LHC FRONT END ELECTRONICS

G. Hall

Blackett Laboratory, Imperial College, London SW7 2AZ, UK

g.hall@ic.ac.uk

Abstract

Progress in selected areas of the front end electronic systems being constructed for the general purpose LHC experiments is reviewed, with emphasis on tracking and calorimetry.

1. Introduction

To review modern front end electronics is an intimidating task, since many developments are under way or in use, employing a wide range of techniques and technologies. It is impossible to do justice to all of them in a single conference review. This paper restricts itself to very topical LHC systems but, even then, the large number of applications must be limited and therefore focuses on ATLAS and CMS which are certainly the largest and possibly the most demanding and furthest advanced. However, the other LHC experiments exploit fascinating detector systems and face some very difficult challenges: for example ALICE has low radiation levels compared to CMS and ATLAS but demands ultra high precision timing, while in LHCb the electronics are more accessible but the experiment will endure some very high radiation levels and a very challenging trigger requirement. Hopefully future reviews will pay more attention to the progress in these systems.

To keep this review manageable, and following advice from the organisers of the symposium, the emphasis was placed on areas requiring radiation hard electronics and therefore discusses in most detail recent progress in trackers and electromagnetic calorimetry. A few selected systems are summarised in Sections 6 and 7.

2. Operation at LHC

A summary is given in Table 1 of the relevant LHC operating parameters for CMS. There are only minor differences for ATLAS, except that heavy ion operation is not part of the plans.

Table 1		
particles	pp	Pb-Pb
Luminosity	10^{34} cm^{-2}.s^{-1}	10^{27} cm^{-2}.s^{-1}
Average integrated luminosity	5×10^{40} cm^{-2}.y^{-1}	10^{33} cm^{-2}. y^{-1} ?
CM energy	14 TeV	5.5 TeV/N
$\sigma_{inelastic}$	~70mb	~6.5 b
Interactions/bunch	~20	0.001
Tracks/unit rapidity interval	~140	3000-8000
beam diameter	20μm	20μm
bunch length	75mm	75mm
beam crossing rate	40MHz	8MHz
L1 trigger delay	≈3.2μsec	≈3.2μsec
L1 trigger rate	≤100kHz	< 8kHz

Some consequences of LHC experiment operation are well known. The high speed signal processing has implications for front end power, signal to noise ratios and overall performance. There are high charged particle and neutron fluences which must be prevented from degrading performance and raise issues of qualification, reliability and maintenance. It is less commonly recognised that radiation concerns exist everywhere, up to the cavern walls where few krad doses are expected along with neutron fluences ~10^{11}cm^{-2} over a ten year operational period. High magnetic fields are to be used (4T for CMS) and significant stray fields are anticipated. The very large data volumes require fast front end data links and efficient processing.

These factors impose the use of radiation hard ASICs, semiconductor opto-electronics and optical fibres. Most components will be employed on a scale not before encountered and some R&D is still underway. Additional constraints are shrinking resources in budgets, manpower *and* timescale as LHC operation approaches, which have implications for the potential to exploit newer technologies.

Clock, trigger & control distribution

Mention should be made of an important common system which will be utilised throughout the LHC experiments for the distribution of clock, trigger and various commands [1]. The TTC system will distribute timing pulses derived from an LHC master clock using high power lasers operating at 1310nm by means of optical fibres with small dispersion and a passive hierarchical network of optical couplers. The system is well advanced and in use in a number of prototype systems. In the experiments typical splitting ratios will be up to 1:1000 from a single source and a timing jitter on clock edges of ~50ps has been achieved. In most cases, this is more than adequate for the sub-detector systems. Commercial components will be used, although some developments have been required to ensure smaller packages than typically found in commercial applications, e.g. for optical connectors.

The high TTC precision permits some sub-detectors to redistribute the clock within their own system using their own components, probably at the expense of slightly degraded performance. In this way, they can avoid the relatively power hungry and bulky TTC receiver components, which would otherwise be required in radiation hard form, and still achieve excellent performance.

3 Sub-detector features

To set the scene, it may be helpful to identify the main parameters driving the definition of the various electronic systems. Descriptive terms like “low power” and “low radiation levels” can mean quite different things in different regions of the experiments.

The tracking requirements are high spatial precision on particle trajectories but with limited energy precision. Thus detectors are highly granular, leading to enormous channel counts. The dynamic range of signals is limited to a few bits since it is rare that highly linear responses are needed for the full range of signal amplitudes; of more importance is the ability to distinguish genuine signals from background, so low noise readout has higher priority. This is not trivial at LHC since it must be achieved using power of a few mW/channel with fast pulse shaping and in the presence of high radiation levels, typically multi-Mrad. Power constraints are imposed by the need for a low material budget to minimise photon conversions and electron bremsstrahlung, so cables and cooling mass should be minimised.

The tracking systems experience some of the most severe radiation environments with, e.g. in CMS, 10 year doses and fluences of about 40Mrad and 8x10^{14} charged hadrons.cm^{-2} in the vicinity of the pixel system, ~15Mrad and 3x10^{14} charged hadrons.cm^{-2} at the innermost silicon microstrips and ~1Mrad and 10^{13} neutrons.cm^{-2} for the outermost layers.

The calorimeter detectors aim for very high energy resolution over a very large dynamic range; the performance must also be very stable with time. Intermediate radiation levels up to ~1Mrad are typical. Power constraints exist since removing heat is still a problem, both in the region of the sensor and from ohmic heating in cables.

Muon systems are spread over large areas. Although the spatial precision required is moderate compared to the trackers, it must be achieved over much larger areas with small systematic effects, so alignment and stability are vital features. Particle counting rates seem small when quoted as a function of area but constitute high values when account is taken of the large area over which signals are integrated. Radiation levels are low but heavy reliance on programmable logic implemented in standard commercial parts raises difficult questions of qualification and maintenance.

Radiation levels inside the experimental halls are nowhere zero and use of Commercial Off The Shelf (COTS) parts means radiation qualification will be mandatory for many commercial systems, such as power supplies and control systems, which are situated in the caverns.

Generic LHC readout system

Two variants of a typical readout system are illustrated in fig.1. All systems require amplification and filtering, analogue to digital conversion (at least one bit!) and storage prior to data readout initiated by a trigger. Pipeline memories are utilised throughout the experiments but they appear in different forms, either analogue or digital, and in different locations, either on or off the front end. In all cases some means must be provided to associate data to the correct beam crossing and the complexity of this requirement is influenced by the point at which data are digitised and compressed or "sparsified". The calorimeter and muon systems have an additional task not required of the trackers, which is to provide fast data for first level trigger generation.

"Deadtime free" operation is a requirement of all systems which, in practice, is taken to mean data losses well below 1% at the maximum trigger rate of 100kHz. In many cases, this is provided by circular pipeline architectures where pointers circulate and mark, and skip over, cells which contain data of interest. Thus the problem becomes one of ensuring a sufficiently long memory to meet the trigger latency requirement and sufficient extra cells to provide the depth of buffering for data awaiting readout. It is a classical queueing problem whose losses can be computed [2] but detailed logic emulation is needed to study in detail. Poisson fluctuations in trigger arrival mean that practical buffer depths cannot be sufficient for unlimited operation without overflows [3]; some examples are shown in fig.2.

Data are lost if buffers overflow but this would normally be either fatal or require action to reset the logic, so action must be taken to inhibit triggers instead. Alternatively, pipelines can be intermittently reset to zero but this strategy has the drawbacks of being time-consuming and, more important, not providing information about real life operation which may turn out to be vital.

4. Radiation tolerance of electronics

For convenience, one can consider two classes of electronics: COTS components, including Field Programmable Gate Arrays (FPGAs), and Application Specific Circuits (ASICs). The radiation tolerance of each strongly depends on technology and application, e.g. current density employed. The principal origins of damage are ionisation charge in oxides and displacement damage to substrates. Further generalisations can be dangerous but some introductory comments.

CMOS circuits are neutron hard but ionising dose sensitive. Typical consequences of oxide damage are gate threshold voltage shifts, increased noise, inter-device leakage currents and latch-up caused by creation of parasitic devices, which can be fatal. Conventional processes, and COTS components made in them, may have very low tolerance, even as low as a few krad.

A further phenomenon in logic circuits which has begun to be increasingly studied as other effects are better understood is Single Event Upset (SEU), in which knock-on ions deposit heavy ionisation in a small volume near to a sensitive circuit node and cause a change of logic state. Even in low flux conditions, it could give rise to unexpected operational errors in control circuits and therefore have magnified importance [4]. Fig.3 shows some recent measurements of SEU cross-sections [5].

Bipolar circuits are affected by atomic displacement which causes carrier recombination in the base of transistors giving rise to gain and matching degradation, but these effects can also result from oxide charging in bipolar ICs. Low energy neutrons diffusing throughout the LHC caverns will be a

significant damage source. Discrete bipolar circuits are hadron sensitive, but usually still withstand high exposures, and the few integrated circuit processes in use for particle physics provide highly tolerant ICs although dose and rate effects are not fully understood.

Electronics on III-V semiconductor materials, including optical components, are also usually hard to high levels.

Not all processes are suitable for every application. CMOS is most commonly used for low power circuits and essential for logic design. A few hardened technologies are available and have been intensively investigated. Military applications have declined in importance in recent years and, in consequence, few specially hardened processes exist and even fewer foundry services, typically expensive compared to standard processes. Some of the companies and technologies available are Harris (AVLSI-RA CMOS and UHF-1X bipolar), Honeywell (RICMOS and SOI CMOS), TEMIC (DMILL BiCMOS), Maxim (bipolar).

Several groups have developed radiation tolerant ASICs for LHC. From personal experience, this is a long term task for complex circuits typically requiring ~1year/cycle of design, fabricate and test, with 4-5 iterations/circuit before finalisation. This must be complemented by much testing to a detailed level in close to operational conditions and radiation qualification, so has turned out to be a very long term enterprise.

Deep sub-micron technologies

Results in the last two years have opened a new avenue which looks very promising. Standard commercial CMOS processes now use feature sizes well below 0.5μm and it has been shown that the thin gate oxides have very high intrinsic hardness [6], apparently due to electron tunnelling into the oxide which neutralises radiation induced trapped charge. Combined with design techniques which protect against inter-device leakage, this has been shown to allow unexpectedly high radiation tolerance (fig. 4) [7, 8].

There are some major benefits of this approach, compared to using conventional hardened technologies: significantly lower price wafers, higher circuit density, state of the art processing,… and also some risks: circuits must be user qualified, possible increased SEU sensitivity, and complications from the fact that very large companies with modern foundries and high volume production do not need small customers.

5. Optical links in LHC experiments

Optical links represent another area where active R&D is still approaching completion. The distributed nature of some LHC electronic systems means the technology is essential. Fibre optic links are attractive compared to electrical transmission because of low mass, absence of electrical interference and high bandwidth, in most cases using less power. Fibres are not individually as robust as electrical cables but parts of the systems will use optical cables with reinforced protective sheaths. There are several applications, in addition to the pioneering TTC system already mentioned: digital control at ~40Ms/s, high speed digital data transmission at ~1Gb/s and, in one case, analogue data transmission at 40Ms/s**.**

Optoelectronics is a fast moving technological area driven by digital applications such as telecommunications and local area networks and analogue transmission for cable TV. However, particle physics requirements are rather different since material, power and cost targets are all lower than most commercial components offer, with the additional need for radiation tolerance. For short distance particle physics applications, the major cost drivers are transmitters and optical connectors.

The ATLAS TRT system chose differential electrical transmission over low mass twisted pairs as most cost effective for the relatively small number (27k) of links required. However, for the inner tracker and the pixel detectors, digital optical links are the baseline, initially proposing LED transmitters. LED degradation during irradiation appeared to be risky, particularly for the high radiation pixel volume. Therefore interest was stimulated in Vertical Cavity Surface Emitting Lasers

(VCSELs) which may be an inexpensive alternative once packages have been finalised [9]. Fig.5 shows recent measurements of VCSEL characteristics before and after irradiation [10].

The CMS tracker solution for both data and control system links is to utilise 1300nm edge-emitting III-V lasers mounted on silicon carriers [11,12]. Transmitters, pin diode receivers and fibres and connectors have shown excellent radiation hardness and prototypes demonstrate good characteristics for both analogue and digital data transmission [13, 14]. Bit error rates below 10^{-12} have been achieved with significant margin of power after irradiation [15].

Semiconductor lasers

Semiconductor lasers [16] are based on p-n diodes in a direct band-gap material; forward biasing creates the population inversion necessary for laser action. Wavelength depends on band-gap and is defined using compound semiconductors, e.g. InGaAs, to match the low attenuation transmission windows in optical fibres. Laser operation requires an optical cavity which is constructed in the Fabry-Perot laser by cleaving the crystal to create optical facets to act as partially reflective mirrors; several other types of edge emitting laser exist. Vertical cavity lasers, as the name suggests, are devices which emit, at ~850nm, transverse to the wafer surface.

A minimum current is required for laser action since photons are lost from the cavity by external emission and internal absorption. Values as low as a few mA and forward voltage drops of 1-2V are now available with output power of many mW. Above threshold light output power is usually highly linear with current in edge emitting devices making them suitable for analogue transmission with 8-9 bit ranges such as the CMS tracker; fig.6 shows a typical laser characteristic.

During irradiation traps are created in the laser material which act as non-radiative recombination centres and reduce optical gain. However modern heterostructure quantum well lasers have naturally small active volumes to minimise electrical power and maximise efficiency and irradiation results on many lasers are extremely good.

Low cost semiconductor devices are manufactured in high volumes with large numbers of die per wafer and automatic testing before dicing maximises yield at the packaging stage. Edge emitting lasers do not match this technology but VCSELs avoid cleaving before test which reduces costs. Packaging, where the emitter is accurately aligned to the core of an optical fibre, is an important cost driver and developments have been required for LHC applications. Single mode fibre is more commonly used commercially, and thus less expensive than multi-mode types, but multi-mode connectors require less precision because of the larger fibre core diameter and thus are cheaper than mono-mode connectors. This is another potential benefit of VCSELs. Fibre radiation results are excellent.

Receiver diodes are commercial parts with material depending on operating wavelength: at 1300nm III-V materials are used with adequate radiation tolerance for low bit error rate links, while at 850nm silicon diodes can act as photodetectors.

High speed digital links

Commercial links have been evaluated for the ATLAS ECAL and CMS RPC systems; these include encoding and decoding electronics to provide error detection. One concern is SEU errors, which could require time-consuming resynchronisation and give rise to potential data corruption, and some custom developments are underway. For the CMS ECAL a GaAs CHFET serialiser is being developed which will be coupled to a VCSEL transmitter. A low power design has been developed [17] and the latest version with improved electrostatic discharge protection and encoding protocol is in fabrication [18]. Another recent development is a serialiser in 0.25μm CMOS which works impressively with commercial transmitters [19] and could form the basis for several LHC applications.

6. Tracking systems

Pixels

The detailed description of the pixel system architectures and circuits is such a large task that it will not be attempted, and the reader is referred to more specialised papers, such as those on recent ATLAS developments [20, 21]. Both CMS and ATLAS are prototyping pixel electronics in DMILL and Honeywell SOI processes with a view to selecting one of them for final production. Both experiments report that the radiation hardness requirements of up to about 30Mrad appear to be obtainable. Having succeeded with "proof of principle" prototypes, both experiments are moving on to larger arrays and control electronics. Issues of adequate noise at low power and control of thresholds and cross-talk have been dealt with successfully.

The principal difference between ATLAS and CMS is the choice of analogue or digital readout[22, 23]. ATLAS has adopted binary readout of pixels 50μm x 400 μm in area, while the CMS system is analogue, with 150μm x 150μm pixels, and charge sharing caused by Lorentz bending in the magnetic field will be used to improve the position measurements. In the endcap region, where ExB effects do not occur, the planes of pixels are tilted to give a similar effect.

ALICE is also pursuing a binary readout of 50μm x 400μm pixels and is now prototyping in 0.25μm CMOS, where space saving allows logic reconfiguration so the same chip can be used by the LHCb RICH detector with a larger cell size of 400μm x 400μm. [24]

The small pixels have low detector capacitance and a favourable signal to noise ratio, compared to microstrips, for power consumption ~50μW/pixel. However, low noise occupancy is required for practical pattern recognition. Since particle flux densities are high, detector radiation damage will reduce signal magnitudes following type inversion, so this must be allowed for. Typical threshold are therefore set at around 3000e and this must be accompanied by small dispersion, ~100-200e. Both process variations and radiation damage tend to make this a challenge so all designs implement trimming DACs, typically with 3bit ranges, to adjust thresholds during operation.

Testability is an important issue, with the high channel density and potentially complex data streams. This is another potential advantage of 0.25μm technology, since logic can be reduced in size and suitable test features incorporated on the chip. The ALICE design includes JTAG boundary scan.

Outer tracking systems

In the outer tracking regions the technological choices of CMS and ATLAS are rather different. ATLAS [25] has an inner system (SCT) based on 6M channels of silicon microstrips, where occupancies below 1% are expected, and a 420k channel outer gas detector (TRT) using straw tubes which must to detect transition radiation in addition to direct ionisation to provide electron-hadron separation. In the TRT occupancies range from 10 to 40%. CMS has an entirely microstrip tracker [23] and has recently taken the important decision to use silicon throughout; occupancies of 1-2% are anticipated.

The ATLAS systems are binary; both SCT and TRT make use of a bipolar front end, but CMOS is required for the pipeline and logic. Two alternatives were followed for each system: a two chip solution with a bipolar chip and a CMOS pipeline and logic chip, or a single BiCMOS chip in TEMIC DMILL. Bipolar amplifiers were believed to offer better noise for low power compared to CMOS, but there are overheads from comparators and drivers required in the two chip solution. The BiCMOS solution has the advantage of eliminating a large number of dense bonds which are required for the signal transfer between the comparator and pipeline stages.

The TRT front end (ASDBLR – fig.7) [26] contains an amplifier with tail cancellation and baseline restoration, which is selectable for CF_4 and Xe gas mixtures, followed by a two level discriminator to provide electron identification. It has a peaking time of 7-8ns to minimise pileup and 8 channels/chip. The 16 channel DTMROC chip provides pipeline and timing measurement and reads out 3 beam crossings of data. Power consumption is ~30mW (ASDBLR) and 10mW (DTMROC) per channel.

The SCT requirements are [25] Equivalent Noise Charge (ENC) below 1500e, 99% efficiency, and single bunch crossing tagging. The noise occupancy should be $5x10^{-4}$ or less and double pulse resolution 50ns after a 3.5fC signal. The derandomising buffer is 8 deep and power consumption is to be less than 3.8mW/channel, for the entire front end. The system uses a bipolar amplifier/discriminator and CMOS pipeline-multiplexer with zero suppression in 128 channel chips. The technologies under study for the two chip solution are Maxim for the bipolar Café-P chip [27] and Honeywell bulk CMOS for the digital ABC chip [28]. The designs require $23mm^2$ and 1.5mW/channel (Café-P) and $33mm^2$ and 0.8-1mW /channel (ABC). The alternative is the DMILL BiCMOS ABCD [29, 30] chip which has an area of $53mm^2$ and consumes 2.8mW /channel [31]. Data are transmitted after sparsification over ~8000 optical links.

Both Café +ABC and ABCD systems are being evaluated in pre-production detector modules and irradiated. Some difficulties have been experienced with radiation behaviour of bipolar components but these now appear to be under control [27, 29].

The CMS tracker contains about 11M channels and uses analogue readout with no on-detector zero suppression. This aims to provide robustness and common mode noise immunity as well as avoiding on-chip thresholds. The front end APV chips provide amplification, pulse shaping and analogue pipeline memory, followed by an analogue signal processor which allows two alternative filters to be selected, and 128:1 multiplexer. The ENC should be below 2000e for the entire system including later contributions from optical transmission and digitisation.

CMS has prototyped APV chips in Harris AVLSI-RA bulk CMOS [32, 33]and TEMIC DMILL [34] processes, which were successful but has recently produced an APV25 circuit [35] in 0.25μm bulk CMOS which has performed impressively up to 20Mrad (fig.8). This will substantially reduce the cost of one of the most expensive elements of the readout system and offer improved noise and power performance, as well as longer pipeline and deeper buffers in a smaller chip. The 160 deep pipeline has been extended to 192 and buffer depths extended from 6 events to 10 and the chip has been shrunk from $77mm^2$ to ~$50mm^2$. Power consumption remains below 2.4mW/channel including a new output driver [36]. The 50k channel analogue optical link drivers add only a fraction of a mW/channel.

Comparison of power consumption between ATLAS and CMS tracker ASICs appears to show no major differences, despite earlier beliefs that bipolar technology would be advantageous. Although the APV25 CMOS amplifier requires a large current (~400μA) in the input stage to maximise transconductance and reduce noise, the amplifier-shaper power consumption is limited to 1.15mW and there is no change after irradiation. In comparison a group of ABCD chips on a module assembled in Master-Slave configuration have an average power consumption of 2.8mW/channel with a possible increase following irradiation.

7 Calorimeter systems

The principal challenges for the electromagnetic calorimeter (ECAL) electronics are the large dynamic range of signals, from 50MeV-2TeV (92dB) requiring 15-16bits, which must be measured with a precision of ≈12bits with excellent stability. The electromagnetic calorimeter performances are usually parameterised by contributions summed in quadrature

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus b \oplus \frac{\sigma_{noise}}{E}$$

representing the stochastic contributions, systematic effects and electronic noise. The electronic noise is constant and could eventually dominate for a high precision system like CMS if care is not taken. Calibration and monitoring are vital to minimise systematic effects (b ~0.5-1%).

The CMS ECAL [17] is constructed from lead tungstate crystals read out by avalanche photodiodes (APD) in the barrel region and vacuum phototriodes (VPT) in the forward region where radiation levels are higher. Scintillation signals are fast, with a time constant of ~10ns. A total of ~80k channels is planned. The worst radiation doses and fluences are expected in the forward regions where

doses of up to 1Mrad are anticipated, with additional high neutron fluxes; in the barrel region several hundred krad doses are experienced. These necessitate radiation hard technologies for ASICs.

The CMS ECAL readout (fig.9) uses a preamplifier [37], now constructed using the Harris UHF bipolar process, close to the photo-detector behind each crystal. It provides several different gain outputs, one of which is selected by the following FPU chip [38] for digitisation in a commercial 40MHz 12bit ADC. The FPU chip provides 2 bits of range information which are combined with the 12bits of amplitude data for transfer to the external data acquisition after encoding on a custom high speed optical link.

The ADC is a commercial part (AD9042) developed for radiation hardness by the USAF in a complementary bipolar process. CMS will follow a stringent procedure to ensure purchases conform to their hardness requirements: from each lot of 5000 parts, 50 will be removed for irradiation with protons. If these are demonstrated to be of good quality, the remaining die will be used for construction, otherwise the entire lot will be rejected.

The ATLAS Liquid Argon sampling calorimeter [39] will be used for both electromagnetic and hadronic calorimetry and is one of the furthest advanced systems under construction. It has 190k channels and common electronics outside the cold volume is used everywhere [40], to avoid heat into the liquid and minimise radiation exposure, except for the very forward hadron system which employs cold GaAs preamplifiers [41]. The front end system consumes ~250kW of power.

The ionisation signals from the argon give a triangular current pulse of ~450ns duration due to the drift time in the liquid. This requires careful shaping (fig.10) to minimise pileup which, because of its random nature, effectively contributes another, luminosity dependent, source of noise. Following multi-gain amplification and shaping, signals are stored in a 144-deep analogue pipeline memory awaiting a trigger decision, before being digitised and transferred to the remote DAQ by 1500 high speed optical links. A local trigger sum is formed from the analogue signals and transferred separately by coaxial cable to the ATLAS trigger cavern 50m away.

The amplifiers, which are constructed as thick film hybrids using discrete bipolar components [42], are coupled to the detector by transmission lines. The hybrid technology was chosen in preference to a custom IC because of lower development costs and lower risk for performance and schedule. Amplifiers have been irradiated with gammas and neutrons to demonstrate that degradation is well within stringent specifications.

A large number of amplifier channels have been constructed and evaluated, including 1 week burn-in tests to identify possible failure mechanisms. Although few failures were found, it was concluded that to weed out all weak components burn-in at 70°C should be used for production. This was accompanied by an analysis following US DoD standards to estimate mean time between failures, which illustrates well the importance of such procedures for electronics which will be largely inaccessible for many years of operation.

8. Conclusions

In most cases, LHC systems are just beginning construction and there are many issues still to be faced before the results of the developments can be properly judged. The last few years have seen much progress but there is still some way to go before LHC switch-on. In particular, there is much important work to be done on issues like power supplies, assembly, services and other practical issues of implementation which must be completed to ensure successful long term operation at LHC.

Acknowledgements

Much of the information presented here has relied heavily on other sources, particularly ATLAS and CMS Technical Design Reports and the annual LHC Electronics Board conference proceedings, which should be consulted for many details. Errors are inevitable in such a review and I apologise in advance for any inaccuracies.

References

LEB 99 = Proc. LHC Electronics Board Workshop. CERN Report CERN 99-09(1999)

LEB 98 = Proc. LHC Electronics Board Workshop. CERN Report CERN 98-36 (1998)

LEB 97 = Proc. LHC Electronics Board Workshop. CERN Report CERN 97-60 (1997)

[1] B. Taylor et al. RD12 Status report. CERN/LHCC 2000-002 (2000).

[2] H. Øverås. Nucl. Instr. and Meths. 104 (1972) 85.

[3] N. Marinelli. *APV logic simulations* CMS Note 1999/028 (1999)

[4] M. Huhtinen & F. Faccio. *Computational method to estimate Single Event Upset rates in the LHC environment.* To be published in Nucl. Instr. & Meths.

[5] F. Faccio et al. LEB 99, 571.

[6] N.S. Saks et al., IEEE Trans. Nucl. Sci., NS-31 (1984) 1249

[7] P. Jarron & A. Marchioro. LEB 97, 139.

[8] A. Marchioro. LEB 98, 40.

[9] M. Pearce. These proceedings.

[10] D. Charlton et al. LEB 99, 185.

[11] V. Arbet-Engels et al. LEB 97, 287.

[12] F. Vasey et al. LEB 99, 175.

[13] K. Gill et al. LEB 97, 277.

[14] J. Batten et al. Proc. 1998 Conf. on Radiation and its Effects on Components and Systems, 75-78.

[15] F. Vasey. LEB 98, 344

[16] G. P. Agrawal. *Fiber optic communication systems*. J. Wiley (New York) 1992.

[17] CMS Electromagnetic Calorimeter Detector Technical Design Report. CERN/LHCC 97-33 (1997)

[18] P. Denes. Private communication.

[19] P. Moreira. LEB 99, 194.

[20] K. Einsweiler. These proceedings.

[21] J. Richardson. LEB 99, 83.

[22] ATLAS Pixel Detector Technical Design Report, CERN/LHCC 98-13 (1998)

[23] CMS Tracker Detector Technical Design Report, CERN/LHCC 98-6 (1998)

[24] K. Wyllie et al. LEB 99, 93.

[25] ATLAS Inner Detector Technical Design Report, CERN/LHCC 97-17 (1997)

[26] B. Bevensee et al. LEB 99, 128.

[27] T. Dubbs et al. LEB 99, 123.

[28] D. Campbell et al. LEB 99, 152.

[29] F. Anghinolfi. LEB 99, 118.

[30] W. Dabrowski et al. LEB 99, 113.

[31] F. Anghinolfi. Private communication.

[32] M. Raymond et al. LEB 97, 158.

[33] L. Jones et al. LEB 98, 185.

[34] U. Goerlach. LEB 98, 108.

[35] L. Jones et al. LEB 99, 162.

[36] M. Raymond. Private communication.

[37] J-P. Walder. LEB 98, 218.

[38] P. Denes. LEB 98, 223.

[39] ATLAS Liquid Argon Technical Design Report, CERN/LHCC 96-41 (1996)

[40] J. Colas. LEB 99, 217.

[41] J. Ban et al. LEB 99, 227.

[42] M. Citterio LEB 99, 237.

Figure captions

1. Two variants of generic LHC front-end readout systems.

2. (a) APV6 pipeline data loss for various latencies. Pipeline length is 160 and buffer depth is 6 triggers (18 samples) [3].

2. (b) APV25 pipeline data loss as a function of trigger latency. Pipeline length is 192 and buffer depth is 10 triggers (30 samples) [3].

3. SEU cross-section as a function of Linear Energy Transfer [5].

4. Threshold voltage shifts measured on 0.25μm transistors [7].

5. VCSEL characteristics measured before and after irradiation with $3x10^{15}$ $n.cm^{-2}$ [10].

6. Edge emitting laser characteristic before and after irradiation [13].

7. ATLAS TRT ASDBLR chip schematic [26].

8. APV25 pulse shapes before and after irradiation [36].

9. CMS ECAL front end readout [17].

10. Simulated pulse shapes in the ATLAS liquid argon ECAL [39].

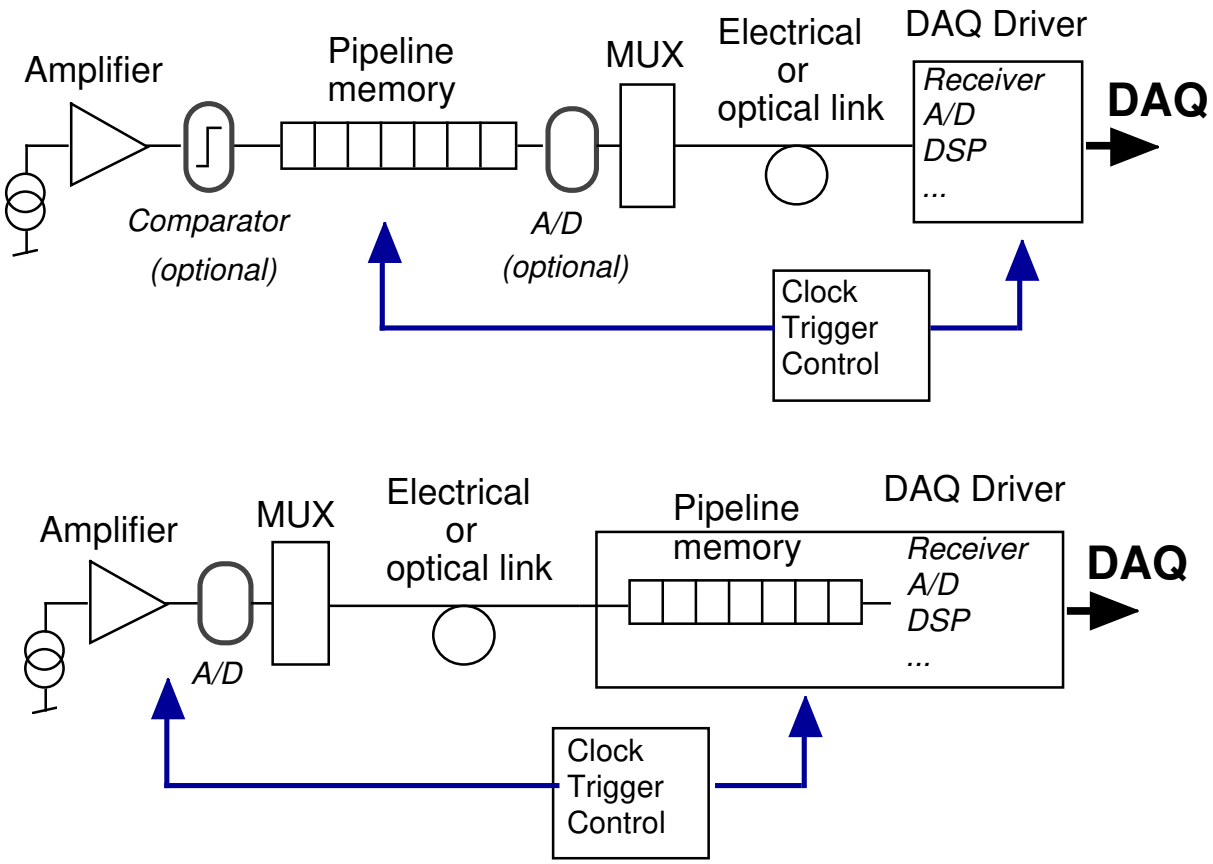

Fig. 1. Two variants of generic LHC front-end readout systems.

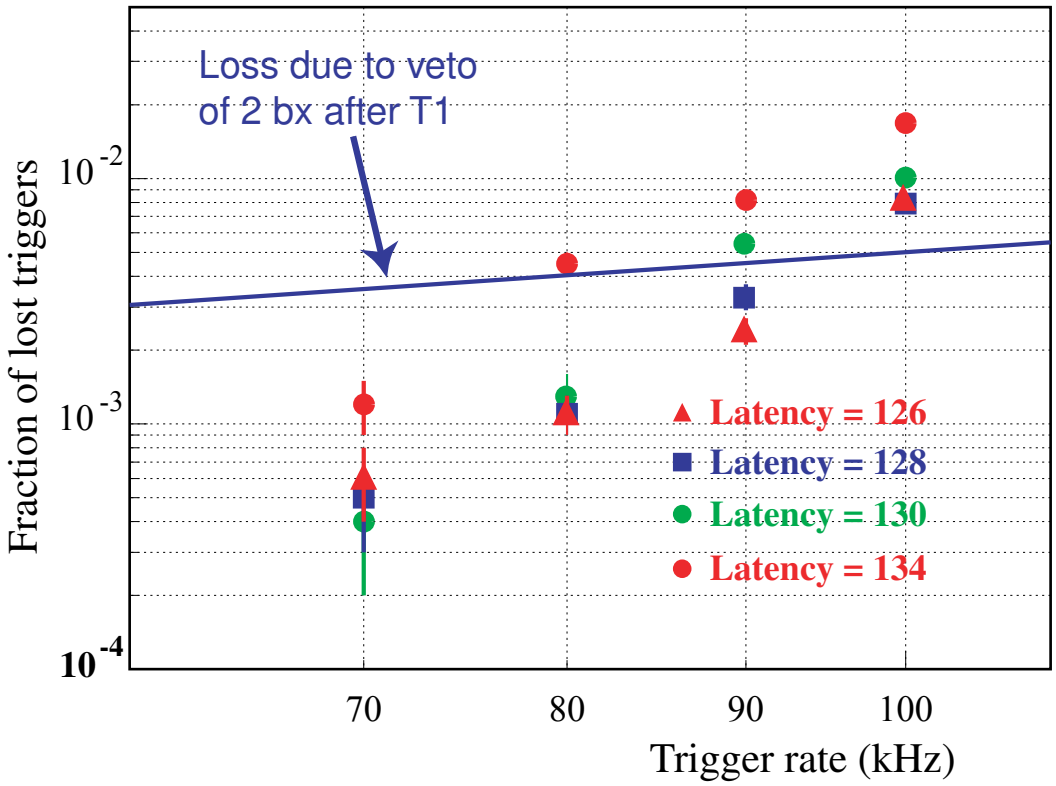

Fig. 2. (a) APV6 pipeline data loss for various latencies. Pipeline length is 160 and buffer depth is 6 triggers (18 samples) [3].

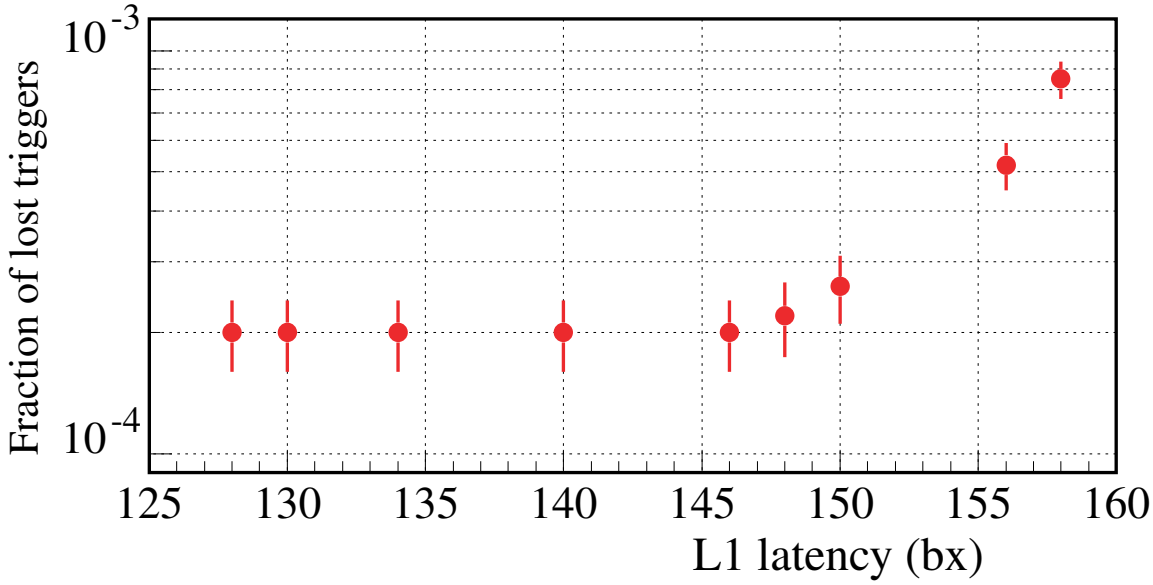

Fig. 2. (b) APV25 pipeline data loss as a function of trigger latency. Pipeline length is 192 and buffer depth is 10 triggers (30 samples) [3].

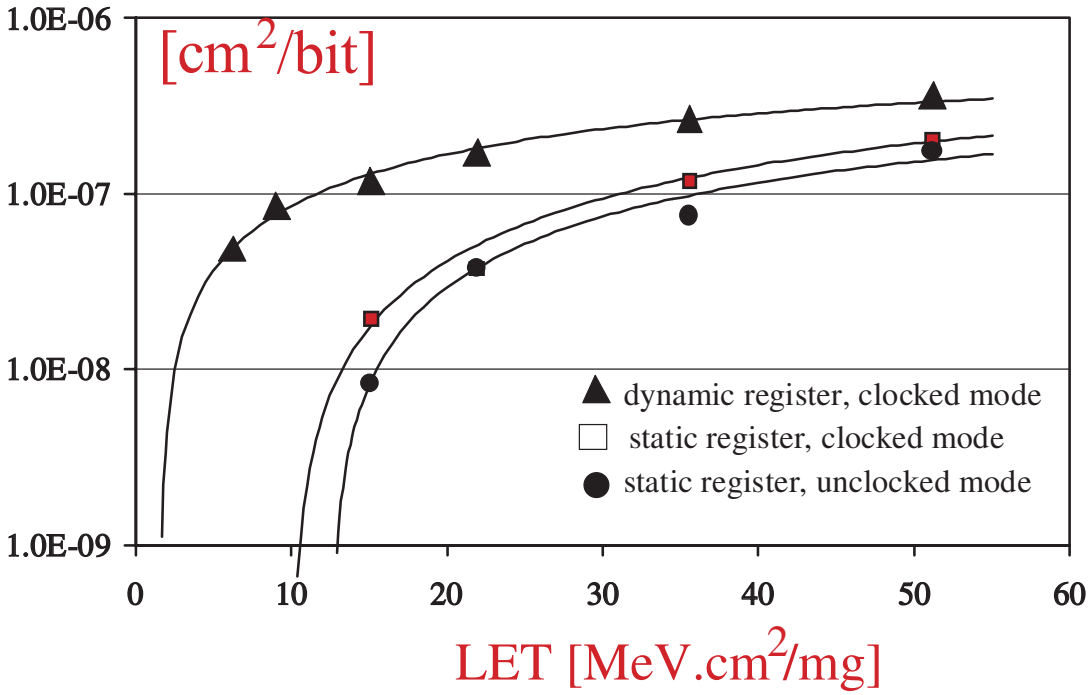

Fig. 3. SEU cross-section as a function of Linear Energy Transfer [5].

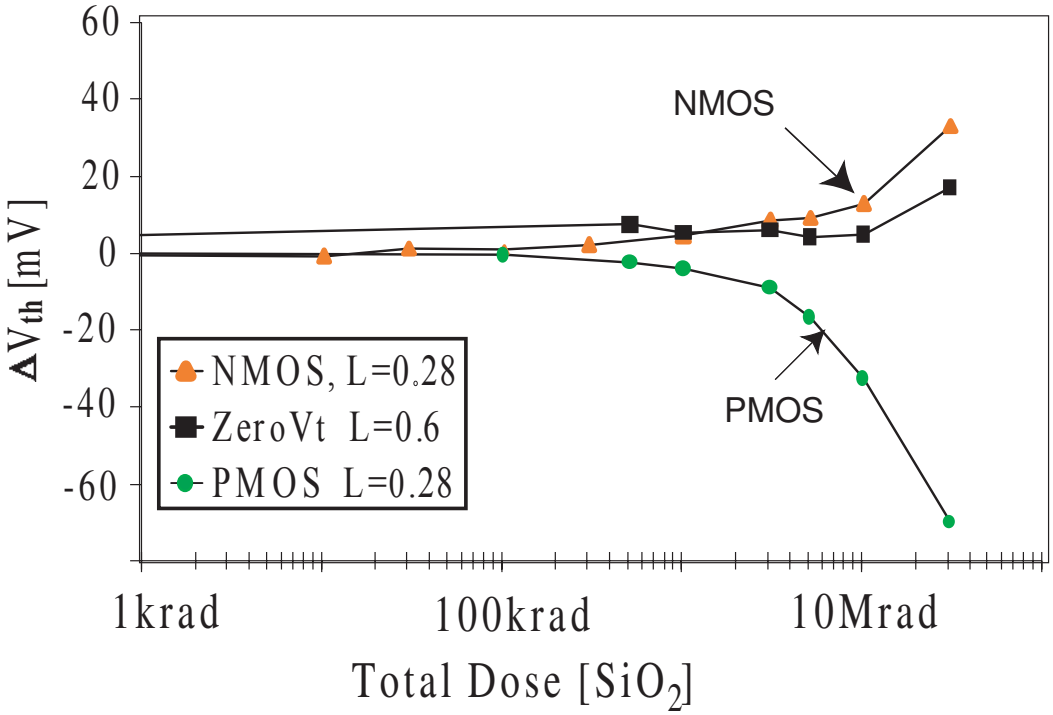

Fig. 4. Threshold voltage shifts measured on 0.25µm transistors [7].

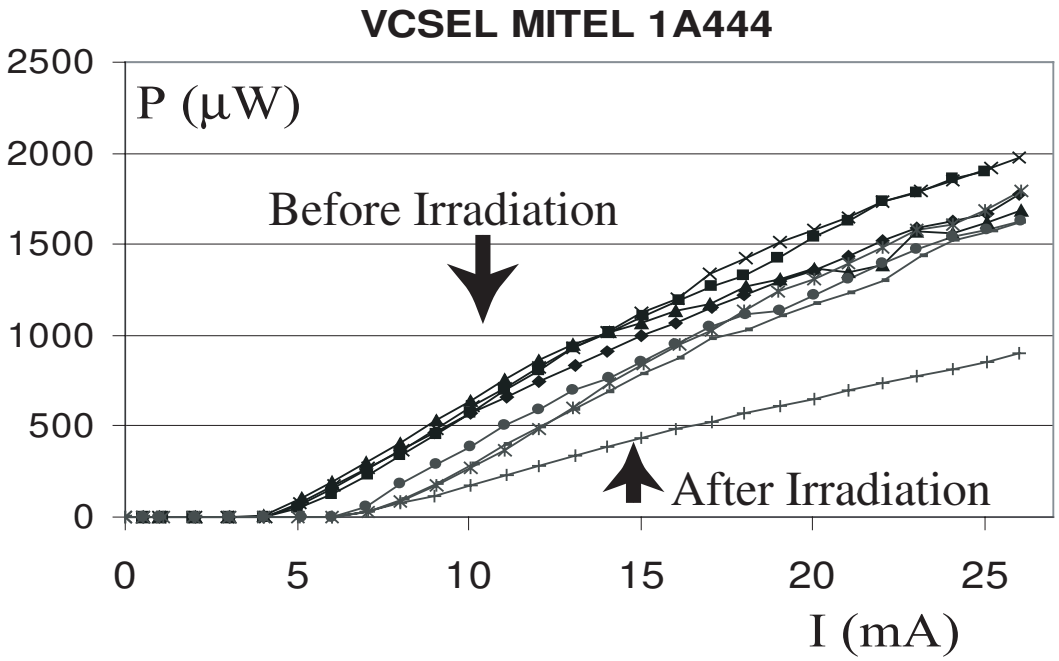

Fig. 5. VCSEL characteristics measured before and after irradiation with $3x10^{15}$ n.cm^{-2} [10].

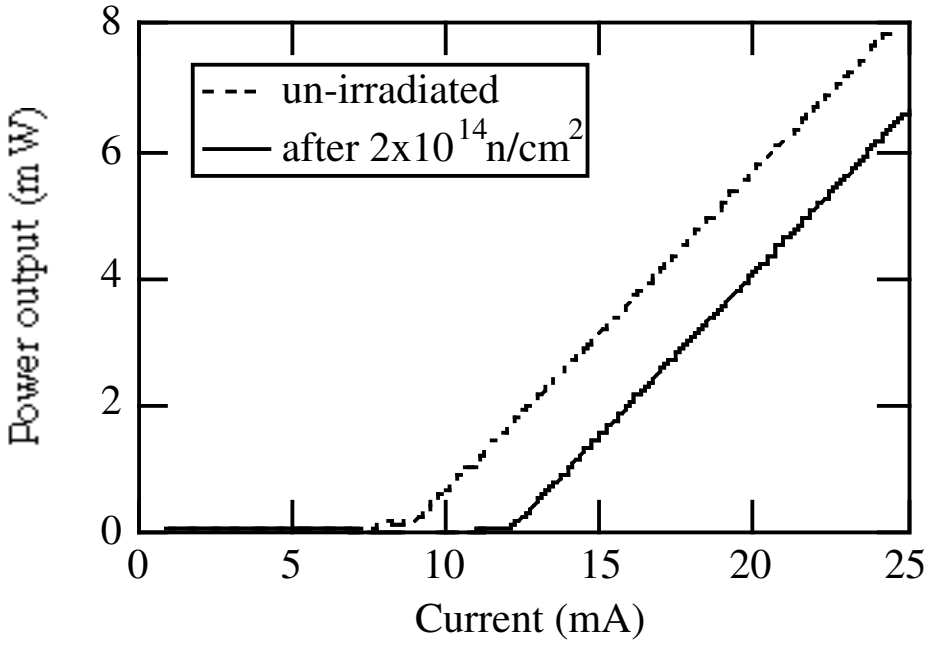

Fig. 6. Edge emitting laser characteristic before and after irradiation [13].

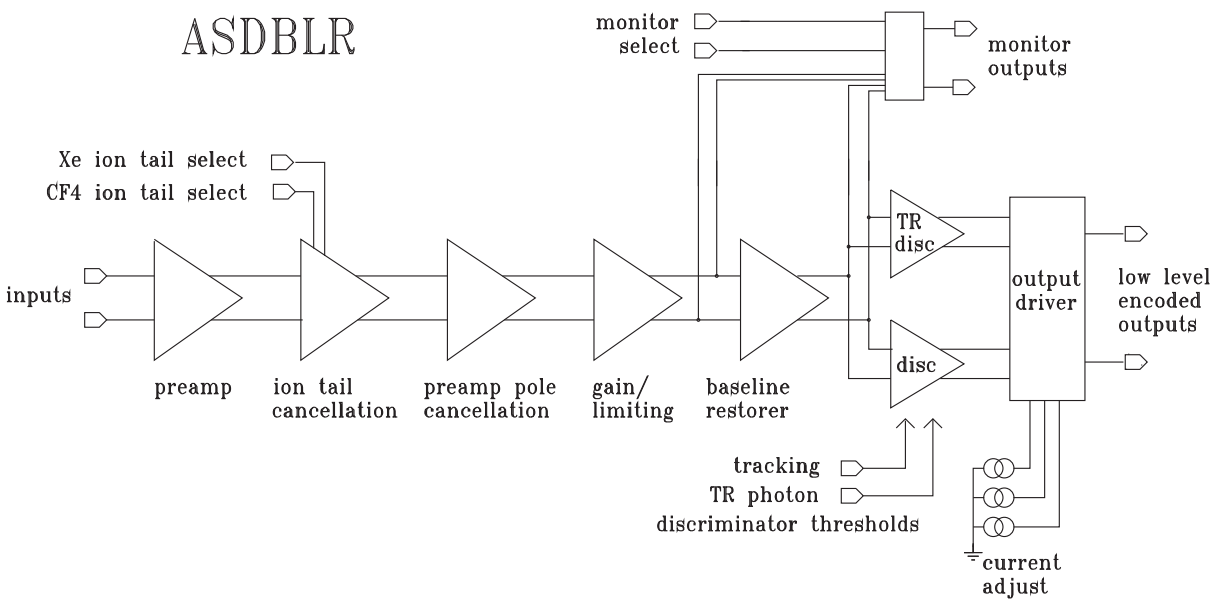

Fig. 7. ATLAS TRT ASDBLR chip schematic [26].

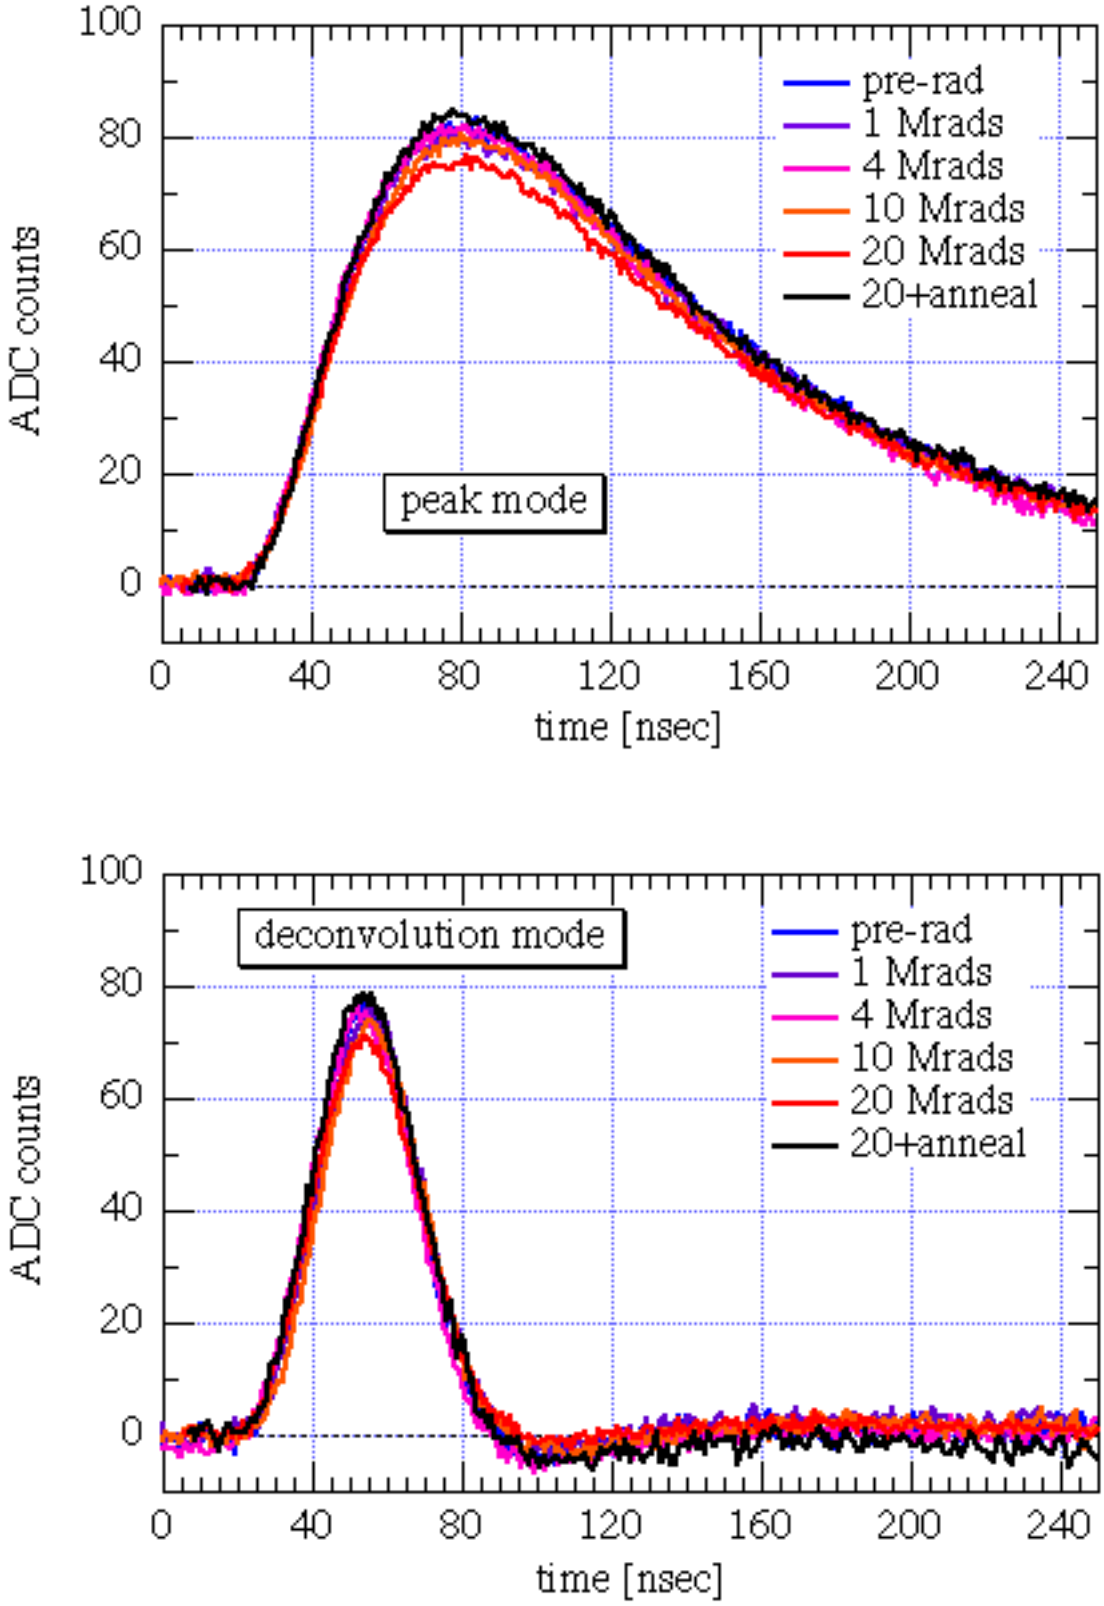

Fig. 8. APV25 pulse shapes before and after irradiation [36].

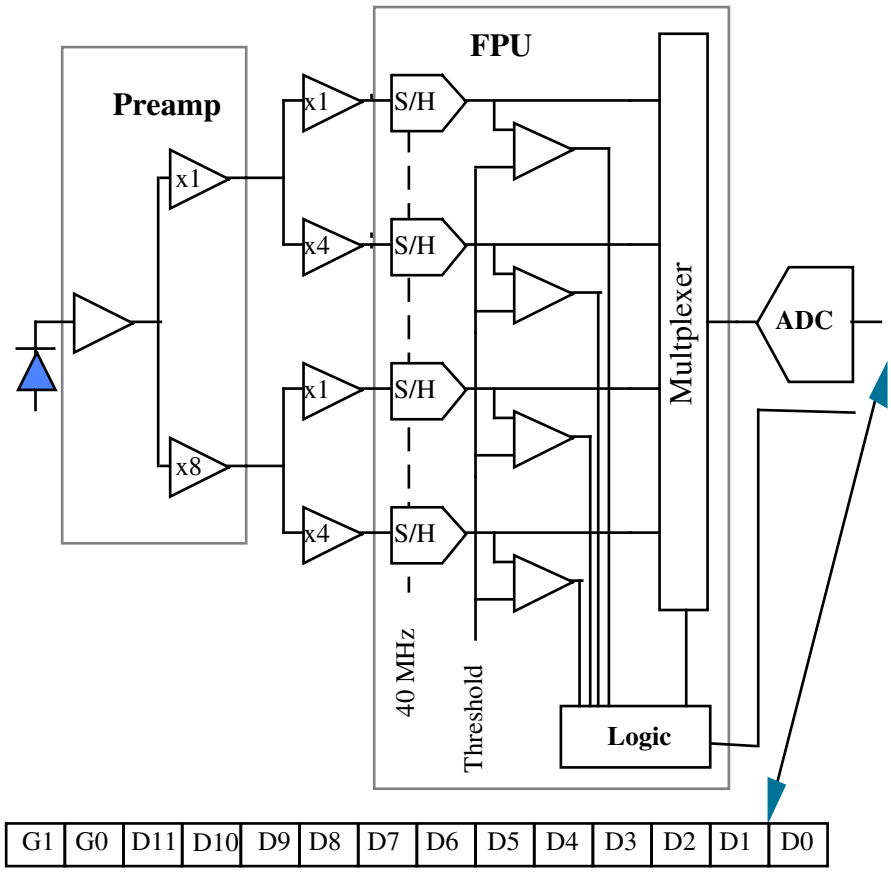

Fig. 9. CMS ECAL front end readout [17].

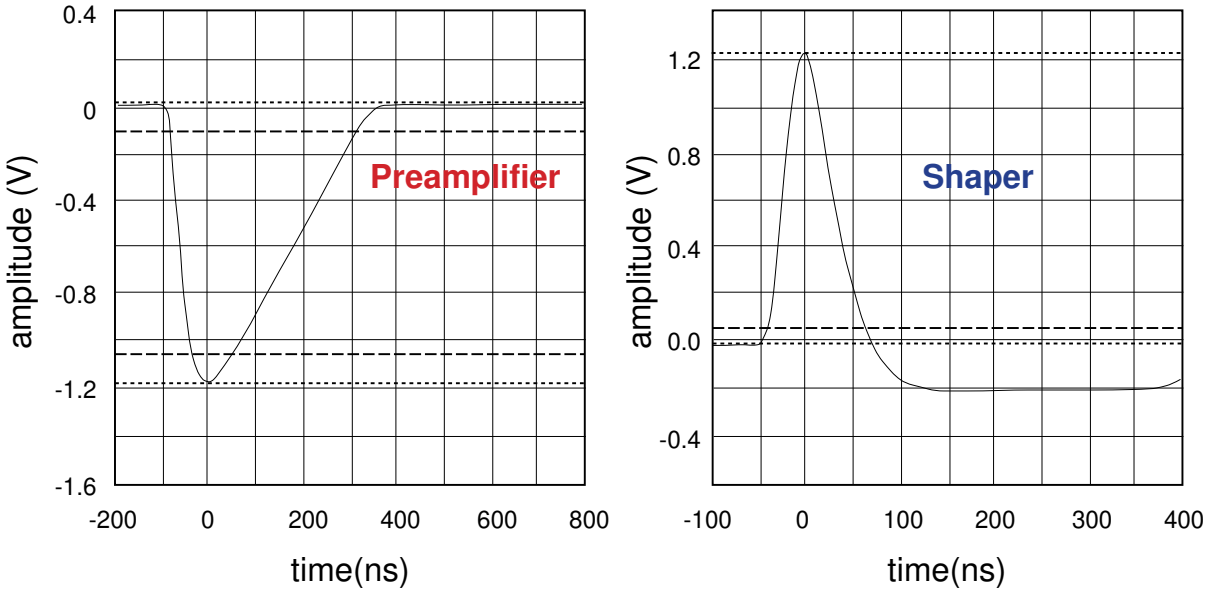

Fig. 10. Simulated pulse shapes in the ATLAS liquid argon ECAL [39].