



The Compact Muon Solenoid Experiment  
**Conference Report**

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland



11 July 2014 (v4, 23 October 2014)

# A Time-Multiplexed Track-Trigger architecture for CMS

G. Hall, D. Newbold, M. Pesaresi, A. Rose

## Abstract

The CMS Tracker under development for the High Luminosity LHC includes an outer tracker based on “PT-modules” which will provide track stubs based on coincident clusters in two closely spaced sensor layers, aiming to reject low transverse momentum track hits before data transmission to the Level-1 trigger. The tracker data will be used to reconstruct track segments in dedicated processors before onward transmission to other trigger processors which will combine tracker information with data originating from the calorimeter and muon detectors, to make the final L1 trigger decision. The architecture for processing the tracker data is still an open question. One attractive option is to explore a Time Multiplexed design similar to one which is currently being implemented in the CMS calorimeter trigger as part of the Phase I trigger upgrade. The Time Multiplexed Trigger concept is explained, the potential benefits of applying it for processing future tracker data are described and a possible design based on currently existing hardware is presented.

Presented at *WIT2014 WIT2014 Workshop on Intelligent Trackers*

# A Time-Multiplexed Track-Trigger architecture for CMS

---

**G. Hall<sup>a\*</sup>, D. Newbold<sup>b</sup>, M. Pesaresi<sup>a</sup>, A. Rose<sup>a</sup>**

<sup>a</sup>*Blackett Laboratory, Imperial College, London SW7 2AZ, U.K.*

<sup>b</sup>*H.H. Wills Physics Laboratory, University of Bristol, Bristol BS8 1TL, U.K.*

*E-mail: g.hall@imperial.ac.uk*

**ABSTRACT:** The CMS Tracker under development for the High Luminosity LHC includes an outer tracker based on “PT-modules” which will provide track segments based on coincident clusters in two closely spaced sensor layers, aiming to reject low transverse momentum track hits before data transmission to the Level-1 trigger. The tracker data will be used to reconstruct track segments in dedicated processors before onward transmission to other trigger processors which will combine tracker information with data originating from the calorimeter and muon detectors, to make the final L1 trigger decision. The architecture for processing the tracker data is still an open question. One attractive option is to explore a Time Multiplexed design similar to one which is currently being implemented in the CMS calorimeter trigger as part of the Phase I trigger upgrade. The Time Multiplexed Trigger concept is explained, the potential benefits of applying it for processing future tracker data are described and a possible design based on currently existing hardware is presented.

**KEYWORDS:** Trigger detectors; Particle tracking detectors; Trigger concepts and systems .

---

\*Corresponding author.

---

## Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. The Time Multiplexed Trigger</b>	<b>2</b>
2.1 Potential pros and cons of Time Multiplexing	2
2.2 Processing challenges	5
<b>3. Design of a Track-Trigger TMT</b>	<b>6</b>
3.1 Hardware	6
3.2 Possible layout of a CMS TM Track-Trigger	6
3.3 The Time Multiplexing period	8
3.4 Algorithms and firmware	8
3.5 A demonstrator	10
<b>4. Conclusions</b>	<b>11</b>

---

## 1. Introduction

A major upgrade of the CMS experiment at the CERN Large Hadron Collider has been under consideration for several years [1, 2, 3, 4]. The objective is for the accelerator to deliver an integrated luminosity of  $3000 \text{ fb}^{-1}$  over a period of about a decade of operation for a physics programme which will include precision measurements of the newly discovered Higgs boson, further searches for new physics extended to higher masses, and deeper scrutiny of data for possible discrepancies compared to the Standard Model.

The High Luminosity LHC will require a new tracker around 2023 because of gradual deterioration of the present detector due to accumulated radiation damage. Its replacement must satisfy several demanding requirements, including a lower material budget and increased granularity compared to the present detector, and enhanced radiation tolerance combined with tolerable power consumption and affordable cost.

A notable feature of the replacement tracking detector is that, for the first time, data should be provided to the L1 trigger to constrain the trigger rate to 0.5-1 MHz with a latency of up to  $\sim 10 \mu\text{s}$ . The baseline outer tracker design has a “conventional” barrel-endcap layout with two types of modules, denoted “2S” and “PS”, which allow selection of hits consistent with a transverse momentum above a chosen threshold  $\sim 2 \text{ GeV}/c$  and thus reduce significantly the volume of data to be transmitted out of the tracker to be used for the L1 trigger decision. The basic concept is to compare the binary pattern of hit strips on upper and lower sensors of a two-layer module [5, 6] to reject patterns that are consistent with a low transverse momentum track. Hit combinations in the two layers consistent with a high- $p_T$  track segment are known as “stubs”.

In the 2S-module [7, 8] two silicon microstrip sensors are separated by a few mm, with the spacing determined by a compromise between transverse momentum precision and the fake stub rate resulting from combinatorial background caused by hits from, e.g., nearby tracks, secondary interactions and photon conversions. In the cylindrical barrel region, high  $p_T$  tracks can be identified if hit clusters lie within a search window in  $R-\phi$  (rows) in the second layer. The sensor separation and search window determine the  $p_T$  cut, where the objective is simply rejection of a sufficiently large fraction of low  $p_T$  candidates in the detector so that trigger primitives can be transmitted within the available bandwidth. In the barrel region, the sensor segmentation in  $z$  (along the beam axis) determines the vertex measurement precision; the outermost tracker region uses 5 cm strips, so dedicated PS-module layers [9] with finer segmentation in one sensor are planned. Similar considerations apply to the end-cap detectors where the same method can be deployed.

In total,  $\sim 15000$  modules, each with a dedicated fibre-optic link, will send pT-stubs to off-detector processors at a 40 MHz rate. The first problem to be solved outside the Tracker is to use the arriving stub data to reconstruct tracks to then be associated with other trigger primitives from the calorimeter and muon detectors.


## 2. The Time Multiplexed Trigger

The fundamental idea is that multiple sources send their data to a single destination (node) for assembly and event processing. It requires two layers with a static switching network between them, which can be a passive optical fibre network even if a large number of connections are needed. The implementation could include information processing in both layers or could simply involve data organisation and formatting at Layer 1, followed by data transmission to Layer 2, where all event processing occurs. The Time Multiplexing (TM) principle is very similar to what is deployed in the CMS High Level Trigger (HLT) [10] which uses a commercial CPU farm as its Layer 2, while the interconnection fabric makes use of an active network to manage traffic between the two layers. However in the HLT the data traffic scheduling is more complex because the volume of data for each event fluctuates greatly, and the latency requirements are very different than for the L1 trigger.

A comparison between a conventional trigger (CT) and TMT is shown in figs 2 and 1 where, for illustrative purposes, data originate from seven regions of CMS. In the CT, the regional segmentation is maintained in the processing stages which operate synchronously, necessitating sharing between the different regional processors as illustrated by the diagonal lines between them. In the TMT data are routed through a multiplexing network which directs all the data from an individual bunch crossing to a single processor. The processors are not synchronous but are arranged to be one LHC clock cycle out of phase with their neighbours so that data should arrive sequentially in a round robin fashion.

### 2.1 Potential pros and cons of Time Multiplexing

A key element of the TM concept is that all the data from a given LHC bunch crossing (BX) pass through at a single node for processing, which avoids boundaries and sharing of information between processors. However, this does not preclude subdivision of the detector into regions within each processing node provided the boundaries are properly handled to minimise data sharing, which

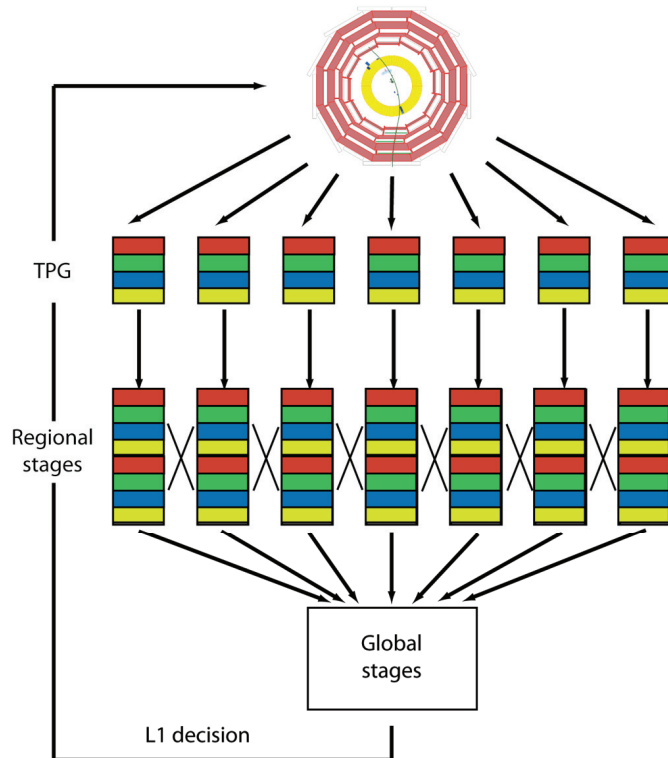


**Figure 1.** In the Time Multiplexed Trigger all data from every TPG are transmitted to a single processing node for each LHC bunch crossing. The processors are therefore out of phase with one another by one LHC clock cycle. The multiplexing fabric may simply be a serial interconnection from each TPG to each TMT node. In this example, each TPG would be equipped with seven serial transmitting links and seven links would arrive at each processor, one from each TPG.

is discussed in more detail later. In fact, for such a large data source as the CMS Tracker it is probably essential to segment the system since the number of links into each processing node is defined by the number of TPG cards, which is in turn defined by the number of detector links, and present FPGAs have insufficient input links to receive all Tracker TPGs into a single device, which is not expected to change.

Actually a more important argument is that the overall processing architecture is well matched to operation of the FPGAs which carry out the processing. FPGAs operate optimally using highly parallel streams with pipelined steps running at data link speed, and it has been found [11] to be very important to adapt the algorithms to the constraints of FPGA operation or else even relatively simple tasks may result in a design which cannot be mapped physically into an FPGA. Many conventional algorithms can overflow the capacity of even a very large FPGA because of timing constraints or routing congestion for two-dimensional algorithms; some examples are given later.

The TMT reduces the requirements on synchronisation, which can be demanding in complex high speed systems. In a CT data can only be reliably exchanged if the processors sharing data are fully in phase. This is achieved by careful serialisation and deserialisation stages, running much faster than the LHC clock frequency, using well defined, high quality clock signals which must be carefully managed throughout the entire system. In contrast, since the TMT processors are



**Figure 2.** A schematic illustration of a regionally organised conventional trigger. Data originate in Trigger Primitive Generators (TPGs) which transmit them to Regional Stages for processing, requiring sharing of data between them to handle boundaries. The regional processors operate in phase with each other. The final trigger decision is taken at a global level, where the results from individual regions can be compared.

essentially operating independently, precise synchronisation is required only between the boards within each node, and not across the entire trigger.

Another advantage of avoiding boundaries or reducing them to a minimum is that only one or two nodes are needed to validate an entire trigger, since each node is carrying out identical processing, but just delayed by one LHC clock cycle compared to its neighbour in a round robin fashion. Additional nodes replicating one of the others can be used for redundancy, or algorithm development. If a hardware failure occurs at one TMT processor, it leads to a temporary loss of efficiency - by a factor 14% in the example of 7 nodes - until the processor is fixed, and a redundant node can be quickly switched in to replace the faulty one. In a CT, the loss of a regional processor affects every bunch-crossing and has more complex ramifications whose impact on physics is not so easy to foresee and the complexity of interconnections could also make replacement a more difficult task.

One possible drawback of the TMT is the time required for transmission of data from Layer 1 to Layer 2; this appears to imply that processing cannot begin in a system with  $N$  nodes until  $N$  clock cycles have elapsed, thus adding to the latency of a time-critical system. However, if the data are properly organised, Layer 2 processing does not need to wait for entire event data to be assembled and the pipelined processing can start as soon as the first cycle's worth of data are present.

As has been observed in the CMS calorimeter trigger, it is easily possible to construct a TMT system using a single type of processing board in both layers, which can be advantageous for procurement and long term maintenance.

Why is it that the Time Multiplexed approach has not been used in previous trigger systems? The reason is mainly that technology limitations have been significantly reduced in recent years with the advent of very large and powerful FPGAs, containing many on-chip high speed serialiser-deserialisers, and the availability of compact, relatively inexpensive, conveniently packaged high speed optical link components.

## 2.2 Processing challenges

The challenge of finding tracks in the high luminosity LHC environment within the L1 trigger latency is formidable and, in any track-trigger processing architecture, must be the subject of considerable investigation in the coming years. For this reason, it is worth looking at previous experience with FPGA algorithm development, even if not directly applicable to track finding, to see where difficulties have arisen. There is so far quite limited experience of developing firmware for the latest generation of FPGAs, such as the Xilinx Virtex-7 family [12], and yet examples of algorithms which appear quite simple and are readily emulated in software but which do not lead to viable firmware code have been identified. They seem to suggest that it is most important from an early stage to focus on algorithms which are well adapted to pipelined parallel processing.

One algorithm which was studied during the development of the CMS calorimeter trigger [13] which led to routing congestion in the logic synthesis was an apparently simple case of a 30 x 36 array of towers ( $\sim 25\%$  of the CMS detector) containing a 10 bit energy measurement where 2 x 2 tower clusters are formed, representing electron candidates, and then the array is searched for 16-cluster entities which could be associated into objects resembling jets. Such a design corresponds to a data throughput of 432 Gbps, excluding encoding, corresponding to about 75% of the available bandwidth into the calorimeter trigger processor board. No firmware apart from the algorithm was defined while, in reality, considerable extra firmware would be needed for sorting, transceiver management, DAQ formatting and transmission, etc. Even in this massively over-simplified case, however, the place and route procedure failed in a Virtex-7 model XC7VX485T because of routing congestion, even though the Look Up Table (LUT) usage was only 29% of the available resources.

In the synthesis of the logic, the routing software identified a dense, highly congested design which it could not overcome. This was addressed in a more realistic example by designing a pipelined jet algorithm for a 56 x 72 tower array in  $\eta$ - $\phi$  space, which is the size of the CMS barrel calorimeter. The 4032 sites were searched individually for a pattern consistent with a circular object above a required energy threshold contained in a 9-tower diameter region. The result with its associated energy was passed to a LUT to determine the status of the object in a subsequent processing stage, e.g. the Global Trigger. This is a more complex problem than the previous example, yet the design was successfully synthesised requiring less than 1% of the FPGA resources and with a very low latency of 9 processing cycles, which would require only 1.5 LHC bunch crossings using an FPGA processing speed of 240 MHz; the design was shown to function successfully at up to 400 MHz clock speed.

The relevance of this example is that by pipelining the data flow in  $\eta$  the algorithm has been effectively reduced from a 2D to a 1D problem, and this reduction in dimensionality by using time



is highly desirable if FPGAs are to be used for track finding.

An even more complex TMT jet algorithm has been developed which searches for a  $9 \times 9$  sum of trigger towers at every site in the calorimeter consistent with a jet signal, including allocating shared signals from overlapping jets and correcting for pile-up. It emulates an algorithm operating in the CMS HLT equivalent to the anti- $k_r$  jet clustering [14] used in CMS full offline event reconstruction. It also provides a pipelined sort of candidates in  $\phi$  and an accumulating pipelined sort of candidates in  $\eta$ , plus energy sums over rings of towers in  $\phi$  with scalar and vector computations of Missing  $E_T$  and  $H_T$ . This is done using less than 45% LUT utilisation including links, buffers, control and DAQ functions in a Virtex-7 XC7VX690T clocked at 240 MHz. The careful planning of the layout of the design and its fully pipelined nature were essential to achieve this. Recent tests allowed the verification of algorithm performance and measurements of the latency, which was well within the maximum allocation of 41 BX.

### **3. Design of a Track-Trigger TMT**

#### **3.1 Hardware**

Rather than speculate too far about what type of processor might be used in the future when this system is actually needed by CMS, we focus on what can be achieved using the most advanced available processors, of which the MP7 [15, 16] is a good example. This board will be used in the current CMS Phase I Trigger upgrade from 2015 [17].

The MP7 is a processing board based on a Xilinx Virtex-7 FPGA and Avago MiniPOD optics in a  $\mu$ TCA format designed to implement Layer 2 of the TMT architecture for the L1 Calorimeter Trigger. The switching fabric is implemented optically via a patch panel. The total optical bandwidth available to each MP7 processor amounts to 0.9 Tbps, provided by 72 links operating at up to 12.5 Gbps for both input and output, and is therefore well matched to a high throughput processing application such as track finding at HL-LHC. A series of MP7 boards have been thoroughly tested during a development and prototyping phase and are now being manufactured for use in CMS where they will be operated using 10 Gbps data links.

#### **3.2 Possible layout of a CMS TM Track-Trigger**

The two stages are envisaged to be comprised of Layer 1 Pre-Processors (PP, or FED in CMS terminology) and Main Processors (MP) at Layer 2. Inputs to the PP layer will be provided by GBT links [18] with data transfer at 3.2 Gbps. The PPs would format event fragments for transmission to the DAQ; format, order and time-multiplex trigger data for onward transfer to the MPs and possibly provide a first stage of trigger processing. Each MP takes links from all PPs as input with the event assembled over a TM period, then algorithms process pipelined data to deliver suitable track information to the further global stages of the L1 Trigger.


Since the CMS tracker will have 15,454 modules, each served by its own optical link,  $\sim 230$  PPs will be required, as the maximum number of input links to the MP7 is 72 which limits the number of Pre-Processor cards which can be connected to the tracker without resorting to an intermediate data compression stage (whose advantages have yet to be evaluated). In fact, 4 bi-directional links are reserved for communications with the DAQ so the maximum number available is 68 for each MP7. Here a data link speed of 10 Gbps is assumed for the conceptual design.



Since each processing node must receive 1 link from each PP, each node must receive at least  $\sim 230$  fibres. This is more than may be received by a single MP7 and so each processing node must be composed of multiple physical cards. It is therefore necessary to geometrically divide the detector into processing regions.

This is done by subdividing into  $\phi$  regions and restricting the problem by estimating the minimum number of required Trigger Regions (TRs), imposing the constraint that one module should not be shared across more than two TRs. This provides a simple solution to the issue of duplicate tracks which may be found in regions which are shared by multiple MPs; if each MP can communicate with only two shared regions a convention can be adopted to accept track candidates which use hits from a shared region only from one side, say “left” or “right”. This avoids a further processing stage which would be required to decide which tracks found in a boundary region should be kept by comparison of the candidates.

Using a software tool which was developed to automate and evaluate alternative layouts of the future tracker [19] it is possible to assign modules to PPs. By conservatively assigning a lower momentum cutoff of 1 GeV/c in the boundary region, it was found that the tracker can be subdivided into  $\phi$  regions using as few as 5 TRs. The 1 GeV/c cut might be tightened, but such a choice might be desirable to allow better reconstruction at 2 GeV/c in case of particle tracks which generate hits outside a simple geometric boundary, such as from bremsstrahlung,  $e^+e^-$  pair production or large multiple scattering events. Fig. 3 indicates, for the barrel region, how the regions were allocated.



**Figure 3.** An illustration of the barrel region of the new Tracker. Modules in blue are allocated to a unique Trigger Region, while data from those in red are duplicated to both of the two neighbouring TRs. This allows the TMTT to handle track candidates which cross TR boundaries and offers a convenient way of assigning the resulting inevitable duplicate tracks, as described in the text.

### 3.3 The Time Multiplexing period

The Time Multiplexing period is not a completely free parameter and a range of values is possible; for this initial evaluation a value of 24 BX has been chosen, which can be optimised following further study; there are possible arguments for larger or smaller values. Note that a longer period implies more processing nodes and a shorter period implies fewer.

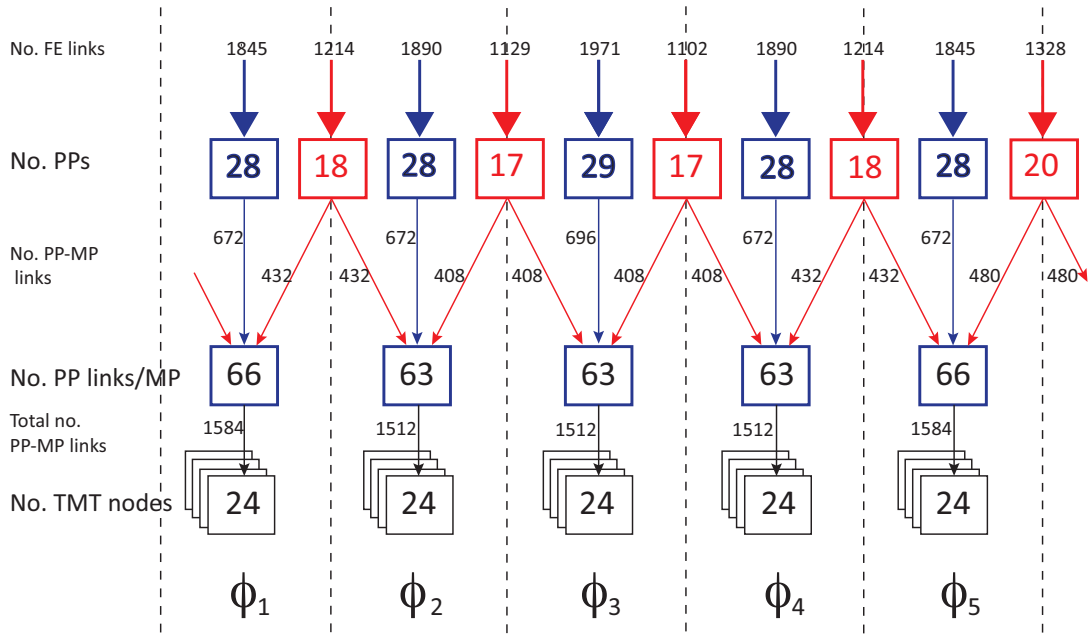
With a small TM period the full event can be quickly assembled into one MP while a longer period could allow more efficient pipelined processing of data into the MP. A small period limits the allowable data volume which can be transferred from PP to MP for each event, which could be increased only by increasing the number of links per PP, which may be physically limited, e.g. by space on the board or power considerations. Clearly increasing the TM period gives rise to converse arguments.

In practical terms, the possible range is from a minimum of  $\sim 15$  BX to a maximum of  $\sim 34$  BX. The minimum value is determined, for a given number of Trigger Regions, by the PP output bandwidth since the total data volume flowing from each TR of the tracker must be delivered to each MP. The calculation should take into account the shared regions as well as the total number of links to each MP, so there is no simple formula and a practical optimisation is required. Similarly the maximum value is mainly driven by the number of links available on each MP7 and the requirement to share them between neighbouring Trigger Regions. It should be borne in mind that the constraints imposed by the MP7 design might easily change with future technology evolution, and the number of Trigger Regions is a choice also decided by realistic factors, such as the handling of duplicates described above. A subdivision into five  $\phi$  regions seems to be a convenient choice for the proposed CMS layout. However, maintaining the maximum of two shared regions convention could still allow more than five sectors, if there were arguments to do so.

The allocation of modules to  $\phi$  regions, or sectors, should be possible independently of the TM period, but the actual distribution of links and their numbers depends on this choice. For the 24 BX value, the result of the allocation of processors and links is shown in fig. 4 where the number of links and boards required to handle 5  $\phi$  Trigger Regions and their interconnections are shown. Typically 28 or 29 Pre-Processors are needed for Trigger Regions which are not shared while shared regions require 17-20 PPs, each of which duplicate their incoming data and transmit to two TRs. The  $\phi$  regions require very similar resources, which is desirable, and it can be seen that 63-66 of the 68 available links are used in each sector. In this scenario the tracker can be read out and trigger data processed using 353  $\mu$ TCA format cards, which can be compared with 440 (much larger) 9U VME FED cards required for a similar number of modules in the present CMS Tracker.

### 3.4 Algorithms and firmware

The purpose of devising a suitable architecture is largely to confront the next major challenge of implementing the CMS track-trigger, which is to establish the best way of finding tracks in the high luminosity LHC environment. Both processing latency and track-finding efficiency are major concerns, so the demonstration of suitable working algorithms successfully programmed into firmware is essential to establish the performance of the system. It is also important since, as previously mentioned, we know from experience with the MP7 and the calorimeter trigger that large FPGAs present serious challenges, including potentially exceeding RAM resources or that the



**Figure 4.** The allocation of links and processors to the five sectors  $\phi_{1-5}$  separated schematically by dashed vertical lines. Data arrive from the Tracker at the top of the figure on 3.2 Gbps links. All other links run at 10 Gbps. PPs which do not share their data are shown in blue; those PPs which share data are in red and hence each have two output paths leading to the MPs. Each of the PPs must transmit their data to all of the 24 MP nodes so the arrows actually represent large numbers of fibres, whose numbers are indicated.

logic place and route task fails to converge within timing constraints after many hours, sometimes for very minor, but nevertheless intractable, reasons. Algorithms devised to work in software may not be easily translatable into firmware logic.

This part of the work is at an early stage and we are currently investigating methods which lend themselves to pipelining and parallel processing. One such example is based on the Hough transform where it is well known that lines in physical space (e.g.  $y = mx + c$ ) can be identified with points  $(m, c)$  in a dual parameter space. This could be applied to incoming data where for each data point  $(x, y)$ , a value of  $m$  is hypothesized and  $c$  calculated. Values of  $m$  and  $c$  are histogrammed into an array, where array elements with significantly more entries than background can be identified with tracks in the two-dimensional space and sent for fitting.

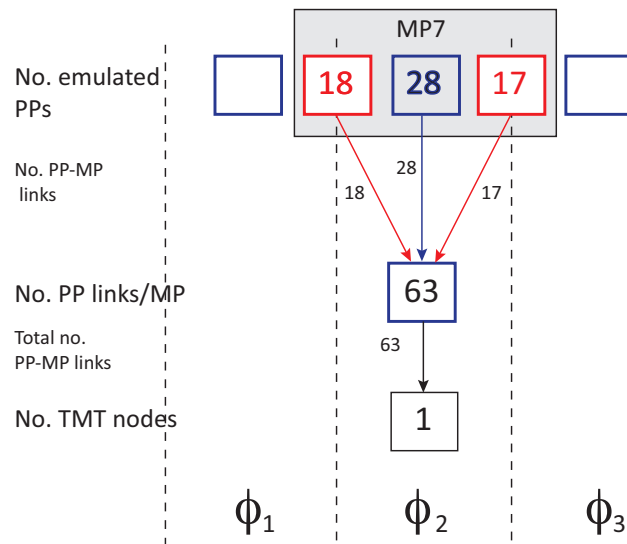
Efficient logic can be defined for this type of problem of populating an array in a fully pipelined manner, with no iterations and only local data transfers. This is then a very realistic method for implementation in an FPGA and would certainly work with sufficient points on a track but it is as yet unclear if the LHC high pile-up conditions will generate too many matching combinations. In any case, this is more than a two-dimensional problem so ways have to be devised to distinguish genuine tracks from increasingly frequent combinatorial background which will occur as the LHC luminosity rises. At this point, it is essential to gain direct experience in a working system, which can readily be done profiting from CMS calorimeter trigger developments.

### 3.5 A demonstrator

One of the very attractive features of the TMT architecture is its great flexibility. This becomes clearer when considering what is required to evaluate a system. Since raw data from the tracker will not be available, the PPs should act as data sources and provide emulated data, which can be stored in on-board memories, to the MPs. The emulated data can range from random bit patterns to simulated Monte Carlo events.

Clearly, a system of  $N$  nodes can be evaluated with  $1/N$ th of the processing cards required for the full system, as every node is performing identical operations on similar but independent data for an individual bunch crossing, where each bunch crossing is uncorrelated with its neighbours. Thus for the 24 TM period system described above, about 10 PPs (using all the output links available on each one) and 5 MPs would be needed. However, this can be further reduced since the MP7 is very flexible and it is desirable to take full advantage of as many of the links available on each module as possible, and the processing power available.

Each sector in  $\phi$  is essentially equivalent and certainly processing data independently of the other sectors. It may eventually be desirable to study in detail the performance of each sector, in case of issues like geometrical acceptance or material budget effects, but at this stage each sector can be regarded as identical. Hence only one fifth of the single TM period system is needed, requiring just 1 MP and 3 PPs, to handle the boundary regions. Since each PP will have enough links to service a single MP, the three logical PPs can be provided by one PP, so the whole TMT system can be studied with as few as two MP7 processors as shown in fig. 5, which can be compared with fig. 4. In a system with regional sharing, this type of simplification is very unlikely to be possible.



**Figure 5.** Three PPs and one MP are sufficient to demonstrate a  $\phi$  slice of the entire system. However, the PPs can be provided by one MP7 since sufficient links and memory are available. Therefore only two MP7s are needed to emulate event data transmission and processing from one out of five Trigger Regions, for one of every 24 BX.

This demonstrator has been built for the CMS calorimeter trigger, including infrastructure

firmware and software, making it easy to redeploy for this application.

## 4. Conclusions

The TMT is now a proven architecture in CMS and will operate in the CMS calorimeter trigger from 2016. It has many attractive features for future trigger applications including great adaptability to external constraints using a minimal number of hardware variants.

The hardware is very flexible and can be deployed for a track-trigger application with only a very small fraction of the final system required to validate the concept. Components already exist using present state-of-the-art technology which have the required performance to build such a system although technological evolution should mean that more powerful systems can be implemented in future.

The next major challenge is to prove that suitable track finding algorithms can be implemented and their performance evaluated.

## Acknowledgments

We gratefully acknowledge financial support from the UK Science and Technology Facilities Council. We thank our colleagues Greg Iles and John Jones for important contributions to the development of the Time Multiplexed Trigger concept and its realisation.

## References

- [1] The CMS Collaboration *CMS Expression of Interest in the SLHC* CERN/LHCC 2007-014 (2007)
- [2] The CMS Collaboration *Technical Proposal for the Upgrade of the CMS Detector through 2020* CERN-LHCC-2011-006 (2011).
- [3] M. Raymond et al. *The CMS binary chip for microstrip tracker readout at the SLHC* 2012 JINST 7 C01033 doi:10.1088/1748-0221/7/01/C01033
- [4] The CMS Collaboration *The CMS experiment at the CERN LHC* 2008 JINST 3 S08004 doi:10.1088/1748-0221/3/08/S08004
- [5] G. Hall, M. Raymond and A. Rose *2-D PT module concept for the SLHC CMS tracker* 2010 JINST 5 C07012 doi:10.1088/1748-0221/5/07/C07012
- [6] M. Pesaresi and G. Hall *Simulating the performance of a pT tracking trigger for CMS* 2010 JINST 5 C08003 doi:10.1088/1748-0221/5/08/C08003
- [7] G. Hall *CBC2: a CMS microstrip readout ASIC with logic for track-trigger modules at HL-LHC* NIM paper doi: tbd
- [8] D. Braga et al. *Beam test performance of a prototype 2S-module for the High Luminosity Upgrade of the CMS Strip Tracker* Proceedings of WIT 2014. To be published in JINST.
- [9] D. Ceresa et al. *Macro Pixel ASIC (MPA): The Readout ASIC for the Pixel-Strip (PS) module of the CMS Inner Tracker at HL-LHC* Proceedings of WIT 2014. To be published in JINST.
- [10] G. Bauer et al. *The CMS High Level Trigger System: Experience and Future Development* Journal of Physics:ConferenceSeries 396 (2012) 012008 doi:10.1088/1742-6596/396/1/012008

- [11] R. Frazier, S. Fayer, G. Hall, C. Hunt, G. Iles, D. Newbold, A. Rose *A demonstration of a Time Multiplexed Trigger for the CMS experiment* 2012 JINST 7 C01060  
doi:10.1088/1748-0221/7/01/C01060
- [12] Xilinx Inc. *The Virtex-7 FPGA Family*  
<http://www.xilinx.com/products/silicon-devices/fpga/virtex-7/index.htm>
- [13] G. Iles *Private communication*
- [14] M. Cacciari, G. P. Salam and G. Soyez *The anti- $k_r$  jet clustering algorithm* JHEP04 (2008) 063  
doi:10.1088/1126-6708/2008/04/063
- [15] K. Compton, A. Rose et al. *The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS* 2012 JINST 7 C12024. doi:10.1088/1748-0221/7/12/C12024  
<http://www.hep.ph.ic.ac.uk/mp7>, Imperial College London.
- [16] G. Iles, J. Jones, A. Rose *Experience powering Xilinx Virtex-7 FPGAs* 2013 JINST 8 C12037  
doi:10.1088/1748-0221/8/12/C12037
- [17] A. Tapper and D. Acosta (eds.) *CMS Technical Design Report for the Level-1 Trigger Upgrade* CERN-LHCC-2013-011 ; CMS-TDR-12 <http://cds.cern.ch/record/1556311>
- [18] P. Moreira et al. *The GBT Project* Proceedings of Topical Workshop on Electronics for Particle Physics (TWEPP 2009) doi:10.5170/CERN-2009-006.342 <http://cds.cern.ch/record/1235836>
- [19] S. Mersi, D. Abbaneo, N. De Maio, G. Hall *CMS Tracker Layout Studies for HL-LHC* Physics Procedia 37 (2012) 1070 doi:10.1016/j.phpro.2012.03.729