Eur Phys J C **34**, s01, s379–s384 (2004)
Digital Object Identifier (DOI) 10.1140/epjcd/s2004-04-038-2

EPJ C direct
electronic only

CMS B physics reach

V. Ciulli, on behalf of CMS Collaboration

Istituto Nazionale di Fisica Nucleare, Sezione di Firenze, Italy, e-mail: Vitaliano.Ciulli@fi.infn.it

Received: 4 August 2003 / Accepted: 3 February 2004 /
Published Online: 13 July 2004 –

Abstract. At the design luminosity of 10^{34} cm^{-2} s^{-1}, about 10^6 $b\bar{b}$ pairs are expected to be produced every second at the Large Hadron Collider. With such huge statistics and the precise tracking of the CMS detector B mesons can be investigated for CP violation, B^0_s-$\bar{B}^0_s$ mixing and rare decays. The trigger, however, cannot retain all $b\bar{b}$ events for a later selection of exclusive B decays. A dedicated trigger strategy, which uses tracking immediately after the first trigger stage, is presented and results on few important benchmark channels are given.

1 B physics at LHC

At the design luminosity of 10^{34} cm^{-2} s^{-1}, about 10^6 $b\bar{b}$ pairs/s are expected in 14 TeV centre-of-mass energy pp collisions at the Large Hadron Collider. A wide B-physics program, including CP violation, rare decays, B^0_s-$\bar{B}^0_s$ mixing, can therefore be covered with the CMS detector, since its precise tracking system allows invariant mass and proper time of b-hadrons decays to be measured with appropriate resolutions. Nevertheless, the rate at which events can be archived for later analysis is only 100 Hz. Because all interesting physics processes must fit in these 100 events per second, the background rate reduction and the signal efficiency of the trigger selections are the most important issues. It is particularly difficult to trigger efficiently on B-physics channels. Only an exclusive reconstruction of few relevant decays can separate interesting events from the huge background due to the $b\bar{b}$ events themselves.

A dedicated trigger strategy, which uses tracking immediately after the first trigger stage, is presented here. First the CMS Trigger and Data Acquisition system is briefly reviewed and the track reconstruction is described. A few benchmark channels are analysed to show the B-decay triggering capabilities and to exemplify the achievable physics results.

2 CMS trigger and data acquisition

The CMS Trigger and Data Acquisition System [1,2] is designed to select events at a maximum rate of 100 Hz for archiving and later analysis, in two steps:

- a Level-1 Trigger built on custom processors reduces the input 40 MHz bunch-crossing rate to 100 KHz;
- the Data Acquisition System (DAQ) reads out the data through an event-builder switch network and performs the remaining analysis, the High Level Trigger (HLT), on a farm of commercial processors.

Compared to more conventional architectures, in which a Level-2 Trigger is present before the switch network, the CMS DAQ design is more flexible. Because any software complexity and any high-level information can in principle be used in the HLT processor farm, the online selection can be changed on demand to accommodate for unexpected situations. Moreover, the use of commercial processors allows the HLT to benefit at low cost from the evolution of computing technology. Finally the CMS DAQ architecture is modular: each slice is able to process events at a rate of 12.5 kHz. At LHC startup only four slices will be deployed, which allows a maximum Level-1 output rate of 50 kHz. To account for unforeseen backgrounds, this rate was further reduced by a factor of three in all trigger studies.

The Level-1 Trigger uses coarse information from calorimeters and muon stations. Events are selected if large energy deposits are present in the calorimeters or if muon candidates are found in the muon stations. The physics objects measured by the Level-1 Trigger are then refined in the High Level Trigger with information from the tracking system, which allow topological cuts such as muon isolation to be applied.

At LHC, most of the QCD-produced b quarks have small transverse momentum, P_T. They are therefore overwhelmed by a prohibitive background and are usually rejected by the Level-1 Trigger. The b-quark decays into muons, however, suffer from a much smaller background, so that relatively low-momentum muons can be used to trigger on these decays. A draft Level-1 Trigger table for the low luminosity period allocates a total of 3.6 kHz to muons triggers. The inclusive trigger at $P_T > 14$ GeV/c yields a rate of 2.7 kHz, actually dominated by mismeasured lower P_T muons. The remaining 0.9 kHz is the rate of dimuon events with $P_T > 3$ GeV/c for both muons.

After the HLT, the inclusive muon rate is reduced to 25 Hz with a P_T threshold of 19 GeV/c, and the dimuon rate is reduced to 5 Hz with a common threshold of 7 GeV/c for the two muons. Contributions from all sources of muons to the HLT rate at high luminosity are shown in Fig. 1 as a function of the transverse momentum threshold, before and after the isolation cut. For the applied threshold the rate is dominated by muons from W decays. Most of b quarks accepted by the Level-1 Trigger are rejected by the HLT muon selection.

In the HLT, however, the full event information is available. A few exclusive B decays can therefore be reconstructed around the muon(s), with little remaining background. The only constraint is the time needed to access and analyse the event data. One single processor of the HLT farm deals with a single event from beginning to end of the analysis. In order to reduce the processing time, events have to be rejected as soon as possible. This fast rejection can be obtained with regional reconstruction. The reconstruction starts from the region of interest defined by the Level-1 Trigger primitive and continues with increasing refinement until the event is discarded or definitely accepted. The average CPU time of the HLT selection, taking into account both signal and background events for all trigger streams, is about 300 ms on a 1 GHz Pentium-III processor.

Charged particle track reconstruction starts from the pixel vertex detector, which provides up to three very precise three-dimensional coordinates. Tracks seeds are built from pairs of pixel hits compatible with a track originating from the interaction region and propagated from inside out, through the silicon microstrips layers. The track reconstruction is very expensive in term of CPU time, both because of the huge number of combinations in seeds generation and because of the propagation time. In this respect, the regional and on demand reconstruction is particularly useful. The number of seeds is significantly reduced if only seeds compatible with a P_T in excess of 1 GeV/c are considered. These track candidates are then used to search for primary vertices, with an efficiency greater than 95%. The signal primary vertex is identified from the trigger lepton impact parameter and the vertex constraint used to further reduce the number of seeds. In addition, only the track candidates compatible with a given (η, ϕ) direction or other conditions are propagated. The propagation can be also stopped as soon as enough precision on the track parameters is reached, but in an inside-out reconstruction, the track parameters at the origin are not known until the information from all hits is propagated back. It is therefore important to know in advance when the propagation can be stopped. The P_T and impact parameter resolutions with a partial reconstruction are shown in Fig. 2. The asymptotic accuracy is almost reached after only five or six hits. After seven or eight hits the fraction of fake tracks is already below 1% for b quarks with 100 GeV transverse momentum.

The CPU time needed for this track reconstruction depends on the event type and on the complexity of the analysis. Typically, less than 300 ms are required to reconstruct tracks from secondary vertices, comparable to the average HLT time. In the following it is shown how this fast reconstruction can be used to trigger a few specific B-mesons decays.

3 Benchmark channels

The B-physics selection is triggered at Level 1 by the presence of one or two muons. The aim of the selection is to reconstruct the decay of the associated b hadron by using invariant mass and displaced vertex cuts. Results are shown for

- $B_s \to \mu^+\mu^-$,
- $B_s \to J/\psi\phi \to \mu^+\mu^- K^+K^-$,
- $B_s \to {D_s}^-\pi^+ \to \phi\pi^-\pi^+ \to K^-K^+\pi^-\pi^+$.

From the results of the first and second channel, an inclusive trigger path for $J/\psi \to \mu^+\mu^-$ is also derived.

3.1 The $B_s \to \mu^+\mu^-$ channel

The $B_s \to \mu^+\mu^-$ decay allows flavour changing neutral currents to be tested. The design of an appropriate trigger is a challenge because branching ratio predicted by the Standard Model is only $(3.5 \pm 1.0) \times 10^{-9}$ [3]. The main motivation to study this channel is the possible enhancement of the branching ratio because of new physics processes.

Events are triggered at Level 1 with at least two muon candidates of $P_T > 4$ GeV/c and with opposite charge. The two muons are required to be isolated, to come from a common displaced vertex and to have an invariant mass close to the B_s mass.

First, all pixel hit pairs compatible with a charged particle track of $P_T > 4$ GeV/c and transverse impact parameter smaller than 0.1 cm are identified. These hit pairs are then used to estimate the longitudinal primary vertex positions. The three most significant primary vertices are retained. Hit pairs are then filtered with the vertex constraints and the regions defined by the Level-1 muons ($\Delta\eta < 0.5$ and $\Delta\phi < 0.8$) and used for the reconstruction of tracks. The reconstruction is stopped when one of the following condition is fulfilled: (a) the track P_T is lower than 4 GeV/c at 5 σ; (b) six hits along the trajectory are found; (c) the relative P_T resolution is smaller than 2%. A B_s candidate is found if two opposite charged particle tracks with an invariant mass within $\pm 150\,\mathrm{MeV}/c^2$ of the B_s mass (Fig. 3) are reconstructed. In order to suppress the combinatorial background, tracks are vertexed and the candidate is retained if the χ^2 of the vertex fit is smaller than 20, and the reconstructed decay length in the transverse plane is larger than 150 μm. The global selection efficiency on signal events is 5% and the timing is on average 240 ms on a 1 GHz Intel Pentium-III CPU. The estimated background rate is below 2 Hz (Table 1).

The algorithm described is identical to the offline selection described in [4] but with much softer cuts. The events selected by the offline analysis (about seven signal

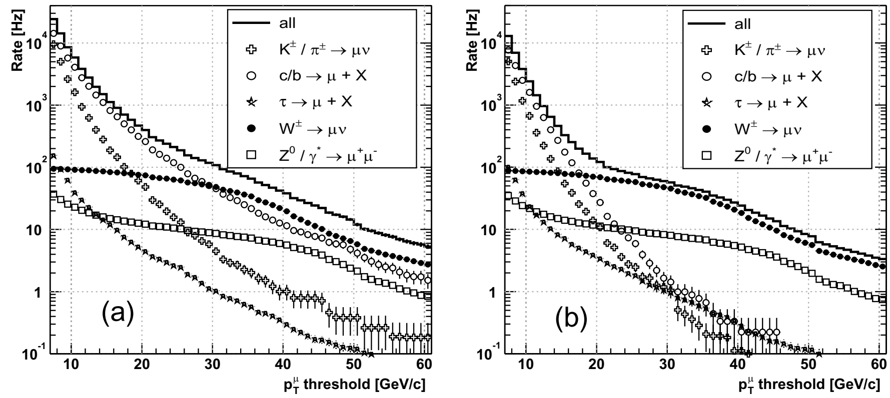

Fig. 1. Contributions to the HLT rate at high luminosity from all sources of muons **a** before and **b** after muon isolation has been required

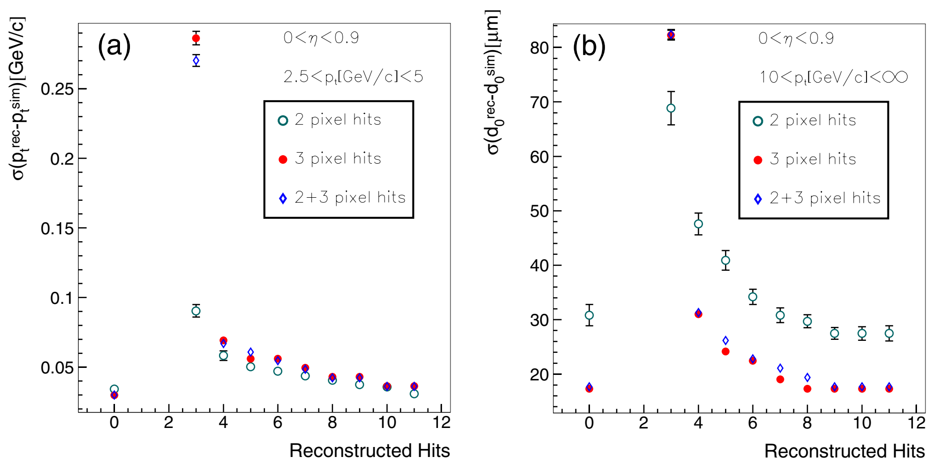

Fig. 2. The resolution on **a** P_{T} and **b** impact parameter for partial track reconstruction as a function of the number of hits along the track. The leftmost point at zero hit shows the full tracker performance

events and less than one background event in $10\,\mathrm{fb}^{-1}$) are therefore expected to all satisfy the HLT conditions. A 5σ observation of the decay $\mathrm{B_s} \to \mu^+\mu^-$ is therefore expected with less than 15 fb^{-1}.

3.2 $\mathrm{B_s} \to \mathrm{J}/\psi\phi$

The "gold-plated" decay, $\mathrm{B_s} \to \mathrm{J}/\psi\phi \to \mu^+\mu^-\mathrm{K}^+\mathrm{K}^-$, is one of the most interesting B-physics channel at LHC. An angular analysis of the final state allows measurements of the $\mathrm{B_s^0}$-$\mathrm{\bar{B}_s^0}$ mixing parameter $\Delta\Gamma_{\mathrm{s}}$ and the weak mixing phase $\phi_{\mathrm{s}} = -2\lambda^2\eta$, from which the Wolfenstein parameter η can be extracted. Moreover, this decay is a sensitive probe to CP-violating contributions from new physics.

This channel is also triggered at Level 1 by the presence of two muons. The HLT reconstruction can be formally separated in two steps (called Level 2 and Level 3 in the following), even if they are both performed in the unique HLT farm.

The reconstruction of the muons from J/ψ decays is done as for the $\mathrm{B_s} \to \mu^+\mu^-$ decays, but slightly tighter cuts on the dimuon mass and the secondary vertex are required to keep the background rate at an acceptable level. The mass window is decreased to 100 MeV/c^2 (Fig. 4) and the secondary vertex has to be at least 200 μm away from the beam axis and must be fit with a χ^2 smaller than 10. This dimuon selection leads to an inclusive rate of about 15 Hz for low luminosity, 90% of which is made of J/ψ from b quarks. This Level 2 processing takes on average 260 ms/event and can be run on every event passing the Level-1 dimuon trigger. The selection is inclusive enough to allow also other analyses involving J/ψ to be performed, such as $\mathrm{J}/\psi\mathrm{K_S^0}$ decays. After this first stage the full reconstruction of the $\mathrm{B_s}$, which is more demanding in CPU time, is run to reduce the rate below 1 Hz. To reconstruct the ϕ, all tracks within a cone of $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 1.5$ around the reconstructed J/ψ direction are reconstructed either until six hits are found or when the relative uncertainty on P_{T} becomes smaller than 2%. Oppositely-charged-particle tracks are paired and retained if the invariant mass is within 10 MeV/c^2 of the ϕ mass. The $\mathrm{B_s}$ candidate is retained if the invariant mass of the J/ψ and ϕ system falls within 60 MeV/c^2 of the $\mathrm{B_s}$ mass. To further decrease the background, the $\mathrm{B_s}$ vertex is required to be fit with a χ^2 smaller

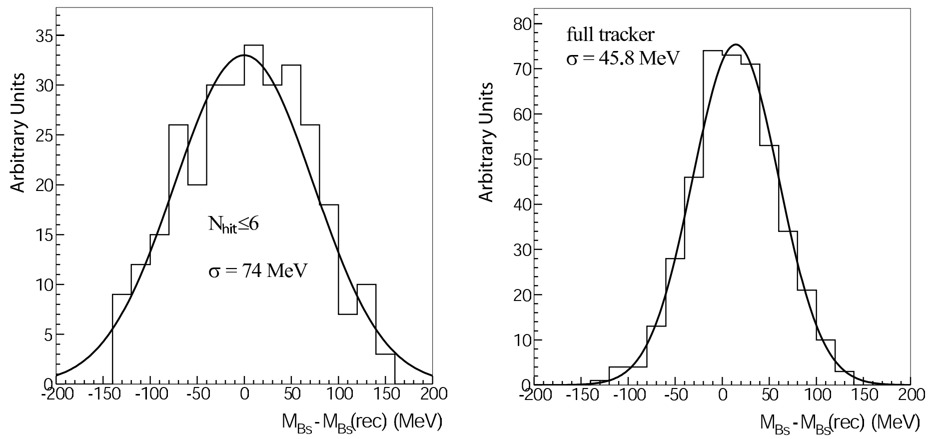

Fig. 3. Dimuon invariant mass distribution in $B_s \to \mu^+\mu^-$ decays after HLT (*left* and offline reconstruction (*right*)

Table 1. Trigger efficiencies, number of signal events expected in 10 fb^{-1} and background rate for the $B_s \to \mu^+\mu^-$ reconstruction. The global efficiency is that of the combined Level-1 and HLT selections. The last two columns show respectively the expected numbers of signal and backgrounds events expected after the offline selection

Efficiency			Events	Trigger	Offline	Offline
L1 Trigger	HLT	Global	per 10 fb^{-1}	Rate	signal	backg.
15.2%	33.5%	5.1%	47	< 1.7 Hz	7	< 1

Table 2. Trigger efficiencies, number of signal events expected in 10 fb^{-1} and background rates for the $B_s \to J/\psi\phi$ channel

Efficiency		Rate	Efficiency	Rate	Events
L1	L2	L2	L3	L3	per 10 fb^{-1}
16.5%	13.7%	14.5 Hz	8.7%	< 1.7 Hz	83800

than 200. The efficiencies and yields for signal events are shown in Table 2 together with the backgrounds rates at low luminosity. The average execution time for signal and background events on a 1-GHz Intel-Pentium-III CPU is 800 ms for the exclusive reconstruction of the B_s (including the Level 2 processing time), out of it about 400 ms is spent for the reconstruction of the ϕ.

The subsequent offline analysis is expected to select 60 to 70% of the remaining signal. About 200 000 signal events are therefore expected in 30 fb^{-1}. From the results of previous studies [3] a relative uncertainty on $\Delta\Gamma_s$ of about 15% and an uncertainty on the weak mixing phase ϕ_s of about 0.025 for $x_s = 20$ can be expected.

3.3 $B_s \to D_s{}^-\pi^+ \to \phi\pi^-\pi^+ \to K^-K^+\pi^-\pi^+$

The $B_s \to D_s{}^-\pi^+$ decay followed by $D_s{}^- \to \phi\pi^-$ and $\phi \to K^+K^-$ decays allows ΔM_s to be measured. The B_s CP state at decay time is tagged by pion charge and the proper time can be measured with about 70 ps resolution from the full B_s reconstruction. These fully hadronic decays can only be triggered by a muon coming from the decay of the other b quark in the event. In addition to the standard single muon trigger, the combination of a low-P_T muon and a low-E_T jet was studied. The expected rates as a function of the different cuts on the muon and the soft jet are shown in Table 3 for the HLT and the Level-1 Trigger. The corresponding yields of signal events are shown in Table 4. In order to reconstruct as fast as possible the decay chain resonances tracks are reconstructed only with three hits. First the primary vertex is reconstructed as in Sect. 2. Tracks compatible with originating from within 1 mm from the primary vertex and with a transverse momentum in excess of $P_T > 0.7$ GeV/c, are searched for with two hits in the pixel detector and one in the first layer of the silicon detector. This procedure allows for a lever arm large enough to determine the momentum with the required precision. A search is then made for ϕ, D_s and B_s candidates with the following three requirements: $\Delta R(K^+K^-) < 0.3$, $\Delta R(\phi\pi^-) < 1.2$ and $\Delta R(D_s{}^-\pi^+) < 3.0$. The mass resolutions for ϕ, D_s and B_s after HLT reconstruction are presented in Fig. 5. Candidates are required to have an invariant mass within three times the resolution for all of the three resonances and transverse momentum greater than 2, 4 and 5 GeV/c for the ϕ, the D_s and the B_s, respectively. Cuts are also applied on ϕ helicity and on the D_s decay vertex. All details

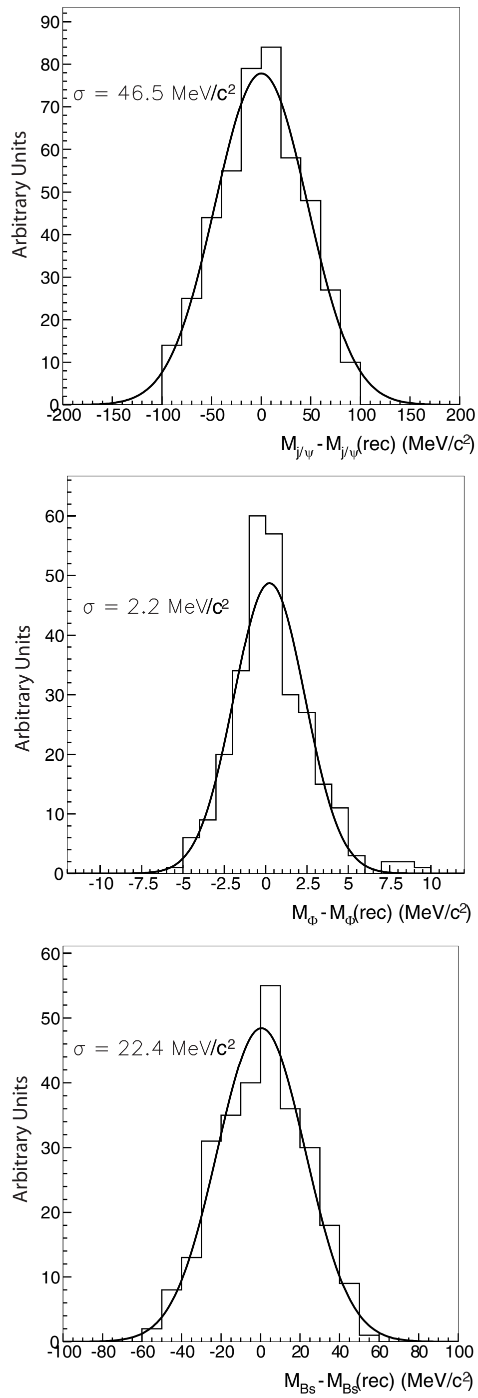

Fig. 4. Mass resolution for the J/ψ, the ϕ and the B_s after HLT reconstruction

can be found in [5]. With respect to the events which pass the Level-1 trigger selection, the signal efficiency after the above cuts is about 9%, while the background is suppressed by a factor 250. On a single 1-GHz Intel-Pentium-III CPU, the average execution time is 640 ms.

The number of selected signal events depends on the bandwidth allocated to this channel. From the above tables even a relatively large threshold of 14 GeV/c gives a rate of 17 Hz on tape after HLT. A combined single-muon

Table 3. Background rates after the HLT (*Level-1 Trigger*) for the $B_s \to D_s{}^-\pi^+$ channel for different cuts on the muon and the jet

Muon P_T (GeV/c)	Background rates in kHz $E_T > 0$ GeV	$E_T > 20$ GeV	$E_T > 30$ GeV
6	0.16(26)	0.082(8.5)	0.055(3.6)
10	0.037(6.4)	0.021(2.5)	0.014(1.3)
14	0.017(3.2)	0.010(1.3)	0.008(0.7)

Table 4. Number of signal events for the $B_s \to D_s{}^-\pi+$ channel after the HLT (*Level-1 Trigger*) for different cuts on the muon and the jet

Muon P_T (GeV/c)	Events in units of 10^3 $E_T > 0$ GeV	$E_T > 20$ GeV	$E_T > 30$ GeV
6	7.4(65)	3.1(24)	1.5(11)
10	2.0(20)	1.0(9.6)	0.6(5.3)
14	1.0(11)	0.5(5.4)	0.3(3.0)

and soft-jet trigger condition would be efficient to decrease the Level-1 rate. It is however helpless because it would also decrease the signal-to-background ratio at HLT. Since the overall possible rate on tape is 100 Hz, the bandwidth allocated to this channel most probably cannot exceed 5 Hz. If the fraction of events written to tape is scaled accordingly, about 300 signal events are expected for 20 fb^{-1}. With this amount of events CMS would be sensitive up to $\Delta M_s \leq 20$ ps^{-1} [5]. In order to fully cover the range allowed by the Standard Model, $\Delta M_s < 26$ ps^{-1}, about 1000 events are needed.

4 Exclusive B-triggers bandwidth

The rate of events selected by exclusive B-triggers algorithms is of the order of few Hz and can be further reduced with refined reconstruction methods. The main limitation is the Level-1 rate which can be processed with these algorithms. Today, the average CPU time used in HLT for each event selected by any Level-1 physics stream on a 1-GHz Pentium-III processor is about 300 ms [2]. This time was used to estimate the computing power needed to process the designed 50 kHz Level-1 output at LHC startup. Any improvement either in processor performance or in the optimization of the reconstruction algorithms will therefore allow

- additional selections to be performed;
- refined reconstruction algorithms to be applied in order to further reduce the HLT output rate.

Present Level-1 thresholds have been chosen to limit the output rate to 16 kHz, a factor of three less than the HLT maximum capability of 50 kHz. If the background conditions were not worse than currently estimated, either 30 kHz of bandwidth would become available or the time per event would be a factor of three larger. As a lowering of

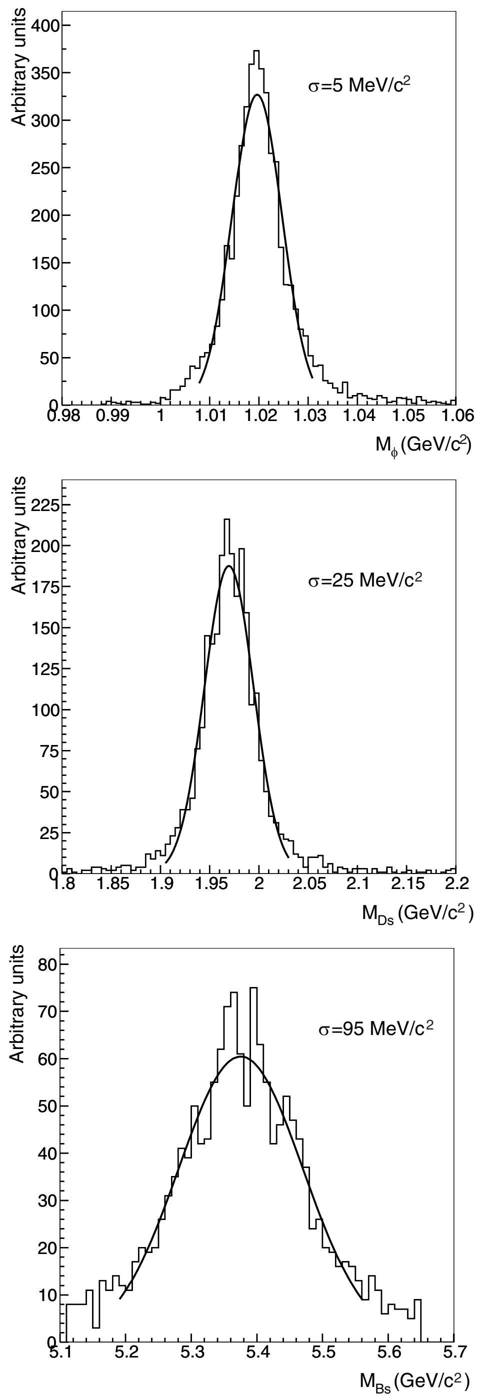

Fig. 5. Mass resolution for ϕ, D_s and B_s after HLT reconstruction

the Level-1 thresholds would not improve the sensitivity on other physics channels, the available free bandwidth can be used for B physics as long as the output rate of the HLT is kept below 100 Hz.

The same holds if luminosity at startup is below the design value of $2 \times 10^{33}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$. In this case the HLT output would be below 100 Hz, and higher background rates for B-triggers would therefore be acceptable.

Finally, as the luminosity drops between the beginning and the end of the fill (by about a factor of two), the trigger thresholds can be changed to allow some B-physics channel to be selected.

5 Conclusions

The flexibility of the CMS Trigger and DAQ system, and in particular of the availability of reconstructed charged particle tracks at HLT, make possible a selection of exclusive B decays, even in the tight timing constraints imposed by LHC rates. Fast tracking, either performed in a selected tracker region or stopped as soon as the required precision is reached, is a key advantage in this respect.

The bandwidth dedicated to the B physics will strongly depend on LHC startup conditions and evolution of the computing performance by then. The implementation of the selection algorithms will have to be addressed in much more details in the coming years. In all studied benchmark channels, however, it has been demonstrated that CMS will be able to cover the full parameter range allowed in the Standard Model.

References

1. CMS Coll.: "The Trigger and Data Acquisition project, Volume I, The Trigger Systems, Technical Design Report", CERN/LHCC 2000-38, CMS TDR 6.1, 15 December 2000
2. CMS Coll.: "The Trigger and Data Acquisition project, Volume II, Data Acquisition & High Level Trigger, Technical Design Report", CERN/LHCC 2002-26, CMS TDR 6.2, 15 December 2002
3. "Standard Model Physics (and more) at the LHC", Proceedings of 1999 CERN Workshop, Ed. G. Altarelli and M. L. Mangano, Geneva, Switzerland, May 25-26, 1999, CERN-2000-004
4. A. Nikitenko, A Starodumov, and N. Stepanov: "Observability of $B^0_{(d)s} \to \mu\mu$ decay with the CMS detector", CMS Note 1999/039
5. A. Giassi, F. Palla, and A. Starodumov: "Study for a High-Level Trigger for the decay channel $B_s \to \pi + D_s \to \phi(KK)\pi\pi$", CMS Note 2002/045