

Synchronization of the CMS Cathode Strip Chambers

G. Rakness^a, J. Hauser^a, D. Wang^b

a) University of California, Los Angeles

b) University of Florida

Gregory.Rakness@cern.ch

Abstract

The synchronization of the trigger and data acquisition systems for the Cathode Strip Chambers (CSCs) in the Compact Muon Solenoid (CMS) detector at CERN is described. The CSC trigger system is designed to trigger CMS on muons with high efficiency (~99% per chamber) and is able to accurately identify its 25ns proton bunch crossing. To date, asynchronous cosmic ray data have been used to define the protocol and to refine timing algorithms, allowing synchronization to be realized within and between chambers to within ± 10 ns. Final synchronization of the CSCs requires timing parameters to be accurate to 2 ns. This goal will be readily achieved from the cosmic ray baseline using data taken with the synchronous beam structure of the Large Hadron Collider.

I. INTRODUCTION

The Compact Muon Solenoid (CMS) is a multi-purpose detector located at the 5th intersection point (IP) in the Large Hadron Collider (LHC) at CERN. Symmetrically designed around the IP, CMS has been instrumented to detect particles emanating from proton collisions with center-of-mass energy $\sqrt{s}=14\text{TeV}$. The collision luminosities at LHC are expected to reach $10^{34}/\text{cm}^2/\text{sec}$ with a bunch crossing frequency of 40MHz. As a discovery detector in a high-energy, high-rate environment, precision is paramount for CMS. It has been instrumented with detector systems covering a large angular acceptance with excellent momentum resolution over the full range of interesting physics.

The Cathode Strip Chambers (CSCs) are part of the muon detection system in CMS [1]. The CSC system comprises 468 multi-wire proportional chambers mounted in eight stations, four on each endcap of the CMS magnet. The stations are interleaved with the magnet return yokes to detect muons produced at polar angles of $0.9 < |\eta| < 2.4$, and cover the full azimuth (ϕ). Each chamber contains 6 layers of orthogonally oriented cathode strips and anode wires. The cathode strips measure the muon's azimuthal position, while the anode wires measure the polar angle. The orientation of the magnetic field in the endcap region bends muons in the r - ϕ plane.

Since many interesting physical processes include muons in the final state, the CSCs are part of the CMS Level-1 (L1) trigger system. The signal from the L1 trigger system which prompts the readout of subdetectors is called the L1-Accept (L1A) signal. In order to enable the large bandwidth readout of the CMS silicon tracker, the L1A is required to arrive at the detectors within $3.2\mu\text{sec}$ of the collision, placing a premium on the reduction of latency of the L1 trigger.

The front end electronics of the CSC L1 trigger and data acquisition (DAQ) systems consist of custom-built boards located both on the chambers and in custom VME crates placed around the periphery of the magnet yokes (peripheral crates). Processing algorithms are contained on FPGAs in a pipelined architecture to process events at the LHC bunch crossing rate. Since the peripheral crates are in a high radiation environment, single-event upsets are expected to occur on the FPGAs during LHC running, after which the programs need to be reloaded from PROMs. The FPGAs are highly configurable, and the parameters are programmed onto "user" PROMs separate from those containing the FPGA program. After the FPGA is programmed, the parameters are loaded from the user PROMs. Inaccessibility to the electronics during collisions requires that remote communication protocol be used to set timing parameters and to monitor trigger performance.

Here we describe the timing alignment of the CSC front end electronics, including the DAQ and the trigger electronics up to the Track Finder crate. The timing of the trigger path further along the CMS L1 trigger is not discussed.

II. SYNCHRONIZATION WITHOUT BEAM

The full complement of CSCs have been constructed and mounted on the magnet yokes in their final positions. Half of the detector has been moved from the assembly building above ground (SX5) to the experimental hall ~100 m below ground (USX55). Until proton collisions commence in the LHC, the CSC system will be in some state of commissioning. To perform this commissioning, some of the chambers have participated in test beams, and cosmic rays have used to test the response of all chambers. Although the spectrum and angular distribution of cosmic ray muons bear little resemblance to those coming from protons colliding at $\sqrt{s}=14\text{TeV}$, they are a reliable source which has been used to measure the response of the chambers to relativistic muons.

Since 2005, a small fraction of the CSC system has been in operation in SX5 detecting cosmic rays to test and commission the trigger and DAQ systems. Since March 2007, a set of 18 chambers in a localized angular bite spread over two stations on a single endcap was studied, called the minus side slice test. Among other commissioning activities, the minus side slice test has been used to define the synchronization procedure for the CSCs. The goal is to determine as many of the delay parameters as possible before LHC collisions in order to minimize the beam time needed and to maximize the potential for quick discovery.

The nominal width of coincidence windows is 3 bunch crossings ($1\text{bx} = 25\text{ns} = \text{LHC beam crossing period}$) to accommodate the high rate of muons from proton collisions at

nominal LHC luminosity. Since the rate of cosmic ray muons is small compared to the coincidence window width, it is straightforward to perform scans over delay parameters to look for timing coincidences in the response of pertinent components in order to synchronize them. For muons which come from the CMS IP, the difference in timing from the closest chamber to the farthest as a result of time-of-flight effects is less than 25 ns. Therefore, effects from time-of-flight are a perturbation on the effects from cable and optical fiber lengths. For cosmic ray muons, the effects from time-of-flight are considerably different, but of the same magnitude for chambers restricted to the minus side slice test.

Since the CSC system comprises 468 chambers with multiple components connected by myriad cables and optical fibers of varying lengths, the synchronization requires a systematic approach. Delay parameters are built in so that there is one delay parameter per feature. An automated procedure to measure the parameters step-by-step is needed to properly include correlations between the parameters. To converge quickly on the final parameters, a model is needed to predict starting points which are close to the final values. Finally, the procedure needs to work with various “sources,” including cosmic rays above ground, cosmic rays in the experimental cavern, test pulses generated on the chamber, and LHC collisions.

III. TRIGGER PATH SYNCHRONIZATION

The CSC data readout is described in detail in Reference [2]. The data is stored in pipeline stages in the front-end until L1A arrives. In order zero-suppress the large amount of data potentially coming from the CSC, a time coincidence is required between these pipeline stages and the L1A time of arrival at the front-end. Therefore, it is important to time in the trigger path first so that the L1A timing is stable.

To reduce backgrounds and accommodate the large rates expected at LHC, the CSC trigger electronics begin analyzing the signals on the front end. Trigger primitives for the anode and cathode signals are determined by comparing hits with patterns that are expected for muon candidates. Since the time development of cathode signals is on the order of 75 ns, and to reduce false triggers from correlated noise, the hits are compared to patterns twice, both before and after a “drift-delay.” The Anode Local Charged Track (ALCT) is formed by a successful match to an envelope which points at the interaction region. The Cathode LCT (CLCT) is formed by a successful match to a set of patterns whose amount of bend enables a coarse measurement of the muon momentum early in the trigger path. The CLCT and ALCT are tagged by their strip and wire group to specify the local position.

The trigger electronics for the CSCs are highly configurable to maximize the efficiency and minimize the latency. These parameters include the thresholds of the anode discriminators and cathode comparators, the number of layers required for a pattern match, comparator threshold, mode, and timing, and the “drift delay.” All of these will affect the timing of the system at some level, and need to be set before the synchronization parameters are determined.

The path of the trigger primitives from the Front-End Boards (FEBs) to the Sector Processor (SP) can be seen in

Figure 1. The synchronization of the CSC trigger requires timing alignment in three sequential steps: alignment of the signals from the anode front end boards, optimization of the coincidence of the anode with the cathode signals to make a Local Charged Track (LCT), and alignment of LCTs from different chambers at the SP.

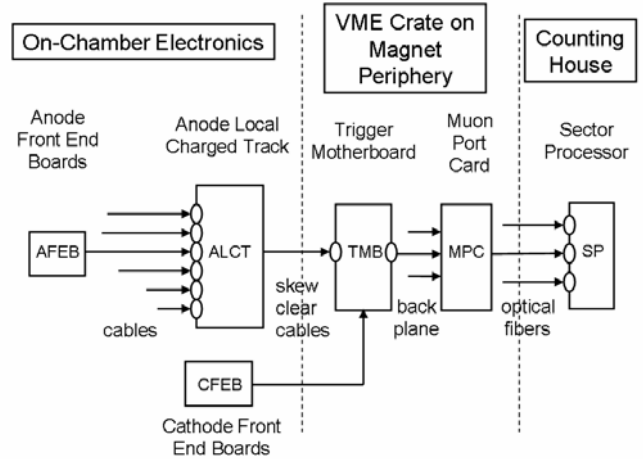


Figure 1: Schematic diagram of the path of CSC trigger signals from the front end boards on the chamber to the Track Finder crate in the counting house. Ovals indicate the locations of the delay parameters needed to synchronize the CSC trigger primitives, and are called (from left-to-right) the AFEB fine delays, the ALCT-CLCT match delay, the MPC output delay, and the SP Alignment FIFO.

A. Anode Front-End Alignment

Signals from the anode wires are grouped together and processed on Anode Front-End Boards (AFEBs) located on the chamber. The AFEB contains an amplifier-discriminator Application-Specific Integrated Circuit (ASIC), which records a “hit” for signals which surpass the threshold in each layer. The time slewing of this ASIC is on the order of 3 ns for the input charge from minimum-ionizing particles. As such, the anode signals are used to define the timing of the LCT in the trigger path and to specify the bunch crossing of the event in the data stream.

Depending on the size of the chamber, there are 12-42 AFEBs for any single chamber. The AFEB signals are sent on cables to the Anode LCT (ALCT)-board located on the chamber. At the input of the ALCT-board, there is a delay ASIC for each AFEB with a step size of ~ 2.2 ns (16 step range). In the CSC system, the “AFEB fine delay” is the single delay parameter with granularity finer than 1 bx. Therefore, the AFEB fine delays are used to correct any misalignments of LCTs at the resolution of ~ 0.1 bx. All other delays in the CSC trigger system are implemented in FPGAs, and have a granularity of a bunch crossing.

The AFEB fine delay is used to align the asynchronous pulses from the AFEB at the ALCT, whose FPGA operates on the LHC clock frequency. On a chamber, timing differences between AFEBs can result from signal propagation through different cable lengths between the AFEBs and ALCT (up to 18 ns difference within a chamber) as well as small chip-to-chip variations in the response of the delay ASIC’s. These effects were measured during the construction and

commissioning of the chambers before mounting them on the disks. Time-of-flight effects from one end of a chamber to the other for muons originating from the CMS IP are estimated to be less than 6 ns.

B. Anode Alignment with Cathode

For the trigger path, the timing of signals from the cathode strips is important insofar as they need to be synchronized with the anode signals. The cathode strip signals are processed in Cathode Front-End Boards (CFEBs) located on the chamber. A CFEB contains an amplifier-shaper ASIC with a shaping time on the order of 100 ns. For the trigger, the signals pass through comparator ASICs that rapidly determine the location of the muon on each layer to a precision of $\frac{1}{2}$ -strip. The strip positions are multiplexed and sent on skew-clear cables to the Trigger Motherboard (TMB) where the FPGA containing the CLCT algorithm is located.

The TMB is a custom 9U VME board which sits in the peripheral crate. There is one TMB for each chamber, on which is mounted a mezzanine board to hold the FPGA which contains the CLCT algorithm as well as logic to perform the time coincidence of the ALCT with the CLCT. The ALCT data is sent on skew-clear cables to the TMB, and is delayed by the “ALCT-CLCT match delay” in the TMB to wait for its CLCT. A coincidence window with configurable width (nominally 3bx for LHC running) is opened when the CLCT is formed. A coincidence between an ALCT and CLCT defines a Local Charged Track (LCT) trigger candidate signal from this chamber. The LCT contains position and timing information of sufficient resolution for the Sector Processor to create tracks using LCTs from different stations to create a Level-1 muon trigger candidate.

The position of the ALCT in the CLCT coincidence window is stored in the TMB header and written in the TMB data. A copy of these data is stored in a RAM which is accessible by VME reads. From these data, the delay needed to center the ALCT in the match window and minimize the latency can be easily extracted.

The ALCT-CLCT match delay depends mainly on the difference in the total signal propagation between the anode signal and the cathode signal to the TMB. The anode latency includes the (equalized) propagation time from the AFEB to the ALCT, the ALCT algorithm, and the TMB-ALCT skew-clear cable length. The cathode includes the amplifier-shaper ASIC, the comparator ASIC chain, the TMB-CFEB skew-clear cable length, and the CLCT algorithm. There is also an effect of the relative clock between the ALCT and the CFEB.

C. Chamber-to-Chamber Alignment

Local Charged Tracks from synchronous muons must be aligned at the Sector Processor. The LCTs are passed from the TMB along a custom backplane to a Muon Port Card (MPC). The MPC selects the highest quality trigger candidates from chambers in a station within a sector. The LCTs are sent together with its clock from the MPC over optical fibers to a Track Finder module to create a muon track trigger candidate. The timing of LCTs at the SP is factorized into two components, the arrival of the LCTs at the MPCs,

and the propagation of the signals from the MPCs to the SPs over different optical fiber lengths.

The variation of optical fiber lengths is corrected for in the firmware of the SP FPGA. At the beginning of every run, a “Resync” signal is sent from the CMS Timing, Trigger, and Control (TTC) system to all triggering subsystems [3]. TTC commands are sent over optical fibers to Clock and Control Boards (CCBs) which distribute the commands to the electronic components in the CSC [4]. Upon receipt of the Resync command, each MPC sends a stream of signals to the SP. The FPGA in the SP contains an “Alignment FIFO,” which automatically lines up these streams of signals from the 5 MPCs in each sector, regardless of the phase of the data. The Alignment FIFO corrects for differences in the combined fiber lengths of TTC-CCB and MPC-SP. There is a delay on the TTCrq mezzanine board on the CCB to correct for the TTC-CCB fiber lengths to synchronize TTC commands at the CCBs. This delay is set according to the measurements of the fibers. Mistakes in this delay will result in mistakes in the LCT arrival at the SP, and will be accommodated by the procedure to align the LCTs at the MPCs.

The alignment of LCTs at MPCs is performed by offline analysis of the Sector Processor data. LCTs at the SP are tagged by the bunch crossing in which they arrive. For every pair of chambers that contain LCTs, the difference in bunch crossing number is histogrammed. The distributions for chamber pairs having a high probability of overlap for cosmic ray muons are analyzed to extract the mean and error on the mean. From the matrix of differences, a global chi-square is constructed, from which the timing shift for each chamber can be obtained simultaneously. Consistent results were obtained using matrix inversion and numerical minimization.

A digital delay on the output of the TMB (called the “MPC output delay”) gives sufficient dynamic range to align all LCTs from all chambers at the MPCs. The coordination of the MPC output delay with the ALCT-CLCT match delay allows LCTs to be aligned at the MPC independently from aligning the ALCT with the CLCT.

The results of applying this method can be seen in Figure 2, plotted versus the chambers in the minus side slice test. The timing alignment between chambers is better than 0.1bx, which is the resolution of the AFEB fine delay step size.

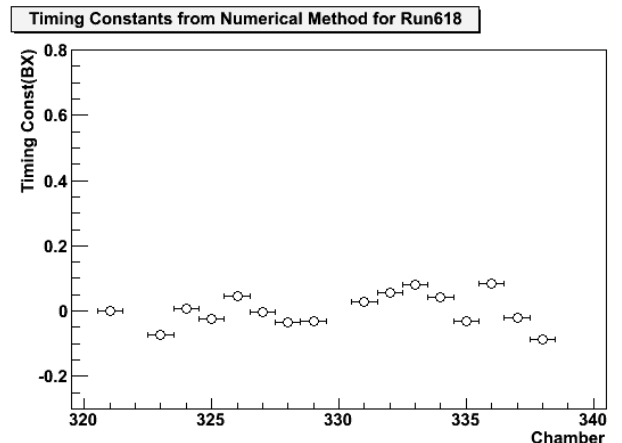


Figure 2: Relative arrival of trigger signals from different chambers at the Track Finder (bx) versus chamber number for a subset of the CSC system. Data are from cosmic ray muons.

The time-of-flight difference between the first and last CSC station is ~ 20 ns for a muon originating at the IP. Within a chamber, time-of-flight effects for IP muons are less than 6 ns. This agrees with the observation that including cable delays as the starting values predicted the relative timing at the SP to better than 1 bx.

IV. DATA ACQUISITION PATH

The CSC data acquisition (DAQ) system is described in detail elsewhere [2]. The number of readout channels for the CSCs is $\approx 400k$, which is greater than the number of channels in the CDF or D0 experiment (excluding the silicon trackers). In order not to burden the Global DAQ with readout from mostly empty chambers, the CSC readout is highly zero-suppressed by sending a chamber's data only when that chamber has created a trigger. The 3 components which send data are the CFEBs, the ALCT, and the TMB. The board which is the conduit into the data stream is the Data Acquisition Motherboard (DAQMB). The DAQMB is a custom-built 9U VME board which is located in the peripheral crate, in the slot adjacent to the TMB. There is one DAQMB for each chamber.

A. Level-1 Accept at Readout Components

In order to readout, the CFEB, ALCT-board, and TMB each require the coincidence of a Level-1 Accept signal (L1A) and the "DAQ-start" signal for that component. A coincidence window is opened upon the receipt of L1A in which the (delayed) DAQ-start signal must arrive to begin the readout for that event. The procedure described in Section III describes the corrections which are made to the trigger path so that the L1A signal arrives synchronously at all peripheral crates. Therefore, it is the DAQ-start signal which must be delayed appropriately for each component.

For the TMB, whose data comprise the CLCT and ALCT-CLCT coincidence information, the DAQ-start is the LCT. Therefore, the amount of time necessary to delay the TMB DAQ-start signal is a constant plus the MPC output delay, since this is the sole delay along the trigger path after the LCT has been defined.

For the ALCT-board, the DAQ-start is the ALCT signal, itself. Therefore, the amount of time necessary to delay the ALCT DAQ-start signal is a constant plus the time for signal propagation through the ALCT-TMB skew-clear cable length, since this is the additional path length the L1A must travel to reach the ALCT.

The TMB contains a set of counters which is updated on every bunch crossing for diagnosis and to record general front-end trigger statistics. The counters contain addresses which count the number of times the ALCT and the number of times the TMB receive L1As coincident with the DAQ-start signals. Scanning over the delay parameters and reading the value of these counters, one easily extracts the delay value needed to center the coincidence of the DAQ-start signal in the L1A window.

The tightest timing requirement for the DAQ path is on the CFEBs. The precision data from the CFEB includes analog pulse height data for each strip stored in Switched

Capacitor Arrays (SCAs), whose readout are in fixed blocks of 400 ns. The DAQ-start signal for the CFEB data is the CLCT pretrigger signal, called the Active-FEB Flag (AFF) for the purposes of readout. Receipt of the AFF at the CFEB initiates storage of the SCA data in a buffer for potential future readout. The AFF is generated by the TMB and sent to the adjacent slot to the DAQMB, where the AFF-L1A coincidence logic is located. The amount of delay needed to align the AFF with the L1A is a constant plus the signal propagation through the CFEB-TMB skew-clear cable length, since these govern the timing of the CLCT creation.

The DAQMB has a counter which starts when the AFF arrives and stops with the L1A. This counter is accessible by VME, so that a distribution of time difference between the most recent AFF and the most recent L1A can be created online. This histogram can be easily analyzed to center the AFF in the L1A window.

B. Data Available Bits at DAQ Motherboard

When the L1A is coincident with the DAQ-start signal at the CFEB, TMB, or ALCT, the component initiates readout of the appropriate data stored in its buffer and sends a Data Available (DAV) signal to the DAQMB. When the DAQMB receives the L1A, the DAQMB waits a fixed time and opens up a single window in which to receive these DAV signals. If the DAV signal from a component arrives in the window, the DAQMB prepares to accept its data.

Since all DAV signals for a particular L1A must arrive in a single window, DAV signals which arrive early must be delayed. Each DAV delay includes the time necessary for the L1A to be accepted by the component and the subsequent propagation of the DAV bit. Since the time it takes to receive the L1A and to send its DAV bit is fixed for each component, the synchronization parameters available on the DAQMB to adjust for DAV only correct for the cable length differences between the chambers. Hence, there is no adjustable parameter for the TMB DAV, but there are parameters for the ALCT and the CFEB.

V. DIFFERENT WAYS OF TRIGGERING

Using cosmic ray muons in the SX5 assembly building, the minus side slice test has been used to establish the synchronization procedures described in Sections III and IV. In moving other parts of the system downstairs to the experimental hall, measurements have confirmed the expectation that the rate of cosmic ray muons 100m underground drops to $\sim 1\%$ [5] of the rate on the surface. In order to increase the rate to a reasonable level underground, it is useful to be able to enable the CSC trigger system to trigger in different configurations.

The essential configurable parameter for ALCT and CLCT operation is the number of layers which are required to satisfy a pattern to declare an ALCT or CLCT found. The standard configuration for collision muons is a minimum of 4 layers out of 6, with a coincidence between ALCT and CLCT. As described in Section III, the standard patterns used for the CSC trigger primitives have been tailored to trigger on collision muons from the LHC. Cosmic ray muons, however, are incident on the chambers according to a $\cos^2(\theta)$

distribution [5]. This means that the fraction of cosmic ray muons which satisfy these LHC collision patterns depends strongly on the orientation of the chamber with respect to vertical. For the minus side slice test, the rate of muons which passed the LHC trigger patterns was $\sim 20\text{Hz}/\text{chamber}$.

The “n-Layer” trigger configuration is defined by removing the pattern restriction for ALCT and CLCT. A hit anywhere within chamber enables a layer, and a coincidence of n-layers fires the trigger. A coincidence between ALCT and CLCT is still required. The angular acceptance of muons accepted by the CSC is increased to any muon which traverses a given number of layers in the chamber. Using $n=4$, the rate of cosmic rays which pass the n-Layer trigger increases compared to those which pass the LHC trigger patterns by about a factor of 6.

The “Radioactive” trigger configuration is defined by reducing the layer threshold to a single layer for both CLCT and ALCT, with a coincidence between the two. The idea is to trigger on the low-energy natural radioactivity within and around the chamber. Approximately 95% of CLCT single-layer triggers are coincident with an ALCT. The rate of “Radioactive” triggers is approximately 100 times the cosmic ray LHC-trigger rate.

Finally, the CSCs have been equipped with a test strip which can be pulsed at an arbitrarily high rate. The test strip is an extra cathode strip, located at the edge of the chamber acceptance, and is connected to a pulsing circuit on the ALCT via a LEMO cable. When a pulse is generated on the test strip, charge is induced both on the strips and on the wires, creating both a CLCT and an ALCT. It should be noted that the test pulses are not synchronized between chambers, but are intended to be used to measure effects due to cable lengths within a single chamber.

The synchronization parameters were determined using each of these triggering methods, and were found to be consistent within ~ 1 bx. This agreement is on the order of time-of-flight effects.

VI. SYNCHRONIZATION MODEL

To expedite the measurements of synchronization parameters for 468 chambers, it is useful to begin with starting values near the final solution. Hence, it is important to use a robust model to predict the values for the 0th order input. Using the data obtained from the minus side slice test, a model was developed to predict the relative values of the delay parameters based on the propagation time of signals through cables and fibers of varying length. The parameters measured at the minus side slice test were found to be in reasonable agreement with the model in all aspects of synchronization, including ALCT-CLCT matching, receipt of L1A, and acceptance of the DAV signal.

The cable-length model can be expanded to include effects due to time-of-flight. Time-of-flight corrections for muons emanating from a point source at the IP can be computed

easily and applied to the 0th order AFEB fine delay values obtained from chamber construction and commissioning. For cosmic ray muons, ad-hoc corrections for time-of-flight effects were added to the model to compare with the data from the minus side slice test. The difference in the predicted and measured values for the AFEB fine delays can be seen in Figure 3. The data are in excellent agreement with the model.

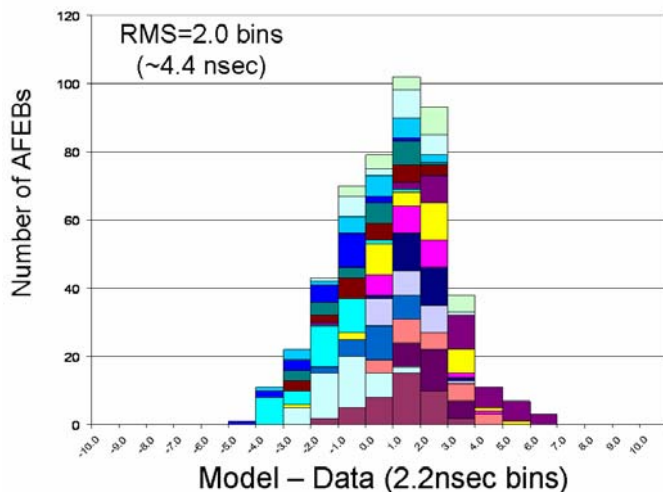


Figure 3: Distribution of the difference between the model and the data for the AFEB fine delays for the 480 AFEBs in the minus side slice test for cosmic ray muons.

Synchronization parameters have been predicted using the cable-length model for 468 CSC chambers, and will be used as the 0th order values from which to begin measurements.

VII. SUMMARY

A procedure to achieve time synchronization of the CSC trigger and DAQ system has been developed using cosmic ray data. All parameters can be measured using trigger configurations with sufficient rate to assure success in the experimental cavern. A model has been developed that is used to predict initial synchronization parameter values. As such, the final synchronization of the CSCs should be readily achieved with a small amount of data taken with the synchronous beam structure of the LHC.

VIII. REFERENCES

1. CMS Muon Technical Design Report, CERN/LHCC 97-32, 15 December 1997.
2. J. Gu, these proceedings.
3. "TTC Distribution for LHC Detectors", by B.G. Taylor, for the RD12 Project Collaboration, IEEE Trans. Nuclear Science, Vol. 45, No. 3, June 1998, pp. 821-828.
4. F. Geurts, M. Matveev, and P. Padley, these proceedings.
5. L.N. Bogdanova *et al.*, Phys. Atom. Nucl. **69** (2006) 1293; arXiv:nucl-ex/0601019v1.