BR-HEP 94-2

Electronics of The ZEUS Central Tracking Detector

A. Cassidy et al.

Invited paper presented by R. Tapper,
IV th International Conference on Electronics for Future Colliders Chestnut Ridge, New York, 11^{th} - 12^{th} May, 1994.

ELECTRONICS OF THE ZEUS CENTRAL TRACKING DETECTOR

A. Cassidy, D. Cussans[1], B. Foster, R. Gilmore, G.P. Heath, H.F. Heath, M. Lancaster[2], T.J. Llewellyn , J. Malos, C.J.S. Morgado, T.L. Short[1], R.J. Tapper, S.S. Wilson
H.H.Wills Physics Laboratory, University of Bristol,UK

D.S. Bailey, D.M. Gingrich[3], N. Harnew, P. Hallam-Baker[4], T. Khatri, I. McArthur, J. Nash, P.D. Shield, S. Topp-Jorgensen, F.F. Wilson
Department of Nuclear Physics, University of Oxford,UK

D. Allen, S.A. Baird, R. Carter, S.B. Galagedera, M.D. Gibson, R. Halsall, P. Hardy, R.S. Hatley, M. Jeffs, S. Jaroslawski, S. Madani, R. Matson, A. Middleton, R. Millborrow, M. Morrissey, A.G. Parham, S.P.H. Quinton, D.J. White
Rutherford Appleton Laboratory, Chilton, Didcot, UK

J.B. Lane, G. Nixon, M. Postranecky
University College, London, UK

A.K. Jamdagni, D.B. Miller, L. Toudup
Imperial College of Science, Technology and Medicine, London, UK

[1] = now at Rutherford Lab. [2] = now at Oxford
[3] = now at University of Alberta [4] = now at CERN

ABSTRACT

We describe the electronics built for the central tracking drift chamber of the ZEUS experiment at the HERA electron-proton collider at Hamburg. The system comprises a 104 MHz FADC system connected to all 4608 sense-wires of the chamber and a z-by-timing system connected to a subset of the wires. The latter provides prompt information to a pipelined trigger processor able to identify tracks originating at the intersection region, and contributes to the experiment-wide first level trigger produced within a few microseconds of the event. A DSP on each sixteen-channel FADC card locates and parametrizes signals on the wires on its card to produce values for the size and time of arrival of any signals. These are made available in a dual-port memory from which they are read by a transputer-based readout controller.

§*Introduction.*

In this paper we describe the electronics built for the central tracking drift chamber of the ZEUS experiment at HERA electron-proton collider at Hamburg. This complex system is now commissioned and is running routinely as an essential part of the experiment. In a number of respects the design goals of this system were more demanding than those of previous systems, and it was required to address, albeit in a much milder form, some of the problems which will face DAQ/trigger systems at LHC. This detector, and its electronics, were designed and built by personnel from a small group of UK universities and from the Rutherford Appleton Laboratory.

§*The HERA machine.*

HERA provides countercirculating bunches of 30 GeV electrons and 820 GeV protons which pass through one another in a beam pipe at the centre of the detector. The bunch crossings occur at intervals of 96 ns although all available "buckets" may not be populated by a bunch of beam particles. This bunch crossing interval is to be compared with 22 μs at LEP, and is one of the critical parameters for the read-out design. The design luminosity was $1.5 \times 10^{31} \mathrm{cm}^{-2}\mathrm{s}^{-1}$ and the intersection region is approximately 25cm along the direction of the beams and $\times(0.5\mathrm{mm})^2$ transverse to them. HERA is a unique machine, and although the problem of background tracks from interactions between the proton beam and residual gas in the beam-pipe was expected to be a major source of difficulty, it was difficult to forsee how severe it would be. It was clearly essential that a highly selective trigger scheme should be used, and that the ability to program and configure the system at a low level should be as far-reaching as possible, to allow the system to cope with possible severe conditions in the chamber.

§*The Drift-Chamber.*

The large-scale organisation of the electronics, as well as many smaller-scale features of the system, reflects the structure of the drift-chamber itself, so a brief description is required. This design of this chamber was based on the CDF central tracker [1].

The central tracking detector ("CTD") is a cylindrical drift-chamber 2 metres in length by 80 cms in diameter containing some 20,000 wires, of which 4608 are small-diameter anode sense-wires read out by the DAQ system, the remainder being cathode and field-shaping wires. The chamber is situated inside a superconducting solenoid maintaining a field of 1.8 Tesla over the whole of the CTD. The drift-chamber is organised as a nest of nine concentric cylindrical "superlayers", each made up of repetitions of a "cell" containing eight sense-wires arranged in a plane making a fixed angle of 45 degrees with the outward radius through the cell. High momentum (i.e. straight) tracks from the intersection region cross two adjacent cells in each superlayer. The detailed layout of the cell is strongly influenced by the presence of the magnetic field, which causes drifting electrons to travel at an angle to the electrical field. This "Lorentz angle", which depends on the gas mixture used and on the size of the electric field was chosen to be 45 degrees when the chamber was designed. This, of course, severely limited the range of possible gases and fields which could be used in the chamber. The size of the cells determined the maximum drift-time, which was approximately .5 microseconds. Sense wires in superlayers 1,3,5,7 and 9 were parallel to the axis of the chamber and provided precise measurements of position in the transverse plane, but no information in the axial (or "z") direction. To provide such information the sense wires in the even-numbered superlayers were inclined at an angle of about 5 degrees to the axis. Information from these "stereo layers" could be combined with that from the axial layers to yield full spatial reconstruction in the off-line analysis. The chamber was divided into sixteen equal sectors each of 22.5 degrees in azimuthal angle for the purposes of readout (see below).

§*Overview of the Electronics*

The main readout from all the sense wires was chosen to be by a flash analogue-to-digital converter ("FADC") system operating at 104 MHz, on the grounds that such a system could provide multi-track capability and pulse-size information, but most importantly that it could be designed to be very flexible in operation if processing intelligence was introduced as near to the front-end as possible.

It was decided at an early stage to provide, in addition, prompt information on the z-coordinate which could be used to reject tracks not originating from the intersection region. To this end, a subset of the axial sense wires was instrumented with a z-by-timing system capable of determining the z-coordinate with a few cms precision within a few microseconds. This data is transferred rapidly to trigger modules, which identify tracks pointing back to the intersection and provide a contribution from the CTD to an experiment-wide global first-level trigger ("GFLT") system which is not described in this paper. A schematic overview of the electronics is shown in figure 1.

§*Crates and Boards*

A large-format ("fast-bus" sized) board of height 9U (365mm) and depth 400mm was chosen for the main part of the electronics, and a twenty-one slot custom crate with integral power supplies was designed. All the 32 crates in the system were provided with a multilayer custom backplane. The lower 3U section consists of a relatively simple readout and control backplane to which connections are made via high-density Teradyne connectors. Half of the crates were used for the FADC system, for which only this readout section is required. The remaining crates, housing the z-by-timing/trigger system, required in addition on the top 6U section of the backplane a very complex pattern of interconnections (also using high density Teradyne connectors) to provide for the rapid transfer of large volumes of information from the z-by-timing modules to the trigger modules.

§*Readout Controller*

It was decided to base the system-wide readout upon the INMOS T425 "transputer" [2]. These processors have four serial links each capable of transferring data at 20 Mbits/sec. and can be interconnected in a complicated network. Program development tools provided by the manufacturers can be used to configure and boot the network, and programs running in the nodes of the network are able to communicate easily using the simple protocols defined in the "occam" language.

Each crate in the entire system is equipped with a transputer-based readout controller ("ROC"). With appropriate buffering, the least-significant 17 address lines of the transputer are extended into the backplane, while five further address lines are decoded to provide a slot select signal. In this way a range of addresses in the readout transputer is associated with each card in the crate, and using the occam language block transfers in or out of the cards can be achieved with single-line commands. The readout transputer is connected by a serial link and a shared dual-port memory to a second transputer for

processing and communications. Each transputer is provided with up to 4 MBytes of DRAM. Further processing power may be provided where necessary to meet the needs of the second-level trigger (see below) by plugging in up to six "TRAM" cards, each of which contains a processor with further RAM. The transputer network was loaded and controlled by software tools provided by INMOS, and further tools developed within the collaboration.

§*Preamplifiers and Postamplifiers*

The main part of the electronics was situated in a counting room some distance from the detector. To provide adequate noise immunity and drive capability for transmission over cables approximately 42 metres long preamplifiers were provided on the end face of the chamber. A simple discrete component circuit using surface-mount devices provided a voltage gain of 2.2 with about 16 mW dissipation. By peaking the amplifier gain it was possible to compensate for some of the high-frequency losses in the cable, thereby producing an approximately flat response up to about 104 MHz. See [3] for more detail. Calibration signals were sent from the counting room along a separate cable and fed into each channel of the readout system at the pre-amp card. In this way as much as possible of the electronic chain was checked, and propagation delays could be measured including the cable from the chamber.

Postamplifiers in the counting room received the signals from the pre-amps and amplified them in order to drive separate outputs to the FADC and z-by-timing systems. The design was based on the NE 592 amplifier, and gave a gain of 50 to the z-by-timing system, and 25 to the FADC system. A pole-zero filter on the output reduced the slow tails characteristic of drift-chamber pulses in order to improve the resolution of closely-spaced signals.

§*FADC System: General Considerations*

For optimal spatial resolution and two-track discrimination the analogue signals from a drift chamber sense-wire should be digitized as frequently as possible. In practice the cost of flash ADC's and suitable memory devices increases rapidly with the digitizing speed, and many effects in the physics of the chamber limit also the spatial resolution, so the cost of digitizing at rates higher than 104 MHz or so cannot be justified. During the development of the system we are describing eight-bit 104 MHz devices cheap enough to be considered became available from Plessey, and we based the final form of the system on their SP97508 device.

This has ECL outputs, and so each FADC was directly connected to a pair of four bit wide by 1024 word deep ECL-compatible memory chips, forming together the 10 microsecond by 8-bit wide "pipeline memory". The high capacitance input of the FADC was driven by a high-speed amplifier (SL9999) designed for this role (also from Plessey). It was decided to isolate this amplifier, the FADC and the ECL-compatible memories on a hybrid circuit 1.1 by 2.4 inches in size. It was expected that this would provide some measure of electrical isolation and facilitate the commissioning of the system by allowing for easy replacement of the most performance-critical components in each channel. In

addition the ceramic substrate of this hybrid would provide for more effective cooling than direct mounting of these components on the motherboard of the FADC module, and reduce thermally-induced mechanical stresses because its thermal expansion more closely matches that of the integrated circuits.

The ECL memories were driven in a non-standard way, with the Write-Enable signal permanently asserted. This allows cheaper lower-speed parts to be used and was found to result in satisfactory data storage provided the transitions on all address lines were arranged to occur simultaneously to within a nanosecond or so. By using Gray code rather than normal binary addresses for both data storage and readout the data was entered in a way which involved only one address line changing at a time.

The system dissipated about 5 watts per channel, mainly in the ECL circuitry on the hybrid, and provision for cooling had to be made (see below).

An essential feature of the system is the provision of front-end intelligence, in the form of a digital signal processor, on each module.

§*The FADC Module, Local and Master Timing Controllers*

Cost and space considerations favoured the highest possible density in terms of channels per crate, but at the same time we decided to avoid the added complexity involved in any scheme which divided the sense-wires from a chamber cell between more than one FADC module. Using the large board-size it was possible to fit sixteen FADC channels (i.e. two cells) into each module. The addresses for all the pipelines in a module were generated by a local scaler counting a 104 MHz master clock and were distributed within the card as ECL signals. Not surprisingly the arrangement of the hybrids on the card and the precise arrangements for driving them were found to be critical, and required several iterations before a satisfactory solution was found.

The master clock was generated by a phase-locked loop locked to a harmonic of the HERA machine clock situated in a single master timing controller ("MTC") and distributed throughout the system as a differential ECL signal to the front-panel of each LTC module. On receipt of a first-level trigger the scalers are stopped and a transfer of a preselected time-window of data from the pipeline memories to "primary buffers" takes place. The latter are 30 MHz CMOS memories shared by a group of four channels, and subdivided into four fixed buffer regions. The addresses for the pipelines and the primary buffers during the copy process are generated by the master timing controller, the addresses from which are replicated throughout the system by local timing controllers in each FADC crate, which distribute them to the modules in the crate via dedicated lines on the backplane. On system initialization, the MTC is provided with the parameters needed to control the transfer, namely the number of clock counts to jump back, and the width of the window to copy, also in clock counts. Under normal circumstances the window size was chosen to provide for the maximum drift-time of 500 ns. and adequate additional time for the full time-development of a signal, while the jump-back distance allowed for the formation of the first-level trigger.

When the MTC receives a first-level trigger the jump-back count is subtracted from the current value of the pipeline address scaler to provide the first pipeline address for the transfer, while the first address in the primary buffer memory is the first address of

the next free buffer. The transfer of the specified number of words takes place at the reduced clock-speed required by the primary buffer. During the transfer each data-byte is compared with a previously chosen threshold by a magnitude comparator chip which is able to set a "hit-flag" bit associated with the channel in question. After the transfer has been completed, the 104 MHz clock is restarted and data-taking recommences, while the data in the primary buffer is now available for processing by the module's digital signal processor.

§*The Digital Signal Processor*

After initial experiments with Texas DSPs, we decided to use the Motorola 56001 DSP in this system. It was regarded as an important objective to be able to download the DSP program, rather than to use ROM-based software. This was achieved by using the "host interface" of the DSP, which was mapped onto a number of lines on the backplane of the crate, and could be accessed by programs running in the ROC transputer. During the early stages of development is was possible to restrict the DSP program's size enough to fit it into the on-chip 512-word program memory directly accessible to the host interface. Anticipating that this would eventually prove to be a limitation, the DSPs were provided with 8K words of external program memory and, as more sophistication was introduced into the data handling in the DSP, it became necessary to write a loader. This was itself downloaded through the host interface, but then copied itself into external program memory and read a further ASCII file in the Motorola .LODfile format which could load any part of the P-, X- or Y- memory spaces of the DSP, including the on-chip P-SRAM. Because of the architecture of the 56001, program running in the external program memory is slower than that running on-chip, but division of the DSP tasks between those which were time-critical and required the internal memory and others minimized the effect of this.

In the simplest situation the DSP simply scans the regions of the primary buffer associated with each channel using a simple pulse-finding algorithm. When pulses have been identified their time of arrival is obtained using a constant-fraction algorithm, and a measure of their area is obtained by recording their amplitude. Within the 1 millisecond average interval between first-level triggers there is time for this operation on all sixteen channels, provided only a few pulses are found. It is possible to avoid wasting time by scanning empty channels if the hit flags are used to guide the DSP. Although this is valuable in the outer superlayers, where the channels are sparsely populated with signals, it is of less value in the inner superlayers where some events may produce several tracks within a single cell. This effect was mitigated by mapping the cells into the hardware in such a way that cells in the inner superlayers shared FADC cards with cells from the outer regions of the chamber. It should be noted that because the primary buffer has room for four events, it is possible for the DSP processing time on a particular event to exceed the mean trigger interval without the system locking up. The status of the primary buffers is signalled on the backplane by the cards, using a "wired-OR" arrangement. When all the cards in a crate have finished processing a particular region of the primary buffer the LTC is informed and communicates the information to the MTC which only re-uses it when it is available throughout the entire system. The output from

the DSP's computation in normal data-taking takes the form of pulse arrival times and sizes for any channels containing signals. These are written to a predetermined area of a dual-port memory in each module which is accessible also to the readout controller.

§*z-by-timing*

The z-by-timing system [5] measures the z-coordinate of a track by means of the small difference in arrival times of the signal at the two ends of the sense-wire. Signals from both ends of the sense wire, having been received by the post-amplifiers mentioned above, are then delivered to constant-fraction discriminators. The discriminated signals from the two ends of the wire switch on and off the charging of a capacitor by a constant-current source. Charge accumulates during a period determined by the difference in arrival time of the signal at the two ends of the wire, with an appropriate cable delay in the stop signal to ensure this time is always positive. The stored voltage is read by an eight-bit 100 MHz FADC with the most significant bit unused (a design similar to that used in the OPAL experiment [4]). Hence one count corresponds to a spatial distance of about 2m/128 i.e. 1.6 cms. The capacitor is then discharged so that it is ready to receive the next event about 50 ns after the first. The data from the FADC is stored in a pipeline store in a similar manner to the data in the main r-phi system, and also sent to the trigger processors.

To avoid the need for individual setting-up and repeated adjustment to allow for thermal effects, the z-by-timing system is calibrated by injecting a pulse at the preamplifiers. The behaviour of the z-by-timing system is sufficiently stable that recalibration is only necessary infrequently, and this is usually done during a beam-off period, however provision has been made to calibrate the system during a group of empty RF buckets which the HERA machine generates every 20 microseconds. The calibration pulse is sent alternately to the forward and rear ends of the chamber so the ADC output should be zero and full-scale in the two cases. This is achieved by adjusting a) one of the two reference voltages controlling the ADC reference chain and b) a programmable variable delay (DDC) in the stop line. If the ADC fails to record the expected value the contents of the DAC and DDC registers are incremented or decremented as appropriate to correct the discrepancy.

The main problem encountered with this system arose from reflections at the ends of the wire. The characteristic impedance of the sense- wires were about 360 ohms, and this was matched to the input impedance of the preamplifiers, so no reflection would have occurred if the sense wire could have been directly connected. Unfortunately it had to pass through a plastic insulator which held the wire under tension and introduced a capacitance of about 7 pF. A further 4 pF came from the preamplifier input connector. The resulting mis-match caused a reflection altering the shape of the leading edge of the signals and causing a systematic error in the reconstructed z-value seen by the system. This effect as largely eliminated by introducing a surface mount inductor of 1μH between the two stray capacitances, thereby forming a "π" section filter. This had the unwanted effect of slightly reducing the bandwidth.

Since more precise z information could be obtained by using the information from the stereo layers (superlayers 2,4,6,8) in the full analysis of each event, the main purpose

of the z-by-timing system was to provide z data for the trigger system. The inner parts of the chamber were more useful for this purpose because they provided a larger angular coverage. Consequently the z-by-timing system was limited to the first three axial superlayers (1,3,5) and in fact was only used on alternate wires in superlayers 3 and 5. In total, 704 sense wires were provided with z-by-timing electronics.

§*First Level Trigger*

The CTD first-level trigger [6] is designed to identify tracks pointing back to the intersection region in the r-z plane. It takes as its input the seven-bit FADC outputs from the z-by-timing system and is based on the technique of histogramming the values of z/r (in effect the polar angle) for wires seeing a signal. Tracks originating from the interaction region will give similar values for this ratio in each layer they pass through, while those from elsewhere, for instance beam-gas interactions, will give inconsistent values (see figure 2). A ROM lookup table in the z-by-timing cards converts the z coordinate as measured by the seven-bit FADC output into a five-bit address selecting one of 32 bins in z/r shown as the horizontal coordinate in figure 2. As mentioned earlier, high momentum tracks cross two cells in each superlayer, so it is necessary to combine the results from these in forming the trigger. In the case of superlayer 1, each segment-finding processor takes inputs from the sixteen wires of two adjacent cells. A further complication is introduced by the large range of drift times produced by a track as it traverses the CTD. Since the z-by-timing system is clocked every 48 ns. the data for a particular track may arrive over about ten successive clock cycles. Data is accumulated in the processor until a "gap" of several clock cycles indicates that all data has been collected from the previous event.

The identification of track-segments is complicated by the lack of precision in the z measurement, which is only accurate to about 3 cms. Because of this, the patterns in the table illustrated in figure 2 for high momentum tracks from the intersection are spread over several adjacent columns. Segments are identified by a lookup table which accepts a spread of four columns around the central value. Segment finding is also performed for superlayers 3 and 5, where only alternate wires have z-by-timing readout, and a restricted range of polar angles is available. Segment-finding is performed in a generally similar way to superlayer 1.

§*Second Level Trigger*

The entire Zeus experiment was designed around a mean GFLT rate of 1 KHz. As described above, the DSPs on each module write a much compressed form of the data into the dual-port memories during the period between GFLTs. When all the modules in a crate have finished this operation for a particular event, the transputer system in the ROC is able to access this data and perform track-finding operations. Since this is primarily a software matter we give no further details here. For a proportion of events this process leads to the generation of a second level trigger ("SLT") which is combined with similar signals from the other ZEUS components in a way analogous to the first level trigger. If this results in a global second level trigger the full readout of the detector

into an event builder is initiated.

§*Layout of Electronics*

The system described above is housed in sixteen racks, each of which handles signals from one sector of the chamber. Each rack contains a crate of post-amplifiers, an FADC crate and a z/trigger crate (see figure 3). The eighteen FADC modules in each FADC crate receive data from 36 cells, of which two are in superlayer 1, two or three in SL2, three in SL3.. and so on up to six in SL 9. An additional rack contains three crates, one for the master control of the 16 crate FADC system, and two for master control of the z/trigger crates and further processing of the FLT data for transfer to the GFLT.

The racks were a custom design with water cooling capable of handling 6 kW per rack. High-current power supplies were mounted on the back of each crate, within the cooled enclosure, and the entire system was continuously monitored for temperature and supply voltages by a slow control system separate from the data readout chain.

§*Present Status*

The system has been running for over a year, except for the first-level trigger processors in SL3 and 5 which have yet to be installed. The system has been operating very reliably. The ZEUS forward tracking detector has adopted many aspects of the system described above. In particular the TRD readout uses the identical FADC system.

References

[1] F. Bedeschi *et al.*, NIM **A268** (1988) 50

[2] INMOS Ltd., Transputer Reference Manual, (Prentice Hall 1988)

[3] B. Foster et al, NIM **A338** (1994) 254

[4] S. Jaroslavski et al. Proc. IEEE 1989 Nuclear Science Symposium, San Francisco, IEEE Trans. Nucl. Sci. NS-37 (1990) 2161

[5] N. Harnew *et al.*, NIM **A283** (1989) 781

[6] G.P. Heath *et al.*, NIM **A315** (1992) 781

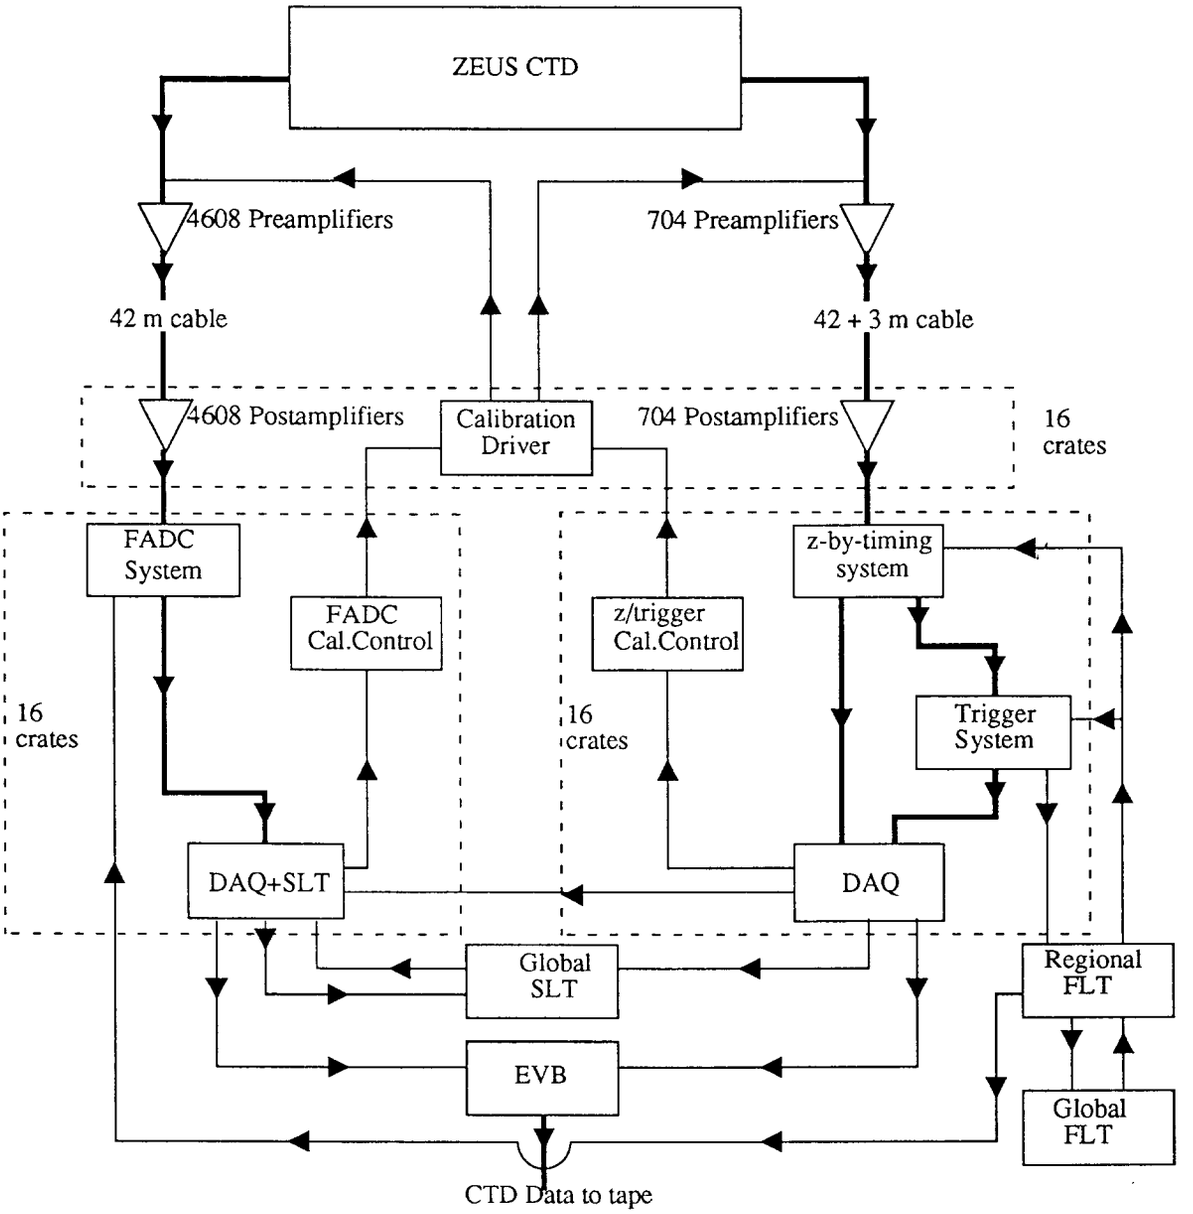

Figure 1: Overview of system

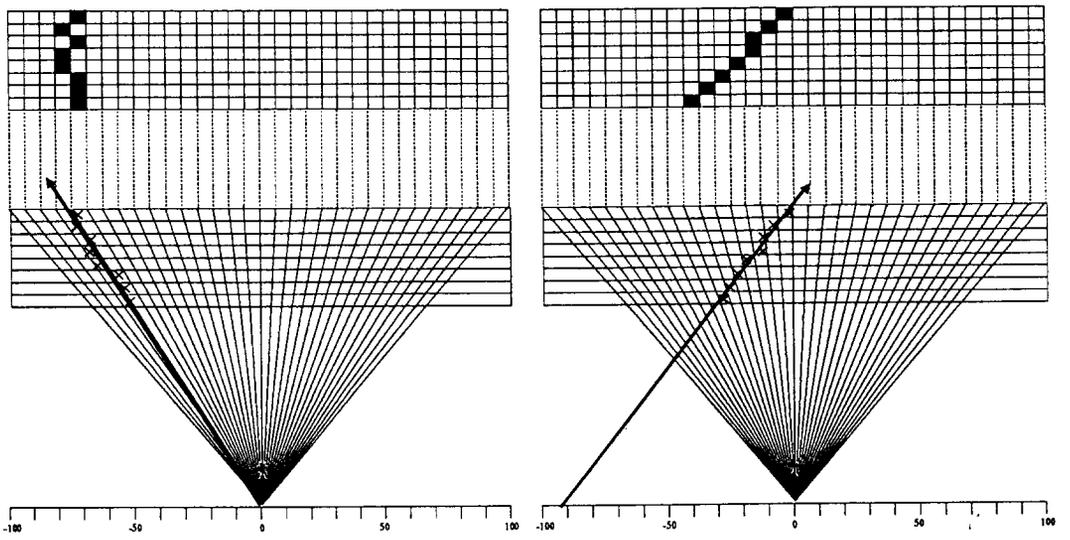

Figure 2: Principle of trigger operation

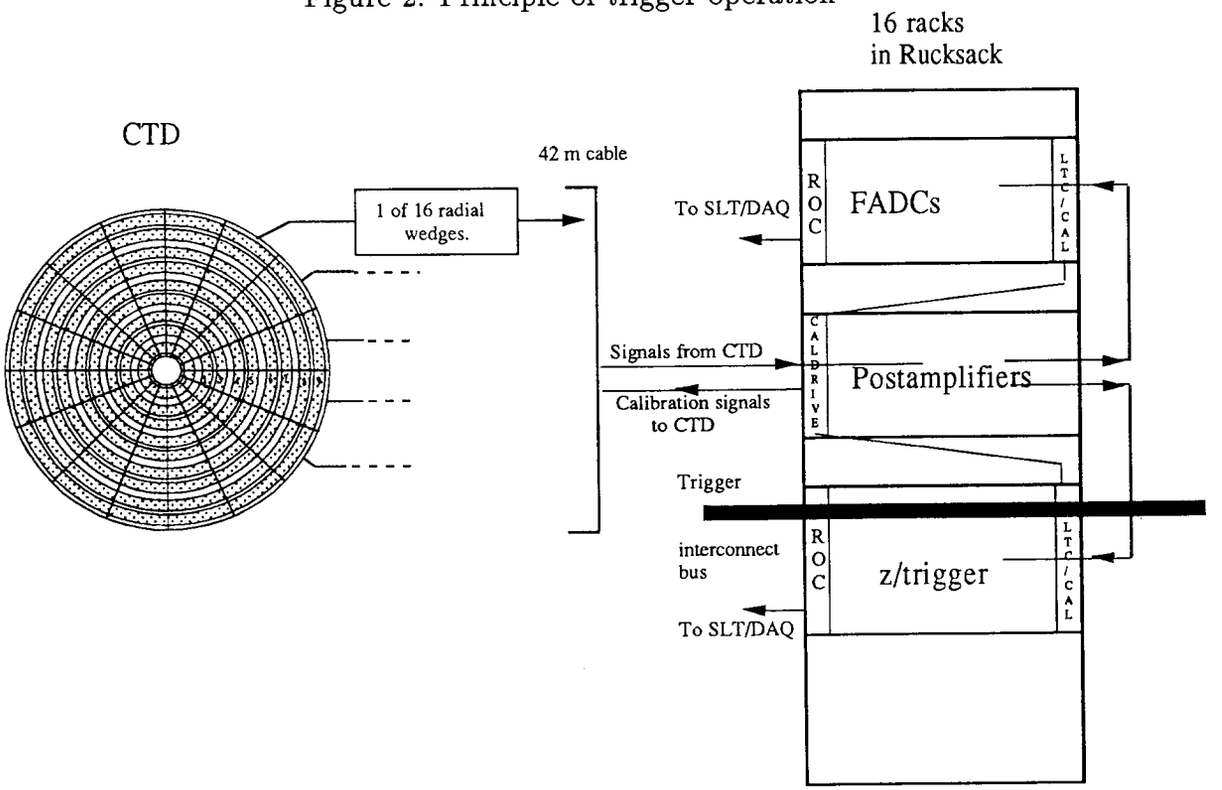

Figure 3: Layout of the Electronics