

Trigger and Data Acquisition for the Super LHC

Wesley H. Smith

Physics Department, University of Wisconsin, Madison, WI, 53706, USA
wsmith@hep.wisc.edu

Abstract

R&D is needed for LHC trigger and data acquisition systems if the LHC luminosity is increased to $10^{35} \text{ cm}^{-2}\text{s}^{-1}$. This upgrade, entitled the Super LHC (SLHC), would occur in the next decade. The physics triggers, the algorithms needed to provide these triggers against the substantially increased backgrounds and pile-up, and the types of electronics solutions that can support these more sophisticated algorithms are discussed. New architectures of Data Acquisition systems exploiting advances in network technology are also presented.

I. INTRODUCTION

The physics scope of the LHC would be significantly extended by the Super LHC (SLHC) luminosity upgrade providing up to $10^{35} \text{ cm}^{-2}\text{s}^{-1}$. This would increase the LHC mass reach by about 20 - 30% and provide the possibility to measure the Higgs self-coupling. This additional capability would be accompanied by a number of experimental challenges such as an increase in pile-up of additional overlapping events and a possible increase in the crossing frequency from 40 MHz to 80 MHz, which would partially offset this pile-up. These conditions at the SLHC would present a substantial challenge to various elements of the ATLAS and CMS detectors.

The ATLAS and CMS trigger [1,2] and data acquisition (DAQ) [3⁴] systems would need significant modifications to operate at the SLHC. Due to the increased occupancy of each crossing, at the SLHC Level-1 trigger systems would experience degraded performance of the LHC algorithms presently planned to select the 100 kHz of crossings from the input rate of 40 MHz. The electron isolation algorithms would experience reduced rejection at fixed efficiency and the muon trigger would experience increased background rates from accidental coincidences. The DAQ system would experience larger event sizes due to greater occupancy. If new, higher channel-count trackers replace the existing ones, then the increase would be greater. This would reduce the maximum Level-1 trigger rate for a fixed readout bandwidth.

In order to meet the challenges of SLHC operation the suggested approach is to hold the overall Level-1 trigger rate at the LHC value of 100 kHz by increasing the readout bandwidth. This approach avoids rebuilding front-end and readout electronics as much as possible since these were designed for an average readout time of less than 10 μsec .

It also permits use of front-end buffers for an extension of the Level-1 Accept (L1A) latency rather than for more post-L1A storage before readout. However, maintaining a 100 kHz L1 rate at the SLHC will increase the burden on the DAQ, which will need to transport more than the anticipated LHC data size of 1 MB per event. Holding the L1 trigger rate to 100 kHz at the SLHC implies raising E_T thresholds on electrons photons, muons and jets, as well as the use of less inclusive triggers. These strategies assume that with several years of data accumulated at the LHC design luminosity before the SLHC upgrade, many of the physics studies requiring lower E_T thresholds would have sufficient statistics and that the physics at the LHC would be sufficiently understood to provide the necessary understanding of efficiency for more complex trigger configurations.

II. SLHC TRIGGER OPERATION

The operation of the SLHC at a bunch crossing frequency of 80 MHz instead of 40 MHz would reduce event pile-up, improve trigger algorithm performance, and reduce data volume for detectors that have time resolution sufficient to identify data associated with individual 12.5 ns bunch crossings. There are some concerns as to whether the SLHC can operate at 80 MHz instead of 40 MHz, but it is important to address the issue of whether the detectors can. Frequencies higher than 80 MHz, such as 160 MHz or greater would require a time resolution of detectors of 6.25 ns or less in order to properly associate event data with specific crossings. The identification of data with crossings is needed for timing in the detectors in the experiment. Since most technologies used or contemplated for use in HEP collider detectors would be unable to efficiently separate event data at a higher crossing frequency than 80 MHz, these higher frequencies can be considered tantamount to continuous beam, which lacks the time structure critical for aligning the large number of geographically distributed detector channels.

At 80 MHz, a number of present LHC detectors would have sufficient time resolution to identify data associated with individual 12.5 ns crossings. Many detectors with readout sampling at 40 MHz can extract 80 MHz bunch crossing identification from those samples. For these, 80 MHz operation does not require rebuilding the front-end electronics, just the trigger primitive calculations that follow. While 80 MHz trigger operation would require rebuilding much of the LHC trigger systems, this appears feasible since much of the ATLAS and CMS trigger

electronics features internal processing at frequencies of 80 or 160 MHz or higher in some cases. Operation at 80 MHz or higher where the logic was 40 MHz before reduces the trigger latency if the processing steps are done at the crossing frequency.

The SLHC trigger systems have three major physics performance requirements. First, they need to be efficient for high- P_T discovery physics. This should not prove a significant rate problem since the thresholds for this are high. Second, they need to provide high statistics for the completion of the LHC physics program, such as precise measurements of the Higgs sector. This requires low thresholds on leptons, photons and jets. This should be addressed by use of more exclusive triggers rendered usable by the previous years of understanding final states observed at the LHC. Finally, they need to provide control and calibration triggers. These are such samples as W, Z and top events. These low-threshold triggers can be held to a low rate by prescaling, which should yield more than adequate size samples.

An initial study of the L1 trigger thresholds for the SLHC[5] suggests inclusive single muon and electron thresholds of 30 and 55 GeV with rates of 25 kHz and 20 kHz respectively. Electron and muon pair thresholds of 30 and 20 GeV respectively are predicted to have rates of a few kHz. Inclusive Jet and missing E_T thresholds of 350 and 150 GeV, respectively would produce rates of about 1 kHz. The actual triggers employed might use less inclusive conditions to access lower thresholds.

III. SLHC TRIGGER PRIMITIVES

The existing trigger primitive information used by the L1 trigger systems needs examination to determine adequacy for use at the SLHC. For the CMS calorimeter, the forward quartz fibre detector is sufficiently fast, but might require finer-grained information to provide a smaller trigger tower size. The CMS HCAL and ECAL have sufficient time and spatial resolution for 80 MHz SLHC operation, using their present 40 MHz sampling without significant modification. However, replacement of the high η calorimetry may be needed due to radiation damage. The ATLAS LAr calorimeter will experience more than a factor of 3 increase in pileup at the SLHC luminosity, which may require a change in the electronics shaping time to optimize noise performance. Some changes would also be necessary for $|\eta| > 2$ to mitigate space charge effects. The ATLAS Tilecal will need additional study of calibration and energy corrections, as it will be difficult to extract a minimum ionizing signal amid the pileup background. It will suffer some radiation damage at high η , which may require partial replacement.

ATLAS and CMS muon systems both use RPCs that may not function at the SLHC luminosity, particularly at high η . The existing ATLAS Muon Cathode Strip Chambers and Thin Gap Chambers will probably continue to be usable for triggering with some improvements and

higher thresholds. The same is true for the CMS Muon Drift Tubes and CSCs.

Another possible source of trigger primitives not presently used in either the ATLAS or CMS LHC L1 trigger systems is the tracker. This would most likely require a replacement of the tracking system and a change in technological solution. However, it is also likely that operation at the SLHC will require replacement of the ATLAS and CMS tracking systems anyway. Therefore it is worth considering what form tracking trigger primitives might take and how they might be used. In fact, CMS has the provision for a type of L1 tracking trigger using the z-vertex positions of pixel clusters of high hit occupancy in $\Delta\eta \times \Delta\phi$ bins. This could be used to reject jets from pile-up events since their z-vertices would not line up with the main event z-vertex. At present the logic for this is not implemented, but retained as an option.

A L1 tracking trigger could provide an inner track and possibly an outer stub. These would be used to combine with the calorimeter at L1 to reject π^0 s and reject jets from pileup. They would be used to sharpen p_T thresholds and reduce accidentals and wrong crossing determinations in the muon system. Implementation would not only require rebuilding the tracker, but also rebuilding the calorimeter and muon trigger systems to various degrees in order to provide outputs with suitable granularity and other information to combine with the L1 tracking trigger.

IV. SLHC TRIGGER ALGORITHMS

The most effective method to significantly improve trigger functionality for the SLHC may be to employ tracking at the earliest stage possible. In order to evaluate the effect of tracking on trigger performance, it is appropriate to examine how tracking is first used in the ATLAS and CMS Higher Level Triggers. As an example, the CMS experiment attaches tracker hits to muon tracks in order to improve the P_T assignment precision from 15% for the endcap muon system stand-alone to 1.5% with the tracker information included [4]. This also improves the sign determination and provides a vertex constraint. In addition, pixel tracks are found within a cone around the muon track and their sum P_T is used as an isolation cut. This is less sensitive to pile-up than calorimetric information if the primary vertex can be determined. The combination of tracking p_T resolution and isolation provide more than an order of magnitude reduction in the CMS L1 muon trigger rate. To implement such a trigger at L1 in CMS would require information on muon track locations on a $0.0125 \eta \times 0.015^\circ\phi$ scale. While finer than the present CMS $0.05 \eta \times 2.5^\circ\phi$ trigger scale, this information is already available but not used.

Tracking information in the HLT also reduces the CMS L1 calorimeter trigger rate. The correlation of an electron trigger with an extrapolated pixel track reduces the rate by a factor of 10. Tracking information is also used as an isolation cut for photon candidates, The CMS L1

calorimeter jet-based τ -lepton trigger is also reduced a factor of 10 by requiring isolation using pixel tracks outside the signal cone and inside an isolation cone. In order to use the information from a tracking trigger at L1, the calorimeter trigger e , γ and τ objects could be used to seed tracks with the full calorimeter trigger tower $0.087 \eta \times 0.087 \phi$ granularity. Candidates could be pre-sorted and limited to a maximum number, such as 32. A single track match within a 3×3 trigger tower region with a coarse p_T resolution (8-bit scale with 1 bit/GeV) could be sufficient to reduce the electron trigger rate. A veto of tracks in the 3×3 trigger tower region would be used for a veto of photon candidates and a single or triple track match would be used for τ candidates.

Other upgrades to the CMS calorimeter would be needed to reduce the jet and missing energy trigger rates. These would include allowing the clustering of jets in multiples of 2×2 trigger towers: 6×6 , 8×8 , 10×10 with a sliding window making one or two tower steps and use of higher resolution scales with more precise geometry for missing energy. All of these changes to the CMS L1 calorimeter trigger would represent a reasonable extension of the present system. Technological advances in FPGAs and data links would permit processing of high speed serial data such as 32 10 Gbit/s links per card with high speed serial output in the 4 - 10 Gbit/s range.

V. SLHC TRIGGER ARCHITECTURE

The implementation of the new algorithms involving tracking discussed above would require modification of the ATLAS and CMS trigger architecture. The 3-level ATLAS Trigger and DAQ system has an opportunity for the use of trigger data at Level-2 in the Region of Interest (RoI). However, CMS, with only two physical levels, would need to provide for the use of tracking information directly in L1. The present CMS architecture provides a flow of data within the muon and calorimeter trigger systems from regional to global components into the CMS Global L1 trigger. For the SLHC the L1 trigger data would need combination between tracking and calorimeter and muon triggers at a regional level with finer granularity than presently employed. After this regional correlation stage, the physics objects made from tracking, calorimeter and muon regional trigger data would be transmitted to the global trigger. The important new feature is that some of the tracking, isolation, and other regional trigger functions would be performed in combinations between regional triggers in a new hardware layer composed of regional cross-detector trigger crates as shown in Figure 1. An advantage of this architecture is that it would leave the present CMS L1 and Higher Level Trigger (HLT) structure intact by not adding additional trigger levels. This minimizes the impact on the CMS readout.

The additional layer of processing for combination of tracking information, increased algorithm complexity and larger trigger data volume due to finer trigger granularity

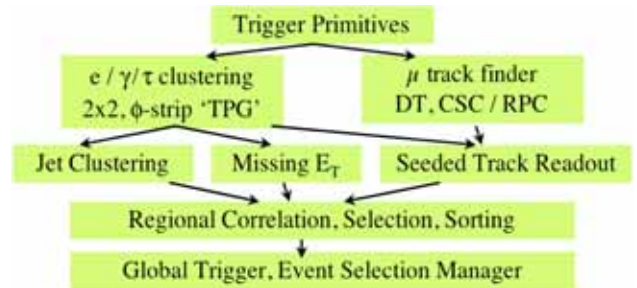


Figure 1. CMS SLHC L1 trigger architecture proposal (from S. Dasu).

suggest an extension of the present CMS 3.2 μsec L1 latency. A longer latency would also be needed for use of FPGA embedded serializers and deserializers, the addition of more serialization and deserialization steps to use high speed serial links or the use of buffers to incorporate commercial serial links running asynchronously with respect to the LHC clock. The CMS L1 latency is limited by the front-end analog storage capacity of the tracker and preshower electronics. Since it is expected that these detectors will be replaced for the SLHC, it is reasonable to assume that their electronics will be replaced also and that this limitation can be removed. The next limitation is the ECAL digital memory depth of 256 40 MHz samples corresponding to time of 6.4 μsec . This is proposed as the CMS SLHC L1 latency baseline.

VI. SLHC DAQ

If one assumes that the L1 trigger rate remains at 100 kHz, the increased channel occupancy and finer detector granularity, leading to a larger channel count, suggest a bandwidth requirement for SLHC DAQ systems at least 5 - 10 times that of the present LHC DAQ systems. One option is to create multiple additional switched slices of the existing DAQ systems to provide an increase in parallel processing capability. However, there are many other ideas. The DAQ upgrade paths for ATLAS and CMS may differ because of their different LHC DAQ architectures. For ATLAS, the RoI-based Level-2 trigger provides an opportunity to put tracking information into the Level-2 hardware to reduce the volume of data into the filter farm. For CMS, if the single scalable hardware level event building architecture were kept, tracking information would be added in the L1 trigger.

The present CMS LHC DAQ uses a network with Terabit/s aggregate bandwidth constructed from two stages of switches and a layer of intermediate data concentrators for optimizing traffic load to the event builder. The capacity of the buffer memories of 100 GB between the front-end readout and the event builder permit a real-time DAQ latency of seconds. A proposed architecture for the CMS SLHC DAQ is shown in Figure 2. The concept is to incorporate as much of the DAQ functionality as possible into a commercial network of the capability one can expect from industry in the next decade. This incorporates a scalable multi-Terabit/s network to interconnect all of the elements. The function of the Event Manager (EVM) is incorporated into the L1 trigger. The EVM updates the list




Figure 2. Proposed CMS SLHC DAQ system (from S. Cittolin).

of available event filter services where events are to be sent for processing. Along with the L1 accept, the trigger transmits additional information to the front ends including the event type for post L1 processing and the destination address of the event filter node where the resident event fragment is to be transmitted. This requires the control logic to process and transmit instructions at the 100 kHz L1 trigger rate to every readout, trigger event filter and other element of the DAQ. The event fragment delivery and the event building itself are provided by the network protocols using the commercial network hardware. This design allows for real-time buffers consisting of Pbytes of storage disks, which would permit storage of events being processed over a period of days. This opens up the possibility of use of non-local compute nodes and GRID tools to maximize access to remote resources in a flexible manner.

VII. SLHC TRIGGER & DAQ R&D

An important development needed for the proposed CMS SLHC DAQ architecture outlined above is a revision of the existing Trigger Timing and Control (TTC) system [6] to transmit and receive added fast control information. A new TTC system should provide the clock, L1 Accept, Reset, Bunch Crossing 0 and trigger type information in real time for each crossing. Should the SLHC operate at 80 MHz, the TTC system will need to deliver signals at this frequency. An R&D effort would be required for the TTC system clock signal so that it can meet the jitter and other requirements to drive the new generation of high-speed serial links, as well as to be capable of functioning at the GHz frequency needed to meet the fast message distribution needs of SLHC trigger and DAQ.

The substantial increase in algorithm complexity and volume of data they need should be met by industrial development of Field Programmable Gate Arrays (FPGA) promising faster devices with higher logic gate counts and increased I/O. Of note is the advent of embedded GHz serializers and deserializers on the FPGA inputs and outputs, enabling very high throughput. However, the latency of these circuits remains a concern and needs study. The use of these devices is becoming more challenging in terms of packaging, routing, mounting, and low supply voltages. R&D on the use and performance of these rapidly

evolving devices is needed for development of SLHC trigger systems. The need to move larger quantities of data at higher speed requires R&D on high-speed serial links and backplane technology, as well as on the clocking systems needed to drive them with low jitter.

The SLHC DAQ systems will also depend on are progress in backplane and data link and technologies, which will need to incorporate the newer frequency 40 GB links and protocols. The much tighter integration of front-end electronics and the links needed suggests that R&D on these be done together. The front-end electronics will have many challenges for R&D to handle the increased processing and channels counts. While improvements in VLSI technologies should provide the necessary components, there are many R&D issues such as radiation tolerance, power, reduction, system complexity, and integrating the commercial data communications developments. The SLHC DAQ system itself also faces a considerable challenge of managing the complexity caused by the increasing numbers of components, operations and stages of processing. This will require R&D on more sophisticated controls and diagnostics. One path is to exploit as much as possible industrial developments in this area.

As the SLHC L1 trigger adopts more sophisticated algorithms that were formerly used in the HLT and the backgrounds for the HLT algorithms increase, the HLT algorithms must become much more sophisticated. This will require considerable study, physics simulation and eventually data analysis of LHC data to determine the initial set of SLHC HLT algorithms.

VIII. ACKNOWLEDGEMENTS

Many of the ideas in this paper come from presentations by and discussions with a number of individuals, including, D. Acosta, S. Cittolin S. Dasu, N. Ellis, E. Eisenhandler, C. Foudas, G. Hall, G. Heath, S. Marchioro, D. Newbold, C. Seez, P. Sharp, F. Taylor, F. Vasey, and A. Watson.

IX. REFERENCES

- [1] ATLAS first level trigger Technical Design Report, CERN/LHCC/98-14.
- [2] CMS level-1 trigger Technical Design Report, CERN/LHCC/2000-038.
- [3] ATLAS high-level trigger, data acquisition and controls, Technical Design Report, CERN/LHCC/2003-22.
- [4] CMS data acquisition and high-level trigger Technical Design Report, CERN/LHCC/2002-26.
- [5] F. Gianotti, *et al.*, Physics potential and experimental challenges of the LHC luminosity upgrade, CERN-TH-2002-078.
- [6] B.G. Taylor, "TTC Distribution for LHC Detectors", IEEE Trans. Nucl. Sci, 45,3, (1998), pp. 821-828.