

Trigger level track reconstruction in CMS with a fully time-multiplexed architecture using a Hough transform implemented in an FPGA

August 3-10, 2016
Chicago

C. Amstutz¹, F. A. Ball², M. N. Balzer¹, J. Brooke³, L. Calligaris^{4,5}, D. Cieri^{2,4,5}, E. J. Clement¹, G. Hall³, T. R. Harbaum¹, K. Harder⁴, P. R. Hobson⁶, G. M. Iles³, T. James³, K. Manolopoulos⁴, T. Matsushita⁷, A. D. Morton⁶, D. Newbold^{2,4}, S. Paramesvaran², M. Pesaresi³, I. D. Reid⁶, A. W. Rose³, O. Sander¹, T. Schuh¹, C. Shepherd-Themistocleous⁴, A. Shtipliyski³, S. P. Summers³, A. Tapper³, I. Tomalin⁴, K. Uchida³, P. Vichoudis⁸, M. Weber¹ **for the CMS collaborators**

¹Karlsruhe Institute of Technology, Karlsruhe, Germany

³Imperial College London, London, UK

⁶Brunel University, London, Uxbridge, UK

⁸CERN, Geneva, Switzerland

²University of Bristol, Bristol, UK

⁴Rutherford Appleton Laboratory, Didcot, UK

⁷Österreichische Akademie der Wissenschaften, Vienna, Austria

⁵Supported by the EU FP7-PEOPLE-2012-ITN project nr 317446, INFIERI, "Intelligent Fast Interconnected and Efficient Devices for Frontier Exploitation in Research and Industry"

The High-Luminosity LHC

The **Large Hadron Collider** will be upgraded to increase its peak instantaneous luminosity up to values between 5×10^{34} and $7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. This upgrade is referred to as **High Luminosity LHC** (HL-LHC). The number of concurrent p-p collisions (**pile-up**) in a bunch crossing will reach values between **140** and **200**, significantly larger than the design conditions for the present detectors.

The **CMS collaboration** plans to upgrade its detector (**Phase-2 upgrade**), to attain good physics performance in the conditions expected at the HL-LHC.




Fig. 1: Doublets of hits detected by the upgraded CMS silicon tracker in a busy event with 140 pile-up collisions.

The current implementation of the first stage (**L1**) of the **trigger** system would experience a large performance degradation in HL-LHC conditions. CMS plans to improve this **performance** by **reconstructing charged particle tracks in hardware**, using this information in the L1 decision.

An overview of the various proposals for such a track trigger implementations is presented in this conference in a poster by Sudha Ahuja titled "Level-1 track trigger for the upgrade of CMS detector at HL-LHC".

Track reconstruction at L1 implemented in FPGA

We propose a **track reconstruction system** implemented in firmware running on **FPGAs**.

The system uses a **Hough Transform** (HT) in the **r-φ plane** for finding tracks, a seeded projective **filter** in the **r-z plane** to reduce combinatorics and remove stubs misassigned to tracks, and further filtering stages to remove **duplicate** tracks identified by the Hough Transform. A final stage uses a **fitting** algorithm to extract the parameters of the trajectory from the collection of stubs assigned to a track. The system employs **time multiplexing** to spread load and latency requirements over multiple copies of the reconstruction system.

A hardware **demonstrator** for the system (Fig. 9) is under development. It is implemented on a chain of **MP7** processor cards and aims to reconstruct tracks in one octant in φ of the CMS Outer Tracker (Fig. 9) with a time multiplex factor equal to 36, that is, processing one LHC event every 36. We expect that in the future more capable FPGAs will allow to fit the entire system on one card. The stubs fed into the system are generated by a Monte Carlo simulation of the conditions and detector configuration expected at the HL-LHC.

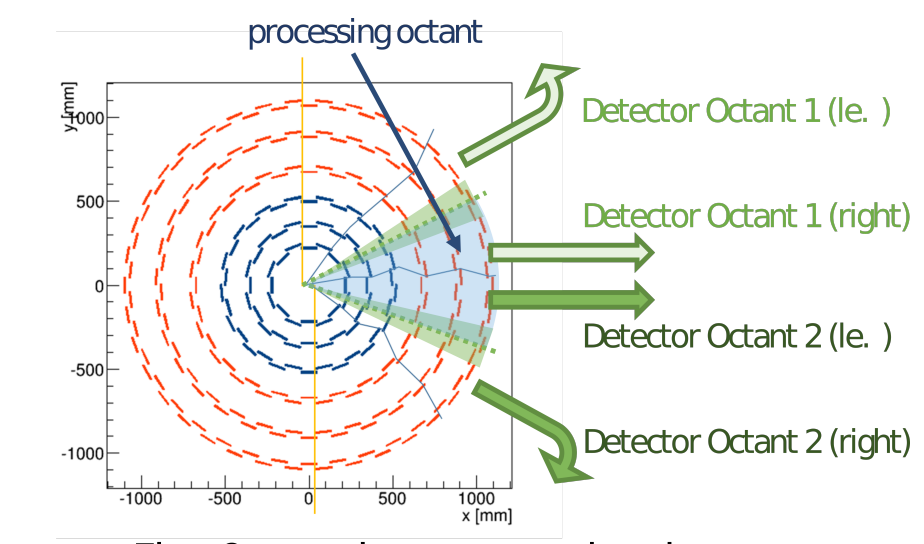


Fig. 9: tracker octant in rφ

Two MP7 cards ("**SOURCE**") act as data source, emulating upstream electronics from two half-octants. An **MP7** implements a geometric processor ("**GP**"), sorting stubs in 36 subdivisions of the octant (4 in φ times 9 in η) and assigns them to independent processing segments running on the 3 MP7 Hough Transform ("**HT**") track finders (each currently fitting 12 segments). One card ("**R-Z**") collects candidate tracks found by the HT and with a seeded projective algorithm filters out stubs that are not consistent with the track in the r-z plane and combinatorial fake tracks. Two or more cards ("**DR/TF**") then remove duplicate tracks, that is, sets of similar tracks sharing stubs and arising from the discrete binning of the HT histogram, and perform a fit of the helix parameters. Finally a card acts as data sink, saving the output for analysis. All connections between cards are established through the MP7 optical infrastructure.





Fig. 8: demonstrator layout

MP7: Imperial Master Processor, Virtex-7

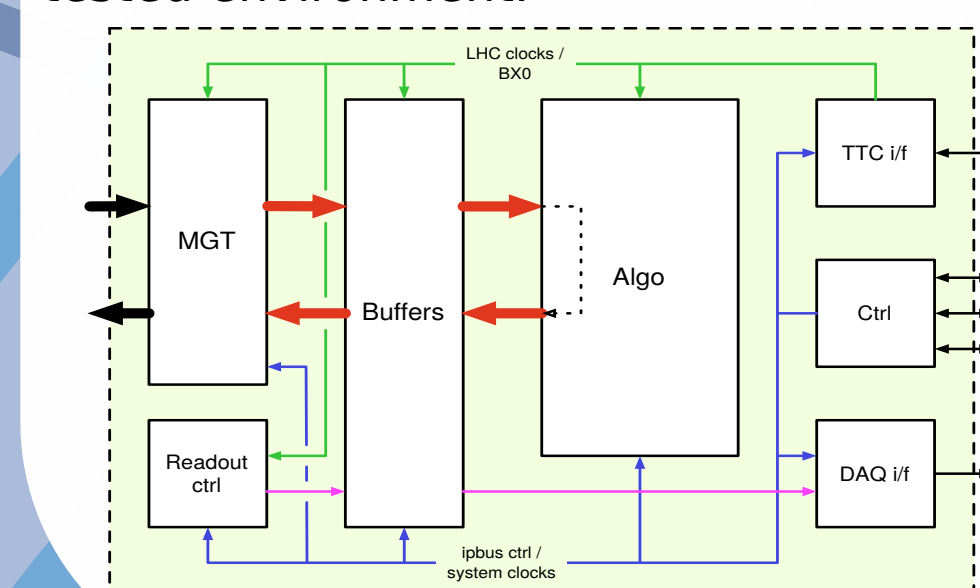
The **MP7** is a generic stream processor card adhering to the **μTCA** standard. The latest version (MP7-XE), used in our project, employs an onboard **Xilinx Virtex-7** XC7VX690T FPGA.



The card is equipped with 72 input and 72 output **optical serial connections**, each running at 10.3 Gb/s. Connection to the μTCA backplane allows communication via **Ethernet** and the programming of the card using the **IPBus** protocol.

The MP7 is a mature and tested product. It is **currently** used to run parts of the **calorimeter** and **muon** triggers in CMS.

An important feature of the MP7 common firmware infrastructure is that it **isolates** the **payload** algorithm away from details of the hardware. This feature eases development of new algorithms by allowing to swap them inside a well-tested environment.



The Upgraded CMS Outer Tracker

The Phase-2 outer tracker is composed of **double layers** of silicon detectors (Fig. 2).

Correlation logic (Fig. 3) in the detector electronics selects pairs of **hits compatible** with the bend of a track with a $p_T > 2-3 \text{ GeV/c}$. These hit pairs (named **stubs**) are sent off-detector for L1 track reconstruction. This selection reduces the bandwidth to L1 electronics.

The stub bend information is sent off-detector, so it can be used as a rough estimate of p_T in downstream filtering stages.

Detector buffers across CMS are designed for a **maximum** L1 trigger **latency** of **12.5 μs**, so the reconstruction of tracks has to take place in a **few μs**.




Fig. 2: r-z quadrant of the silicon outer tracker




Fig. 3: Stubs are pairs of hits compatible (green band) with a $p_T > 3 \text{ GeV/c}$ track.

Hough Transform track finding

The **trajectory** of a charged particle in the magnetic field of CMS (which is aligned to the z-axis) is bent in the **r-φ plane** (Fig. 5). If the radius of curvature is large compared to the size of the tracking volume, as expected for high p_T tracks, the following relation holds for stubs:

$$\phi_T \approx \phi_{\text{stub}} - C \cdot (q/p_T) \cdot (r_{\text{stub}} - T)$$

where T is a reference radius (in our case 58cm), ϕ_T is the azimuthal angle of the track at radius T , (q/p_T) is the charge to p_T ratio, C a constant proportional to the magnetic field, r_{stub} and ϕ_{stub} are the stub radius and azimuthal angle.

The linear relation means that **stubs** can be represented by straight **lines**, as shown in Fig. 6. If a set of stubs is consistent with a real track, they will **meet** at the coordinates of this track.

Track finding uses a **histogram** that **counts** the number of stub lines that have crossed each bin. Parts of a stub line which have a (q/p_T) ratio not compatible with the strip distance measured by the double layer (see Fig. 3) are excluded. Bins associated with stubs in at least **five** distinct tracker **layers** are identified as candidate tracks.

Three different firmware architectures have been tested to implement the histogram, with increasing improvements in FPGA resources utilization. The latest design, which we refer to as "**Daisychain**", currently used in development, is able to fit **up to 18 segments** in one MP7 card, meeting timing requirements.




Fig. 5




Fig. 6




Fig. 7

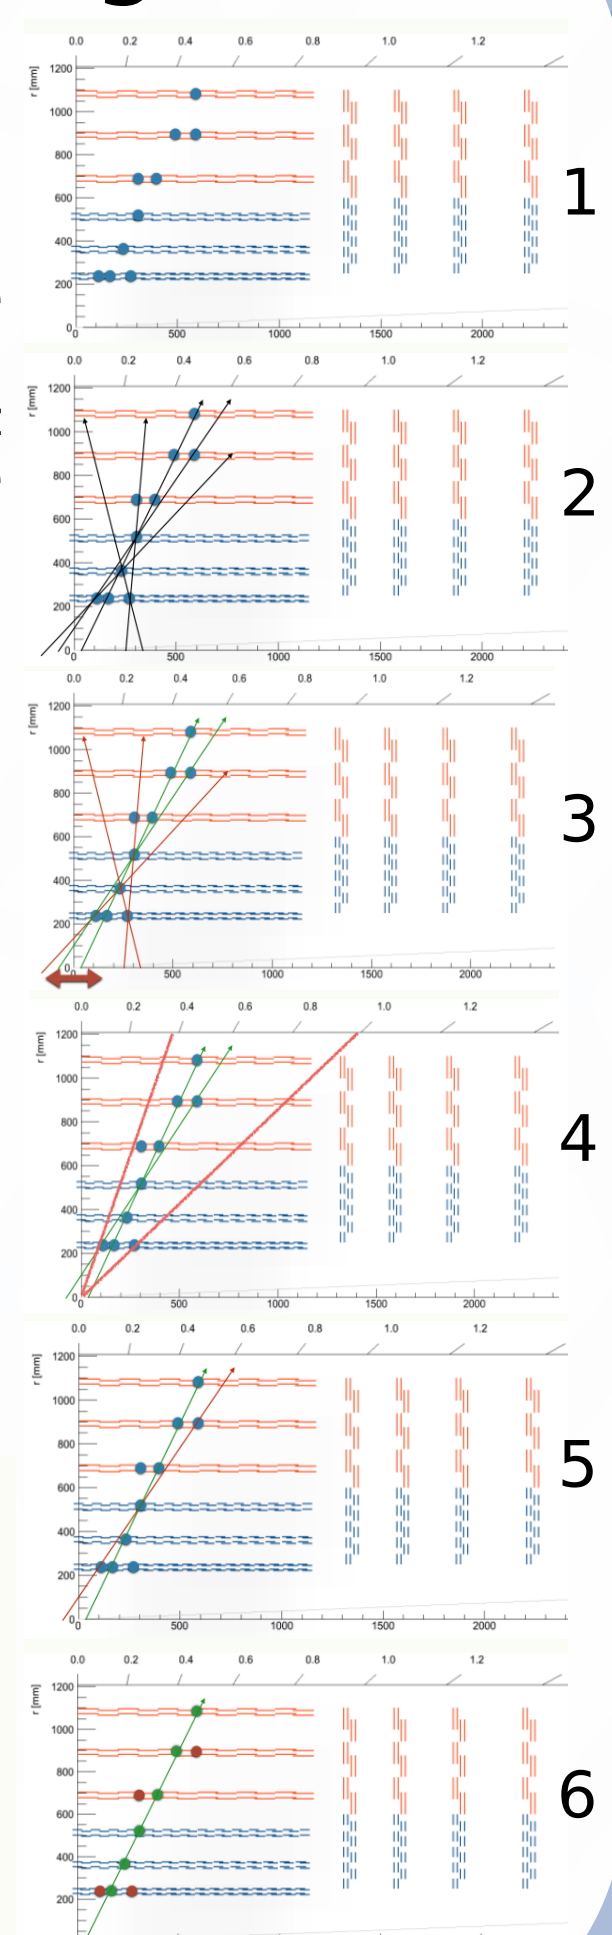

Projective track filtering in rz

The Hough Transform in r-φ is capable of finding tracks with high efficiency, but is not robust against fake tracks arising from combinatorial stub background in busy environments.

An **r-z filter** downstream of the HT checks the consistency between **stubs** and their corresponding track by verifying if they are **aligned** in the r-z plane. We chose to use a seeded projective algorithm for the filter implementation:

Once all stubs are received (1) all pairs of stubs in the innermost layers act as seed for candidate tracks (2). Candidates not compatible with the beam spot (3) and pseudorapidity boundaries of the segment (4) are pruned. Among the remaining track candidates (5), those which are built from stubs in at least 5 layers are accepted (6).

In simulation, the purity of tracks (i.e. tracks matching to Monte Carlo truth, including duplicates) before and after the r-z filter has been studied. Results show that the filter is **successful** at reducing the number of **fake tracks** in top pair events, both with 140 and 200 pile-up collisions. Firmware implementation is now in development.



Hardware Demonstrator

We are currently running **two setups** employing MP7 cards: a smaller setup is based at **Rutherford Appleton Laboratory** and is used for firmware development, a larger one is based at **CERN** and will be used to implement in hardware a complete demonstrator for our approach.

Currently, the firmware components of the demonstrator are being integrated on a total of 11 MP7 cards, as shown in Fig. 10 below. The cards can be managed and run **remotely** via the Ethernet connection. Seven cards are being cascaded to work together, following the scheme shown in Fig. 9 up to the Hough Transform track finder stage, with the last card acting as data sink. The remainder of the cards are currently run in isolation to develop firmware for the filters and the track fitter and will be connected in the processing chain later on.




Fig. 10

Data read out from either the sink cards or the internal MP7 data buffers is compared with the expected values from a **simulation** implemented in CMS Software (**CMSW**) and a firmware **emulation** implemented in the **CIDAF** emulation framework.

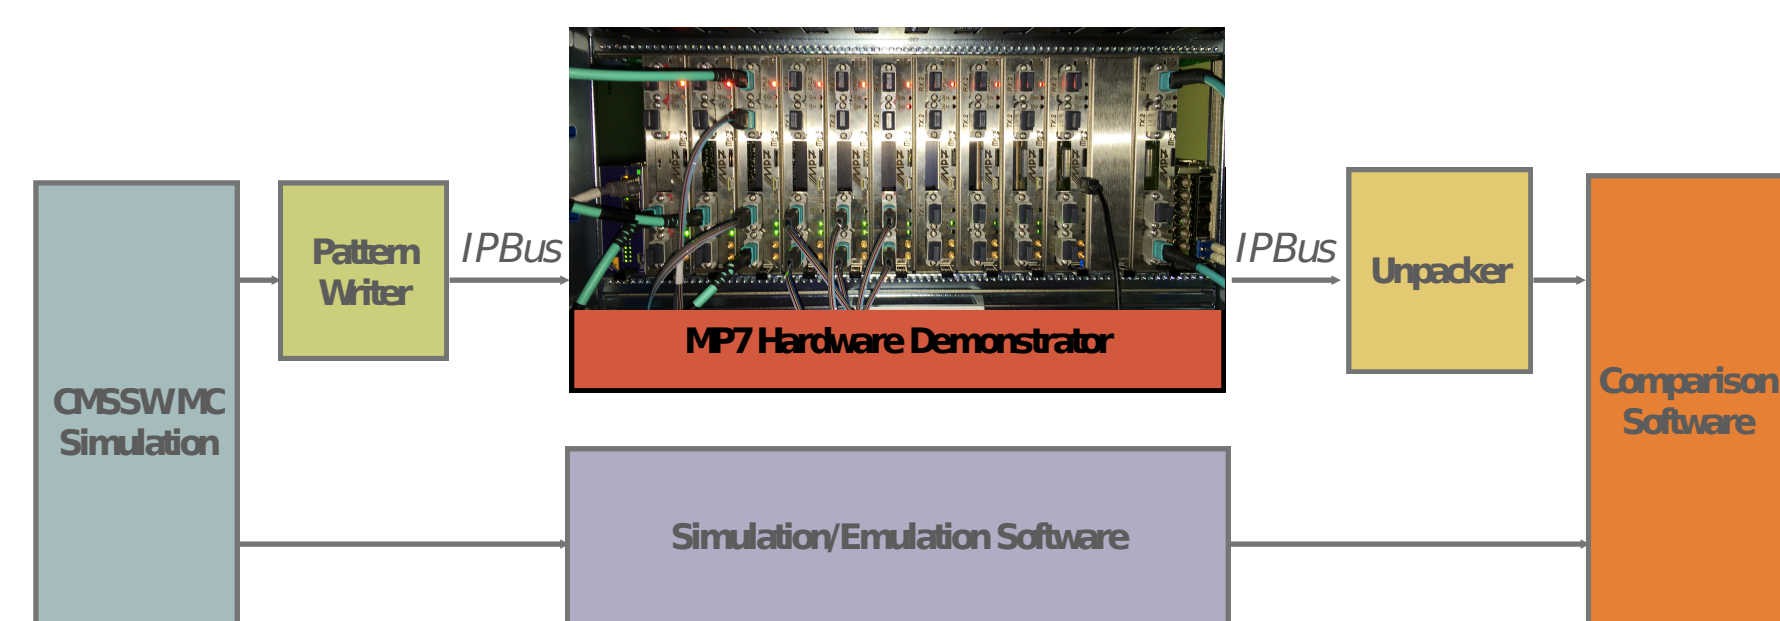


Fig. 11

Results

On the top row, the figures show the number of tracks found per q/p_T bin (left), the tracking efficiency per unit of p_T (center) and pseudorapidity (right) - as measured from Monte Carlo truth - in the case of an earlier version of the HT track finder. An **improved version** (bottom row) of the HT reduces the resources/latency load on the filtering stages, by reducing the **number of candidate** tracks by approximately a **factor of two**. This is accomplished by halving the size of the bins on the q/p_T axis for tracks $p_T > 6 \text{ GeV/c}$. The number of candidate tracks is successfully reduced, while the **tracking efficiency** does not change appreciably. The simulation software generally reproduces behavior in hardware with a good degree of agreement.

