



ATLAS Internal Note
MUON-NO-179
14 May 1997

Requirements and Specifications of the TDC for the ATLAS Precision Muon Tracker

Yasuo Arai

KEK, National High Energy Accelerator Research Organization

Institute of Particle and Nuclear Studies

yasuo.arai@kek.jp, <http://www-atlas.kek.jp/~arai/>

and

Jorgen Christiansen

CERN, ECP/MIC

christia@sunvlsi.cern.ch

Contents

1. INTRODUCTION	3
2. REQUIREMENTS FOR THE ATLAS TDC	4
2.1 Requirements from Detector	4
2.2 Requirements from Beam	7
2.3 Requirements from first level trigger	8
2.4 Requirements from readout	9
2.5 Requirements from in system test and loading of setup parameters	12
2.6 Radiation hardness requirements	12
3. BASELINE ARCHITECTURE	13
3.1 Time measurement	15
3.2 Channel buffer	19
3.3 First level buffer	20
3.4 Trigger interface and trigger FIFO	22
3.5 Trigger Matching	23
3.6 Readout	26
3.7 Bunch count reset	28
3.8 Event count reset.	29
3.9 Error monitoring	29
3.10 Technology and design methodology	29
3.11 Summary of Specifications	30
4. SIMULATION STUDY	30
4.1 Simulation condition	31
4.2 Hit and trigger generation	31
4.3 Channel buffer occupancy	32
4.4 Double track separation.	34
4.5 First level buffer occupancy	35
4.6 Trigger Search time	36
4.7 Trigger FIFO occupancy	36
4.8 Readout system	37
5. SCHEDULE AND PLAN	42
6. SUMMARY	43
Appendix A : Coding of TDC control signals	45
Appendix B : Definitions of read-out types and sub types	45

1. Introduction

One of the guiding principles in the design of the ATLAS detector is its high-quality measuring capability of muons with precision tracking detectors and super-conducting air-core troids. The barrel part of the muon tracking detectors use high-pressure drift tubes attaining 60 μm intrinsic resolution. More than 300k tubes are needed to select muon tracks effectively under the large background rates. To get the required resolution and handle the large number of channels, front-end electronics must be built with integrated circuits. The small analog signal from the drift tubes must be amplified/ shaped / discriminated by an analog IC [1] followed by a Time to Digital Converter (TDC) performing the required drift time measurement.

A TDC with sub-nano second resolution is required to obtain the necessary spatial resolution of the drift tubes. It must be capable of measuring closely spaced pulses without introducing any additional dead time to the Monitored Drift Tubes (MDT). The high rates expected in the ATLAS detector requires that the TDC is capable of continuously accepting new hits while selectively extracting hits related to a trigger after the first level trigger latency. To do this several levels of data buffering is required to keep up with the trigger rates and hit rates expected in the experiment.

TDC's with sub-nano second resolution have been required in many high-rate experiments using drift chambers/tubes, and several developments have been done of which many have been used in experiments. One of pioneering work was started by one of the authors (Y.A.) in 1986 [2] for the SSC experiments followed by extended versions [3, 4] which have been used in several experiments [5]. These TDC's have mainly been aimed at high rate applications and a simple pipelined buffer have been used to store the measurements during the first level trigger latency. The memory of the first level trigger buffer have in these designs occupied a major part of the chip area and an integration level of four TDC channels per chip have been obtained.

Independently, one of the authors (J.C.) has developed highly integrated TDC's for low to medium rate experiments. Using a data driven architecture TDC's with 16 channels [6] and lately with 32 channels [7] have been developed in the microelectronics group at CERN. An experimental high resolution TDC using an array of DLL's [8] has also been developed.

Common for these chips is the use of a chain of CMOS gates as a fine time measurement within a clock cycle and a clock synchronous ~ 40 MHz counter as a coarse time measurement. This scheme prevents the use of very high-speed clocks in the circuit and results in a low power device (~ 10 mW/channel). The gate delay of CMOS devices normally have very large variations as function of process, voltage, and temperature. By using voltage controlled delay elements as a part of a Delay Locked Loop (DLL) or a Phase Locked Loop (PLL) a self-calibrating TDC can be built.

The hit rates and trigger rates in the ATLAS muon detector are such that a data driven architecture similar to the 32 channel TDC [7] is appropriate to obtain a well integrated TDC. The Verilog simulator [9] will be heavily used to optimize the TDC architecture specifically for the ATLAS MDT detector. A strong emphasis will be put on the flexibility and programmability of the chosen architecture such that changes to the detector (drift time, hit

rates, trigger latency, matching window, leading/trailing edge, etc.) can be accommodated without having to completely redesign the TDC.

Although the basic techniques and methods for building such TDC's are well established, further studies are still needed to surmount the hard conditions of the LHC experiments, minimizing cost, and match it to the ATLAS Trigger and Data Acquisition (DAQ) system. In addition, radiation hardness of the chip must be confirmed. CERN and KEK have started a collaboration pursuing a common TDC architecture for the ATLAS MDT muon detector to make full use of common experience and resources.

This document is a first draft for the path to a final TDC chip for the MDT. Several options are still not well defined and rapid technology changes seen in the VLSI field may have a significant impact on the architecture and implementation of the final TDC. We believe this document will give us a firm starting point and also be useful for the ATLAS muon community. At first, we summarize the requirements on the TDC from the ATLAS muon spectrometer. Then we discuss about several different approaches and technologies to fulfill the requirements. Several simulations of the chosen data driven architecture are also shown.

2. Requirements for the ATLAS TDC

2.1 Requirements from Detector

The MDT front-end electronics consists of 8 channel Amplifier/Shaper/Discriminator (ASD) chips [1] and a 24 channel TDC chip as shown in Fig. 2-1. The primary function of the TDC is a leading edge time measurement. The leading edge timing must be measured with a sufficient accuracy to get a good track resolution. In addition to that, several options have been proposed to get more information from the detector signal. These options are still under study and a short description of each option and the effects on the TDC architecture are given below.

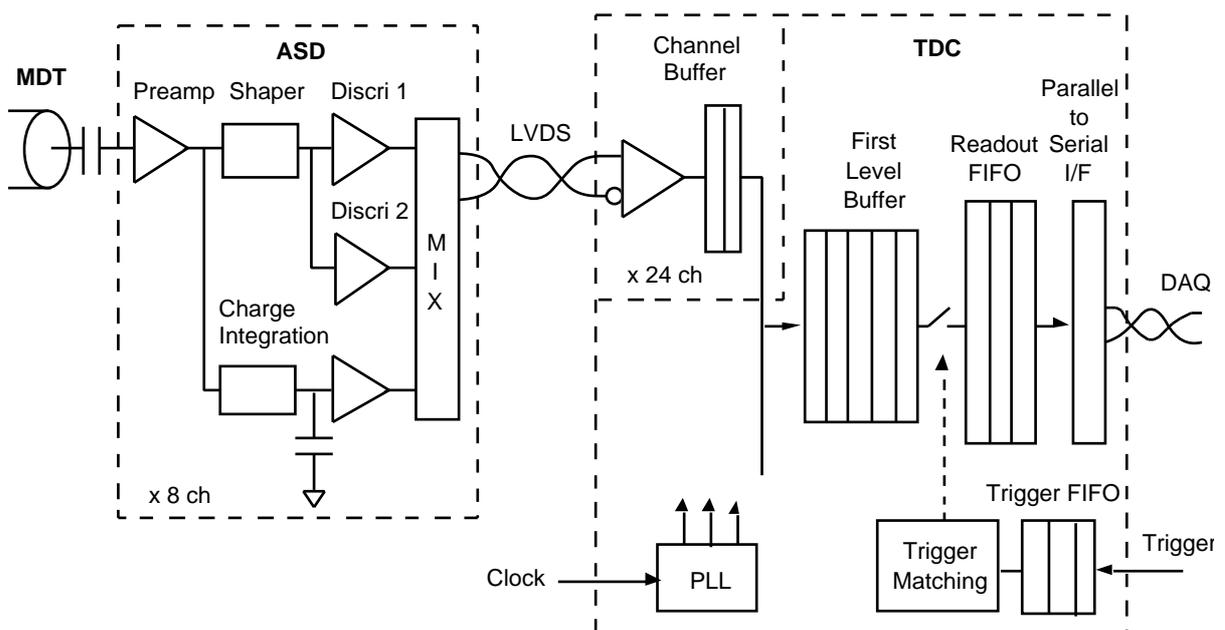


Fig. 2-1. MDT front-end electronics with charge measurement option.

(i) Charge measurement :

The MDT signal is latched at the beginning of the signal and a charge or voltage measure is converted into a pulse width using a Wilkinson type voltage-to-time converter within the ASD. This information can be used to improve resolution by compensating for the signal slew rate, discriminate neutron and γ -ray background, and/or monitor the tube operation. A resolution of 7 - 8 bits is considered sufficient for this purpose. Assuming a minimum pulse width of 30 ns and a bin width of 0.78 ns, the pulse width will be in the range of 30 - 150 ns thereby not adding any additional dead time to the detector (tube + analog front-end electronics have an estimated minimum dead time of 150 ns).

The pulse width can be mixed with the leading edge signal as shown in Fig. 2-1 and Fig. 2-2, so only one connection is needed to send the required information from the ASD to the TDC. When measuring charge via the trailing edge the analog front-end should only generate one pulse (hit) per track.

The TDC must measure both leading and trailing edge time which increases the information to be stored in the channel buffer and the first level buffer. The increase of chip area is estimated to be less than 30%. The increase of the output data size is limited if the charge information is combined with the leading edge information.

(ii) Trailing edge time measurement :

In a round drift tube, the last electron arrives to the wire with a fixed latency from the track. This information can be used to exclude hits coming from other bunch crossings. A direct data reduction within the TDC is though very risky since the efficiency of the bunch identification is not so high. The trailing edge information is only useful in later stages of the data acquisition system such as in the 2nd/3rd level trigger or off-line analysis.

(iii) Two track separation :

Two-track separation capability will increase the track finding efficiency. It is however well known that two-track separation is inherently poor in a round drift tube caused by the large drift distance in the tail of the pulse. Nevertheless, intensive studies have been done to distinguish a second hit from the first.

A second discriminator on the output of the shaper using a higher threshold have been found to obtain a double track separation efficiency of the order of 75 %. The resolution of a second track is though reduced by a factor 2 - 4 with a fixed dead time of 150 ns. The second discriminator will not only trigger on a second track but also on primary ionization clusters from the first track. It has been found to generate in average 1.6 hits per track. The output from the second discriminator can be mixed (OR'ed) with the signal from the first discriminator + charge measurement in such a way that the original signal with high resolution is maintained. The second discriminator in this way only generates additional output pulses on a second track plus the fake hits from the clusters.

In this mode of operation the TDC will have to handle an increased number of hits per channel and the data acquisition system will need to process an equivalent increase in data volume. The use of the second discriminator will in principle not require any changes to the basic architecture of the TDC. The buffering requirements in the TDC will though increase by 50 - 100 % which must be taken into account in sizing the buffers in the TDC. The rejection of secondary hits by the TDC is not of prime importance as the efficiency of

these are only 75%. A rejection of secondary hits of the order of 10% is considered acceptable. An effective support of this mode will only be implemented if the additional buffering required in the TDC does not significantly increase the required silicon area (cost).

The measurement of pulse width is included in the current baseline specification. The pulse width can be used to carry the charge information or the trailing edge of the detector signal depending on the programmed settings of the ASD front-end chip. In case of a charge measurement a maximum pulse width of 150 ns will be measured with a time bin size of 0.78ns. In case of a trailing edge measurement a maximum pulse width of 1500 ns will be measured with a reduced bin size of 6.25 ns in order to limit the data size.

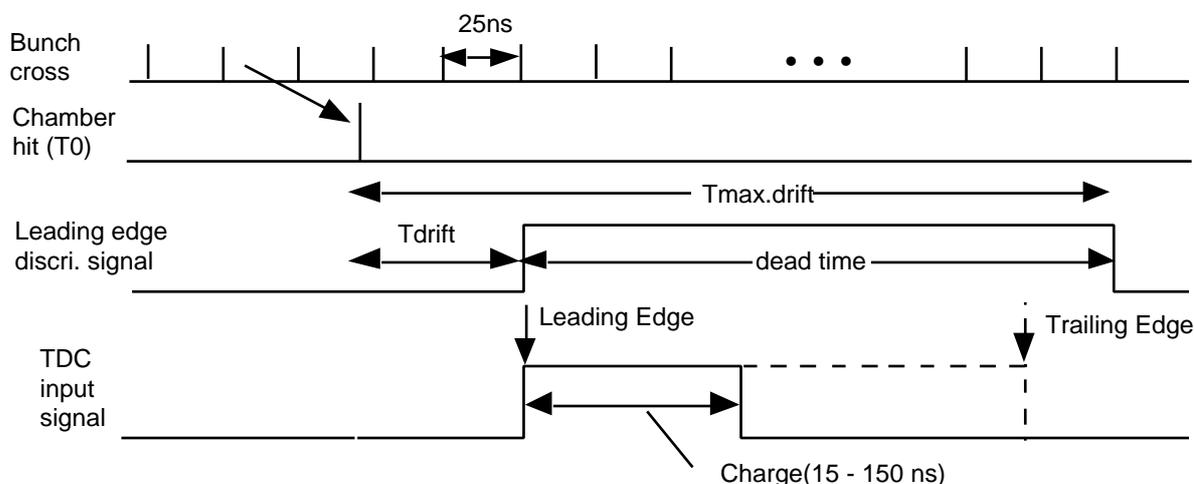


Fig. 2-2. MDT and TDC signals with a charge or a trailing edge measurement option.

Time Resolution

The final gas mixture of the MDT is not yet fully specified but a drift time of approximately $30 \mu\text{m}/\text{ns}$ is expected. To achieve a $60 \mu\text{m}$ single-tube resolution, a timing resolution of better than 1 ns is required in the TDC.

A time bin size of 0.78 ns can be obtained by sub-dividing the 25 ns clock period (40 MHz) into 32 intervals. An RMS time resolution of 250 ps has already been obtained using this method in existing chips [4,7].

Masked Hits

As mentioned previously, the signal from a round drift tube has a long tail. A second track might therefore be masked by the signal tail of a preceding track. In the TDC all leading edge information is stored in the first level buffer. The existence of a preceding hit can be detected by checking the data before the trigger matching region. This 'masked hit' information can be sent out with one word per event using a flag bit for each channel. The masked hit detection increases the trigger matching time and volume of data kept in the first level buffer. However, the trigger matching still has enough time at the present trigger rate (100 kHz) and the increase

in number of gates is insignificant. This function is therefore included in the baseline design.

Hit Rate

The maximum hit rate in a wire is assumed to be 100 kHz. This number is 5 times higher than expected from detector simulations in the barrel region at the $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity. We assume two thirds of the hits are coming from neutron and γ -ray background and distributed randomly in all channels. The other one third of hits originates from charged tracks which generates correlated hits in several channels (3 or 4). The correlated hits have important effects on the buffering scheme used in the TDC since these hits arrive at the TDC at almost the same time. The behavior of the TDC at increased hit rates must also be verified such that unexpected high hit rates does not have serious effects on the system.

Drift Time

The drift time of the ionized track in the drift tube is considered to be a flat distribution from 0 ns to the maximum drift time of 600 ns without any magnetic field present. In the magnetic field of the air core troid the maximum drift time increases to ~ 800 ns in certain regions of the detector. The dead time of the tube without magnetic field is assumed to be a flat distribution from 150 ns to 650 ns. In the magnetic field the dead time period is expected to increase to 150 ns - 850 ns.

Number of channels

The number of channels implemented in a chip depends on the required silicon area per channel and the technology used. This must be optimized before going into the final design. This is also related to packaging, PC board area, and power consumption.

Since the number of tube layers is expected to be 3 and the ASD chip will have 8 channels per chip, a 24 channel TDC matches well this configuration. However we still keep the possibility open to make a 16 or 32 channel TDC which is a practical number to implement in a digital system. A 32 channel TDC may be more practical in selected muon chambers where 4 layers could be required to obtain a sufficient double track resolution.

2.2 Requirements from Beam

The LHC beam has a complex timing structure as shown in Fig. 2-3 [10]. There are 3564 clock periods and 2835 real bunch crossings per beam revolution. Bunch crossings are identified with a 12 bit Bunch Crossing Identifier (BCID) [11]. To correctly identify hits within the bunch structure a 12 bit coarse time counter is required in the TDC.

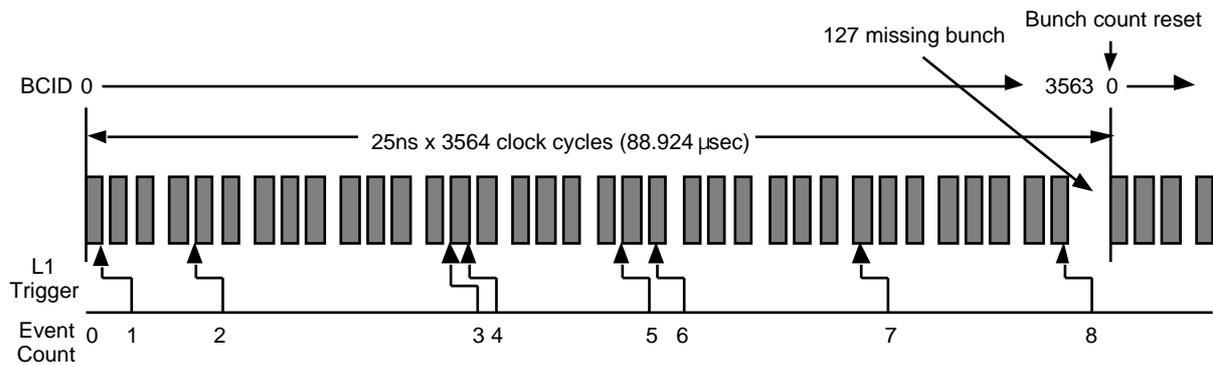


Fig. 2-3. LHC bunch structure and the bunch count reset.

TDC Clock

A TDC clock with the same frequency as the bunch crossing rate must be used. In this way the coarse time count can be correlated directly with the bunch count. In some internal parts of the TDC, such as the PLL or the serial interface, a 2~4 times higher clock frequency may be used to improve performance. A high speed clock will in this case be phase locked to the beam clock.

Bunch Count Reset

There will be a bunch count reset at the beginning of each revolution. Although there are 127 missing bunch periods before the bunch count reset, there may still be data stored in the internal buffers of the TDC from the previous cycle and these can therefore not be cleared at this time. Therefore special care must be taken to synchronize the reset and data.

2.3 Requirements from first level trigger

The first level trigger is a signal generated at the system level signaling that an event of interest have been detected. This signal is generated based on data from several sub detectors in the experiment and it arrives to the first level trigger buffers with a constant latency (delay) from the time when the event actually occurred. When a positive trigger signal is given to the first level buffers, measurements related to that event must be extracted from the buffer and sent to the data acquisition system for further analysis. The trigger is delivered to the TDC as a clock synchronous yes/no trigger signal where a logic high signals the occurrence of a positive trigger.

Trigger conditions in the ATLAS experiment are given in ref. [11]

- maximum level 1 trigger rate is 100 kHz
- level 1 trigger latency is 2.5 μ s
- minimum interval between two triggers is 75 ns
- there are no more than 16 triggers in any given 16 μ s period.

Each event accepted by the first level trigger is identified by a 24 bit event ID at the system level. The event identifier count is reset at the same time as the bunch counter but not for all bunch count resets. At the level of the front-end electronics a reduced event ID of 12 (or 8) bits is considered sufficient. Currently it is not known if the front-end electronics should empty

all their internal buffers when a event count reset is generated (reset of all front-end electronics at regular intervals).

Since the drift time of the MDT tubes is longer than the minimum interval between two triggers, some hits may be shared by several triggers. This is illustrated in Fig. 2-4 where some hits are shared for event N and N+1. Randomly accessed memory is therefore required instead of a FIFO for the first level buffer to handle shared data. The algorithm to find hits belonging to triggers is described in section 3.5.

The time region where relevant data for a trigger may exist is called the ‘Matching Window’. The ‘Mask Window’ is the time region in which hits may mask a hit in the matching window.

The TDC must receive its clock and all real time control signals (trigger, bunch reset, event reset, global reset) on two twisted pairs using LVDS signals. Currently it is envisaged to use one pair for the clock and the other pair for a simple clock synchronous encoding of the control signals. This encoding scheme requires that there is a minimum spacing of 3 clock cycles between any control command.

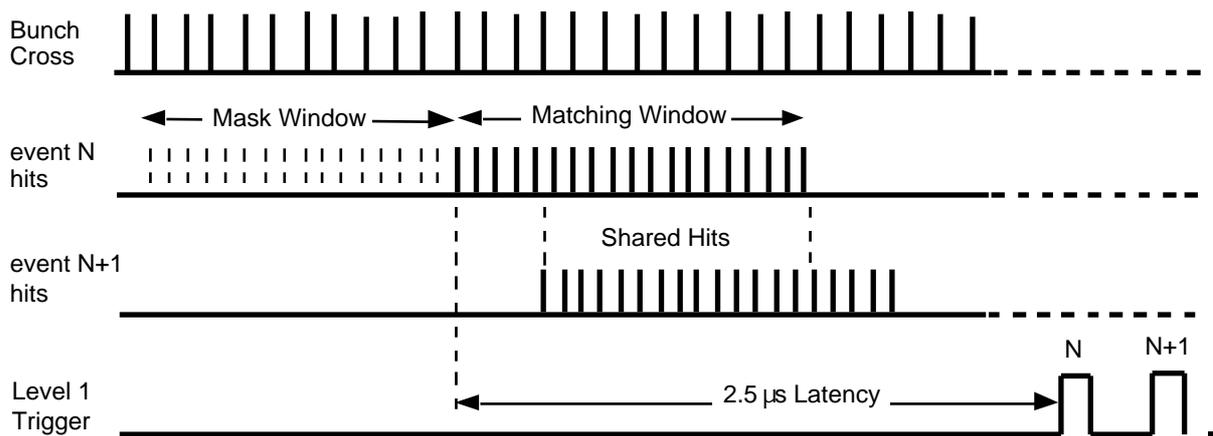


Fig. 2-4. Timing relations between hit and trigger signals.

2.4 Requirements from readout

Configuration and Signal levels

The TDC chip is intended to be mounted on a small board together with the ASD chips to avoid large number of cables between the TDC and the ASD. This front-end board will be located inside the detector directly at the ends of the MDT tubes as shown in Fig. 2-5.

It is of vital importance that noise from digital signals on the front-end board is kept to an absolute minimum so the small analog signal from the tubes are not corrupted. Differential low voltage signals are preferable for all connections to/from the TDC which are actively running while taking data (ASD to TDC connection, Clock, Trigger, Bunch reset, Event reset, Readout, etc.). One of the present candidates is Low Voltage Differential Signaling (LVDS) defined in the IEEE 1596.3 SCI-LVDS standard. LVDS is a high speed standard which can possibly also be used directly for a high speed serial readout link. Signals not actively changing state while taking data can possibly be implemented with standard CMOS levels (3.3 volt).

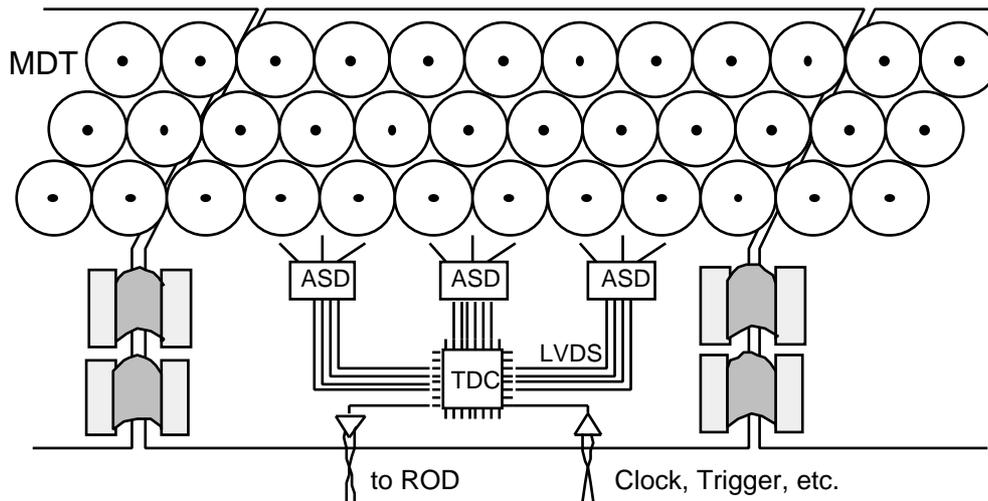


Fig. 2-5. Layout Image of the front-end board and the TDC.

Hit data selected by a trigger must be transferred from the TDC to the data acquisition system located outside the detector. The relative large physical dimensions of a MDT chamber (several meters) makes it difficult to read out the TDC data over a shared parallel bus. The use of a shared media for the readout also increases the risk of losing all the TDC data from a chamber in case of a hardware failure. The distance from the TDC's to the first level of the data acquisition system (ROD: Read Out Driver) is of the order of 10 meters.

The connection from the TDC's to the ROD is assumed to be performed on serial links using thin cables with shielded twisted pairs. This reduces the physical dimensions of the cables required inside the detector to a minimum. 80 Mbits/s (or 160 Mbits/s) links using low voltage differential signaling (LVDS) is the most likely candidate technology. The cable driver may possibly be external to the TDC to be capable of optimizing the driver to the selected type of twisted pair. The serial line may be daisy chained between several TDC's to further reduce the number of cables inside the detector. When daisy chaining several TDC's the possibility of bypassing a failing IC must be considered to minimize the number of lost channels in this case.

Data Packet

Data packets of 32 bit are assumed in the serial data stream. The first four bits of the packet are allocated to a type identifier which specifies the type of data contained in the 32 bit packet. The following four bits must be allocated for a TDC chip identifier (allowing a maximum of 16 TDC's per serial readout link) to implement an efficient readout protocol where only TDC chips having hit data for an event actively transmits on the shared links. Event data from different events must be separated by the optional use of a global event header and a global event trailer. The use of local headers and trailers for each TDC on a daisy chain must also be supported to enable extensive data checking of TDC data during debugging of the muon electronics system.

A header must contain the event ID and bunch ID of the event being readout. A trailer must contain the event ID and a word count (or possibly a CRC check sum). All required information from a hit: type identifier, TDC ID, channel ID, leading edge time plus pulse width must be contained in one 32 bit word. The conditional mask flags for the 24 TDC channels

must also be contained in a single 32 bit word. In case a TDC have detected an internal error condition (e.g. buffer overflow) it must send a special error status package for all events which may have been affected by the error.

Data Rate

At a 100 kHz hit rate per TDC channel and a trigger matching window of 800 ns an average of 1.87 hits per trigger + 0.85 masked hit word = 2.72 words of 32 bits must be transferred from a 24 channel TDC to the ROD per trigger when not using local headers and trailers (see section 4.2). With a trigger rate of 100 kHz this requires an average serial bandwidth of 8.7 Mbits/s per TDC plus 6.4 Mbits/s for the global header/trailer. In case local headers and trailers are used the bandwidth required per TDC increases to 15.1 Mbits/s.

To prevent excessive long event delays it is advisable only to utilize half the maximum bandwidth available. This translates into a maximum of 2 (3) TDC's per serial link when using local headers and trailers and to a maximum of 4 TDC's per link when only global headers/trailers are used. For other parts of the detector where the hit rate is low the maximum number of TDC's per serial link is mainly limited by the transmission of headers and trailers sent for each event. When using local headers/trailers the bandwidth required per TDC is ~6.4 Mbits/s and 6 - 8 TDC's can be daisy chained on one serial link. Using only global headers/trailers per event the number of TDC's are limited to 16 by the 4 bit TDC ID field in the data packets.

Serial Link Protocol

Many serial link protocols have been developed actively in these days. It is desirable if we can use such an industry standard, but usually these protocols are rather complex and we need a simple one. At the current moment in time it is not determined what kind of protocol to use.

A potential coding scheme of the serial data is the two wire signaling scheme used in transputer links. This transputer link protocol is called the DS protocol (IEEE P1355). The coding scheme keeps the required bandwidth to an absolute minimum and do not require any complicated PLL or phase alignment circuitry in the receiver.

To limit the number of cables in the detector it is not considered to have a back propagation signal from the DAQ to the TDC telling it to stop sending data in case buffers in the DAQ system are running full. The buffering capability in the TDC (32 words deep readout FIFO) is limited in comparison to what is going to be implemented in the input buffers of the DAQ system (~ Kbytes). If the DAQ stopped the readout of data from the TDC's they would very fast get overflows in their internal buffers. The DAQ system must handle input buffer overflows locally without stopping the data transmission from the TDC's.

Error Recovery

Since the chosen architecture is data driven, there is always a risk of a buffer overflow when a large fluctuation in the hit or trigger rate occurs. Extensive simulations of the TDC architecture have been performed and the buffers are sized such that this should in principle never occur. In a system of the size and complexity of the ATLAS MDT detector there is always a potential risk of some local hot spots (in space and in time) where an internal buffer

in the TDC may overflow. When this happens the TDC must properly mark events which have lost data. The TDC must under no circumstances lose complete events as this may result in data from different events to be mixed when several TDCs are daisy chained on one serial link. The TDC must be capable of recovering event synchronization locally without any higher level intervention from the DAQ system.

System monitoring and debugging.

The TDC being a data driven architecture means that the occupancy of the internal buffers will have quite some variations over time. To perform efficient system monitoring and debugging finding potential hot spots it must be possible to obtain statistics on the buffer occupancies in the TDC's. To get this information in real time the monitoring information must optionally be included in the readout data stream.

2.5 Requirements from in system test and loading of setup parameters

The fact that the TDC is going to be embedded inside the detector requires special attention on monitoring and in system testing capabilities. The cause and place of hardware failures must be detectable without opening the detector such that the effect of a hardware failure can be minimized and such that a repair can be performed fast and effectively. A relatively slow bi-directional communication path is also necessary to be capable of loading setup and calibration parameters before data taking and to monitor the system during operation.

The standard IEEE 1149.1 JTAG protocol [12] is today an accepted standard for in system chip and module testing, and standard debugging tools are available. The use of full boundary scan enables efficient testing of TDC module failures (shorts and opens between ICs) while located in the system. Testing the functionality of the chips themselves are also supported by JTAG. To insure fast and effective testing of embedded memory and data path structures in the TDC chip special scan path registers must be implemented for this.

The analog front-end chip (ASD) is assumed to have a full implementation of JTAG for in-system testing and loading of setup parameters. The connection from JTAG on the front-end boards to the DAQ/slow control system is still not determined (CAN bus, I²C, IEEE 1149.5 MTM bus, IEEE 1394, or equivalent.).

2.6 Radiation hardness requirements

The radiation levels in the muon detector is orders of magnitude smaller than seen in the central parts of the ATLAS detector. The radiation level in local hot spots in the muon detector may though be so high that normal commercial CMOS processes may be sufficiently degraded to provoke hardware failures. The exact radiation levels in the locations of the TDC chips still needs to be verified. Under all circumstances the functionality of the TDC as function of radiation levels must be seriously investigated before accepting a final design.

Radiation may also introduce single event upsets by accidentally changing the state of a memory element in the TDC. The effect of such an event must be minimized and should be detected by the TDC itself possibly using error detecting/correcting codes in all internal memory structures.

3. Baseline Architecture

A block diagram of the proposed baseline architecture is shown in Fig. 3-1. The hit signal coming from the analog front-end chip via differential LVDS is used to store the fine time and coarse time measurement in individual channel buffers. The fine time measurement is obtained as a fine interpolation of the basic clock cycle using a chain of delay elements being a part of a self-calibrating Delay Locked Loop (DLL) or a Phase Locked Loop (PLL). The timing of both leading and trailing edge of the hit signal can be stored as a pair to enable a pulse width measurement to be performed. Optionally the leading edge only or the trailing edge only (or both independently) time measurement can be performed in case the pulse width measurement is not required.

The time measurements from the channel buffers are stored during the first level trigger latency in a common first level buffer. First level triggers converted into trigger time tags and a corresponding event ID are stored temporarily in a trigger FIFO waiting to be matched with the hit measurements from the first level buffer. Hits matching triggers are written into a readout FIFO waiting to be transferred to the DAQ system via a serial link.

The signal path from the trigger and hit inputs to the read-out contains several buffers so multiple events can be processed at the same time (Fig. 3-2). One trigger in the process of being received, One event in the process of being trigger matched and a third event in the process of being read out.

A major choice in the design of the basic architecture was the selection of buffering scheme to store measurements during the first level trigger latency. Basically two different buffering schemes exist (Fig. 3-3).

In the synchronous (or pipelined) architecture [4] the fine time is stored into a dual port memory at each clock cycle irrelevant to the existence of a hit. In this architecture there is almost no hit rate limitation and it is relatively easy to extract triggered events from the memory for a fixed trigger latency. An overlapping trigger can be handled with a pointer and additional memory locations. This architecture requires a relatively large amount of memory per channel.

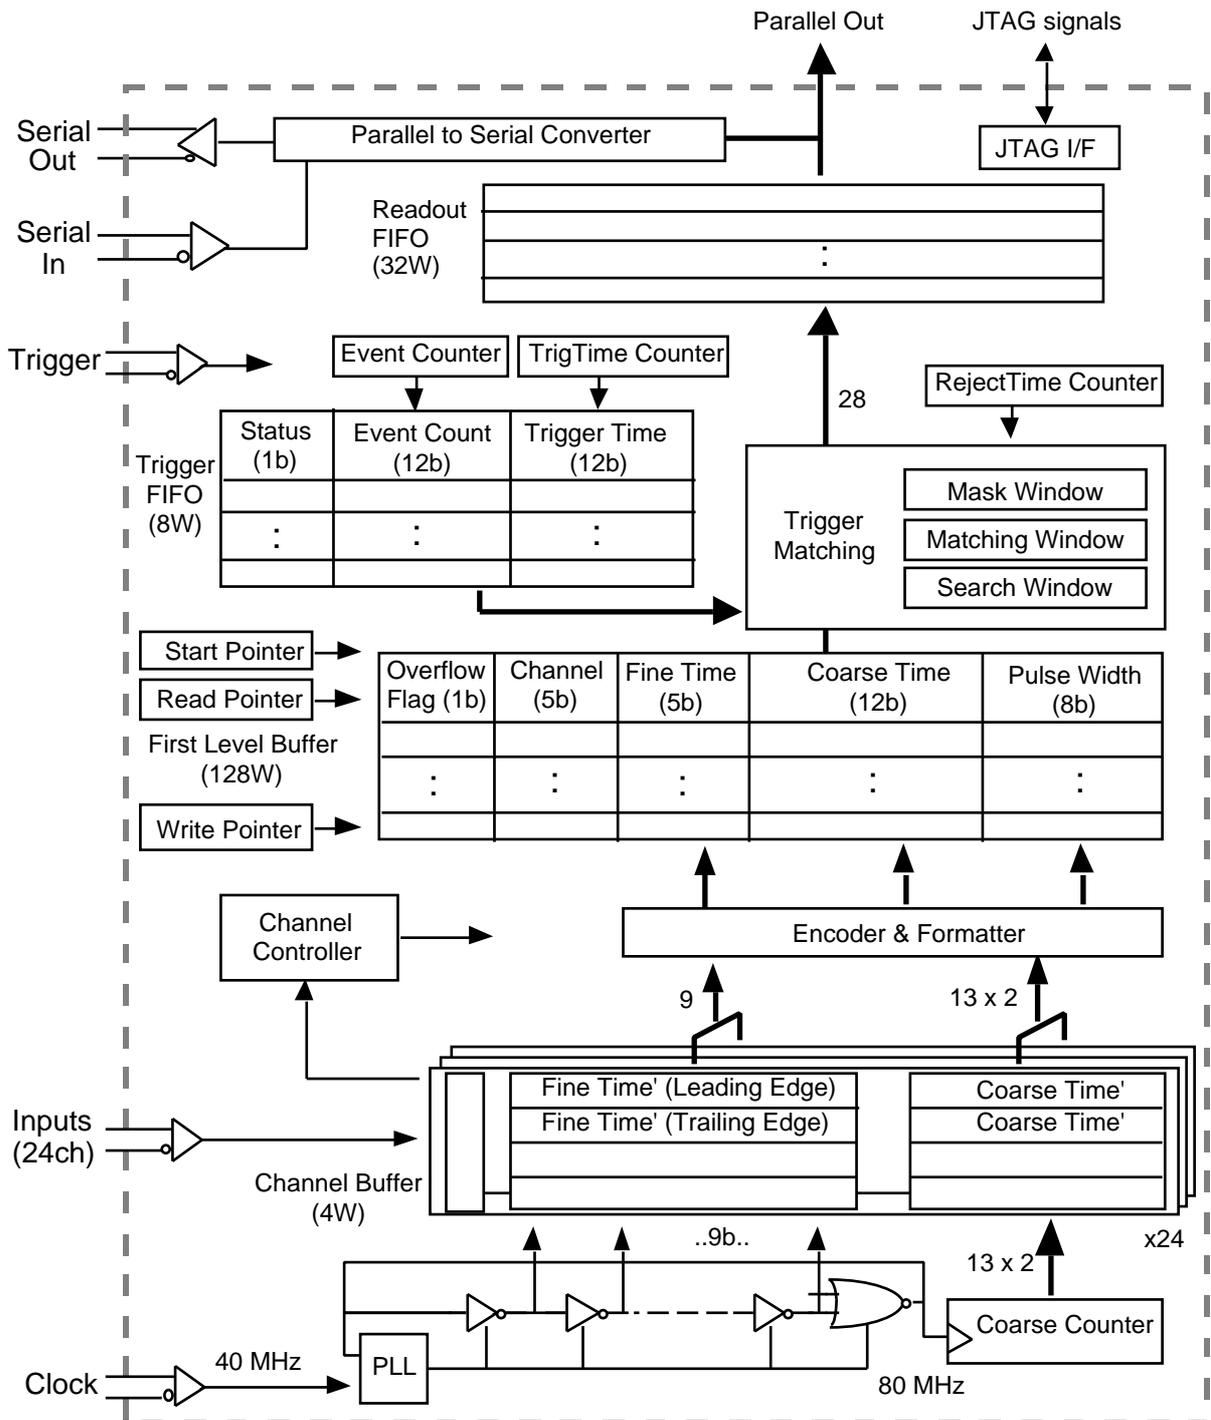


Fig. 3-1. Block diagram of the TDC for the ATLAS MDT detector.

In an asynchronous (or data driven) architecture [6,7] a fine time and a coarse time is only stored when a hit has been detected. The required memory size in this case depends on the hit rate. An additional possibility using this buffering scheme is to have one common buffer shared by many channels giving additional savings in the use of memory. The use of a common resource shared by many channels may though introduce a significant dead time if not carefully designed. The use of a small buffer (2 hits) per channel before the common L1 buffer have for the ATLAS MDT case removed this problem (the MDT detector has itself a dead time of 150 ns - 650 ns). The data driven architecture requires slightly more complicated circuitry to extract hits belonging to a trigger and the trigger latency is not a hardwired property.

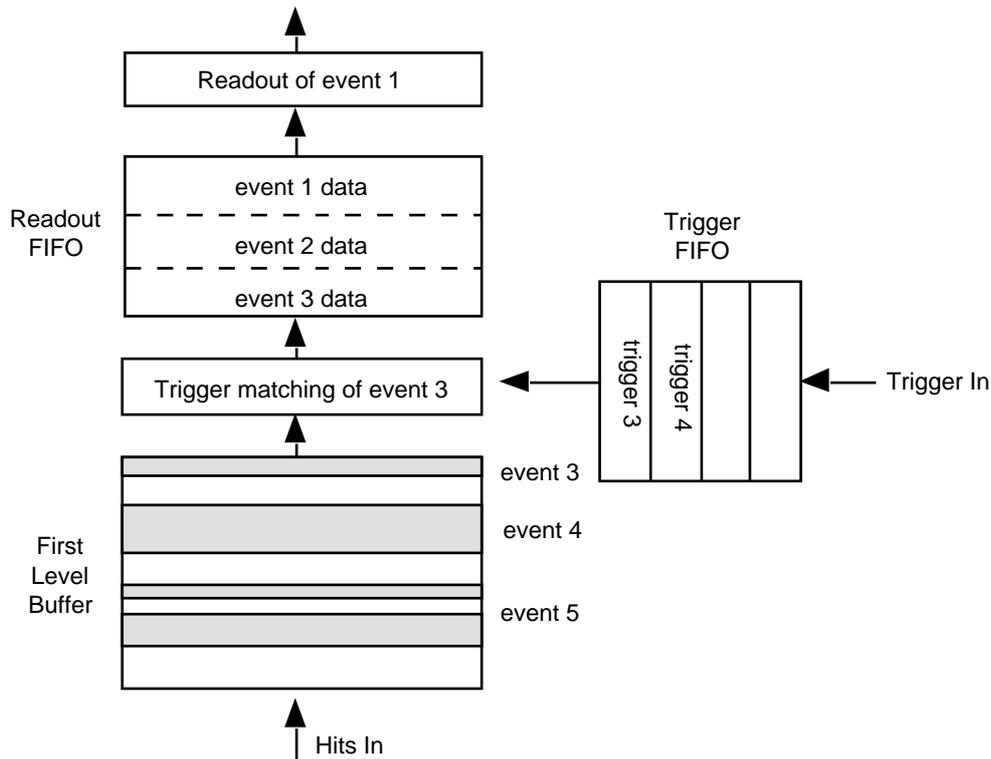


Fig. 3-2. Processing of several events in parallel.

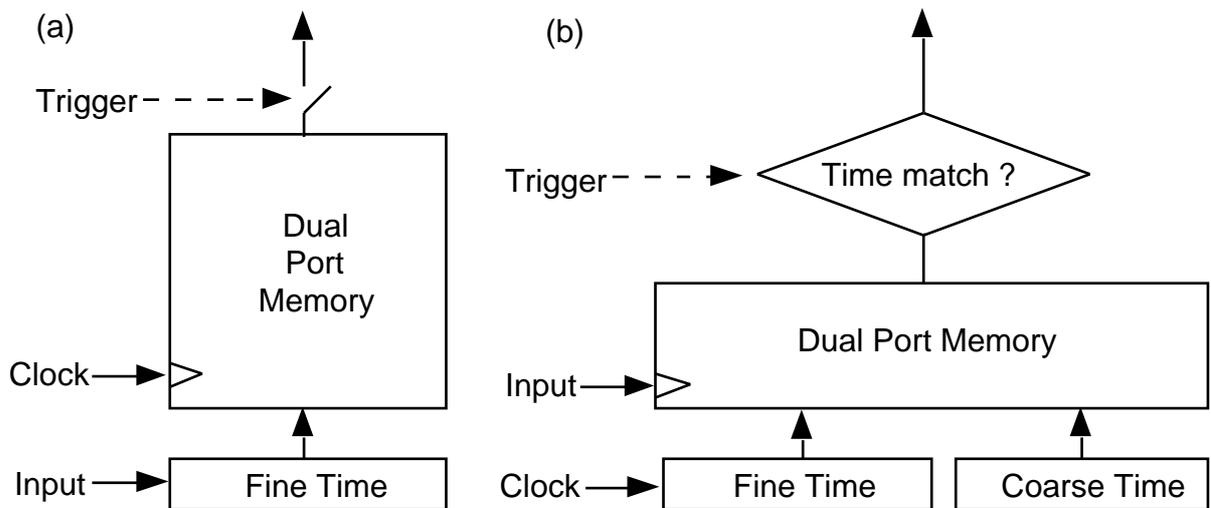


Fig. 3-3. (a) Synchronous and (b) Asynchronous recording architecture.

3.1 Time measurement

A CMOS gate delay is used as the base for the fine time measurement. The delay of a CMOS gate has a large variation as function of process parameters and is also very sensitive to changes in temperature and supply voltage. To get a high resolution time measurement with sufficient stability voltage-controlled delay elements, used as a part of a self calibrating feedback loop, must be used. The feedback loop can be made as a Delay Locked Loop (DLL) or as a Phase Locked Loop (PLL). The two feedback schemes are very similar in construction and the final choice will depend on details of the final IC technology to use. A PLL can potentially reduce jitter in the input clock and can if needed generate on-chip high speed clocks. By dividing the basic 25 ns clock period into 32 intervals a time bin size of 0.78 ns is obtained

(~250 ps RMS resolution).

The dynamic range of the fine time measurement is expanded by storing the state of a clock synchronous counter. The hit signal may though arrive asynchronously to the clocking and the coarse counter may be in the middle of changing its value when the hit arrives. To circumvent this problem two count values, 1/2 a clock cycle out of phase, are stored when the hit arrives. Based on the fine time measurement one of the two count values will be selected, such that a correct coarse time value is always obtained.

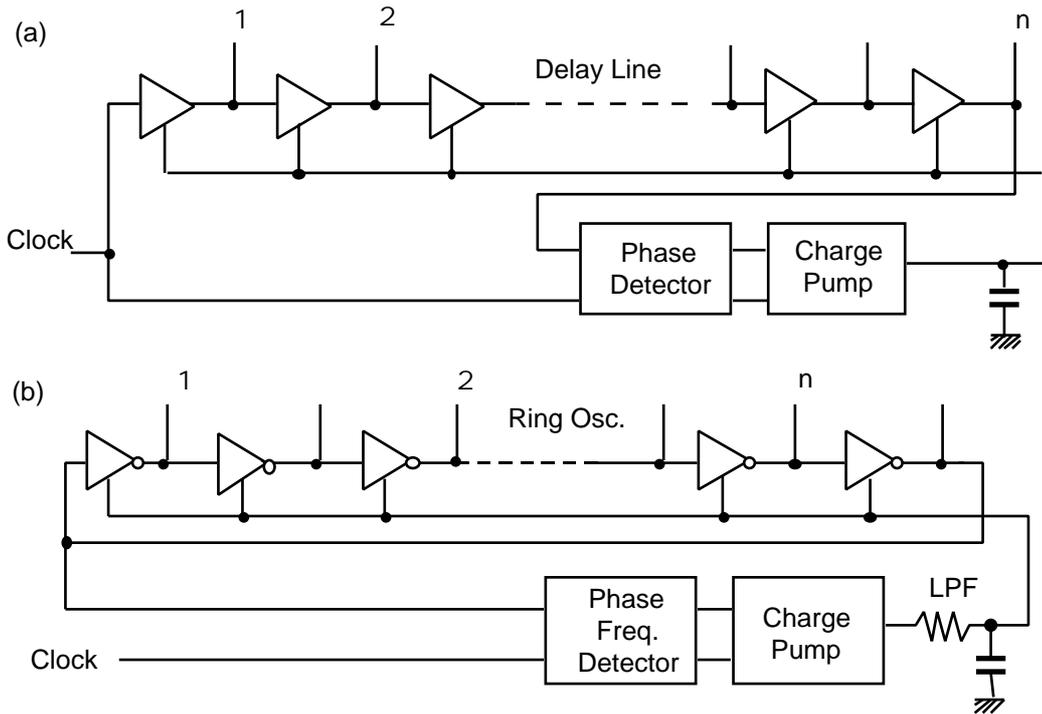


Fig. 3-4. (a) DLL, and (b) PLL circuit.

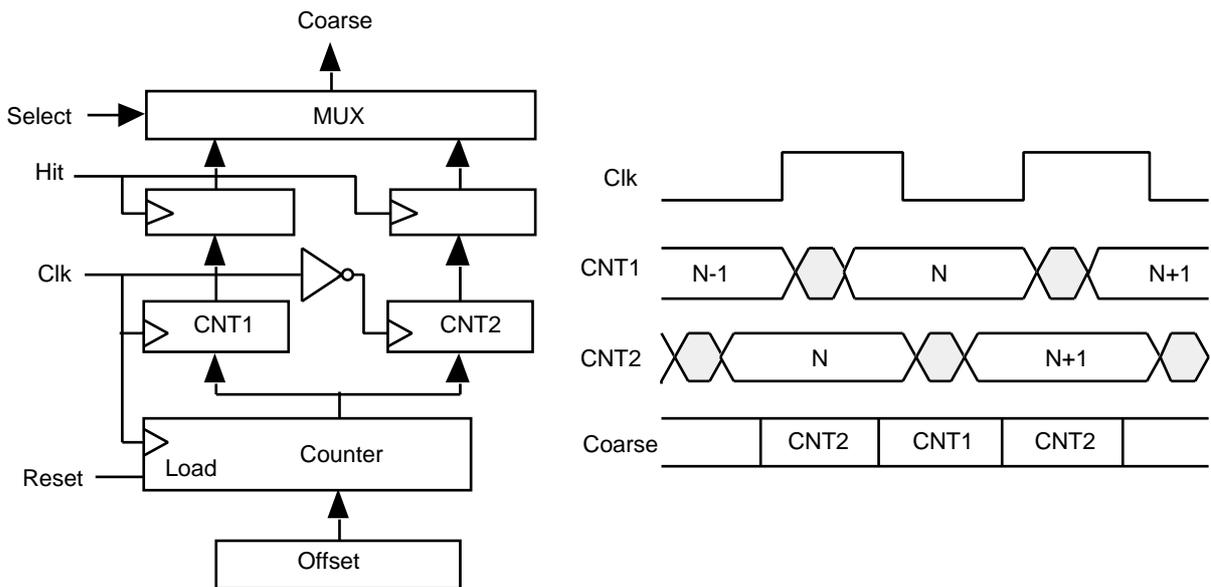


Fig. 3-5. Phase shifted coarse time counters loaded at hit.

At reset the coarse time counter is loaded with a programmable coarse time offset. In this way time differences between TDC chips in the system can be partly compensated for (see also channel adjustment: 3.3).

Using the basic scheme of a fine time made from 32 time intervals and a 12 bit coarse time count, $32 \text{ bits} + 2 \times 12 \text{ bits}$ are required to capture a time measurement. Since each channel buffer contains 4 time measurements (leading and trailing timings for two hits), $(32 + 2 \times 12) \times 4 \times 24\text{ch} = 5376$ bits of memory is required in total. This memory is in a time critical part of the TDC and it will be advantageous to reduce its size especially if using a gate-array technology where memories are not implemented very efficiently.

Several methods to reduce the number of required bits are shown in Fig. 3-6. Since the delay chain has both rising and falling edges within the chain, one can reduce the number of taps to almost half (not exactly half, 17 taps are required from the characteristics of the asymmetric ring oscillator [4] which is assumed to be used) if both edge information's are used. This OPTION A though has a significant risk to degrade the differential linearity since the characteristics of the two edges are very different in CMOS technologies.

In OPTION B, the oscillating frequency of the PLL is increased to 80 MHz and the coarse counter is clocked at 80 MHz. The number of taps in the fine time is reduced to half while the coarse time count is only incremented by one. By combining the OPTION A and B, one could possibly reduce the number of taps to 9, and the memory requirements for the channel buffer now decrease to 3360 bits.

An alternative scheme would be to let a PLL generate directly a ~ 1 GHz on chip clock driving a Gray code counter[13]. In this case the required memory in the channel buffer would be reduced even further to 1632 bits. To implement this scheme requires a modern $\sim 0.3 \mu\text{m}$ CMOS technology which should be commercially available within the time frame of the ATLAS experiment. More detailed studies and careful comparisons are needed for this option.

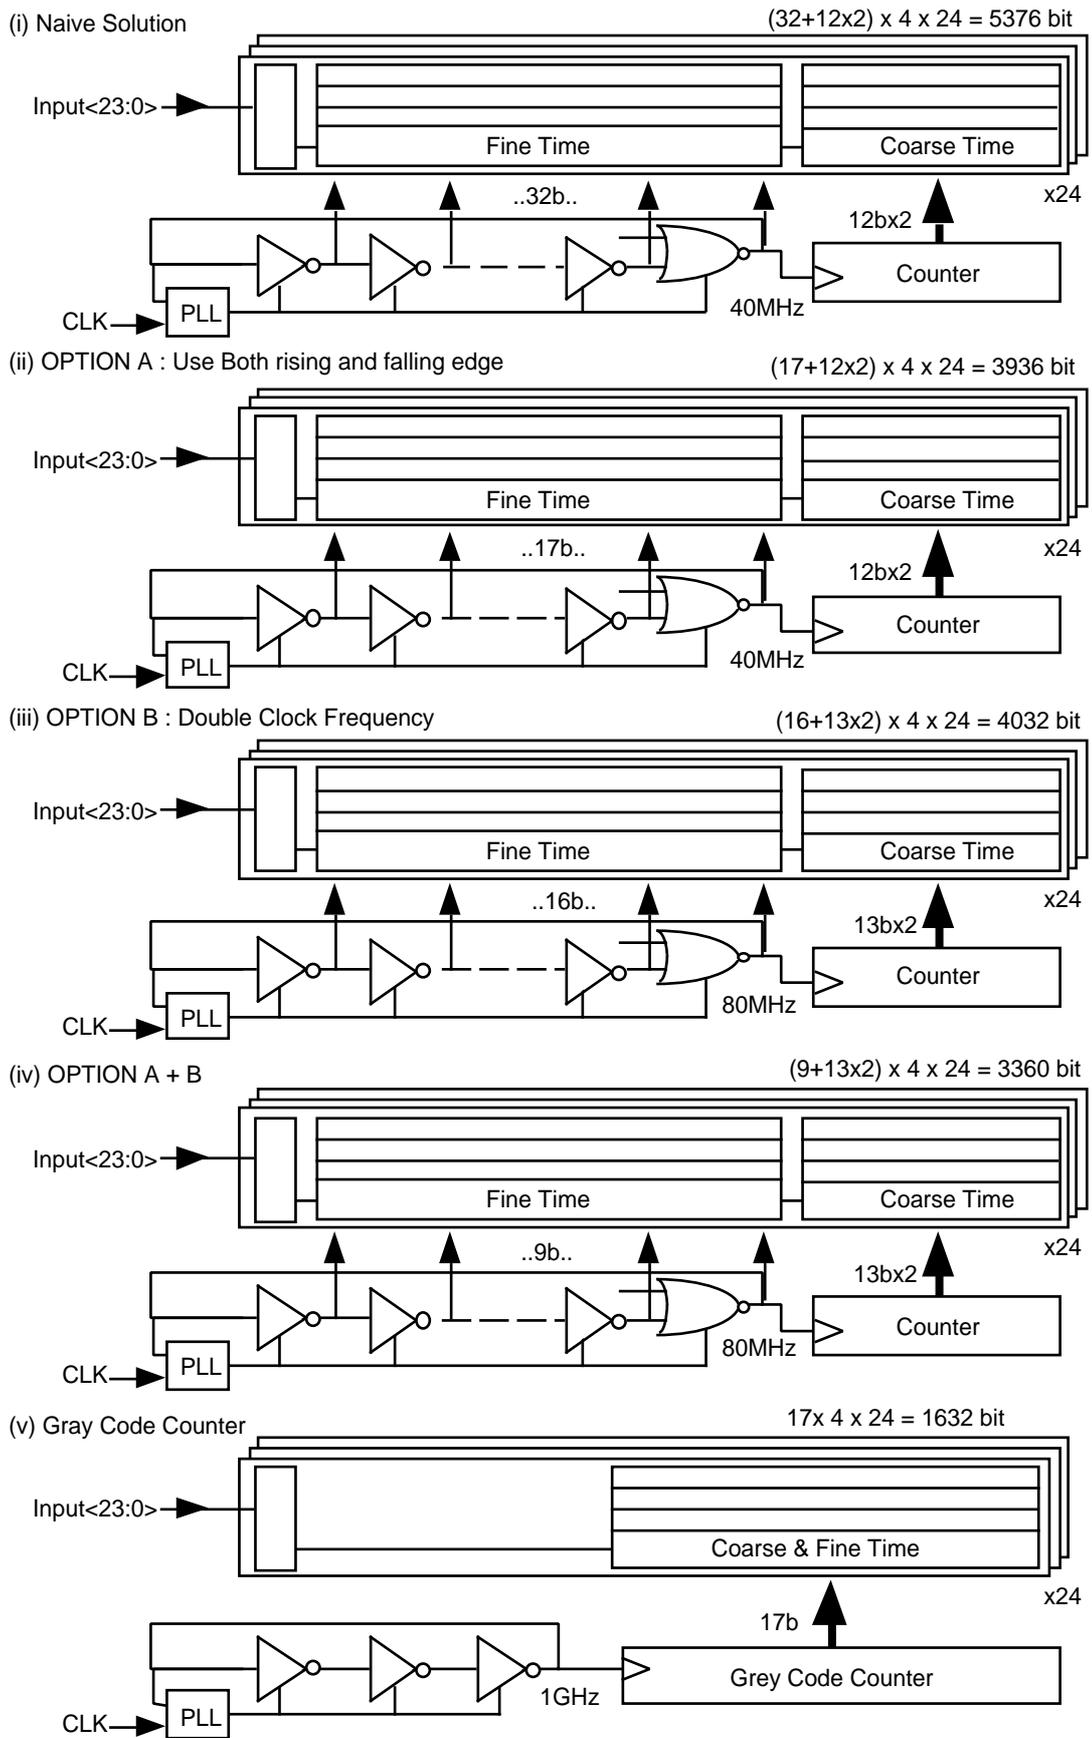


Fig. 3-6. Various configurations for the fine time measurement.

3.2 Channel buffer

Each channel has a small buffer where measurements are stored until they can be written into the common on-chip first level buffer. The channel buffer is implemented as a FIFO and must contain at least two complete time measurements (leading and trailing edge) as show in Fig. 3-7.

The channel buffer must be controlled such that it is capable to measure the minimum pulse width of ~ 10 ns. To measure the pulse width of the hit signal a time measurement pair consisting of a leading and a trailing edge must be assembled. The paring of the two measurements can be performed directly at the input of the channel buffer (Fig. 3-7 (a)). Alternatively the leading and trailing edge measurements can be performed as separate measurements written into a common buffer and then paired at the output of the channel buffer (Fig. 3-7 (b)). Paring of measurements at the output of the channel buffer has the advantage that in case no pulse width measurement have to be performed the channel buffer can store 4 single edge measurements.

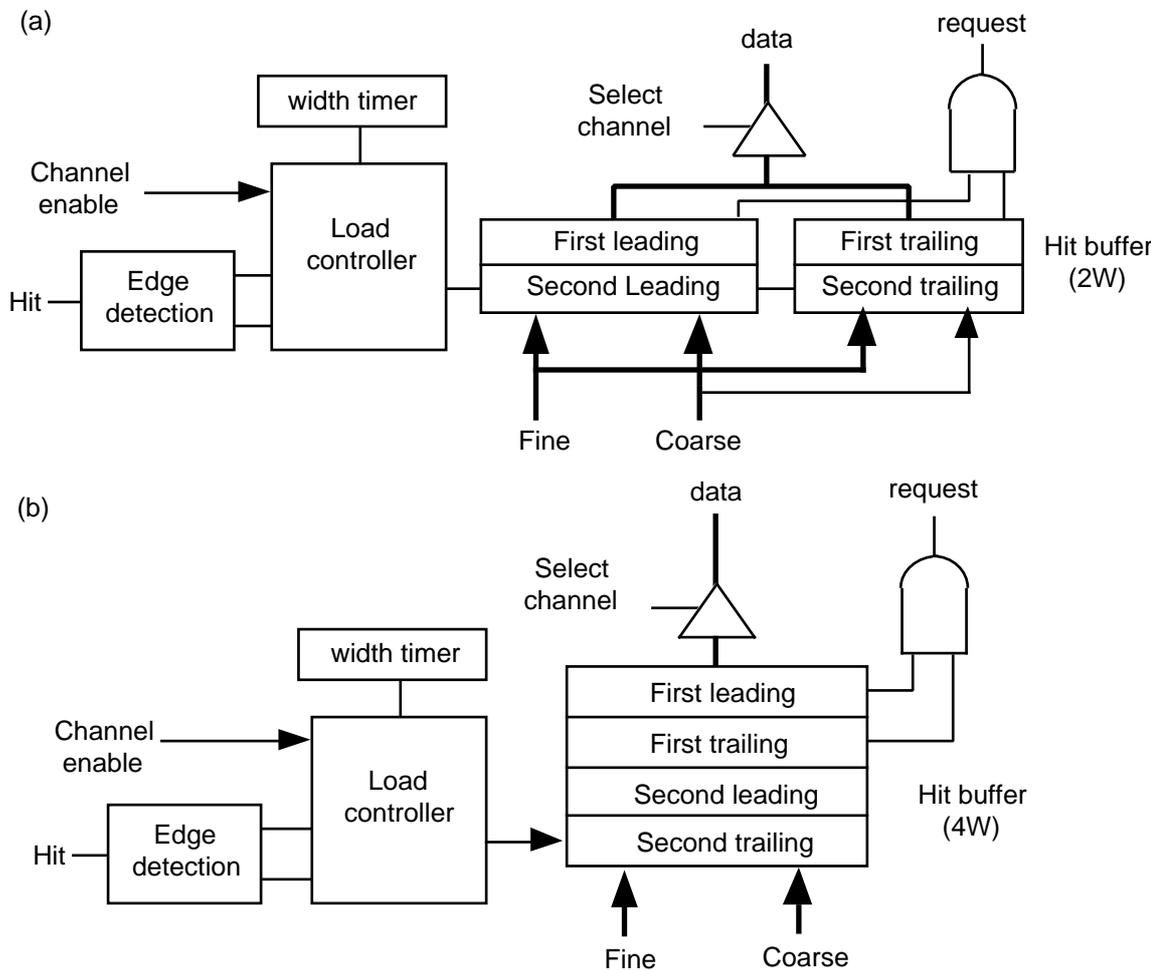


Fig. 3-7. The channel buffer control scheme. (a) separate buffers, and (b) a common buffer for the leading and trailing edge timing.

If a hit occurs when the channel buffer is full, it will be rejected and no time measurement will be performed. The loss of hits due to overflow of the channel buffer is extremely small $< 10^{-7}$ at 100 kHz hit rates partly because of the intrinsic dead time of the MDT tube and analog front-end electronics (see simulations).

When the channel buffer have performed the required measurements a request to be written into the clock synchronous first level buffer is issued. This request signal is synchronized to avoid possible meta-stable states and then processed by the channel arbitration logic.

3.3 First level buffer

The first level buffer is 128 (256) hits deep and is written into like a circular buffer when a hit have been detected on a channel. The size of the buffer must be determined before the chip is designed according to simulation results. Reading from the buffer is random access such that the trigger matching can search for data belonging to triggers.

The layout of the first level buffer is optimized for the ATLAS MDT detector by having a complete leading edge measurement plus a reduced pulse width measurement per word. In case individual leading and/or trailing edge measurements are required one of the pulse width bits will be used to mark data as being a leading edge or a trailing edge measurement and the remaining pulse width bits will be left unused. When performing a charge measurement the pulse width is stored as 8 bits with a resolution of 0.78 ns. In case of a trailing edge measurement the pulse with is stored with a 6.25 ns resolution to obtain the required dynamic range using a 8 bit representation (limited by readout considerations).

Channel Arbitration

When a hit has been detected on a channel the corresponding channel buffer is selected, the fine time measurement is encoded into binary form, the correct coarse time count value is selected and the complete time measurement is written into the first level buffer. With each time measurement a channel number is appended.

When several hits are waiting to be written into the first level buffer an arbitration between pending requests is performed. A registered arbitration scheme illustrated in Fig. 3-8 is used to give service to all channels equally. Each channel have a hardwired priority but new requests are only allowed into the arbitration queue when all pending requests in the queue have been serviced. The request queue is serviced at a rate equal to the clock frequency.

The time measurements are not written into the first level buffer in strict temporal order. A request to be written into the first level buffer is not made before the corresponding pulse width measurement has been finalized and the channel arbitration does not take into account which hit occurred first. This have important effects on the following trigger matching as will be described later.

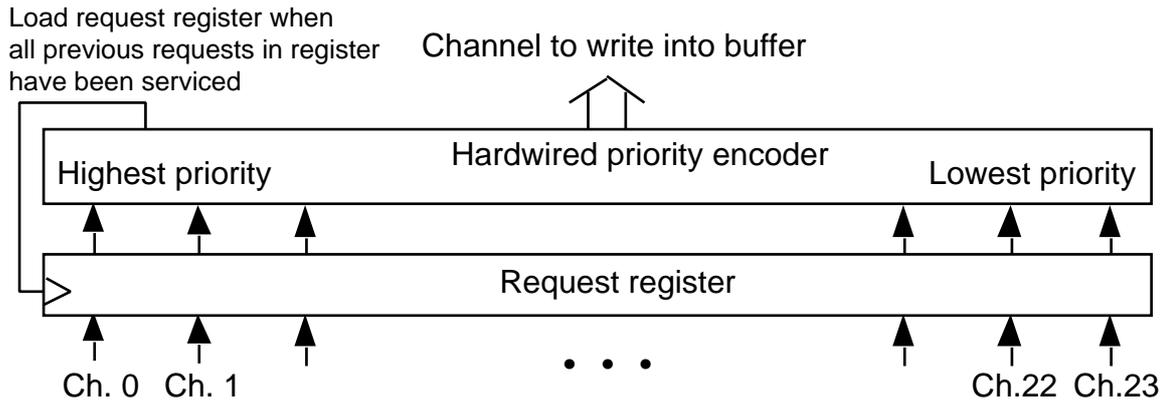


Fig. 3-8. Registered hardwired priority arbitration of channels.

Channel adjust

Delay adjustment for each channel might be required within the TDC. However, the delay difference between TDC channels will be less than a few ns and the trigger matching is performed with a 25 ns resolution. The channel adjustment is probably not a mandatory requirement as the final adjustments will be done in the DAQ system and off-line.

The channel adjustment circuit if implemented will be located before the first level buffer. The hit measurements can be adjusted with a programmable 8 bit channel dependent constant and separate adjustment constants can be available for the leading edge and the trailing edge measurements.

First level buffer overflow

In case the first level buffer runs full, the hit measurements from the channel buffers will be rejected. A special buffer overflow detection scheme takes care of handling the buffer overflow condition and properly marking events which may have lost hits.

When the first level buffer only have space for one more hit (full - 1) and a hit measurement arrives from the channel buffers the hit is written into the buffer together with a special buffer full flag. After this the buffer is considered full and following hits are rejected. When 4 measurements have been removed from the buffer by the trigger matching (full - 4) and a new hit arrives the buffer full status is cleared and the hit is written into the buffer with the buffer full flag set. This situation is illustrated in Fig. 3-9 where the two hits with the buffer full flag set is shown. These two hits define a time window where all hits have been rejected and this is used in the trigger matching to detect if an event potentially have lost some hits.

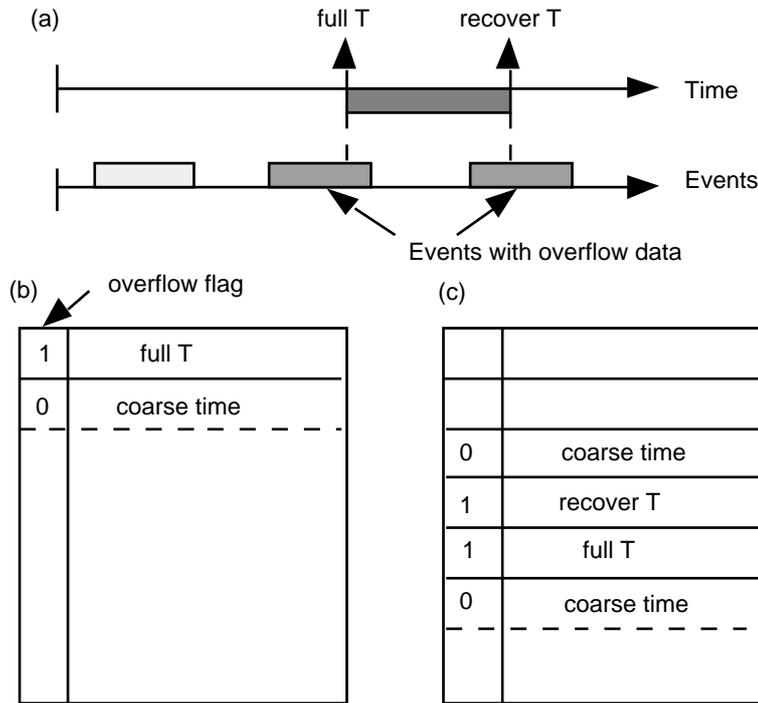


Fig. 3-9. First level buffer overflow handling. (a) Events of which matching time overlap with the period of overflow are marked. (b) Full time mark in the first level buffer. (c) Recover time mark.

3.4 Trigger interface and trigger FIFO

A positive trigger is in the trigger interface translated into a 12 bit trigger time tag together with a corresponding 12 bit event number. The trigger time tag is simply generated by loading the value of a synchronous counter when a positive trigger is signaled. The trigger time tag counter is at reset loaded with an offset different from the coarse time offset of the hit measurements. The difference between the two offsets determines the effective first level trigger latency in number of clock cycles. The trigger information is stored temporarily in a 8 words deep trigger FIFO to accommodate several triggers arriving within a short time interval.

In case the trigger FIFO runs full, triggers are rejected and complete events from the TDC are potentially lost. This would have serious consequences for the synchronization of events in the readout of the TDC. A full flag in the trigger FIFO together with the event number of the triggers are used to detect the loss of triggers and to make sure that the event synchronization in the readout is never lost. If a positive trigger is signaled when the trigger FIFO is full the trigger interface stores the fact that one (or several) triggers have been lost. As soon as the trigger FIFO is not any more full the event number of the latest lost trigger is written to the FIFO together with a trigger lost flag. The trigger matching can then detect that triggers have been lost when it sees the trigger lost flag. The trigger matching can also in this case determine how many triggers have been lost by comparing the event number of the previous trigger and the current event number stored together with the trigger FIFO full flag. This scheme is illustrated in Fig. 3-10 where it can be seen that triggers for event 11,12,13,14 have been lost. For each lost trigger the trigger matching will generate “empty” events (header + trailer) marked with the fact that it contains no hits because the trigger was lost.

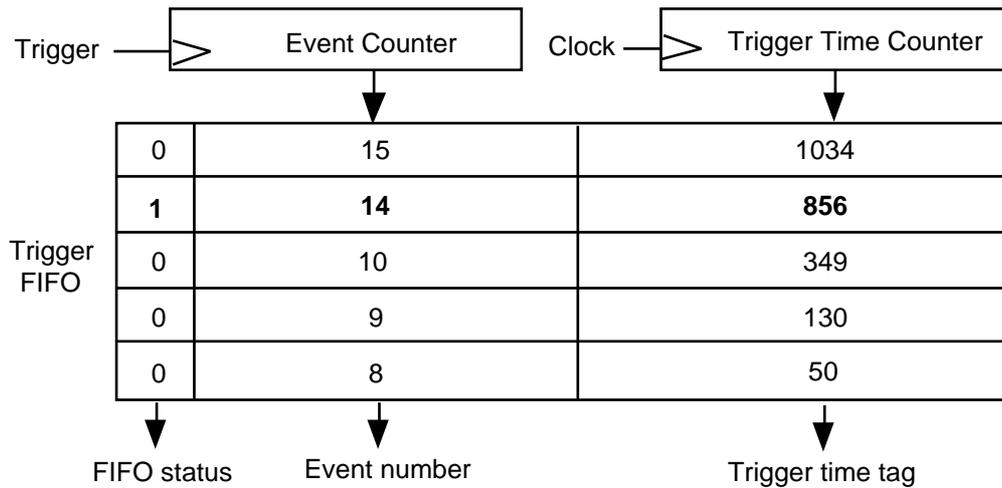


Fig. 3-10. Trigger FIFO overflow handling.

3.5 Trigger Matching

The basis for the trigger matching is the trigger time tag locating in time where hits belong to an event of interest. The trigger matching function selects the data related to the trigger from the first level buffer, and stores them into the read-out FIFO together with a header and a trailer. The header and trailer contains event ID, bunch ID, number of hits matched to trigger plus a set of flags.

A match between the trigger and a hit (leading edge) is detected within a programmable time window (matching window) to accommodate the maximum drift time of the detector (Fig. 3-11). The trigger is defined as the coarse time count when the event of interest occurred. All hits from this trigger time until the trigger time plus the matching window will be considered as matching the trigger. The trigger matching being based on the coarse count means that the “resolution” of the trigger matching is one clock cycle and that the trigger matching window is also specified in steps of clock cycles.

A mask window before the trigger time is optionally checked to search for possible hits which may have masked a hit in the matching window (section 2.1). The masked hit information is written to the readout buffer as one single word with individual flags associated to individual channels only if a masking hit has been found.

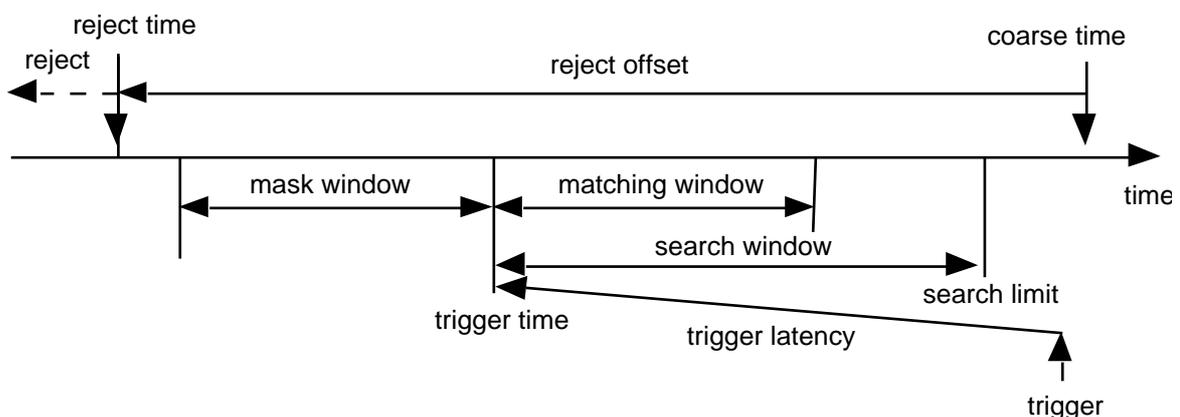


Fig. 3-11. Time relations of trigger matching and data rejection.

Trigger matching algorithm

The search for hits in the first level buffer, matching a trigger, can not be performed in a simple sequential manner for several reasons. As previously mentioned the hits are not guaranteed to be written into the first level buffer in strict temporal order. In addition a hit may belong to several closely spaced triggers. A fast and efficient search mechanism which takes these facts into consideration is implemented using a set of memory pointers and a set of programmable time windows;

- Write pointer: Memory address where new hit data is written.
- Read pointer: Memory address being accessed to look for a time match.
- Start pointer: Memory address where search shall start for next trigger.
- Mask window: Time window before trigger time where masking hits are to be found.
- Matching window: Time window after trigger time where matching hits are to be found.
- Search window: Time window specifying how far the search shall look for matching hits. (extends searching range to compensate for the fact that hit data are not perfectly time ordered in the first level buffer).

The pointers are memory addresses being used during the search as illustrated in Fig. 3-12 and the windows are measures of time differences from the trigger time to the time of the leading edge of the hit signal.

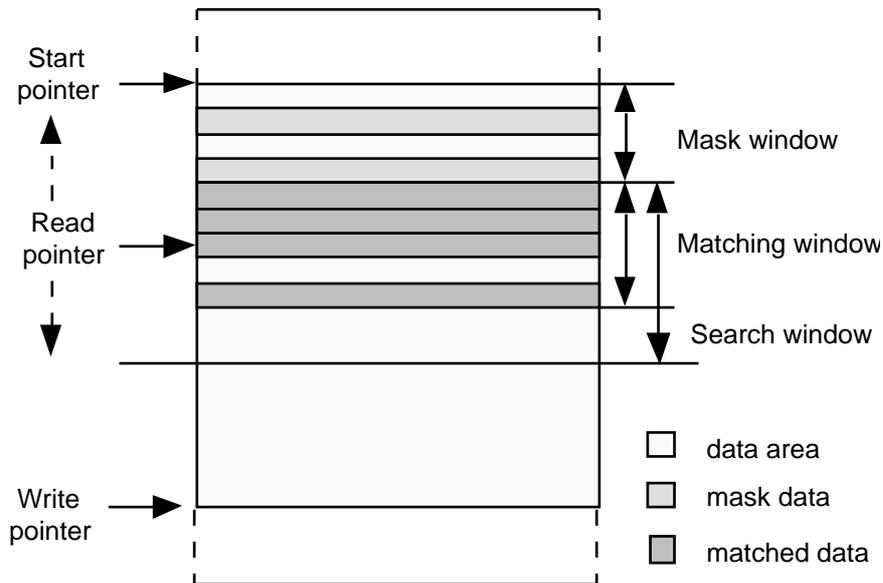


Fig. 3-12. First level buffer pointers and time windows.

The trigger matching search starts from the location pointed to by the start pointer. When the first masked hit or matched hit is found the start pointer is set to the new location. The search continues until a hit with a coarse time younger than the search limit has been found or there is no data in the buffer. The start pointer is also incremented while data rejection is being done as described below.

Data Rejection

To prevent the first level buffer to overflow, when no triggers have occurred for extended

periods of time, hits are automatically rejected from the buffer when getting older than a programmable reject limit (Fig. 3-11). The reject function is only active when the trigger FIFO is empty to prevent the removal of hits belonging to triggers which are waiting to be processed by the trigger matching.

The detection of hit data being older than the reject limit is based on a reject counter. This counter is at reset loaded with an offset different from the coarse time offset and the differences between the two offsets determines the effective reject time in number of clock cycles. A hit is rejected simply by incrementing the start pointer.

Rejecting matched hits when readout fifo full.

The trigger matching can optionally be programmed to reject matched hits when the readout fifo is full. This rejection can be made conditional on the fact that the first level buffer is more than 3/4 full. This option will prevent event data to pile up inside the TDC. Event data that have piled up inside the TDC will take very long time before arriving to the second level buffers in the DAQ system. Here they will probably be discarded as they arrive to late. In case the TDC is not programmed to reject matched hits, the trigger matching function will stop when the readout fifo is full and the l1 buffer and the trigger fifo will start to fill up.

Counter roll over

The trigger matching algorithm is strongly based on matching the time tag of a trigger to the time measurements of the hits. A potential problem may occur in such a scheme when one of the time tag counters (coarse time counter, reject counter, trigger time tag counter) overflows. This can relatively easy be taken care of when the overflow occurs at a “natural” binary limit [7]. At the LHC experiments the bunch structure defines that bunches are numbered from 0 - 3563 which is not a nice binary range. Several schemes of coping with this is currently being considered and they are explained in a separate paragraph dealing with the specific problems related to the bunch count reset (section 3.7).

Trailing edge matching

In case the TDC is programmed to perform paired leading edge - trailing edge measurements without extracting the pulse width measure, the trigger matching could in principle perform a trigger matching on the trailing edge and pass the corresponding leading edge measurement to the readout buffer in case of a match. This function will only be implemented if time allows and it is considered as an interesting option by the ATLAS muon community.

Disable matching

The trigger matching function may also be completely disabled whereby all data from the first level buffer is passed directly to the read-out FIFO. In this mode the TDC have an effective FIFO buffering capability of $128 + 32 = 160$ measurements. This mode is necessary for calibration and debugging of the detector and its electronics.

When writing hit data to the readout FIFO the trigger time tag is subtracted from the

leading edge measurement such that it is referenced to the time (bunch crossing) when the event of interest occurred. This reduces the number of bits in the event data to read out and makes it easier to analyze in the DAQ system. Optionally, the TDC can be programmed not to subtract the trigger time tag if a charge measurement is not required.

3.6 Readout

The read-out FIFO is 28 bits wide and 32 words deep and its main function is to enable events to be read out while others are being processed in the trigger matching. Data belonging to different events are separated by an event header and an event trailer.

Data from the readout FIFO is serialized and sent out via a serial link running at 80 Mbits/s. Several TDC's can share one serial link using a token based protocol to assemble data belonging to one event. The serial data from the TDC's in the daisy chain is sent from one TDC chip to the next together with a token finally arriving at the master TDC driving the serial link to the DAQ system. When the token returns to the master TDC and it has sent all its event data the transmission of the next event can begin. The serial data and the token are synchronized to the transmission clock (80 MHz) in each TDC in the chain to prevent excessive delays to build up when passing signals through the TDC's in the chain. The master TDC sending the optional global event trailer can if required send a common check sum (CRC) for all data belonging to an event.

When trigger matching is disabled no event headers and event trailers are generated and all hit measurements on the channels have to be read out. In this mode each TDC reads out one TDC measurement and immediately passes the token to the next TDC in the chain, even if it has more measurements waiting. This enables each TDC to have equal access to the shared readout link preventing a TDC with a noisy channel to monopolize the available read-out bandwidth.

The use of a daisy chain controlled by a single token is an architecture sensitive to failures in any of the TDC's in the chain. The possibility of bypassing the serial data and the token around a failing TDC increases the reliability of the system significantly. Bypassing of the master TDC is though difficult as it directly drives the serial link unless an external multiplexer is used to select a master chip.

An extended bypassing scheme is under consideration such that data to be send on a failing front-end link can be transmitted on a neighbor link. In this case the loading of the shared link will be doubled and a read-out bandwidth bottle neck may be the result if the background rates are higher than estimated.

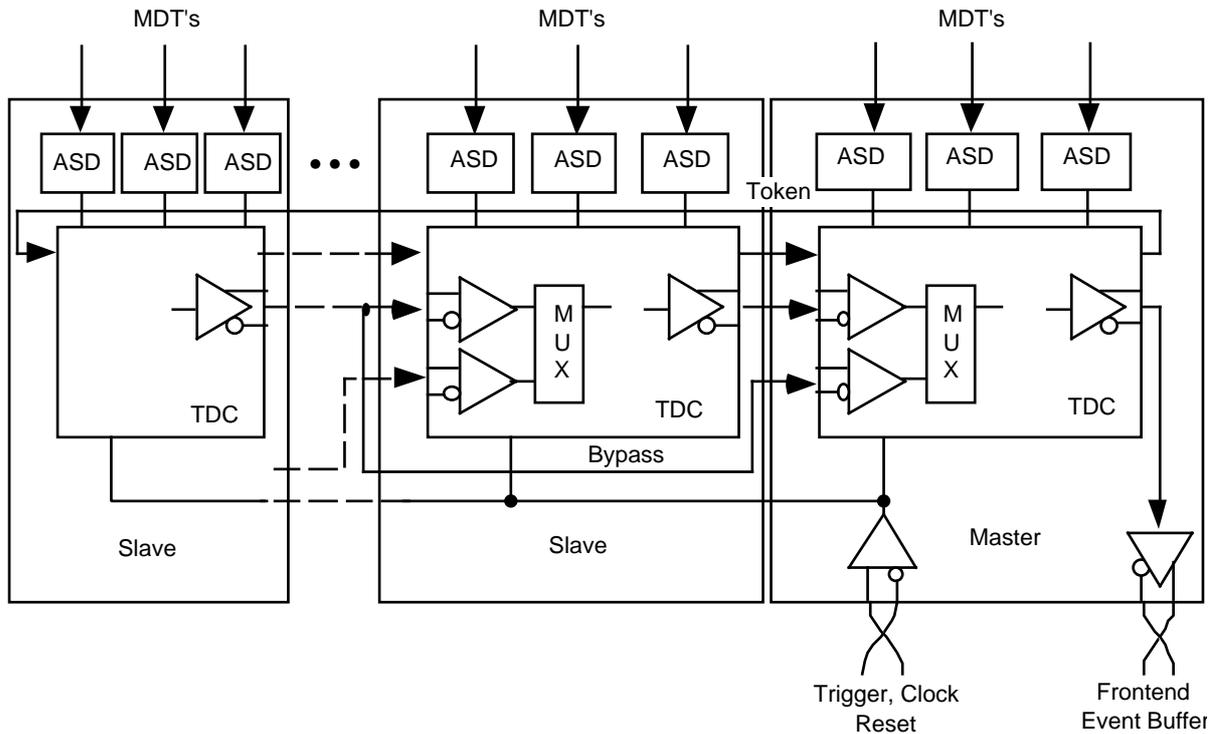


Fig. 3-13. TDC Readout in daisy chain connection with a bypass option.

Data on the serial links are composed of 32 bit words as illustrated in Fig 3-14. The first 4 bits are a type identifier used to distinguish between headers, trailers, and different kinds of event data. The following 4 bits are assigned to a TDC identifier. Each 32 bit word is preceded by a start bit (high) and followed by a parity and a stop bit (low). The serial front-end link is likely to use a two wire data encoding scheme used by the commercial transputer links. This encoding requires two wires but removes the need of any clock recovery circuitry in the front-end receiver (possibly implemented in a FPGA). No special encoding is used on the serial data sent from one TDC to the next in the chain.

The use of a global header, a global Trailer and the use of local headers and trailers is going to be individually programmable such that the readout protocol can be optimized for debugging or data taking purposes.

An alternative direct parallel readout from the readout FIFO will probably be implemented to ease the debugging and testing of the TDC chips. This parallel readout bus can in this case be completely disabled and the signal pads for the bus does not have to be bonded to the package during the final production of the TDC's for ATLAS.

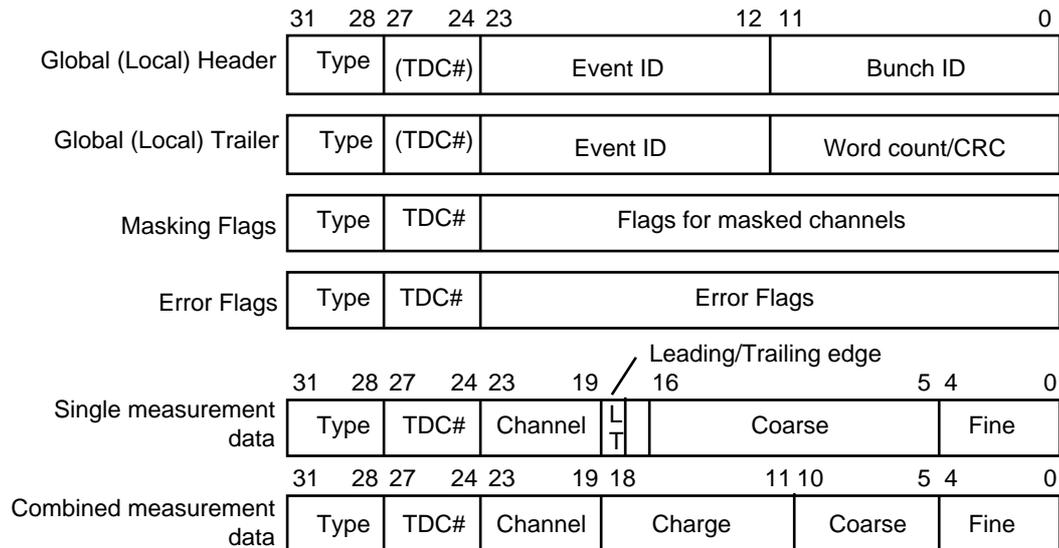


Fig. 3-14. An example of output data format.

3.7 Bunch count reset

As previously mentioned the overflow/reset of the time tag counters requires special attention. For the LHC experiments the bunch ID counters will be reset by a dedicated bunch count reset signal asserted every 3563 clock cycles. No bunch collisions will occur in the 127 clock cycles preceding the bunch count reset but it is possible that special calibration triggers will be generated during this window. Three different solutions to this problem is currently being considered.

Natural binary overflow in time tag counters: As previously mentioned the time tag counter overflows are handled straight forward [7] when it occurs at binary limits. All time tag counters in the TDC is in this solution only reset at the beginning of data taking and allowed to naturally overflow independently of the bunch count reset signal. To insure a correct bunch count ID associated to the triggers a special bunch ID counter is included in the trigger interface which value is stored in the trigger FIFO together with the event ID and the trigger time tag. This requires the word size of the trigger FIFO to be expanded by 12 bits. The time tag counters in the TDC will in this way only be reset at very infrequent intervals and one may fear that they get out of synch caused by a glitch on the clock or a radiation induced single event upset.

Forced roll-over at 3563 clock cycles: All time tag counters are in this approach loaded with their programmed offsets when the bunch reset is asserted. The time tag comparing logic in the trigger matching is changed to automatically take into account the roll-over at 3563 (must be programmable). This scheme requires changes to the otherwise simple time comparison logic.

Bunch reset separators: Special separation markers are inserted in the first level trigger buffer and the trigger FIFO to separate hits and triggers belonging to different LHC cycles. These bunch reset separators will occupy a small part of the buffers but will insure a clear distinction between hits and triggers belonging to different LHC machine cycles.

3.8 Event count reset.

The main function of the Event counter is to ensure that all event fragments from different subdetectors can be assembled into one complete event during event building in the DAQ system. The event count reset must reset all event counters at the same time so all event fragments gets assigned a unique event ID. The Event count reset will only be activated infrequently (possibly a few times per second). It will potentially be required that the front-end clears all their data buffers. The TDC will in all cases reset its event counter when an event count reset is detected. The used of a special buffer separator (equivalent to the bunch reset separator) or the reset of all state machines and buffers will be programmable.

3.9 Error monitoring

All the main functional blocks in the TDC must be continuously monitored for error conditions. The error status of the individual parts shall be accessible via JTAG. A special error output signal will be asserted when a fatal error condition that the TDC can not cope with itself has occurred (e.g. PLL losing lock, illegal state or data format detected). When this kind of error occurs a self test via JTAG should be performed as soon as possible to detect if a hard failure has occurred or if it is a soft error caused by a noise glitch in the system or a radiation induced single event upset.

The TDC being data driven means that its function depends strongly on data stored in its internal buffers to be correct. A soft error or a hard error in any of the memory elements may cause the TDC to mal-function until reset (event reset). It is therefore considered important that the content of all buffers are parity checked. In case a parity error is detected the corresponding data can be ignored and the TDC can continue to work correctly.

3.10 Technology and design methodology

Many different VLSI technologies and VLSI design methodologies are available on the commercial market place today. The most elaborate design methodology is full-custom which have the highest level of freedom and can achieve the most dense and high performance chips. A simpler design method is the use of gate arrays. Here all the transistors are pre-fabricated, and only the wiring is customized for each user. In gate arrays a library of basic blocks are used as the basis for the design. An intermediate design methodology is the use of a standard cell library for basic logic functions combined with full custom designed blocks for critical functions.

New high performance technologies for the development of Application Specific Integrated Circuits (ASIC) are normally available first as gate arrays compensating to a high degree for their reduced flexibility. The final choice of technology and design methodology will depend on their price and availability at the point in time when the final chip is going to be submitted.

Using modern Hardware Description Languages (HDL: Verilog or VHDL), to specify and simulate the architecture of the TDC, and finally use logic synthesis tools to translate the HDL description into gates, the final decision of which technology to use can be made at a very late stage in the development.

3.11 Summary of Specifications

The main specifications of the TDC are summarized in Table 3-1.

Table 3-1. ATLAS MDT TDC Specifications

Technology	0.3 - 0.5 μm CMOS gate array
Master gate size	\sim 200 k gates
Number of channels	24 channels (16 - 32 channels)
Clock frequency	40 MHz
Time bin size	0.78 ns
Time resolution	$<$ 300 ps RMS
Differential non linearity	$<$ 100 ps RMS
Integral non linearity	$<$ 100 ps RMS
Variation with temperature	$<$ 100 ps RMS
Offset between channels	$<$ 2 bins
Dynamic range	17 bit
Double pulse resolution	15 ns
Hit rate	$<$ 300 kHz (*)
Trigger rate	$<$ 200 kHz (*)
First level buffer	128 words x 31 bit
Readout FIFO	32 words x 28 bit
Trigger FIFO	8 words x 25 bit
Power supply	3.3 volt
Power consumption	\sim 10 mW/channel
Temperature range	0 - 70 $^{\circ}\text{C}$
Signal input level	IEEE 1596.3 SCI-LVDS standard
Other I/O level	CMOS level
Package	0.8 mm pitch, \sim 100 pins plastic QFP ?

(*) Maximum rate where loss of data is almost negligible.

4. Simulation Study

Particle hits in detectors are of statistical nature and it is very difficult to estimate the performance of a given TDC architecture in a specific application based on a few specified parameters. The baseline TDC has in its internal architecture a merging of hits from 24 channels into one common first level buffer. The two hit buffer per channel acts here as a limited de-randomizing buffer. The merging of the 24 channels is performed at the full clocking speed of the TDC and is not such a narrow bottleneck as one may tend to believe. The next bottleneck in the system is the read-out and merging of data from several TDC's. Here the data rate may though have been significantly reduced if an efficient first level trigger mechanism is used.

To get a feeling of the performance of a TDC in a real system it is required to perform simulations of the architecture using hits signals with realistic characteristics. In this section, several simulations of the baseline TDC architecture under different conditions are presented based on about one million hits per TDC chip by using the Verilog simulator [9].

The performance of the TDC at high hit rates is sensitive to the statistical properties of the hit inputs. If they are highly correlated, many hits may arrive within a small time window. This occurs if a TDC chip is connected to several detector planes behind each other and a large fraction of particles in the detector are contained in highly concentrated jets.

4.1 Simulation condition

The baseline simulation conditions are the following:

- 100 kHz hit rate (2/3 random hits and 1/3 correlated hits)
- 100 kHz trigger rate
- 24 channels per chip
- 40 MHz clock rate
- drift time = 0 ~ 600 ns
- pulse width = 30 ~ 150 ns
- dead time = 150 ~ 650 ns
- number of clock cycles per beam revolution = 3564
- trigger latency = 2.5 μ s
- mask & matching windows = 800 ns
- search window = 1200 ns
- reject offset = 4 μ s
- data format = 32 bit / word (Global header, hit data, conditional masked hit, conditional error, and global trailer word).

Drift time values without magnetic field is used in the above conditions, since a shorter drift time and short dead time are the most demanding conditions for the TDC. The effect of sharing a serial link for the readout is not included in the basic simulations of the TDC architecture. The TDC is considered the only driver of a 80 Mbits/s link. The effects of several TDC's sharing a serial readout link is shown in a separate paragraph.

The TDC architecture is also simulated under 300 kHz and 20 kHz hit rate conditions. The 300 kHz hit rate condition shows how the TDC behaves under much higher hit rates than expected in the MDT detector and the 20 kHz hit rate condition shows the behavior without any safety factor.

4.2 Hit and trigger generation

Fig. 4-1 shows the distribution of generated hit intervals and trigger intervals at 100 kHz. Each distribution has an exponential distribution with an average interval of 10 μ s as expected. The detector dead time rejects about 0.9% at 20 kHz, 4% at 100 kHz and 11 % at 300 kHz of the hits before reaching the TDC. About 0.3% of the triggers are rejected before coming to the TDC by the trigger conditions described in section 2.3.

Fig. 4-2 shows the distribution of number of hits matched with triggers with an average of 1.87 for the baseline conditions. This agrees well with the estimated value of :

$$\frac{\text{Matching Window (800 ns)}}{\text{Hit Interval}(10\mu\text{s})} \times \text{Tube Efficiency (96\%)} \times 24\text{ch} = 1.84 \text{ hits/trigger.}$$

At 20 kHz hit rates the average value is 0.38 compared to an estimated value of 0.38. At 300 kHz hit rates the average value is 5.10 compared to an estimated value of 5.13.

In addition to the matched hits, there is a masked hit word if hits are found in the mask window. The fraction of events having one or more masking hits at baseline conditions have been found to 85% which must be compared to the calculated value of:

$$1 - 1 - \frac{\text{Mask Window (800 ns)}}{\text{Hit Interval (10}\mu\text{s)}} \times \text{Tube Efficiency (96\%)} = 0.85$$

For 20 kHz hit rates 32% of all events will have a mask flag and for 300 kHz hit rates 99% will have a mask flag. The average total event size (without header/trailer) for the 24 channel TDC will be;

$$\begin{aligned} 0.38 \text{ data} + 0.32 \text{ masked hit} &= 0.70 \text{ words/event (@20 kHz)} \\ 1.87 \text{ data} + 0.85 \text{ masked hit} &= 2.72 \text{ words/event (@100 kHz),} \\ 5.10 \text{ data} + 0.99 \text{ masked hit} &= 6.09 \text{ words/event (@300 kHz).} \end{aligned}$$

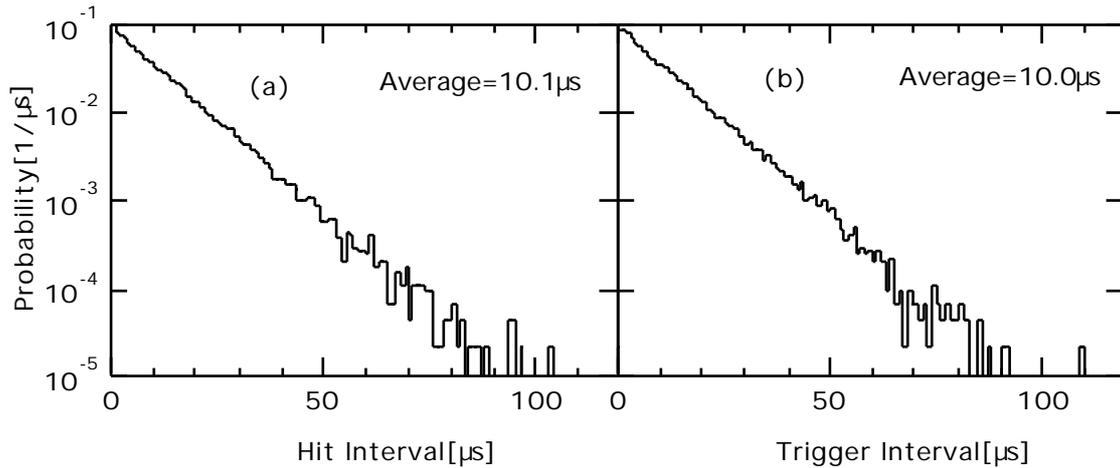


Fig. 4-1. Distribution of hit and trigger interval at 100 kHz.

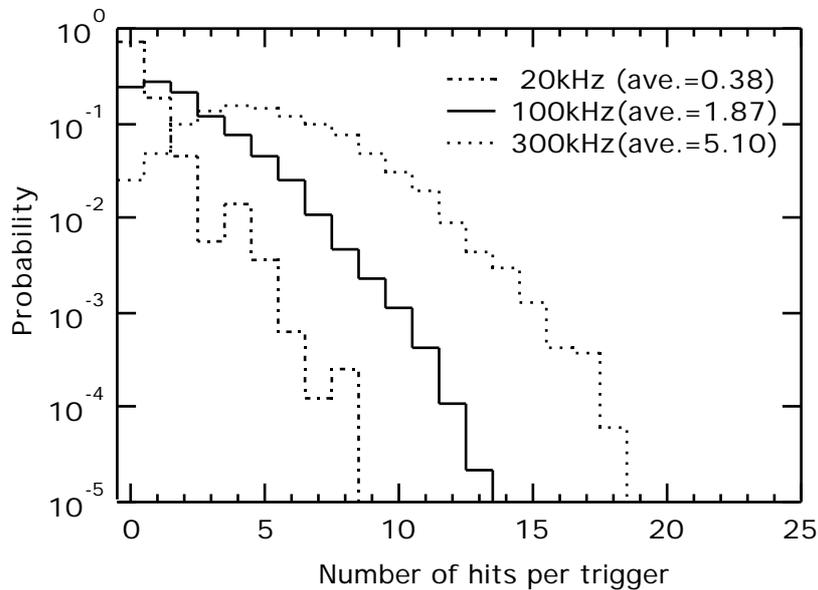


Fig. 4-2. Number of matched hits per trigger for 20, 100, and 300 kHz hit rate.

4.3 Channel buffer occupancy

The data transfer from the channel buffer to the first level buffer is an essential part of this TDC architecture. In present Verilog models it takes 4 clock cycles to transfer a single hit into the level one buffer. When several hits are waiting, the first takes 4 clock cycles to transfer and

the following are handled with a rate of one hit per clock cycle. To make a conservative simulation an additional clock cycle has been added to the transfer time of the first hit.

Fig. 4-3 shows the occupancy of the channel buffer for the baseline condition and a for hit rates a factor 3 higher. For each arriving hit, the status of the hit registers are recorded. If the two hit registers are full at this moment, it means that the hit is rejected. The average occupancy is only 2.1×10^{-4} and 8.3×10^{-4} for 100 kHz and 300 kHz hit rates respectively. No missing hits have been seen in a one million hits simulation. The main reason of this low (zero) rejection rate is the relatively large dead time (150 - 650 ns) of the MDT tube itself. There is plenty of time before the next hit arrives to the TDC, where the hit can be moved to the first level buffer. In other words, closely spaced track signals are already rejected in the detector and analog front-end.

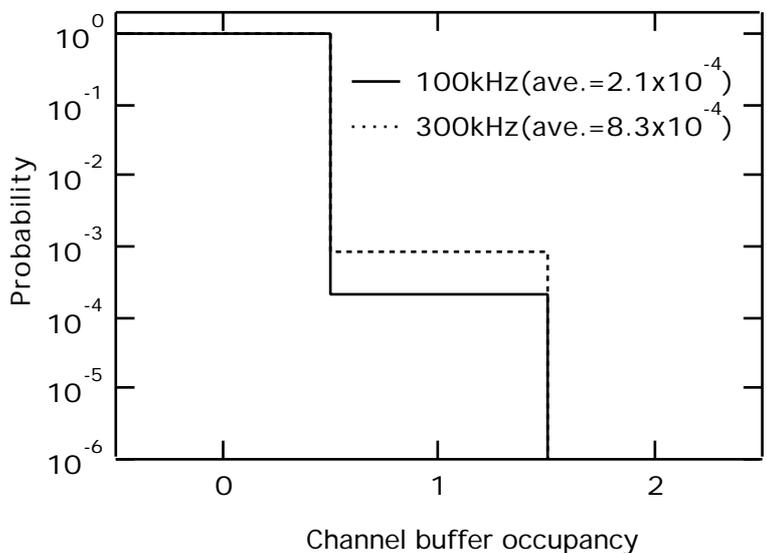


Fig. 4-3. Occupancy of channel buffer. Buffer occupancy 0 means there is no hits in the buffer when a new hit arrives; 1 means one waiting, and finally 2 means that a hit is rejected because the channel buffer is full.

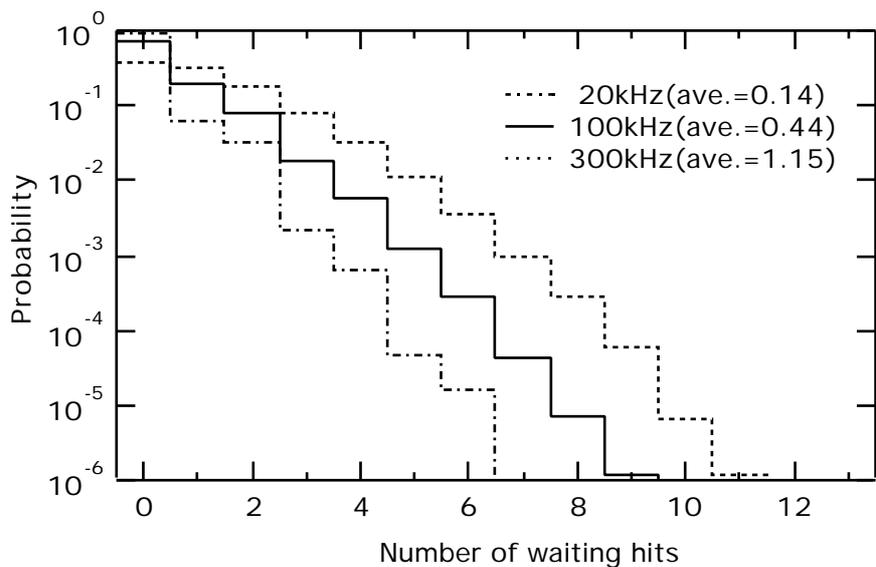


Fig. 4-4. Number of waiting hits in 24 channel buffers.

The number of waiting hits in the 24 channel buffers when a new hit arrives is shown in Fig. 4-4. The average number is only 1.15 at 300 kHz hit rates. This figure also implies that the search window needs to be several clock cycles longer than the matching window to insure to find all matching hits. Simulations show that a search window of 1200 ns insures that all matching hits are always found.

4.4 Double track separation.

In case double track separation is required in critical parts of the muon detector, the amount of hits (primary, secondary and fake hits) the TDC must accept is significantly increased. To verify an acceptable functionality of the TDC in this mode a large series of simulations have been performed. The major difficulty of the TDC in this mode is to accept fake hits plus a potential secondary hit during the drift time of the primary hit. The critical part of the TDC for this mode of operation is the limited size of the channel buffers (two pairs or 4 edges).

The hit rejection has been simulated for a series of primary hit rates largely covering the rates encountered in the MDT detector. The acceptance of hits has been divided into three different sub groups. Primary hit: single track hit, Secondary hit: hit from double track, Fake hit: Noise hits from clustering of tube signal. The acceptance of the different hit types are shown in the table below. For the comparison the hit acceptance in single hit mode is also shown.

In the first four cases the TDC performs paired measurements to measure the pulse width (charge) of primary hits. In the last case the TDC only performs leading edge measurements doubling the effective size of the channel buffer. The effective number of hits per track for the multi hit mode is shown as a histogram in Fig. 4-5.

Table 4-1. The acceptance of hits.

Primary hit rate	Single hit mode	Multi hit mode		
	prim. hit	prim. hit	second. hit	fake hit
50 kHz (pair)	100%	100 %	94.1 %	94.2 %
100 kHz (pair)	100%	99.99 %	93.1 %	93.0 %
200 kHz (pair)	100%	99.97 %	92.6 %	92.0 %
400 kHz (pair)	100%	99.86 %	90.4 %	90.1 %
400 kHz (leading)	100%	99.98 %	99.8 %	99.7 %

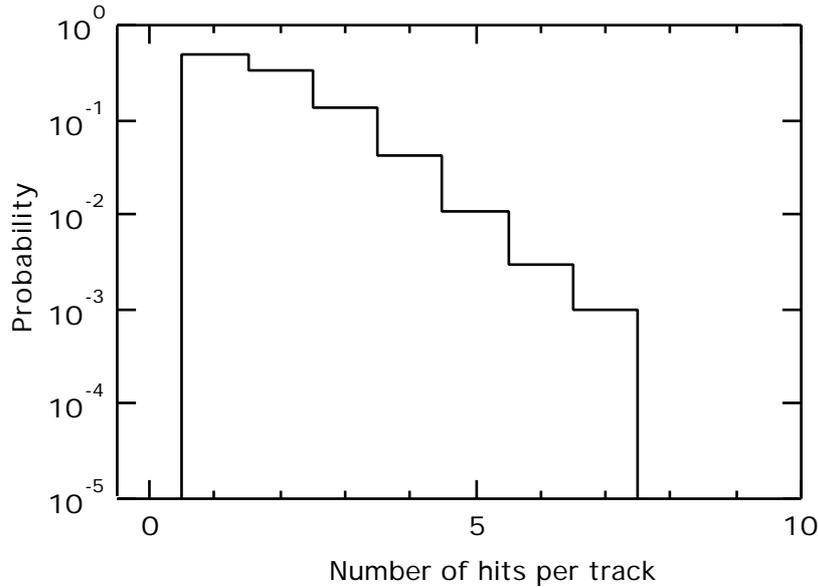


Fig. 4-5. Histogram of the effective number of hits per track for the multi hit mode.

It can be seen that the acceptance of secondary hits are in all cases better than 90 %. It should here be kept in mind that the effective hit rate seen by the TDC is 400 kHz * 1.6 = 640 kHz at the highest rate.

At an effective hit rate of 640 kHz per channel the amount of data to store and to read out is significantly increased. Simulations show that a first level buffer depth of 128 words is still sufficient. Read-out of the trigger matched hits via the serial link requires in this case the full bandwidth for one TDC chip only (no sharing possible at this rate).

4.5 First level buffer occupancy

The average level 1 buffer occupancy when not taking into account the processing needed for trigger matching can be found from:

$$24 \times \frac{\text{TriggerLatency (2.5}\mu\text{s)} + \text{Mask Window (800 ns)}}{\text{Hit Interval (10}\mu\text{s)}} = 7.9 \text{ (@100KHz)}$$

The simulated first level buffer occupancy is shown in Fig. 4-6. In the baseline condition the average buffer length is 8.9 (calculated 7.9) and with a maximum of 34 (at the level of 10^{-6}). At 300 kHz hit rates the average value increases to 24.9 (calculated 23.7) and the maximum reaches 60. The simulated average occupancies are as expected slightly larger than the calculated occupancies as hit data have to be stored for an extended period while the trigger matching searches for matching hits. Based on these simulations a 128 words deep level 1 buffer is estimated to be sufficient for the ATLAS MDT detector.

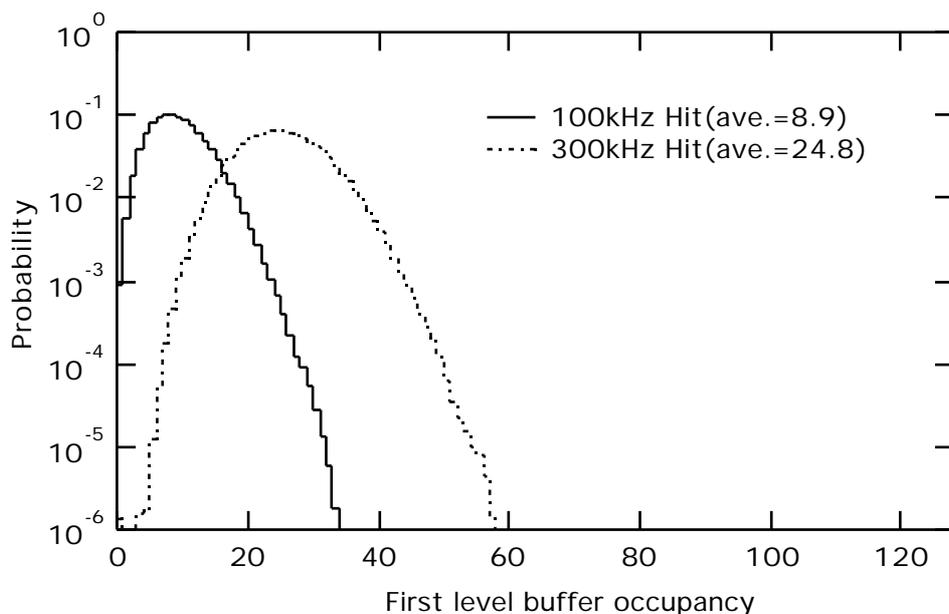


Fig. 4-6. Occupancy of the first level buffer..

4.6 Trigger Search time

Fig. 4-7 shows the time needed to search for all hits corresponding to a trigger. In the baseline condition the average search time is 9.6 clock cycles and at 300 kHz hit rates 19.9 clock cycles.

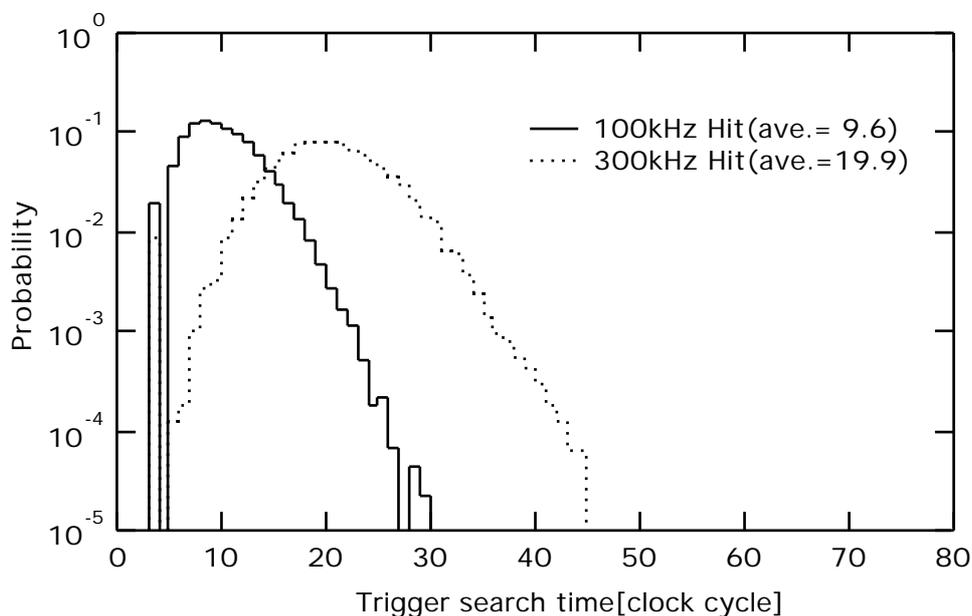


Fig. 4-7. Trigger search time for 100 kHz and 300 kHz hit rates.

4.7 Trigger FIFO occupancy

Fig. 4-8 shows the occupancy of the trigger FIFO under the baseline conditions, the 300 kHz hit rate condition, and a 200 kHz trigger rate and 100 kHz hit rate condition. Normally only one or two words are used, and there is little dependence on both the hit rate and the trigger rate. A trigger FIFO depth of 8 seems more than sufficient under these conditions.

When creating a readout bottleneck by sharing a serial readout link the occupancy of the trigger FIFO though increases (see readout simulations).

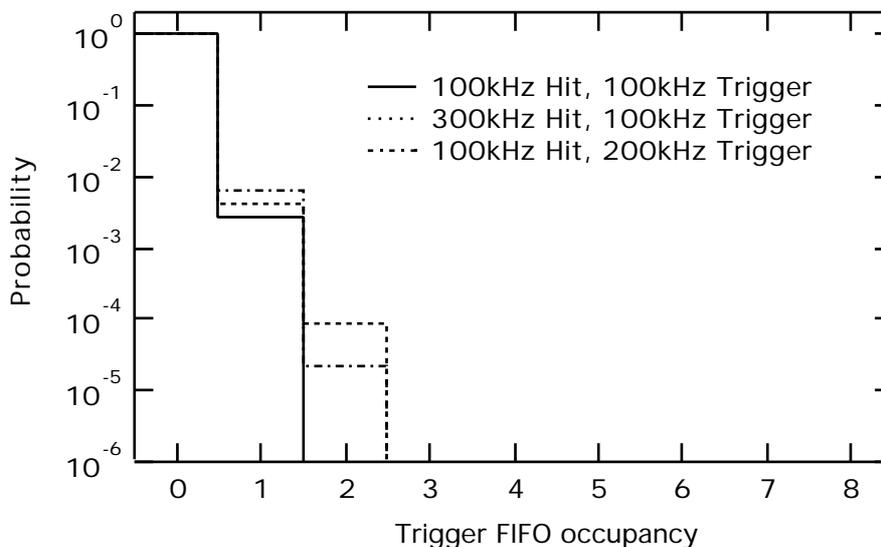


Fig. 4-8. Trigger FIFO occupancy for various conditions.

4.8 Readout system

The readout architecture of the TDC is based on the sharing of a serial readout link by several TDC's connected together in a daisy chain. The performance of this readout architecture depends heavily on the number of TDC's sharing a serial link and the required amount of data which have to be read out. The readout of TDC data is decoupled from the rest of the TDC by the 32 words deep readout FIFO. When this FIFO does not run full the readout protocol will have no effect on the performance of rest of the TDC chip. In case the readout FIFO runs full the trigger matching of triggers will be stopped and the occupancy of the level 1 buffer and the trigger FIFO will be directly affected.

The amount of event data from each TDC per event is directly proportional to the hit rate per channel and the size of the trigger matching window. The readout of mask flags and the use of local event headers and trailers may add significant data overheads especially at low hit rates. The total bandwidth required on a shared serial link is directly proportional to the trigger rate. The trigger rate is considered a well defined parameter and only simulations with the trigger conditions specified in paragraph 2.3 is presented. The hit rates have a high level of uncertainty and variability from chamber to chamber. A hit rate of 100 kHz is used as a worst case scenario and a hit rate of 20 kHz is used as a realistic case.

The readout FIFO could in principle be made very deep to be capable of utilizing the available readout bandwidth nearly 100 %. This will though have serious effects on the latency from a trigger is given until all event data have been transferred to the DAQ system. The maximum allowed event delay will be limited by requirements from the 2nd level trigger system. The exact requirements from the 2nd level trigger system is though at the moment not yet well defined.

The use of local event headers and trailers for each TDC are not included in the simulation cases as this mode of readout is only intended to be used while debugging the readout system

when the LHC machine runs at reduced luminosity. The readout of global headers and trailers and the conditional mask flags are included in all simulations shown. The data bandwidth on the serial link is supposed to be 80 Mbits/s using both edges of the 40 MHz system clock (or a 80 MHz clock from a PLL). No data encoding (Manchester or alike) is included in the serial protocol.

The results of the readout simulations are shown as a histogram of all the readout FIFO occupancies and as the total event delay from the trigger until event data from all TDC's have been transferred to the ROD. In case the merging of data from several TDC's generates a bottleneck such that the level 1 buffer and the trigger FIFO occupancies are effected the corresponding histograms of these are also shown. (Data is sent from TDC1 at first, then TDC2 .. follows)

2 TDC's per serial link

Fig. 4-9 shows readout FIFO occupancies and event delay for 2 TDC's per serial link. At 100 kHz hit rate the readout FIFO of the TDC2 become full. Maximum event delay is 10 μ sec and 20 μ sec for hit rates of 20 kHz and 100 kHz respectively.

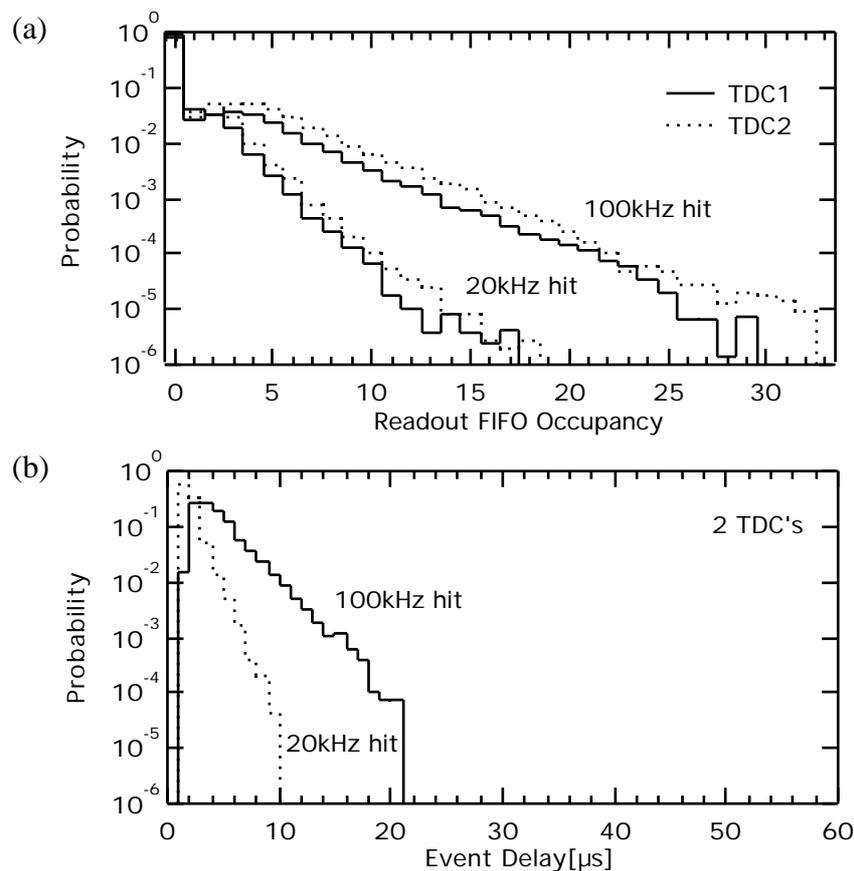


Fig. 4-9. 2 TDC's per serial link. (a) readout buffer occupancies, and (b) event delay.

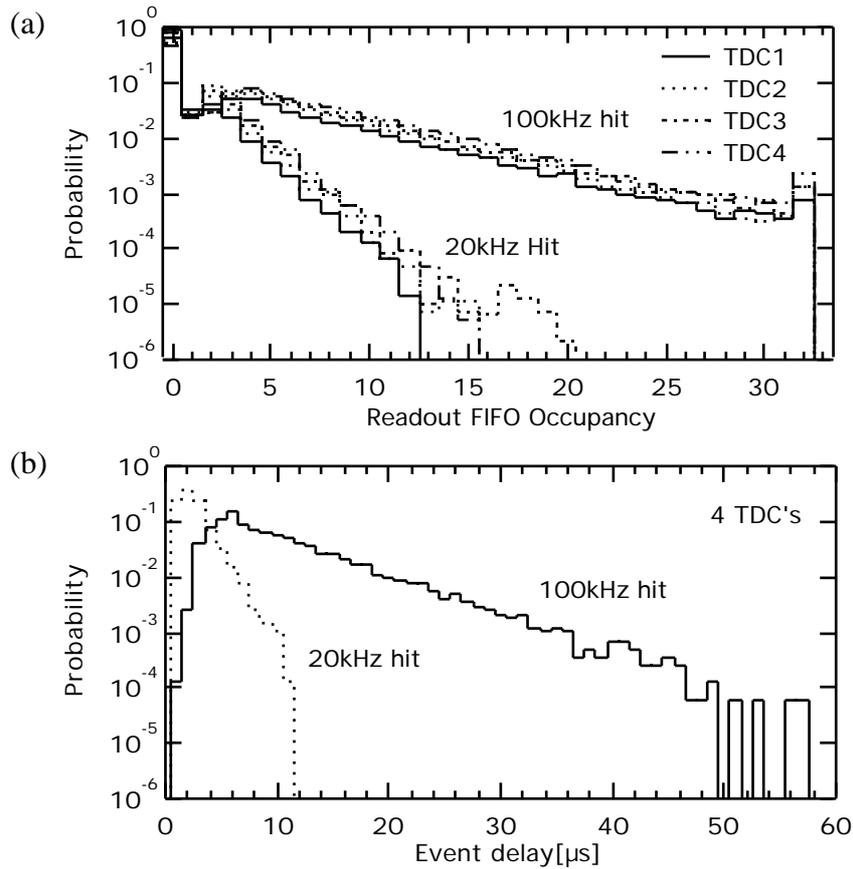


Fig. 4-10. 4 TDC's per serial link. (a) readout buffer occupancies, and (b) event delay.

4 TDC's per serial link

For 4 TDC's per serial link readout FIFO occupancies and event delay are shown in Fig. 4-10. At 100 kHz hit rate the readout FIFO become full at the probability of 10^{-3} .

Fig. 4-11 and Fig. 4-12 shows level 1 buffer occupancies and trigger FIFO occupancies at 100 kHz hit rate respectively. The depth of these buffers are still enough in this condition.

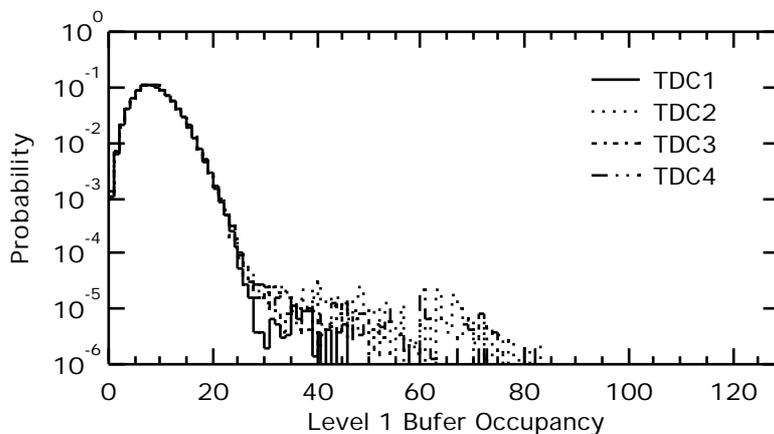


Fig. 4-11. Level 1 buffer occupancies with 4 TDC's per serial link at 100 kHz hit rates.

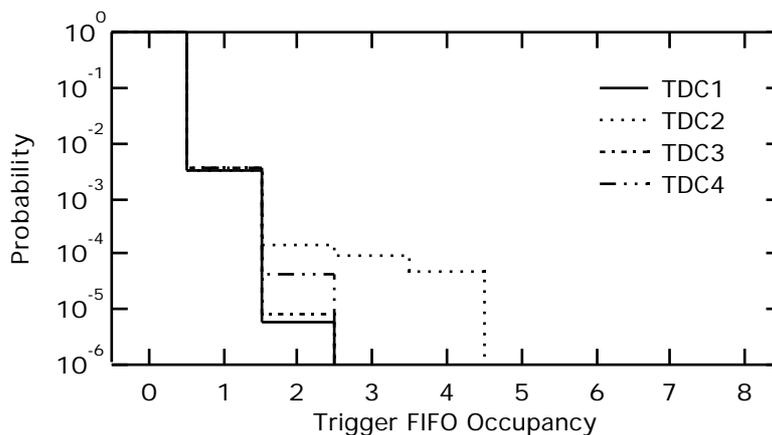


Fig. 4-12 Trigger FIFO occupancies with 4 TDC's per serial link at 100 kHz hit rates.

8 TDC's per serial link

When daisy chaining 8 or more TDC's it is not possible to sustain hit rates of 100 kHz as the average required bandwidth of: $100 \text{ kHz} \times (2 + 8 \times (1.87 + 0.85)) \times 32 = 76 \text{ Mbits/s}$ leaves too small margins for fluctuations in hit and trigger rates. At 20 kHz hit rate, the readout FIFO occupancies and the event delay are shown in Fig. 4-13

16 TDC's per serial link

When 16 TDC's are connected in a serial line, the readout FIFO becomes full even at 20 kHz hit rate as shown in Fig. 4-14. In addition, the event delay becomes about 100 μsec . Therefore 16 TDC's per serial link can be used in very low hit rate regions only.

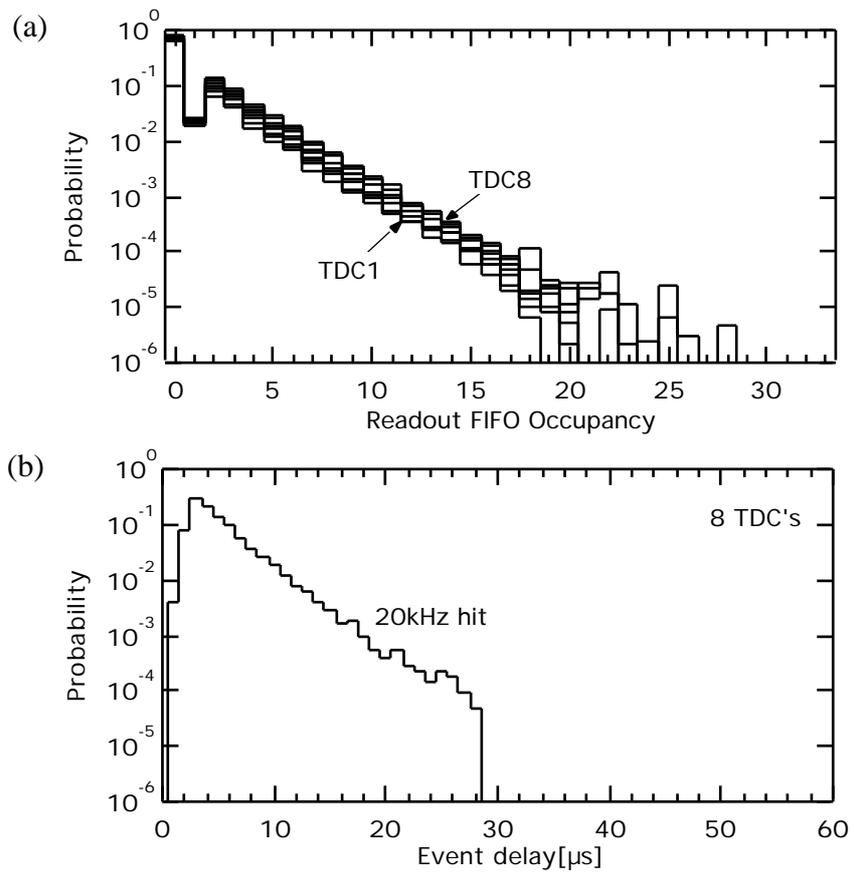


Fig. 4-13. 8 TDC's per serial link at hit rate of 20 kHz. (a) readout buffer occupancies, and (b) event delay.

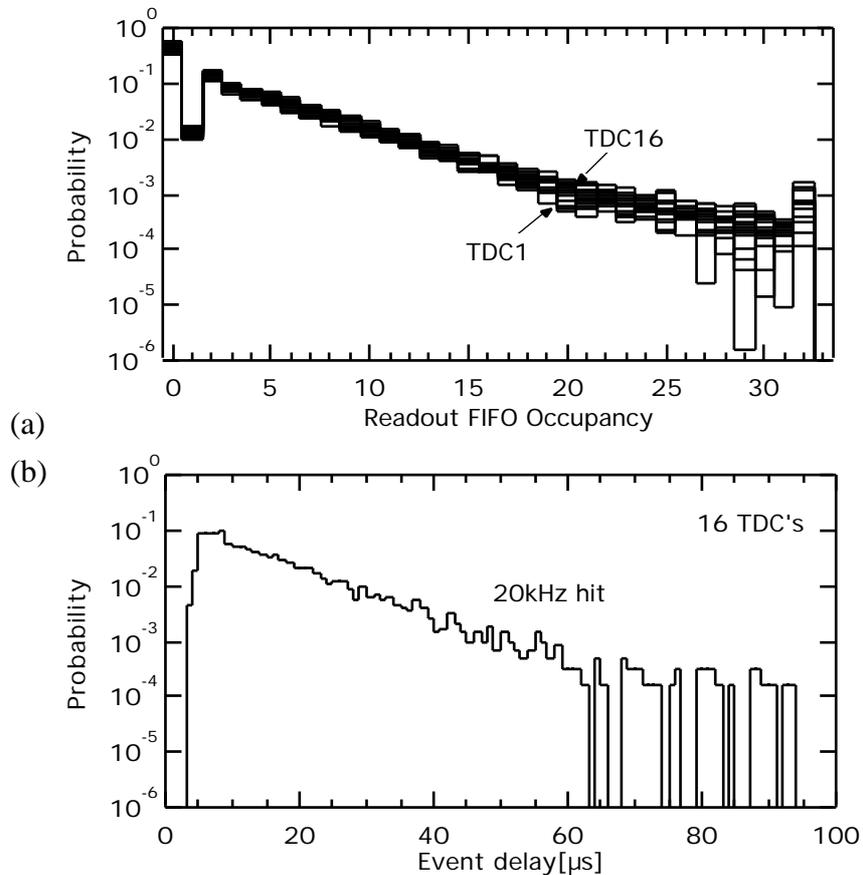


Fig. 4-14. 16 TDC's per serial link at hit rate of 20 kHz. (a) readout buffer occupancies, and (b) event delay.

5. Schedule and Plan

A general view of the schedule is shown in Fig. 5-1 and a milestone list of the Atlas Muon TDC (AMT) is shown in Table 5-1.

The first one or two years will be used to perform extensive simulations of the TDC and to design and verify a test chip with all the basic functions required. This chip must be used for full chamber tests to verify its correct function when embedded in the system. At the same time, radiation hardness of the candidate process will be measured by using the test chips.

Based on the experience obtained with this chip, the final mass production TDC will be designed in next two years or so.

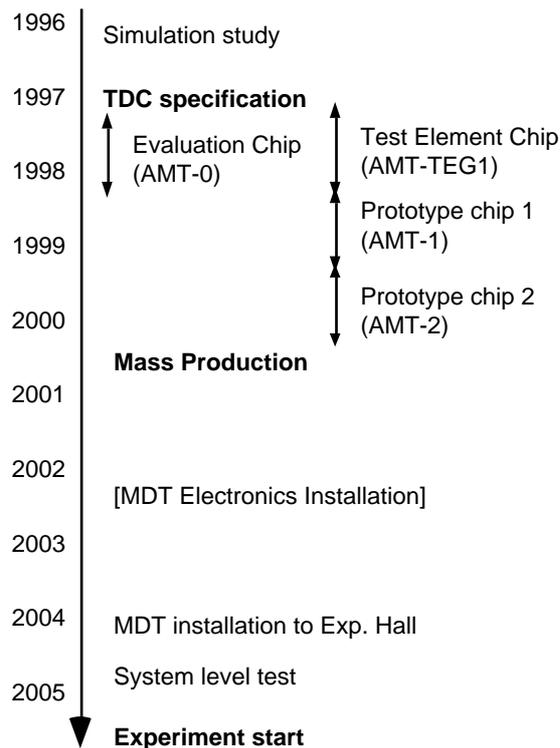


Fig. 5-1. General view of the TDC development schedule.

Table 5-1. Milestone list of Atlas Muon TDC development.

First evaluation TDC in 0.7 μm CMOS standard cell technology	
'97.5	Verilog behavioral model ready
'97.7	Verilog register level model ready
'97.8	Adaptation of DLL macro from 32 channel TDC ready
'97.10	submission of evaluation TDC (AMT-0)
'98.2	Test of AMT-0
'98.4	AMT-0 ready for system tests (limited quantity)
Mass Production TDC chip schedule in 0.3 μm CMOS gatearray technology	
'97.9	Finish Test Element Group (AMT-TEG1) chip design.
'97.11	AMT-TEG1 available.
'97.12	AMT-TEG1 gamma-ray irradiation test.
'98.3	Complete AMT-TEG1 tests.
'98.9	Finish 1st production chip (AMT-1) design.
'98.11	AMT-1 available.
'99.1	AMT-1 gamma-ray irradiation test.
'99.3	Complete AMT-1 test.
'99.9	Finish 2nd production chip (AMT-2) design.
'99.11	AMT-2 available.
'00.1	AMT-2 gamma-ray irradiation test.
'00.3	Complete AMT-2 chip test.

- AMT-0: Quick implementation of basic TDC functionality to enable early system tests to be performed. (Noise, cross talk, MDT tube + electronics resolution, Slew rate compensation by charge measurement, trailing edge measurement, serial readout, etc.). Implemented in a 0.7 μm 5.0 volt CMOS standard cell process from ES2.
- AMT-TEG1: Test chip of PLL, first level of hit registers, LVDS driver/receiver, and test structures for radiation test in CMOS 0.3 μm 3.3 volt gate array technology from Toshiba. (not useable for system tests).
- AMT-1: First complete implementation of MDT TDC in Toshiba process.
- AMT-2: Final production TDC (only required if changes must be made to AMT-1).

6. Summary

The requirements of the TDC for the MDT chamber have been summarized, and a baseline TDC specification has been presented. A TDC, using a DLL or a PLL for a fine time measurement and having several internal buffers is shown to be adequate for the ATLAS MDT.

Several simulations of the baseline TDC under realistic conditions shows that the present TDC architecture works satisfactory under 100 kHz hit rates per channel and a trigger rate of 100 kHz.

Results of a multiple TDC connection using a serial link is also shown. By using the shared link the number of cables and connections are significantly reduced. Furthermore a mechanism to bypass failing TDC chips have been presented.

Acknowledgments

We would like to acknowledge the people in microelectronics group at CERN for their continuous support of our work. Especially, we thank M. Turala, M. Letheren and A. Marchioro for giving us a chance to start our collaboration.

We also special thank to C. Fabjan and T. Kondo for their strong encouragement's and support of our work.

References

- [1] E. Hazen, J. Shank, T. Giordano, J. Huth, and J. Oliver, "Status of the Front End Electronics for the MDT System", ATLAS Internal Note MUON-NO-111, March 1996.
- [2] Y. Arai and T. Ohsugi; "An Idea of Deadtimeless Readout System by Using Time Memory Cell", Proceedings of the Summer Study on the Physics of the Superconducting Supercollider, Snowmass, 1986, p.455-457. KEK Preprint 86-64.
- [3] Y. Arai and T. Matsumura and K. Endo; "A CMOS 4 ch x 1 k Time Memory LSI with 1 ns/bit Resolution", IEEE Journal of Solid-State Circuits, Vol.27, No.3, March 1992, p359-364.
- [4] Y. Arai and M. Ikeno; "A Time Digitizer CMOS Gate-Array with a 250 ps Time Resolution", IEEE Journal of Solid-State Circuits, Vol. 31, No. 2, February, 1996, p212-220. KEK preprint 95-75.
- [5] H. Shirasu, Y. Arai, M. Ikeno, T. Murata and T. Emura; "A VME 32 ch Pipeline TDC Module with TMC LSIs", IEEE Trans. on Nucl. Sci., Vol. 43, No. 3, pp1799-1803(1996).
- [6] J.Christiansen, O. Klingsheim, C. Ljuslin, A. Marchioro, An integrated 16 channel CMOS Time to Digital Converter, IEEE trans. on Nuclear Science, Vol. 41, No. 4, pp1104 - 1108, August 1994.
- [7] J. Christiansen, "32 channel general purpose time to digital converter".
- [8] J. Christiansen, "An integrated high resolution CMOS timing generator based on array of delay locked loops", IEEE Journal of Solid State Circuits, Vol. 31, No. 7, July, 1996.
- [9] Verilog-XL, Cadence Design Systems, Inc.
- [10] "The LHC Conceptual Design Report", CERN/AC/95-05.
- [11] "Trigger and DAQ Interfaces with Front-End Systems: Requirement Document", Version 1.0, Atlas Trigger-DAQ Steering Group, January, 1996.
- [12] IEEE 1149.1 standard. Available from IEEE.
- [13] M. Passaseo, E. Petrolo, S. Veneziano; "A TDC integrated circuit for drift chamber readout. Nuclear Instruments Methods Phys. Res. A 367 (1995) 418-421.

Appendix A: Coding of TDC control signals

The TDC control signals: trigger, bunch reset, event reset and global reset is encoded on one clock synchronous line. Three clock cycles are used for the encoding as it is guaranteed that the minimum trigger interval is 75 ns.

clock period	N	N+1	N+2
Trigger:	1	0	0
Bunch reset:	1	1	0
Event reset:	1	1	1
Global reset	1	0	1

Appendix B: Definitions of read-out types and sub types

(This definition should be the same as seen in NIMROD specification.)

Type

1001	TDC global header
1010	TDC local header
1011	TDC global trailer
1100	TDC local trailer
0110	TDC error flags
0010	TDC mask flags
0100	TDC combined charge measurement (relative leading + charge)
0101	TDC combined trailing measurement (relative leading + pulse width 6.25 ns res.)
0111	TDC monitoring

TDC global header:

bits:

31:28	Type = 1001
27:24	TDC ID of master TDC in daisy chain
23:12	Event count
11:0	Bunch count

TDC local header:

bits:

31:28	Type = 1010
27:24	TDC ID
23:12	Event count
11:0	Bunch count

TDC global trailer:

bits:

31:28	Type = 1011
27:24	TDC ID of master TDC in daisy chain
23:12	Event count

11:0 Word count (including all headers and trailers)

TDC local trailer:

bits:

31:28 Type = 1100

27:24 TDC ID

23:12 Event count

11:0 Word count (including local headers and trailers)

TDC error flags:

bits:

31:28 Type = 0110

27:24 TDC ID

23:7 Reserved

6 Channel reject error

5 Overflow data flag

4 Self check error detected

3 Hit error detected (parity error in internal buffers)

2 Hit data rejected because readout fifo full

1 Event lost by trigger fifo overflow

0 Data lost by l1 buffer overflow

TDC masked flag channels:

bits

31:28 Type = 0010

27:24 TDC ID

TDC in 24 channel mode: 23:0 mask flags

TDC in 32 channel mode: 23 0 = channel group 15 - 0, 1 = channel group 31 - 16

22:16 zero

15:0 channel flags (channel 31 - 16 or channel 15 - 0)

TDC single measurement

bits

31:28 Type = 0011

27:24 TDC ID

23:19 channel

18 0 = leading, 1 = trailing edge

17 Reject flag (reject flag to mark that hit have been rejected on channel)

16:5 Coarse time

4:0 Fine time

TDC combined charge measurement

bits

31:28 Type = 0100

27:24 TDC ID

23:19 channel

18:11 charge (resolution 0.78 ns)

10:5 Relative leading edge coarse time

4:0 Relative leading edge fine time

TDC combined trailing measurement

bits

31:28 Type = 0101

27:24 TDC ID

23:19 channel

18:11 pulse width (resolution 6.25 ns)

10:5 Relative leading edge coarse time

4:0 Relative leading edge fine time

TDC monitoring

31:28 Type = 0111

27:24 Sub type 0000 = bunch reset or event reset separator

11:0 bunch count

0001 = buffer status

7:0 first level buffer occupancy

8 readout fifo full