Available on CMS information server

CMS CR -2017/361

The Compact Muon Solenoid Experiment

Conference Report

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

11 October 2017 (v3, 15 October 2017)

Boosted Decision Trees in the CMS Level-1 Endcap Muon Trigger

Darin Acosta, Andrew Brinkerhoff, Elena Busch, Andrew Carnes, Ivan Furic, Sergei Gleyzer, Khristian Kotov, Jia Fu Low, Alexander Madorsky, Jamal Rorie, Bobby Scurlock and Wei Shi for the CMS Collaboration

Abstract

The first implementation of Boosted Decision Trees (BDTs) inside a Level-1 trigger system at the LHC is presented. The Endcap Muon Track Finder (EMTF) at CMS uses BDTs to infer the momentum of muons in the forward region of the detector, based on 25 different variables. Combinations of these variables are evaluated offline using regression BDTs, whose output is stored in 1.2 GB look-up tables (LUTs) in the EMTF hardware. These BDTs take advantage of complex correlations between variables, the inhomogeneous magnetic field, and non-linear effects such as inelastic scattering to distinguish high-momentum signal muons from the overwhelming low-momentum background. The LUTs are used to turn the complex BDT evaluation into a simple look-up operation in fixed low latency. The new momentum assignment algorithm has reduced the trigger rate by a factor of 3 at the 25 GeV trigger threshold with respect to the legacy system, with further improvements foreseen in the coming year.

Presented at *TWEPP 2017 Topical Workshop on Electronics for Particle Physics 2017*

PROCEEDINGS OF SCIENCE

Boosted Decision Trees in the CMS Level-1 Endcap Muon Trigger

Darin Acosta, Andrew Brinkerhoff, Andrew Carnes, Ivan Furic, Sergei Gleyzer, Khristian Kotov, Jia Fu Low*, Alexander Madorsky and Bobby Scurlock
University of Florida
E-mail: darin.acosta@cern.ch, andrew.wilson.brinkerhoff@cern.ch, andrew.mathew.carnes@cern.ch, ivan.kresimir.furic@cern.ch, sergei.gleyzer@cern.ch, khristian.kotov@cern.ch, jia.fu.low@cern.ch, alexander.madorsky@cern.ch, bslock@phys.ufl.edu

Elena Busch, Jamal Rorie and Wei Shi
Rice University
E-mail: elb8@rice.edu, jrorie@rice.edu, ws13@rice.edu

on behalf of the CMS Collaboration

The first implementation of Boosted Decision Trees (BDTs) inside a Level-1 trigger system at the LHC is presented. The Endcap Muon Track Finder (EMTF) at CMS uses BDTs to infer the momentum of muons in the forward region of the detector, based on 25 different variables. Combinations of these variables are evaluated offline using regression BDTs, whose output is stored in 1.2 GB look-up tables (LUTs) in the EMTF hardware. These BDTs take advantage of complex correlations between variables, the inhomogeneous magnetic field, and non-linear effects such as inelastic scattering to distinguish high-momentum signal muons from the overwhelming low-momentum background. The LUTs are used to turn the complex BDT evaluation into a simple look-up operation in fixed low latency. The new momentum assignment algorithm has reduced the trigger rate by a factor of 3 at the 25 GeV trigger threshold with respect to the legacy system, with further improvements foreseen in the coming year.

Topical Workshop on Electronics for Particle Physics
11 - 14 September 2017
Santa Cruz, California

*Speaker.

1. Introduction

The CMS Level-1 (L1) trigger system is responsible for selecting 100k interesting events for physics studies out of 40 million proton-proton collisions occurring every second at the LHC. These interesting events are typically characterized by the presence of leptons, photons, or hadrons with momentum as low as 10-20 GeV. The goal of the L1 trigger is to maintain at least 90% efficiency for such events by keeping low trigger thresholds while rejecting the overwhelming background. The L1 trigger has to make the trigger decision in less than 4 μs to cope with the rate of the incoming collision data. In order to meet this latency requirement, the CMS L1 trigger system is made of custom hardware processors such as FPGAs that are capable of massively parallel data processing at high speed [1].

The Endcap Muon Track Finder (EMTF) is part of the L1 trigger system that reconstructs muons in the forward region, using only hits from the endcap muon chambers. There are two types of detectors: Cathode Strip Chambers (CSC) and Resistive Plate Chambers (RPC), located in 4 separate stations as shown in Fig. 1. We assign the muon transverse momentum (p_T) based on the track curvature in the magnetic field, with low-p_T muons having a larger bend. A typical interesting event that we want to trigger has a real muon with p_T of 25 GeV or higher. However, for every 25 GeV muon, there are 1000 muons with p_T of 5 GeV. It is critical to distinguish the low-p_T muons from the high-pt muons in order to keep the trigger rates under control.

During the Phase 1 upgrade of the CMS L1 trigger, we introduced the first implementation of Boosted Decision Trees (BDTs) at L1 for the muon p_T assignment [2]. The BDTs combine multiple discriminating variables to perform a more accurate regression. We use the TMVA software package [3] to train the BDTs offline. Large look-up tables (LUTs) are used to store the response of the BDTs for every possible combination of the input variables. Each combination is encoded in 30 bits and is used as the address in the LUTs to retrieve the output p_T, which is stored as a 9-bit word. Using the LUT online allows us to perform the p_T assignment very quickly, as the complex BDT evaluation has been turned into a simple look-up operation.

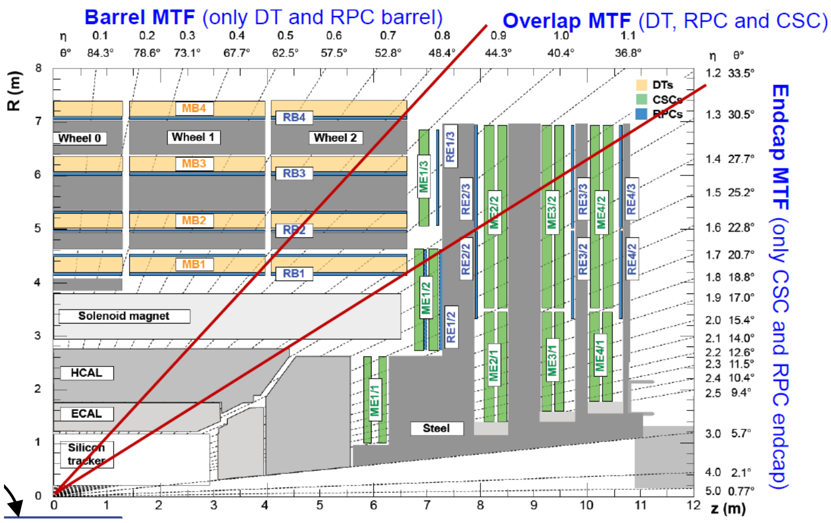

Figure 1: The $R-z$ view CMS muon trigger geometry [4]. The EMTF uses hits from the CSC detectors (green) and the RPC detectors (blue) in the 4 forward stations.

2. Momentum Assignment using BDTs

The EMTF reconstructs muon tracks using hits in the CSC and RPC chambers that are correlated in θ and ϕ between stations. The trajectory of a low-p_{T} muon is bent more by the magnetic field compared to that of a high-p_{T} muon. Typically, the bending in ϕ ($\Delta\phi$) is less than $12°$, while the bending in θ ($\Delta\theta$) is less than $1°$. These angles provide the main discrimination power to determine the muon p_{T}. However, this relationship is complicated by several factors. At low p_{T}, muons can experience significant multiple scattering and energy loss; and at high p_{T}, muons can go through electromagnetic showers. In addition, the CMS magnetic field strength varies with pseudorapidity (η), so the muons of similar momenta can have different behavior in the more central region ($|\eta| < 1.55$) and in the more forward region ($|\eta| > 2.1$).

The complicated dependencies make this an ideal case for machine learning. We choose the BDT regression technique to take advantage of multiple discriminating variables and their correlations to improve the accuracy of the p_{T} assignment. The BDTs are trained on Monte Carlo simulation of single-muon events. Different variables are used depending on the track "mode" that describes which stations have matching hits to the track. (For instance, the mode 1-3-4 denotes a track that has matching hits in the 1st, 3rd and 4th stations.) The variables are chosen to provide as much information as possible to the BDTs for each mode, such as $\Delta\phi$'s, $\Delta\theta$'s, detector types, track θ, etc. For the training, we use the inverse of p_{T} (p_{T}^{-1}) as the regression target and the least squared difference, $L = \left(x_{\mathrm{predict}} - x_{\mathrm{true}}\right)^2$, as the loss function. The event sample is re-weighted to fall as p_{T}^{-2}, such that overestimation of the p_{T} of the low-p_{T} muons is penalized more strongly. The BDT parameters are optimized based on both the trigger rate and efficiency as the figures of merit, as our goal is to minimize the trigger rate while maximizing the efficiency.

At L1, the trigger system must operate in fixed low latency. However, the BDT evaluation would require traversing every decision tree and evaluating the "decision" at every leaf node in the tree. This would require a lot of FPGA logic resources as well as using many clock cycles, thus it is infeasible to implement the BDTs directly into the FPGAs. Instead, we implemented large LUTs that store the response of the BDTs. The input variables are discretized and compressed into a 30-bit word, which is then used as the look-up address to retrieve the BDT p_{T} assignment, which is stored as a 9-bit word that represents p_{T} in unit of 0.5 GeV. Given the 30-bit address space, there are $2^{30} \approx 10^9$ possible addresses, which mean 1 billion combinations of the input variables. We loop through all of these combinations and evaluate them offline and store the results into custom designed 1.2 GB memory modules in the EMTF hardware that allow low random-access latency. The compression scheme of the input variables is explained in Table 1.

3. Performance

Using the data collected during 2017, the EMTF performance has been analyzed [5]. The unprescaled muon trigger is operating at the p_{T} threshold of 25 GeV with a tight quality selection. The tight quality selection requires each track in the endcap to contain hits in at least three stations. In the forward region ($1.2 \leq |\eta| \leq 2.5$), we have obtained a factor of 3 rate reduction compared to the legacy system that ran through 2015 (Fig. 2 left). The BDT-based momentum assignment has significantly reduced the number of low-p_{T} muons passing the trigger threshold which constitute

Four-Station Modes													
Mode	**Feature**	$\Delta\phi_{12}$	$\Delta\phi_{23}$	$\Delta\phi_{34}$	sign	$\Delta\theta_{14}$	B_1	B_2	B_3	B_4	FR_1	θ	Mode
1-2-3-4	**Bits**	7	5	4	2	2	2	1	1	1	1	3	1

Three-Station Modes												
Mode	**Feature**	$\Delta\phi_{12}$	$\Delta\phi_{23}$	sign	$\Delta\theta_{13}$	B_1	B_2	B_3	FR_1	FR_2	θ	Mode
1-2-3	**Bits**	7	5	1	3	2	1	1	1	1	5	3
Mode	**Feature**	$\Delta\phi_{12}$	$\Delta\phi_{24}$	sign	$\Delta\theta_{14}$	B_1	B_2	B_4	FR_1	FR_2	θ	Mode
1-2-4	**Bits**	7	5	1	3	2	1	1	1	1	5	3
Mode	**Feature**	$\Delta\phi_{13}$	$\Delta\phi_{34}$	sign	$\Delta\theta_{14}$	B_1	B_3	B_4	FR_1	FR_3	θ	Mode
1-3-4	**Bits**	7	5	1	3	2	1	1	1	1	5	3
Mode	**Feature**	$\Delta\phi_{23}$	$\Delta\phi_{34}$	sign	$\Delta\theta_{24}$	B_2	B_3	B_4	FR_2	–	θ	Mode
2-3-4	**Bits**	7	5	1	3	2	1	1	1	–	5	4

Two-Station Modes									
Mode	**Feature**	$\Delta\phi_{XY}$	$\Delta\theta_{XY}$	B_X	B_Y	FR_X	FR_Y	θ	Mode
X-Y	**Bits**	7	3	3	3	1	1	5	7

Table 1: The compression scheme of the variables used for the BDT-based momentum assignment for each mode. The variables are compressed into a 30-bit word. $\Delta\phi_{ij}$ ($\Delta\theta_{ij}$) denotes the bending angle in ϕ (θ) between stations i & j; "sign" encodes the signs of the later $\Delta\phi_{ij}$'s relative to the first $\Delta\phi_{ij}$ for the mode. the "B" variables indicate whether the hits come from CSC or RPC; the "FR" variables indicate whether the CSC hit is in the front- or rear-positioned chamber; and θ is the estimate of the track polar angle. There are 6 two-station modes: 1-2, 1-3, 1-4, 2-3, 2-4 and 3-4 that use the same scheme.

the major source of background. At the same time, we have also achieved about 95% efficiency at $p_{\mathrm{T}} = 25$ GeV, essentially unchanged with respect to the legacy system (Fig. 2 right). The trigger efficiency is measured in a sample of dimuon events using the tag-and-probe method, where the tag muon fires the trigger and the efficiency is measured with the probe muon. In 2016, the BDTs achieved a factor of 2 rate reduction with CSC-only tracks, at a cost of 5% reduced efficiency. But with the inclusion of RPC hits in the EMTF in 2017 (as per the original Phase 1 upgrade design), we regained the 5% in efficiency, and further reduced the rate. Therefore, the rate reduction effectively comes at no cost to the efficiency.

4. Conclusions

To improve the momentum assignment of the CMS L1 endcap muon trigger, we introduced the first implementation of BDTs inside a L1 trigger system at the LHC. The BDTs are capable of taking advantage of 25 different variables and their complex correlations, the inhomogeneous magnetic field, and non-linear effects such as inelastic scattering. The output of the BDTs is stored in 1.2 GB large LUTs, which turn the complex BDT evaluation into a simple look-up operation in fixed low latency. Using 2017 data, we have shown that the new momentum assignment algorithm reduced the rate by a factor of 3 at the 25 GeV trigger threshold with respect to the legacy system at effectively no cost to the efficiency.

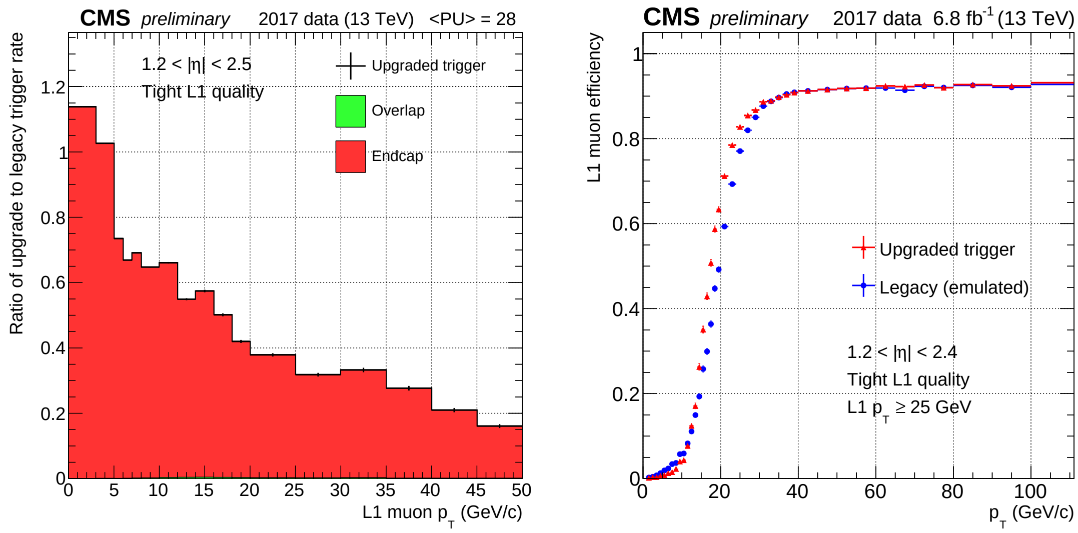

Figure 2: **(Left)** The ratio of the upgraded EMTF trigger rate over the the legacy trigger rate for the muon trigger with tight L1 quality as a function of p_T in the forward region [5]. A factor of 3 rate reduction at the p_T threshold of 25 GeV is achieved. **(Right)** The efficiency of the muon trigger as a function of p_T in the forward region. The upgraded EMTF retains the same efficiency as the legacy system.

References

[1] CMS collaboration, *The CMS trigger system*, *JINST* **12** (2017) P01020, [1609.02366]. https://cds.cern.ch/record/2212926.

[2] CMS collaboration, *CMS Technical Design Report for the Level-1 Trigger Upgrade*, CERN-LHCC-2013-011, CMS-TDR-012 (2013). https://cds.cern.ch/record/1556311.

[3] A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne and H. Voss, *TMVA: Toolkit for Multivariate Data Analysis*, *PoS* **ACAT** (2007) 040, [physics/0703039].

[4] CMS collaboration, S. Chatrchyan et al., *The performance of the CMS muon detector in proton-proton collisions at sqrt(s) = 7 TeV at the LHC*, *JINST* **8** (2013) P11002, [1306.6905]. https://cds.cern.ch/record/1558674.

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