Available on CMS information server CMS CR -2018/073

The Compact Muon Solenoid Experiment
Conference Report
Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

11 June 2018 (v3, 29 June 2018)

The algorithm of the CMS Level-1 Overlap Muon Track Finder trigger

Karol Bunkowski for the CMS Collaboration

Abstract

The CMS Level-1 trigger, based on custom electronics built around FPGA devices, was upgraded in 2016 to achieve the required performance with almost two times higher LHC luminosity than originally designed. The upgraded Level-1 muon trigger merges data from the three muon detectors in the CMS (DT, CSC, and RPC) in the track reconstruction stage – contrary to the legacy trigger which comprised three separate systems each based on one muon detector. This approach allows better use of the detector redundancy and improved measurement of the muon transverse momentum. However, it is particularly challenging in the barrel-endcap transition region, where up to 18 muon chamber layers are present, the detector geometry is complex, and the magnetic field bending the muon tracks is heterogeneous. For this region the Overlap Muon Track Finder (OMTF) uses a novel algorithm based on a naive Bayes classifier. The algorithm identifies the muon tracks and measures their momentum by calculating the probabilities of matching the detector hits to defined transverse momentum hypotheses. The algorithm was tailored to the needs of muon measurements based on the detector data (e.g. to cope with missing hits or multi-muon events) and implementation in FPGA technology. The algorithm details, performance, and optimization methods are discussed in this report.

Presented at *ELBA2018 14th Pisa meeting on advanced detectors*

The algorithm of the CMS Level-1 Overlap Muon Track Finder trigger

K. Bunkowski[a,*], on behalf of the CMS Collaboration

[a] *Institute of Experimental Physics, Faculty of Physics, University of Warsaw, Poland*

Abstract

The CMS Level-1 trigger, based on custom electronics built around FPGA devices, was upgraded in 2016 to achieve the required performance with almost two times higher LHC luminosity than originally designed. The upgraded Level-1 muon trigger merges data from the three muon detectors in the CMS (DT, CSC, and RPC) in the track reconstruction stage – contrary to the legacy trigger which comprised three separate systems each based on one muon detector. This approach allows better use of the detector redundancy and improved measurement of the muon transverse momentum. However, it is particularly challenging in the barrel-endcap transition region, where up to 18 muon chamber layers are present, the detector geometry is complex, and the magnetic field bending the muon tracks is heterogeneous. For this region the Overlap Muon Track Finder (OMTF) uses a novel algorithm based on a naive Bayes classifier. The algorithm identifies the muon tracks and measures their momentum by calculating the probabilities of matching the detector hits to defined transverse momentum hypotheses. The algorithm was tailored to the needs of muon measurements based on the detector data (e.g. to cope with missing hits or multi-muon events) and implementation in FPGA technology. The algorithm details, performance, and optimization methods are discussed in this report.

Keywords: Trigger, FPGA
PACS: 29.85.Ca

1. Introduction

Compact Muon Solenoid (CMS) [1] is one of the experiments at the Large Hadron Collider (LHC) at CERN. A crucial element is the trigger system which analyses the detector data in real time and selects only potentially interesting events for permanent storage. The trigger is composed of two stages: Level-1 is implemented in custom electronics based on FPGA (field-programmable gate array) devices, and reduces the event rate from 40 MHz to approximately 100 kHz. The High Level Trigger is dedicated software running on a computing farm, which selects $O(1k)$ of these events per second. The Overlap Muon Track Finder (OMTF) is part of the upgraded Level-1 Muon Trigger [2], and processes muon data from the detector region $0.8 < |\eta| < 1.24$. The main motivation of the upgrade was to achieve purer event selection (compensating the higher LHC luminosity) while keeping high efficiency for the single and double muon triggers, with similar transverse momentum (p_T) thresholds as the legacy trigger - which is crucial for the physics performance of CMS. In practice this requires more accurate measurement of the muon p_T and a low fake rate from non-muon backgrounds.

2. Naive Bayes classifier in the OMTF algorithm

The original idea of the OMTF algorithm was described in [3] from the perspective of implementation in FPGA devices. In this report we focus on the statistical approach of the OMTF algorithm and applied optimizations.

The OMTF algorithm performs the track identification and muon p_T measurement in one step. It can be considered as a naive Bayes classifier - a classic machine learning algorithm [4]. The muon p_T spectrum is divided into 52 bins (or classes) k, 26 for each charge, and the track candidates are assigned to the p_T^k class that is the most likely. The logic unit that performs the likelihood calculation for a given p_T^k is called the “golden pattern”.

According to Bayes theorem, the conditional probability that a muon has a given p_T^k provided detector hits x_i in L layers is:

$$P(p_T^k \mid x_1,\ldots,x_L) = \frac{P(x_1,\ldots,x_L \mid p_T^k)P(p_T^k)}{P(x_1,\ldots,x_L)} = \frac{\prod_{i=1}^{L} P(x_i \mid p_T^k)P(p_T^k)}{P(x_1,\ldots,x_L)}$$

Since the denominator in the above equation is the same for all p_T^k classes (if the same hits are used), it is not important for selecting the highest $P(p_T^k \mid x_1,\ldots,x_L)$ and can be dropped.

The $P(x_1,\ldots,x_L \mid p_T^k)$ can be written as $\prod_{i=1}^{L} P(x_i \mid p_T^k)$ only if the x_i are independent (not correlated). The magnetic field produced by the CMS solenoid bends charged particles in the $r-\phi$ plane (i.e. perpendicular to the beam axis), so precise measurement of the muon hit ϕ positions allow the p_T to be determined. The absolute ϕ position of the muon hits are highly correlated and cannot be used as x in the above equation. Therefore the distance between each hit and one chosen reference hit is used: $x_i = \phi_i^{dist} = \phi_i - \phi^{ref} - \Delta\phi_i^{mean}(p_T^k)$. The $\Delta\phi_i^{mean}(p_T^k)$ is

*Corresponding author
Email address: kbunkow@cern.ch (K. Bunkowski)

the average ϕ distance for a given p_T^k in a layer i, so adding this term centers the $P(\phi_i^{dist} \mid p_T^k)$ at 0 and thus reduces its width, which is important for storing it into the look-up table (LUT), as described below.

The first problem that arises with this approach is that more than one hit can happened in a given event in some layers (due to noise or photon emission by a muon). Therefore in each "golden pattern" layer among all available hits only one with the smallest ϕ_i^{dist} is selected.

It is also possible that in some layers there are no hits at all (because inefficiency or geometrical acceptance of a chamber). For that reason the $P(\phi_i^{dist} \mid p_T^k)$ is included only for the detector layers which have a hit. Then, the "golden patterns" with the highest number of the fired layers are selected, and finally the one with the highest total likelihood is chosen. Candidates with less then 3 active layers are rejected – this reduced significantly the rate of fake triggers or candidates with overestimated p_T.

The next issue is related to the reference hit: if the reference hit is a fake (e.g. noise) then the ϕ_i^{dist} are not correct and the algorithm either finds no muon or the measured p_T is incorrect. Additionally, more then one muon can appear in the same event in the detector region handled by one processor. To handle such cases the track finding for each event is executed four times, with a different reference hit in each iteration (the reference hits are chosen from 8 set layers according to a predefined priority). Of course this can produce duplicate muon candidates. They are removed in the last stage of the algorithm by comparing the reconstructed ϕ of the candidates.

Calculation of the $\prod_{i=1}^{L} P(\phi_i^{dist} \mid p_T^k) P(p_T^k)$ requires many floating point multiplications, which in FPGA would be slow and consume excessive resources. Instead of multiplication, the sum of the likelihood logarithms is used. The values of $\log P(\phi_i^{dist} \mid p_T^k)$ are stored in LUTs implemented in the FPGA block-RAM (BRAM) modules. The BRAMs of the Xilinx Virtex-7 FPGA (which is used in the OMTF) can be used as look-up tables with 10 bits of address and 18 bits of output values. Just 7 bits are enough to encode the ϕ_i^{dist} (which is used to address the LUT), and the remaining 3 bits are used to encode the number of the reference layer, since $\log P(\phi_i^{dist} \mid p_T^k)$ depends on the layer of the reference hit in a given iteration.

For the output value, 6 bits are enough to encode the $\log P(\phi_i^{dist} \mid p_T^k)$ – another advantage of using the logarithm, as it reduces significantly the required range of values. Thus it is possible to store in one BRAM up to three $\log P(\phi_i^{dist} \mid p_T^k)$ distributions of the same layer of different "golden patterns" with neighboring p_T values – of course then the same address (i.e. ϕ_i^{dist}) is used for them – but this is next advantage, because then the logic selecting the hits with minimal ϕ_i^{dist} can be common for these "golden patterns". These optimizations reduce significantly the amount of the FPGA logic and BRAMs that are needed to implement the algorithm.

The distributions $P(\phi_i^{dist} \mid p_T^k)$ for each reference layer were obtained from the Monte Carlo simulation.

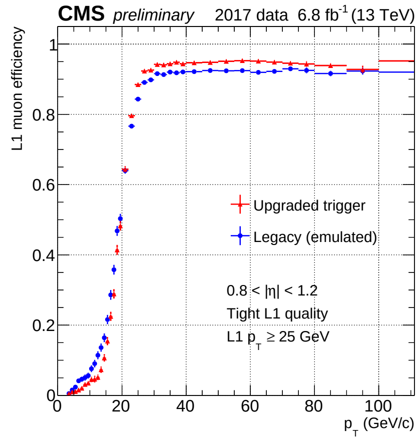

Figure 1: Trigger efficiency vs. offline muon p_T in the overlap region, comparing the upgraded (2017) and legacy (2015) systems. The plot shows the performance for the main single muon trigger threshold used in CMS i.e. $p_T > 25\ GeV$ [5].

3. OMTF performance and conclusions

The L1 Muon Trigger efficiency was measured with the tag-and-probe method based on proton-proton collision data collected in 2017. Events were triggered by the single isolated muon trigger, but contained two reconstructed muons; the muon which fired the trigger was used as a tag, the second muon as a probe to measure the L1 trigger efficiency. The efficiency for a trigger requiring Level-1 muon $p_T > 25\ GeV$ in the OMTF region is presented on Figure 1. It is compared with the efficiency curve for the legacy L1 Muon Trigger in the same η region. This curve was obtained by running realistic emulation of the legacy trigger on the same set of 2017 data. Above the threshold the OMTF efficiency is about 2% higher then the legacy trigger. Additionally the efficiency below the threshold is lower for the OMTF – this results in 25% lower trigger rate for p_T thresholds above 15 GeV [5]. The results show that the OMTF provides significantly better performance than the legacy trigger in the same region.

Acknowledgments

Work supported by National Science Centre, Poland, grant UMO-2015/19/B/ST2/02861

References

[1] CMS collaboration, The CMS experiment at the CERN LHC, 2008 JINST 3 S08004. Also published by CERN Geneva in 2010.

[2] CMS collaboration, CMS Technical Design Report for the Level-1 Trigger Upgrade, CERN-LHCC-2013-011, CMS-TDR-12 (2013).

[3] W.M. Zabolotny and A. Byszuk, Algorithm and implementation of muon trigger and data transmission system for barrel-endcap overlap region of the CMS detector, JINST 11 (2016) C03004.

[4] I. Rish, Irina, An Empirical Study of the Nave Bayes Classifier, IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence 3.

[5] CMS Collaboration, Level-1 muon trigger performance in 2017 data and comparison with the legacy muon trigger system, CMS-DP-2017-041, (https://cds.cern.ch/record/2286327).