CERN/LHCC 95-72

LERB Workshop, 11-15 September 1995

Panel discussion summaries

F.Bourgeois / CERN-ECP,
27 October 1995

Note : The summaries of the discussions held at the end of each session of the First Workshop on Electronics for LHC Experiments have been written by F.Bourgeois/CERN-ECP on the basis of notes taken by M.Campbell/CERN-ECP and himself.
F.Bourgeois claims full responsibility for the content of this document which should be read in conjunction with the proceedings of the Workshop (CERN/LHCC/95-56). This document is for limited distribution to those who spoke, LERB and LHCC.

Summary of the Radiation-Hardness Session, Chairman Erik HEIJNE/CERN-ECP

Spoke (in alphabetical order of their initials) :
D.Camin/INFN-Milano, H.Spieler/LBL,
J.Marine and J.Redolfi/Thomson TCS, M.Dentan/Saclay, M.French/RAL,
M.Liu/Honeywell, N.Van Vonno/Harris, P.Jarron/CERN, V.Radeka/BNL,
W.Dabrowski/Cracow

The discussion was started on the subject of obsolescence.

PJ asked what is the future of small foundries given that the cost of a new foundry will be a few B$?

JM replied that the big ones would absorb the small ones.

In the opinion of NVV there would be fewer large facilities and Harris might buy capacity in such facilities. There would be a move towards fabless semiconductor companies that farm out production in foundries like e.g. UTMC in Taiwan. JM agreed but said that in such a case full technology control would not be possible. NVV stated that such foundries now often require equity investment from the user semiconductor companies.

MC asked whether there would be any possibility to use slightly modified standard processes to achieve Radiation Hardness (R/H)?

NVV said that this would not happen because manufacturers would be unwilling to modify volume processes.

ML said that out of the 20 Rad-Hard developers which existed in 1990 only a few might be strong enough to survive.

MD said that such a scheme was possible for a Bipolar technology but that for CMOS one cannot start from a non rad-hard process.

VR asked whether Harris considered using SiGe transistors?

NVV replied that Harris would wait for commercial tools to be available before deciding. He mentioned IBM might sponsor such tools.

The discussion moved to the subject of the assessment and the monitoring of processes. EH wondered how long we have to continue measuring individual transistors and what resources (staff) would we have to foresee in our labs?

HS reported on his experience with bipolar technologies. He said that circuit behaviour followed very closely what would be expected from transistor parameters. However with CMOS this was not the case and he indicated as an example variations in noise figures.

ML said that Sandia had done a lot of work to characterise 1/f noise in CMOS but EH remarked that the Sandia measurements cover a frequency range up to 10 kHz which is where our measurements usually start, going up to 10 MHz.

DC made a plea for a standard set of fundamental parameters to be quoted to allow different circuits to be compared e.g. noise density/$\div$Hz..

WD agreed but pointed out that some unexpected behaviour was encountered with circuits which could not be explained by conventional models e.g. excess parallel noise in PJ's ICON amplifier.
WD expressed his worry that much of the data on R/H electronics was based on very low statistics. This was particularly worrying as one observes sometimes large variations from lot to lot.

JM noted that most of the processes being used are still under development and therefore less stable. JR added that 2 years were needed to stabilise a process.

NVV said that even with a stabilised process a way would have to be found to guarantee R/H (Radiation Hardness) during volume procurement.

HS said that there was a need for each design to correlate the characteristics of the reference transistors with those of the complete circuit.

In reply to EH MF said that a standard test vehicle was used by RAL on all runs to help monitor the process. In particular this helps to distinguish between the effects of process variations and new designs.

EH pointed out that our designs required some parameters which were not explicitly guaranteed by the foundries.

MD said that we had to extend the set of parameters and test devices we require to include noise. JR replied that TCS will guarantee process compliance with respect to Process Control Monitoring PCM (device characterization), but not to specific test circuits.

MD repeated that with mixed mode circuits we had to find a way of achieving these goals.

NVV confirmed that the customer was no longer in control of the means used at the foundry to monitor component performance. One had to find a low cost and streamlined approach that focuses on chip performance. JM agreed completely. ML agreed too but pointed out that in some cases e.g. SOI, circuits would be necessary to guarantee a process as inter transistor effects had to be accounted for.

Summary of the TRACKING Session, Chairman H.J. HILKE/CERN-PPE

Spoke (in alphabetical order of their initials) :
G.Hall/IC London, H.Spieler/LBL,
J.Redolfi/Thomson TCS, M.Campbell/CERN, M.Dentan/Saclay,
M.French/DRAL, M.Turala/CERN, P.Jarron/CERN, S. Klein/LBL,
V.Polychronakos/BNL, V.Radeka/BNL

HJH opened the discussion asking engineers whether they understood the specifications of the trackers as given by physics for the front-end electronics?

MF replied that through the RD20 collaboration the Si strip behaviour was well understood. However signals from the other detectors e.g. MSGCs, were less clear and the best processing algorithms were not understood for the other types of detectors.

VP said that he did not agree for gaseous detectors for which one had in general a very good idea of the behaviour.

MC said that the physics requirements for the pixels were unclear partly because these new detectors might provide new physics.

PJ raised the question of specifications with regard to cost. In particular this may be strongly affected by yield especially for Rad-Hard processes.

HJH remarked that GH had given figures for the yield which appeared higher than those actually achieved at LEP.

GH replied that in his cost breakdown he assumed the yield as given by the manufacturers for chips. As for the hybrids and for a 9 chip design a 90% yield is achievable but did not include the cost of rework.

HJH asked how safe was it to assume a figure of less than 3 CHF per channel?

SK remarked that in his experience cost overruns came mainly from unanticipated assembly problems.

HS emphasized the necessity of making chips which were testable on the wafer. Rework was expensive and Known Good Dies (KGD) must be identified before mounting them on hybrids.

HJH asked where the various steps would be executed e.g. would bonding and testing be done in institutes? His guess was that testing would be done in the Institutes.

HS reported on his experience at LBL which is well located close to Si valley. They used the same test set-up for testing both chips and hybrids in the institute (plug and software compatible). Most of the basic assembly is done in industry and he admitted that other institutes were not in such favourable locations as compared to LBL. He thought that the assembly of hybrids with the detector would be done in the institutes.

GH noted that LBL was two years ahead of Europe as a result of their participation to SSC experiments. He said that it would be unlikely that CMS would have hybrids assembled in industry. In his opinion this work is less challenging an therefore snould be done in house. Complete module assembly must definitely be done in the Institutes.

U.Kotz/DESY asked the panel if the testing of the analogue chips could be done in industry?

MF pointed out that testing of the APV5 currently required 30' per chip and to test the whole lot in a two year time frame would need 20 probe stations. He said that although progress had been made on making chips testable there was still a long way to go to achieve high module yield. Features should be added to the design which make the identification of KGDs (Known Good Dies) more efficient.

SK advocated the use of automated testers operated by unskilled labour. Without such an approach the time required for testing would be prohibitive.

MT asked how to verify Quality Assurance of chips?

PJ said that one has to get rid of infant mortality.

HS remarked that some tests increase the failure rate and that they should be avoided. In reply to a question on the need for burn-in cycles HS said that this had been done for SVXs and no chip had failed at that stage nor in the experiment. In his experience, digital testing is well covered by industry whereas analogue testing had to be done in the institute.

VP said that Amplex chips had been tested and used for several years and that one could give some yield figures.

PJ said that automatic wafer testing of the Amplex was carried out at IMEC and took 30s per chip. Rejected chips were out of specs as given by the test vectors e.g. gain window.

MD insisted that foundry should provide the appropriate QA service in particular analogue, digital and R/H aspects.

GH remarked that little analysis had been done in our community to understand failures and therefore where to put the effort. He said that some of his colleagues still expected to service the detector every year. In his opinion this may not be possible. He felt that failures due to connectors may be dominant whether optical or electrical.

PJ replied to MD that QA would be costly.

HJH emphasized that he had heard many times that one had better not ask industry what is not done already.

MD said that a trade-off had to be found between standard digital monitoring and full monitoring of rad-hard and analogue properties of individual wafers.

JR pointed out that the discussion was mixing different aspects and that industry was used to detailed cost analysis and was in a position to supply KGD so as to avoid rework. It is mainly a problem of test coverage and a trade-off between cost and risks.

SK advocated the adoption of industry standards such as JTAG.

HJH moved to the problem of the many options that were still being considered for the various detectors. Could some parameters be removed from the decision process? He gave the example of noise figure of the preamplifiers which were similar for all designs. He also felt that power requirements were becoming an important factor which might favour the bipolar solution.

GH said that CMS had set an upper limit of 2 mW/channel which was not very different from the figures for the bipolar solution.

HS said that power dissipation was a problem and that local hot spots existed whatever care was taken in the designs. Keeping electronics cool is essential to minimize reliability problems. He noted that the community was often designing on implementation specs rather than on performance specs. He gave the example of the trade-off between occupancy and efficiency. He said that physicists were infinitely inventive and that, at the test level, they emphasized wanted modes of operation rather than unwanted ones.

PJ quoted as an example the bug found in the Pentium.

MD questioned the attitude of the community if radiation resistance was not to be guaranteed above 1 Mrad.

MF admitted that the Harris AVLSI-RA process was qualified up to 300 Krad. However verbal indications were given that the process could withstand 10 or even 100 Mrad. In his opinion, delivered batches should be subject to radiation resistance tests on a sample basis and rejected in case poor resistance was to be found. He felt that variations in noise were more important MD said that one could not return 50 wafers so easily. MF replied that Harris was aware and that they want to succeed.

HJH asked where the community stood with respect to the R&D phase and for how long it would be extended?

MC pointed out in his opinion many problems still existed and that engineers and physicists did not agree on this issue.

HJH asked when could system tests be carried out on the chosen solutions - may be in 2 years (in fact the milestones for the experiments)?

PJ said that, at this time, he did not know of final designs that had been thoroughly tested before and after irradiation.

WD said that according to his knowledge nobody had succeeded to make a production run which had a satisfactory yield and satisfied the required specifications. One should first get the circuit with full functionality and then do the R/H test. He felt that more time was needed. This was particularly true for R/H technologies as one iteration takes about a year.

GH said that the situation had vastly improved over the past three years and that impressive progress had been made in R/H electronics: cycle times are improving. He felt that the situation was more worrying for electro-optical devices.

There was some disagreement on the problem of knowing how long it would take. It was said that the question was more to find how long we could do R&D. HJH said : 3 years?

SK said that instead of asking how long should we continue R&D, we should ask how long can we continue? The longer you wait, the better detector you get at lower cost.

FB asked the panel for a list of items -e.g. preamplifiers, SCA- that were considered to be understood and no longer subject to further R&D? No reply was given.

VP said that he knew of 150 ASIC designs under development at any one time, but only a few of these had been used in an experiment. Only 5-6 were actually used after several years of development. Could the development time be shorter?

SK said that for NA49 (in a non R/H technology) 18 month were needed.

PJ agreed but noted that people had been working on SCAs since 1986 and that one should clarify how the development time was calculated.

VP was convinced that design time had shortened but the problem is clearly to agree on a deadline when to go to production : 7 years before the start of LHC ? 6, 5?

HJH suggested that 2-3 years were still required

VR said it may be better to look from the installation date and work backwards to see how much time is available. He suggested that all subsystems should be ready two years before the start and that another two years would be necessary for production.

GH agreed that assembly time was poorly defined and that there was an agreement that designs would have to be ready 5 years before beam time. In his opinion 1998 was a serious deadline for the start of the production.

Summary of the ELECTRO-OPTICS Session, Chairman W. SMITH/U.Wisconsin

Spoke (in alphabetical order of their initials) :
A.Weidberg/Oxford, B.Taylor/CERN, E.Monnier/CPPM,
E.Petrolo/INFN Rome, F.Vasey/CERN, G.Hall/IC London,
K.Borer/U.Bern, M.Campbell/CERN, M.Turala/CERN,
Ph.Farthouat/CERN, S.Klein/LBL, V.Radeka/BNL

WS opened the discussion asking 1) what were the critical points for R&D in the field and 2) whether one could agree on a smaller variety of E-O systems in order to achieve economies of scale?

FV said that for RD23 there were two main concerns : second sourcing and sufficient R&D resources.

EM shared FV's views but felt that all projects were facing similar packaging problems. The subject had to be pushed and progress should be made within the next two years. We are at the beginning of the story. In his case mock-ups expected next July for the MQW laser arrays. Industrialisation is an issue.

KB said that their LED activities had been driven by the wish to use a mature technology and have a fall back solution in case the current R&D would not be successful. He confirmed that packaging was the major problem and that he was disappointed by the lack of willingness of the manufacturers to cooperate. He quoted his unsuccessful experience with two particular firms. He is worried by the price levels and the timescale. He felt that the community could not wait too long.

BT said that he was happy with the RD12/TTC. On the connector side a well known coaxial connector manufacturer has developed a new miniature fibre connector which has a profit making niche market. This connector was developed without funding from CERN. In his view things are looking well and there is a fast evolution of the market. He felt that ATLAS and CMS had no reason to have more than one TTC distribution system. He felt that the present technology may be the one to be used in particular if R/H must be satisfied.

AW agrees with KB. He is however more optimistic as a result of his cooperation with GEC-Marconi. Developments need be paid and a R/H LED based analogue link should cost 100 CHF/fibre. It is a link candidate for the ATLAS SCT. Short timescales/deadlines are set by the tracker electronics. Things must be fixed in '96 leaving little time for R&D.

GH said that it was wise to rely on standard components but that it may not be completely true in particular when long lived products may disappear from the market. He therefore said it might be better to rely on a promising R&D.

FV said that in the case of HEP our packaging requirements made R&D necessary.

AW agreed that packaging was the major problem especially for the on-detector components.

FB noted that the panel on tracking had not shown any urgency in that there did not exist fully developed electronics.

EP advocated the need to concentrate the effort on a limited number of link technologies. In his opinion Gb link costs are not diminishing. He also reminded the group that the bandwidth

required by ATLAS was equivalent to 2x the world traffic in voice communications... Standardization would be needed to reduce costs.

WS said that he did not completely agree and that the market was more competitive. He added that the HP G-link would be available in a rad tolerant version.

MC asked why RD23 were concentrating on an analogue link when a digital link was clearly viable and may speed up R&D.

FV said that the emphasis was on the analogue transmission because it was requested. Also if analogue performance is demonstrated one can relax on some parameters and use the same technology for digital transmission.

VR expressed the worry that the wide spreads exhibited by some parameters may indicate deeper problems.

FV agreed. The main problem is related to the growth of layers and one cannot afford several attempts to build up statistics. He said that the yield was good but that confidence had to be gained. This cannot be done with the resources of the present R&D.

SK mentioned recent price reductions in the USA. A fair fraction is spent on the fibres. What about faster fibres to achieve higher levels of multiplexing, reduce the number of fibres and hence save on the installation costs. what is the appropriate breakpoint?

BT replied that the RD12/TTC was specific and that little could be done to further reduce the total number of fibres and connections.

EM said that distances in experiments were below 100m and that there was a high level of multiplexing. One can use multifibre connectors. Pigtailing should be avoided. He repeated that in his view packaging and connectors were the key issues.

SK saw the economics differently e.g. cost of laying fibres.

EM added that there were tighter constraints for monomode fibres.

AW said that in any case the modularity of the inner tracker of ATLAS imposed a certain fibre configuration if large quantities of copper were to be avoided.

WS questioned again what was appropriate to reduce the number of systems?

FV said that ideally there should be only one. Not necessarily RD23. Indeed the packaging techniques developed in RD23 could equally be applied to lasers and photodiodes. Whichever system is selected should benefit from what has been learned in the various R&Ds.

KB agreed that one system was enough in principle. However, if analogue and digital links were needed 2 systems may be necessary as the digital solution would probably be cheaper.

EM said that constraints may differ and that it was a problem of cost. Systems could be single or multi channel. Development should be shared in particular on packaging and receivers.

BT said the there should be only one TTC system. However, it may be sensible to have an alternative solution in case of unforeseen circumstances.

AW argued that there was a number of specific constraints for the trackers. Ideally zero mass and zero power. One cannot insist blindly and a compromise has been found in ATLAS.

PhF said that there was little information on the cost and that a number of patch panels would be needed to allow for the opening of the detectors. One or two more connections are needed.

BT said that the TTC had a tree structure and only five levels of interfacing. Only the last connection counts.

EM felt that any new development would soon or later be accounted for in the final cost of the components of the system.

FV said that connectors have insertion loss and hence affect the optical power budget. He would not recommend more than two connectors.

PhF agreed but said that the ATLAS TRT had decided for copper because they did not have a sure number for the cost of the optical link. Design could not be frozen if the cost of the complete fibre connection was not known.

FV referred to the calculations of GH for RD20 and said that the cost and the associated uncertainties were known.

Summary of the CALORIMETRY Session, Chairman V.Radeka/BNL.

Spoke (in alphabetical order of their initials) :
A.Marchioro/CERN-ECP, B.Lofstedt/CERN-ECP, C.Bohm/U.Stockholm, C.de la Taille/LAL, F.Bourgeois/CERN-ECP, H.Peek/NIKHEF, I.Brawn/DRAL, J.Parsons/U.Columbia, L.Wiggers/NIKHEF, M.Campbell/CERN-ECP, PJarron/CERN-ECP, W.Dabrowski/U.Cracow,.W.Smith/U.Wisconsin

VR presented a list of questions to be discussed (ADC, preamp noise, cost and power, packaging etc...). He opened the discussion by asking the panel why an 80MHz ADC was necessary?

JP said that in his opinion more bits (12) at 40MHz was more interesting and had a higher priority than 80MHz. Also the implications on system complexity, cost and power were non-trivial.

BL agrees on most of these points. However, there was an official request from CMS for 80MHz. He thought the important thing was to consider the whole system. He added that he would be glad to stay at 40 MHz. He admitted that more bits were needed and said that each problem cannot be seen in isolation.

IB said that in his opinion 8bits and 40MHz were sufficient for the liquid calorimeter trigger and he feared that 80MHz may hinder bunch crossing (BCO) identification.

WS remarked that the CMS ECAL people would not agree possibly for pile-up reasons.

FB then asked JP why they were exploring a digital solution if the Nevis analogue pipeline appeared to be adequate?

JP replied that because the analog solution was mature it was interesting to spend the remaining time exploring the digital alternative. There could be advantages in using the digital approach although he admitted that the timescale was aggressive.

VR said that in his opinion the faster shaping time was interesting for the 16ns BCO situation, but he would adopt a more conservative approach now that we have a 40 MHz BCO frequency and a higher luminosity.
He then turned the discussion to the question of noise as expressed in ADC counts. Preamplifier noise should dominate and not the ADC noise.

BL agreed and said the ADC noise should be much less than 1 LSB.

CdlT remarked that if there are 2x12 bits ranges the noise should be equivalent to 14 bits dynamic range. However, one should take care when summing towers that coherent noise does not dominate.

JP said that using multiple sampling weighted means gave (for 5 samples) a noise reduction of a factor 2. However ADC offset could still be a problem.

PJ turned the discussion to the comparison of the two schemes:
1) Multi range , linear
2) Compression, non-linear (Fermi)

BL said that there was not a true 14-bit resolution with the switching scheme because of the overlap between ranges. The Fermi scheme has the advantage that there is no switching in the analog part.

VR said that the NA48 scheme involved switching the gain which is not the case for the new Nevis design and questioned the virtues of such systems.

JP replied that the resolution was close to 14 bits as the overlap regions were small. Also calibration is greatly facilitated by these regions. In the Fermi scheme the 1st knee is in a very interesting range and this would require very careful calibration.

BL agreed that calibration would be more difficult for Fermi. A general problem was that he didn't know how to generate calibration signals which were precise enough. Maybe they would have to calibrate on known events. In any case calibration was a serious problem for any system.

JP said that using say Z^0s would be part of the answer, but this was why ATLAS had chosen the liquid calorimeter option. It was felt that in that case calibration is understood.

HP asked if the same signal could not be applied to all channels and the responses averaged.

JP said the number of channels was too large and also the required 0.2% was difficult to achieve with the averaging method.

CdlT said that the linear solution was easier to calibrate as the compressor solution required precise calibration throughout the entire dynamic range.

BL said that only the knee regions presented a serious problem as other parts of the transfer curve are linear. VR expressed doubt that all channels could be calibrated with particles and that ultimately the calibration problem is strongly coupled to the stability of the system. He thought that the compressor solution would be less sensitive to charge injection because of the absence of switching.

FB asked if all collaborations were referring to the same multi chip modules MCM when they showed the Fermi block diagram?

BL replied that both approaches were based on the same scheme.

WS said that ATLAS and CMS would use the same functional blocks but that the MCMs would be different because of the detector designs.

FB remarked that this would have a serious implication on cost.

WS said that most of the cost would come from obtaining good components (Known Good Dies) for the MCMs.

BL said that Fermi was trying to develop an architecture which would be suitable for both experiments. However, he agreed that the implementations would probably be different.

LW asked how the pile-up problem was addressed?

JP said that with multiple samples pile-up could be detected. At Zeus pile-up candidates were written out but these events were never analysed. This problem would be worse with 200K channels but that the option of writing out the events should be available.

BL thought that the Fermi architecture had a slight advantage here. The 1st event is passed to the LV1 trigger while the data is stored in the pipeline. Subsequent filtering of the pipelined information would enable these events to be flagged.

IB thought that pile-up would not be a major problem for the LV1 trigger due to the 1% occupancy.

VR asked when the systems would be available and how much they would cost?

BL said that all the required ASICs for Fermi (version 1) were there and that a VME system would be available by the end of '95. A new ADC (10-bit 40MHz) would be sent to the foundry in November '95 (ready April '96) and the first MCM (version 2) would be ready for test in the beam in '97.

VR asked which ADCs are available from industry. Two were mentioned, one of which is used by NA48. One manufacturer had also announced a 12-bit 40MHz device, but parts are not yet available for testing.

VR then raised the question of packaging, noise and power density.

JP said that this was an urgent matter particularly in ATLAS where the space given on the detector is very limited. As well as this the analog pipeline had to be proven to operate correctly in read write R/W mode in a large system. With the digital approach one has to take care of crosstalk. His major worry for Fermi was the 50x increase in channel density which they had to achieve soon.

BL said that the Fermi people had spent the last 6-8 months investigating digital noise issues. The key was to take care of power distribution and transmission lines. Technology advances such as 3.3V logic and submicron should help. MCMs have the advantage that line impedances can be controlled. Wire bonds should be avoided. Modern simulation tools help also.

VR warned that experience with a 15-bit dynamic range system of this size does not exist. BL agreed.

HS said that he would worry about effects such as substrate coupling and asked whether reliable simulation tools were available.

CdlT said that cables and grounding could be problematic.

MC said that experience with pixels showed that crosstalk effects were difficult to model. Problems are not always seen with the available simulation tools.

BL said that they are aware of such problems and that EMC will be watched at each stage of the development.

AM pointed out that development times in industry for such complicated circuits are very long and that many designs are never used. He quoted 100 man-years for the design of an advanced HDTV circuit and that 20% of the circuits worked after the second iteration, 40 % after three iterations and that the remaining 40% never worked !

WD made the point that the NA50 system proved that it was possible to run a 50MHz clock while remaining sensitive to analog signals, but they were far from the 15-bit level (7-8 bits). He would certainly worry.

VR asked about the power density of the various developments.

BL replied that the version #2 of Fermi should have a 4cm x 4cm MCM dissipating 5-7W at 40MHz (6 channels). He thought that using submicron technologies would help to reduce the power e.g. 0.6 micron technologies may reduce the power consumption by x2-3.

FB pointed out that this would still mean 200 KW on the calorimeter.

JP agreed and said that the point had been made by VR. JP felt that the final design would have to be ready around 1998 and that one should be careful not to rely on new technologies.

BL felt that for digital circuits VHDL descriptions should ease the problems of transfer between technologies enabling us to take advantage of new technologies with little redesign or risk. 0.35 micron would be available in 1998 and he felt that it would be wrong not to benefit from it.

VR reminded the panel that the electronics would have to be installed before switching on the beam as maintenance afterwards would be extremely difficult.
He then asked how to select from the various approaches to level 1 trigger generation?

WS said that CMS and ATLAS were collaborating through RD27 and that CMS had selected already one approach. ATLAS had still 2 options.

CB said that ATLAS would take a decision by mid- to end-'97, but that he was worried by practical problems such as the number of ASICs and connectors.

Summary of the MUON Session, Chairman A.Nisati/INFN Roma

Spoke (in alphabetical order of their initials) :
A.Kluge/U. Vienna, I. Kudla/Warsaw, J.Oliver/Harvard, L.Wiggers/NIKHEF, L.Zanello/INFN Roma, R.Breedon/UC Davis, R.Martinelli/INFN Padova, S.Veneziano/INFN Roma, V.Polychronakos/BNL, W.Smith/U. Wisconsin, Y.Arai/KEK

AN opened the discussion by reminding the audience that the radiation environment for the muon electronics as considerably less harsh than that for the other detectors (less than 10**11 n/cm2 fluence).
As a first question he said that the current muon electronics appeared to be adequate for the LHC therefore what work would be needed to come to a design which was ready for production?

SV replied that the ATLAS muon trigger needed to concentrate on the design of the system, in particular the interface to the TTC system.

YA said that for TDCs he did not see any fundamental problem and that the main target should be to reduce the cost by integrating many channels.

JO said that they had simulated a latency of 6-7 BCOs. This could be reduced to 2 by using full custom. For ATLAS they will spend 12-18 months to optimize the sub-circuits before finalizing the design of the full system.

RM said that the drift tube technology was well developed but that the main problem remaining was to enhance the reliability of the system.

LZ said that the solutions proposed appeared to be adequate. However she worried about the maintenance of such home-designed products in the long term.

IK remarked that their system was very simple to understand (look-up table memory with pipelined architecture).

AK remarked that it was clear that the LHC experiments would have to take care that experts passed on their knowledge to new people before leaving.

RB said that they were negotiating already with industry for the SCA.

VP asked why there were so many competing efforts (10 different groups)?

WS remarked that collaborations usually look for a commercial solution before starting a design. Normally this did not exist and full custom or gate arrays were designed. VHDL and Verilog allow such designs to be well documented and transferred between technologies which is an improvement on the past. He also pointed out that commercial products have a certain lifetime as well.

FB asked the panel if their devices were testable and how they avoided bias in their triggers?

SV replied that the individual devices were testable.

YA said that in his case this was an important parameter of the design.

FB then asked how they could guarantee the timing at the level of the system?

YA replied that this depended on the hierarchy of the system.

EP said that in ATLAS the trigger system itself was treated like any other sub-detector and is monitored continuously.

IK said that boundary scan techniques would allow a static test but wondered if a dynamic (on-line) test could be envisaged.

LW asked if test patterns had been foreseen for the front end systems. IK replied positively.

RM said that test points had been foreseen for the drift tube electronics.

AK said that an on-line monitoring system was necessary and that test of the entire system should be carried out.

RM said that an on-line monitoring facility was needed in order to test the overall functionality and connectivity.

WS remarked that both collaborations were looking at prescaled samples of events to verify trigger efficiency.

AN said that in his opinion there was not yet a clear strategy for testing. He suggested that one could take muon triggers at low threshold, measure the muon tracks with the tracking system and check two things 1) whether the trigger muons are found by the tracking system, 2) that all muons found by the tracking system are seen by the trigger.

AN introduced a short discussion on cooling. Was the price of cooling included in the system costing?

RB said that the cost of cooling had been included in the estimation of the FE costs.

SV said that in ATLAS the power budget was a few mW/channel and that the electronics was distributed throughout the detector.

Summary of the DAQ Session, Chairman A.Lankford/UC Irvine

Spoke (in alphabetical order of their initials) :
C.Kiesling/MPI Munich, G.Rubin/KFKI, H.Peek/NIKHEF,
J.Elias/FNAL, K.Muller/KFA, K.Noffz/Mannheim, L.Wiggers/NIKHEF,
L.Zanello/INFN Roma, M.Mojaver/UC Santa Clara, Ph.Farthouat/CERN-ECP,
S.Falciano/INFN Roma.

AL started the discussion by focussing first on triggering. He commented that in the technical proposals of ATLAS and CMS there is a so called data driven back-up that is focussing on commercial RISC processor farms. In the talk of SF she discussed the algorithms that would be used for the muon trigger level 2 and then there was some discussion we of neural networks, data driven architectures and so on. Could SF comment on how you see those machines?

SF replied that first of all for the moment they are very much open to consider eventually a hybrid solution and Enable++ offers a solution which is very interesting. Probably it can be used somehow to accelerate a part of the algorithm which is the simplest part and which is the input to the most complex which is for example the determination of the momentum. So she would be very much still in favour to look at this possibility. From her past experience she felt that to follow a fully programmable approach should be easier to use when the algorithm has to be tuned at the beginning of the experiment. However, algorithms are very complex and a hybrid solution may be needed. She was not sure about neural networks because she had the impression that this kind of technique could be used better for energy clusters rather than with track finding for which it looked less performant.

CK made a general comment concerning the neural network algorithm experience in H1. It was to look for high dimensional correlations for which, with general purpose processors, they could not apply simple cuts on a fast enough timescale. In his opinion, neural networks should be used when massive parallelism and large computations are required. Neural networks proved to be useful, successful, and well studied in calorimetry applications. This was confirmed in H1. One should really study other application areas. Tracking and more generally, pattern recognition, should be studied. A first example based on tracking information, and without calorimeter information, will be studied in H1. Results to be reported in a year from now.

AL said that in his view the major role of the L2 was to sharpen Et and pt cuts. He asked the panel whether neural networks were a good tool for this type of problem?

CK replied that the focus had been put on more general aspects such as event topology. He said that to sharpen the pt cuts on single tracks one needed full pattern recognition. Neural networks have been successful for their speed namely : a quick answer with good precision. He said that it is not good to try to find the mass of the W from jets and that in general neural nets need not be used to try to get high dimensional information.

KN said that the Enable++ was able to execute algorithms for sure. He did not know of limiting parameters. FPGA processing is good for many low level algorithms of preprocessing tasks. Neural nets had been implemented for the global decision task. There are many possibilities and Enable++ is well adapted to the first data reductions and to produce objects that will be processed in the higher levels.

AL noted that CK did not speak of any particular computational model.

CK replied that there were several ways for constructing neural networks. They had focused on the feed forward neural net which in itself is basically a multiply add machine. The determination of the weights is the delicate part in that one has to train the network with the correct examples and/or know what had to be thrown away e.g. background.

LZ asked KN why floating point was not supported in his parallel C. She felt that it was a major limitation in that one had to execute some functions which could not do without it e.g. find radius and centre of a circle.

KN replied that floating point was not as effectively implemented in FPGA as integer was. It was excluded from their C language because it was too complex for automatic generation. He mentioned a few attempts in the USA for which the speed up factors were limited to 10 or more as compared to the factors of 100 or more achieved on integer arithmetic.

LZ added that in her opinion such processors were better adapted to preprocessing rather than to feature extraction which required 3 dimensional calculations, hence floating point calculations.

KN said that he was not in agreement with her because several feature extraction algorithms had been implemented and none required floating point arithmetic. He felt that such processors were particularly well suited for the execution of the inner loop of an algorithm involving simple calculations on the full data set e.g. centre of gravity. In addition one could benefit from parallelism.

JE questioned the training risks associated with neural nets. He gave as an example the problem of distinguishing quark jets from gluon jets which involve 12 or more non independent variables. In such a case the steepest descent algorithm may lead to a local minimum rather than to the true minimum.

CK replied that he has never tried to answer that question as it is not related to triggering. He agreed that neural nets were not the best tool to address such problems. He said that the approach should be more phenomenological and that it was basically to set a goal like "I have a class of events for which I want to suppress background to a sufficiently good level". The problem was then to define the tolerable failure rate. If once the net is trained one gets 85% of physics and a rejection factor of 200 in background one may be satisfied but it is then that one begins to worry. For H1 many different sets of parameters were exercised. Very comparable rejection and efficiency rates were obtained after long enough training. He admitted that there was not a general solution to the problem of the local minimum in particular when there were shallow valleys.

AL said that he agreed with SF on the flexibility offered by truly general purpose processors. However he felt that this might not be realistic if one could not execute the required algorithm in the finite amount of time available. He felt that this was especially the case for level 2 and 3 algorithms. He questioned the audience on what was the state of the art in terms of level 3 processing for running experiments and how one would get the overall factor of 10**7 required at LHC.

JE said that a factor of 4 was achieved at level 3 for CDF. In his opinion, if LV1 and LV2 have large factors of reduction, L3 has obviously less to do.

PhF wondered whether the processing time was really the limiting factor. In his opinion too large a fraction of the time is used in data transfers. He gave the example of the ATLAS TRT for which 125 microseconds are needed to transfer data associated with a region of interest ROI whereas 700 microseconds of a PowerPC running a non- optimized code are currently

needed for the processing. Thus, at most, a factor of 6 could be gained over standard CPUs by the use of special purpose processors. He expressed doubts on the real need for special purpose processors which would soon be superseded by progress in the technology and require long term support.

HP said that the real problem was that the algorithm had to be defined beforehand and tested off-line to assess the required processing power. In his view, algorithms are often finalized when the experiment is running.

LW quoted reduction factors of 10 and 4 at Zeus for the level 2 and 3 respectively. In fact, level 3 involves the same tracking algorithms as the off-line ones.

LZ quoted two cases analyzed by by S.Cittolin and the CDF team which rejected 2 orders of magnitude less than expected. The reason was mostly related to the quality of the data. She said that if triggers were to be 10 times less efficient at LHC, experiments might be in trouble. It was not considered to be a matter of computing power.

SF referred to L3 at LEP and said that with time the trigger level 3 became a good place where to monitor the entire detector. One has full granularity and readout errors can be promptly signalled. The rejection factor was 60% without any reduction at the level 2. In her opinion, contacts with the off-line teams should be established at an early stage because extra effort is needed to taylor the algorithms to the permitted latency.

AL proposed to move to the subject of event assembly and asked PhF whether buses would be used in LHC experiments?

PhF said that the bandwidth of most bus standards was vastly insufficient to cope with the throughputs at LHC. He felt however that there would be room for VME in control activities even though it was an overkill.

AL suggested that the use of VME could be envisaged at the level of the readout buffers to connect to the links to the level 2 trigger.

PhF felt that ROIs have limited amount of data when there is sufficient granularity and that ad hoc connections would be better suited. In addition, if the level 2 is efficient one has 2 orders of magnitude less data between LV2 and LV3. Automatic scan of the buffers should be preferred.

AL said that several commercial and custom switches had been presented. He asked the panel members what would be the most widely accepted commercial solution at LHC?

MM said that he would first comment on the discussion on the use of buses. He felt that even though VME will be used for control purposes in the years to come, the bus concept is now antiquated. It will be replaced by switching networks which are basically distributed buses. One of the problems is that as clock speed is raising, electrical performance on a tapped system cannot be controlled to a sufficient level to maintain the high clock rates. It remains that the VME packaging format is nice. As for the best technology for switching one should consider that videophone and conferencing will dominate and drive the market. Video switching will have to make it possible. It is widely accepted that ATM is made for this type of application and it will be the winner. However, other variances may be used in LANs e.g. FibreChannel (FC) or even synchronous switches (they are being adopted in LANs). Video transport is very similar to event transport.

GR agreed with MM for video applications. He however felt that FC would find a market in mass storage technology. Mass storage has a significant share of the FC market and there exist switches with 64 ports at 1 Gb/s. One limiting factor for event building applications is the maximum size of the switch. Cascading is possible at the expense of an increased latency and the loss of a few ports for the cascading.

In reply to AL, CF said that progress had been made on SCI. Dolphin has announced vast improvements on the latencies and he felt confident with the future performance of SCI.

HP said that VME was being extended from 32 to 64 bit and that PCI was being introduced. He asked how long one would wait for the next upgrade.

MM felt that his previous comment had implied an answer to the question. He did not see that a "Futurebus++" would emerge in that its share of the market would be limited to defense and HEP. In his view the future of buses will be limited to VME and PCI over the next 5 to 10 years..

AL asked the opinion of the audience on the organization of the DAQ session?

CK felt that trigger level one talks had taken place in the detector related sessions and that it might be good in future to have them together with the TLV1 and TLV2 talks.

LZ said that the problem of error recovery at all levels of the system should be addressed. She felt that there was a need for a comprehensive study. AL replied that he had a number of transparencies on this issue and that by lack of time he had not included them in his review talk.

KM regretted that after the review talks given on the subject one had not come back in the course of the workshop on the issues of R/H and quality control. He said that in future one should seriously consider i) how one would merge the fine developments that were presented during this workshop and ii) how the transition to R/H would take place.