

An Initial Look at a CMS Level-1 Trigger for an Upgraded LHC.

J. J. Brooke, D. G. Cussans, G. P. Heath, D. M. Newbold

H.H. Wills Physics Laboratory, University of Bristol, UK
David.Cussans@bristol.ac.uk

Abstract

A possible upgrade to the LHC has been proposed. Two potential architectures for an upgraded CMS level-1 trigger are discussed, as are some ideas for a possible generic trigger processing module.

I. INTRODUCTION

Possible upgrades to the LHC resulting in a “Super LHC” (SLHC) have been discussed[1], which would result in increasing the luminosity ten-fold to $10^{35} \text{ cm}^{-2} \text{ s}^{-1}$. It has been proposed to double the LHC bunch-crossing frequency to about 80MHz to reduce the number of interaction vertices per bunch-crossing. However it is currently thought that reducing the bunch-crossing interval will not be technically possible[2].

At level-1 (L1) the CMS trigger[3] generates decisions based on information from the muon tracking and calorimetry systems. Information from the central tracking system is only read out of the detector when a L1 trigger (L1A) has been generated. In the high-level trigger(HLT) system[4] matching tracks reconstructed from the central tracker with energy deposits in the calorimeter or tracks in the muon tracker are a powerful tool in reducing the trigger rate. At the SLHC, it will be necessary to match tracking, calorimetry and muon information in the low-level, L1, trigger rather than the HLT [5]. Including tracking information in the level-1 trigger would allow it to gain much of the track-related rejection power currently only available at the HLT level. This matching is probably the only way to reduce the L1 rate in the harsh triggering environment of the SLHC to an acceptable level.

II. ASSUMED CONSTRAINTS.

In the following discussion it is assumed that in the CMS detector for SLHC the following constraints are in place:

- 1) On-detector electronics is not modified. In particular, to avoid dismantling the electromagnetic calorimeter (ECAL) the very-front-end (VFE) electronics[6] will stay in place. This limits the L1A rate to less than 200kHz (assuming five time samples per crystal are taken for each trigger). Using the current VFE also limits the latency between a bunch-crossing and receipt of an L1A to less than 200 bunch-crossings.
- 2) The central tracker will be completely replaced. The requirement for the smallest possible amount of material remains. This results in strong pressure to minimize the number of readout fibres, but also minimize the power consumption and hence minimize the volume of cooling pipes and power supply infrastructure.
- 3) The off-detector readout can be modified as required.

III. POSSIBLE TRIGGER ARCHITECTURES

There are a number of possible ways of combining calorimeter, muon and tracker information in the low-level trigger. Two of them are considered here. The first involves forming tracker “trigger primitives” (in this case tracks) on the detector and the second forming trigger primitives off the detector. In both scenarios, a subset of data is used to form tracker trigger primitives which are then used in the low-level trigger.

A. Off-detector Tracker Trigger Primitives

Perhaps the simplest change to the tracker architecture is to run the L1 trigger at its maximum rate, on the receipt of each L1A read out a subset of tracker data needed for triggering, then combine tracker with calorimeter and muon information in a “level-1.5” trigger. The tracker primitives would be generated off-detector, probably in a highly parallel track-builder farm. If a L1.5-accept is generated the full tracker information would be read out and sent to a largely unmodified readout chain. Calorimeter and muon information would be read out on receipt of a L1A as in the current CMS trigger design, but not combined with the full tracker data and not sent further up the readout system until receipt of a L1.5A.

Events with more activity in the tracker are likely to require more processing to produce trigger primitives than largely empty events. In the L1 trigger system the latency between a bunch crossing and the issue of a L1A is fixed by design. This does not have to be the case for a L1.5 trigger. Events could be stored in event buffers time-stamped with the bunch-crossing they correspond to and readout in response to a L1.5A with the same time-stamp. Using time-stamp information would allow the L1.5 trigger to have variable latency, and minimize the total processing power required by the track builder.

The advantage of an off-detector tracker trigger primitives generator (TPG) is that the average data volume transmitted out of the tracker would be lower than if all the data were readout on receipt of a L1A. In this way the granularity of the tracker could be increased without increasing the tracker data volume transmitted off-detector. However, the maximum L1A rate, set by the ECAL VFE, is limited to approximately 200kHz (this rate is only achievable if it is acceptable to read out only five time-samples per crystals per L1A, rather than the ten or twelve currently foreseen). This L1A rate limit implies that either the thresholds for calorimeter and muon trigger objects have to be high enough to reduce the L1A rate at SLHC to below 200kHz or some triggers have to be pre-scaled.

B. On-detector Tracker Trigger Primitives

In this scenario tracks are formed inside the detector and then transmitted to the L1 trigger. Tracker information is then used directly in the formation of the level-1 trigger. In this approach, unlike forming trigger primitives only on receipt of a L1A, the maximum L1A rate is not a bottleneck. In addition, the volume of data that needs to be transmitted off-detector is smaller, since only the tracker primitive objects are passed out of the detector, rather than the data needed to form them.

To form tracker primitives it is necessary to correlate data from different parts of the tracker. Fortunately, in the simplest case, attention can be limited to reconstructing high transverse momentum tracks, since it is only high p_t objects that are considered by the trigger. These particles leave data in a narrow radial “wedge” - a 20GeV/c charged particle is contained in a 4.2mrad wedge. Since data only need to be transmitted radially from layer to layer, with only a narrow spread in ϕ , it may be possible to use free space optics for data transfer between layers. Data transmitted from a sensor in one layer would have to be transmitted to more than one sensor in the next layer out. This could be achieved by shining the light from a transmitter onto more than one receiver in the next layer. Measures would have to be taken to ensure that data was only transmitted to the intended sensors. For example, coarse wavelength division multiplexing could be used, which in this context means having different colours for adjacent transmitters and filters on the receivers to pick out the desired signal.

C. Power Dissipation in Upgraded Tracker

Power dissipation is a critical parameter in the tracker system: An increase in power will increase dead material in the form of cabling to provide power and cooling to remove it again. Naively, doing the smallest amount of processing in the tracker (the off-detector tracker TPG option) will result in the lowest power dissipation and hence dead material, but this might be offset by a higher volume of fibre needed to transmit the higher data-rate off-detector. For the on-detector TPG option, even transmitting the data needed to make trigger primitives from layer to layer will need a significant amount of power. At a luminosity of $10^{35}\text{cm}^{-2}\text{s}^{-1}$ there will be approximately 2.5 charged tracks per cm^2 at a radius of 10cm [7]. Taking a simple model of a three-layer pixel detector, with an inner radius of 10cm and 1cm^2 sensors with $50\mu\text{m} \times 50\mu\text{m}$ pixels (i.e. 8-bit position information) this implies a total data volume of approximately 75 Tbit/s to pass data between layers. Currently available serializer/deserializer devices (serdes), for example [8],[9],[10], consume approximately 100mW per Gbit/s. This would imply a power consumption of $\sim 7\text{kW}$ just to serialize the data for transmission via optical links, This compares to just 3kW for the entire CMS pixel detector as currently designed. Transmitting this volume of data off-detector to form L1 trigger primitives off detector would require 7500 fibres, assuming a data rate of 10Gbit/s per fibre.

IV. POSSIBLE TRACKER GEOMETRIES

The current tracker[11] is optimized for the best possible momentum resolution and has more-or-less equally spaced

layers. To provide information for use in a low-latency trigger it might be preferable to have **pairs** of layers, with layers being closely spaced (a few millimetres) and tightly electrically coupled. This would allow identification of pairs of hits that point almost radially away from the beam-pipe (i.e. from high p_t particles). Identifying these pixel-doublets would greatly simplify track-finding – a FPGA-based method for identifying tracks from pixel-doublet information has already been proposed[12], which would be straight-forward to implement in an ASIC. Indeed, low-latency track finding is already performed in the muon trigger – albeit with lower multiplicity and channel count.

V. GENERIC TRIGGER PROCESSING MODULE

The CMS Global Calorimeter Trigger[13] has to perform a number of different functions: It must sort out the highest ranked electron trigger objects passed to it by the regional calorimeter trigger crates, form jets from region energies, calculate total E_t and missing- E_t from region energies and from jets. It must also collect and format data for transmission to the data-acquisition system and collect bunch-by-bunch histograms for on-line luminosity monitoring. Most of these functions must be done in a fixed (and small) latency in the region of $0.25\mu\text{s}$ (depending on the function), which precludes the use of general purpose microprocessor based boards.

Given the large number of different functions to be performed it was decided to use a generic trigger processing module(TPM) for all functions, based on the use of programmable logic devices[14]. Later, CDF also decided to use a general purpose trigger-module[15] for run-II to replace a number of different designs. Considering the advances in programmable logic devices it seems plausible that in any CMS L1 system for SLHC the large number of different modules used in the current CMS L1 trigger could be replaced by a much smaller number of more generic designs. Further studies would be needed to decide the parameters required for such a general processing module, but a reasonable starting point would be to extrapolate by an order of magnitude from the current GCT TPM.

D. Current GCT Trigger Processing Module

The current TPM is a 9U x 400mm VME module. It has a VME64 J1 backplane for set-up and control, but data I/O is implemented using 1.44Gbit/s serial-links on the front panel over short copper cables. Data sharing between modules is implemented using point-to-point 3.2Gbit/s serial links transmitted over copper twin-ax cables in a custom “cable backplane”. Each module has four large FPGAs, containing a total of 12M system gates, for data processing and smaller FPGAs for data routing and interface. Power is distributed at 48V with on-board DC-DC converters.

E. System Architecture

The VME system is unlikely to be the best architecture for any upgrade to the L1 trigger. The data throughput on the backplane is very low compared to high speed serial links and is shared between all modules in the crate. High bandwidth requires the use of custom backplanes. Power is distributed at

low voltage (12V, 5V, 3.3V) which for high power modules results in heavy gauge wiring and difficulties in precise voltage regulation.

It seems probable that a newer architecture such as the recently developed AdvancedTCA (ATCA) [16] will be a better platform. ATCA “front boards” have a 6HP wide front panel (1.5 times as wide as a VME module) and have a 8U x 280mm form factor. Power is distributed at 48V with DC-DC converters on the modules. There are a number of defined backplane architectures, implemented using Tyco “Zpack HM-ZD” connectors[17] which with careful backplane design are capable of transmitting data at 10Gbit/s per pair. Set-up and control is handled by a single 10/100/1000base-T Ethernet link to each module, taking up 2mm of vertical backplane space, compared to 90mm for a VME J1 connector. Figure 1 show an example of a ATCA crate.



Figure 1: Photograph of a Schroff Advanced TCA crate.

F. FPGA Choice

Although “Moore's Law” may well have broken down[18] on the SLHC time-scale, there are already new FPGA devices announced that would enable an order of magnitude increase in performance over the current GCT TPM. Moving to devices with 90nm feature size (e.g. Altera Stratix-II , Xilinx Virtex-4) rather than the current 150nm feature size devices (Xilinx Virtex-II) would allow the algorithm clock to be increased from 4xLHC-clock (~ 160MHz) to 8xLHC-clock (~ 320MHz) and the inter-FPGA bus speed to be increased from 160Mbit/s per line to either 640Mbit/s or even 1.28Gbit/s. At these speeds it is likely that a source-synchronous rather than system-synchronous clocking scheme will be necessary. Many of the recently announced FPGA devices include multiple 10Gbit/s serdes allowing higher system density than the current TPM, which uses discrete serdes. However, the latency of these built-in serdes will be a critical factor in their usefulness.

G. Input/Output Links

The TPM uses 1.44Gbit serial links connected using Infiniband 1x and 4x connectors, giving a total I/O throughput of 37.5Gbit/s. The 4x connectors provide up to 20Gbit/s total throughput (if run at 2.5Gbit/s.pair) in 40mm of front panel space. It is probable that connectors offering higher bandwidth

over copper cable will become available, however a more generic solution would be to use a pluggable connector system such as XFP[19]. This standard provides for a bi-directional link up to 10Gbit/s in each direction. Figure 2 shows an exploded view of an XFP connector and its plug in module.




Figure 2: Isometric Sketch of an XFP connector.

An interesting feature of the XFP system is that there is a carrier attached to the PCB, which is the same for all implementations, with a application specific plug-in module into which the cable connects. Hence, with the same generic module, copper cable could be used for short links and fibre-optic for long links. The selection of cable types is made by plugging in the appropriate modules. XFP connectors can be placed at intervals of 25mm. For even higher I/O density connectors to be placed both sides of an ATCA board. This allows up to 120 Gbit/s front-panel I/O per module (240 Gbit/s if connectors mounted both sides of the PCB). The XFP connectors could be driven directly by internal serdes in the module FPGAs.

H. Inter-Module Links

Many trigger functions require high connectivity between different parts of the system. The current GCT TPM has up to twenty-four 3.2Gbit/s links per module, giving a total system bandwidth of approximately 60Gbit/s. A standard “mesh” architecture ATCA backplane allow for up to 120 differential pairs per module. Data communication at 10Gbit/s per pair using serdes inside FPGAs has already been demonstrated [20], giving a total system bandwidth of approximately 9.8Tbit/s. It should be noted however that transmitting data over copper links at these speeds will require skills in electromagnetic modelling that are not commonly available within the HEP community at the moment.

VI. CONCLUSIONS

The very high multiplicity of interactions in each bunch crossing at the SLHC make triggering a difficult task. In this environment using central tracker information in the level-1 trigger will be an important tool in keeping the L1 trigger rate down to an acceptable level. Of the different possible architectures, two have been described an “off-detector TPG”

where the tracker information is integrated with muon detector and calorimetry trigger information after a level-1 accept and an “on-detector TPG” where track trigger primitives are generated in the detector and combined with other sub-detectors at level-1. Unfortunately, the most straightforward and flexible option, the off-detector TPG, does not improve the trigger performance much with respect to the existing trigger design. This is because the L1 accept rate can not be increased much beyond its current rate unless the existing calorimeter front-end electronics are replaced. Forming trigger primitives within the tracker would give full integration of tracking information at L1 without dismantling the calorimeter. However, designing such an on-detector tracker TPG would be challenging, if for no other reason than the difficulty of keeping power dissipation to an acceptable level.

In any upgraded trigger nearly all of the trigger components would have to be replaced. With the rapid progress of programmable logic devices it would be advantageous to design a common generic trigger processing module that could be reconfigured to perform many different roles. This concept has been successful not only within the CMS GCT but also on other experiments. Although it is not possible to accurately predict technologies that will be available on the SLHC time-scale, even with currently available components a successor to the CMS GCT TPM could be designed with an order of magnitude higher performance. Hence, we can be confident that it will be possible to construct such a generic trigger processing module for the SLHC.

VII. REFERENCES

- [1] S. Tapprogge, “Physics Potential and Detector Implications of an Upgraded LHC”, *Atlas Conference Report, ATL-CONF-2001-003*
- [2] O. Bruning “Options for future high luminosity upgrades of the LHC machine”, *Proceedings of the 10th Workshop on Electronics for LHC Experiments, In preparation*, Boston, 2004
- [3] CMS, “The Trigger and Data Acquisition project, Volume I: The Level-1 Trigger Technical Design Report”, **CERN/LHCC 2000-038**
- [4] CMS, “The TriDAS Project Technical Design Report, Volume 2: Data Acquisition and High-Level Trigger”, **CERN/LHCC 02-26**
- [5] W. Smith, “Trigger and Data Acquisition for the Super LHC”, *Proceedings of the 10th Workshop on Electronics for LHC Experiments, In preparation*, Boston, 2004
- [6] M. Hansen “The New Readout Architecture for the CMS ECAL”, *Proceedings of the 9th Workshop on Electronics for LHC Experiments*, pp78-82, Amsterdam, 2003
- [7] C. Foudas, “Initial thoughts for a First Level Tracking Trigger at SLHC”, Second CMS Workshop on Detectors and Electronics for SLHC, Imperial College, London, 2004 <http://agenda.cern.ch/fullAgenda.php?ida=a043080>
- [8] LSILogic VSC1053
- [9] Vitesse VSC7226
- [10] Lattice Semiconductor XPIO 110GXS
- [11] CMS, “The Tracker System Project : Technical Design Report”, **CERN-LHCC-98-006**, 1997 Geneva : CERN, ISBN: 929083-124-3
- [12] Jinyuan Wu, “(1) Tiny Triplet Finder (2) A Pattern Recognition Scheme for Large Curvature Circular Tracks and an FPGA Implementation Using Hash Sorter”, *Proceedings of the 10th Workshop on Electronics for LHC Experiments, In preparation*, Boston, 2004
- [13] J.J. Brooke et al., “Hardware and Firmware for the CMS Global Calorimeter Trigger”, *Proceedings of the 9th Workshop on Electronics for LHC Experiments*, Amsterdam, 2003, **CERN-LHCC-2003-055** pp. 226-229
- [14] J.J. Brooke et al., “An FPGA-based implementation of the CMS Global Calorimeter Trigger”, *Proceedings of the 6th Workshop on Electronics for LHC Experiments*, Cracow, 2000, **CERN-LHCC-2000-041**, pp. 363-367
- [15] Tiehui Ted Liu, “The Pulsar Project”, <http://hep.uchicago.edu/~thliu/projects/Pulsar/>
- [16] PCIMG, “AdvancedTCA PICMG 3.0 Short Form Specification”, 2003 http://www.picmg.org/pdf/PICMG_3_0_Shortform.pdf
- [17] Tyco/Amp, “Z-PACK HM-Zd connectors.” <http://hmzd.tycoelectronics.com/default.asp>
- [18] LB Kish “End of Moore's law: thermal (noise) death of integration in micro and nano electronics”, **2002 Physics Letters A 305 144**
- [19] SFF Committee, “INF-8077i 10Gigabit Small Form Factor Pluggable Module.”, 2004 http://www.xfpmsa.org/XFP_SFF_INF_8077i_Rev4_0.pdf
- [20] Suzanne Deffree, “SuperComm: Xilinx hits 10Gbit/s”, *Electronics Weekly*, June 2004, <http://www.electronicsweekly.com/Article36548.htm>