

5 Front-End Electronics

This chapter describes the on-detector electronics as well as relevant work on prototyping that has already been performed. Concerning the TRD trigger, its functionality is outlined in Chapter 6 with emphasis on the implementation of the *tracklet* search and electron candidate identification. This Chapter details the implementation of the front-end electronics (FEE) with its real-time constraints, together with its integration into the ALICE trigger system.

5.1 Electronics overview

In this section the requirements for the FEE are reviewed and the general architecture and basic building blocks are introduced. Because we are interested in both the identification of the transition radiation (TR) signal and also in the TRD online tracking, momentum, and invariant mass reconstruction capability (see Chapter 6), the FEE is rather complex and more TPC-like involving a sampling ADC, tail cancellation, detection of overlapping hits etc. The trigger will generate a Level-1 accept (L1A) and therefore has to occur on a time scale of $6 \mu\text{s}$. This requirement drives the over all architecture, clock speeds and limits the extent of multiplexing possible.

5.1.1 General requirements

As detailed in Section 4.4 the FEE is used to read out and analyze for the Level-1 trigger the charge induced on 1.156.032 pads located in 540 individual readout chambers arranged in 6 layers in the TRD barrel. Most of the front-end electronics sits directly on the readout chambers. For the trigger, however, information from the 6 layers has to be combined at a convenient point close to all readout chambers. The readout chambers deliver on their pads a current signal with a very fast rise time and a long tail due to the slow motion of the Xe ions (see Fig. 4.14). The typical current for a minimum ionizing particle is of the order of $0.2 \mu\text{A}$. The pad on which the signal is induced can be viewed as a pure capacitance of 10-20 pF.

The main requirements for the front-end electronics are summarized in Table 5.1 and briefly discussed below.

- The space point resolution in the bending direction (y) is derived by charge sharing between 3 adjacent pads. The pad response function is chosen such that for a hit centered on one pad each neighbour still sees 10 % of the signal (see Fig. 4.11 and 11.9). This means adequate space point resolution can be reached for a signal to noise ratio of at least 30:1. Also, it was shown that digitization errors contribute visibly to the space point resolution if the channel number of the peak pad is below 30.
- For a minimum ionizing particle typically a charge of $3 \cdot 10^4$ electrons contributes to the signal on the maximum pad for each time bin. The requirement of signal to noise equal or larger than 30 defines the goal for an upper limit for the noise of 1000 electrons (r.m.s.).
- In order to not waste dynamic range, one aims to keep the noise amplitude within the ADCs LSB. For an ADC with 1 V dynamic range (see below), this fixes the conversion gain of the preamplifier-shaper (PASA) to 6.1 mV/fC .
- Our main interest is the detection of the TR signal superimposed on the normal ionization. As shown in Fig. 11.4 the TR photon energies reach with noticeable probability up to 20-30 keV. Simulations have shown that the electron-to-pion separation improves with dynamic range and for a minimum ionizing signal amplitude at ADC channel 30, a 10 bit ADC is desirable.

Table 5.1: Front-end electronics requirements.

Parameter	Value
Number of channels	$1.156 \cdot 10^6$
Signal-to-noise (MIP)	30:1
Dynamic range	1000:1
Noise (ENC)	1000 e
Conversion gain	6.1 mV/fC
Time bins in drift region	≥ 15
Separation of time bins	≤ 133 ns $\cong 2$ mm
Sampling frequency	(8-)10 MHz
Shaping time (FWHM)	$\cong 120$ ns
Cross talk	≤ 0.3 %
Bandwidth readout	15 TB/s
Bandwidth detector to GTU	216 GB/s
Bandwidth DDL	1.8 GB/s
Trigger latency at TRD	6.0 μ s
Trigger dead time (L0/L1 reject)	1.7 ... 7.0 μ s
Trigger dead time (L1 accept)	20 ... 40.5 μ s
Power consumption	≤ 50 mW/channel

- In Chapter 11 it is shown that in terms of tracking efficiency and momentum resolution it is sufficient to sample the drift region in 15 points (time bins).
- As shown in Section 4.5, for non-perpendicular angles of incidence the resolution is limited by the long ion tail of the Xe leading to a correlation of the individual time bins. This effect gets worse as the distance between time bins gets shorter or as the drift velocity is increased and the total drift time decreased. This constrains the drift time to be not smaller than 2 μ s and to an corresponding distance of two consecutive time bins of 133 ns for 15 time bins. That would correspond to a sampling frequency of 7.5 MHz. Since it is for other reasons convenient if the time intervals are multiples of the bunch separation a frequency of 8 MHz would be a good lower limit. Of course higher frequencies combined with a larger number of time samples would be possible and would slightly reduce the trigger latency due to the faster draining of the ADCs pipeline.
- In order to keep the correlation between the consecutive time bins of a track segment minimal to optimize resolution one would like a shaping time as short as possible. This is however connected with a loss in signal and also the existing long ion tail makes very short shaping times useless. A shaping time of 120 ns, comparable to the separation of the time bins, was found to be a good choice.
- The position and angular resolution can be improved by unfolding the time response function as demonstrated in Fig. 14.38 and 14.39 using a tail cancellation. Since this also noticeably improves the trigger performance (see Section 6.4.2.1) it is desirable that this deconvolution is done on the digital chip before the processing of the trigger information.
- The channel-to-channel cross talk is limited by the pad-to-pad capacitance which is between neighboring pads in one pad row 6.5 pF. This will lead to a cross talk of about 5 % for the present PASA design. The cross talk within the PASA chip itself and in the cable should be negligible as compared to this. It turns out that a value below 0.3 % was easily achievable in the existing PASA

prototype and this is the number quoted in Table 5.1.

5.1.2 System overview

The front-end electronics for the ALICE TRD consists of $1.156 \cdot 10^6$ channels. The basic building blocks are shown for one channel in Fig. 5.1. They are: a charge sensitive PreAmplifier/ShAper (PASA), the analog chip, a 10 Bit 10 MHz low power ADC, and digital circuitry where data are processed and stored in event buffers for subsequent readout. The data processing is performed on one hand during the drift time at digitization rate by the *Tracklet* Pre Processor (TPP) in order to prepare the information necessary for the *Tracklet* Processor (TP). On the other hand at the end of the drift time the Tracklet Processor, a micro CPU implemented as Multiple Instruction Multiple Data (MIMD) processor, operating at 120 MHz, processes the data of all time bins in order to determine potential *tracklets*. These *tracklets* are shipped to a Global Tracking Unit (GTU), which combines and processes the trigger information from individual TRD readout chambers.

Upon receipt of a L1 accept, the MIMD processor also ships the zero suppressed raw data from the event buffers on the front-end chips to the GTU, where they are stored in a large memory until read-out (see Chap. 7).

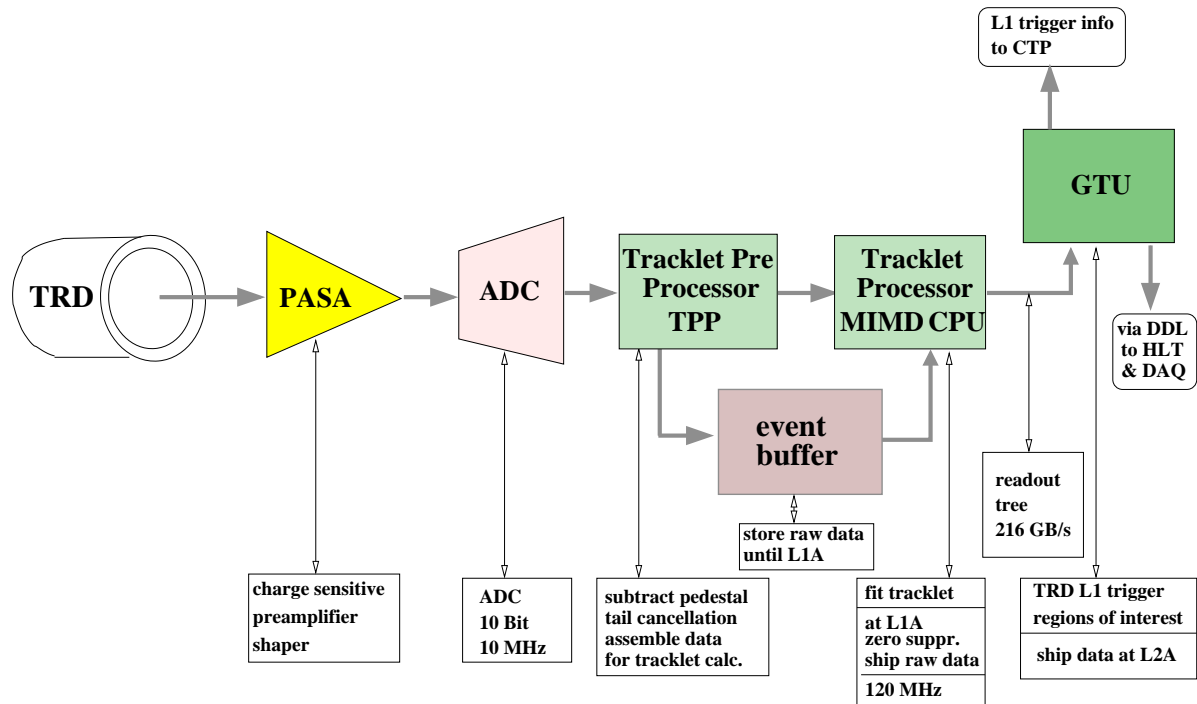


Figure 5.1: Basic logical components of the TRD front-end electronics. Everything but the GTU is located directly on the readout chamber. The ADC, digital circuitry, event buffer and MIMD CPU are combined into one digital chip. This chip determines the *tracklets* and is therefore also referred to as local tracking unit (LTU).

The requirement for minimal radiation length, power and cost drive the integration density as high as possible. In order to support mass production of the electronics, 18 channels are grouped together on one multi-chip module (MCM), housing both the preamplifier and the digital back-end (Section 5.8). The particular choice of 18 channels per MCM is a compromise of die size, MCM count and trace length of the analog pad signals. Figure 5.2 indicates the components on one MCM. Basically this module is targeted to contain just those two chips, and possibly the addition of minimal miscellaneous components, such as blocking capacitors. As sketched in Fig. 5.2, the 18 entities labelled 'PS' are on one chip, the PASA, everything enclosed by the thin grey rectangle is on a second digital chip. The logic integrated

on the digital chip includes the ADCs, the *tracklet* preprocessor, and the high-speed multi-threaded processor (MIMD CPU). Therefore this chip is called the local tracking unit (LTU). The pad plane itself carries the readout signals and they are routed to the PASA input via short cables. The MCMs, which are implemented as Ball Grid Arrays (BGA), are soldered directly onto the readout mother boards. The only additional circuitry required on the readout boards are the drivers for the clock fan-out and additional power filtering circuitry. All signals connecting the MCMs are routed on the readout mother board. There are 64224 MCMs mounted on the detector, making the MCM one of the most crucial electronics components, which have to be mass produced.

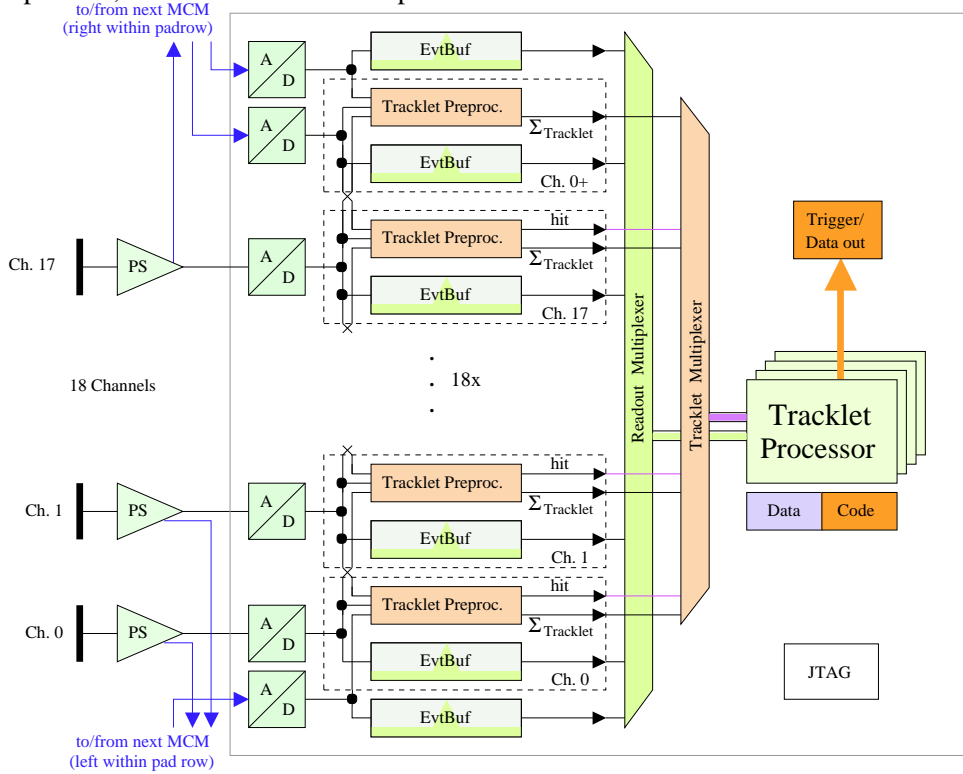


Figure 5.2: Overview of the electronics for 18 channels on one MCM.

Given the high (digital) clock rates and the low duty cycle of the trigger system of less than 3%, the digital part of the electronics is operated with gated clocks, allowing for the disabling of the clock to any part of the circuitry that is inactive. This method also permits the reduction of digital noise during digitization. All clocks are synchronous to the LHC clock.

In order to avoid granularity effects at the MCM border, some data need to be exchanged among the neighboring multi-chip modules. For a detailed description of the *tracklet* preprocessing architecture and the *tracklet* merging within the MCM and among neighboring MCMs, refer to Section 6.3. Since it is sufficient to merge *tracklets* only in ascending pad number direction, a total of three additional channels (two left and one right) is required to be processed, as indicated in the figure. Therefore, a total of 21 ADC channels is required for each 18-channel *tracklet* processor. Consequently, three out of 18 preamplifier outputs are required to drive two ADC inputs. In order to avoid any non-linearities, those channels implement two independent output stages, driving one ADC input each. The preamplifier outputs are analog differential signals.

All digitized ADC outputs, including the redundant channels, are stored in 32x10 Bit deep event buffers. During the digitization, the *tracklet* preprocessor identifies candidates and prepares them for later processing by the MIMD CPUs, the *tracklet* processor (Chapter 6.3). During that time, however, the digital back-end is operated at exactly the ADC clock rate.

At the end of the drift time, the fast digital clocks are enabled starting the MIMD processor. Any additional digital noise produced here is irrelevant as the relevant data already sit in the internal event buffers. During stand-by, all digital clocks are disabled. Figures 5.5 and 5.6 sketch which digital clocks are active at what point. The MIMD processor is capable of processing up to four *tracklet* candidates simultaneously. If a *tracklet* is identified and matches the required deflection cut requirements, its *tracklet* parameters are projected onto the global reference plane which is in the middle of the six planes and then forwarded to the readout tree.

The MCM output is a single 16 Bit differential data link, implementing Low Voltage Differential Signals (LVDS). There are additional 5 bits for correcting one bit errors and detecting two bit errors per data word. This format is used as standard link everywhere within the readout tree.

The readout tree terminates the differential output (LVDS) of all MCMs into four 16 Bit data links on either side of the detector per layer and sector, thus merging up to 304 MCMs into one high-speed data link to the global tracking unit. The readout is performed in a strictly ordered fashion to support consecutive readout and highly parallel global tracking. Any of the readout signals is kept inactive during acquisition or stand-by in order to minimize the electronic noise contribution. For a detailed description of the readout tree, refer to Chapter 7.

5.2 Chip technology

In general, there are five major components of the front-end electronics chain as summarized below:

- shaping preamplifier
- 10 Bit analog to digital converter
- digital filter for tail cancellation
- event buffer and *tracklet* preprocessor operating at ADC clock rates
- high-speed *tracklet* processor and filter
- high-speed readout tree

The shaping preamplifier is a full custom analog design tailored towards low noise and low power (in this order of priorities). The last three components are purely digital systems running at clock rates ranging from 10 MHz to 120 MHz. These clock rates can well be implemented using standard cell designs. The only requirement for full custom design are some special cells, such as the quad-port memories (refer to Chapter 6, Section 6.3.3.3). Although the first implementations of the digital circuits were designed for the AMS¹ 0.35 μm process, they can be ported to basically any silicon process. All three components, the *tracklet* preprocessor, filter and readout, are purely digital. They can all be implemented on the same die without presenting any particular technological challenges. In order to separate the analog and digital circuitry, the preamplifier will be designed as a separate chip in AMS 0.35 μm technology. One channel requires about 0.3 mm² in area, making this a fairly small chip.

The ADC, however, is a combination of analog and digital components and is acquired as an external cell. In principle, it could be implemented on both dies. The TPC design will integrate the ADC together with the ALTRO digital readout chip. A similar choice was made for the TRD. This choice represents a compromise with the advantages and disadvantages outlined below.

¹Austria Micro Systems, www.amsint.com

5.2.1 ADC Integration

Here we discuss the arguments to integrate the ADC on either the PASA or the digital chip.

Advantages of merging ADC with digital tracklet processors:

- clear separation of the purely analog and mixed signal parts; no potential coupling of the ADCs digital state machines noise into the sensitive preamplifier front-end
- use of available space as the digital chip alone would be pad-limited
- separation of the preamplifier design cycles from the ADC selection process
- lowest digital/analog interconnection pin-count (one differential pair per channel instead of 10 signals per channel)

Disadvantages of merging ADC with digital tracklet processors:

- coupling of the digital design to the process chosen for the ADC, making retargeting of the digital *tracklet* processor difficult
- 21 ADCs required for 18 channels ($\lesssim 2 \text{ mm}^2$ per ADC channel)
- some analog outputs have to drive two ADC inputs, and thus require two individual output stages with corresponding matching problems
- functional chip testing requires some additional logic on the analog front-end

Implementing a 21-channel ADC on a third chip on the MCM is not desirable as this chip would be either pad-limited or the readout would have to be multiplexed, resulting in higher (2x or 4x of digitization rate) clock rates on the ADC die. In any case, the additional number of wire bonds per MCM (336 if ADC readout is not multiplexed) would increase the cost. However, this issue will be revisited when the final size of the ADCs and digital circuitry, and thus the yield of the resulting chip, is determined.

5.2.2 ADC technology choices

The choice of ADC silicon technology is critical as it also drives the choice of the process to be used for the digital back-end. The majority of digital design is based on standard cells, providing for easy retargeting to another process, particularly if it implements a smaller feature size, and thus is inherently faster while using less power. However, there are a few special components required, such as LVDS I/O, PLLs for high-speed clock generation, temperature sensors, etc. These components are likely to be available for modern processes or will be easy to procure. However, there are also multiple instances of a quad-port memory that are required, which are implemented as full custom design and therefore have to be retargeted as well. Although the required clock rates are comparably slow, this retargeting is basically a redesign of the quad-port memory as the optimization parameters, area, speed and size, that drive a certain memory architecture depend on the available number of metal layers, via stacking, minimum spacing, size of contacts and vias, etc., typically change enough. For example, the AMS 0.35 μm process supports three metal layers while all deep sub-micron processes support a minimum of five metal layers. On the other hand, it should be noted that a first prototype quad-port memory, which was taped out in June 2001 in the AMS 0.35 μm process, already supports access times of about 3 ns, which are much faster than required (refer also to Chapter 6, Section 6.3.3.3).

The ADC chosen for the TPC (ST TSA1001) is adequate for the TRD as well (for ADC requirements, refer to Chapter 5). This ADC is a commercial product and available as intellectual property core. As the preamplifier output stage is designed to deliver both a 1 V differential voltage swing and the capability

to drive high capacitive loads, the particular choice of ADC is rather independent of the design of the preamplifier. The final choice of ADC depends on a variety of parameters, such as price of both chip real estate and licensing, power consumption, latency, the long-term availability of its silicon technology (which is relevant as the TRD production schedule is different from that of TPC, particularly when taking into account that one scenario is a staged production), and many other similar issues. Currently, several options are evaluated. However, in order to have a credible architecture, the ST TSA1001 ADC was chosen as baseline. This ADC is implemented in the ST HCMOS7 0.25 μm process.

5.3 Preamplifier / Shaper

The preamplifier/shaper (PASA) is the first block of the front-end electronics, receiving the signals from the detector pads.

5.3.1 Requirements

The current signals of the detector pads are first amplified by a charge-sensitive preamplifier. It is followed by a pole-zero cancellation circuit and two second-order shaper-filters, assuring a shaped output pulse with about 120 ns FWHM. The last functional element of the preamplifier/shaper chain is an output amplifier, which delivers differential output signal according to the ADC requirements concerning driving capability and output levels.

The overall gain of the preamplifier/shaper is 6.1 mV/fC and the shaping type is CR – RC⁴. The differential outputs of the preamplifier/shaper drive a 10 Bit differential 1 V range ADC input.

The functional block diagram of the preamplifier/shaper is shown in Figure 5.3.




Figure 5.3: Block diagram of the preamplifier/shaper.

From the point of view of the implementation, 18 channels of preamplifier/shaper will be integrated on one chip with a core area of about 7.7 mm².

The main requirements of PASA for the TRD front-end electronics and readout are given in Table 5.2.

5.3.2 Implementation

The final implementation takes into account the experience achieved from previously developed preamplifiers built with discrete components and from the first version of the preamplifier/shaper chip. Important input was also derived from a design review of the preamplifier/shaper circuit which took place at CERN on January 24-25, 2001.

The preamplifier is built around a NMOS input transistor folded cascode circuit. The NMOS input transistor allows achievement of a greater transconductance parameter than a PMOS input transistor and also enables a design with a single power supply. A greater transconductance leads to lower input impedance of the preamplifier and consequently to lower crosstalk. Also, it enables the main gain to be distributed towards the front of the preamplifier/shaper chain (preamplifier and pole-zero circuit). For the given short shaping time, the advantage of a PMOS input transistor concerning $1/f$ noise is not important.

Table 5.2: Preamplifier/shaper requirements.

Parameter	Value
Gain	6.1 mV/fC
Shaping time (FWHM)	~ 120 ns
Shaping type	CR – RC ⁴
Max. equivalent input noise (on the bench)	500 e
Max. equivalent input noise (in system)	1000 e
Input dynamic range	164 fC
Output pulse level	1 V differential
Max. internal chip crosstalk	0.3%
Max. power consumption/channel	10 mW

The preamplifier is followed by a pole-zero cancellation circuit and two second-order filters. The addition of two more poles, relative to the first version of the chip, translates into a more symmetrical response at the output of the preamplifier/shaper.

The output amplifier, as a differential-output type, drives a 10 Bit ADC. The differential output structure is less sensitive to perturbations.

The simulated main outputs of the preamplifier/shaper chain are shown in Figure 5.4.

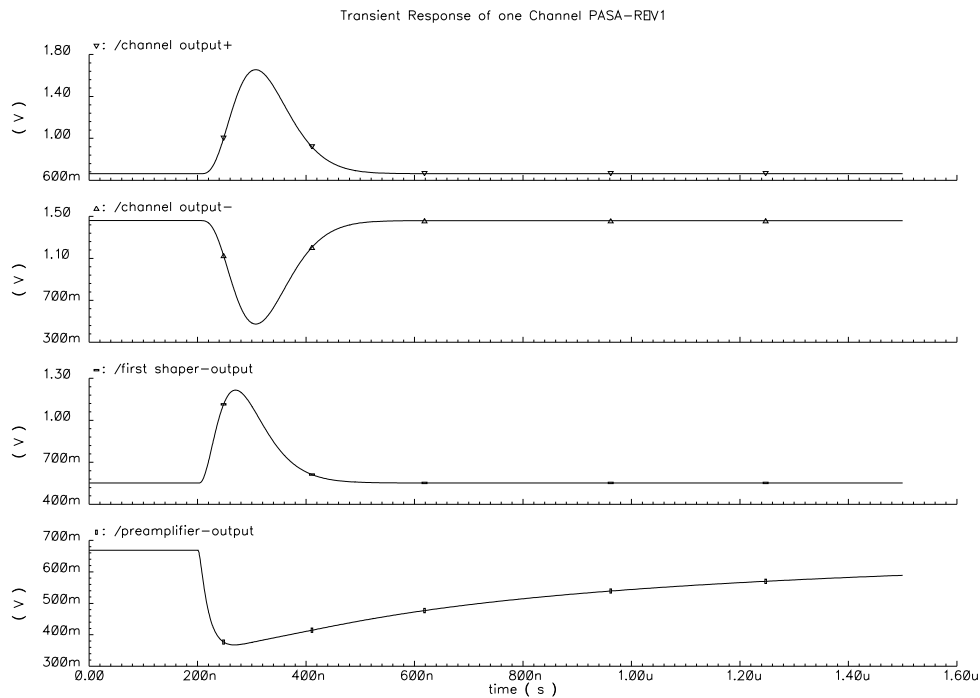


Figure 5.4: Simulated preamplifier-output, first shaper-output and channel output+/output- signal of the preamplifier/shaper. They correspond to the block diagram in Figure 5.3. As stimulus an equivalent input charge of 165 fC is used.

Consideration concerning input protection

The classical protection circuit of the chip I/O pads avoids effects of electrical over-stress (EOS). There are three types of electrical over-stress [2]:

- electrostatic discharge (ESD)
- electromigration: slow wear-out mechanism caused by high current densities
- antenna effect: charge accumulation on gate electrodes during etching or ion implantation

From all of these, ESD protection must correspond to a human body model (HBM) and machine model (MM). For the TRD PASA, additional stress may come from abnormal detector signals (like sparks). For a normal protection to ESD, verified I/O standard pads from AMS were used. For negative input surges, two diodes are tied to ground. One diode tied to analog supply protects against positive input surges. To limit the peak transient currents and electromigration, a resistor of 10.6 Ω is added in series with the pad input. The value of this resistor is limited by noise consideration. For example, for a 25 pF detector capacitance, a 10.6 Ω resistor increases the overall noise by 8%. For the next version of the chip, an additional array of resistors to limit and dissipate the positive input surges will be added.

Some considerations concerning latch-up protection:

Mixed PMOS-NMOS transistor structures are present in many parts of the PASA circuits. For example, in a simple CMOS inverter, parasitic structures of both transistors form an inactive PNP sandwich, having inverse polarized junctions. Due to a parasitic current into the substrate or to a parasitic electrostatic coupling, the PNP structure can accidentally become conductive from VDDA to GND, like a thyristor. The thyristor may be latched up and the whole chip may be destroyed due to high currents.

To avoid latch-up, two classical methods are used [4], [5]:

- Electrostatic protective structure for I/O pads, which allows low resistance paths for accidental currents like transient-type currents;
- Diffusion-type low resistance rings around MOS transistors.

For the TRD PASA, the latch-up is prevented by:

- The use of dedicated, verified I/O standard pads from AMS
- Guard rings for each MOS transistor
- A separate guard for each analog channel
- A complex guard of the type 'P diffusion -n well -P diffusion' is placed in between channels
- For each channel, the different functional blocks are separated by guard circuits

5.3.3 Prototypes

The first three models of preamplifier/shaper were built at GSI-Darmstadt with discrete components. Also, all were tested in beam with detector prototypes.

As the main component, the first one had the current feedback-type MAXIM² MAX4182 operational amplifier. It was used to design an eight channel preamplifier module. Having the capability to change the input impedance, it was also useful in determining the optimum input impedance for a good signal/noise ratio and crosstalk. The main specifications for 1600 Ω input impedance are: gain 0.7 mV/fC, noise about 11000 e, and crosstalk between adjacent channels 10%.

The second preamplifier/shaper was also built around the MAX4182 operational amplifier, but in a current-type configuration. Having a low input impedance of about 160 Ω , it exhibited low crosstalk for adjacent channels (only 2%). It had a gain of 1.3 mV/fC, and noise about 7000 e.

²Dallas Semiconductor MAXIM, www.maxim-ic.com

The third preamplifier/shaper was a charge-sensitive type. With a gain of 2 mV/fC, CR – RC shaping and a noise of only 1500 e and a crosstalk of 8% between adjacent channels. It was used for most of the measurements with detector prototypes (see Chapter 14).

The first chip, with 21 channels, was submitted at the end of October 2000. The 21 analog channels are basically identical. The only difference is in the value of the input pad resistance. There are eight channels with 0 Ω , eight channels with 50 Ω , two channels with 200 Ω , and two channels with 500 Ω , to estimate the influence of the pad input resistance to the overall noise. One additional channel with 50 Ω input impedance allows monitoring of the signals at each stage of the preamplifier/shaper.

Each of the 21 channels is implemented as a charge sensitive amplifier. The main specifications are: gain about 5.2 mV/fC; output pulse FWHM = 125 ns; shaper type CR – RC²; input dynamic range 0 to 330 fC for 2 V output signal; and noise about 1500 e. It was part of a multi-project run, together with the TPC preamplifier and a digital multi-port memory

The second prototype of the PASA chip has the characteristics presented in Table 5.2 and was submitted to AMS in June 2001. The photo of the layout of this second version (18 channels) is shown in Color Fig. 5. The evaluation of its performance is underway.

5.4 ADC

The requirements of the ADC for the TRD are summarized in Table 5.3. It should be noted that the whole system, including each MCM, will be actively cooled in order to guarantee enough temperature stability. Therefore, no particular requirements are presented with respect to temperature stability.

Table 5.3: ADC requirements

Parameter	Value
Resolution	10 Bit
Digitization rate	10 MHz
Max Power consumption	20 mW
Input	2 V differential (+/- 1 V)
Input bandwidth	5 MHz
Max. differential non-linearity	0.7 LSB for channels [0, 511] 1.5 LSB for channels [512, 1023]
Max. integral non-linearity	1.0 LSB for channels [0, 511] 2.0 LSB for channels [512, 1023]
Effective Number of Bits	> 9 Bit
Max. latency	5.5 clocks
Min. input impedance	100 k Ω
Max. input capacity	7 pF
Max. area	2 mm ²
Max. channel to channel variations on same die	0.5 %

The ADC cores will be operated with individual power and the digital chips floor plan arranged such that the ADCs are geometrically isolated from the digital back-end. One of the ADCs will be used for detector control and readout independently from the data readout channels.

5.5 TRD trigger states

The TRD trigger operates in different states corresponding to the different tasks it performs. An overview of these states, together with the associated external stimuli, is sketched in Figure 5.5. The TRD default state is in stand-by with all digital clocks switched off. A pretrigger starts the archival of the ADC's raw data and the *tracklet* preprocessor (TPP) and computes the appropriate sums (see Chapter 6). The ALICE Trigger system issues a Level-0 (L0) trigger at a fixed latency (about 900 ns) after the interaction. This L0 trigger is the first confirmation of the TRD pretrigger. Should the central trigger processor (CTP) have decided not to issue a trigger, the missing L0 trigger (which constitutes a L0 reject for the TRD as it starts early) will lead to the TRD being cleared, aborting the trigger sequence as indicated in Figure 5.5.

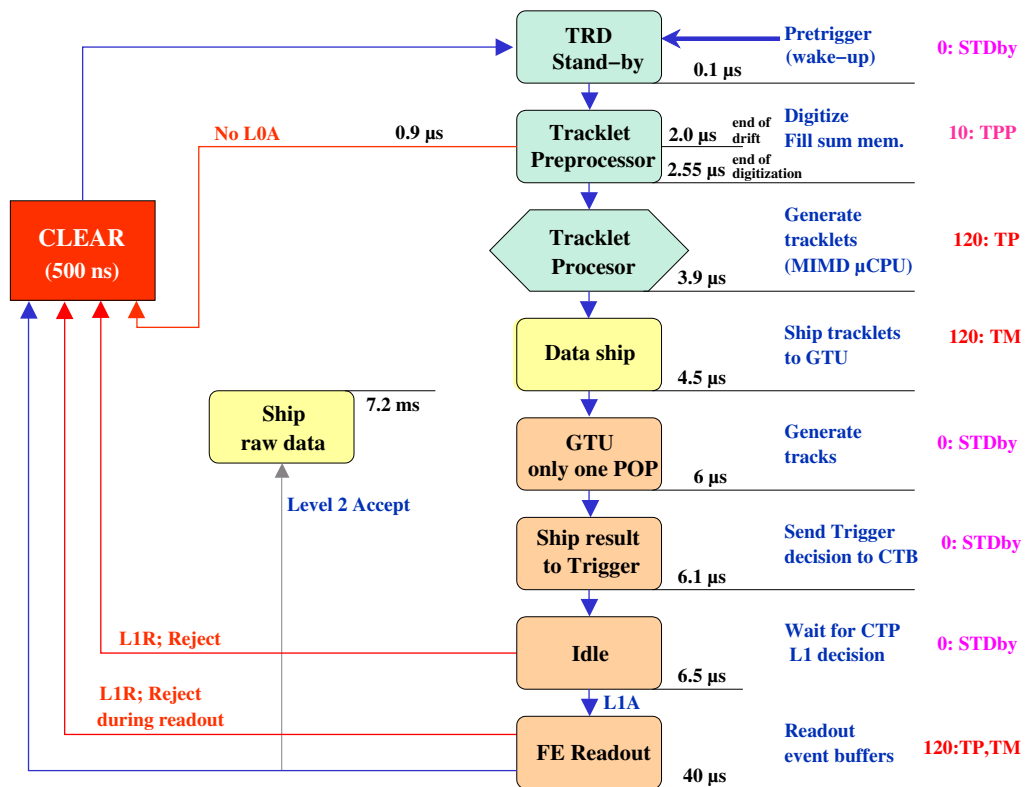


Figure 5.5: The various TRD trigger states from pretrigger to readout. Note that data shipping through the DDL is done concurrently and independent of the TRD front-end electronics. The various functions (TPP, TP, TM) are labeled together with their associated operating clock frequencies in MHz.

As it is described in Chapter 6, Section 6.3 at the end of the drift time and when the preprocessor has finished its task, the MIMD processor calculates the *tracklet* parameters and applies the configured selection cuts. After identification of the *tracklets* they are transposed into the TRD reference plane and formatted for shipping to the GTU, which is completed $3.9 \mu\text{s}$ after the interaction. Data shipping concludes at the $4.5 \mu\text{s}$ mark, assuming a maximum of 40 *tracklets* per chamber. Excess *tracklets* are ignored. The readout is performed in an ordered way, such that the global tracking unit can already start processing the first *tracklets* once they have arrived (see Chapter 7). The result of the GTU processing is a potential trigger and a 36 Bit vector, which defines the regions of interest for readout. This information is shipped to the CTP at the $6 \mu\text{s}$ mark.

After delivering the trigger to the CTP, the TRD trigger awaits the response as Level-1 (L1) accept or reject. Note the ALICE CTP does not implement specific L1 accept and reject signals, but delivers a L1 trigger at a defined time slot after the interaction (about $6.4 \mu\text{s}$ mark), like in the case of L0 triggers.

However, for improved legibility and less redundancy, a missing L1 trigger will hereinafter be referred to as L1 reject (L1R) and a L1 trigger at the appropriate time slot will be referred to as L1 accept (L1A). The TRD electronics operates in stand-by mode, with all fast clocks disabled to avoid excess noise, while waiting for the CTP L1 trigger decision. A L1 reject will abort the pending event, placing the system back into stand-by mode. However, a L1 accept will trigger the readout of the event buffers through the same data path that was used to ship the *tracklet* candidates to the GTU. The GTU implements appropriate readout buffers to absorb the 216 GB/sec data stream. Should the activation of the fast readout clocks generate any noise problems, for example within the TPC, the L1 accept signal can purposely be delayed transparently within the TRD to the trigger system.

The completion of the front-end readout leads to clearing the TRD electronics and putting them back into stand-by without further outside interaction. The given event resides now in an appropriate event buffer, which is implemented as part of the GTU. A Level-2 accept (L2A) will schedule the event for transmission off the detector. A L2 reject (L2R) will free the appropriate buffer space. The data transfer functionality is independent of the TRD state sequence (refer to Chapter 7).

It should be noted that the TRD trigger electronics is not pipelined. Once enabled, it cannot process any other event until it is cleared, which, in the case of a L1 accept, can be as late as $40 \mu\text{s}$ after the interaction. For details of the TRD readout, refer to Section 7.1.2. However, assuming a 200 Hz accept rate, which is the maximum TPC Pb-Pb gate opening rate, the corresponding dead time is 0.8 %. The handling of the corresponding TRD busy is discussed in Section 5.7. For a detailed discussion of the timing relationship between the various trigger states, refer to Section 5.6.

Some of the activities do not depend upon each other and are executed in parallel. For example, as soon as the first data words arrive at the global tracking unit (GTU), they are processed rather than waiting for the complete delivery of all *tracklets* from the front-end. Further, the data shipping to the high level trigger or event builder system is done in parallel upon a L2A while the TRD front-end may already be operating in stand-by, thus increasing its lifetime.

5.6 Trigger timing

For Pb-Pb running, the TPC trigger rate is limited to about 200 Hz. In order to inspect a larger number of events, the TRD has to derive its decision prior to the TPC gate opening. On the other hand, the TPC drift begins with the interaction. Therefore, any trigger latency effectively reduces the active volume of the TPC. Given a drift time of $80 \mu\text{s}$, an overall TPC pretrigger latency of $6.5 \mu\text{s}$ corresponding to 8 % of the drift time is defined as an acceptable baseline.

Figure 5.6 outlines the resulting system timing. A very fast minimum bias TRD pretrigger, which is gated with the TRD BUSY, is used to wake up the TRD electronics. This pretrigger bypasses the ALICE CTP and is expected 100 ns after the interaction at the TRD point of presence (POP), from where it is fanned out to all the individual detector modules (see also Sections 5.6.1 and 5.6.2). Given the large surface area of the TRD detector, the definition of such a reference point (POP) is required in order to allow unambiguous definition of the required timing relationships. The pretrigger is also copied as L0 input into the CTP. The distribution of the TRD pretrigger to the various MCMs requires another 200 ns, corresponding to a total of 10% of the TRD drift time for pretrigger distribution as indicated in Figure 5.6. However, the first 250 ns are not crucial to be read out, as they contain the ionization of the primary track from the amplification region.

Low-power ADCs typically implement an internal pipeline. The particular device chosen as baseline implements a 5.5 clock pipeline, effectively storing 5.5 analog samples in a kind of analog memory, thus implementing an equivalent pretrigger history of 550 ns at a digitization rate of 10 MHz, which corresponds to 1/4 of the LHC clock. This is indicated as visible drift region in Figure 5.6. However, the same latency has to be added at the end of the drift time in order to drain the ADC's pipeline.

As outlined in Chapter 6, the computation of the sums required for the linear fit can already be performed during the drift time. For each time bin, every ADC channel is checked to match the criteria

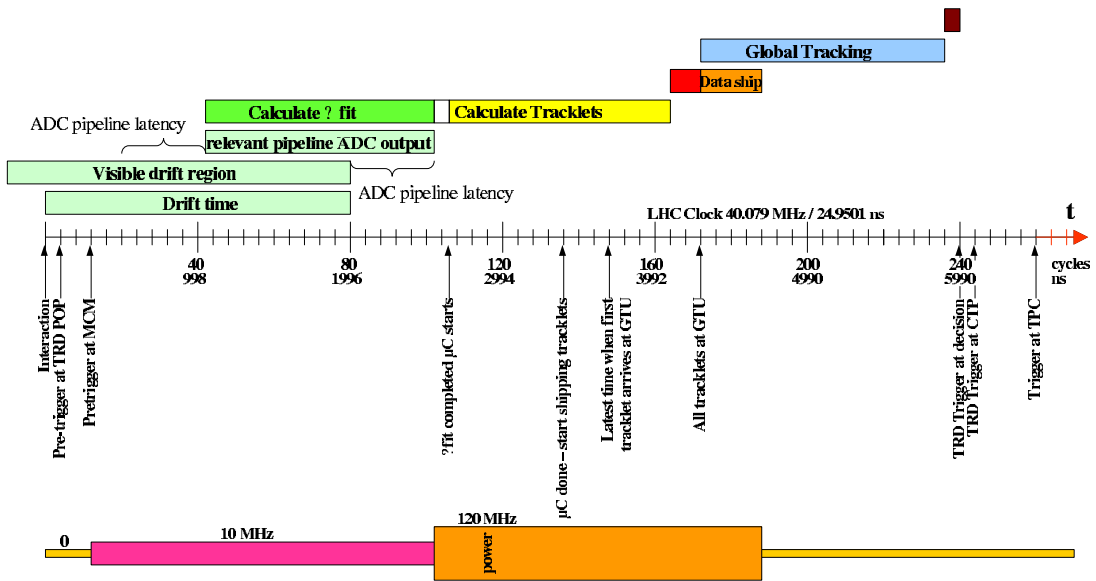


Figure 5.6: TRD Timing. The time axis is calibrated in units of LHC clocks, where each tick corresponds to four LHC clocks or about 100 ns.

for a cluster centroid. In such case, the appropriately derived entries for the sums are calculated and stored in the channels sum-memories using read-modify-write cycles. In order to influence the ADC performance as little as possible with the associated digital noise, the appropriate logic is run at the ADC digitization speed. One more ADC clock cycle is required at the end of the digitization period by the preprocessor in order to provide it's results. To shorten the latency, the digital clocks are switched to full speed operating mode, amounting to 120 MHz, at the end of the digitization. This reduces the preprocessor pipeline latency to 67 ns. Note that at this point, (about 2.55 μ s after the interaction) any digital noise produced will affect neither TRD nor TPC, as the TRD data is already stored in its event buffers and the TPC has not yet started digitizing.

The preprocessor pipeline is fully drained at the 2.55 μ s mark, at which the embedded MIMD micro-processor starts analyzing the various *tracklet* candidates. Up to four *tracklet* candidates are assigned automatically, one each to a processor thread, at the conclusion of the preprocessor task. Therefore, the *tracklet* processing time is independent of the number of *tracklets*. The available time for *tracklet* processing and selection is 1.5 μ s.

Each identified *tracklet* is forwarded to the global tracking unit using the high-speed TRD readout tree. Since the readout of each chamber is ordered, the global processing of the first regions of a chamber can happen while other parts are still being read out. However, there is a minimum readout latency, which corresponds to the worst case readout time of the first *tracklet* candidates. No pipelined processing can be done during this time, amounting to 200 ns as indicated in Figure 5.6. The remainder of the data shipping, which corresponds to a maximum of 40 *tracklets* per chamber, is overlapped with the global tracking of the GTU.

The first and last *tracklet* arrive at the GTU at latest 4.3 μ s and 4.7 μ s after the interaction, respectively, allowing 1.3 μ s for the global tracking functionality. It should be pointed out that this functionality will be implemented in FPGAs, running at a target clock rate of 40 MHz.

The TRD trigger decision has to be determined 6 μ s after the interaction, allowing a total of 500 ns for shipping it to the CTP and back to all involved detectors in case of a L1 accept.

5.6.1 Clock distribution and clock domains

In order to reduce clock noise, all TRD clocks are derived from and synchronized to the LHC clock. It is distributed using the RD48 Trigger, Timing and Control (TTC) system. Each detector implements one TTC receiver module as mezzanine card together with the appropriate slow controls functionality, which fans out the system clock to about 200 clock nodes per chamber using the IEEE 1596.5 LVDS standard [1]. The signal fan-out is implemented such that the individual skew between the various clock nodes is minimized. Clock rates higher than the LHC clock are generated on the detectors using PLLs.

The readout tree runs at 120 Mwords, using a 120 MHz clock. The digital processor also operates at 120 MHz, while the readout and *tracklet* preprocessor operates at 10 MHz or 1/4 of the LHC clock, using the same clock as the ADC.

In order to keep the digital electronics as quiet as possible and to save power, all digital clocks are gated and switched off when the TRD is idle. The only exception is the differential clock fan-out and the PLLs, which cannot be started quickly. The ADCs and digital filter for tail cancellation, however, are kept running during stand-by like the preamplifier circuits as they cannot be enabled quickly enough.

After a pretrigger, the 10 MHz clock to the *tracklet* preprocessor is enabled for the duration of the drift time. After $2.55 \mu\text{s}$, the fast 120 MHz clocks are enabled, starting the multiprocessor and the *tracklet* readout to the GTU. Upon completion of the *tracklet* readout, the various chips fall back into stand-by operation. They are re-enabled by either a L1 accept or a L1 reject, performing the necessary cleanup in order to get ready for the next pretrigger. The clock usage and corresponding power consumption is indicated in Figure 5.6. In this context it should be noted that the lifetime of the gated digital circuitry is as short as a few microseconds per activation. The required energy for this activity will be stored in filter capacitors, such that there will not be large currents switching on the power distribution lines if the TRD is activated.

5.6.2 Distribution of fast signals

Each TRD MCM requires the following fast logical input signals:

- Synchronized clock reference
- Pretrigger at TRD point of presence within 100 ns after the interaction
- L0 accept/reject at configured fixed LHC clock
- L1 accept/reject at configured fixed LHC clock (only after L0 accept)
- L2 accept/reject at undetermined time in chronological order for each L1 accept

All these signals are defined with respect to the LHC clock. In order to guarantee the correct phase, all LHC clock-related signals are routed together with the clocks for the given device.

The system default state is stand-by, operating at minimum power. The first TRD trigger is the pretrigger starting the system. The front-end chips will continue processing according to the time line as sketched in Figure 5.6 until a trigger decision is delivered to the CTP. During this process, the logic can be aborted any time, which is done by asserting the TRDTrigger signal for at least two consecutive clocks.

The pretrigger is the most time-critical signal. Table 5.4 shows a breakdown of all latencies involved in transmitting the pretrigger signal to all MCMs. Negative latencies are defined as signals arriving early. The most efficient way to avoid additional trigger cabling is utilizing the TTC system, which is going to be used for the distribution of the clock signals. The TTCvi module will forward minimum bias pretriggers as TTC L0 accepts on its A channel only if the TRD electronics are operating in stand-by. It should be noted that the pipeline latency of the ADC chosen as baseline allows for much larger pretrigger latencies. However, the ADCs digitization latency is technology dependent and can be as little as one

clock. Therefore, in order to allow for other ADC technologies (refer to Section 5.4), this requirement is not relaxed. On the other hand, in the event the final ADC does not implement a pipeline and the pretrigger results in an unavoidable larger latency as specified here, an appropriate digital pipeline can be implemented, which would have the advantage of not adding latency at the end of the drift time as the pipeline ADC does. Such a digital pipeline, however, comes at the price of more digital circuitry continuously being operated.

Table 5.4: A breakdown of all latencies involved in distributing the TRD pretrigger signals. The same parameters are applicable to the other fast input signals.

What (source)	t/ns	[t _{min} , t _{max}]/ns	clocks
Interaction to TRD point of presence input (ALICE requirement)	+100	[100, 150]	4...6
TTC system latency (RD12 measurement)	+68	[65, 100]	3...4
Signal propagation including fan-out (20 m, each TTC fan-out counts for 1 m)	+100	[50, 150]	4
Clock/Trigger signals fan-out on detector (2 stages estimated)	+15	[10, 25]	
Signal propagation on detector (3 meter estimated, periodic signals can be adjusted to compensate latency - trigger signals cannot; all signals are relative to reference clock)	+15	[10, 20]	
Sum total	+298	[235, 445]	12
Pipeline ADC @ 10 MHz (5.5 clocks)	-550	[0, 600]	25
Ignore beginning of drift time	-250	[200, 300]	8...12
Total TRD Pretrigger latency (negative means early)	-502		

The TRD Trigger system is not pipelined and is therefore BUSY starting with the pretrigger until the readout of the front-end buffers completes or the event is aborted. This allows for the use of the same trigger input (hereinafter referred to as TRDTrigger) for different functions depending on the state of the TRD trigger. Different inputs at one state can be encoded in pulse length as multiple back-to-back triggers are not possible. The signal TRDTrigger is fanned out to all front-end systems as TTC L0A trigger on the TTC A channel. For example a pulse of the TRDTrigger during the TRD idle state is considered a pretrigger, a TRDTrigger pulse at the 900 ns mark (L0A time slot) is a L0 trigger. Longer TRDTrigger bursts can be used to encode other functionality, such as clears.

The fixed latency of the ALICE Trigger system's L0 trigger allows implementation of a L0 reject as a missing L0 trigger. This condition is detected at the TTCvi root module. In the case of a missing L0 trigger, two consecutive L0A triggers are transmitted through the TTC system to all chambers. The logic required to generate the this pulse length clear code of the TTC L0A signal is required only once for the entire detector.

After delivering the TRD trigger decision, the system enters a wait state (idle state) while awaiting receipt of the CTP's L1 decision as another TRDTrigger pulse, which is now interpreted as a L1 accept. This L1 accept allows to start the readout of the TRD front-end buffers at any time after the TRD entered this idle state. Thus this readout can purposely be delayed past the TPC drift time in case of TPC coincident triggers in order to keep the TRD electronics chain quiet during the TPC drift, should this become a noise problem. However, such functionality would be implemented at the TTCvi root module like the L0 clear functionality, which is entirely transparent to the CTP. Note that the CTP's decision can be completely independent of the TRD's trigger suggestion, resulting in TRD L1 accepts after a TRD reject and vice versa. However, no L1 accept is expected by the TRD without having received an appropriate pretrigger.

The L1 accept results in the readout of the front-end system, and thus allows release of the TRD BUSY as soon as this function completes. The system returns to the stand-by state while awaiting the next pretrigger.

The data readout through the detector links is triggered upon a L2A. The L2A and reject signals are not required at the detector front-end and are shipped to the appropriate functional units of the global tracking unit. The Level-2 (L2) decisions are in isochronous order, thus simple accept/reject signal pairs are adequate. A L2R simply frees the appropriate event buffer space within the global tracking unit.

The architecture outlined above allows the implementation of the local tracking units (LTU) as simple state machines that operate after a pretrigger up until they receive a clear. Only two fast signals (TRDTrigger and clock) are required to be distributed to all MCMs. The LTUs implement the additional feature handling the assertion of the TRDTrigger signal for two or more contiguous LHC clocks as clear. No specific clear signal is required. This scenario operates the TTC in a simplified mode, using the L0A channel for all synchronous triggers. However, given the short latency budget for the pretrigger, the TTCvi root module would have to be located close to the TTC point of presence. Should this turn out to be a problem, then the pretrigger has to be distributed individually. The rest of the signal coding would remain unchanged. The implementation of this coding can be done in a simple programmable logic device (PLD) as part of the TRD trigger logic.

5.7 Interface to the ALICE trigger system

The TRD requires a fast pretrigger as a wake-up signal. The sole purpose of this signal is to allow the operation of the digital components within the system in low-power mode while the system is in stand-by. The timing requirements for this signal are discussed in Section 5.6. The pretrigger has to be issued before the ALICE CTP has issued a L0 trigger. It is implemented as a minimum-bias trigger. Further, in order to have clean events within the TRD, particularly for Pb-Pb running, the TRD requires the pretrigger to be pre-history and pile-up protected. Future protection is implemented by rejecting appropriate pile-up events at L1 time. All TRD related triggers have to be counted before and after dead time by the CTP in order to allow proper calibration.

Given those requirements, the integration of the TRD is more complex than a canonical, stateless, dead-time free trigger detector or a generic detector, which is triggered by L0A or L1A, such as the TPC. Figure 5.7 sketches the architecture. The critical path timing of the pretrigger is designated by the thick line (the signal). The pretrigger is issued by a fast minimum-bias trigger detector, which is routed directly to the TRD using the shortest possible path in order to minimize its latency. A second independent copy of the signal, which is less time critical, is routed to the CTP. In order to avoid unnecessarily waking-up the TRD electronics, the pretrigger is to be issued only in case of a clean history. This functionality has to be implemented by the TRD system as it is in the critical path of the pretrigger and the time to route signals to and from the CTP would far exceed the maximum allowable latency. The clean TRD minimum bias signal can be recreated by the CTP. In general, past protection is easily implemented by using a retriggerable one-shot, which is triggered with the minimum-bias trigger and which has a decay time corresponding to the TRD drift time. The resulting pretrigger signal is relevant only in case of the TRD being idle, which is sketched in 5.7, by gating the clean minimum bias signal with the TRD BUSY status. The TRD BUSY itself is started by each valid pretrigger and cleared either by rejecting the event or after the L1A related readout has completed. The valid TRD pretrigger wakes up the TRD digital electronics and starts the TRD state machine as sketched in Figure 5.5. This signal is forwarded to the CTP, where it is treated as regular L0 trigger input. All trigger classes, including the TRD, require this signal to be present.

A TRD pretrigger may or may not result in an appropriate TRD L0 trigger. Not receiving a L0 trigger at the specified time slot will be handled as an abort. Should the TRD receive a L0 trigger at any other time, an error is flagged and the trigger is ignored. Such a scenario would most likely be caused by L0 trigger classes involving the TRD, but without requiring the TRD pretrigger as input.

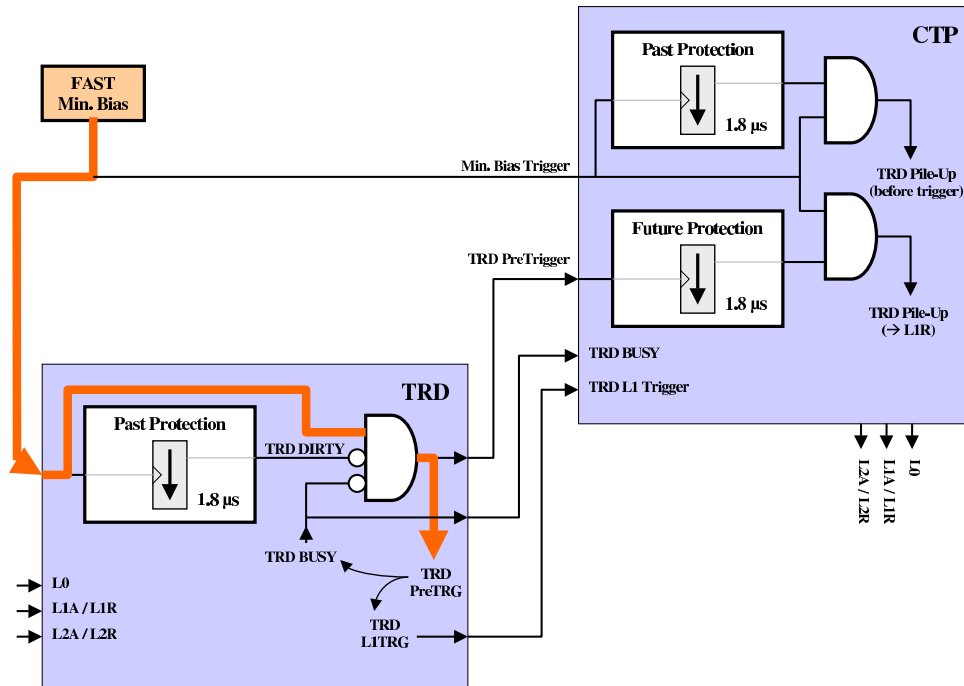


Figure 5.7: TRD Pretrigger architecture. Note: the TRD past/future protection is implemented as programmable counters, and such can be configured within a range of 1...100 μs . The past protection located on the TRD detector is logically part of the CTP and configured by the CTP in order to guarantee coherent configuration. In case of coincident running with another detector requiring a larger past protection, such as the TPC, the TRD past protection will be adjusted accordingly.

After receipt of the L0, the TRD trigger will proceed to determine its trigger decision, which it forwards to the CTP at the 6 μs mark, and which may result in either a L1 accept or reject, depending on the trigger class or classes.

The TRD can be aborted at both L0 and L1 time. It should be explicitly noted that the TRD will abort only upon an appropriate CTP decision and never by itself. Any pile-up related aborts have to be issued by the CTP.

Past and future protection is standard circuitry, which is implemented by the CTP. Another detector requiring such logic is the TPC. The only variation is the different time constant of 2 μs instead of 80 μs in case of the TPC. The fact that the past protection circuitry is mirrored by the TRD should be considered as an implementation detail solely driven by the pretrigger being in the critical path. The appropriate logic within the TRD front-end will be connected to the Trigger DCS in order to ensure coherent configuration. However, only a very small number of parameters is concerned here, which do not change often. Architecturally, the past and future protection logic has to be part of the CTP in order to allow for proper cross-section calibration. All future protection is implemented by the CTP and results in rejecting pile-up events at L1 time. However, in order to reduce the overall TRD dead time and power consumption, TRD pile-up should also be used as qualifying input to the appropriate L0 trigger. This results in all pile-up events happening during the first half of the drift time being rejected before the high-power digital circuitry is even enabled.

The TRD BUSY signal is not required for normal trigger operation as it is already included in the TRD pretrigger. This is also driven by the long roundtrip delays to and from the CTP. However, in order to allow for proper counting of before/after dead time, this signal is sent to the CTP.

It should be noted that the TRD BUSY must not be qualified for trigger selection as this would always prevent TRD L0 triggers due to the nature of the TRD starting early with the pretrigger and, thus, also asserting its BUSY early.

In summary, the TRD will deliver to the trigger system its dead-time state (BUSY) plus the TRD trigger bits, consisting of one bit for each trigger type plus the region of interest bit mask selecting sector and hemisphere (2x36 Bit). All other trigger signals received from the CTP (L0A, L1A, L1R, L2A, L2R) are received centrally at the TRD point of presence and distributed appropriately within the TRD.

5.8 Multi-Chip Module (MCM) overview

The number of channels per MCM is driven by various parameters, resulting in the choice of 18 channels per MCM. On one hand, the *tracklet* preprocessor architecture requires only processing of neighboring pads of one pad row. The preamplifier inputs are the direct pad signals. In order to minimize the pad capacity, signal crosstalk and pad-to-pad variations, the maximum length of any given pad trace is limited to about 100 mm. Further, the number of channels per MCM or local tracking unit (LTU) and preamplifier chip (PASA), respectively, drives up the chip size, and thus drives down the yield. However, this additional cost is offset by the production cost of the MCM itself, which does not scale much with the number of channels, as one channel adds only one analog input and a few bonding wires because most of the additional circuitry is consolidated into the LTU on the MCM. The resulting optimum is 18 channels per MCM and 8 MCMs per pad row.

A number of scenarios was iterated with respect to the architecture of the MCM. The original approach of rather large readout boards required the MCMs to be mounted using mezzanine connectors. Therefore, the only components on the motherboards would have been such connectors, simplifying the production. However, even in that scenario, the cost for just the connectors was rather high.

The baseline scenario (refer to Chapter 4) implements small enough readout boards, so that they could be mass-produced using standard production facilities. This architectural choice enabled the implementation of the MCM as Ball Grid Array (BGA), which can be produced and soldered without requiring expensive mezzanine connectors. The disadvantage is the increased complexity for replacement of a given MCM. However, taking into account the effort required to remove a chamber for repair, the additional BGA soldering to replace an MCM becomes a minor issue. On the other hand, the number of I/O pins per module is now a small contribution to the overall cost. The choice to use soldered BGAs as opposed to MCM mezzanine cards mounted via connectors resembles a trade-off between overall cost, material budget and maintainability.

5.9 MCM prototypes and performance

5.9.1 Prototype Motherboard

The first digital chip that was designed is a prototype of the *tracklet* preprocessor (TPP) in the AMS 0.35 μm process, which also implements appropriate readout circuitry. One of the goals for this chip was to better understand the noise introduced by the close proximity of fast digital clocks and sensitive analog preamplifiers. In order to test the *tracklet* preprocessor prototype, together with the well-understood existing discrete preamplifier, an appropriate motherboard was designed, hosting both the *tracklet* preprocessor and the digital readout chip. Figure 5.8 shows the device. It hosts eight ADCs, the digital chip in the ceramic package, and some additional glue logic for generation of miscellaneous signals, such as clocks.

5.9.2 Prototype MCM

The block structure of the MCM reflects the already discussed connection diagram of the *tracklet* preprocessor prototype 1 with eight channels: Preamplifier Chip with 8 analog inputs/outputs, 10 ADCs (including two neighbouring channels, 8 Bit NSC³ ADC08351), the *tracklet* preprocessor itself, and the

³National Semiconductor, www.national.com




Figure 5.8: Prototype Motherboard.

connectors. There are two possibilities to deliver the data of neighbouring analog channels of MCMs: analog or digital. In the first case, there are two extra ADCs on the MCM that digitize the analog signal, coming from the neighboring MCMs. In the second case, the digital outputs of the boundary ADCs from adjacent MCMs are fed in parallel to the MCM. The *tracklet* preprocessor prototype 1 has 5×8 Bit inputs for ADC data; each of the two ADCs are combined together with a common readout bus and multiplexed in time. This is possible, as the ADC sampling rate is 10 or 20 MHz, while the *tracklet* preprocessor works at the $4 \times$ speed (40 or 80 MHz). The two ADCs, belonging to the same readout bus, have 18 common pins and only the OE (output enable) and Vin (analog input) lie on different networks. We decided to solder the second ADC directly onto the first one, while the two pins mentioned above are connected to the board via small vertically-mounted SMT 0Ω resistors. This topology saves a lot of space and vias on board. Some technical details: the MCM is a twin-layer board, and with minimal distances/route widths 6 mil/6 mil ($152 \mu\text{m}$), the size of the board is $51 \times 40 \text{ mm}^2$. There are two FPC connectors (30 pins) for inter-MCM communication, one FPC connector (18 pins) for the command bus and one FPC connector (18 pins) for the analog inputs. All FPC connectors are commercial 0.5 mm pitch connectors (eg. HARWIN⁴).

The first MCM was mounted on a small universal board. Color Fig. 7 shows a photo with both the preamplifier and digital back-end chip integrated. The ADCs are implemented as discrete chips with two stacked on top of each other in order to save space. This carrier board contains the voltage regulators (two 3.3 V, one 1.65 V), a quartz oscillator, and normal connectors for easier tests. The digital control/readout was made by an universal PCI I/O board (already used for tests of the *tracklet* preprocessor). Due to difficulties at bonding of the digital chip, some inputs from one pair of ADCs were accidentally shorted to ground and therefore the corresponding two ADCs were not soldered on the MCM. The primary aim of the MCM was not to test the *tracklet* preprocessor as a digital chip, but to test the MCM technology and to estimate the performance when we put in close proximity a high-speed digital chip (TPP), several pipelined ADCs, and a very sensitive analog chip (PASA).

Figure 5.9 shows the digital output of one ADC with reduced reference voltage. The input of the corresponding preamplifier channel was open. If there is a signal applied to an adjacent channel of the preamplifier (so that we have maximal amplitude at the preamplifier output), there appears a disturbance for 1-2 time bins, with amplitude 1-2 LSB of the ADC. If we short the input of the preamplifier to ground and apply the same signal to the adjacent preamplifier input, we do not see any change at the output of the ADC. In this case, however, the first stage of the preamplifier with grounded input is out of DC stability, while the second stage of the preamplifier is still DC stable and delivers normal voltage to the output. In Fig. 5.10 the shaped pulses measured from six ADCs are shown. There is a slight shape variation in one of the channel.

⁴HARWIN Components, www.harwin.com




Figure 5.9: ADC response on MCM with digital clock enabled. For this test, the ADC reference voltage was reduced in order to increase its gain. Only the least significant bit changes.

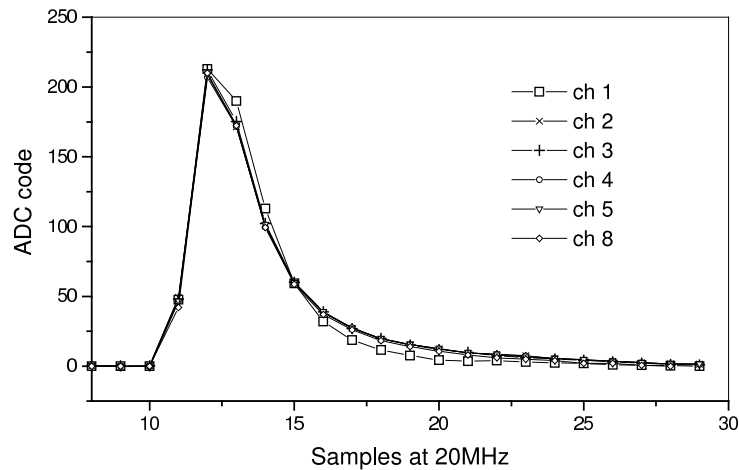


Figure 5.10: Superposition of all analog channels digitized with the ADCs on the MCM and readout by the *tracklet* preprocessor. The ADCs are implemented as discrete chips.

5.10 Design for test

For a high yield in the production of multi-chip modules, it is essential to verify the chips before bonding. There are two chips on the MCM, which will have to be tested independently prior to assembly. However, this testing requires only simple functionality testing as it already will identify most of the broken chips.

5.10.1 Preamplifier

A cheap and fast solution for analog functional testing is the use of a factory standard automatic digital tester. A 4 Bit DAC will be implemented on the preamplifier die together with the appropriate means to inject different charges, defined by the DAC, into the preamplifier front-end. The highest DAC setting would correspond to a PASA output pulse which has an amplitude close to the ADC full scale. This pulse can be easily measured by the digital tester if its readout thresholds are adjusted appropriately. Care must be taken in order to prevent this circuitry from increasing the preamplifier's channel-to-channel crosstalk and input capacitance. A simple internal state machine is programmed through an external single-ended two-wire serial interface, such as Philips⁵ I²C. This state machine implements one enable bit per channel, thus allowing any combination of channels to be activated. The DAC is programmed through the same interface. The clock required for this test controller is held low during normal operation, thus keeping the logic in stand-by and not generating any digital noise. This circuitry allows simple functional verification of the preamplifier chip while being operated on a digital tester and using its digital inputs with an adjustable threshold.

⁵Philips Semiconductors, www.semiconductors.philips.com

5.10.2 Local Tracking Unit (LTU)

All major building blocks of the LTU will be encapsulated by JTAG boundary scan logic (IEEE 1149.1), allowing for the isolation and diagnostics of errors. However, since the LTU implements a multi-processor, in-situ self-testing is implemented, which includes the testing of all internal memories and registers. The event buffers can be uploaded with simulated events. The *tracklet* preprocessor can be configured to process this data, instead of reading the ADC outputs and filling the event buffers. The test routines can be uploaded quickly via the readout tree when configured in upload mode. Given the available four processor kernels, four test instructions can be executed per clock cycle, for example, simultaneously testing four regions of the data memory, thus allowing shortening of the test time. This mode results in the test vectors basically uploading the test program and data, providing the clock, and expecting the test results.

5.10.3 MCM testing and verification

After assembly, the MCMs require testing and burn in. The test infrastructure required shall be as simple as possible in order to allow a large number of modules to be tested. Allowing six months for the testing of all MCMs for the entire detector requires, for example, the completion of one MCM per minute, assuming an eight-hour work day. These tests are expected to be performed periodically in-situ when the detector is idle. The MCM shall be able to complete such testing with a minimum of external additional logic. The list below specifies the MCMs self-test functionality.

- verification of checksum on internal code and data RAM
- memory read/write testing on all internal memories (code, data, event buffer, look-up tables, configuration registers)
- processor configuration, synchronization
- test of core register file
- test of *tracklet* preprocessor by uploading simulated events first into event buffers, then configuring the *tracklet* preprocessor to accept input from the event buffers rather than from the ADCs, and finally by performing a regular trigger and verifying the results in the sum memories
- measurement of supply voltage while switching on fast clocks
- measurement of *tracklet* processor chip core temperature during burn-in and in-situ
- injection of test charge (6 Bit granularity) into any individual or group of preamplifier channels, allowing the measurement of crosstalk and linearity of each individual channel.
- measurement of analog supply voltage during acquisition

Each MCM will carry an unique ID. Test and repair cycles will be archived in a database and stored for the lifetime of the experiment, based on this ID. The MCM test software running on the *tracklet* processor will utilize its readout bus in order to forward status and progress messages to the test environment. This scenario allows a large number of MCMs to be tested simultaneously.