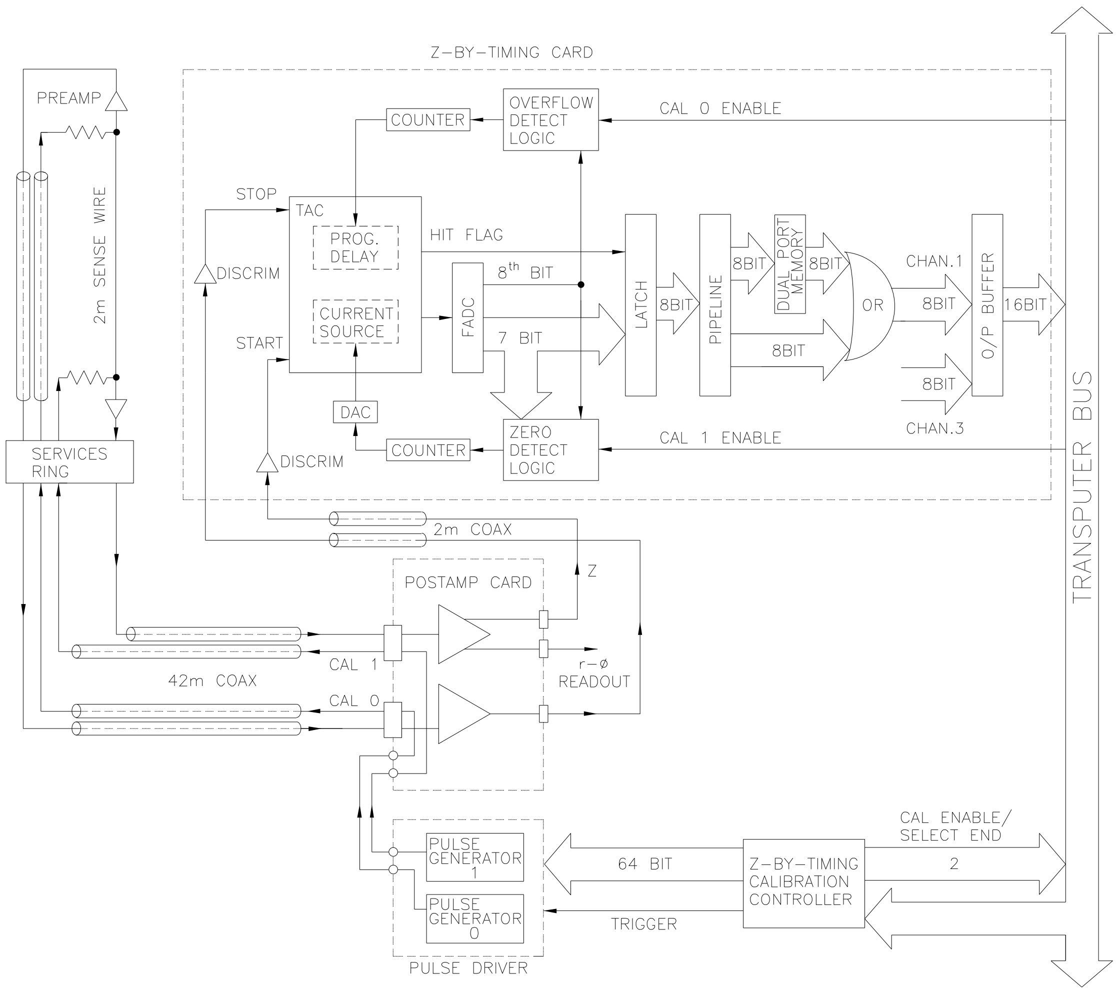

Figure 8: A block diagram showing the calibration control of the z-by-timing read-out system.

Each DAC and DDC is controlled by its own 8-bit up/down counter which is incremented or decremented according to the status of the FADC bits. Naturally the sensitivity of the DAC and DDC bits have to be much greater than one equivalent FADC bit; a transition of 6 and 4 bits of DAC and DDC respectively correspond to the transition of a single FADC bit.

Generally the z-by-timing calibration process is performed during non data-taking periods. A calibration 'run' is usually made up of a sequence of several hundred pulses to each cell of the chamber. Histogrammed values for typical DDC and DAC counters during calibration are shown in Figs. 9(a) and (b). Their mean values are calculated in real time during the calibration run and the subsequent values are down-loaded to preset the DACs and DDCs via the backplane bus prior to data-taking. The standard deviations of the DAC and DDC distributions shown in Fig. 9 are a measure of the noise and timing instability of the channel. The average standard deviations over all DAC and DDCs in the system are 1.9 and 1.1 counts respectively.

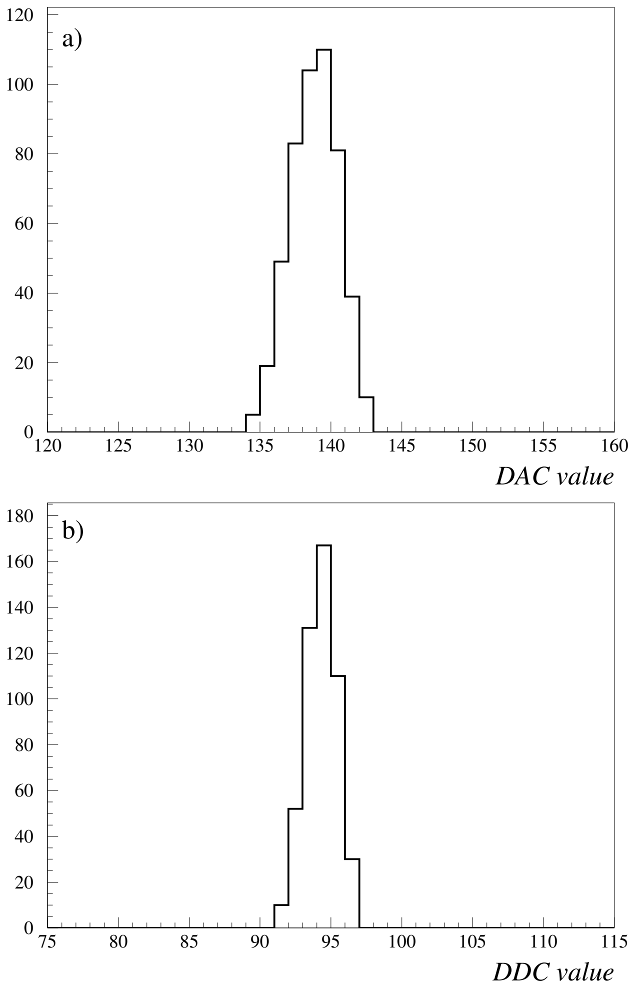

Figure 9: (a) Histogrammed values for typical DAC and (b) DDC counter values during a calibration run.

We have found that the time dependence of the DAC and DDC calibration points is extremely stable over the course of a HERA data-taking period. This is demonstrated in Figs. 10 (a) and (b). Here the mean values of typical DAC and DDC counters are shown as a function of time, covering a three year period. We have calculated the mean value of each individual DAC and DDC counter in the system over the course of this three year period. Considering how the DAC and

DDC values differ from their mean values over this time period, we subsequently obtain an RMS spread of 3.3 and 1.1 counts respectively for all DAC and DDC counter values in the system. Given this intrinsic stability, in principle there is need to recalibrate the system only infrequently.

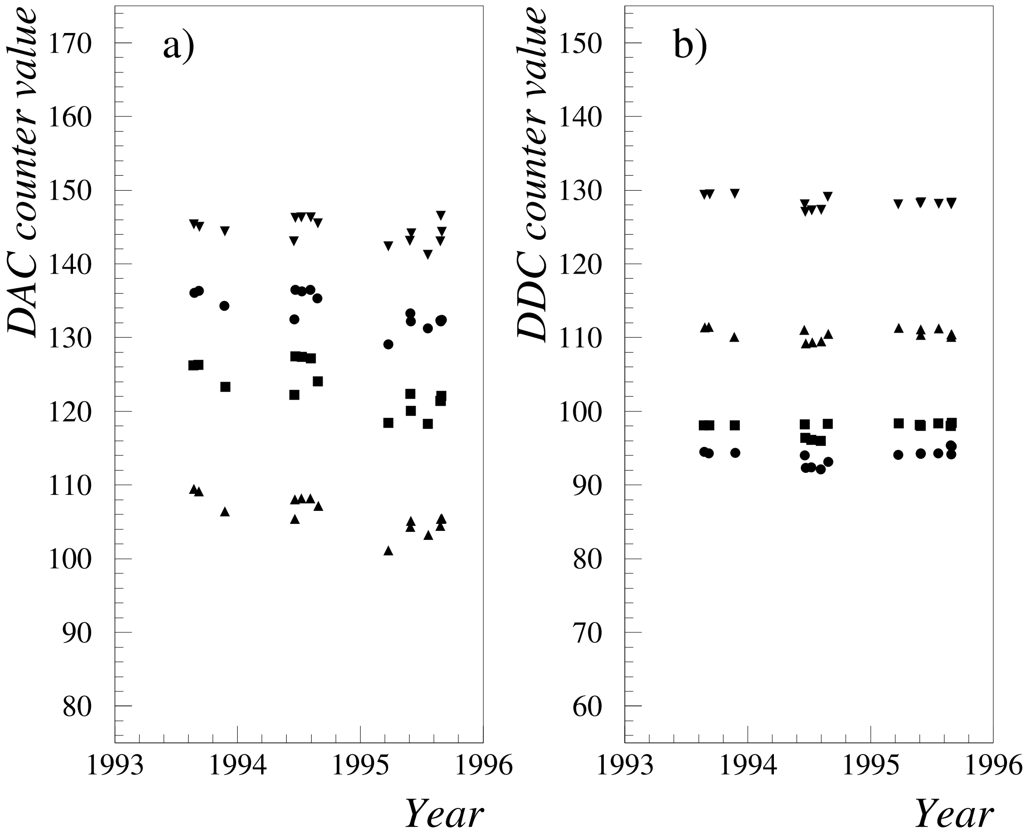

Figure 10: The mean values of (a) the DAC and (b) the DDC from the four channels of a typical z-by-timing readout card as a function of time. The data cover a three year period.

5.3 The Digital Readout

The mother-board contains four 8-bit RAM pipeline memories (PLMs) which store the data during FLT processing, four 8-bit dual port memories (DPMs) which act as a data store prior to backplane readout, and all the associated data transfer logic and backplane interface. All four channels operate synchronously and in parallel during data-taking. The pipelines, DPMs and calibration registers are all accessible (read and write) from the backplane data bus. The module can be operated in various modes by presetting a function code register (e.g. read or write of DPM, PLM; select CAL etc).

The 8-bit z-by-timing data are clocked from the FADC output latch (shown in

Fig. 5) into sequential locations of the PLM every 48 ns, the strobe being derived from the HERA 96 ns beam crossing clock. The length of the pipeline RAM is 1024 locations of memory which allows data storage for considerably longer than 5 μs, the time for the GFLT to reach a decision. However, during normal data-taking only the first 440 locations of the PLM are used, the data being overwritten every 220 beam crossings. This means that the pipeline address always corresponds to an equivalent HERA bunch crossing number, this feature simplifies the DAQ architecture and testing. The LTC module, described in section 7.4, provides the 48 ns clock pulses and the appropriate addresses for the PLMs via the crate backplane.

Events which are accepted by the GFLT have their z-by-timing data transferred from the pipelines into the DPMs, again under control of the LTC which supplies the appropriate addresses. During transfer, if a valid hit flag is encountered, all subsequent hit-flag locations in the DPM for that channel are set high. This speeds up the DAQ process by testing for valid hit data merely by reading the *last* word in the DPM pertaining to that event and channel. The size of each DPM is 8$\times$2k bits which is more than sufficient to buffer a maximum of 10 events which await readout (again see section 7.4).

The final stage of the readout is performed by a Transputer-based DAQ network. A single Readout Controller in the crate initiates the transfer of data from all DPMs in that crate to on-board Transputer memory, and subsequent Second Level processing.

6 The CTD Calibration System

6.1 Overview

The CTD calibration system serves a number of purposes :

- To provide a general diagnostic tool to check the electrical continuity of every instrumented wire from wire end to readout.
- To provide shaped pulses with risetimes of $\sim$1 ns to calibrate the z-by-timing system.
- To provide square pulses of varying amplitude for dE/dx calibration, required to compensate for channel-to-channel variation in gain.
- To provide square or shaped pulses of fixed amplitude to determine wire-to-wire time offsets (t_0's).

The possibility of crosstalk between pulsed cells means that only a selected pattern of cells within each sector are pulsed at a single time. We have therefore designed flexibility into the system, allowing *any* combination of cells to be pulsed independently or simultaneously.

The components which form the calibration system are as follows :

- A 'Calibration Controller' module, one located in each of the z and $r-\phi$ Crates. This is a programmable module whose purpose is to supervise the calibration sequence and provide timing synchronization.
- A 'Calibration Driver' module, one located in each Postamplifier Crate. This module generates pulses of pre-specified amplitude and shape which are then transmitted via the postamplifiers to the chamber. There is an individual pulse generator channel for every cell with readout.

A block diagram showing how the calibration modules are interfaced to the z-by-timing system is shown in Fig. 8.

6.2 The Calibration Controller

The calibration sequencing is supervised by 32 Calibration Controller modules. Two Controller modules service each $\frac{1}{16}^{th}$ sector; a Controller in the $r-\phi$ Crate supervises dE/dx and t_0 calibration runs, a Controller in the z-Crate supervises the z-by-timing calibration. Each Controller shares a 9U card with the respective LTC. The two Controllers in the z and $r-\phi$ Crates of a single azimuthal sector share mastership of that sector's Driver module; this mastership is mutually exclusive.

A schematic diagram of the Calibration Controller is shown in Fig. 11, to which much of this section refers. The Controller contains its own logic and memory which must be preloaded with a calibration 'run' definition. This is then successively triggered (either under software control or externally from the GFLT) to enable the specified pulse sequence. A 64-bit word, loaded in the Controller memory, contains the bit map of the cells in that sector which will be pulsed simultaneously from a single trigger, along with bits to preset the pulse shape and amplitude. This single pulse definition word is latched into the Calibration Driver situated in the Postamplifier Crate of that sector; the Driver subsequently generates the appropriate calibration pulses on receipt of the trigger from the Controller. The pulses are then sent to the chamber via the postamplifiers. When the Driver is triggered and the pulses sent, the next calibration pulse definition is automatically down-loaded from the Controller to the Driver. The total size of Controller memory is 64×65k bits consisting of four 16×65k fields, however only the first 16k locations are used. Hence 16k independent calibration sequences can be sent, which then wrap around on subsequent triggers. The wrap-around length is set by the 'Sequence Restart Pointer' and can take any value between 2 and 16k. The Controller uses a page system to communicate with the four fields of memory; which field of memory is selected for access by the address bus is controlled by the 'Internal Page Register', a 4-bit control word (CWF) points to the required field and this must be preset prior to memory access. Control and status of the Calibration Controller is provided via a 16-bit read/write register, the 'Control/Status Register'. All registers are readable to scrutinize the instantaneous calibration status.

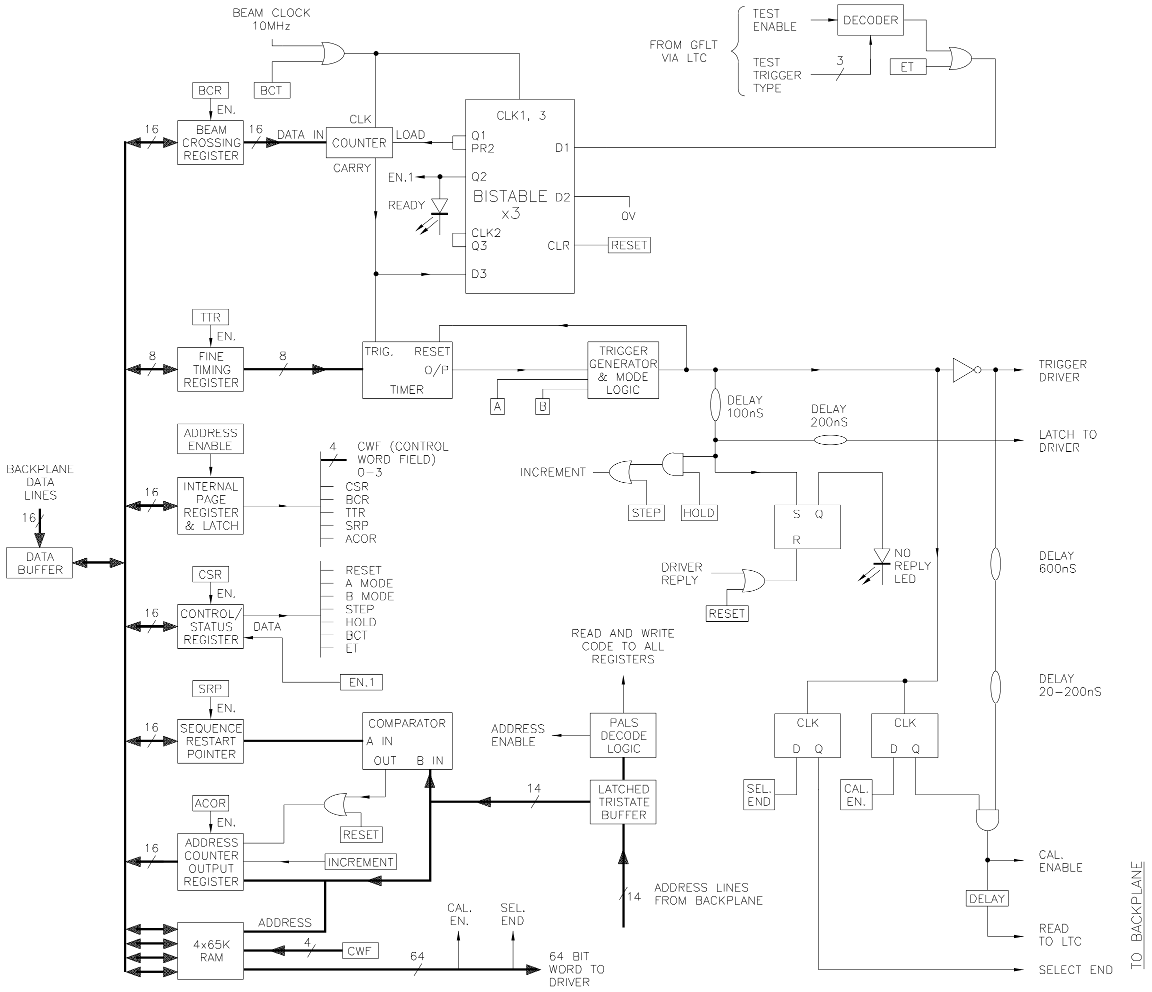

Figure 11: A schematic diagram of the CTD Calibration Controller.

During a z-by-timing calibration run the Controller sends calibration control signals to the z-by-timing cards via the backplane. These signals enable the calibration DAC and DDC counting logic ('Cal. Enable') and define the end which is being pulsed ('Select End'). The Cal. Enable signal is a gate which is made wide enough to accept the full range of arrival times of the analogue chamber pulses at the input of the z-by-timing cards, but short enough to disable the DAC and DDC counting logic when outside this range (i.e. to protect against calibrating on spurious noise). The relative timing of this signal is defined by a lumped delay (typically set to 600 ns) on each Controller, which allows for the return time of the pulse from the chamber.

Three trigger modes are possible, selected by setting appropriate 'Mode' bits in the Control/Status Register. The first mode is 'one-shot', used in stand-alone calibration runs. Here the trigger is under software control and is used when calibration is performed prior to a run or during beam-off. The second mode is 'free-run'

(100 Hz), used for diagnostic testing as well as stand-alone calibration. The 'Trigger Generator' contains the free running trigger generator and the one-shot, also the logic to decode the 2-bit Mode levels from the Control/Status Register. The third mode allows calibration during data-taking and here the overall sequencing of CTD calibration is derived from the GFLT which initiates a prescheduled 'Test Enable' sequence, synchronized to the beam clock. In this case calibration can occur during the empty beam buckets of the HERA machine (HERA has 220 bunches, of which 10 are sequentially empty). The timing diagram for such a calibration run performed during data-taking is shown in Fig. 12. The calibration run definition is down-loaded into the Calibration Controller memory before the run starts. A total of 220×100 beam crossings before a calibration trigger Accept, the GFLT sends a 'Test Enable' signal to the CTD. The programmable 'Timer' delay is set such that when the Controller triggers the Driver and calibration pulses are subsequently sent to the chamber, their z-by-timing digitizations are timed to appear in the pipelines just before the GFLT issues its Accept.

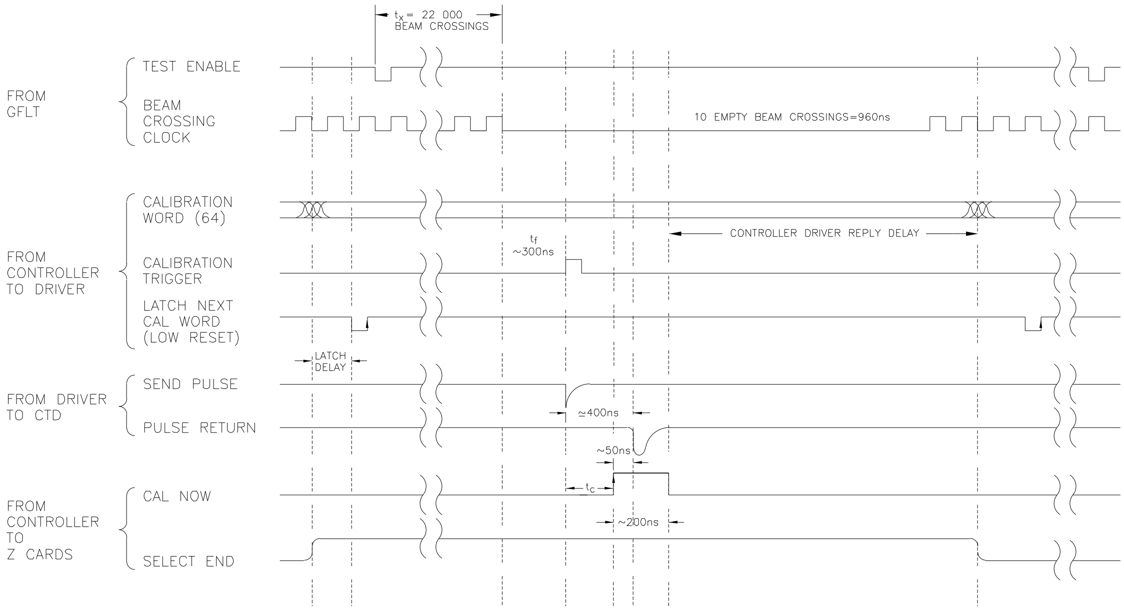

Figure 12: A timing diagram of the CTD Calibration Controller sequence during a data-taking run. The delays t_x, t_c and t_f are either set in hardware or are programmable.

On the Controller, a selectable number of beam clock cycles (contained in the 'Beam Crossing Register') and a fine adjustment delay (contained in the 'Fine Timing Register') can be preselected to set the delay between the GFLT Test Enable signal and Calibration Driver triggering. The timing of the calibration sequence is controlled by the action of three bistables. The Test Enable arrives at D1 of the first bistable and the next beam clock latches Q1 high which loads the contents of the Beam Crossing Register into the counter. This register contains the number of beam clocks after which a calibration trigger is sent. Q1 also presets the second bistable (PR2) high which lights the 'Ready' LED. When the Test Enable pulse goes low, on the next beam clock the counter is enabled to count all subsequent beam clock pulses. When the specified count is complete, the carry bit goes high. This then triggers the calibration sequence after a preprogrammed delay, defined by the Timer device (which has been previously preloaded from the Fine Timing Register). The carry bit is also input to D3 of the third bistable and through Q2 the Ready LED is extinguished on the next beam clock. The three bistables can be reset from software via the Control/Status Register.

On receipt of an output pulse from the Timer, the Trigger Generator sends a trigger pulse to the Driver as well as a delayed 'latch' pulse. This latches the memory output into the Driver in preparation for the *next* trigger. During a z-by-timing calibration run, 'Select End' oscillates alternately high and low on every trigger pulse.

6.3 The Calibration Driver

A single Calibration Driver module resides in each Postamplifier Crate. The function of this module is to generate pulses which are then sent passively via the postamplifiers down the ~45 m of coaxial cable to the preamplifiers. Charge is thus injected onto the preamplifier inputs, as close as possible to the physical ends of the sense wire. The Driver contains 46 pulse generators, each of which can produce a fast shaped pulse or a square pulse. Each pulse generator represents a single channel, and sends pulses to a single cell. Hence all eight wires of a cell are pulsed simultaneously. During a z-by-timing calibration run, fast chamber-like pulses are generated and sent to each end of the chamber alternately.

A schematic diagram of the Calibration Driver is shown in Fig. 13. The Driver receives the 64-bit control word from the Calibration Controller and on receipt of a trigger sends out a shaped or square pulse of a specified amplitude to the preselected channels. The 64-bit word comprises 46 channel select bits, one pulse-shape bit, 8 pulse amplitude bits and additional control bits. Alternatively the Driver can be used independently from the Controller by means of front-panel switches. The variation of pulse amplitude (up to a maximum of 1.2 V) is controlled via an 8-bit DAC.

Each Driver circuit is in the form of an individual hybrid (46 per card). A circuit diagram is shown in Fig. 14. Two DC inputs to the hybrid (V_Z and V_R), one for shaped pulses and one for square, determine the ambient current and hence define the amplitude of the output pulse. The shaped pulse is generated by switching a differential pair of high-speed transistors from one state to the other. When switching occurs, current flows through Q2 causing a voltage drop at the collector of Q1, this step voltage is then differentiated by C2. This results in a fast output pulse with a risetime of 1 ns. Dispersion in the 45 m cable increases this to ~2 ns at the preamplifier input. The maximum pulse duration is 20 ns (FWHM 5 ns) and the absolute timing jitter is less than 25 ps. The channel-to-channel amplitude variation is of the order of 5%. The square pulse is similarly generated, however here

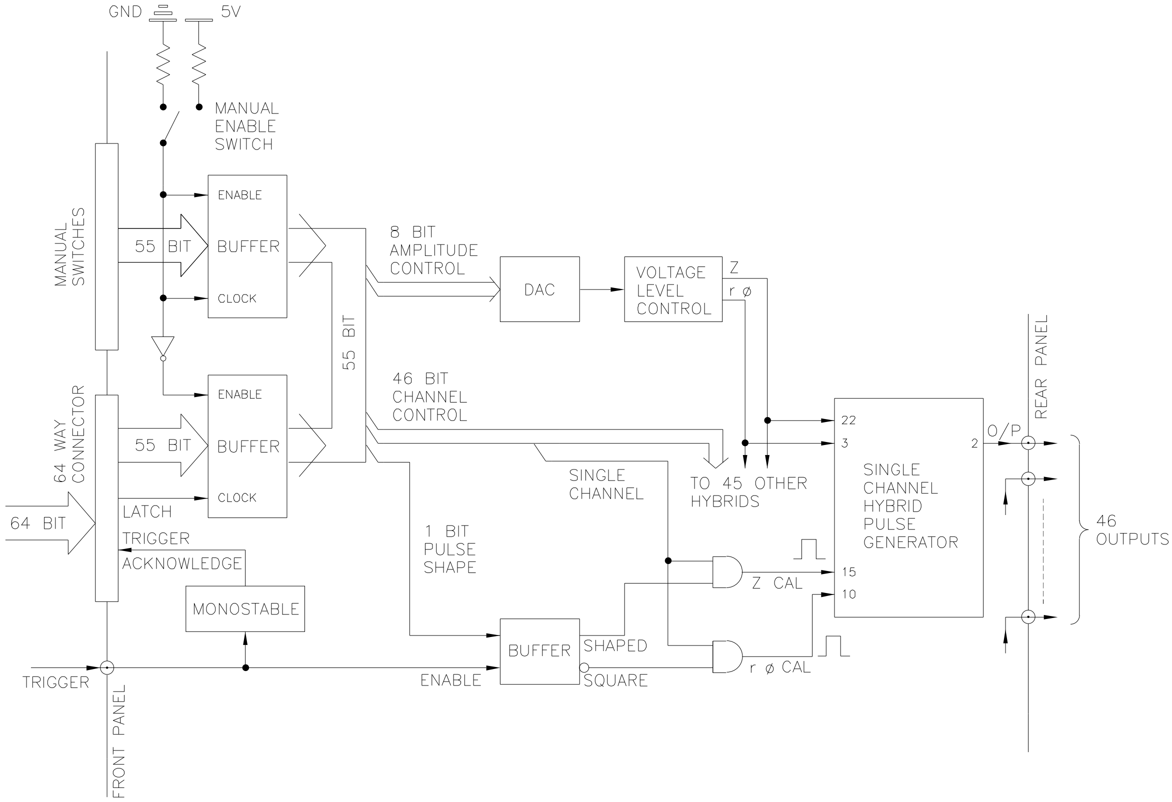

Figure 13: A schematic diagram of the CTD Calibration Driver.

the switching transition is not differentiated and the negative-going output pulse has a width defined by the trigger signal (approximately 250 ns) with a risetime of 10 ns at the Driver output under no load. *RC* effects considerably increase the pulse risetime from the chamber and hence the trigger pulse has to be chosen wide enough so a flat top is reached. The channel-to-channel amplitude variation of the square pulses is controlled within one card to about 2%, between cards to 5%. The pulses from the hybrid circuits typically have peak-to-peak noise levels less than 5 mV; those with levels greater than 10 mV were rejected.

7 The z-by-Timing Data Acquisition System

7.1 Overview

Local data acquisition of the z-by-timing data is supervised by a Readout Controller, one per crate. Its purpose is to read out and zero-suppress the data, perform calculations on the data (in particular SLT processing) and to pass the data onto the ZEUS Event Builder for subsequent processing. Timing control is provided by a single Master Timing Controller interfaced directly to the GFLT, and 16 Local

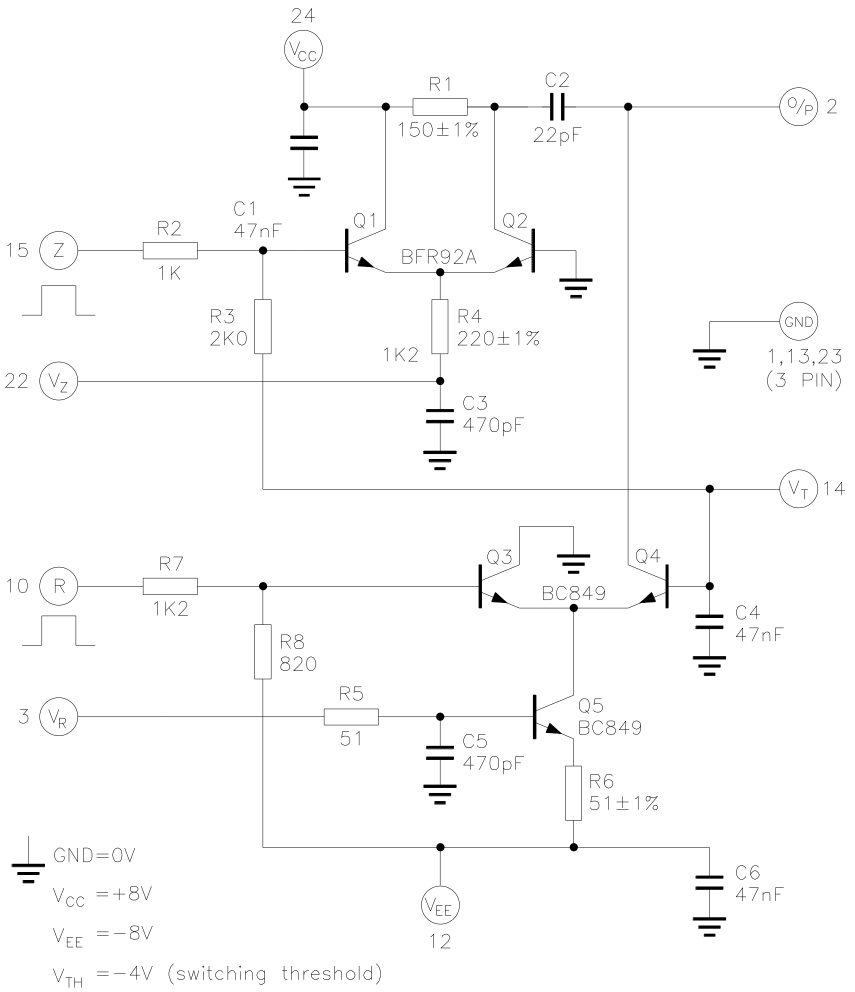

Figure 14: A circuit diagram of one of the 46 Calibration Driver hybrid circuits.

Timing Controllers, one per z-Crate.

The Timing Controllers have the following functions :

- To synchronize data from the 16 z-Crates.
- To interface with the GFLT.
- To control the pipeline and buffers of the front-end electronics.
- To interface with the DAQ software.

Fig. 15 is a schematic diagram showing the functionality of the timing system and readout control. The Master and Local Timing Controllers provide timing to both the z and the CTD FLT systems. The design philosophy of MTC/LTC communication is one of simplicity and robustness; all timing is derived from a single clock and the only control signal returning from the LTC to the MTC is a 'Busy'.

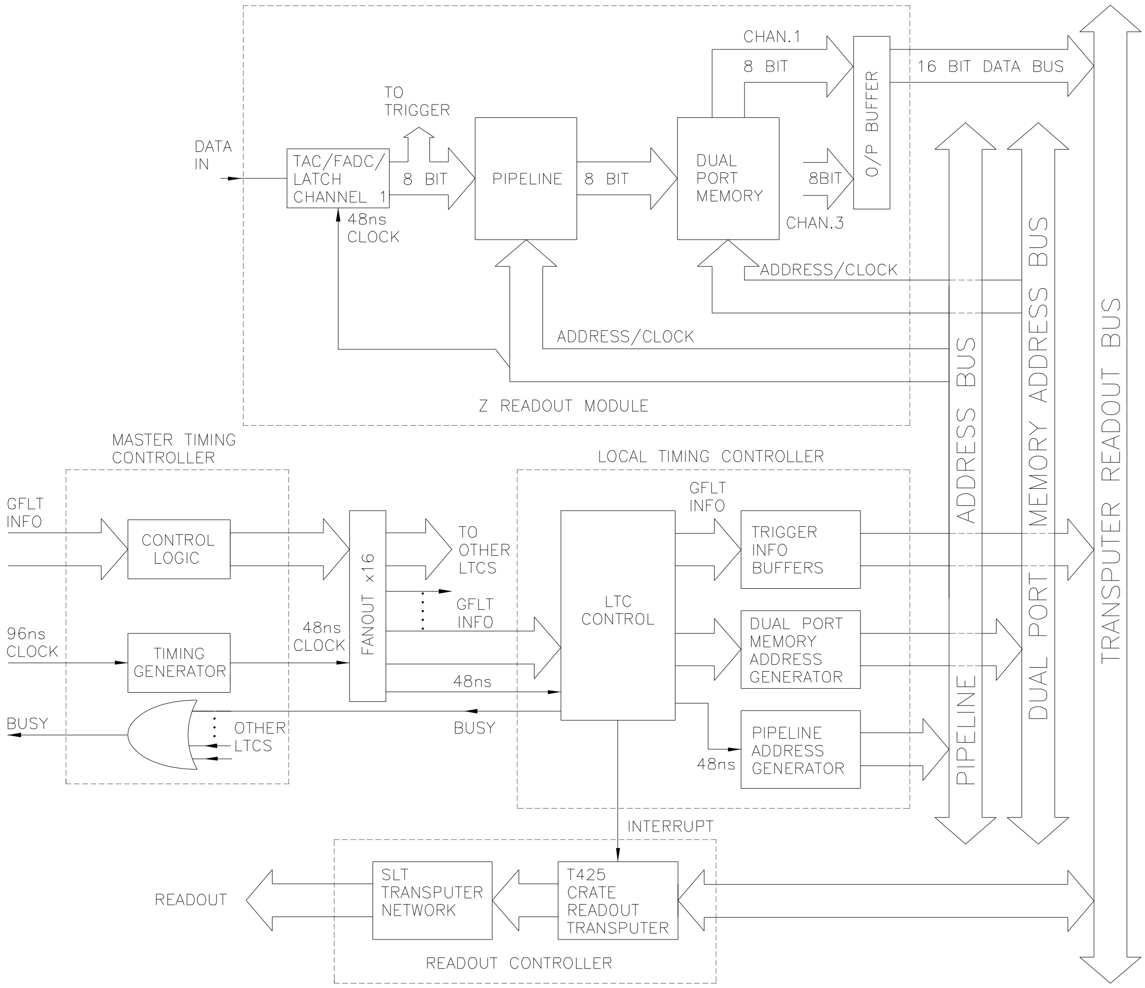

Figure 15: The local z-by-timing DAQ system, demonstrating the functionality of the Timing Controllers.

7.2 The Readout Controller

The CTD DAQ and SLT processing system is based on Transputer-based ROCs [16, 17]. The ROC is a single 9U card, one per crate. The z-Crate ROC contains two T425 Transputers with a total of 3 Mbytes of external memory and three 'TRAM' modules, each containing a T800 Transputer with 1 Mbyte of external memory [21]. The ROC operates a backplane bus to down-load parameters to and read data from the front-end cards. Each Transputer is equipped with a 10 MIPS processor, on-chip memory, and four bidirectional 20 Mbit/s serial links which make it ideal for the application of SLT parallel processing. In response to an interrupt from the LTC, readout of the crates is carried out by one of the T425 Transputers on each ROC, known as the 'Readout Transputer'. The Teradyne backplane is mapped directly onto the Readout Transputer memory, hence one can down-load and read data from each of the front-end cards by simply reading from and writing to locations in memory. The second T425 performs data transport. The three T800 Transputers are dedicated to running the software-based SLT algorithm [11, 22]. Fig. 15 shows the interface between the front-end z-by-timing cards, the LTC and the Readout Transputer via the address and Transputer buses of the crate backplane.

7.3 The Master Timing Controller

The main purpose of the MTC is to synchronize the CTD readout and FLT processing to the 96 ns HERA beam crossing clock (and hence to the rest of the experiment), and to provide the interface by which the DAQ system is able to recognize a trigger Accept. The MTC resides in a dedicated crate (the so-called 'RBOX' Crate) which also houses MTC fanout units and the cards responsible for the final stages of the First Level track trigger processing and the interface to the GFLT. Although its operation during normal data-taking is relatively simple, its test capabilities make the circuitry rather complex.

A functional diagram of the MTC is shown in Fig. 16. The MTC receives the 96 ns beam crossing clock from the GFLT, derives a 48 ns clock using a phase-locked loop, and this clock is subsequently fanned out to each of the 16 LTCs. In addition to the clock, the MTC receives a number of signals from the GFLT:

- 'Beam Crossing Zero' (BCN0) is a signal which defines which clock pulse corresponds to the beam crossing numbered zero. Its role is to synchronize the pipelines throughout the system by resetting the LTC pipeline addresses to zero (see section 7.4). Timing of this signal is important, hence the MTC synchronizes this signal to the 48 ns clock before it is fanned out to the 16 LTCs. Although the pipeline generation is done on the LTCs, there is also a pipeline counter on the MTC which serves two important functions. Firstly it enables the MTC to detect a 'pipeline jump' error if the BCN0 signal from the GFLT goes out of synchronization (in which case the MTC will flag an error bit). Secondly, when running in stand-alone test mode, it enables the MTC to generate its own BCN0 signal which is then sent to the LTCs.

- 'Accept' and 'Abort' are the signals that control the system. The GFLT Accept signifies an event trigger and hence initiates a pipeline copy. The Abort initiates an abort of the pipeline copy and on receipt of this signal the system becomes free to restart clocking data into the pipelines. The MTC decodes and time-synchronizes these signals with the 48 ns clock before passing them on to the LTCs. Likewise the Test Enable signal passes through the MTC and thence to the LTCs and Calibration Controllers. This signal instigates a calibration cycle.