CMS MICROSTRIP TRACKER ELECTRONICS

1. OVERVIEW

This summary of the CMS tracker readout provides an introduction to the system, updated from the information in the Technical Design Report to reflect changes in the tracker since then. The main differences are the use only of silicon detectors and utilisation of 0.25µm CMOS technology for all integrated circuits.

In this description, there are links to various other documents, accessible via the Web where possible. Many papers related to the Tracker Readout System can be found at the following regularly updated site:

http://pcvlsi5.cern.ch:80/CMSTControl/

Requirements

The electronic challenges are principally to ensure sufficiently low noise during the full operational lifetime and adequate bunch crossing identification, ultimately limited by the speed and magnitude of detector signals.

The system performance target has been to achieve an Equivalent Noise Charge below 2000 electrons for a front-end chip power consumption close to 2mW/channel. Silicon signals will be localised in time to a precision of a single bunch crossing. There is a range of strip parameters in the system so the target reflects the likely variations in input capacitance and resistance, as well as radiation induced changes in leakage currents.

There are several reasons why CMS preferred an analogue system, among which are optimal spatial resolution from signal charge sharing, material budget, operational robustness, ease of monitoring, less custom radiation-hard electronics and possible future upgrades. They are discussed in the Tracker Technical Design Report in some detail.

System summary

Each microstrip is read out by a charge sensitive amplifier with a time constant of 50ns whose output voltage is sampled at the beam crossing rate of 40MHz. Samples are stored in an analogue pipeline for the first level trigger latency of 3.2µs and, following a trigger, are processed by an analogue circuit. This confines the silicon signal to a single beam crossing interval and provides measurement of signal amplitude and assignment of the hit bunch.

Pulse height data are multiplexed from pairs of 128 channel front-end chips on the detector hybrid onto a differential line over a short distance to a laser driver. Electrical to optical signal conversion follows, then transmission over a 100m fibre optic cable to the counting room adjacent to the cavern. The optical link employs edge-emitting semiconductor laser transmitters operating at a wavelength of 1300nm transmitting through single mode fibres.

The tracker data acquisition is based on a VMEbus system housed in the underground counting room outside the cavern. Pulse height data from the front-end chips, with no zero suppression, are converted back to electrical levels matched to the range of a 10bit ADC. Approximately 2 bits of the range allow for baseline level variations within the system. The remaining 8 bits are sufficient for adequate resolution over the range of signals expected. The Front End Driver (FED) digitises the data, performs some signal processing, including reordering and pedestal subtraction, and stores the results in a local memory until required by the CMS data acquisition. In high luminosity conditions at maximum trigger rate, cluster finding will reduce the data volume to be transmitted.

A VMEbus module in the counting room, the Front End Controller (FEC) controls and monitors the electronics system. It distributes the machine master clock and first level triggers received from the global LHC Timing Trigger and Command (TTC) system. Digital optical links send and receive trigger, clock and control data at 40MHz which are recovered by photodiodes and amplifiers and distributed electrically by a Communication and Control Unit (CCU) to detector modules. Clock signals are processed by Phase Locked Loop (PLL) chips on each module to ensure high reliability and minimum phase jitter. The FECs are located close to the shortest cable route to the cavern to minimise the contribution to the trigger latency.

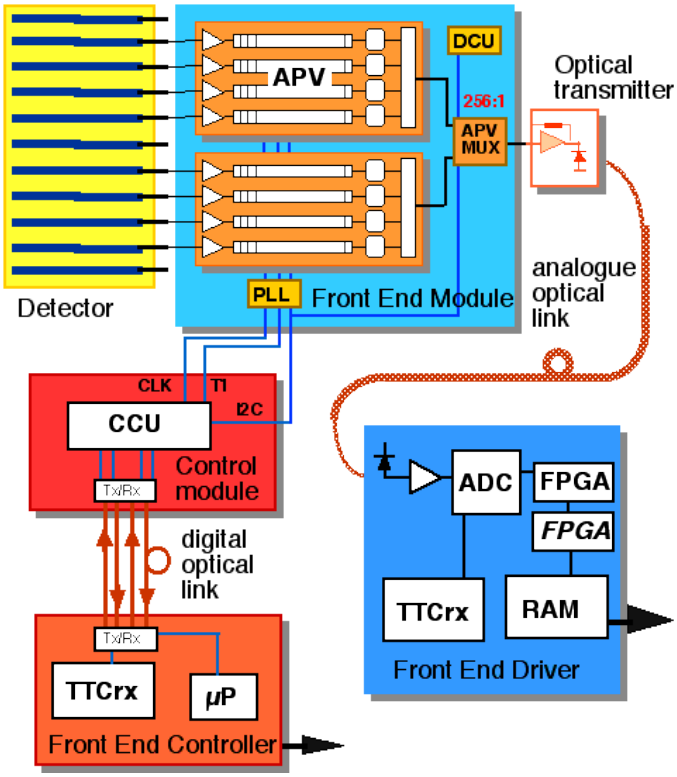

Fig.1.1. Schematic diagram of tracker readout and control system

Table 1.1 Summary of the components used in the system with their functions

Component	**Function**	**Number (very approx!)**
APV25	128 channel silicon front-end chip	90k
APVMUX	2 APV: 1 optical link multiplexer	45k
CCU	Distribution of clock, trigger, commands and monitoring data inside tracker	2k
PLL	Precision clock recovery and programmable delay	30k
DCU	Interface to monitor slowly varying parameters	25k
Laser driver	Bias and modulate laser current for analogue or digital transmission (4 channel)	15k
Laser diode	Electrical-optical converter for outgoing analogue or digital data	50k
Photodiode	Optical-electrical conversion of digital or analogue data	50k
Analogue receiver	Transimpedance amplifier for analogue signals	45k
Digital receiver	Transimpedance amplifier for digital control, clock and trigger	2-5k
LVDS Tx/Rx	converts digital electrical levels to LVDS standard	3k
FEC	distribution of clock, trigger and control data and reception of monitoring data, all via digital optical links	20-100
FED	reception and processing of analogue data and interface to CMS DAQ	550

2. RADIATION HARDNESS

The tracking system experiences one of the most severe radiation environments in the whole of CMS with 10 year doses and fluences up to about 15Mrad and $3x10^{14}$ charged hadrons.cm^{-2} at the innermost microstrips and ~1Mrad and 10^{13} neutrons.cm^{-2} for the outermost detectors.

Few tracker electronic components can be purchased formally qualified as radiation hard. A modern 0.25µm bulk CMOS process has been selected on the basis of demonstrated hardness at the levels required, as well as cost and performance. MOS devices rely on conduction in shallow channels very close to the gate oxide interfaces. They are intrinsically very hard against displacement damage and very high neutron levels can be tolerated without degradation. Ionising particles cause threshold voltage shifts, increased noise and reduced transconductance or gain through oxide charging. The measured noise, gain and threshold voltage changes are insignificant in the 0.25µm process to well beyond 20Mrad (200kGray).

Two other mechanisms for significant damage exist: latch-up and single event effects. Latch-up is caused by creation of parasitic bipolar devices by irradiation. Single event effects result from local heavy ionisation changing the state of circuit nodes. Both need to be measured experimentally. The CMS tracker circuits do not exhibit single event upsets at a high rate.

CMS optoelectronic components have shown impressive performance after irradiation and a brief summary is given below. A great variety of optoelectronic devices, fibres and connectors have been irradiated. Obviously, levels of radiation induced damage are very much dependent on component position inside the detector.

Semiconductor lasers are sensitive to displacement damage only. A shift in laser threshold current under irradiation will correspond, in operation, to a shift in the DC-bias point and will eventually have to be compensated for by the driving electronics to avoid falling into the sub-threshold regime of the laser. For a linear device, a varying bias point does not affect the modulated AC-signal. In contrast, the very small efficiency drop observed during irradiation will change the link AC-gain and will therefore have to be corrected for by recalibrating the system.

P-i-n photodiodes are also sensitive to displacement damage. The leakage currents increase non-linearly with fluence by up to 6-7 orders of magnitude above the pre-irradiation values of ~10pA. In contrast, there is only a small decrease in photocurrent up to high fluences.

Ionisation is the dominant damage source in optical fibres subject to irradiation. The induced loss measured in standard telecom fibres is ~0.12dB/m or less after 80kGy of total dose. This low level of damage indicates that the few meters of fibre exposed to high radiation levels in the tracker will not degrade the link performance significantly.

Extensive connector insertion-loss measurements and repetitive mate-demate tests have been performed on MT ferrules exposed to gamma and neutron radiation levels expected in LHC experiments. No significant degradation of the optical and mechanical performance of connectors could be detected, even after 100 mating cycles.

Details on radiation measurements of tracker components can be found in many technical reports and published papers.

3. FRONT-END ELECTRONICS

http://www.te.rl.ac.uk/med/apv25_web/user_guide.html

The front-end chip is based on a design denoted as the APV series. It is designed to operate with minimal deadtime up to trigger rates of 100kHz.

Each APV25 channel contains a preamplifier and shaper, with a 50ns peaking time, followed by a 192 location memory into which samples are written at 40MHz. Locations of data awaiting readout are flagged so they are not overwritten. Following a trigger, three samples from the memory are processed with the APSP deconvolution filter (http://pcvlsi5.cern.ch:80/CMSTControl/documents/GeneralDoc.htm), which re-filters the data with a shorter time constant.

The chip can be operated in three modes: **peak** mode, in which the output sample corresponds to the peak amplitude following a trigger, **deconvolution** mode, in which the output corresponds to the peak amplitude of the APSP filter, or **multi-mode** where, following a trigger, three samples are read out. After the APSP data are held in a further memory buffer prior to switching through an output analogue multiplexer. This is required so that one event can be multiplexed out while another is prepared for transmission.

The APV also contains system features including programmable on-chip analogue bias networks, a remotely controllable internal test pulse generation system and a slow control communication interface.

Analogue performance

Figs. 3.1 shows average amplifier pulse shapes in peak and deconvolution modes. The peak mode pulse is a good approximation to ideal 50ns CR-RC pulse shaping. The deconvolution mode data illustrate the effectiveness of the technique in achieving a short pulse shape for single bunch crossing time resolution.

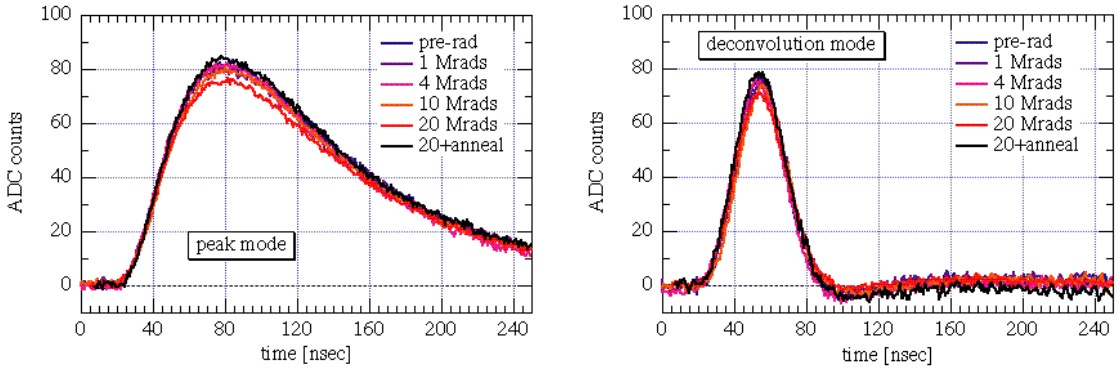

Fig3.1. APV25 Analogue pulse shape in peak and deconvolution mode

Fig. 3.2 shows typical measurements in both peak and deconvolution modes using the first version of the APV25 chip. The noise performance of the final version will improve further on this. The increase in noise after deconvolution is a simple consequence of the fast filter.

Analogue stages

The preamplifier is a charge sensitive amplifier with a PMOS input transistor of dimensions 2000µm/0.36µm and current of 400µA. It consumes 0.9mW and is the dominant contribution to the total APV25 power budget of 2.3mW/channel. The power supplies are +1.25 and –1.25V.A switchable unity gain inverter is used to allow signals of either polarity to be measured. The shaper is an effective CR-RC filter with shaping adjustable over a limited range. The total front end gain of the amplifier is approximately 100mV/MIP (25000e).

The pipeline is a 128 by 192 array of switched capacitor cells. Each cell comprises two transistors, to perform the read or write operation, and a storage capacitor. The pipeline is read out by the APSP processor which is an amplifier with a switched capacitor network in the feedback loop. The ratios of capacitors define the weights used by the deconvolution algorithm.

A 128:1 multiplexer drives the analogue output from the chip. A nested three level architecture results in analogue data emerging in non-consecutive channel order at 20MS/s following a digital header sequence. A

programmable resistor converts the voltage into a current output. Data from two APV chips are interleaved at the APVMUX chip to arrive at the final transmission speed of 40MS/s.

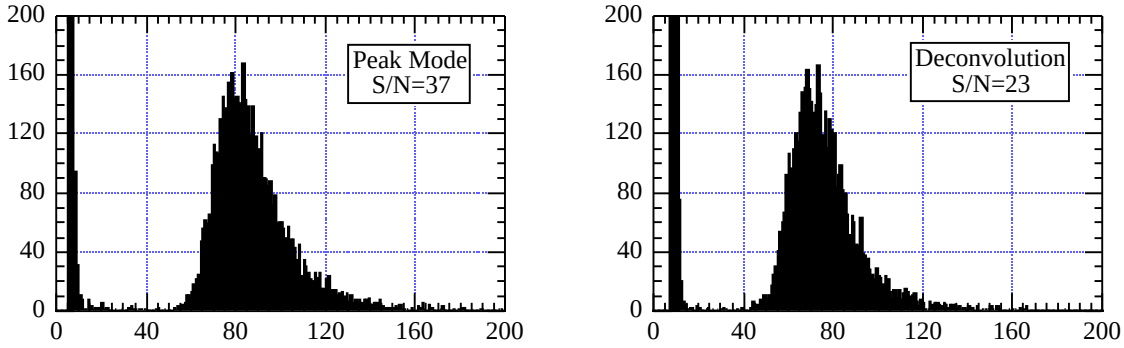

Fig. 3.2 APV25 signal to noise measured in peak and deconvolution mode with a 5cm, 50µm microstrip detector using a Strontium beta source.

Control interface

The configuration of the APV25 is handled with a serial interface based on the Philips I2C standard. Up to 15 APV25 chips may share the same controller. A command register defines which variable or register is to be accessed and specifies the direction of data transfer. These allow setting of bias currents and voltages in the amplifier and shaper, choice of operation mode, calibration, latency adjustment and error reporting, etc.

Error conditions

The tracker readout is fully synchronous so synchronisation errors must be avoided and faulty chips identified. There are a few ways in which errors could arise. For short periods the APV25 may be triggered faster than it can output data. If so, a queue develops. As the queue grows the FIFO, which stores addresses of columns awaiting readout, could become full or all available memory locations be allocated. Either case leads to pipeline failure and must be avoided. This will be done by real time emulation of the state of the APV chips, and trigger inhibition if buffers are almost full.

Operation

The APV25 requires only one control line on which three commands can be sent: TRIGGER, Test Pulse Request or RESET101. The RESET101 signal clears the pipeline and initialises pointers. After power-up and programming of control registers a RESET101 is required before TRIGGER or test pulse requests can be sent.

The trigger latency is the time between a detector signal at an input and a TRIGGER signal applied to output that signal. It corresponds to the separation of read and write pointers in the pipeline memory.

Output data stream

The APV25 multiplexer voltage output is converted to a differential current. When there is nothing to transfer the output of the chip is at the logic “0” level with single logic “1”s (referred to as ticks) every 1.75µs. When an event has been triggered, a data frame made up of a digital header, containing the pipeline address and an error flag, and an analogue data series is output. The analogue baseline may be adjusted to give optimal dynamic range.

Calibration

An on-chip calibration system is provided to simplify testing and monitoring. It applies a voltage step to the preamplifier inputs in 8 groups of 16; the circuit delay can be adjusted in steps of one eighth of a clock period (3.125ns).. One application is adjusting the amplifier pulse shape. Applying a calibrate signal to the APV produces a pulse at the amplifier output which can be measured by sampling the pulse at a fixed time and progressively shifting the timing of the test pulse, building the shape from those samples.

APVMUX

Matching of laser transmitters to detector modules, which contain differing numbers of APVs, is necessary to utilise fully the optical channels. The laser drivers are mounted a short distance from the front-end hybrids and low mass cables transfer APV output streams differentially.

Savings in link costs are achieved by multiplexing data from two adjacent APVs onto each optical channel, with 256 detector channels to each fibre. Multiplexing is done by interleaving data from pairs of APVs transmitting at 20Ms/s into a 40Ms/s stream. Each APVMUX chip contains 8 APV inputs and 4 differential outputs. Programmable termination resistors are provided.

To control the switching, the clock and trigger signals used by the APVs are also received on the APVMUX. For optimal performance the data stream from each APV is skewed by 25ns if its local address is an odd number, with even and odd APVs providing the two inputs, so that maximum settling of the analogue levels at the APVMUX inputs is achieved.

Testing and qualification

Automatic wafer testing of each chip will identify “known good die” which will be cut from wafers and assembled onto the hybrids and re-tested prior to module assembly. All of this must be tracked and catalogued. The time required to test each APV die on the wafer is ~2 minutes. Some radiation qualification will be required, so sample die will be irradiated using x-rays at regular intervals.

Once known good die have been identified, batches of wafers will be assembled and sent for cutting. Chips will be distributed to other CMS centres for assembly onto hybrids. Detector module assembly will be undertaken by other centres within the collaboration.

Front-end hybrid

The detector module front-end hybrid supports the APV chips, the APVMUX and PLL chips, together with power supply decoupling capacitors and the DCU. The number of APV chips per hybrid varies throughout the tracker. The physical layout of the hybrid must enable efficient transfer of heat to the tracker cooling system, most of the power being dissipated in the APV chips.

4. OPTICAL LINKS

http://www.cern.ch/cern/divisions/ecp/cme/opticallinks/

Optical data transmission is essential for the CMS tracking system to transfer data from the tracker volume with minimal contribution to the material budget. An optical system will also be immune to electrical interference. CMS has chosen analogue data transmission at 40MS/s. The transfer rate is well below the full bandwidth of the fibre link and more than sufficient linearity (<2%) and dynamic range (up to 9bits) for tracking requirements have been demonstrated.

The transmitter is an edge emitting semiconductor laser; several devices have been identified with the required characteristics. They are mounted with fibre pigtails in a low mass package. A minimum threshold current is required before laser action commences; presently available devices have sufficiently low thresholds and forward voltage drops and high enough electrical to optical conversion efficiency that the power consumption is a small fraction of the tracker power budget. Above threshold light output power is highly linear with current over a wide range.

The optical link under development for the CMS tracker is shown schematically below. Three connector breakpoints (two inside the detector and one on the backplane of the readout or control electronics) allow easy testing, installation and maintenance of the system. The total length of the link is approximately 100m, of which about 10m is within the high radiation environment.

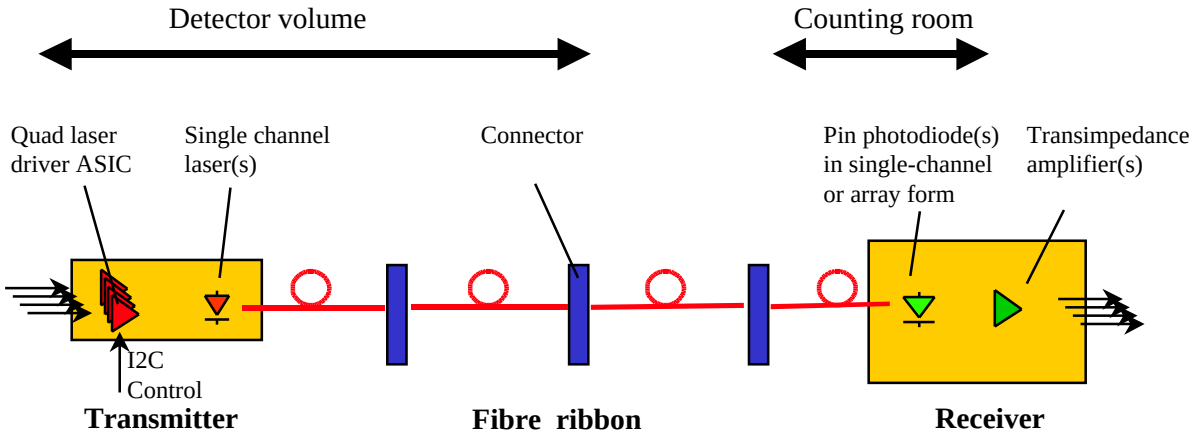

Fig. 4.1 Generic optical link block diagram.

The optical link development effort relies, as much as possible, on commercially available products. Laser and pin-photodiode die, fibre and multi-way connectors are all off-the-shelf components selected after evaluation and irradiation. Custom developments are required for radiation hard front-end electronics as well as for non-magnetic, low mass packages and connectors.

System description

The projected layout for the optical readout system is shown in fig 4.2. To increase cable and patch panel density, fibres are merged to 12 way ribbons at the first patch panel (1) situated at the edge of the tracker mechanical support structures; ribbons are then packed into 8-ribbon cables (96-way) at the second patch panel (2). The 96-fibre cable runs out of the detector towards the counting room. It fans out at the Front End Driver (FED) readout board.

The projected layout for the optical part of the digital control system is illustrated in fig.4.3. It is a bi-directional link with control data, clock and First Level Trigger (T1) signals sent from the back- to the front-end on two fibres, as well as status data and clock signals returned from the front- to the back-end, also on two fibres. The data transfer rate is 40Mb/s and the clock frequency is 40MHz.

The similarity between the readout and control system layout is evident. The number of links needed for the control system will be about two orders of magnitude smaller than for the readout, but the technologies and components used will be identical as far as possible.

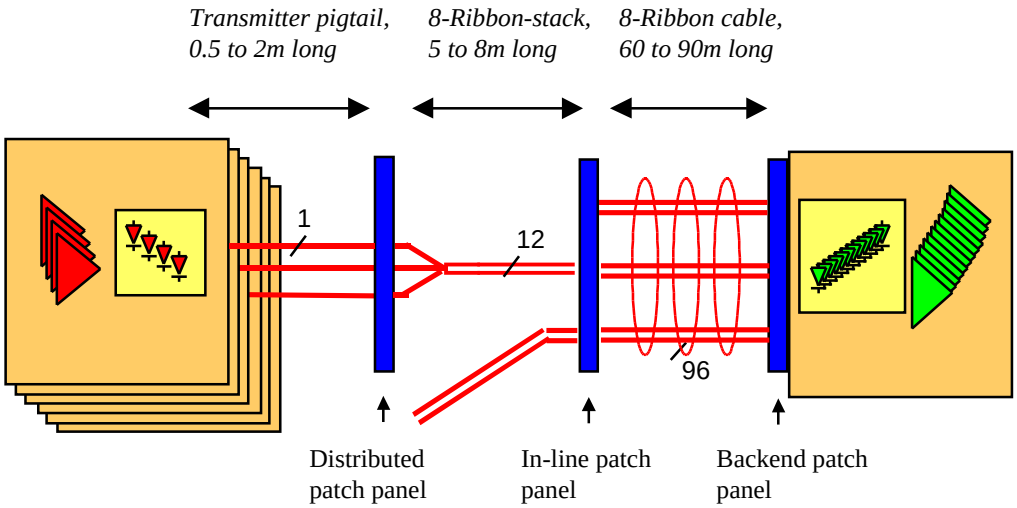

Fig4.2. Layout of the analogue readout optical system for the CMS-tracker:

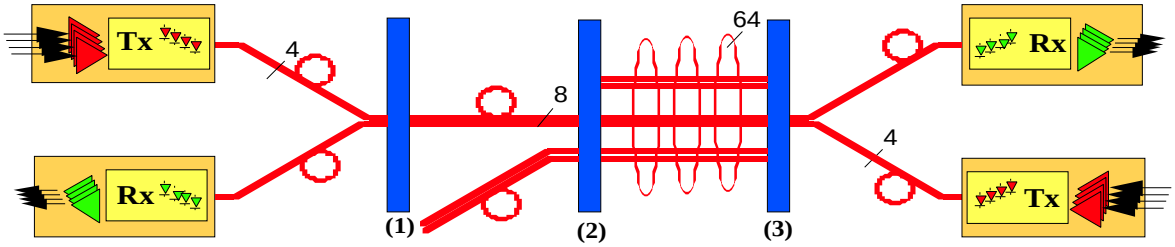

Fig. 4.3 Layout of the digital control and timing distribution optical system.

Electronics

A 4-channel linear laser-driver ASIC, composed of four linear drivers and an I2C control interface, is used in both the analogue and digital links. Each channel in the driver takes as input a differential voltage - originating from one of the tracker detector hybrids - and converts it into an unipolar current that is used to modulate the laser-diode. Besides signal modulation, each driver in the IC also generates a DC current that serves to bias the laser-diode above threshold in the linear region of its characteristic. This bias current is programmable through the I2C interface to compensate for threshold current variations due to device ageing and radiation induced performance degradation. Additionally, a "power-down" function is implemented for each channel allowing reduction in the power consumption of non-used or defective channels.

A 12 channel linear amplifier array is integrated into the analogue receiver module. The ASIC features a current output, allowing precision gain adjustments via the external resistor terminating the line. It is realised in a standard BiCMOS process and interfaces directly to the 10bit A/D converters located on the FED.

A radiation-hard digital receiver is used at the front-end side of the control link. The ASIC operates at 40MHz for the clock line and 40Mb/s for the signal line. It accommodates a wide dynamic range with its automatic gain control capability, and complies with large DC input currents.

Optoelectronics

The laser diode chips are commercially available Multi-Quantum-Well (MQW) InGaAsP edge-emitting devices selected for their good linearity, low threshold current and proven reliability (lifetime well in excess of 10^6 hours).

Photodiodes are epitaxially grown, planar InGaAs devices of small active volume.

Optical fibre, cables and connectors

The optical fibre used is a standard, single-mode, non dispersion shifted telecommunication fibre.

The 96-way cable (64-way for the digital system) features a small diameter (<10mm) and a low bending radius (8cm). Each cable serves one Front End Driver (FED) board.

The three connectorised breakpoints allow for easy testing, installation and maintenance of the system. Connectors based on single-channel, small form factor ferrules, or multi-channel angle-polished MT ferrules are used throughout, achieving the highest patch panel density currently possible with commercial components.

OPTO HYBRIDS

Custom hybrids are needed for the transmitters and laser drivers for the analogue data. (http://www.pg.infn.it/cmsweb/optical.htm). The digital optoelectronic interface for the control system between the FEC and CCU also requires a custom development matching the local topology of the tracker.

5. FRONT-END DRIVER

The final Front End Driver module is under development and will not be produced before 2001. Meanwhile, an inexpensive PCI Mezzanine Card (PMC) has been produced for use in small scale systems and beam tests. It will be complemented in the near future by an interface card to be used in conjunction with tracker optical links. Information on the FED can be found at: http://hepwww.rl.ac.uk/CMS_fed/Default.htm

Final FED

The Front End Driver module (FED) receives 96 analogue signals (compared to 64 in the TTDR) each of which is digitised by a fast ADC, digitally processed and packed with event header information before transmission to the next level of the CMS Data Acquisition system. The major features of the module, many of which will be implemented in FPGAs for flexibility, are summarised below.

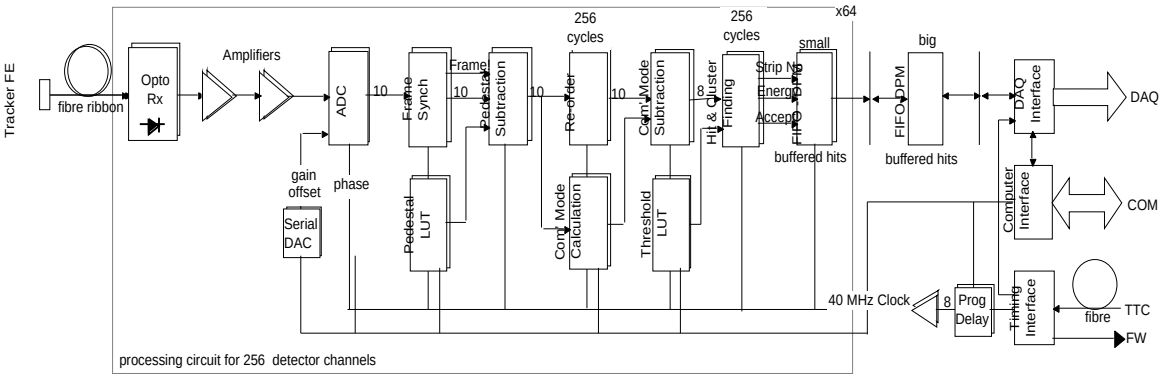

Fig.5.1 . Front End Driver Block Diagram

Optical input stage

The FED receives analogue data on 12-way ribbons, each fibre carrying data from a pair of APVs. The optical signal is converted to electrical levels by the opto-receiver unit described earlier mounted in sockets in the mother board close to the ADC with fibre ribbons routed through a front panel connector. The ADC is a 40 MHz, 10 bit commercial component. The data from each ADC is processed by its own independent digital parallel-pipelined processing logic running at 40 MHz.

Synchronisation

The start of an APV data frame is recognised, producing a start frame signal which triggers the next processing stage. It employs an auto-synchronisation method which locks directly to the APV data stream requiring no external trigger. (There is no simple method of predicting the arrival of APV data frames relative to the trigger, except by state machine emulation, due to the complex logic of the APV pipeline.) The circuit also extracts the pipeline address from the data header and performs synchronisation checks.

Re-ordering

The APV outputs detector channels in non sequential, but fixed, order so that strips with shared signals are not adjacent at the digitisation stage. The cluster finding logic requires data to be in channel order and thus the APV frame must be re-ordered before hit finding. Re-ordering is triggered by a start pulse from the synchronisation circuitry. A re-ordered data stream is clocked out 256 cycles later, along with a start signal to indicate the first value to the next stage.

Pedestal correction

Each APV chip will have a unique static offset level associated with each detector channel. Although these variations should be small it is necessary to compensate for them due to their effect on subsequent processing. This is done by storing pedestal values in an SRAM Look Up Table and subtracting them from

the data stream as they are clocked through. The pedestal-corrected data frame is passed onto the next stage with a latency of only a few clock cycles .

Common mode correction

At present common mode is assumed to be a common DC offset, captured on every APV25 input channel. It could be different for every APV and each trigger. Thus each APV requires a separate calculation for every trigger.. Any faulty channels will be excluded from the summations.

Cluster finding

The basic principle of the hit finding circuit is to identify which channels are hit and write them into the output buffer. The hit occupancy is 1-3% so an average data frame will contain a few hits. Only these are written to the output buffer while the remaining data are discarded. Each hit value must be tagged with an address as it is no longer uniquely identified by its position in the frame. Finally the sparsified data, data length and processing history, are clocked into the output buffers ready for readout by the FED DAQ interface. The sparsified frame data will require a header consisting of processing channel identity, status, type and processing history.

Output

Approximately 512 cycles after the first analogue sample is clocked into the processing chain, the output buffers will contain the sparsified hits. The model assumes three FIFOs, for data, length and header data, which can buffer multiple events. The number of hits in the data FIFO for each trigger is random so an associated length buffer is required so that the readout interface knows how much data to expect from each triggered event.

Calibration and Monitoring Modes

All calibration and monitoring on the FED is expected to be performed through the FED computer system interface under the direction of the external computing system. Post-ADC processing requires set-up of the processing algorithm, such as pedestal correction and threshold LUTs to be determined during calibration.

Interfaces

The function of the DAQ interface is to create on the FED the output data block for each event and transmit this to the DAQ system via the FED-RPDM data link. The data block is created by merging 96 individual data blocks from the post-ADC processing with TTC and other header information for each triggered event.

The Readout Dual Port Memory (RDPM) is the first stage of the DAQ system designed to receive and merge data from 1 to 8 FEDs, depending on the particular detector sub-system, at up to 400Mbyte/s via the FED-RDPM data link.

The Computer System Interface provides access to all internal FED processes to the external computer system for purposes of set-up, control and monitoring. The circuit will probably be implemented in an FPGA for flexibility with a single board computer in the FED crate providing a bridge via the VME backplane.

The Trigger, Timing and Command interface is based around the TTCrx ASIC to decode all TTC broadcast information, particularly the clock, trigger, event and bunch counters. The clock is fed into programmable delay lines to de-skew the clock in steps over a 25ns range. The delayed clocks supply groups of ADCs receiving signals from the same fibre ribbon to compensate for the lengths of different ribbons.

Event sizes and rates

The event size output to the DAQ is a product of the number of detector channels per FED (96), the average occupancy of detector channels served by the FED(<3%) and the number of bytes used to code the data. For a simple zero suppression scheme each hit would be coded as 2 bytes, namely strip address and signal amplitude. Additionally the event size must include the fixed length header and trailer data. The data rate is the product of event size and trigger rate.

In heavy ion running the occupancy of hit detector channels is much higher, up to 18% in the outer layers (the inner layers are expected to be switched off), which gives a FED event size of around 8Kbytes. However the trigger rate is reduced to ~8kHz for Pb-Pb collisions, assuming full acceptance, which gives each FED an output rate of 64Mbyte/s. This is equivalent to p-p running with a 1% occupancy and a 100kHz trigger rate.

6. CONTROL AND MONITORING

The control and monitoring of the tracker can be divided into two areas:

the **internal** control of the readout system and monitoring of some tracker specific components installed in the tracker volume itself, which requires custom radiation hard components.

the operation of **external** elements, such as power supplies and the computer interface to users who will operate CMS, managed by a Detector Control System (DCS) common to the entire experiment.

The **internal** tracker control system distributes the LHC clock and trigger signals and supervises embedded electronics. It consists of three main functional elements:

- a **network** based on a simple token-ring architecture between control room and embedded electronics; the long sections operate over optical fibres, the short ones on low mass copper cables,
- a **Front End Controller** (FEC) card in a control room VME crate manages the network and interfaces it to the CMS slow control system,
- a **Communication and Control Unit** (CCU) linking the network and on-detector ICs.

Communication between electronics on the CMS tracker and outside will use a ribbon of four optical fibres. Two fibres are used for sending the timing and data signals to the front-end and two to transmit a return clock and the return data.

To minimise costs, it is possible to connect a number of control modules serially in a ring-like arrangement. The tracker will have up to 20,000 front-end modules so, grouping 10-20 modules on each CCU and 1-4 CCUs per ring, the number of network rings is expected to be between 250 and 1500.

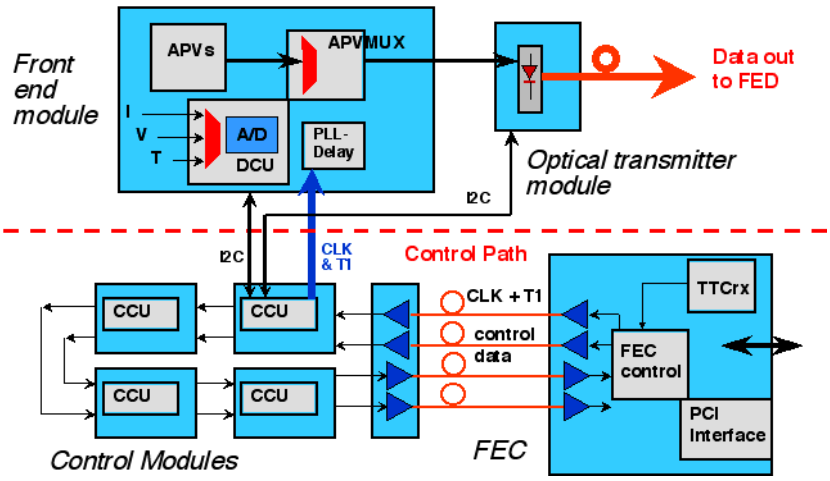

Fig. 6.1. Control system schematic

Components

Communication and Control Unit (CCU) and CCU Module (CCUM)

The CCU ASIC is the special purpose integrated circuit which implements the dedicated control link in the CMS tracker for control and monitoring of embedded front-end electronics and distribution of the time critical trigger and of the low jitter clock to the front-end. A block diagram of the CCU is shown in fig.6.1.

The CCU interfaces the ring network to front-end chips using the industry standard I2C bus, which was chosen for its simplicity. It is also completely silent (no clock running) outside transfers, avoiding switching noise close to the front-end chips.

Other interfaces supported by the CCU are a simple 8 bit wide memory-oriented, non-multiplexed bus, a programmable parallel interface, which can be used to control or read switches or other semi-static elements, and a decoded trigger port, for trigger information encoded in the FEC.

One CCU ASIC mounted on a CCUM is dedicated to a set of tracker front-end modules.

Phase Locked Loop chip (PLL)

A highly reliable clock will be ensured by a custom PLL on embedded control modules and front-end modules providing a very low jitter, phase adjustable clock signal to the local electronics. The clock and trigger signals are coded on the FEC with a very simple protocol as shown in fig. 6.2. A clock pulse is removed from the clock train to indicate the presence of a L1 trigger signal. The missing clock pulse and the corresponding trigger signals are regenerated by the PLL.

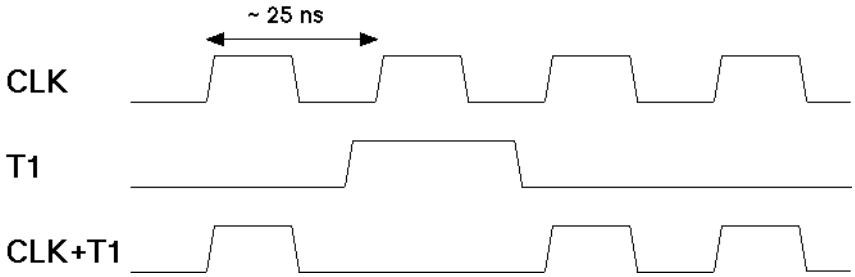

Fig.6.2 Encoding of the clock and trigger signals

The chip measures ~2 by 2.5 mm^2 and requires less than 50mW power. The output clock phase can be adjusted from 0 to 25ns.

Detector Control Unit (DCU)

This ASIC is used to measure and monitor slowly varying analogue signals such as temperature, low voltages, detector currents etc. It will be mounted on the detector hybrid. The DCU is connected to the external FEC via the CCU. It will be a 2mm x 2mm chip containing a 12 bit ADC with a power consumption of less than 50mW.

Front End Controller

The FEC controls operation of the slow control ring network. Some vital network monitoring and initialisation functions are generated only on the FEC with most of them implemented in software.

The FEC could be built around a commercial VME CPU module with a user definable PCI based optical adapter (on a PMC card) where the actual serialising hardware - specific to the ring protocol - and the optical translation are implemented. This will allow us to adapt to the most modern commercial cards, by pushing the implementation of most functions onto software, and to minimise construction of special hardware. A FEC card is expected to support several 4-way optical links. The FEC modules will interface via VME to a single board computer in the crate and to the Detector Control System.

The FEC receives the timing and trigger from the TTC system and encodes it a way suitable for the CCUs. This encoding is already implemented in the current version of the TTCrx ASIC and therefore a special PMC housing a TTCrx ASIC will be present on the FEC.

Trigger Supervisor

An important function is emulation of the APV front-end pipeline logic to ensure that buffers do not overflow by inhibiting the trigger for the few events which would overfill the pipeline. The trigger supervisor, which is being defined, is required to monitor all buffers in the system to avoid such problems.

Partitions

The readout and control system will be grouped into sub-systems, referred to as partitions. Around 4-8 partitions are foreseen, with each unit served by its own branch of the TTC network. It is important, both electrically and logically, that the control and readout elements are common to a single partition and not shared.

7. LOW VOLTAGE POWER SUPPLY

The very large number of front-end electronics channels requires stable, low-noise power to be supplied to the front-ends with overall currents in the range of 15kA at the ±1.25V nominal voltage, with a total power dissipation of about 60kW. The cooling system is required to evacuate the heat generated with negligible rejection into the ambient environment, and to keep the silicon tracker at a temperature of around -5°C.

The granularity of the low-power supply system is determined by detector configuration, material budget, and cost. The design should not include any voltage or current regulators in the front-ends, where the radiation levels are so high that the feasibility of such power devices is unproven even with radiation hardened processes. The baseline configuration is a highly distributed system with individual power supplies feeding groups of detector modules.

Care will be needed in the design because of the implications of the unbalanced positive and negative supply currents to the APV, which is a design feature to save power.

Constraints: magnetic field, radiation background, cable cooling

A key issue is whether power supplies should be installed in the cavern or in the counting room. Generally power sources are placed close to the load, to simplify regulation and minimise the risk of noise pick-up. However, this approach is made difficult by a relatively strong magnetic field and significant radiation levels. The magnetic field is estimated to be around 50gauss near the walls, and can reach several hundred gauss on the balconies where electronics racks can be installed. Any inductor or transformer using ferromagnetic material would need adequate shielding if placed close to the detector.

The background radiation level in the cavern may impose more constraints. Current estimates point to an integrated dose and fluence of about 300rad and 10^{11} $n.cm^{-2}$ over 10 years. The uncertainty on these calculations makes it necessary to adopt a conservative approach. Total dose effects can be dramatic on some commercial components. Neutron fluences can lead to total failure of commercial bipolar devices, such as those in power circuits, particularly in the low dose-rate conditions of the cavern.

An option is to place the power supply system in the counting room. In this case LV cables must be water cooled to remove power losses since the capacity of the cavern air conditioning system is limited. With power supplied at distance, and without local regulators, electromagnetic compatibility (EMC) is of primary importance. Power cabling from the relay rack to the patch panel should be as compact as possible to reduce the cross-section for services and minimise cracks. For this purpose, the insulator/conductor ratio in the cables should be kept to a minimum.

8. SYSTEM OPERATION

Early LHC operation requires careful planning and verification. An LHC-like test beam using modest scale prototypes is vital during the next few years to develop procedures. It is possible to foresee broadly how this should work.

Initialisation

This will be a data intensive operation and experience from prototyping will be important. The 100,000 front-end chips and 50,000 optical links must be set to their operating points accurately and quickly.

It is expected that a data base of provisional parameter settings will exist. Settings and bias values for each APV chip must be downloaded via the serial link followed by trigger delay scans, using calibrate signals, to ensure that pulses are set for the correct trigger latency. Peak mode operation is foreseen in initial low luminosity LHC conditions, so timing precision to about 10-20ns may be adequate.

Synchronisation

Once each module has demonstrated satisfactory response, timing with proton-proton interactions can be verified. Minimum bias triggers will be the first test of adjustments for time of flight and cable delays. This will be simplified by the slow shaping (50ns peaking) of the APV amplifier and absence of pileup. Simple on-line track finding will be the test.

Different sub-detectors must assign the same event to identical bunch crossings. A possible solution is to use a physics trigger, such as a high p_T jet, which generates a large number of tracks in a small volume and correlate track segments in the microstrips with energy deposits in the ECAL.

Eventually verification of the system will be possible using correlations between the LHC beam structure and observed hits or tracks. As LHC luminosity increases, timing precision should gradually improve, eventually switching to deconvolution mode operation.

Operational performance

The principal contributions to the system noise are from the APV, which is the largest source and dominated by the input transistor. Other contributions come from thermal noise due to the resistance of the microstrip itself (taking into account the distributed resistance and capacitance), leakage current, affected by radiation damage and cooling, thermal noise in the bias resistor, noise due to the optical link, which depends on the gain of preceding elements as well as intrinsic noise from components, and quantisation error in digitisation, also gain dependent.

Calibration

This should be undertaken regularly because of possible radiation damage. The tracker is self-calibrating since the Landau distribution should allow reliable monitoring. Regular electronic calibration should take place in dedicated runs. Internal calibrate levels are not known precisely, except by reference to an energy measurement, and optical fibre disconnection and reconnection can change the link attenuation. Amplitudes, pulse shapes and laser thresholds can be compared and monitored. Calibration within physics data taking runs is not anticipated and the control network will be inactive during running, except for alarms.

Long term maintenance

Maintenance will be difficult given the inaccessibility for long periods of time. Even during early periods, when induced radiation will be low, it will be difficult to exchange internal components without major dismantling, which should be avoided as far as possible. Prototype system operation will be important to diagnosis potential weaknesses in components, assembly or handling procedures. Simulated LHC operation will be essential to learn some of the pitfalls to be avoided.

Reliability will be the key to successful long term operation. Although much attention has been given to radiation damage, component failure may arise for mechanical reasons or ageing for other unexpected reasons.

GLOSSARY

APV	Front-end readout chip: Analogue Pipeline (Voltage mode)
APV25	Silicon APV in 0.25µm CMOS
APVMUX	APV Multiplexer
ASIC	Application Specific Integrated Circuit
CCU	Communication and Control Unit
CCUM	CCU Module
DCS	Detector Control System
DCU	Detector Control Unit
FEC	Front End Controller
FED	Front End Driver
FIFO	First In First Out buffer
FPGA	Field Programmable Gate Array (=PLD)
I2C	Philips two wire serial command protocol http://www-us.semiconductors.philips.com/i2c/
LUT	Look Up Table
LVDS	Low Voltage Differential Signalling http://www.national.com/appinfo/lvds/
MCM	Multi-Chip Module
PCI	Peripheral Component Interface
PLL	Phase Locked Loop (chip)
PMC	PCI Mezzanine
RDPM	Readout Dual Port Memory
TTC	Timing Trigger and Command
TTCrx	TTC receiver
TTDR	Tracker Technical Design Report
T1	First Level Trigger