FRONT END ELECTRONICS DEVELOPMENTS FOR PARTICLE PHYSICS: PARADIGM OR PARADOX?

G. Hall

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

g.hall@ic.ac.uk

Abstract

Recent years have seen much dedicated work on front end electronics for hadron colliders, with an strong emphasis on radiation hardness and low cost. This has been challenging for a number of reasons, some of which are discussed further below. The developments also suggest opportunities and constraints for the development of such electronics in future

1. Introduction

Reviewing progress in front end electronics is intimidating because it covers a wide range of components, techniques and technologies. Some recent developments, particularly those for the Large Hadron Collider, have been under way for almost a decade and are still not complete, which gives an indication of the magnitude of some challenges. It is interesting to ask why this is so.

It is sometimes valuable to look backwards as well as forwards. During the last decade many different technologies have been investigated, and some have been discarded while others have perhaps matured faster than expected. This article surveys some selected areas where a lot of attention has been focussed and offers comments on what has been achieved and where some developments may go in future.

Do developments focussed mainly on hadron collider operation offer any relevant lessons for other experiments in particle physics? Similar systems are now also often applied in other fields of science, so lessons learned from particle physics may be relevant elsewhere.

This paper addresses general issues, without providing details of solutions; many of these can be found in [1-4].

2. Hadron collider electronic requirements

Since most developments presently underway can be traced back to the special needs for electronics for the SSC and subsequently the LHC, the most relevant LHC operational parameters as seen by one general purpose detector, CMS, are summarised in Table 1. There are only minor differences for ATLAS, except that heavy ion operation is not planned.

Almost completely hermetic general purpose detectors operated with high efficiency to detect rare events were recognised to be essential for TeV physics, and it was understood very early that radiation tolerance would be one of the biggest problems, supplemented by very high readout rates. The innermost regions of the experiments are subject to high particle fluxes and thus radiation levels. The large majority of the sensor technologies in use for particle physics in the 1980s were known either to be problematic for application in tracking systems and electromagnetic calorimeters, e.g. silicon, crystalline and plastic scintillators, or completely excluded, e.g. gas detectors such as TPCs or large volume drift chambers.

These challenges motivated dedicated R&D programmes; most emphasised sensors, although several projects recognised that compatible electronics would have to be developed in parallel to match operational conditions.

Table 1		
particles	pp	Pb-Pb
Luminosity	10^{34} cm^{-2}.s^{-1}	10^{27} cm^{-2}.s^{-1}
Average integrated luminosity	5x10^{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/crossing	~140	3000-8000
beam crossing rate	40MHz	8MHz
L1 trigger delay	≈3.2µsec	≈3.2µsec
L1 trigger rate	≤100kHz	< 8kHz

The LHC experiments are enormous in comparison with previous generations and most sub-systems have huge numbers of readout channels. The sub-detector systems are highly distributed and remote and inaccessible for long periods, demanding high reliability. The very large data volumes and minimal deadtime requirement require fast front end data links and efficient processing. Many secondary specifications are *extremely* demanding, such as temperature stability, magnetic field operation, radiation tolerance, and material budget reduction.

High speed signal processing has major implications for front end power, signal to noise ratios and overall performance. The high charged particle and neutron fluences must be prevented from degrading performance but also raise issues of qualification, reliability and maintenance. On-line reduction of first level data is mandatory.

Radiation concerns apply throughout the experiment caverns where, even at the walls, low ionising doses but significant neutron fluences are expected. These are a concern for support electronics, like power supplies, or computing and control components. High magnetic fields are used and significant stray fields are anticipated.

3 Requirements for different detectors

Colliding beam experiments are all constructed in a generally similar way and it may be helpful to identify the main parameters defining the various electronic systems. Descriptions like “low power” and “low radiation levels” can mean quite different things in different sub-detectors.

For the benefit of the general reader, tracking systems surround the beam pipe to measure charged particle trajectories, calorimeter systems follow, with high density, large volume absorbers as detectors, and, finally, detectors on the outside measure the locations of penetrating muons. The surface area of each system therefore usually grows with distance from the beam but the density of electronics usually decreases.

Requirements for tracking are high spatial precision on particle trajectories but limited energy precision. Detectors are highly granular, leading to enormous channel counts. The dynamic range of signals is limited to a few bits since highly linear responses are not 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 high priority. This must be achieved using a few mW/channel with fast pulse shaping and in multi-Mrad radiation levels because of proximity to the beam. Power constraints are imposed by the need for a low material budget to minimise photon conversions and electron bremsstrahlung, so cables and cooling mass must be minimised.

Calorimeter detectors aim for very high energy resolution over a very large dynamic range; with very stable performance over time. The detector absorbs most hadrons so usually shields the electronics and radiation levels up to ~1Mrad are typical. 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 are high when integrated over the large areas over which signals are sensed. Radiation levels are low but heavy reliance on programmable logic implemented in standard commercial parts raises difficult questions of qualification and maintenance.

Generic LHC readout system

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 widely utilised but they appear in different forms, either analogue or digital, and in different locations, either on or off the front end. In all cases data must be associated 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”.

Two variants of a typical LHC readout system are illustrated in fig.1.; other variants are possible. They demonstrate that many LHC systems are rather complex. Not only do front end chips embody amplification and bandwidth limiting filters but many of them include digital control logic and, often, other analogue or digital processing functions.

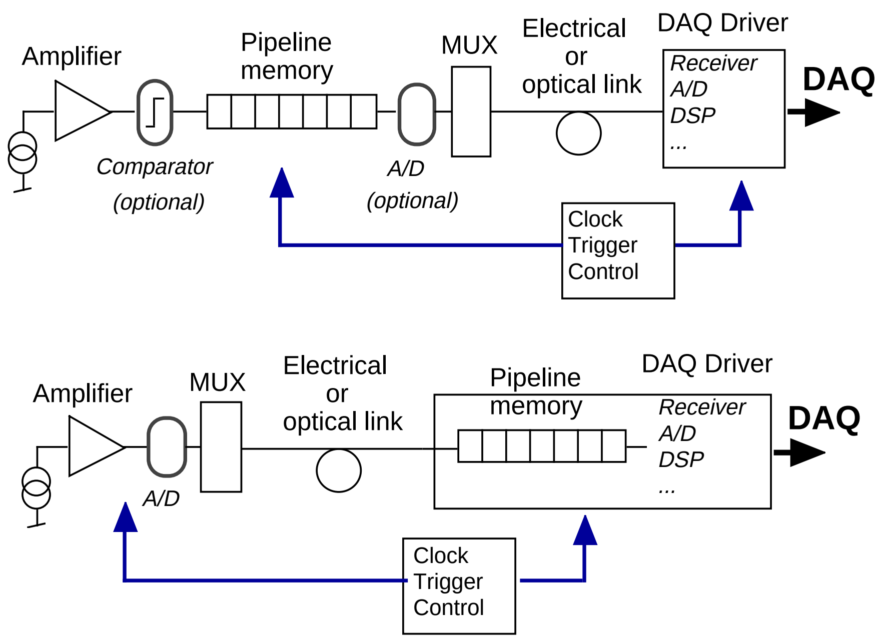

Fig. 1. Two variants of generic LHC front-end readout systems.

Systems should be able to handle incoming signals while processing data from previous events at up to the maximum trigger rate. This is provided by circular pipeline architectures which sample signals at the machine clock frequency and pointers circulate, marking and skipping cells which contain data awaiting readout. This is a classical queueing problem but logic emulation is needed to compute losses in detail. Poisson fluctuations in trigger arrival mean that practical buffer depths cannot guarantee unlimited operation without overflows; in practice data losses of ~1% are tolerated.

First level triggers are generated by fast readout of limited precision data from calorimeter and muon systems. This is the first level of data rejection. Subsequently increasingly selective rejection of “uninteresting” events will take place using more precise data from all sub-systems so that eventually data can be written to permanent storage devices at the maximum permitted rate of ~100Hz.

4. Radiation tolerance of electronics

Radiation hardness is one reason for lengthy LHC electronic developments, since at the outset few people had experience of radiation hard electronic technologies. This applied also to detectors, but much detector manufacturing was done in-house (gas detectors) or with relatively small, specialised companies (silicon or scintillators) who were willing to prototype and even modify existing technologies at modest cost. This was much less possible for electronics.

For convenience, one can classify electronics as COTS (Commercial Off The Shelf) components, including programmable logic (FPGAs), or Application Specific Integrated 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 atomic displacement damage to substrates.

CMOS circuits are neutron hard but ionising dose sensitive. Charge created by ionisation builds up in the oxide near to the silicon-oxide interface and induces shallow surface charge layers or creates interface states which influence device operation. Typical consequences 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 a few krad.

A further phenomenon in logic circuits, 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 a sensitive circuit node and cause a change of logic state. Even in low flux conditions, operational errors in control circuits could magnify the importance of SEU.

Bipolar circuits are affected by atomic displacement in the body of the device which increases carrier recombination in the base of transistors giving rise to gain and matching degradation. 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, as well as hadrons from beam collisions. Discrete bipolar circuits usually withstand high exposures, and the few integrated circuit processes in use for particle physics provide highly tolerant chips 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. In part, this is because the materials are not usually such perfect crystals as silicon but devices such as MESFETs are not too sensitive to the details of bulk imperfections, so some damage is tolerable. In optical devices, such as semiconductor lasers, the very small active volume of modern components protects against big changes in performance.

Junction FETs are also relatively hard against ionising damage. These are often interesting for input devices of amplifying circuits because of potentially very low noise, since carrier conduction is not close to the interface as in MOS devices, but usually where long time constants can be used. However JFETs are not available in modern commercial integrated circuit technologies, where CMOS dominates; three quarters of the world market is for CMOS products [5].

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 companies and technologies available are Harris (AVLSI-RA CMOS and UHF-1X bipolar), Honeywell (RICMOS and SOI CMOS), TEMIC (DMILL BiCMOS), Maxim (bipolar). However, almost none of them are formally qualified to multi-Mrad levels, where export regulations can bring problems.

Because of the decline of military markets, there has been a big change of focus for commercial manufacturers who now look to space applications as one of the growth areas. However, foundry investments are now so huge - ~$2 billion for state of the art modern facilities – that only the largest markets justify such investments. In addition, some markets, such as for memory, are highly competitive with modest profit margins, so emphasis is on high volume production with long term contracts, and careful scheduling so the foundries are never idle. Particle physics, and other areas of science, do not come close to high volume. The largest LHC project requirements are equivalent to only a few days production for a single manufacturer.

Deep sub-micron technologies

Despite the small volumes of particle physics circuits, there is a motivation to attempt to follow commercial trends - cost. The semiconductor industry continues to deliver larger chips at lower prices and to prevent this trend from declining, although it becomes ever more difficult. The well known "Moore's law" predicts that minimum feature sizes should halve every 6 years and the low cost of consumer products is evidence of the industry's success.

Standard commercial CMOS processes now use feature sizes well below 0.5µm, with 0.25µm and 0.18µm processes in large scale operation. Wafer costs increase with time only slightly despite their larger area; 200mm diameter is the present norm, so even large chips are produced at lower price, once development costs have been amortised. As the feature size decreases, so does the gate oxide thickness. This is associated with lower voltage operation, reducing power, but it has had another important consequence, which is improved immunity to radiation damage.

It was well known that immunity to oxide damage increases rapidly with decreasing oxide thickness, roughly as the square of the thickness. However very thin gate oxides typical of deep sub-micron processes perform even better than this scaling would predict and have very high intrinsic hardness, apparently due to hole tunnelling from the oxide which reduces radiation induced trapped charge[6]. Combined with design techniques which protect against inter-device leakage, this has been shown to allow radiation tolerance in the multi-Mrad range [7] (figs. 2 & 3).

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 thus high yield and quality. There are 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 encourage small customers.

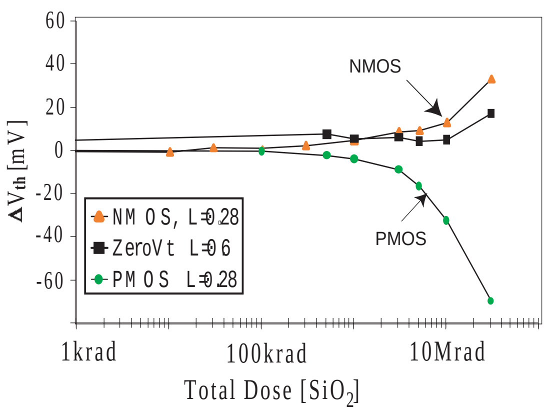

Fig. 2. Threshold voltage shifts measured on 0.25µm transistors [7].

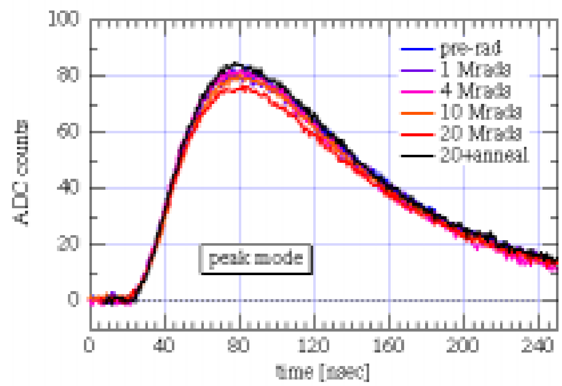

Fig. 3. APV25, for the CMS tracker, pulse shapes before and after irradiation [8].

5. Optical links in LHC experiments

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. The distributed nature of LHC electronic systems means the technology is essential for some sub-detectors. R&D in the last few years with industry has begun to deliver components matching particle physics applications.

Optoelectronics is a fast moving area driven by telecommunications, local area networks (digital0 and cable TV transmission (analogue). However, particle physics requirements on 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.

There are several applications: the LHC machine clock and triggers, as well as control commands, will be distributed using a dedicated system used by all experiments, then custom systems matching the front end electronics for digital control distribution at ~40Ms/s, high speed digital data transmission at ~1Gb/s and analogue data transmission at 40Ms/s.

Fibres are not individually as robust as electrical cables but optical fibre-ribbon cables with reinforced protective sheaths can be used, although customisation, e.g. to avoid magnetic material, has been required. For short links (~100m) used for particle physics, signal attenuation is not a problem, unless fibre or connector radiation damage were to be large, and this has been measured. At the telecommunications wavelength of 1.3µm, used for the CMS tracker analogue data transmission, standard single-mode fibre will be adequate, and digital requirements at all wavelengths are of less concern.

A number of different transmitters have been investigated for LHC, including various optical modulators, where light output can be varied by electrical signals. However, most progress has been made using components closer to the commercial mainstream.

Light Emitting Diodes are not suitable for analogue data transmission due to their non-linearity but appeared possible for digital links. However LED degradation is significant in high radiation regions. Meanwhile, semiconductor laser diodes began to become available in larger quantities offering many features required: small size, compact packages (in some cases, with custom developments), low power and good radiation resistance.

5.1 Semiconductor lasers

Semiconductor lasers are based on p-n diodes in a direct band-gap material, under forward bias to create 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 at 1.3µm and 1.5µm. 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. Vertical Cavity Surface Emitting Lasers (VCSELs) emit, at ~850nm, transverse to the wafer surface, with the cavity constructed epitaxially by multiple dielectric interfaces.

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.4 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 (i.e. engineered multi-element materials, such as InGaAsP) quantum well lasers naturally have small active volumes to minimise electrical power and maximise efficiency. Irradiation results on most lasers are extremely good.

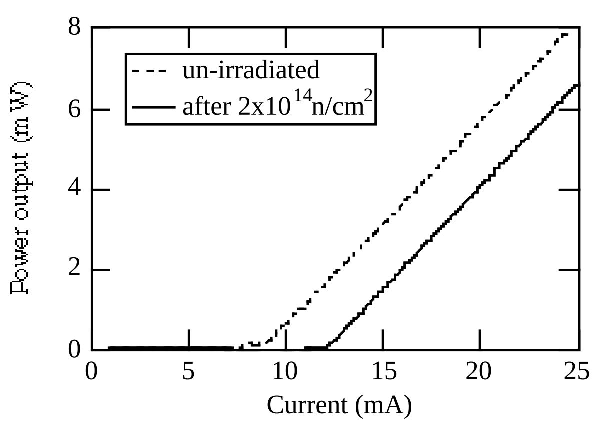

Fig. 4. Edge emitting laser characteristic before and after irradiation [9].

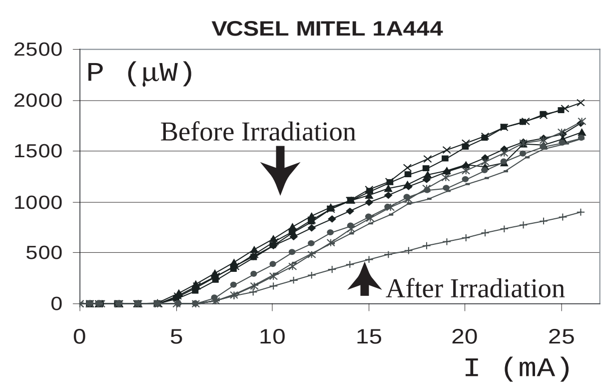

Fig. 5. VCSEL characteristics measured before and after irradiation with $3x10^{15}$ n.cm^{-2} [10].

Low cost is achieved in semiconductor manufacturing by high volume production with large numbers of die per wafer. Although optoelectronic wafers are smaller, so are die, and the same methods apply. Automatic testing before dicing maximises yield at the packaging stage, which is essential for cheap consumer products. Edge emitting lasers do not match this requirement, which is one barrier to increased use of optoelectronic technologies. VCSELs avoid cleaving before test, which reduces costs, and explains much of the interest in them. They do not at present have the linearity which would make them suitable for many analogue applications. This may change, but in any case most applications are digital.

5.2 Packaging

Packaging, where the emitter is accurately aligned to the core of an optical fibre, is an important cost driver and developments have been required to meet the requirements for LHC. Usually in experiments space is at a premium whereas commercial applications are often more concerned about robustness and the ability to exchange a single component without delicate handling. Single mode fibre is used for commercial long distance links, 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.

Receiver diodes are commercial parts constructed using material matched to operating wavelength. Only for diodes located inside the detectors is radiation damage a concern; at 1300nm III-V materials are used and low bit error rates are demonstrated, while at 850nm silicon diodes can act as photodetectors, probably with adequate radiation tolerance

5.3 High speed digital links

For digital links, data should be transmitted with the ability to confirm data quality for asynchronous transmission. For systems like LHC where a well defined clock is available everywhere some saving in complexity can in principle be achieved by reducing error correction overheads. However, loss of synchronisation between transmitter and receiver caused, for example, by SEU errors could then cause severe problems of data corruption.

Commercial links which include error detection encoding and decoding electronics have been evaluated for ATLAS and CMS and some custom developments are under way. For the CMS ECAL a GaAs CHFET serialiser is being developed to be coupled to a VCSEL transmitter while in CERN a 0.25µm CMOS serialiser has worked impressively with commercial transmitters and could form the basis for several LHC applications.

Technology trends and constraints

From work in the last few years, some trends can be discerned.

6.1 Industrial developments

Many developments are dependent on co-operation with industry, with companies who are very large compared to those traditionally involved in sensor or mechanical engineering projects. Entry costs for electronic technologies are usually very high and commercial interest is frequently small for customers with distant projects and modest final orders. Some manufacturers were initially attracted by potential large LHC purchases but many quickly realised that particle physics projects were demanding and unconventional, perhaps offering little profit.

Nevertheless, access to modern technologies has been achieved, sometimes based on individual contacts, some of which have been enhanced by the opportunity to carry out funded R&D for future projects. However, the risk of dependence on a single vendor can be increased, and commercial strategies and technologies do sometimes change at very short notice.

6.2 Custom circuits

Development of complex ASICs is a long term task typically requiring ~1year/cycle of design, fabricate and test, with 4-5 iterations/circuit before finalisation. This must be complemented by much evaluation to a detailed level in close to operational conditions, and radiation qualification, so has turned out to be a very long term enterprise [1,11]. Testing, meaning much more than mere demonstration of functionality, is often a weak point in circuit development [12]. In most detectors the electronic readout systems, including control logic, are technically much more complex than sensors in which, in contrast, there are often heavier investments.

Developments are expensive, with typical ASIC foundry costs of ~$150k for a few wafer engineering run, but usually to reach this point some expectation of sizeable future orders is needed. Most particle physics projects could be satisfied by a small number of wafers so, in purely commercial terms, ASIC development is unjustified and this trend is increasing. Multi-project runs have been utilised successfully but are not usually a service offered by foundries. Small projects are dependent on sharing costs, and time-scales, with larger projects, which may not always be feasible.

The trend to finer feature sizes will continue to increase entry costs and technical difficulties, as each new technology must be characterised for analogue performance. Although digital circuits dominate, front end amplifiers and filters are essential and critical components, whose success is intimately linked to the overall circuit design.

Given the large scale and investments in these projects, it seems important to eliminate competition early, building multi-group teams, and gradually evolving new designs from proven ones where possible.

Programmable logic

FPGA developments have been rapid in recent years and rely heavily on the success of the integrated circuit industry. In contrast to ASICs, they offer flexibility and programmability, based on standard components. Costs are falling, size is growing and power is being contained [13]. As with IC technology in general, this trend will be difficult to maintain but it seems assured for some years to come. In systems where space and radiation constraints do not dominate they are already heavily used.

Optical technologies

This is perhaps the fastest changing area, where most new developments for the future may be expected. Activity in fibre optic links, which is subject to many of the same constraints as ASICs, continues to grow. There is continuous development of other optoelectronic components to provide new functions, like optical switches or holographic memories, which promise eventually to revolutionise some particle physics systems.

6.5 Tools

Design tools are essential to undertake circuit developments; they are usually expensive but collaborative ways have been found to make them accessible to a large community. However, modern circuits also require fast, high speed and expensive test instruments, which can only be purchased by some large laboratories and need qualified professionals to support. Some labs carry out testing of analogue circuits which often require greater flexibility and customisation closer to the end use. Both aspects are important for success.

Conclusions

The opportunity to be innovative in the R&D stage stimulated some immature, or even unproven, technologies. Ambitious use of novel technologies is not new to particle physics but the scale of the LHC experiments, and possible future projects, introduces a new factor, which is the discipline of industrial scale production and assembly. In electronics, this discipline is already widely felt in the development stage.

There is still some way to go before LHC switch-on and there is much important work to be done which was not covered here on electronic related issues like power supplies, assembly, services and other practical issues which must be completed to ensure successful long term operation at LHC.

The SSC and CERN Detector Research and Development programmes were very beneficial in evaluating potential technologies and hastening progress.

References

[1] G. Hall *LHC front end electronics.* Paper presented at INSTR99, IC/HEP/99-9
http://www.hep.ph.ic.ac.uk/preprints.html

[2] Proc. LHC Electronics _Board Workshop. CERN Report CERN 99-09(1999)
http://hep.physics.wisc.edu/LEB99/

[3] Proc. LHC Electronics _Board Workshop. CERN Report CERN 98-36 (1998)

[4] Proc. LHC Electronics _Board Workshop. CERN Report CERN 97-60 (1997)

[5] *The National Technology Roadmap for Semiconductors.* Semiconductor Industry Association. (1997). http://www.semichips.org/

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

[7] P. Jarron & A. Marchioro. Ref [4], 139.
RD49 Status Report CERN/LHCC 2000-003 (2000)
http://www.cern.ch/RD49/Welcome.html

[8] http://www.te.rl.ac.uk/med/apv25_web/apv25_top.html
http://pcvlsi5.cern.ch:80/CMSTControl/documents/Marcus/APV25_Dec99.pdf

[9] K. Gill et al. Ref [4], 277.

[10] D. Charlton et al. Ref [2], 185.

[11] A. Lankford. *Developing and Commissioning Babar Electronics.* Ref [2], 41

[12] J. Christiansen. *Testing LHC Electronics.* Ref [2], 41

[13] P. Alfke *The future of Field Programmable Gate Arrays.* Ref [2], 36.