



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Nuclear Instruments and Methods in Physics Research A 534 (2004) 441–485

**NUCLEAR
INSTRUMENTS
& METHODS
IN PHYSICS
RESEARCH**
Section A

www.elsevier.com/locate/nima

Bunched beam test of the CMS drift tubes local muon trigger

P. Arce^a, M. Bellato^b, M. Benettoni^b, A. Benvenuti^c, D. Bonacorsi^c,
M. Bontenackels^d, J. Caballero^a, V. Cafaro^c, P. Capiluppi^c, L. Castellani^b,
F.R. Cavallo^c, M. Cerrada^a, P. Checchia^b, N. Colino^a, E. Conti^b, M. Corvo^b,
B. de la Cruz^a, F. Dal Corso^b, G.M. Dallavalle^c, M. De Giorgi^b,
F. Fabbri^c, A. Fanfani^c, C. Fernandez^a, J. Fernández de Trocóniz^c,
M.C. Fouz^a, P. Garcia Abia^a, F. Gasparini^b, U. Gasparini^b, P. Giacomelli^c,
V. Giordano^c, F. Gonella^b, C. Grandi^c, L. Guiducci^c, M. Gulmini^f,
T. Hebbeker^d, J.M. Hernandez^a, K. Hoepfner^d, I. Josa^a, S. Lacaprara^b,
I. Lippi^b, R. Mameghani^d, S. Marcellini^c, G. Maron^f, R. Martinelli^{b,1},
S. Maselli^g, G. Masetti^c, A.T. Meneguzzo^b, G. Meng^b, V. Monaco^g,
A. Montanari^c, F. Montecassiano^b, F. Navarria^c, F. Odorici^c,
M. Passaseo^b, M. Pegoraro^b, C. Peroni^g, A. Perrotta^c, A.J. Ponte Sancho^{b,2},
J. Puerta^a, H. Reithler^d, A. Romero^g, L. Romero^a, P. Ronchese^b,
A. Rossi^c, T. Rovelli^c, R. Sacchi^g, A. Staiano^g, N. Toniolo^f,
E. Torassa^b, G. Torromeo^c, R. Travaglini^c, S. Vanini^b, L. Ventura^b,
S. Ventura^b, C. Villanueva^a, C. Willmott^a, M. Zanetti^b, L. Zangrando^f,
P. Zotto^{b,*}, G. Zumerle^b

^aCIEMAT, División de Física de Partículas, Avenida Complutense 22, E-28040 Madrid, Spain

^bDipartimento di Fisica dell'Università e Sezione INFN di Padova, Via Marzolo 8, I-35131 Padova, Italy

^cDipartimento di Fisica dell'Università e Sezione INFN di Bologna, Viale Berti Pichat 6/2, I-40127 Bologna, Italy

^dIII. Physikalisches Institut der RWTH Aachen, D-52056 Aachen, Germany

^eUniversidad Autónoma de Madrid, Ctra. de Colmenar km 15, E-28049 Madrid, Spain

^fINFN, Laboratori Nazionali di Legnaro, Viale dell'Università 2, I-35020 Legnaro (PD), Italy

^gDipartimento di Fisica dell'Università e Sezione INFN di Torino, Via Giuria 1, I-10125 Torino, Italy

Received 9 June 2004; accepted 17 June 2004

Available online 12 August 2004

*Corresponding author. Tel.: +39-49-827-7052; fax: +39-49-875-6233.

E-mail address: zotto@pd.infn.it (P. Zotto).

¹Until December 31st, 2000.

²Until August 31st, 1999.

Abstract

The 40 MHz bunched muon beam set up at CERN was used in May 2003 to make a full test of the drift tubes local muon trigger. The main goal of the test was to prove that the integration of the various devices located on a muon chamber was adequately done both on the hardware and software side of the system. Furthermore the test provided complete information about the general performance of the trigger algorithms in terms of efficiency and noise. Data were collected with the default configuration of the trigger devices and with several alternative configurations at various angles of incidence of the beam. Tests on noise suppression and di-muon trigger capability were performed.

© 2004 Elsevier B.V. All rights reserved.

Keywords: CMS; Muon trigger

1. Introduction

The purpose of the muon trigger in the CMS experiment at LHC [1] is to provide the identification of muons and a good measurement of their curvature in the magnetic field to enable a sharp cut on their momenta for rate reduction. These tasks are separated in the drift tubes trigger: a local algorithm provides muon track segments (trigger primitives) inside each chamber and a regional algorithm links these segments in the different chambers crossed by the muons extracting their vector parameters.

The local trigger must perform the identification of a muon's parent pp collision. Thus each track segment is uniquely assigned to a LHC bunch crossing as soon as it is found. This process is dead time free in order to avoid event losses. In addition the trigger dead areas were required to be negligible, leading to the concept of a redundant design. The CMS barrel muon detector layout [2] was designed around these stringent trigger requirements.

The bunch crossing assignment mechanism relies on a good determination of the drift times and the precise positioning of the anode wires inside each drift cell. Since the momentum cut at the trigger primitive level is more effective if higher resolutions are available, the resolution of the trigger algorithm was pushed to a value comparable to the muon chambers resolution. The inclusion of corrections used for the conversion of trigger data from the local reference system to the CMS one (computed by the hardware alignment system and stored in look-up tables) is therefore required.

The existence of punchthrough muons and the need to find di-muons requires the ability to identify more than one track within a single chamber. A zone of a few centimeters width being insensitive to di-muons is considered to be still acceptable inside each chamber. In fact two muons usually have different momenta and directions: hence even if they are very close inside one chamber they separate while moving through the magnetic field of the detector.

The implemented algorithm is flexible and designed to reduce background and noise. Therefore coincidences, thresholds and selections based on quality codes are available and configurable to easily modify the algorithm flow according to different experimental conditions.

The test and validation of the correct performance of the barrel muon chamber readout and trigger electronics need a bunched beam with a time structure similar to the LHC where proton collisions occur every 25 ns. The first complete test of the electronics was performed in May 2003 at the CERN SPS when an LHC-like beam was set up.

2. System overview

The CMS barrel muon drift chambers were already described extensively elsewhere ([3] and references therein). They are composed of three groups of drift tubes called Superlayers (SL). Each SL is built of four layers of drift tubes, staggered by half a tube width. Two SLs measure muons in the transverse bending (r, φ) plane of CMS

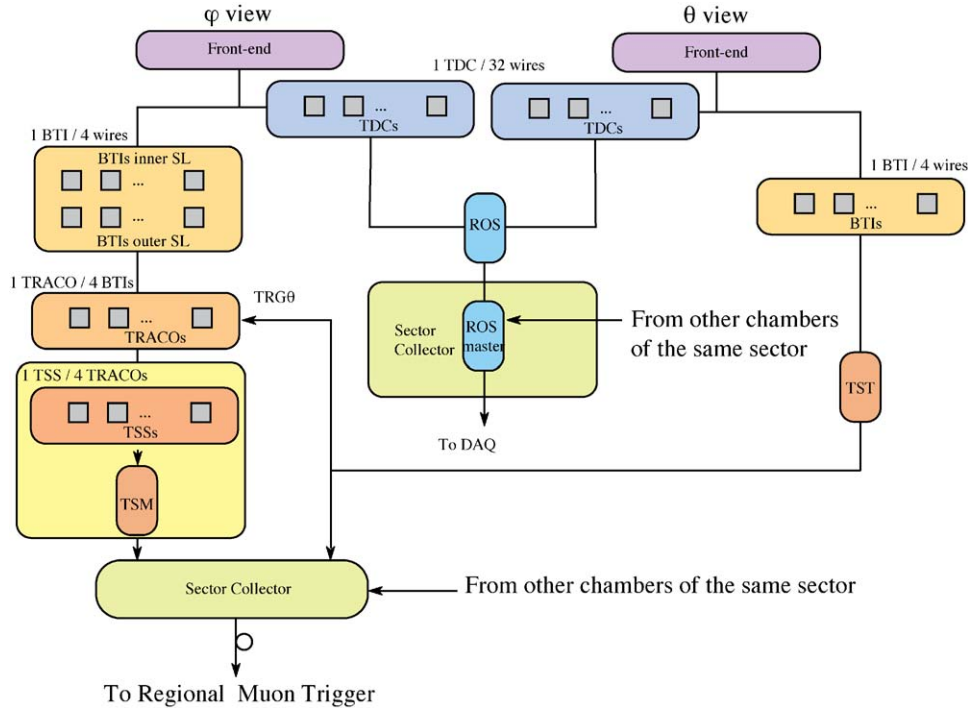


Fig. 1. Overview of the electronics layout of a barrel muon station.

(hereafter ϕ view) and one SL detects them in the longitudinal non-bending (r, θ) plane of CMS (hereafter θ view). A honeycomb panel, providing stiffness to the chamber, and a θ SL separate the two ϕ SLs and provide a lever arm for a better track fitting. Inside CMS muons will cross four chambers at different depths. Each chamber traversed by muons is called a measurement station (labeled MB1, MB2, MB3 and MB4, respectively, starting from the centre of CMS).

They are instrumented with readout and trigger electronics both in ϕ view and in θ view. The block scheme of the drift tubes electronics mounted on each muon station is shown in Fig. 1.

Signals from each drift tube are amplified by a charge preamplifier followed by a shaper with baseline restorer, then fed into a fast discriminator whose output is stretched and translated into LVDS compatible signal levels suitable to drive twisted pair cables [4]. Four of such chains are integrated in a full custom ASIC designed in $0.8\mu\text{m}$ BiCMOS technology by AMS. The threshold, the width of the output pulse and other

features are controlled by an external monitoring system.

Drift times are digitized by 32-channel High Performance General Purpose TDCs [5], hosted on Readout Boards (ROBs) [6]. Each board is equipped with four TDCs for a total number of 128 readout channels. The data of the ROBs are sent to a dedicated Readout Server Board (ROS) which forwards the data to the CMS Data Acquisition System.

The front-end trigger device which is directly interfaced to the wire front-end electronics, is called Bunch and Track Identifier (BTI): it performs a rough track reconstruction within each SL and uniquely assigns the parent bunch crossing to a track candidate. The BTI computes track candidates at each LHC bunch crossing and is therefore intrinsically dead-time free. Only one track candidate per bunch crossing is allowed inside a BTI.

The output of a BTI is sent to a Track Correlator (TRACO) that associates track segments in the same chamber by combining the

information delivered by groups of BTIs of the φ view. The TRACO enhances the angular resolution and produces a quality hierarchy of the triggers.

The trigger data of at most two tracks per bunch crossing reconstructed by each TRACO are transmitted to the chamber Trigger Server (TS) whose purpose is to perform a track selection in a multitrack environment. The TS of the φ view is composed by two devices (Track Sorter Slave (TSS), Track Sorter Master (TSM)) and selects two tracks (which from the physics point of view should correspond to the two muons with highest transverse momentum) among all tracks transmitted by the TRACOs. The TS of the θ view (Trigger Server Theta (TST)) sends the hardwired OR of the BTI outputs to the TRACOs for trigger qualification purposes and codes the triggers in a 16 bits string giving the position and the quality of all tracks pointing to the vertex.

Trigger and readout data in each CMS sector (each wheel of the CMS detector is divided into 12 wedge shaped sectors, corresponding to $\Delta\varphi = 30^\circ$) are sent to Sector Collector (SC) units, where the trigger information (track position, bending angle, quality bits) is coded and transmitted to the Regional Muon Trigger (running an algorithm called Drift Tubes

Track Finder: DTF) [7] using optical links. The DTF links the track primitives provided by each muon station in all CMS sectors and assigns to each detected muon its φ coordinate (using the position measured by the TRACO), η coordinate (using the TST information) and transverse momentum (from the displacement of the muon track segments in different muon stations).

All devices are implemented in custom integrated circuits. The flexibility needed to adapt the system to different environmental conditions is provided by a set of configuration parameters for each device that can be modified in situ by software to change the selection criteria of the implemented algorithms. The most important ones will be discussed in detail while the full technical details of the CMS first level trigger actual hardware implementation are available in Ref. [7].

2.1. Drift tube local trigger layout

Chamber trigger and readout electronics are lodged in a Mini-Crate (MC) mounted on the front side of each chamber inside the C profile surrounding the honeycomb layer. Fig. 2 shows a MC with the front panel removed. Inside a MC a

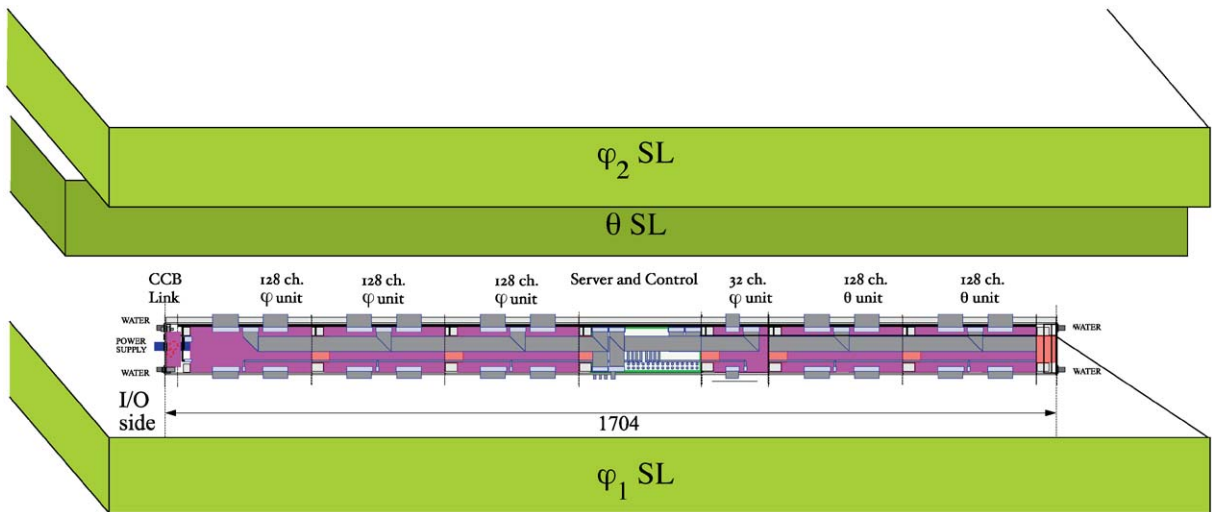


Fig. 2. Schematic layout of the Mini-Crate of an MB1 station. Its location on the chamber and its composition for the part of the trigger electronics are shown.

layer of trigger boards (indicated in Fig. 2) is mounted on top of a layer of ROBs while a central Server and Control Unit is available for communication with the external units (Detector Control System, CMS Data Acquisition and CMS Trigger System).

The trigger electronics (like the readout electronics) is grouped in 128-channel Trigger Boards (TRB) containing 32 BTIs assembled in eight multichip modules (BTIM) connected to four TRACOs and one TSS. An additional 32-channel unit is needed in the first and fourth CMS stations to match the channel number. The θ view is instrumented with 128-channel units containing 32 BTIs assembled on eight BTIMs. The TSM, consisting of three chips, is lodged on a central Server Board (SB). A Drift Tubes Chamber Control Board (DTCCB) takes care of electronics setup, monitoring, control and test functions. Parts of the control electronics are hosted on the SB. A Microcontroller (Motorola MC68HC16) processes and distributes commands issued by the Detector Control System to the trigger system and reports system faults to it. This board is connected via a dedicated optical fibre to the Detector Control Crate in the counting room.

Front-end signals are received from the TRBs through the ROBs where LVDS to CMOS translators are located. Trigger data are sent by each TRB to the SB via fine-pitch flat cables. Lateral fine-pitch connectors at each TRB, distributing detector signals and chamber control signals coming from the central Server and Control Boards, are used for communication between neighbouring units.

All units are powered by separated 3.3 V lines in order to maintain the highest possible redundancy. They are protected against overcurrents on the power supply lines by precise current monitors with fast shut off capability.

One TTCrx [8] receiver chip, used for the distribution of clock and broadcast signals, is lodged on a CCB Link Board. The clock is distributed using Pseudo-ECL signals on twisted pair cables. One clock input with low skew is available for each TRB. A low-skew clock distribution tree, consisting of PLLs and multiple drivers, is implemented on the TRBs. The clock

phase of the SB is adjustable with respect to the TRB clock to ensure the synchronization of the signals. The maximum allowed skew from TTC output to ASIC input is about 1 ns.

Further details on the characteristics of the boards can be found in Ref. [7].

The MCs are cooled by water flowing in tubes extruded in the MCs aluminum profiles along the full length.

2.2. Bunch and track identifier

The BTI is directly interfaced to the front-end electronics of the muon chambers. Using the signals from the wires it generates a trigger at a fixed time after the muon crossed the chamber, allowing bunch crossing identification, and it computes the position and direction of the muon. A complete description of the algorithm implemented in this integrated circuit was already given in a previous paper [9], but we shortly recall the most important logical steps of the algorithm and the information relevant to the tests performed in May 2003.

2.2.1. Working principle

The BTI chip exploits the generalized Mean-Time method [10]. It was explicitly developed to extend this technique to work on groups of four layers of staggered drift tubes. The aim is the identification of tracks detected in at least three out of four measurement planes.

The method relies on the fact that the path of a particle inside a SL is a straight line and that the measurement planes are equidistant. Considering the drift times of any three adjacent planes of staggered tubes (e.g. cells in layers A, B and C of Fig. 3) it can be shown that the relation

$$T_{\text{MAX}} = \frac{T_A + 2T_B + T_C}{2} = \text{const.}$$

holds independently of the impact point and angle of incidence of the track if there is a linear relationship between track position and drift time to the wire. The quantity T_{MAX} is called Mean-Time and corresponds to the maximum drift time in the ideal case of wall-less drift tubes.




Fig. 3. BTI geometric layout showing the channels allocation and important parameters.

Actually the BTI digitizes the shift time T_s after particle detection at 80 MHz frequency and computes the apparent drift time $T = T_{\text{MAX}} - T_s$ at every clock cycle. This calculation gives the real drift time only at the time T_{MAX} after particle crossing. Therefore the digitized times have values satisfying the relation only at that clock count while the relation does not hold true at any other time. The constant time difference between the particle crossing and the detection of the validity of the relation allows the identification of the parent bunch crossing. Furthermore since only at the time T_{MAX} after muon crossing the apparent drift times are the actual drift times, they are said to be *aligned*, i.e. the hits form an image of the muon track, thus allowing the extraction of the full track information (track impact position and direction).

The extension of the method to four layers has the advantage that a bunch crossing identification is possible even if the drift time of one tube is missing or wrong, for instance due to the emission of a δ -ray masking the good hit, since there are still three useful cells giving the minimum requested information. The Mean-Time method is also insensitive to uncorrelated single hits and is therefore well-suited to a high radiation environment.

2.2.2. Algorithm description

Each BTI is connected to nine wires allocated as shown in Fig. 3. All SLs are equipped with one BTI every four wires and therefore the BTIs overlap by five wires, i.e. the next BTI starts from cell labeled 5 in Fig. 3. This overlap assures the

redundancy needed to limit the inefficiency in case of a BTI failure.

The evaluated parameters are the position, computed for the SL centre, and the angular k -parameter $k = h \tan \psi$ where ψ is the angle of the track with respect to the normal to the chamber plane in the transverse projection; $h = 13$ mm is the distance between the wire planes.

The actual BTI algorithm for track candidate finding computes in parallel several track pattern hypotheses: a pattern is identified by a sequence of four wire numbers labeled L or R for a track crossing a drift tube on the right or on the left side of its wire (e.g. in Fig. 3 the track corresponds to the pattern 5L3R6L4R). Any given pattern includes six pairs of planes (AB, BC, CD, AC, BD, AD) each one providing a measurement of the position (through an x -equation) and of the k -parameter (through a k -equation) of the track.

The equations are computed at every clock cycle using the time after hit arrival digitized with 12.5 ns resolution. At any clock cycle the value of a k -equation corresponds to a rough measurement of the track direction at that cycle and it is time dependent. The k -equations are used to identify the bunch crossing. As an example the k -parameter is shown in Fig. 4 for the pair 3L2R. Its value is

$$k_{32} = T_{\text{MAX}} - T_3 - T_2.$$

Since the BTI can only use the shift times computed after hit detection on the wire the corresponding k -equation is

$$E_{32}(t) = T_{s3} + T_{s2} - t = (T_{\text{MAX}} - T_3) + (T_{\text{MAX}} - T_2) - t.$$




Fig. 4. Definition of the k -parameter.

This equation gives the correct value k_{32} only when computed at the right time $t = T_{\text{MAX}}$. Therefore each pair included in a pattern gives its own measurement of the track direction at every clock cycle: the hits are *aligned* when, after applying a pair dependent proportional factor correcting for its lever arm, the values of the k -parameter computed for each pair are equal.

Hence at every clock cycle the whole set of k -equations is computed and a BTI trigger is generated if at least three of the six k -parameters associated to any of the patterns are in coincidence. The coincidence of the k -equations values is verified within a programmable tolerance window. This tolerance is defined according to the resolution of each pair that in turn depends on the distance between the wires and was chosen to allow a maximum cell linearity error equivalent to 25 ns. The coincidence allows the bunch crossing identification owing to the time-dependence of the k -equations values.

If there is a coincidence of all six k -parameters, the trigger corresponds to the *alignment* of four hits and is marked as High Quality Trigger (HTRG). In any other case, with a minimum of three coincident k -parameters, the trigger is due to the *alignment* of only three hits and is marked as Low Quality Trigger (LTRG). The angular resolution is track pattern dependent and is generally worse for LTRGs.

If several track patterns give a response, the HTRG is chosen as the triggering track pattern. If there is more than one HTRG or the triggers are all LTRGs, the first one in an arbitrarily defined order is selected.

Only one track can be computed per clock cycle inside a BTI: its k -parameter, its position and the quality bit (H/L) are forwarded to the TRACO connected to the triggering BTI.

The request of the *alignment* of any three hits is a substantial source of background since it introduces effects creating false triggers. It is possible that the *alignment* of four hits at some clock count produces the *alignment* of only three hits at the preceding clock count or after the HTRG signal, thus generating *ghost* LTRG track candidates. It is also possible that a random LTRG occurs at an arbitrary clock count with a

valid k -parameter due to the left–right ambiguity, that is duplicating the possible choices for each hit. Finally δ -rays produced inside a cell will provide wrong time measurements enhancing the probability of an out-of-time trigger.

The noise production mechanisms and the implemented noise reduction methods will be discussed in Section 3.

With the present geometric parameters of the chambers the BTI equations are fully covering the angular range up to $\psi_{\text{MAX}} = \pm 45^\circ$.

The total latency of the BTI depends on the drift velocity and the drift path length and it is 22 LHC clock cycles at the nominal drift velocity (56 $\mu\text{m}/\text{ns}$) for CMS chambers.

The hardware implementation of the algorithm was realized in a 64-pin ASIC with CMOS 0.5 μm Standard Cell technology.

2.3. Track correlator

The BTI is followed in the electronics chain by a TRACO. The TRACO interconnects the two SLs of the φ view. It receives the information from the BTI devices connected to it and tries to find the couple of BTI track segments that fit the best track, linking the inner layer candidates to the outer layer ones.

The introduction of this device is necessary since the BTI is intrinsically a noisy device and therefore a local preselection and a more sophisticated quality certification of the BTI triggers is required. Furthermore the number of BTIs per chamber is around a few hundred and it is not possible to connect all channels together to perform a preselection at chamber level.

2.3.1. Algorithm description

The number of BTIs connected to a TRACO is limited by the size and by the number of pins of the chip and it is determined by the acceptance requirements. In the chosen layout four BTIs of the inner SL are connected to 12 BTIs of the outer SL as shown in Fig. 5, assuring a full coverage up to ψ_{MAX} .

The block diagram of the TRACO operations is given in Fig. 6. We describe the TRACO algorithm referring to the flow sketched in this diagram.




Fig. 5. Track Correlator layout.

There are four *data flows* inside a TRACO: two track calculation flows and two track Preview flows. Inside a TRACO there are in fact two parallel flows delayed by one cycle: the first path computes the First Track chosen from all BTI candidates while the delayed second path computes the Second Track from all unused candidates. The duplication of the data flows is imposed by the requirement of di-muon identification inside the same correlator. The preferences, described later in detail, for the choice of the First Track and the Second Track can be programmed independently, although we believe that basically the same criteria should apply.

The communication between the TRACOs and the TS is done using a dedicated Preview data bus for each track, in order to minimize the time needed for calculations of the whole trigger chain. The Preview is the copy of the k -parameter of one of the candidates chosen for the correlation and is selected according to programmed selection flags (H/L identifying trigger quality and IN/OUT choosing between inner and outer SL candidates). The TS selection is based on the quality of the Preview of the track candidates sent from each TRACO. The Preview data are coded in 9 bits: 5 bits for the modulus of the k -parameter, 1 bit for the track quality (H/L); 1 bit identifying First/Second track; 1 bit identifying Inner/Outer layer; 1 bit identifying Correlated/Uncorrelated track candidate.

The *Input Register* (16×8 bits) receives and latches the data values and the qualification flags from the 16 BTI chips connected to it. The input

data bus from each BTI provides the k -parameter and the position in the BTI coordinate system, multiplexed at 80 MHz on the same lines (6 bits wide). Two extra flags are provided: the trigger quality (H/L) and the strobe.

The *Angle and Position Converter* module receives the k -parameter input word from the BTI and converts it into local radial coordinates. The converted angle is used for internal calculations and sent on the Preview bus to the TS for further track selection. Each position calculated by a BTI is converted into a position in the TRACO coordinate system by adding an appropriate offset according to the geographical region of the SL the BTI is connected to. An additional SL shift parameter is provided to correct for possible construction misalignments between the two φ SLs.

The *Sorter* module receives the converted angles and selects among all candidates the track segment with the smallest angle, i.e. the closest angle to the local radial direction to the vertex. This choice corresponds to the selection of the BTI candidate with the highest transverse momentum. There is one sorter for the four inner BTIs and another for the 12 outer BTIs. Hence the choice is done twice on the two φ SLs independently. The sorting operation can be programmed to give preference to candidates tagged with the HTRG quality flag. Two other sorters are used in the Second Track path.

The *Calculator and Comparator* module computes the k -parameter and the position of a correlated candidate. It transforms the inner and

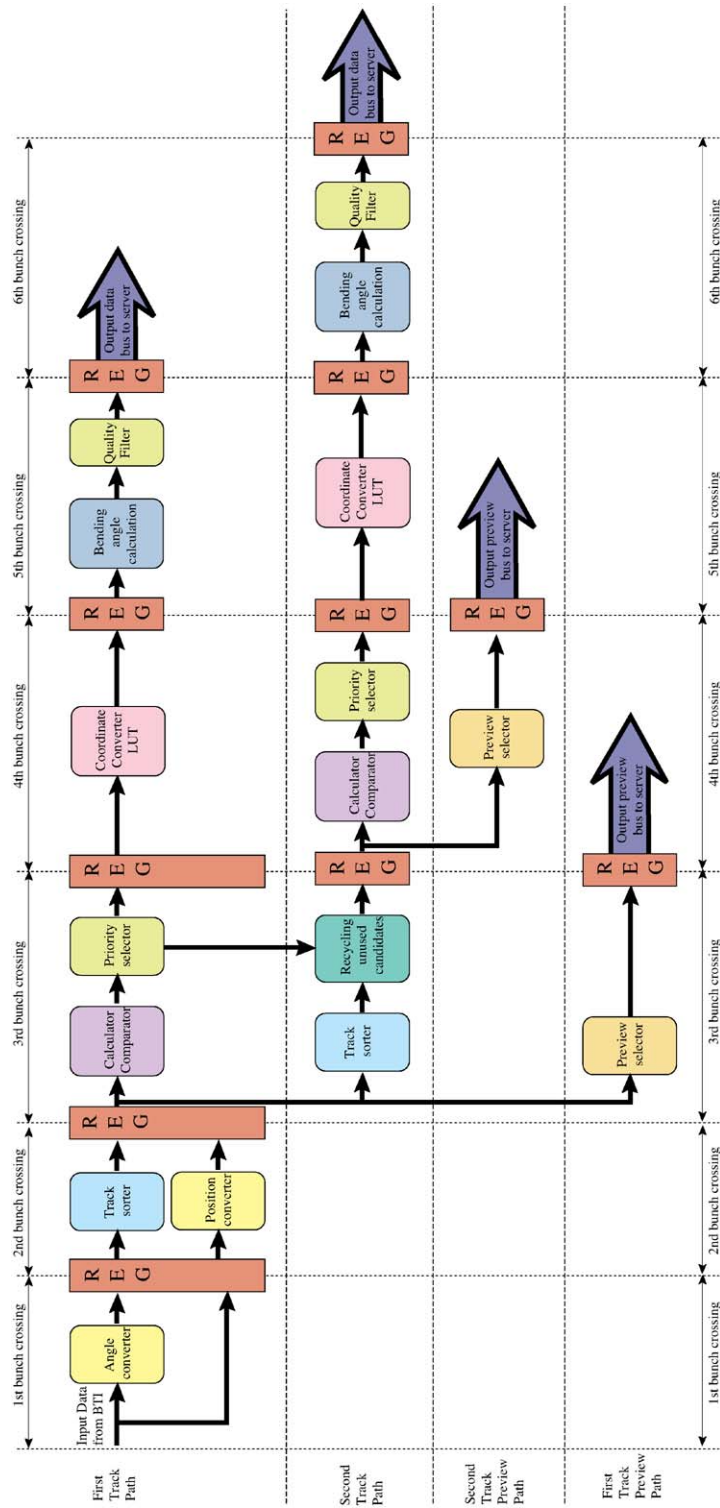


Fig. 6. TRACO block scheme.

outer k -parameter of the two independently selected track segments into the correlator coordinate system and computes the correlated track parameters. The internal parameters computed for the correlated tracks are:

$$k_{\text{COR}} = \frac{D}{2} \tan \psi = x_{\text{inner}} - x_{\text{outer}}$$

$$x_{\text{COR}} = \frac{(x_{\text{inner}} + x_{\text{outer}})}{2}$$

where D is the distance between the φ SLs and ψ is the angle of incidence as shown in Fig. 5.

Due to the long lever arm between the two φ SLs the angular resolution of a correlated track candidate at normal incidence is ~ 6 mrad for the nominal drift velocity, thus improving the BTI angular resolution (~ 50 mrad) while the position resolution is not significantly changed. The compatibility between the k -parameters of the selected track segments in the inner and outer SLs and the correlated track is checked against a programmable tolerance. If the correlated track fits inside the programmed acceptance window a correlation flag is raised.

The *Priority Selector and Preview Selector* module selects one of the candidates according to programmed criteria. If the correlation was successful the priority selector chooses the correlated candidate and forwards its parameters to the further stages. If the correlation fails the correlator creates an uncorrelated track following a preference list that includes the parent SL (IN/OUT) and the quality bit (H/L) of the two track candidates. A further selection can be activated to connect the trigger generated in the φ view to the triggers generated by the BTIs in the θ view. A programmable coincidence between the two views can be activated to certify the uncorrelated triggers. In particular, since the noise generated by the BTI algorithm is concentrated in the LTRGs, this coincidence is requested by default for the LTRGs and it is optional for the HTRGs. If no correlation is possible since there is no candidate in a SL, the existing uncorrelated track is still accepted.

The priority selector sends only one candidate towards the output bus, and generates a 3 bits qualification code identifying the HH, HL, LL,

H_i, H_o, L_i, L_o track candidates: the first three types identify correlated tracks, while the others identify uncorrelated tracks in either outer or inner SLs with obvious symbol meaning. One code value is reserved as a data strobe to mark null tracks.

Should the correlated track not satisfy the acceptance value, one out of the two track candidates selected to try the correlation is forwarded as the First Track choice and the other one can be reused for the Second Track calculations. This task is performed in the *Recycling Unused Candidates* module. This feature can be software disabled.

The two selected tracks are output on the same bus at consecutive bunch crossings. Therefore it is possible that a Second Track from bunch crossing n is computed at the same time of a First Track from bunch crossing $n + 1$. A First Track choice has always priority on the output bus and therefore overlaps the Second Track from bunch crossing n . The *Mixer* block performs this choice and activates a flag if an overlap occurs.

Inside the *Coordinate Converter and Bending Angle Calculation* block the internally calculated position and k -parameter are converted into the chamber reference system: the position is transformed into the radial angle φ and the k -parameter is converted into the bending angle φ_b , as defined in Fig. 7. The task is performed with direct access to two programmable look-up tables. The first look-up table is used for the conversion of the local correlator position (coded in 9 bits) into the track radial angle φ (coded in 12 bits). The second one performs the conversion from the k -parameter into the angle ψ , both coded in 10 bits. A further block performs the computation of the bending angle $\varphi_b = \varphi - \psi$.

Some filtering functions are performed in the *Quality Filter* block to select the output value passed to the TS. These functions include an uncorrelated Low Trigger Suppression and a programmable tolerance window for the bending angle output value. All filters will be discussed in detail in Section 3.

Track data (10 bits for the bending angle, 12 bits for the radial angle and the three quality bits) are output on a bus. The data output bus provides one track at each clock cycle with up to two tracks per




Fig. 7. Definition of TRACO output parameters.

bunch crossing at consecutive clock cycles. In order to avoid arbitration conflicts a TRACO is allowed to forward its track candidates to the DTF only upon authorization by the TS system which performs the final selection.

The latency for the First Track is five clock cycles while the Second Tracks are output after six clock cycles.

The TRACO was realized in a 240-pin ASIC with CMOS 0.35 μm Gate Array technology.

2.4. Trigger server

The TS has to select the two best trigger candidates among the track segments selected by all TRACOs in a muon station and has to send them to the Sector Collector where the trigger candidates will be forwarded to the Regional Muon Trigger.

The TS is composed of two subsystems: one for the transverse view (TS_ϕ) and the other one for the longitudinal view (TS_θ). The first one processes TRACO outputs while the second one works on the output of the BTIs of the θ view. It must be

noted that there is only one TS for each station: therefore it represents the bottleneck of the on-chamber trigger devices and it should have built-in redundancy.

The number of TRACOs in a station will be as much as 25 for the biggest station. Each TRACO transmits its two best tracks serially to the TS_ϕ , ordered in quality, during two consecutive bunch crossings. In order to minimize the latency of the TRACO- TS_ϕ system the TRACOs send Previews of track segments. While the TS_ϕ makes its selection (in pipeline) the TRACOs compute the full track parameters (in absolute coordinates with higher precision). Then the TS_ϕ serially reads the full track parameters of the two best tracks from the corresponding TRACOs and sends them to the Sector Collector. This mechanism allows to gain two bunch crossings and to limit the total latency of the TRACO- TS_ϕ system to six bunch crossings.

Since important physics processes involve a pair of close-by muons that can hit the same muon station, some emphasis should be put on the efficiency and purity of both selected segments. We define *bunch1* and *bunch2* as the first and the second bunch of tracks arriving from the TRACOs connected to the TS_ϕ . The sorting algorithm could be simple if it was just selecting independently the best track of *bunch1* and the best track of *bunch2*. However, it is not assured that the best track of *bunch2* represents the second-best track among all *bunch1* and *bunch2* tracks. The TS_ϕ must identify the truly second-best track among all *bunch1* and *bunch2* tracks. In order to achieve this the TS_ϕ selects among the tracks of *bunch1* the first-best-track (FBT) and the second-best-track (SBT). At the following bunch crossing the search for the best track is done among *bunch2* tracks and the SBT of previous bunch crossing (called *carry*). Therefore the sorting algorithm is applied in pipeline for each bunch crossing. In this way the truly two best tracks among all the possible tracks are found. For normal triggers the TS_ϕ is able to find the two best tracks within two bunch crossings. In case of *pile-up* events the TS_ϕ can provide at least the FBT data resulting from the sorting of *bunch1* to the Sector Collector. In case of two close-by muons they likely produce two track segments in *bunch1* by different TRACOs in the

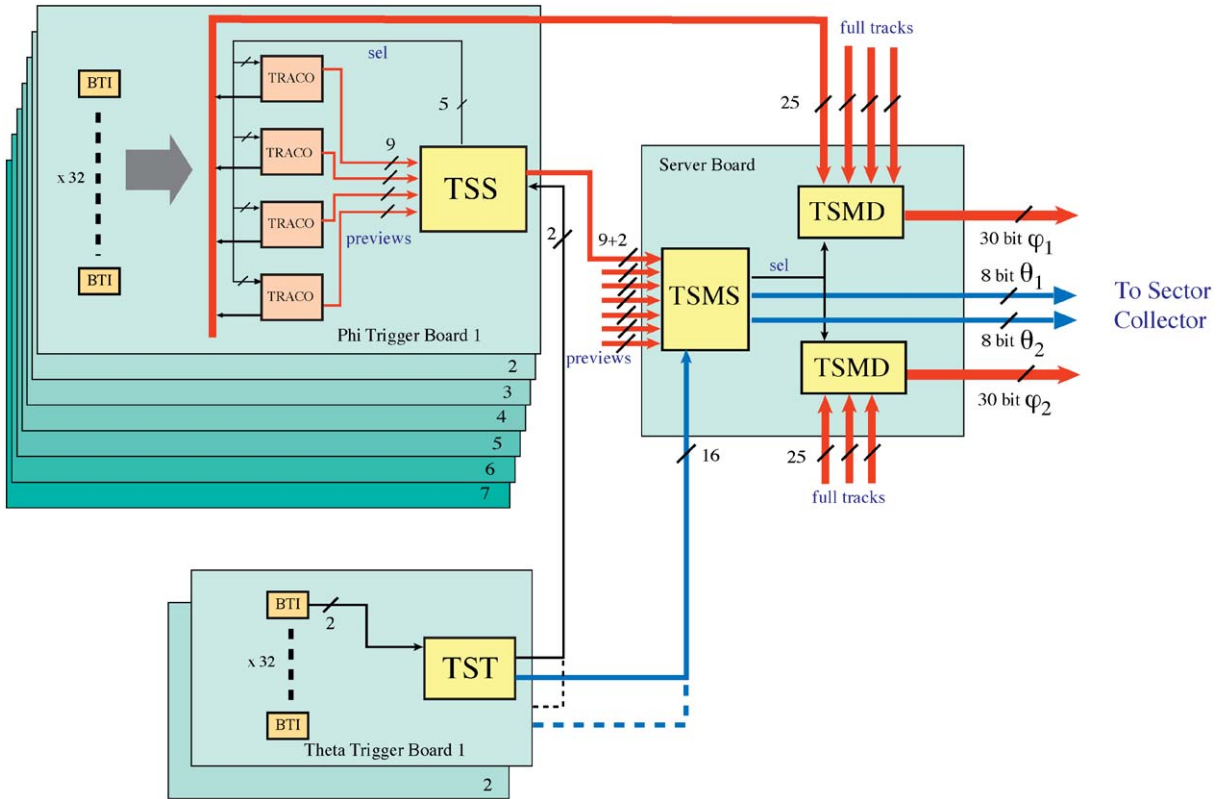


Fig. 8. Trigger Server architecture. The Trigger Server is composed by two sub-systems. The $TS\phi$ includes TSSs, TSMDs and TSMS and the $TS\theta$ includes TSTs.

same station which are correctly picked up by the $TS\phi$ algorithm through the *carry*.

The TS has to reject ghosts generated by the TRACOs, using configurable noise rejection algorithms, and should be able to treat *pile-up* events. A *pile-up* event occurs when a TRACO in the station delivers at two consecutive clock cycles two First Track candidates associated to two consecutive bunch crossings while other TRACOs deliver the normal sequence of First and Second Tracks, both associated to the same bunch crossing. In this case a mixture of First Tracks and Second Tracks is delivered by the TRACOs at the same clock cycle, corresponding to two consecutive bunch crossings. The TS has to decide whether to perform the scheduled sorting of the Second Tracks of *bunch2* or to drop the Second Tracks and to do another *bunch1* sorting among the incoming new First Tracks.

The $TS\phi$ logic diagram is shown in Fig. 8. The selection algorithm uses a two-layer cascade of processing units. This architecture was chosen in order to minimize the number of logic cells within a unit and the amount of I/O between blocks. In each unit a parallel minimum and next-to-minimum search is performed over a small group of input words, using 2×2 fast 9-bit comparators. The fully parallel approach guarantees a fixed time response independent of the number of TRACOs in a station. Each unit of the first layer (TSS) is connected to four TRACOs while the second layer unit (TSM) can be connected to up to seven TSSs.

The $TS\theta$ is formed by two identical units (TSTs) which form the OR of groups of BTIs. The information about the presence of a trigger in the θ view and its quality is sent to the TRACOs via the TSSs and can be used as trigger validation in the ϕ view. Besides a pattern of trigger hits is

sent to the TSM and forwarded to the Sector Collector.

As shown in Fig. 8 the hardware partitioning of the system matches with the logical blocks. Each TSS device is mounted on a TRB which contains four TRACOs and 32 BTIs. The TSM system is composed of three devices mounted on a SB which receives the output of at most seven TSSs (for the largest chamber). The two TST devices are mounted on two separate boards including 32 BTIs each.

2.4.1. Track sorter slave

The main tasks performed by the TSS are the sorting of Preview data coming from the four TRACOs placed on the same board and the suppression of noise generated by them.

The best track is the one with best quality and smallest angle (which means higher transverse momentum). If the quality bits are coded in an adequate way, the search for the best track is a search for a minimum. During one bunch crossing the TSS is able to activate a select line addressing the TRACO which sent the best Preview: then the TRACO will be allowed to send the corresponding full track parameters to the TSM for further processing. At the same time the best Preview is also sent to the TSM for the second stage processing.

The ghosts that should be treated from the TSS are generated either by superpositions between TRACOs or by correlation failures. In these cases duplicates of the same track are sent to the TSS. The ghosts due to superpositions can involve nearby TRACOs belonging to different boards and therefore the final ghost suppression can only be done by the TSM.

The functionality of each TSS is performed in two consecutive cycles, called *sort1* and *sort2*. The *sort1* processing status is recognized when at least one TRACO gives a non-null track of *bunch1* type while the *sort2* status simply corresponds to the cycle following *sort1*. The *sort2* cycle can be aborted in case of *pile-up* triggers. In the *sort1* cycle each TSS unit analyses four Preview data words and transmits the minimum to the TSM unit in the second layer while the next-to-minimum is stored locally and carried over to the *sort2* cycle.

At the same time a local select signal is output to enable the transmission of the full data from the selected TRACO to the TSM. In the *sort2* cycle each TSS unit analyses the four input words of *bunch2* together with the *carry* word of the *sort1* cycle. If the *carry* is the best trigger candidate in the *sort2* cycle, a post-select line is used to inform the TRACO that it has to transmit its best track segment of the previous cycle to the TSM.

The TSS functionalities are implemented in a single 120-pin ASIC chip built in CMOS 0.5 μm technology. Because of the severe speed requirements (sorting of two out of four 9-bits words with *carry* in one bunch crossing) it is not possible to use programmable ASIC devices like it is done for the TSM. The main building blocks of TSS are shown in Fig. 9.

The processing unit of the TSS is called *Sorting Core*. First the four 9 bits preview words from the TRACOs are filtered: they can be masked individually or depending on their quality and the priority of their quality for the sorting can be changed. Then they enter a battery of ten 2-words comparators together with the *carry* track of the previous cycle. The First Best Track is sent to the TSM while the Second Best Track is kept as *carry*. At the same time the select lines to TRACOs are activated. The other blocks are used for configuration, control and monitoring of the integrated circuits. Using the internal configuration registers it is possible to steer the sorting algorithm. In fact each of the input Previews from the TRACOs can be masked in case of noisy channels. The priority order of the quality bits in the sorting stage can be swapped. During normal operation and in order of decreasing importance they are: correlation, quality of the trigger, position of the trigger (i.e. IN/OUT SL), angular deviation with respect to the radial direction. Furthermore the *carry* mechanism can be disabled as well as the algorithms for ghost suppression.

2.4.2. Track sorter master

The TSM unit in the second processing layer of TS system analyses up to seven Preview words from the TSSs. There is one TSM per muon station. It behaves similarly to a TSS unit but its processing starts two bunch crossings later. The

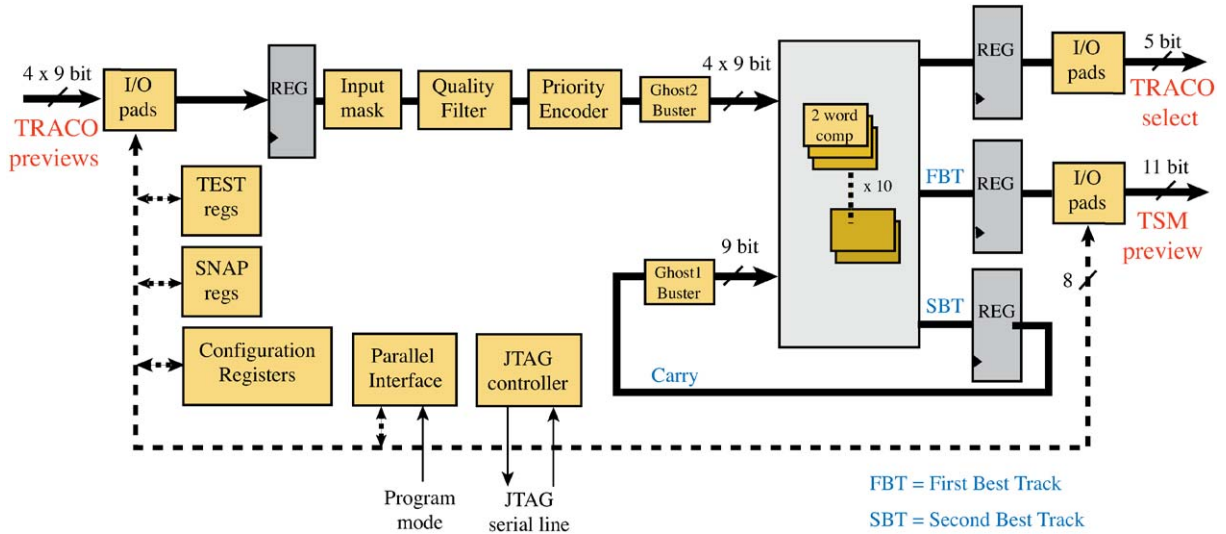


Fig. 9. Main blocks of TSS design.

information handshake between the TS ϕ system and the TRACO devices allows data from up to 14 tracks to be stored in the TSM unit. The selected output signal from the TSM, corresponding to the FBT in the first processing cycle and to the SBT in the second cycle, is used to enable the transmission of full track data to the Sector Collector for two out of the 14 possible candidates stored in the TSM unit. The TSM also receives information from the θ view. The hit pattern received from the TS θ is synchronized with the ϕ view track data (signal processing is much faster in the θ view) and forwarded to the SC.

The TSM system has two different logic components (shown on the right side of the scheme in Fig. 8): a sorter block (TSMS) which performs the sorting of the Previews from the TSSs and a data multiplexing block (TSMD) which outputs the full data from the TRACOs, corresponding to the selection done in the TSMS.

Robustness is the focal point in designing the TSM system since the TSM operates both as last functional element of the track-segment sorting tree of a drift chamber and as nodal point for distribution of the monitoring and configuration information to the various elements of the chamber local trigger. From the point of view of

functionality the main constraint comes from di-muon physics: this means the capability of the system to maintain a good efficiency also in case of hardware failures. In order to achieve this the system was segmented in blocks with partially redundant functionalities. Care has been taken that the segmentation does not deteriorate the expected performance, in particular the latency assignment.

The TSM is split into three parts: a TSMS block and two TSMD blocks (TSMD0, TSMD1). The TSMS inputs are the Preview data from the TSSs. The TSMDs inputs are the data of the selected track candidates of half a chamber each.

The TSM can be configured in two processing modes. In the default processing mode the TSMS sorts the TSS Preview data and issues select signals that the TSMDs use to choose between the track candidates coming from TRACOs. The TSMS can select two tracks in TSMD0 or two tracks in TSMD1 or one track in each of them. In the back-up processing mode the TSMS processing is bypassed: each TSMD sorts the best track candidate among the data of a half chamber and outputs one track. In case of failure of one TSMD the Previews of the corresponding half chamber are disabled in the TSMS sorting, so that full

efficiency is maintained in the remaining half chamber. A similar technique is used in case of damage of the data lines from TRACOs.

The default processing implements the full performance and guarantees that di-muons are found with uniform efficiency across the chamber. The back-up processing is activated in case of TSMS failures or in case of damages in the Preview lines. It guarantees full efficiency for single muons and for open di-muon pairs (one track in each half-chamber).

Besides performing the second stage of the sorting algorithm begun in the TSSs the track selection procedure also applies data masking and ghosts rejection in a way consistent with the TRACO–TSS system. In the TSM the rejection of the ghosts generated by TRACO superpositions is completed. In fact the ghosts showing up in TRACOs at the border of contiguous TSSs cannot be suppressed by the TSSs. Since each TSS serves four TRACOs and this particular kind of ghosts can only appear in two contiguous TRACOs, each TSS forwards to the TSM 2 bits giving the relative address of the selected TRACO. The TRACO with address 00 of the TSS_i is adjacent to the TRACO with address 11 of the TSS_{i-1} . If such adjacent segments are found, one is cancelled if it is an outer segment, i.e. a segment observed only in the outer SL: no duplication of inner SL segments is possible in the TRACOs by construction.

The TSM logic is implemented using three identical Actel A54SX32 chips which, tested for space applications, have shown a good tolerance to high radiation doses up to 10–50 krad and high thresholds for Single Event Effects. They belong to a new generation of FPGAs also called pASICs (programmable ASICs) based on the 0.35 μm silicon antifuse technology: once programmed the chip configuration becomes permanent making them effectively ASICs.

2.4.3. Trigger server in the longitudinal view

The TST has to group information from the 64 BTIs in the θ view and to send a suitable pattern to the Regional Muon Trigger. Studies on the DTTF algorithm suggest that the minimum pattern resolution, without loss of efficiency, corresponds to the area covered by eight BTIs. Hence it receives

2 bits from each BTI (trigger strobe and H/L quality) and, for both bits, it performs a logic OR of groups of eight BTIs. The output is formed by 2 bits for each group: in total 8 bits for the trigger position (1 bit every 16 cm) and 8 bits for the corresponding quality. These signals are sent to the SB where they are synchronized with the track segments found in the φ view and sent to SC which transmits them to the DTTF. Moreover it sends a 2-bit signal (64 BTI wired OR + 64 BTI Quality wired OR) to the TSSs where they are forwarded to the TRACOs to support the noise reduction. The TST is made of two identical devices, each one located on the TRBs of the θ view and connected to 32 BTIs. TSTs are built using commercial ICs.

2.5. Sector collector

The SC will be lodged on the side of the cavern containing CMS. All high speed optical links either for readout or for trigger are placed on the SC. Readout data, coming from each ROB on serial medium speed links (LVDS at 240 Mbps) on twisted pairs, are grouped and formatted to form sector readout packets and sent, via a high speed optical link, to the DAQ. Trigger data, coming from each SB on serial high speed links (LVDS at 480 Mbps) on twisted pairs, are grouped to form sector trigger packets and sent, via high speed optical links (Gigabit Ethernet at 1.6 Gbps), to the DTTF.

The link connecting each SB to the SC consists of eight chipsets of 10-bit LVDS Serializer/Deserializer connected by two variable length Category six Ethernet cables, each consisting of four screened twisted pairs. The data will be packed inside the SC and sent to the DTTF modules over optical fibres, separately for φ view and θ view information of each sector and multiplexing First and Second Track data.

3. Noise reduction

The design of the trigger devices was done with the purpose of providing a robust and efficient system. Unfortunately the way to meet these requirements introduces a certain number of

redundancies in the system causing an important fraction of false or duplicated triggers. Therefore noise reduction mechanisms must be implemented.

3.1. Noise generation mechanisms

The BTI trigger algorithm can actually work requiring only three layers of staggered tubes. The drawback of a three-layers algorithm is that an inefficiency or a bad measurement on any of the cells yields an inefficiency or a wrong trigger. The introduction of the fourth layer with the minimal request of an *alignment* of three out of four hits improves the efficiency and reduces the rate of wrong measurements. Still spurious *alignments* of three hits can occur at any bunch crossing, depending on the actual track crossing position and direction. Most of the bad *alignments* are generated by the unavoidable left–right ambiguity. The effect is that fake *alignments* of hits which do not correspond to the actual direction of the incident muon are found at wrong bunch crossings. An example of the mechanism is shown in Fig. 10 where a real track orthogonal to the chamber is displayed and the hit positions are

marked with small circles on the track line. The BTI is able to find the *alignment* corresponding to the real track, but other two tracks are detected as well. These tracks, called *ghost tracks*, correspond to *alignments* of a mixture of real hits and their mirror images. Indeed the BTI finds a false *alignment* at the time Δt_1 after the correct bunch crossing, supposing that wire 2 is inefficient and that the signal of wire 4 comes from the right side of the tube. In the same way, supposing that wire 5 is inefficient, the BTI finds another ghost track formed by the signals of wires 2 and 4 and by the mirror image of the signal from wire 3 at the time Δt_2 before the correct bunch crossing. This effect occurs inside the BTI at different bunch crossings and therefore generates temporal noise.

Furthermore, in order to be fully efficient the trigger system provides an overlap between adjacent devices: each BTI is overlapped by five cells to its neighbours and BTIs in the outer SL are always assigned to three consecutive TRACOs. The overlap is foreseen to minimize the trigger inefficiency due to the loss of a device since the remaining one can be programmed to partially cover the dead area switching on some redundant




Fig. 10. Illustration of the temporal ghost generation mechanism inside BTI.

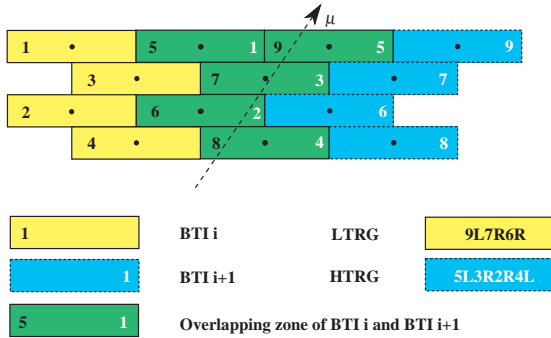


Fig. 11. Illustration of the irreducible redundancies between overlapping BTIs. A valid HTRG pattern in BTI_{i+1} (cells marked by white labels) is seen as a valid LTRG pattern in BTI_i (cells marked by black labels). In fact for BTI_{i+1} the LTRG pattern 9L7R6R is common to HTRG patterns 9L7R6R4R (valid) and 9L7R6R8L (redundant). Even if the redundant pattern is switched off by configuration choice, the LTRG pattern remains valid for cases where cell 4 is inefficient.

patterns. It is possible to define a set of non-redundant HTRG patterns, but the underlying LTRG patterns may be common to more non-redundant patterns and therefore some of them are available in two consecutive BTIs: in fact there are five patterns generating LTRGs on the devices close to the one generating the HTRG at the same bunch crossing. In Fig. 11 we see a case where a valid HTRG pattern in one BTI is also seen as a valid LTRG pattern in the adjacent one. Therefore the TRACO may have to choose between track candidates in adjacent BTIs that are images of the same track, carrying exactly the same information. This is not a problem for the First Track sorting since both are equivalent. But the result may be that the TRACO forwards the same track twice with a chance of losing other available candidates. This duplication of the trigger at the same bunch crossing generates spatial false triggers.

There is another situation where spatial noise track candidates are generated. The BTIs in the outer SL are assigned to three consecutive TRACOs, being in the left, the central or the right group of the outer SL. The BTI data are sent to each TRACO through a dedicated port, each programmed with a different angular acceptance window, depending on the TRACO group (left, central or right) the port is communicating to.




Fig. 12. Illustration of the double track selection due to overlapping TRACOs. The solid lines are the acceptance window of the i th TRACO and the dashed lines are the acceptance window of the $(i+1)$ th TRACO. The diagram on the right draws the acceptance windows at the same origin to evidence their intersection (shaded region). A muon falling in this intersection region is assigned to both TRACOs.

These tolerance windows partially overlap. Hence a track candidate falling in the intersection zone is forwarded by the BTI to more TRACOs as shown in Fig. 12. Hence, as in the case of adjacent BTIs, adjacent TRACOs may forward the same track twice to the TS introducing a bias in the Second Track selection at the TS level.

3.2. Noise reduction methods

Some filters were provided to reduce the overall importance of the spatial and temporal noise. Different noise reduction algorithms run independently in the BTI, the TRACO and the TS.

The BTI can only act on temporal noise. The first mechanism provided to perform this task is the Low Trigger Suppression algorithm (LTS) which cancels the LTRGs occurring in a programmable time range (at most -1 bunch crossing, $+8$ bunch crossings) around a HTRG. The drawback of this filter is an efficiency loss when a HTRG is issued at the wrong bunch crossing and a good LTRG is issued at the right bunch crossing since the right LTRG is cancelled by the nearby wrong HTRG.

We introduced also an optional LTS mechanism inside the TRACO: the lower quality tracks (LL, L_o , L_i) are cancelled if a HTRG occurred within neighbouring bunch crossings. It is possible to suppress triggers at bunch crossings from -1 to $+4$ with respect to any HTRG without any

latency addition. This algorithm is complementary to BTI LTS mechanism since it can suppress ghosts occurring in neighbouring BTIs.

Another algorithm available for noise reduction is the validation of φ view triggers by θ view triggers. Since the two views are independent the probability of ghost generation at the same bunch crossing in both views is smaller than in a single view. The θ view trigger is forwarded to the TRACOs and uncorrelated triggers (LTRGs by default and HTRGs optionally) are discarded if they are not associated to a valid θ view trigger. The θ view trigger validity is defined by the TS system and can optionally request just the presence of any trigger or the presence of a HTRG.

Spatial noise affects only the Second Track selection. It is possible to avoid sending twice the same track using geometrical suppression filters.

If a HTRG was selected in the First Track sorting operation of the TRACO flow, all LTRGs in the neighbouring BTIs are removed from the Second Track sorting list. This filter is always active on ghosts associated to redundancies in the BTI equations.

A ghost generated by the BTI acceptance ports has to be an Outer segment. This kind of ghosts can be cancelled by removing *carry* tracks of Outer type in the TRACO close to the one that sent the best track inside the chamber TS, both at TSS and TSM level. This kind of ghosts are called *ghost1* in the TS system and the suppression filter is optional.

Another potential source of spatial background is the correlation failure of a TRACO. In this case the two track segments of the Inner and the Outer SLs belonging to the same track are sent as a First Track (the Inner segment) and as a Second Track (the Outer one). This ghost can be removed by requiring that an Outer Second Track sent by the same TRACO which gave the best track in the previous bunch crossing is not valid. This kind of ghosts are called *ghost2* in the TS system and the suppression filter is optional.

A programmable tolerance window implemented for the bending angle can be optionally applied to clean the TRACO output. The average bending angle, as defined in Fig. 7, for some low momentum muons at all the measurement stations

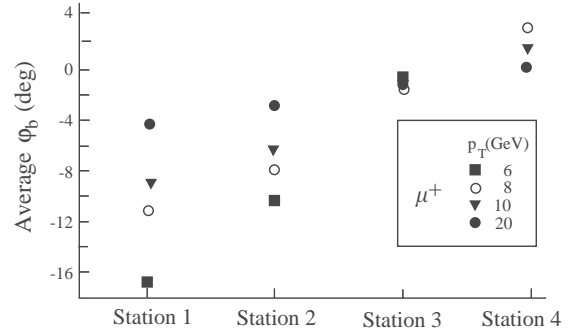


Fig. 13. Average bending angle at the different CMS muon measurement stations, calculated for some low transverse momentum muons.

is shown in Fig. 13. There is a large spread for the average bending angle values at stations 1 and 2 while the bending angle is very small at stations 3 and 4. Although we cannot safely apply any cut in the first and second station a tolerance on the bending angle can be used for station 3 and/or 4. This cut could have some importance in particular for station 4 where the θ view drift tube SL is not available. The cut is an 8-bit value to be downloaded into the TRACO.

4. Experimental layout and analysis tools

4.1. Detector arrangement

An MB3-type chamber [11] was installed in the H6 zone of the CERN SPS North Area and exposed to a secondary muon beam with 120 GeV/ c momentum. In contrast to future LHC conditions the chamber was not exposed to a magnetic field. The tests were carried out in May 2003 when the SPS radio-frequency structure was similar to the one foreseen for the LHC. A proton beam delivered by the SPS hit a primary target in narrow bunches (about 2 ns long, separated by 25 ns) generating muons. From these muons a secondary beam (having the same timing structure as the proton beam) has been set up by momentum selection. Trains of 48 bunches occurred every 23 μ s during a spill of 2.7 s length (a so called slow-extraction cycle of the SPS). Since a secondary beam was used the mean occupancy by muons of a

bunch was rather low, at the order of 10^{-2} – 10^{-3} , i.e. muons were separated in time by several microseconds on average.

The 40 MHz signal, synchronous with the accelerator RF signal, was distributed in the experimental area via a TTC system [8] through optical links. It was used as clock signal for the readout and trigger electronics in the MC. The MB3 chamber had the ϕ view wires in the vertical direction and was placed on a support which allowed rotations around its central vertical axis. The chamber was equipped with a MC foreseen for MB1 (i.e. some channels were not connected) with complete readout and trigger electronics.

The chamber was flushed with an Ar/CO₂ (85%/15%) gas mixture with a gas flow of 21/min. The oxygen concentration was measured by an oxygen monitor placed downstream before the gas exhaust and was always below 100 ppm during data taking.

The high voltage values of the electrodes were: $V_{\text{wire}} = 3600$ V, $V_{\text{strip}} = 1800$ V, $V_{\text{cathode}} = -1200$ V. The discrimination threshold of the front-end readout electronics was set to 5 fC. The setup discrimination threshold and HV values were those foreseen for operation at the LHC.

An external beam trigger was performed by the coincidence of two plastic scintillators defining a 10×10 cm² area. The scintillators were placed upstream in the muon beam approximately 4 m in front of the centre of the chamber (corresponding to the position of station 1 in the CMS detector). The time jitter, including the intrinsic jitter of the particles in a bunch and the time resolution of the scintillators, was measured with an oscilloscope using the 40 MHz clock signal from the TTC as time reference. It turned out to be around 3 ns.

The trigger to the MC (i.e. the start signal for the ROB) and to the DAQ system to read out and store an event was given by a Level 1 Accept (L1A) signal which was generated by the synchronization of the beam trigger signal with the 40 MHz clock from the TTC system. This was achieved by means of a custom PCI board, designed and manufactured by the CMS DAQ group which accommodated a TTCrx receiver mezzanine board and an FPGA, hosting a PCI bridge and the whole synchronization logic. The board implemented

also a trigger handshake mechanism with the DAQ system, in such a way that the broadcast of the L1A was conditioned to the ability of the DAQ to read a new event.

The trigger data were recorded inside a Pattern Unit [12] (a very fast deep FIFO) storing information at 40 MHz. The Pattern Unit was started by the L1A signal. Data available in the first 40 locations (slots) were readout. Each slot corresponds to a 25 ns clock count (i.e. the time distance between LHC bunch crossings) and therefore uniquely identifies the equivalent LHC bunch crossing assigned to each trigger.

The main components of the DAQ were: the Readout Unit (RU), the Builder Unit (BU), and the Run Control (RCMS). The Purpose of the RU was to collect all data from the ROB) and trigger electronics, buffer them locally and serve them to the BU. The RU block was implemented and run on a dedicated PC where a PCI/VME interface provided access to the ROS VME boards. The purpose of the BU was to receive buffered events from the RU through the Event Builder Network (a fast Ethernet Switch) and to write them into a file. The DAQ behaviour was controlled and monitored by the RCMS software through XML/SOAP messaging.

4.2. Detector performance

The actual behaviour of the chamber during the test was checked using the methods outlined in Ref. [6]. The results were perfect according to the expected performance: the drift velocity was 55 $\mu\text{m}/\text{ns}$; the single hit detection efficiency was better than 98%; the single hit position resolution was found to be angle dependent as shown in Fig. 14. In the ϕ SLs this dependence is due to the increasing deviation from linearity of the space–time relationship with larger angles of incidence of the particles, while the slight improvement seen in the θ SL is due to the longer paths inside the cells that increases the primary ionization statistics.

4.3. Trigger emulation algorithm

The predictions of the performance of the CMS muon barrel detector heavily rely on the precise




Fig. 14. Measured single wire resolution as a function of the angle of incidence. The resolution quoted for the ϕ view is the average of the two ϕ SLs.

emulation of the trigger flow and its correct description is therefore compulsory. The emulator was developed as a tool to be used in the offline simulation of the CMS detector. It is supposed to reproduce exactly the way front-end signals are treated in the hardware devices, implementing all features described in the previous sections. For this purpose a dedicated interface was developed which generates trigger input signals from the drift times recorded by the TDCs on the ROB. These signals are sent as input to the emulator in order to check the trigger response on an event by event basis. The algorithm is included in the official CMS reconstruction package. The validation of the trigger emulator was one important goal of the tests since it is the only way to derive the actual performance of the trigger in CMS. In fact this test was intended only to qualify the trigger algorithm because of the lack of background (neutrons, δ -rays and electromagnetic showers are expected in CMS) during the test beam. It was also done in the absence of magnetic fields which are not negligible at some positions of the CMS detector. Therefore the results cannot be transferred to CMS without a validated algorithm emulating the system performance.

5. Data selection

The selection of the data was done trying to avoid any possible bias on the trigger performance

evaluation. Different cuts were applied in order to select a clean single muon sample and a genuine di-muon sample.

5.1. Single muon event selection

Since the triggering device was a scintillator the main request is the presence of its signal whose time was recorded by two dedicated TDC channels. The tolerance on this signal was set to 2 ns. Defining t_0 as the TDC time corresponding to hits crossing the tube at the wire position, the existence of recorded signals within a 500 ns window after t_0 on at least two cells in the region illuminated by the beam was required in order to reject beam halo triggers.

The BTI has a programmable dead time that protects the computations in case of multiple hits on the same cell (due to afterpulses, electromagnetic background or multiple beam tracks). The first detected hit on each BTI channel raises a flag that rejects signals from late hits within a programmable time window (by default set equal to the maximum drift time T_{MAX}). Therefore only hits detected in the range $-T_{\text{MAX}} \leq t_0 \leq 2T_{\text{MAX}}$ may affect the BTI calculations: hits arriving before the good one will mask it modifying the trigger decision, while hits coming afterwards will restart the BTI calculations adding noisy triggers. Therefore we can clean the sample rejecting events with out-of-time muons by asking for less than three hits outside this time window without introducing any bias.

A relevant number of events with two muons are still left in this selection, mainly due to the large time window allowed. A smaller window cannot be applied without biasing the analysis, therefore we selected a single muon sample counting the number of hits recorded in the TDCs falling in the selection time window. We chose to select on the number of hits independently on each SL: we asked for at most six hits in at least one of the ϕ SLs (in order to keep possible electromagnetic showers generated inside any of them and absorbed in the aluminium honeycomb) and at most six hits in the θ SL (this cut rejects double muons hitting the same ϕ tube). The results of the selection on the normal incidence sample is shown

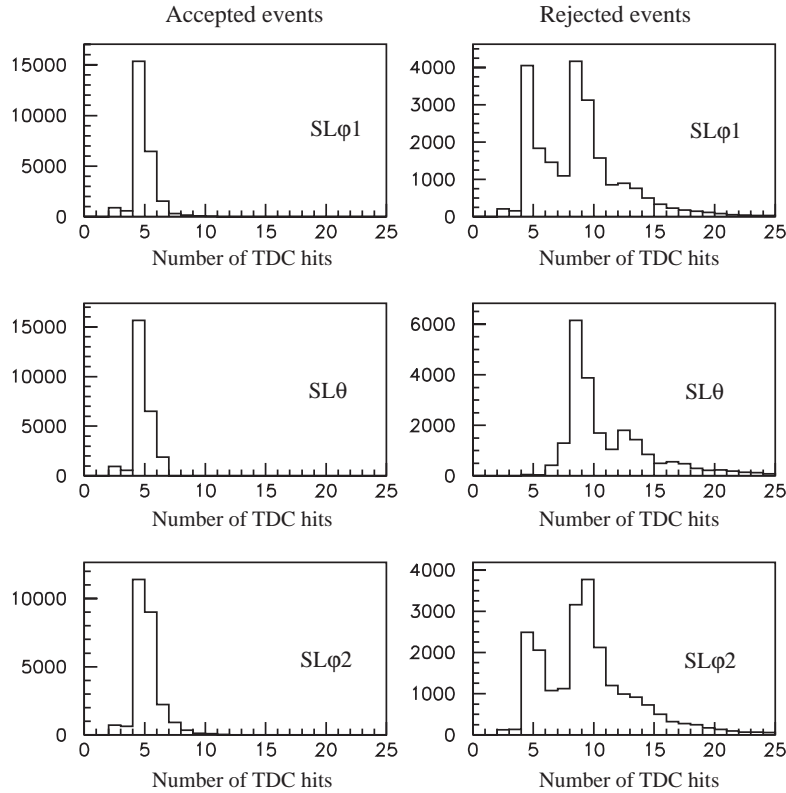


Fig. 15. Number of hits recorded by the TDCs inside the selection window ($-400 \text{ ns} \leq t_0 \leq 800 \text{ ns}$) for normal incidence muons. The plots on the left show this quantity for data in the single muon selection independently for each SL. The plots on the right report the same information for the events rejected by the applied cut on number of hits. The different peaks on these plots correspond to multimueon events. Data at other angles of incidence show the same behaviour.

in Fig. 15 where the distributions of the number of hits (for accepted and rejected events) in each SL are shown. The multimueon background in the selected sample is clearly negligible and the rest is a mixture of multimueons: the apparent single muon peak in the φ SL plots is due to multimueons identified by the cut on the θ SL and classified as muons traversing the same tube in the φ view. Evident peaks corresponding to two, three and four muons are visible. Similar results hold at any angle of incidence.

The final cuts identifying the single muons selection sample were therefore set as follows:

- (a) scintillator trigger time with $\pm 2 \text{ ns}$ tolerance around the average time,
- (b) \geq three cells with recorded hits in the beam region in the φ view,

- (c) \leq two hits recorded outside the time interval $-400 \text{ ns} \leq t_0 \leq 800 \text{ ns}$ in the φ view,
- (d) \leq six hits in at least one φ SL,
- (e) \leq six hits in θ SL.

After the application of all these cuts about 50% of the events are classified as clean single muons. The largest reduction was caused by cut (c).

5.2. Di-muon events selection

The beam is contaminated by halo muons and in addition the triggering muon can produce an energetic δ -ray. Hence it is not rare that the chamber is crossed by two tracks. Unfortunately most of the events with two muons are due to particles belonging to different beam cycles

(equivalent to LHC bunch crossings) and are separated in time by multiples of 25 ns. The track fitting which is based on the drift times recorded by the TDCs can still reconstruct them as double track events, but from the trigger point of view the two muons correspond to two different bunch crossings. In order to test the trigger performance in the case of di-muons, only events with track pairs crossing the chamber simultaneously (i.e. at the same bunch crossing) have to be selected. The selection has to be obtained from the information delivered by the chamber itself (the TDC hits on the three SLs and trigger tracks) as no external device was available to trigger directly on di-muon events. Therefore special care has to be taken to avoid a bias in the event selection. The following cuts were used to select di-muon events:

- (a) scintillator trigger time with ± 2 ns tolerance around the average time,
- (b) \geq one cell in the φ view of the chamber, with recorded hits in the beam region, in the time range $0 < t_0 < 400$ ns. We will refer to the hits in this time range as *in-time* hits,
- (c) \leq two hits not *in-time* recorded anywhere in the chamber,
- (d) \geq seven hits *in-time* in at least one of the two φ SL,
- (e) \geq seven hits *in-time* in the θ SL,
- (f) \leq 11 hits *in-time* in the θ SL,
- (g) two fitted tracks in the φ view, each one with at least four points, and $\chi^2/\text{d.o.f.} < 10$,
- (h) two fitted tracks in the θ view, both with $\chi^2/\text{d.o.f.} < 5$,
- (i) No HTRGs in the θ view, out of the correct time slot.

Cuts (a), (b) and (c) are also employed in the single muon analysis, although in that case the *in-time* hits are defined in the time window $\sim -400 \text{ ns} \leq t_0 \leq \sim 800 \text{ ns}$. The latter definition is more correct from the trigger point of view as hits occurring in this time range can affect the trigger performance. On the other hand it is mandatory to select a clean sample of simultaneous di-muons. Thus a looser time window for *in-time* hits, although more correct in principle, would select

too many events with two muons in different slots and therefore it is not used in this case.

Cuts opposite to (d) and (e) are used in the single muon analysis to reject di-muon events. In this case they are used to enrich the sample of such events.

The effect of cut (f) is to reduce the number of events with splashes on the chamber. The other cuts are used to improve the selection of events with two muons coincident in time and thus belonging to the same bunch crossing as discussed below. The χ^2 was computed with a fixed single hit resolution $\sigma = 290 \mu\text{m}$ without including the observed angular dependence of Fig. 14.

The Mean-Time computed by using the drift times of the hits associated to a fitted track can be used to test whether the muon pairs belong to the same bunch. We use the hits associated to each track fitted in the φ SLs separately to compute all the possible Mean-Times. Their average is called $\langle \text{MT} \rangle$. The value of $\langle \text{MT} \rangle$ corresponds, in the ideal

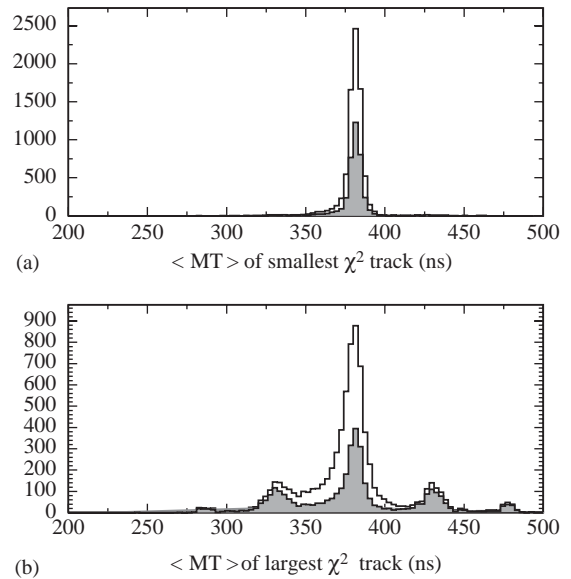


Fig. 16. The distribution of the average Mean-Time (MT) computed using the hits associated to a track in the φ view (white histogram), superimposed to the distribution of the same quantity for events which also have a HTRG in the θ view at a time slot different from the correct one. Plot (a) refers to the selected track with the smallest χ^2 and plot (b) to the selected track with the largest χ^2 . The peaks due to off-time muons are almost entirely accounted by these cases.

case, to the maximum drift time, namely ~ 380 ns. The track with lower $\chi^2/\text{d.o.f.}$ is more likely to be at the correct bunch crossing and contributes to the histogram in Fig. 16a while the other track contributes to Fig. 16b and is more likely to be at a wrong bunch crossing. Fig. 16 shows a large quantity of events clustering around the expected value, but there are also clusters separated from the main one by multiples of 50 ns, mainly in Fig. 16b. This shift corresponds to twice the bunch crossing time separation and comes from muons at wrong bunch crossing. As all events are assumed to occur at the correct bunch crossing, for off-time muons which belong to a bunch different from the correct one, the t_0 is subtracted with a wrong offset of 25 ns or a multiple of it. Due to the expression for the Mean-Time reported in Section 2.2.1 they contribute to Fig. 16 with a value of $\langle \text{MT} \rangle$ shifted by multiples of twice the bunch crossing separation. Such tracks generally have a poor fit quality corresponding to high values of $\chi^2/\text{d.o.f.}$. Therefore the cut $\chi^2/\text{d.o.f.} < 10$ is on one hand very efficient in rejecting these off-time muons but it is still quite inclusive for tracks occurring at the correct time.

The effect of cut (i) is also apparent in Fig. 16b. When a track crosses the muon station at a given bunch crossing the trigger in the θ view delivers a HTRG track segment in about 87% of all cases. Therefore θ HTRGs at a given bunch crossing can be used to tag the existence of tracks crossing the muon station at that bunch crossing. The shaded histogram shows the quantity $\langle \text{MT} \rangle$ for tracks in events in which there was a HTRG in the θ view at wrong bunch crossing. When such events are rejected essentially all off-time muons are removed from the data sample.

The event selection is rather loose concerning the information from the ϕ view and more strict on the θ view. Although there is obviously a correlation between ϕ and θ quantities, they are actually rather decoupled except for muons which undergo a large electromagnetic shower through the whole chamber. Their fraction is negligible since there is no material in front of the detector. This guarantees that the event selection does not introduce a strong bias on the trigger performance in the ϕ view which we want to study.

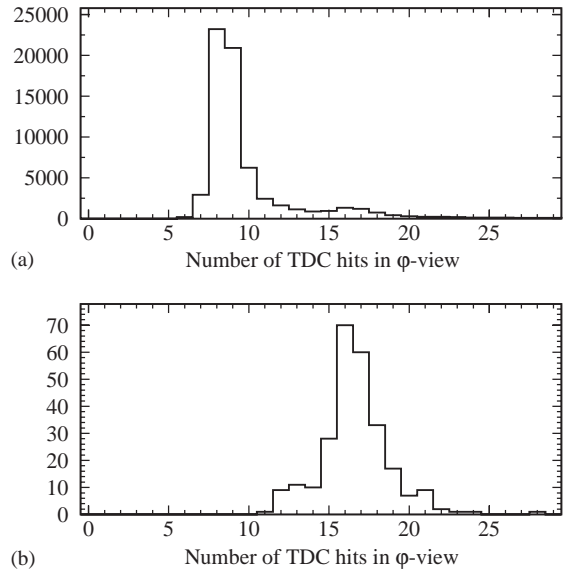


Fig. 17. The distribution of the total number of hits per event in the ϕ view, in the correct time window and before the di-muon event selection, is shown in plot (a). The main peak, placed at eight, corresponds to single muons which cross both ϕ SLs. The peak due to di-muons is enhanced after the event selection (b).

Fig. 17 shows the number of *in-time* hits in the ϕ view before and after the event selection. After the selection cuts roughly 0.4% of the events are tagged as di-muon candidates suitable for analysis.

5.3. Collected data

The goal of the tests was to provide complete information about the general performance of the trigger algorithms. The first set of data was collected with the default configuration setup at various angles of incidence of the beam. This setup is the one currently proposed for the data taking at LHC and it is the basis of all the predictions of the trigger performance. Its main special features are: BTI LTS mechanism enabled; BTI time-independent equations disabled; BTI redundant patterns disabled; uncorrelated LTRGs validated from a θ trigger of any type; preference to HTRGs in TRACO and TS sorting; *carry* enabled in TS; *ghost1* and *ghost2* rejection enabled in TS.

Data were collected in several other configurations changing one by one the default values of the

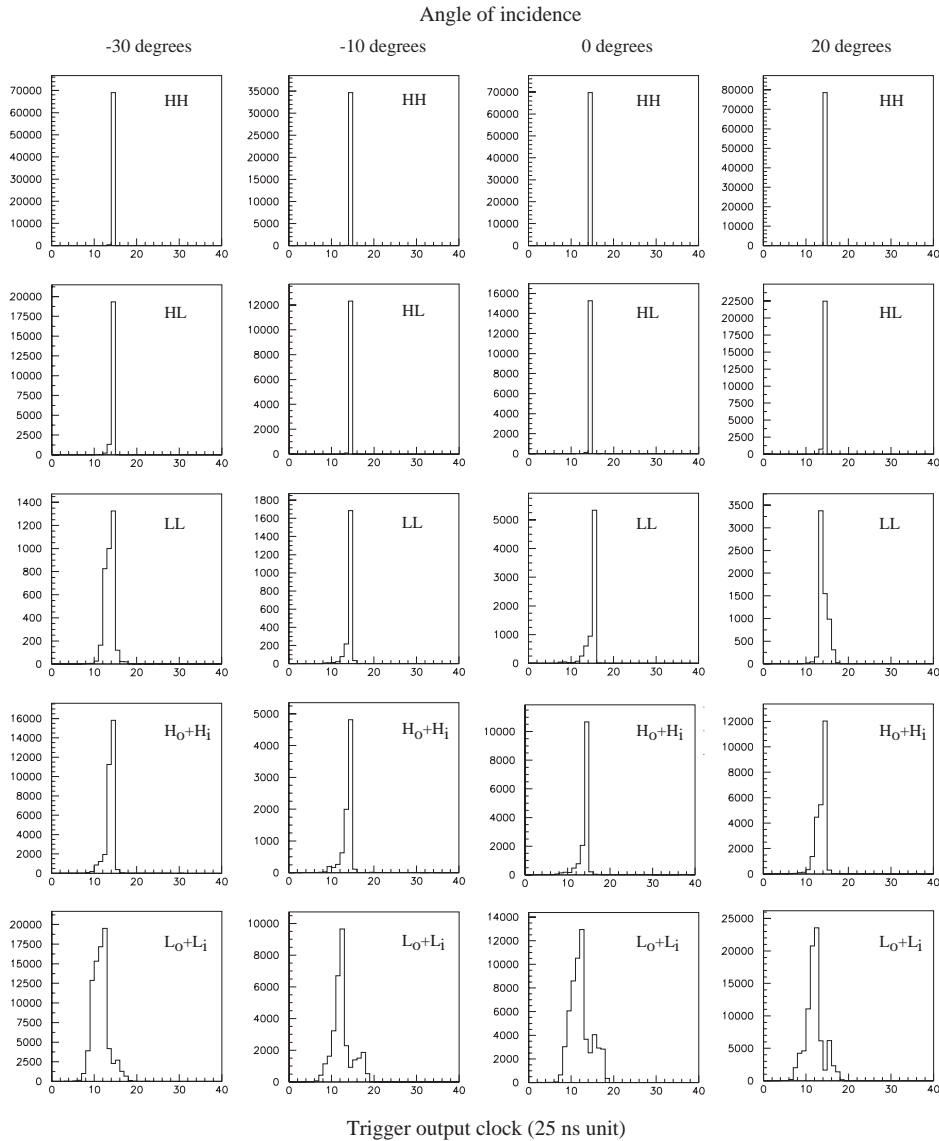


Fig. 18. Distribution of the trigger timing recorded by the Pattern Unit for different trigger types at several angles of incidence.

most important parameters of the devices constituting the trigger system. These parameters should modify the trigger response, i.e. change its efficiency and its noise rate. Data were collected in the following relevant configurations (apart from the default one):

1. Single LTRGs filtered asking for a HTRG in θ view (TS parameter acting on TRACO),
2. Single LTRGs always rejected (TRACO parameter),
3. BTI LTS algorithm disabled in both views,
4. BTI LTS algorithm disabled in θ view only,
5. TRACO LTS algorithm enabled,
6. Carry mechanism disabled on TS system,
7. TSM recovery of H_i triggers enabled,
8. TS *ghost2* rejection disabled,
9. TS *ghost1* and *ghost2* rejection disabled.

In each configuration only the listed parameter was changed while all others had their default value. Most of the parameters (cases 1, 2, 3, 4, 5, 8, 9) were introduced to provide means of control and of reduction of the background of false triggers. They are relevant for the single muon trigger performance. In contrast some of them are related to the correct muon selection (cases 6 and 7) and are relevant for the di-muon trigger efficiency.

Moreover the TSM device can be programmed to follow a default mode or a backup mode. Since this is one of the most important features of the whole system special runs were taken to test the performance of the trigger in the backup mode.

6. Results on single muons

6.1. Default configuration

6.1.1. Timing response

The trigger system is designed to give a late trigger upon detection of the *alignment* of the hits recorded in the drift tubes. The requirement of a minimum of three *aligned* hits causes many false *alignments* and therefore the correct trigger is accompanied by several ghost triggers issued at different bunch crossings. The spectra obtained for few cases are reported in Fig. 18 where the trigger timing recorded by the Pattern Unit is shown for some relevant angles of incidence (normal, close to maximum, average positive and average negative angles). Plots for different trigger types defined by the TRACO quality codes are shown. The cleanest trigger types are HH and HL triggers. Their time distribution allows the determination of output slot 14 as the correct one (remember that this clock cycle is relative to the startup time of the Pattern Unit defined by the L1A signal). Thus triggers recorded in other slots of the Pattern Unit are classified as ghosts. Most of the false triggers appear in the uncorrelated categories with a non negligible fraction of correlated LL triggers.

Within each chamber the trigger system provides up to two track candidates per clock cycle in order to allow multimMuon detection and to provide

a recovery mechanism for the case that the best choice selects a ghost track. The second track is marked and transmitted one clock cycle after the

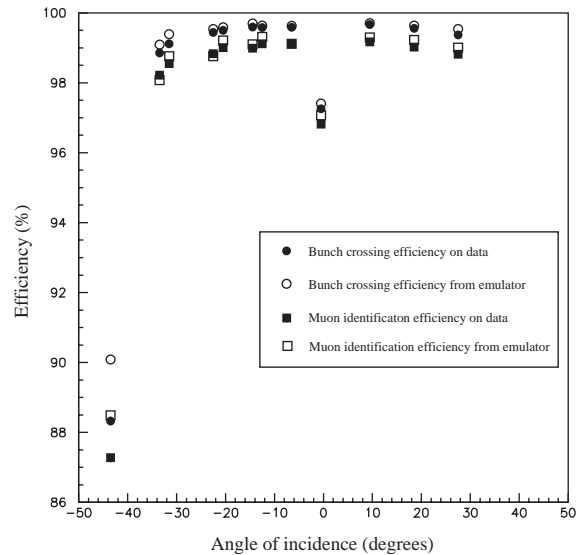


Fig. 19. Bunch crossing efficiency as a function of the angle of incidence. The efficiency from recorded data is compared with the efficiency computed by the software emulator. The bunch crossing efficiency is also compared with the muon identification efficiency defined at the end of Section 6.1.4.




Fig. 20. Contributions to the bunch crossing efficiency of correlated (HH, HL and LL) and uncorrelated (H_0 , H_i , L_0 and L_i) triggers.

first track if no other first track was found. The time spectra for first and second tracks at normal incidence are similar but, as expected, the correct output clock step of second tracks is 15.

A quantitative discussion of ghost triggers is done in Section 7.2 when checking the trigger performance of alternative configurations.

6.1.2. Efficiency

The overall efficiency is measured counting triggers occurring at the correct clock step. The efficiency computed in this way is called bunch crossing efficiency since no attempt is made to check if the output parameters correspond to the actual muon track parameters. The bunch crossing efficiency is shown in Fig. 19 as a function of the angle and it is compared to the results of the emulation algorithm.

The overall trigger efficiency is good since it is almost flat in the range of 98–99% and starts to drop at $\sim 35^\circ$. The 3% drop at normal incidence is due to tracks hitting two I-profiles in both ϕ SLs

so that no hits occur in the drift tubes separated by these profiles. Below 35° the emulator gives the correct answer within 0.1%. The efficiency drop at the largest angle is not harmful since it corresponds to very low transverse momentum muons ($< 5 \text{ GeV}/c$) hitting the chamber close to its edges.

A deeper insight into the algorithm performance can be obtained by checking the fraction of trigger types as a function of the angle. The contribution of each trigger type to the efficiency is shown in Fig. 20 and in Fig. 21. There is a progressive reduction of correlated triggers with increasing inclination, with a fast drop at very large angles where most of the triggers are uncorrelated.

The bunch crossing efficiency is plotted as a function of the fitted position (intercept at the chamber centre) in Fig. 22. The inefficiency at normal incidence has a regular pattern and is located at the position of the drift tube cathodes. The inefficiency is randomly scattered in the other cases but at 45° angle of incidence where a complicated structure is found. This result shows

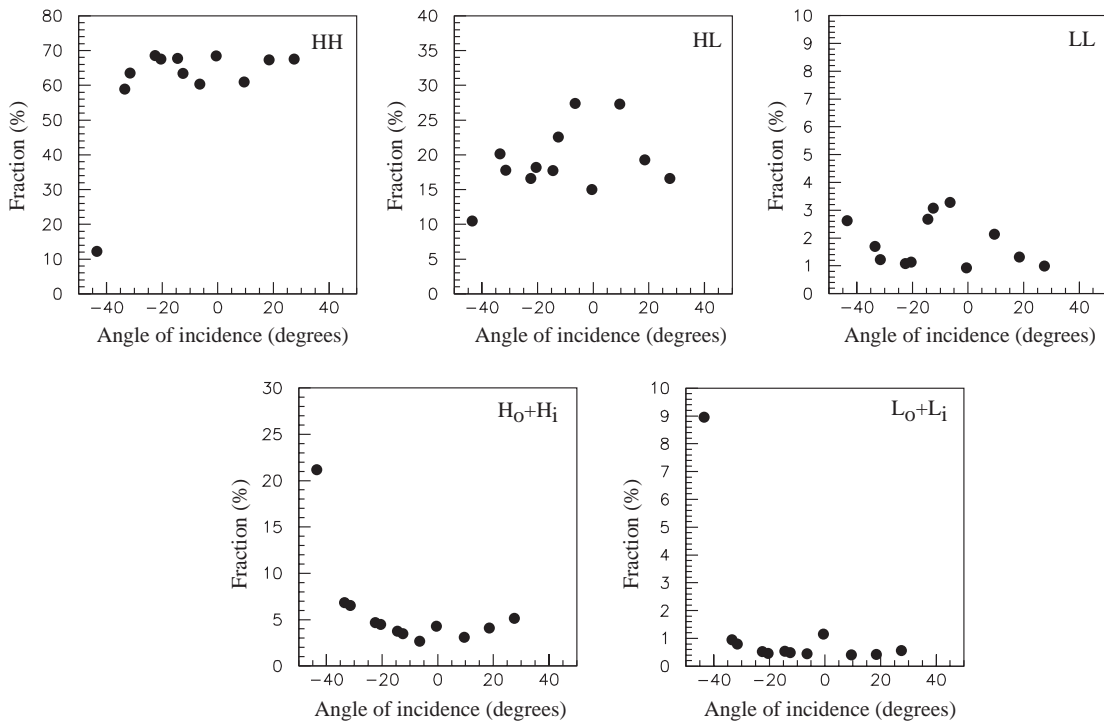


Fig. 21. Relative contributions to bunch crossing efficiency of different trigger types.

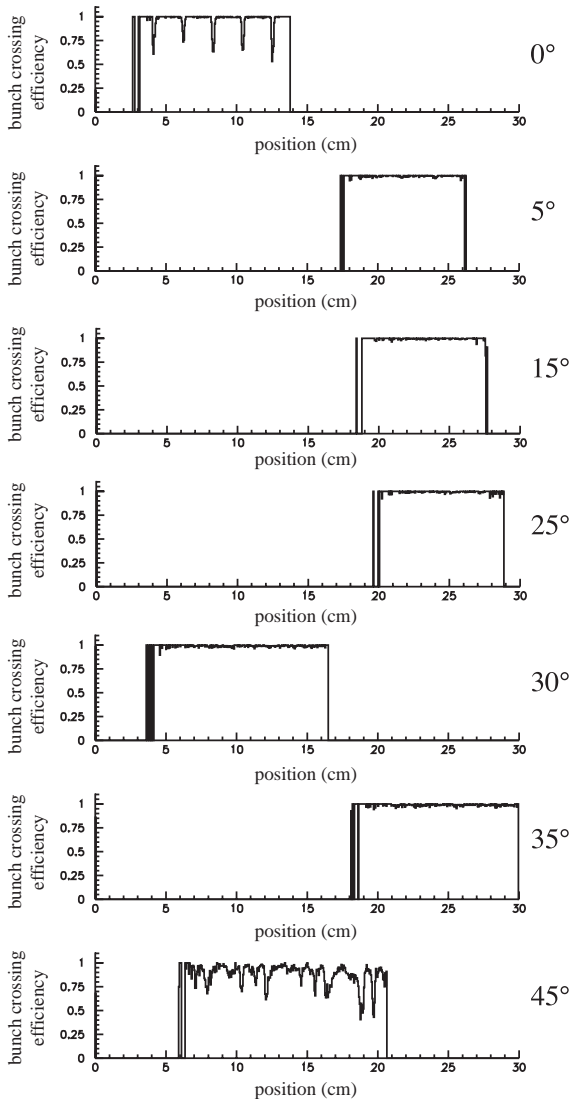


Fig. 22. Bunch crossing efficiency as a function of the position of the track intercept at the central plane of the chamber at different angles of incidence. The beam hits different cells because the chamber was moved several times.

that the Mean-Time algorithm is really effective and that the loss of linearity in the space–time relationship is important only at very large angles of incidence.

6.1.3. Correlations

The trigger output parameters can be cross-checked with the results of the track fitting done

using the drift times recorded by the TDCs. The comparison is easier if we transform the TRACO radial angle and the fitted position inside the chamber in the same units and reference frame (in fact using a polar coordinate system) because in this case a diagonal line is expected in a corresponding correlation plot. After transformation the correlation between the radial angle measured from the fit and the radial angle calculated by the trigger algorithms for HH quality triggers is shown in Fig. 23 where data from all files recorded with the default configuration at any angle are superposed: the correlation is rather impressive. Looking at the other data no dependence on the trigger quality is observed.

The angle of incidence can be computed using radial and bending angle. It can be compared to the same angle measured from the inclination of the fitted track. This comparison is done in the plot of Fig. 24 showing a nice correlation between the two quantities.

A deeper analysis of the single muon sample provides a comprehensive control of the correct internal performance of the trigger algorithm. Any event can be classified depending on the quality of the fitted track in the φ SLs. In particular we selected a sample of events where four hits are found in each SL. Then we compared the output trigger quality with the χ^2 of the fitted track. In the φ view we expect that HH triggers correspond to low χ^2 fits while a lower trigger quality should be found when the fit quality for one or both segments is worse. Fig. 25a–c shows for some angles of incidence the χ^2 distributions of the track segment with the worst fit quality out of the two available segments in the φ SLs (distinguished between events associated to HH triggers and events associated to a lower trigger quality). As already remarked the angular dependence of the spatial resolution (c.f. Fig. 14) was not included in the χ^2 computation, causing shape variations with angle. The procedure was repeated for the θ view where only one track segment is available: the χ^2 of the track fit in the θ SL is shown in Fig. 25d, separately for BTI H and L triggers. There is a clear correlation between the trigger quality and the χ^2 value: higher quality triggers are

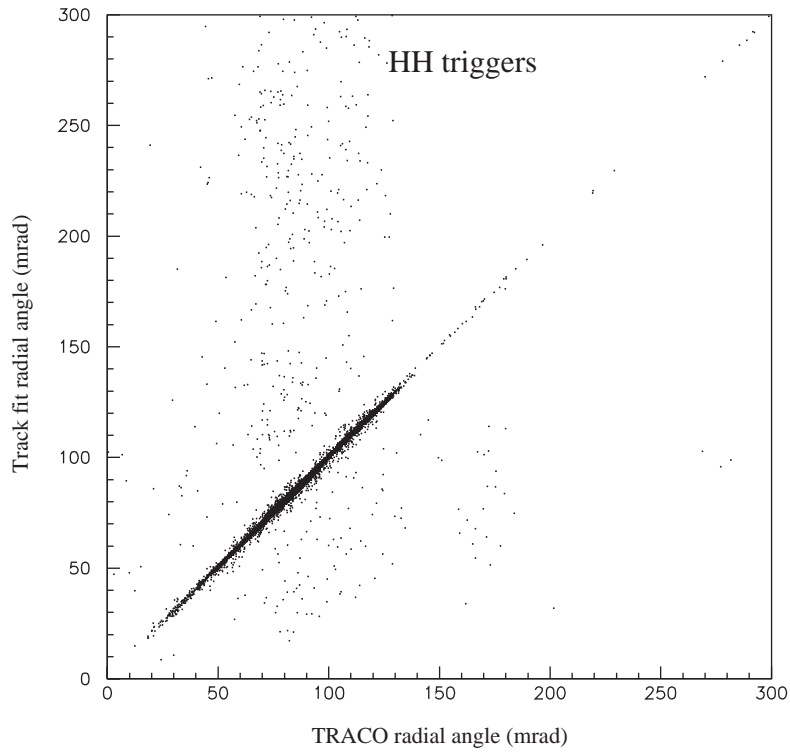


Fig. 23. Correlation between the radial angle output by the TRACO and the radial angle computed by the track fit for HH triggers.

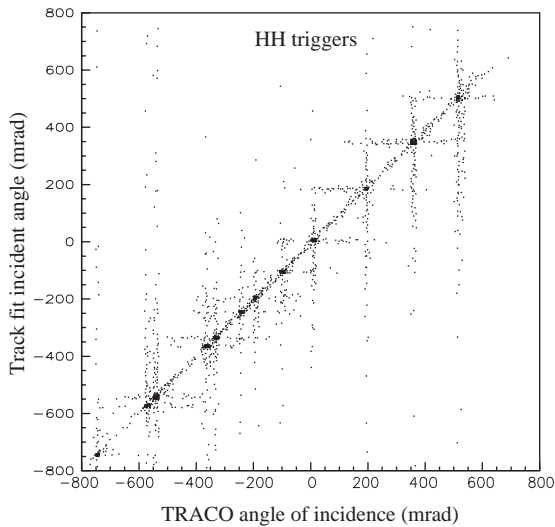


Fig. 24. Correlation for HH trigger type between the angle of incidence computed using the TRACO bending angle and the angle of incidence computed from the track fit.

concentrated at low χ^2 while lower quality triggers dominate the high χ^2 region.

We have defined several subsamples according to the expected quality of the triggers delivered from the system. The four-hit tracks with $\chi^2 > 50$ were downgraded to three-hit segments refitted after rejection of the worst hit. The selected subsamples were:

- Sample 4+4: a four-hit fit in both φ SLs,
- Sample 4+3: a four-hit fit in one φ SLs and a three-hit fit in the other one,
- Sample 3+3: a three-hit fit in both φ SLs.

The 4+4 sample is expected to provide mainly HH triggers. The 4+3 sample included downgraded three-hit fits and true three-hit fits and should provide mainly HL triggers. The 3+3 sample should be dominated by LL triggers. The results collected in Table 1 show that this is indeed

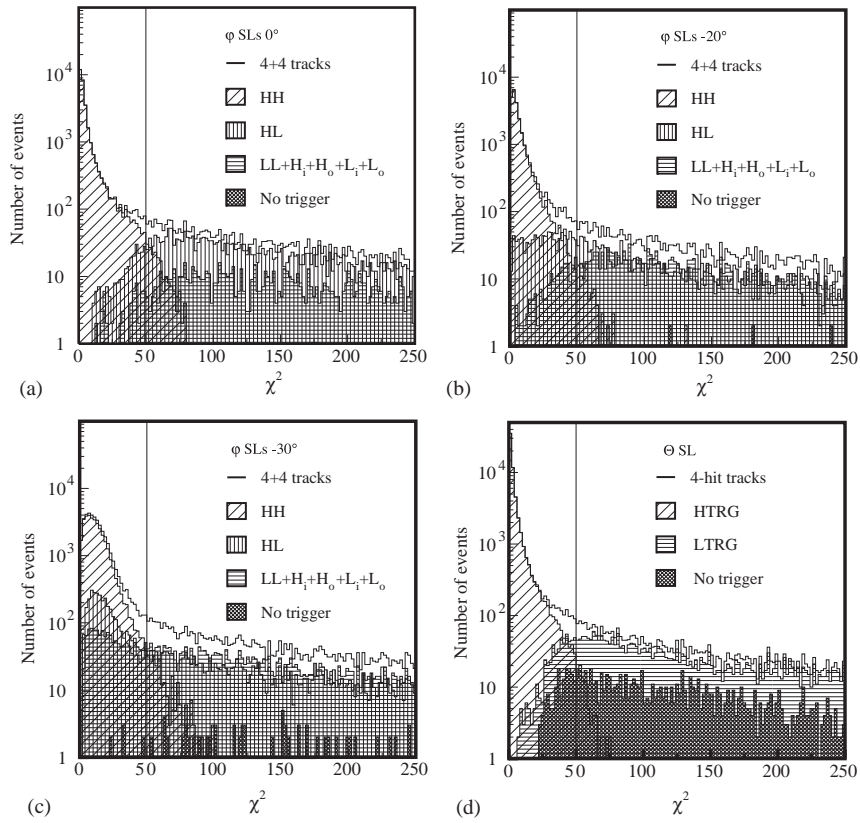


Fig. 25. Distribution of the worst χ^2 for tracks having two fitted four-hit track segments in the ϕ view (a–c) and a fitted four-hit track segment in the θ view (d). The vertical line at $\chi^2 = 50$ shows the cut applied for track refitting to define the muon subsamples.

the case. Sometimes the trigger delivers a better than expected quality, due to its built-in tolerances which are larger than the tolerances accepted for the track fits. On the other hand the fits in the lower quality samples are sometimes better since inside the TRACOs the correlation between the two available track segments sometimes fails.

6.1.4. Resolutions

The determination of the resolution of the track parameters of the trigger algorithm is complicated by the quantized nature of the output. The internal calculations done by the BTIs and TRACOs should provide the same resolution on the radial angle for all kind of triggers while the resolution on the angle of incidence should be rather different for correlated and uncorrelated triggers.

This structural difference has its origin in the actual device which provides the result of the computation. The radial angle's nominal least count is ~ 0.25 mrad whereas the nominal least count of the angle of incidence is ~ 2 mrad. In the case of a correlated trigger the calculation is done by the TRACO that works at the full resolution while for uncorrelated triggers the calculation is done by the BTI that has a lower resolution. There is also a slight difference between uncorrelated H and uncorrelated L triggers due to the different method used to compute the parameters. As a construction choice the position resolution is kept constant in the correlator. Therefore we expect the resolution of radial angles to be independent of the trigger type. Instead the track k -parameter calculation is improved by the TRACO causing a difference between uncorrelated and correlated

Table 1
Fraction (%) of trigger quality types recorded for different muon samples

Quality		Angle of incidence			
		0°	10°	20°	30°
Sample 4 + 4	HH	99.0	96.8	95.0	87.5
	HL	0.8	2.6	4.2	8.8
	LL	0.0	0.0	0.0	0.2
	H _o + H _i	0.2	0.4	0.8	3.3
	L _o + L _i	0.0	0.0	0.0	0.2
	No trigger	0.0	0.0	0.0	0.0
Sample 4 + 3	HH	2.6	1.0	1.5	2.7
	HL	72.6	74.0	60.6	48.5
	LL	0.4	2.2	1.2	2.3
	H _o + H _i	24.3	22.3	35.5	43.2
	L _o + L _i	0.2	0.5	1.3	3.3
	No trigger	0.1	0.0	0.2	1.3
Sample 3 + 3	HH	0.6	0.4	0.5	0.0
	HL	3.9	1.7	4.6	4.3
	LL	50.8	70.6	45.6	31.8
	H _o + H _i	1.2	0.6	2.1	1.9
	L _o + L _i	43.6	26.6	47.1	62.0
	No trigger	9.5	7.3	13.3	21.8

The error varies depending on statistics of the sample and is about 0.1% in the best case. See text for other details.

triggers. Furthermore the least count of the k -parameter is not constant in angle: at normal incidence it corresponds to ~ 50 mrad for the BTI (i.e. for uncorrelated triggers) and ~ 6 mrad for the TRACO (i.e. for correlated triggers).

We measured the resolution of the track parameters obtained from the trigger by comparing them to the results of the track fitting. Fig. 26 shows the difference in the radial angle between trigger output and fit results for different trigger types. Table 2 reports the variance of a gaussian fit to these and further plots. The radial angle resolution is as expected almost constant for any trigger type (in fact slightly worse for single LTRGs) and compatible with the least count of the parameters. Taking into account the distance from the chamber to the interaction point (i.e. the scintillator in the test beam setup) we can estimate the resolution in the coordinate measured by the drift chamber. The plots in Fig. 27 show an r.m.s.

around $700 \mu\text{m}$ not too far from the chamber resolution ($\sim 100 \mu\text{m}$). This is an unusual situation for a first level trigger, which normally has resolutions that are orders of magnitude worse than the detector resolution. It is infact even more impressive, since the BTI is actually computing its equations with a single hit resolution of about $700 \mu\text{m}$ to be compared with the resolutions plotted in Fig. 14.

The same simple situation does not apply to the angle of incidence. In this case the expected resolution is rather different for correlated and uncorrelated triggers. But this is not important since in the DTTF the muon transverse momentum is evaluated from the position, while the bending angle is used only for extrapolation between chambers. As an example Fig. 28 shows the resolution of the angle of incidence for some angles and trigger types. The variance of the correlated trigger distributions is about 3 mrad. The variance of the uncorrelated trigger distributions is close to 30 mrad for H triggers and 40 mrad for L triggers. There is no strong dependence on the angle of incidence.

We can apply a cut on the difference in the radial angle to get the actual efficiency of identifying a muon correctly. A so called muon identification efficiency can be evaluated applying a window of ± 4 mrad in the radial angle difference. We did not try to cut also on the angle of incidence because of the large difference in resolution among trigger types that we have shown before. The muon identification efficiency is compared to bunch crossing identification efficiency in Fig. 19. It is rather close to the latter and fairly constant with angle.

6.1.5. Trigger noise

The drift tubes local trigger produces false triggers due to the intrinsic behaviour of the algorithm (temporal ghosts) and to the superposition between adjacent devices (spatial ghosts) as already described in Section 3.

A temporal ghost is generated at wrong bunch crossings and is related either to the allowed *alignment* tolerances or to *alignments* associated to the left–right ambiguity. Several temporal ghosts can be created from the same true muon hits. The

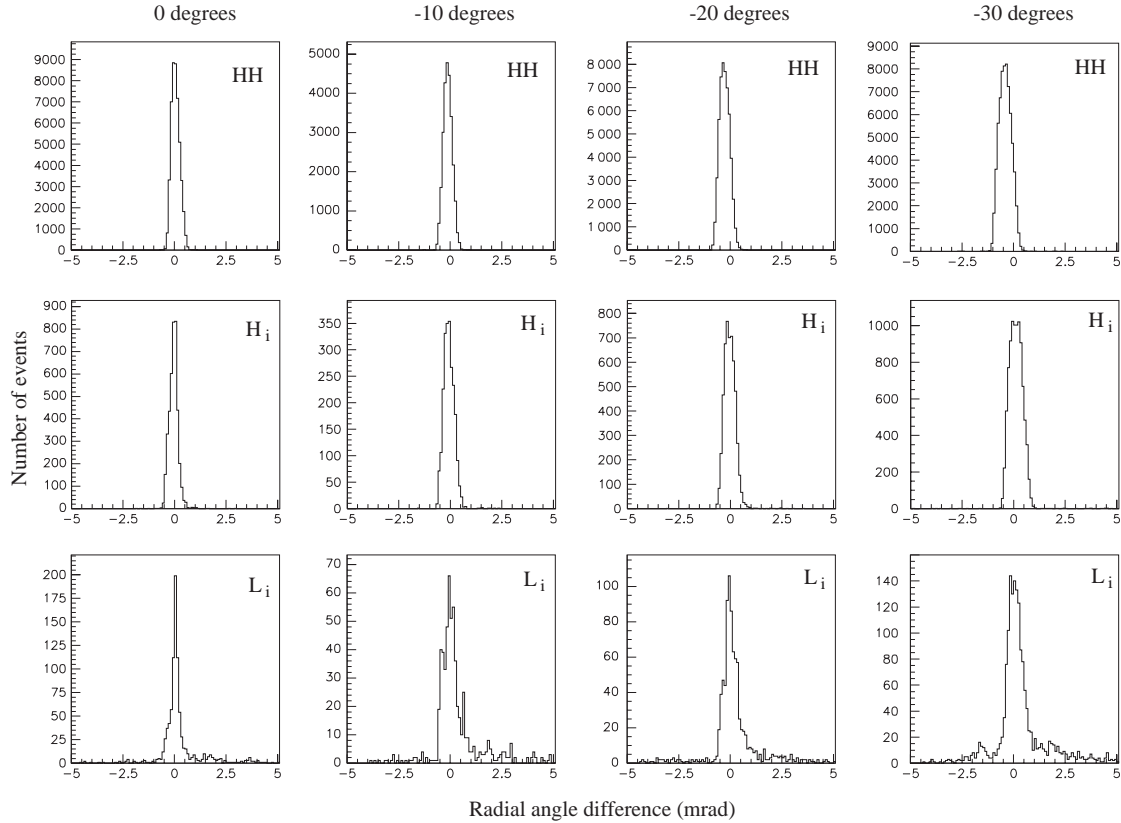


Fig. 26. Difference between the radial angle as measured by the track fit and as output by the trigger at the correct clock cycle for different angles of incidence and trigger types.

Table 2

Variance of a gaussian fit to the difference in radial angle computed from fit and hardware trigger for various angles and different trigger types

Angle(degrees)	-42	-30	-20	-10	0	10	20	30
HH resolution (μrad)	226	281	216	214	191	258	258	257
Hi resolution (μrad)	321	303	256	245	204	274	279	304
Li resolution (μrad)	507	372	369	314	282	280	376	394

fraction of events with more than one trigger is shown in Fig. 29: the average fraction of wrong triggers increases with the angle. This is consistent with the availability of more BTI patterns for alignment.

A spatial ghost is either a copy of a good trigger associated to the overlap of TRACOs or a lower

quality trigger due to the overlap of BTIs. These ghosts are output as second tracks in the same bunch crossing as the good trigger. The output of two trigger candidates per chamber per bunch crossing is a mandatory choice needed to allow di-muon detection within the same chamber. In the single muon selection we expect all second tracks to be noise triggers.

The fraction of events with two triggers at the same bunch crossing, shown in Fig. 30, is rather small and it is again a function of the angle. Almost all second tracks are uncorrelated L trigger types. We observed a correlation between the probability of generating temporal ghost tracks and the track position. Fig. 31 shows the positions where ghost tracks are found for normal incidence and 30° incidence. A clear peak structure is found. Each peak corresponds to a border between cells

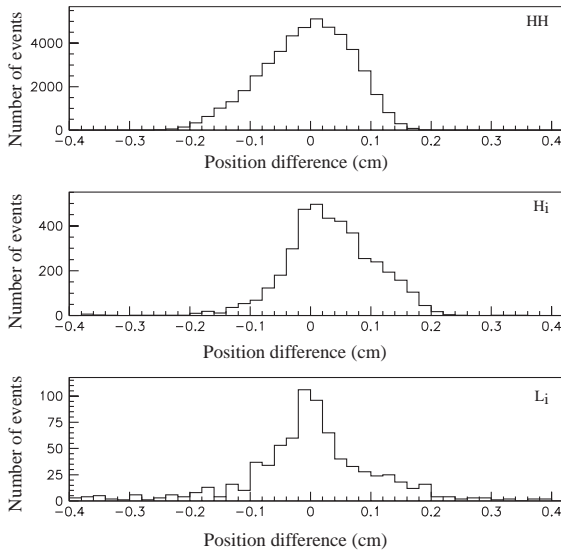


Fig. 27. Trigger position resolution for some trigger types at normal incidence. The plots are obtained by conversion of the radial angle as measured by the TRACO and correspond to the histograms of Fig. 26. Therefore the resolution at other angles of incidence is very similar.

crossed by the muon track. These triggers are a LTRG duplication of the HTRG trigger due to BTI overlap as described in Fig. 11.

We show in Fig. 32 two different ways of quoting the trigger noise: Fig. 32a reports the number of ghost triggers over the total number of triggers, while Fig. 32b shows the fraction of events with at least one ghost trigger. The difference of course reflects the fact that more than one ghost trigger can be found in an event. The first definition is related to the increase of the trigger rate while the second one shows only how frequently the problem happens. The most relevant plot is probably the first one since the DTF treats each ghost trigger as an independent track and tries to build a trigger primitive connecting ghost tracks in different chambers. One should remember that most of these ghosts are uncorrelated L triggers and the DTF algorithm uses them only as targets and never as origins of tracks when linking track segments in different stations. This fact, together with a generally poor position and direction resolution of the ghosts, ensures that most of them will not contribute to fake track

reconstruction. In the CMS detector the distribution of the angle of incidence is basically flat below 20° , with tails extending to higher angles. The design of the trigger primitive generator for the CMS first level muon trigger required the rate of ghost triggers to be at most equal to the true muon rate. This requirement is clearly met by the found trigger performance.

Looking at the out-of-time triggers there is a particular sample of events where the trigger is output only at the wrong bunch crossing. Indeed we have verified that almost all inefficiencies shown in Fig. 19 are identified as muons triggered out of time. The possible reasons for this behaviour are: a penetrating δ -ray very close to the real position of the muon; a local distortion of the electric field or of the space–time relationship; an unusual fluctuation of the primary electron generation. Except for the highest angle of incidence where probably the major contribution comes from deviations from linearity of the space–time relationship, a mixture of these problems is the actual cause. This conclusion is also supported by the distributions of out-of-time triggers shown in Fig. 33, separately for the different trigger types. Apart from observing that the bulk of wrong triggers is of single L type, we see that the higher quality wrong trigger fraction increases with the angle of incidence as expected from deviations from linearity of the space–time relationship.³ The single L triggers are a less dangerous background since the DTF uses this kind of triggers only as target and not as track origin when connecting the different stations.

6.2. Alternative configurations

As already mentioned several configurations of the trigger algorithm, listed in Section 5.3, were tested. Changing the configuration enables alternative algorithm flows that were provided as handles to cope with unexpected background or noise levels. The test of the alternative flows was intended to validate the design choices and to see

³This is not true at normal incidence where the small bump in L triggers is correlated with the smaller probability of four hits due to the presence of consecutive drift tube cathodes.

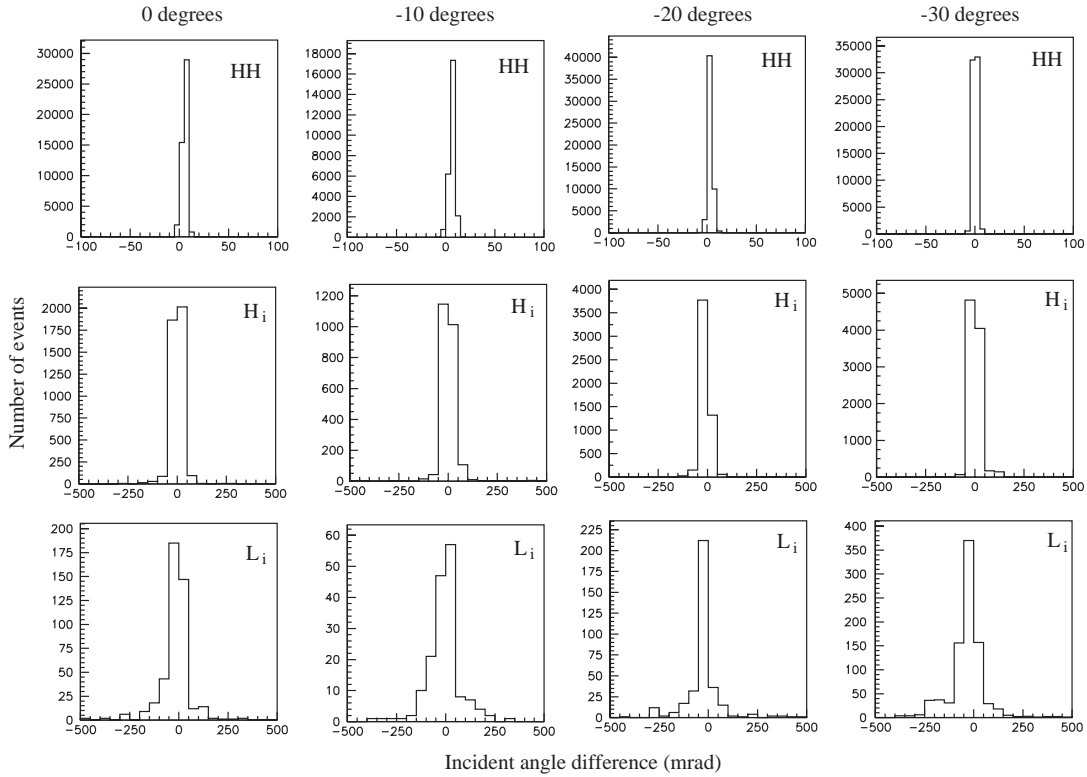


Fig. 28. Difference between the angles of incidence as measured by the track fit and as output by the trigger at the correct clock cycle for different angles of incidence and trigger types.

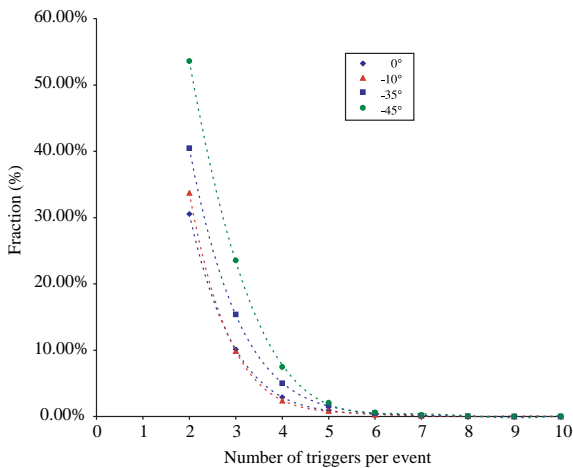


Fig. 29. Fraction of events with more than one trigger for some angles of incidence. The result is symmetric in angle. The events without wrong triggers are not included for sake of clarity.

how far the noise suppression is extended when using a non-default algorithm flow. All the relevant configurations are explained in the corresponding figure captions. A few of them (#6 and #7) were not relevant for the single muon selection. The important quantities to be checked are efficiency and noise (both spatial and temporal).

We grouped the alternative flows depending on their actions. Configuration #6 changes the way the selection of the second track is done by the TS disabling the *Carry* mechanism. Configuration #7 recovers H_i triggers marked as second tracks if the first track of the current event is an uncorrelated L trigger. These features were introduced in order to make a correct selection of di-muons and should have no effect on single muons (apart from modifying the actual parameters of ghost tracks)

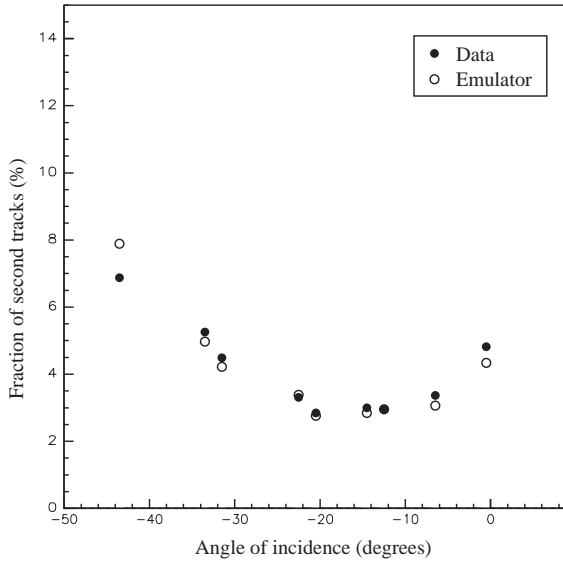


Fig. 30. Fraction of second tracks recorded at the same bunch crossing in the trigger data compared to the emulator response.




Fig. 31. Muon incident position at $SL \phi_1$ for events having a spatial ghost trigger.

but in the case of important electromagnetic background. Since in this test there was no material in front of the chamber this source of background is negligible and no effect should be seen. We therefore ignored them once checked that the change of the default values had no effect.

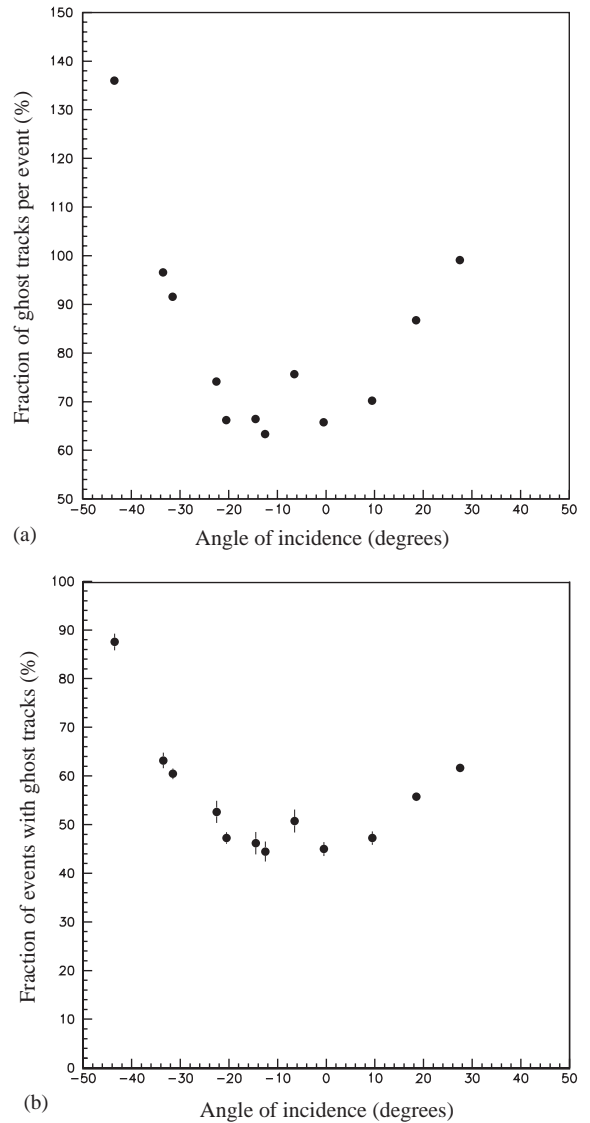


Fig. 32. Fraction of out-of-time ghost tracks as a function of the angle of incidence. Plot (a) reports the fraction of ghosts per good trigger, while plot (b) reports the fraction of events with at least one ghost trigger.

Firstly we consider possible ways of validating the uncorrelated L triggers. In the default configuration these triggers are validated if any type of BTI trigger (LTRG or HTRG) is found in any BTI of the θ view. We tested only configurations that could reduce the background, namely the possibility of accepting uncorrelated L triggers only if

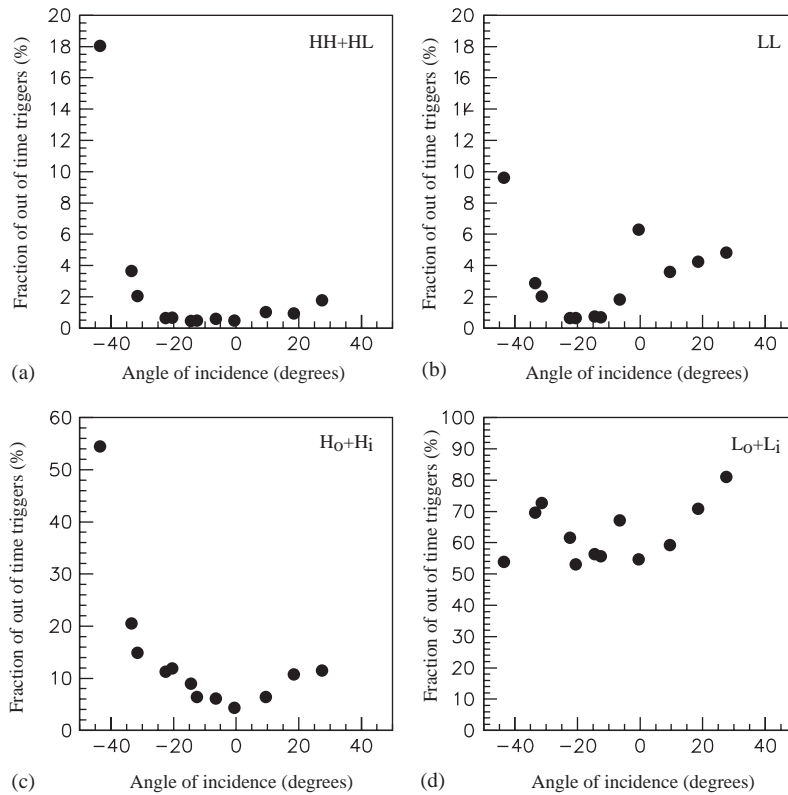


Fig. 33. Relative contributions of each trigger type to the out-of-time ghost rate.

there was a HTRG in any BTI of the θ view and the rejection of any uncorrelated L trigger.

The relevant plots are shown in Fig. 34. The first choice reduces the background significantly almost preserving the efficiency while the second one actually does not have a large impact on ghost rate but reduces the efficiency. The impact on second tracks is marginal.

The second group of configurations considers different application of the available LTS algorithms. Results are plotted in Fig. 35. It is clear that the BTI LTS algorithm must be applied otherwise the ghost rate would be unacceptably high. If needed the TRACO LTS could be applied since a tiny efficiency reduction is associated with a good noise suppression.

The last group modifies the TS ghost suppression algorithm. Fig. 36 reports the effect of the available options. Disabling spatial ghost suppression should modify the fraction of second track

ghosts. This is indeed the case: many false di-muons appear in the relevant plots. Therefore ghost suppression inside the TS system is mandatory.

7. Results on di-muons

7.1. Efficiency and purity

We define the *di-muon trigger efficiency* as the fraction of identified di-muon events which have two triggers at the correct bunch crossing in the φ view. The efficiency is studied as a function of the distance of the two muons determined by the track fitting in the φ view.

The distribution of the muon distances of such triggers at an angle of incidence of 20° is shown in Fig. 37a, superimposed to the distribution of the two track separation for all di-muons. Most of the

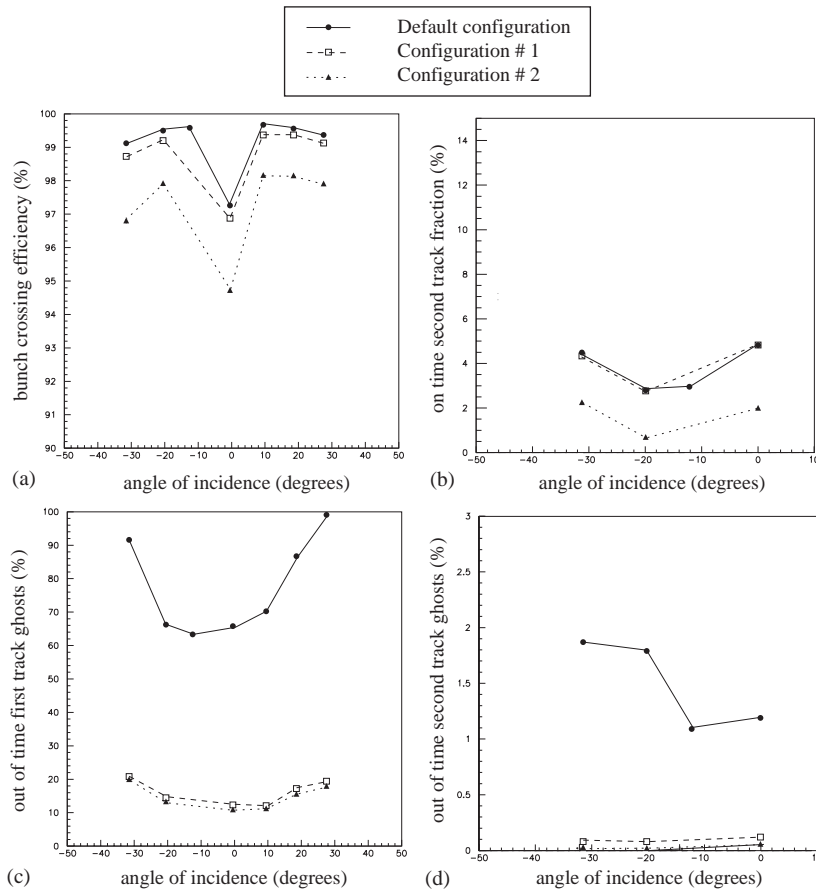


Fig. 34. Performance of the trigger for different selections of the single L triggers: default configuration means that L triggers are accepted if validated by any kind of trigger in θ view; configuration #1 means that L triggers are accepted if validated by a H trigger in θ view; configuration #2 means that L triggers are rejected. The plots show the effect on (a) bunch crossing efficiency, (b) second track ghosts at correct bunch crossing, (c) first track out-of-time noise, (d) out-of-time second track ghosts versus the angle of incidence of the track. The point at -10° for the alternative configuration is not reported due to bad data collected (no scintillator signals available).

tracks are close in space as they both belong to the beam spot, but muons from the beam halo are also present. The lack of muons closer than a few centimeters is not real but it is due to the event selection. The requirement of two fitted tracks rejects events with overlapping hits. The ratio of the two distributions is shown in Fig. 37b.

When the system does not output two triggers a single trigger is always delivered.

It is important to know whether the two triggers at the correct bunch crossing also reproduce the characteristics of di-muons in terms of position and angle. Fig. 38 shows the correlation between

the distance of the two fitted tracks in the φ view as a function of the distance of the two triggers (in count units). This is called a *purity plot* since it shows how good the correspondence between triggers and fitted tracks is. A clear correlation is observed. Events on the diagonal have trigger tracks which represent the two incoming muons properly. Ghost tracks are expected to be localized mainly in the region of very small distances between two triggers, independent of the track separation. In the data such ghosts are almost totally absent, only some scattered points are observed. After a visual inspection of the hits in

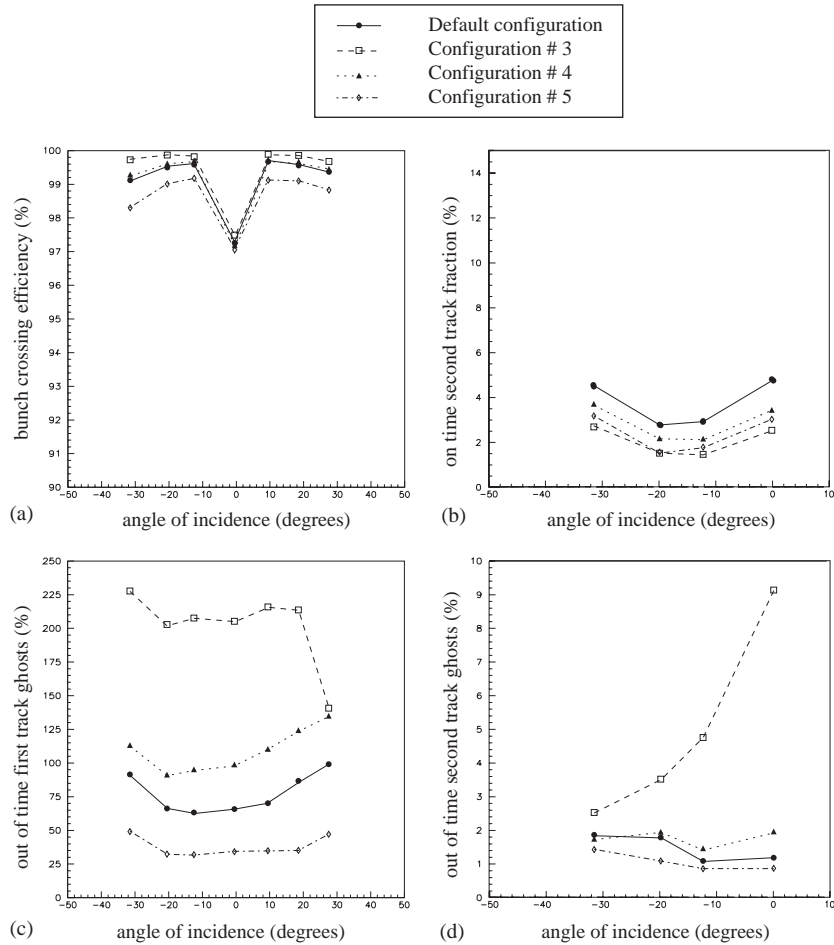


Fig. 35. Performance of the trigger for different ways of application of the LTS algorithm: default configuration means that BTI LTS is enabled in both views; configuration #3 means that BTI LTS is disabled in both views; configuration #4 means that BTI LTS is disabled only in θ view; configuration #5 means that TRACO LTS is enabled. The plots show the effect on (a) bunch crossing efficiency, (b) second track ghosts at correct bunch crossing, (c) first track out-of-time noise, (d) out-of-time second track ghosts versus the angle of incidence of the track.

the chamber they are interpreted as particles splashing in the detector which passed in some way the di-muon event selection and represent therefore failures of the di-muon offline selection algorithm.

The muon trigger candidate is reconstructed by the DTTF out of the segments delivered by the local trigger in the muon stations. The bunch crossing associated with the track is determined by the bunch crossing associated with the segments which are matched to form the track itself. Actually a track candidate can be made using also

triggers at different bunch crossings. When the segment matching is performed the DTTF opens a window of ± 1 bunch crossing centred at the time slot under consideration. If no track segments are delivered by the local trigger at a given bunch crossing in a muon station, the DTTF can use segments at the nearby bunch crossings, if they match with the segments from other muon stations. The bunch crossing of the resulting track will be determined by the bunch crossing associated with the segment with the best trigger quality.

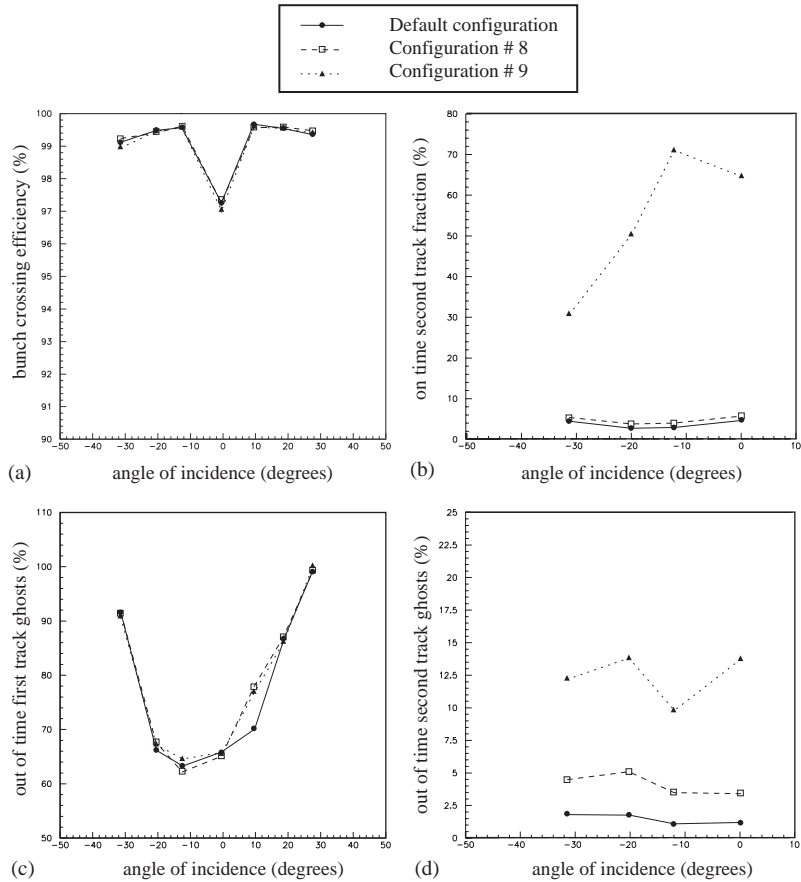


Fig. 36. Performance of the trigger for different ghost suppression levels of the TS system: the default configuration rejects ghosts of type 1 and 2; configuration #8 disables rejection of ghosts type 1; configuration #9 disables rejection of ghosts type 1 and 2. The plots show the effect on (a) bunch crossing efficiency, (b) second track ghosts at correct bunch crossing, (c) first track out-of-time noise, (d) out-of-time second track ghosts versus the track's angle of incidence.

We have therefore also determined the efficiency of finding two muon segments in a station from the DTF point of view. Segment pairs have been considered to be good pairs if both segments reproduce the position of the incoming muon tracks (if the two triggers belong to the diagonal band of Fig. 38) either in the case of two triggers being delivered at the correct bunch crossing or when one trigger is delivered at the correct bunch crossing and the other one at a nearby bunch crossing (either before or after the correct one).

Fig. 39a shows the distance of muon tracks in the ϕ view in the case of two triggers at the correct bunch crossing (dotted histogram) and at two consecutive bunch crossings (shaded histograms)

superimposed to the distribution of all muon pair candidates.

The sum of the first two distributions is then divided by the third one. The ratio is shown in Fig. 39b for the case of 20° track inclination. The result shows a higher efficiency compared to the case of Fig. 37b. This means that, even when a trigger segment related to one of the two muons is lost at the correct bunch crossing, it is still present in the nearby bunch crossing in most of cases.

In case of first and second trigger segments belonging both to the correct bunch crossing their quality is generally good with a large preference for HH triggers and a negligible amount of single L triggers. This assures that the characteristics of




Fig. 37. The distribution of the two muon track candidates separation (as determined by the fit of the TDC hits in the ϕ view of the chamber) for all the di-muon candidates (white histogram) and for events which also have two triggers at the correct bunch crossing (shaded histogram) (a). Data for 20° track inclination are shown as an example. The ratio of the two plots is the efficiency of the local trigger to deliver two triggers, both at the correct bunch crossing (b).




Fig. 38. The distance between the two muon track candidates (as determined by the fit to TDC hits in the ϕ view of the chamber) as a function of the difference between the radial angle of the two triggers, when both occur at the correct bunch crossing. The diagonal band contains triggers which represent the two muon candidates correctly. One count is angle dependent and roughly corresponds to 1.4 mm at normal incidence.

the muon tracks are well reproduced by the two triggers. In the case of two track segments belonging to nearby bunch crossings the trigger at the correct bunch crossing has a very good quality, whereas the one at the nearby bunch crossing is generally of poorer quality with about 50% probability of being a single L trigger. This might imply a worse resolution of the bending angle of such segments. Anyway it should be reminded that such cases occur only with a probability of a few percent.

The *di-muon trigger efficiency* and *purity plots* are shown in Figs. 40 and 41, respectively for different angles of incidence of the beam. The results are rather similar in all cases but at normal incidence where the efficiency is about 10% lower than at other angles.




Fig. 39. The distributions of the two muon track candidates separation (as determined by the fit of the TDC hits in the ϕ view of the chamber) for all di-muon candidates (white histogram), for the events which also have two triggers at the correct bunch crossing, both identifying the correct position of the two muons (dotted histogram), and for the events which have one trigger at the correct bunch crossing and one at the nearby bunch crossing, both providing a correct position of the two muons (shaded histogram) are shown in figure (a). The sum of the dotted and shaded histograms, divided by the white one, represents the efficiency of the local trigger to deliver two triggers, both with correct identification of the muon position (b).

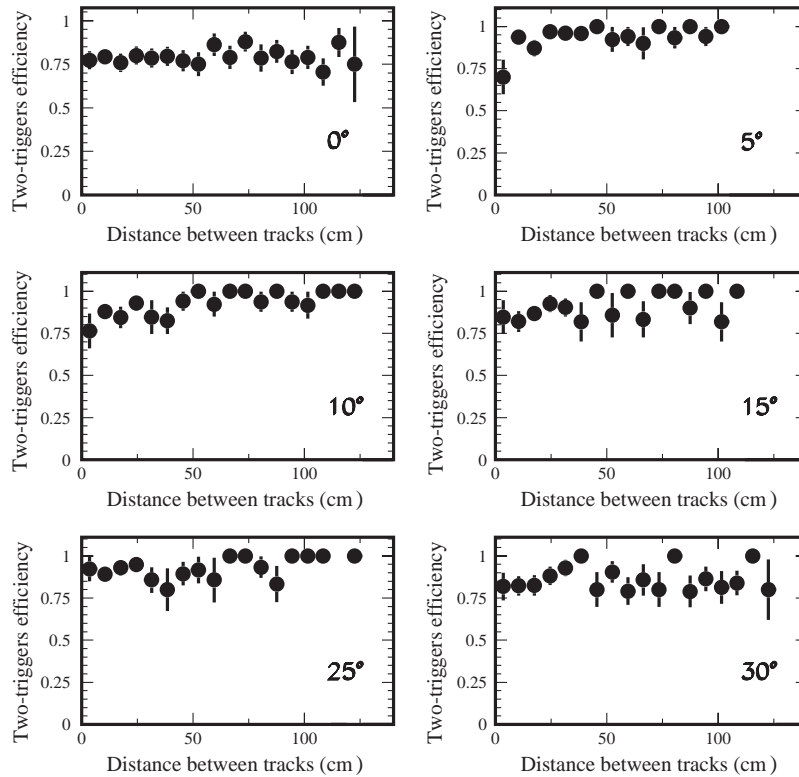


Fig. 40. Efficiency to detect two triggers at the correct bunch crossing as a function of the track separation, as determined by the fit to the TDC hits on the ϕ view, for different angles of incidence of the beam.

At angles of incidence $>0^\circ$ the efficiency is sometimes lower for tracks very close to each other (distance <20 cm) because the ghost rejection mechanism discards some of the potentially good triggers, misinterpreting them as false triggers. For larger track separations (distance >17 cm, corresponding to more than two TRACOs) the efficiency reaches a kind of plateau which is generally well above 90%.

The lower efficiency at normal incidence of the beam is due to the presence of ghost triggers at a bunch crossing following the correct one, causing the loss of one of the two triggers at the correct bunch crossing. These cases fall in the category of pile-up events. About 80% of such ghost triggers are a duplication, typically with lower quality, of the trigger at the previous bunch crossing. They are peculiar for an angle of incidence of 0° and

occur when tracks cross a chamber through the border of the drift tube cells leading to more ambiguities in the patterns detected by the BTIs. This is the same effect as the one reported in Section 6.1.5 for ghosts in single muon events. At angles different from zero this specific ghost rate becomes negligible. The remaining 20% of ghosts occurring at the bunch crossing following the correct one are due to triggers representing tracks which were lost at the correct bunch crossing. This specific ghost rate is fairly constant with angle and the lost triggers can be recovered by the DTF track matching if a window of ± 1 bunch crossing centred around the current one is considered.

At 30° or more the *purity plot* shows a worsening of the resolution of the track position.

By summing up all collected data recorded with the default configuration at angles of incidence

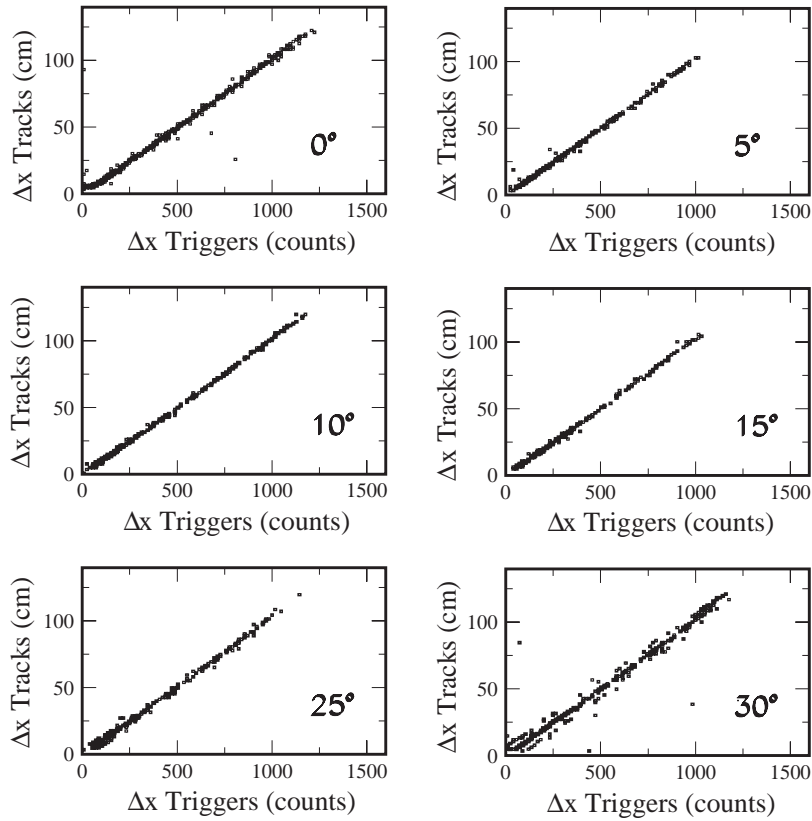


Fig. 41. Purity plots for different angles of incidence of the beam. One count is angle dependent and roughly corresponds to 1.4 mm at normal incidence.

between 0° and 30° the plot in Fig. 42 is obtained. The data are compared to the results obtained using the emulator. An overall efficiency above 90% is observed for tracks separated by more than 30 cm (outside the region with disconnected cells). The emulator reproduces the data correctly.

An analysis of all data considering also nearby bunch crossings (i.e. from the DTF point of view) was performed and the result is shown in Fig. 43 where the plotted quantities are the same as in Fig. 39. The efficiency is above 90% for well separated muons.

7.2. Alternative configurations

Among the alternative configurations listed in Section 5.3 only those are considered to be

relevant for the di-muon trigger which are mainly related to a TS specific parameter.

In configuration #6 the *carry* option is disabled. This does not affect single muon events but is very important for muon pairs: in this case the local trigger would miss most of the di-muons, triggering only on one of the two. This can be seen in Fig. 44a which shows the *di-muon trigger efficiency* as a function of the distance between the two muons comparing the default configuration with the configuration where the *carry* option is disabled. The algorithm is almost completely inefficient for tracks separated by more than 20 cm (i.e. the case when the trigger system would deliver two first TRACO tracks). The tracks belonging to the vertical band with null separation of the two triggers in the corresponding purity plot (Fig. 44b) are ghosts and account for about half of the




Fig. 42. Efficiency to detect two triggers at the correct bunch crossing, obtained by summing up all data taken in the default mode at 0° , 5° , 10° , 15° , 20° , 25° and 30° angle of incidence of the beam. The black dots are the results of the emulator, whereas the white dots are the data. One count is angle dependent and roughly corresponds to 1.4 mm at normal incidence.




Fig. 43. Efficiency to detect two triggers, either both at the correct bunch crossing or one at the correct bunch crossing and the other one at the nearby bunch crossing. Both triggers correctly reproduce the position of the two incident tracks. All data collected at 0° , 5° , 10° , 15° , 20° , 25° and 30° angle of incidence of the beam were summed together. The plotted quantities are the same as in Fig. 39.

events. Hence we conclude that the *carry* option is mandatory to trigger on close-by muon pairs.

The option enabled in configuration #7 is related to the treatment of *pile-up* events. When this option is enabled and a *pile-up* event occurs, the TS compares the second track at the current bunch crossing with the incoming first track at the next bunch crossing and selects the best one according to its quality. More in detail the first track of the incoming bunch crossing is taken (and the second track at the current bunch crossing is discarded) if it contains a H quality segment in the inner ϕ SL. In case of di-muons this functionality gives slightly better results in terms of efficiency and purity as shown in Fig. 44c and d.

Configuration #8 disables *ghost1* rejection in TS while still keeping *ghost2* rejection active. The corresponding results are shown in Fig. 44e and f. The efficiency is rather similar to the default configuration. A slightly higher efficiency for very close-by muon pairs is observed, as in principle expected, when the hits of two muons are processed by the same TRACO. This does not seem to result in a corresponding excess of ghosts. Both *ghost1* and *ghost2* are accepted in configuration #9. Results are shown in Fig. 44g and h. The increase in efficiency at short muon distances compared to the default condition is fake since many close ghosts appear. It should be reminded that when the last two options are disabled, the ghost rate for single muon events increases considerably.

7.3. The back-up mode

If the TSMS fails the *back-up mode* can be activated. It allows partially to recover the functionalities of the local trigger by performing the final sorting of the trigger candidates in the TSMD instead of in the TSMS. The sorting is done in a simplified way, using only the information of the quality of the full tracks. In addition each TSMD performs an independent sorting on half a chamber and delivers a first track and a second track. Therefore the task of selecting the absolute best and the absolute second best track will be done by the Sector Collector.

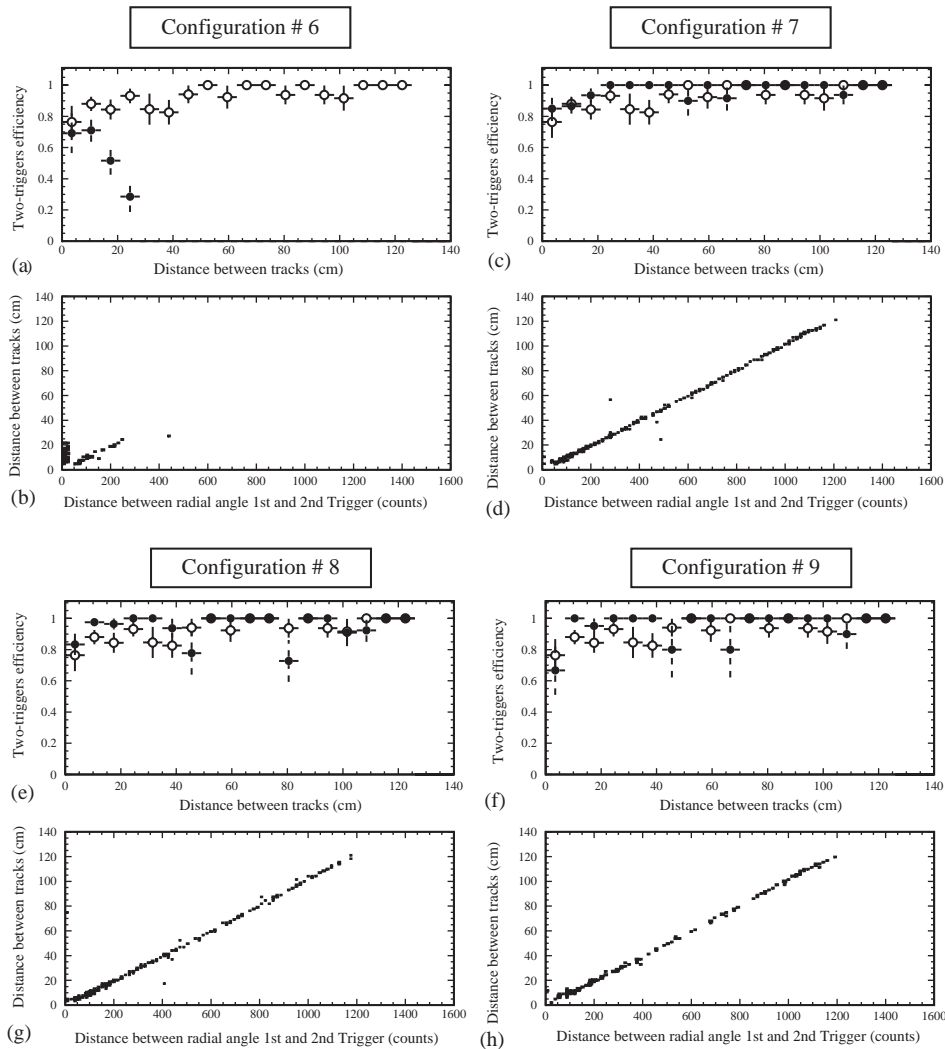


Fig. 44. Comparison with the default configuration (white dots) for trigger efficiency for di-muons at the correct bunch crossing: (a) when the *carry* option is disabled (configuration #6); (c) when the H_1 *trigger recovery* option is disabled (configuration #7); (e) when the *ghost2* options are disabled (configuration #8); (g) when the *ghost1* and *ghost2* options are disabled (configuration #9). The related *purity plots* for all configurations are shown right below each efficiency plot. Results refer to a beam inclination of 10° .

A peculiar aspect of the *back-up mode* is that the TSMD does not know whether a trigger candidate was labelled by the TRACO as First or Second Track since this tag exists only in the preview tracks which are not used by the TSMDs. Therefore if the two selected tracks, after the whole sorting, come from the same TRACO, in the *back-up mode* they are labelled to belong to two consecutive bunch crossings. Another feature of

the *back-up mode* is that uncorrelated L triggers are suppressed. The results for single muons are therefore similar to the results of configuration #2.

During the data taking special runs were taken to test the sorting algorithm of the *back-up mode*. The whole chamber was connected to a single TSMD to test the performance of the modified sorting. Results are shown in Fig. 45. The meaning of the plotted quantities is the same as in Fig. 39.




Fig. 45. Efficiency of the local trigger when the *back-up* mode is activated. The distribution of the two muon candidates track separation (as determined by the fit of the TDC hits in the ϕ view of the chamber) for all di-muon candidates (white histogram), for the events which also have two triggers at the correct bunch crossing, both identifying the correct position of the two muons (dotted histogram) and for the events which have one trigger at the correct bunch crossing and one at the nearby bunch crossing, both providing a correct position of the two muons (shaded histograms) are shown on the top (a). The sum of the dotted and shaded histograms, divided by the white one, represents the efficiency of the local trigger to deliver two triggers, both with correct identification of the muon position (b). The efficiency measured in the default configuration is superimposed (white dots).

Close-by muon candidates mostly have correct triggers and the overall efficiency is quite high and definitely comparable with the default operating mode.

When the *back-up mode* is activated the quality of both track segments delivered by the local trigger is good, as it is in the *default processing mode*. In the *back-up mode* the quality of the triggers is also good, if one trigger is delivered at the correct bunch crossing and the other at the next bunch crossing. This is contrary to the *default mode*, when usually the track at the wrong bunch crossing has a poorer quality. The better performance of the *back-up mode* in this special case is explained by the fact, that the assignment of the second trigger to a bunch crossing next to the

correct one is not due to an incorrect bunch crossing identification in the BTI but to the typical treatment of the trigger candidates in this mode as described above.

8. Conclusions

A complete test of the drift tubes local trigger system was performed on a muon beam with the same time structure as at LHC allowing the calculation of efficiencies and noise in several situations. The trigger system showed a very good detection efficiency for single muons (flat at 99% up to 35° angle of incidence) and dimuons (almost 100% efficiency for tracks separated by more than 30 cm). The noise in the system was within the design value. We have verified that some room is available for further noise reduction at the price of an acceptable efficiency reduction if the CMS environment will be worse than expected. The test of the emergency system that could be activated in case of major failures in the trigger electronics provided satisfactory results.

The results obtained fully validate the design choices and show that we meet the trigger requirements with the foreseen default operations flow.

Acknowledgements

We acknowledge the support of the staff of the electronics and mechanical workshops of INFN Bologna, INFN Padova and CIEMAT. In particular we are grateful to Dr. Eng. I. Dantone and acknowledge the valuable help of L. Barcellan, M. Boldini, M. Buldrini and L. Maselli.

Furthermore we acknowledge the support by the German “Bundesministerium für Bildung und Forschung” (BMBF).

References

- [1] CMS Technical Proposal, CERN/LHCC 94-38, LHCC/P1, 1994.
- [2] CMS, The Muon Project, Technical Design Report, CERN/LHCC 97-32.

- [3] M. Aguilar-Benitez, et al., Nucl. Instr. and Meth. A 480 (2002) 658.
- [4] M. Pegoraro, in: Proceedings of the Fourth Workshop on Electronics for LHC Experiments, CERN LHCC 98-36, 1998, p. 257.
- [5] J. Christiansen, High Performance Time to Digital Converter, CERN/EP-MIC, 2002.
- [6] C. Albajar, et al., Nucl. Instr. and Meth. A 525 (2004) 465.
- [7] CMS, The Level-1 Trigger, Technical Design Report, CERN/LHCC 2000-038.
- [8] B. Taylor, IEEE Trans. Nucl. Sci. NS-45 (1998).
- [9] M. De Giorgi, et al., Nucl. Instr. and Meth. A 438 (1999) 302.
- [10] F. Gasparini, et al., Nucl. Instr. and Meth. A 336 (1993) 91.
- [11] F.R. Cavallo, et al., Test of MB3 Muon Barrel Drift Chambers with Cosmic Rays, CMS Note 2003/017.
- [12] G.M. Dallavalle, et al., in: Proceedings of the Fourth Workshop on Electronics for LHC Experiments, CERN LHCC 98-36, 1998, p. 291.