Search for exotic long-lived particles with the LHCb detector

Pieter Norbert Yvonne David

“Zoektocht naar exotische langlevende deeltjes met de LHCb-detector”

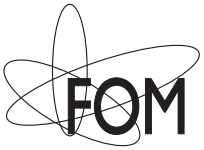

Dit werk maakt deel uit van het onderzoekprogramma van de Stichting voor Fundamenteel Onderzoek der Materie (FOM), die deel uitmaakt van de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO).
Het werk is gefinancierd door de NWO VIDI-subsidie getiteld “A search for long-lived heavy particles” van dr. W.D. Hulsbergen.

VRIJE UNIVERSITEIT

Search for exotic long-lived particles with the LHCb detector

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. V. Subramaniam,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de Faculteit der Exacte Wetenschappen
op donderdag 7 juli 2016 om 11.45 uur
in de aula van de universiteit,
De Boelelaan 1105

door

Pieter Norbert Yvonne David

geboren te Lokeren, België

promotor:	prof.dr. M.H.M. Merk
copromotor:	dr. W.D. Hulsbergen

Introduction

Despite the success of the standard models of particle physics and cosmology, a number of open questions about the fundamental structure of nature and the history of the universe remain. Searching for physics beyond the standard model is therefore one of the main goals of the ongoing experimental high-energy physics efforts. One of the pursued strategies is the detection of new particles produced in high-energy particle collisions at the large hadron collider, the accelerator with the highest collision energy that has ever been built. The production of new particles with a lifetime that is long enough to decay after a few cm into known particles, is an experimentally attractive signature. This possibility is also theoretically well motivated: it features in several classes of theories for physics beyond the standard model, often connected with the presence of a particle dark matter candidate. The main subject of this thesis is a search for such a new long-lived particle. The theoretical motivation, a number of specific models, and the experimental status of the search for such signatures, focusing on hadronic final states, will be discussed in chapter 1.

For long-lived particles that have a comparatively small mass and lifetime, below about $50\,\mathrm{GeV/c^2}$ and 1 ns, respectively, the LHCb experiment provides particularly interesting opportunities. The decays of heavy long-lived particles produce several jets or leptons with high transverse momenta, in which case the general purpose detectors at the large hadron collider, CMS and ATLAS, can exploit this signature and take advantage of the large data samples that they can record. Also the decays with very large displacements of $\mathcal{O}(1\,\mathrm{m})$ can be searched for by these experiments, using dedicated reconstruction strategies. In the region with low mass and lifetime, however, they have difficulties to efficiently identify collisions containing such a signature, especially at the level of the trigger system. This needs to reject most recorded collision events, except for a tiny fraction that can be stored, and that needs to arrive at a decision in a short time. The LHCb experiment, designed for the study of beauty and charm hadron decays, suffers less from this problem. Its trigger system is built for selecting tracks and low-multiplicity vertices originating from a point that is only slightly displaced from the collision point, down to low track momenta. In addition, the smaller acceptance and the lower number of interactions

per bunch crossing lead to a smaller amount of information to be stored for each collision event, which allows to retain a larger fraction. A more detailed description of the LHCb detector, and the properties that make it suitable for such searches, will be given in chapter 2

The experimental search described in chapters 3 to 5 aims at finding long-lived particles that decay into hadronic final states that can be reconstructed as a pair of jets. As few further assumptions as possible are made about the production process: even though pair production of long-lived particles in the decay of a scalar resonance is assumed, like in most other experimental searches for similar signatures, only one long-lived particle decay is required to be found. The lifetime range covered is as large as achievable with the reconstruction and trigger setup in the run 1 data taking period, which ended in 2012. After a description of the search strategy and selection in chapter 3, the effects that cause a systematic uncertainty on the signal efficiency are discussed in chapter 4. The obtained results are presented in chapter 5.

This analysis builds on the efforts on long-lived particle searches that have been ongoing in the LHCb collaboration. The first LHCb search conducted was aimed at pairs of displaced vertices in the 2010 data sample [1], and was based exclusively on vertex reconstruction. Later on, a procedure for jet reconstruction was developed, and adapted for jets produced at a displaced point. The work carried out in the context of this thesis contributed to a publication using the 2011 data sample [2] through studies of the backgrounds, a correct description of the signal model in simulation, and the statistical interpretation of the results. In addition, improvements were made to the trigger selection for 2012 data taking, mostly aimed at large displacements, the offline reconstruction strategy was modified to exploit the trigger candidates and the jet reconstruction as early as possible. Finally, for the results described here, the analysis of the combined 2011 and 2012 data set was performed, taking advantage of these improvements for the selection and for the understanding of systematic effects.

1

Exotic long-lived particles

1.1 The standard model and beyond

The standard model of particle physics provides a description of the interactions through the electromagnetic interaction and weak and strong nuclear interactions between elementary particles. Although it is in good agreement with experimental observations, it leaves various fundamental questions unanswered. A review of the theoretical and experimental status, including historical references, can be found in [3]. Here only a brief overview is given as an introduction to new physics models predicting the existence of exotic long-lived particles and to the experimental search presented in this thesis.

Formally, the standard model is a relativistic quantum field theory with an $\mathrm{SU}(3)_\mathrm{C}$ $\otimes\mathrm{SU}(2)_\mathrm{L} \otimes \mathrm{U}(1)_\mathrm{Y}$ gauge symmetry. The electroweak $\mathrm{SU}(2)_\mathrm{L} \otimes \mathrm{U}(1)_\mathrm{Y}$ symmetry is spontaneously broken to $\mathrm{U}(1)_\mathrm{EM}$, which makes three of the four associated gauge bosons massive through the Brout-Englert-Higgs mechanism. The massive gauge bosons are the $\mathrm{W}^\pm$ and Z^0 bosons that mediate the weak interaction. The remaining massless boson is the photon that mediates the electromagnetic interaction. The left-handed and right-handed components of the fermion fields have different $\mathrm{SU}(2)_\mathrm{L} \otimes \mathrm{U}(1)_\mathrm{Y}$ quantum numbers: the left-handed components are in $\mathrm{SU}(2)_\mathrm{L}$ doublets, while the right-handed components are $\mathrm{SU}(2)_\mathrm{L}$ singlets, with the same electric charge $\mathrm{Q} = \mathrm{T}_3 + \frac{\mathrm{Y}}{2}$, where T_3 is the quantum number corresponding to the third component of the weak isospin and Y the weak hypercharge.

The fermions can acquire mass in a gauge-invariant way through Yukawa couplings to an $\mathrm{SU}(2)_\mathrm{L}$ doublet scalar field ϕ that acquires a vacuum expectation value

$$\mathcal{L}_\mathrm{Yukawa} = -\mathrm{U}_{ij}\delta^{ab}\overline{\Psi}_{\mathrm{L}ia}\phi_b\psi_{\mathrm{R}j} - \mathrm{U}'_{ij}\epsilon^{ab}\overline{\Psi}_{\mathrm{L}ia}\phi_b\psi'_{\mathrm{R}j} + \mathrm{h.c.} \tag{1.1}$$

where $\Psi_{\mathrm{L}i}$ denotes the $\mathrm{SU}(2)_\mathrm{L}$ doublets of left-handed fermion fields, $\psi_{\mathrm{R}j}$ and $\psi'_{\mathrm{R}j}$ the accompanying right-handed fermion fields, δ^{ab} and ϵ^{ab} represent the symmetric and antisymmetric sum over the $\mathrm{SU}(2)_\mathrm{L}$ representation indices, respectively, and U_{ij}, U'_{ij} are the Yukawa coupling matrices, with i and j generation indices. Neither for the $\mathrm{SU}(3)_\mathrm{C}$ singlet lepton nor for the color-charged quark sector, both containing

three generations of two fermions each, can the U and U′ matrix simultaneously be diagonalised. This leads to mixing between the flavour and mass eigenstates. Recent measurements of the properties of the resonance with a mass around 125 GeV/c^2 observed by the ATLAS and CMS collaborations [4] indicate that it is the massive Brout-Englert-Higgs boson H^0 associated to the scalar doublet that is responsible for both electroweak symmetry breaking and the fermion masses [5].

Each generation of leptons contains an $SU(2)_L$ doublet with weak hypercharge −1, such that the upper component, the neutrino, is electrically neutral and the lower component carries a negative unit electric charge. The charged leptons, *e*, μ and τ, are Dirac fermions and have a corresponding right-handed component with Y = −2. Neutrinos are much lighter than any other fermions, but were shown to have non-zero mass by the observation of oscillations between the different neutrino flavours. The precise nature of neutrinos, however — Dirac, like the other fermions, or Majorana spinors — the mass hierarchy and the mechanism that generates their masses, much smaller than those of any other fermion, remain open questions.

The quarks, described by another set of fermion fields, in addition transform according to the fundamental representation of $SU(3)_C$. There are again three generations with two quarks each — where one quark denotes all fields in the colour triplet. The quarks in the first, second and third generation are called up and down, charm and strange and top and beauty, respectively. Each up-type quark carries electric charge 2/3 and each down-type quark −1/3. The names reflect the position of the left-handed components in the $SU(2)_L$ doublet; the right-handed components of both are again $SU(2)_L$ singlets.

The strong interaction, described by the $SU(3)_C$ gauge theory, shows a dependence of the effective coupling constant on the energy scale opposite to the screening of electric charges through vacuum polarisation. The coupling constant of quantum chromodynamics decreases for higher energies and increases for lower energies: at low energy quarks are confined in color-neutral hadrons with binding energy related to the nonperturbative scale $\Lambda_{QCD} \sim 200$ MeV, while at high energy they are asymptotically free particles.

1.1.1 Shortcomings of the standard model

The standard model is consistent with almost all experimental tests. A few notable and recent tensions and hints of deviations are the anomalous magnetic moment of the muon [6], the hint of an excess in the search for the $H^0 \to \mu^+\tau^-$ decay [7], the ratio between the branching ratios of the decays $B^+ \to K^+\mu^+\mu^-$ and $B^+ \to K^+e^+e^-$ [8] and $\overline{B}^0 \to D^{(*)+}\tau^-\overline{\nu}_\tau$ and $\overline{B}^0 \to D^{(*)+}\mu^-\overline{\nu}_\mu$ [9], the angular distributions of the decay products in the decay $B^0 \to K^{*0}\mu^+\mu^-$ [10] and the excess in the diphoton channel

seen by the ATLAS and CMS collaborations in the 2015 LHC data [11, 12].

There are also theoretical and phenomenological arguments suggesting that a more fundamental theory exists, and that first signs of this new physics may be observable at the energy scale currently probed by the LHC accelerator.

Structure of the fermion sector

Most of the free parameters of the standard model, for whose values the theory provides no explanation, are associated to fermions: twelve masses and two mixing matrices between the three generations. The masses, which are shown in fig. 1.1, span a large range, more than five orders of magnitude even if the smallness of the neutrino masses would be explained by a different mechanism. As can be seen in fig. 1.2, the quark mixing matrix shows a strong hierarchical structure, with suppressed mixing between the generations, while the current best fit lepton mixing matrix elements show much larger mixings.

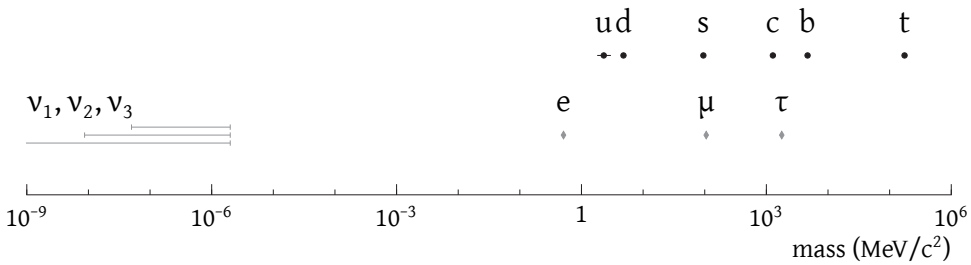

Figure 1.1: Illustration of the masses of the standard model quark (black circles) and lepton (gray diamonds) masses

Figure 1.2: Graphical illustration of the size of the matrix elements of the CKM matrix, which relates the down-type quark flavour eigenstates with the mass eigenstates (left), and the PMNS matrix, which relates the neutrino mass eigenstates ν_j and flavour eigenstates ν_l. The area of each circle is proportional to the size of the corresponding matrix element

Also the fact that there are three generations is intriguing: it is the minimum

number that allows for a CP-violating phase to be present in the CKM matrix. In the PMNS matrix, one or three such phases are allowed, depending on whether the neutrinos are Dirac or Majorana fermions, respectively.

Another remarkable fact is that the lepton and quark charges in a generation are related, and thus that the number of generations should be the same for quarks and leptons. This is because of the need, for consistency of the theory, to cancel the chiral anomaly produced by the loop diagrams in fig. 1.3, when summing over all possible fields. The combinations that do not vanish due to symmetries give a set of constraints on the hypercharges — or, equivalently, the electric charges — of the fermion fields that are fulfilled by the charge assignments in the standard model, but in no way explained by the theory.

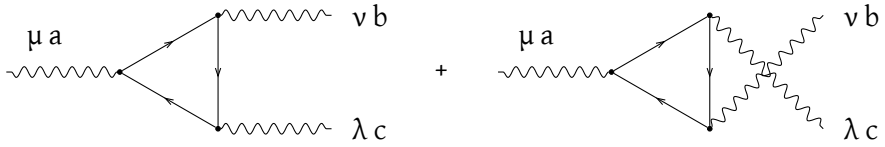

Figure 1.3: Feynman diagrams contribution to the chiral anomaly (μ, ν, λ and a, b, c are four-vector and gauge group indices). The sum over all allowed diagrams should vanish, for each combination of standard model gauge group generators, and gravity. Left-right symmetric processes cancel automatically and, due to the properties of the $\mathrm{SU(N)}$ generators, only the cases $(\mathrm{SU(2)_L}, \mathrm{SU(2)_L}, \mathrm{U(1)_Y})$, $(\mathrm{SU(3)_C}, \mathrm{SU(3)_C}, \mathrm{U(1)_Y})$, $(\mathrm{U(1)_Y}, \mathrm{U(1)_Y}, \mathrm{U(1)_Y})$ and $(\mathrm{U(1)_Y}, \mathrm{gravity}, \mathrm{gravity})$ remain, leading to the requirements $\sum_{\mathrm{L}} \mathrm{Y_L} = \sum_q \mathrm{Y}_q = \sum_f \mathrm{Y}_f = \sum_f \mathrm{Y}_f^3 = 0$, where the indices L, q and f denote sums over all left-handed fields, all quark fields and all fermion fields, respectively, *cf.* chapter 20 of [13].

The Hierarchy Problem

At the Planck scale, about 10^{19} GeV, a quantum description of gravity is required, so the standard model can only be an effective theory, valid at most up to that scale. In case no additional degrees of freedom between the electroweak and the Planck scale exist, it is considered unnatural that the scalar potential vacuum expectation value is only about 246 GeV: when running the coefficient of the quadratic term of the potential up to the Planck scale it acquires radiative corrections due to fermion loops that are proportional to $-|\lambda_f|^2 \Lambda_{\mathrm{UV}}^2 + \ldots$, where λ_f is the fermion Yukawa coupling and Λ_{UV} the ultraviolet cutoff scale [14]. The vacuum expectation value is thus very sensitive to the details of physics at the Planck scale if these are not counteracted. This is the main argument for expecting new physics to exist around the TeV scale.

Dark matter and dark energy

Some observations of the gravitational interactions within and between galaxies are not explained by the measured distribution of visible matter. Also in big bang cosmology, more matter is required to describe the evolution of the universe than is observed, and the observed structure of visible matter indicates that this component should be non-relativistic. Both arguments support the existence of dark matter, an invisible component of the universe that interacts only gravitationally and constitutes a contribution to the content of the universe that is about five times larger than the contribution from baryonic matter. The nature of dark matter is however not known. One of the still viable options is a new particle that is stable, massive, electrically neutral and does not or only weakly interact with the particles of the standard model.

The standard model of cosmology also assumes a cosmological constant term in the evolution equations of the universe, which corresponds to a component with constant energy per unit volume. Its nature remains an open question.

Matter-antimatter asymmetry

Matter and antimatter are generally assumed to have been produced in equal amounts in the early universe, and the currently observed asymmetry to be produced because of their imperfect annihilation. Three conditions need to be met for a net amount of baryons to be formed: an interaction that violates the discrete C and CP symmetries, an interaction that violates baryon number and an out of thermal equilibrium state. The amount of CP violation present in the quark sector of the standard model is however not sufficient — see *e.g.* [15] for recent reviews — so additional sources of CP violation are expected from physics beyond the standard model, and possibly also in the neutrino sector.

The strong CP problem

Due to the complicated structure of the QCD vacuum, an effective term that violates CP symmetry is included in the Lagrangian. Only a specific choice for a parameter value, with a precision of at least 10^{-10}, can remove this effect. This is known as the strong CP problem. A possible solution is the introduction of an additional $\mathrm{U}(1)$ symmetry that is spontaneously broken. The associated pseudo-Nambu-Goldstone boson is called the axion. It is a light particle that interacts electromagnetically through a vertex with two photons, which has been searched for experimentally, but that has not been observed yet.

1.1.2 Proposed extensions of the standard model

In order to solve the problems mentioned in the previous paragraphs, many possible extensions of the standard model have been proposed. Various approaches are pursued, *e.g.* adding weak interactions and particles at low energy, adding a strongly interacting, more fundamental theory, of which the standard model particles are bound states, or postulating new symmetries or different properties of nature at high energy in order to arrive at a more natural description of nature as we observe it.

The following paragraphs list some of the ingredients of these theories. Concrete examples of predictions of new particles with a long lifetime are given in section 1.2. A more complete overview of the diverse landscape of the proposed models for physics beyond the standard model can be found in the comprehensive set of reviews provided by the Particle Data Group [3].

The seesaw mechanism

The seesaw mechanism explains the smallness of the neutrino masses by introducing degrees of freedom at a mass scale much higher than the electroweak scale, that do not couple to the standard model gauge interactions and that mix with the neutrinos. By diagonalising the mixing matrix, with a mixing term at the electroweak scale and a zero Majorana mass term for the isospin $+\frac{1}{2}$ neutrinos, a small mass for the known neutrinos is obtained, proportional to the electroweak scale and the ratio between the electroweak and the high mass scale.

Grand unification

Grand unified theories try to explain some of the unexplained features of the standard model, *e.g.* the cancellation of anomalies, with symmetries of a theory at higher energy that unifies the strong and electroweak interaction. The fermions in one generation then form one representation of the larger gauge group, *e.g.* $\mathrm{SU}(5)$ or $\mathrm{SO}(10)$ or even larger symmetry groups. These theories are strongly constrained by the experimental measurements of electroweak precision observables and proton decay. Still viable scenarios require more complicated constructions, or the combination with supersymmetry at low energy.

Extra dimensions

The picture of the unification of gravity with the interactions described by the standard model could be drastically altered if spacetime were to have more than three

spatial dimensions. If only gravity propagates in the extra dimensions, for example, this could explain the difference in the observed strength of the interactions, lower the unification scale, and solve the hierarchy problem.

Supersymmetry

Supersymmetric extensions of the standard model solve the hierarchy problem in an elegant way by introducing scalar superpartner particles for the fermions and fermion superpartner particles for the bosons in the standard model. The quadratic term of the radiative correction to the quadratic term in the Brout-Englert-Higgs potential has a different sign for fermions and bosons, such that no net correction is produced if the symmetry between fermions and bosons is exact. Such a procedure is natural when building a quantum field theory that is also invariant under transformations with fermionic generators, in addition to the Poincaré group of spacetime transformations and the bosonic gauge symmetry generators. Standard model and superpartner fields then together form superfields, which form representations of the enlarged symmetry group. This also means that they have the same quantum numbers under the standard model gauge groups, and the same mass.

Supersymmetry needs to be a broken symmetry, however, because no additional light particles are observed, but it needs to be broken in such a way that no additional quadratic divergences arise. Different mechanisms to achieve this exist: typically supersymmetry is spontaneously broken in a hidden sector, and mediated to the visible sector through gravitational interactions, gauge interactions coupling to both the visible and the hidden sector or symmetry-breaking quantum loop effects.

In supergravity models, supersymmetry is a local symmetry. Gravitational interactions are then included in the theory through couplings to the massless spin-2 graviton. Its superpartner, the spin-$\frac{3}{2}$ gravitino, acquires mass by absorbing the goldstino associated to the spontaneous breaking of supersymmetry.

Due to the large number of additional fields, the phenomenology of supersymmetric extensions of the standard model is very rich. Typically a number of constraints are applied in order to obtain a testable model with manageable parameter space: no flavour-changing or CP-violating processes in addition to those described by the CKM and PMNS matrices, unification of the gauge couplings at a high energy scale, where many superpartner masses are also taken to have the same value *etc.* A second scalar doublet is needed to give masses to all fermions, such that the Higgs sector is extended with four particles.

Usually also R-parity is imposed, *i.e.* the conservation of $R = (-1)^{3B+L+2S}$, which

forbids the lepton number and baryon number violating terms

$$W_{RPV} = \frac{1}{2}\lambda_{ijk} L_L^i \cdot L_L^j E_R^k + \lambda'_{ijk} L_L^i \cdot Q_L^j D_R^k + \frac{1}{2}\lambda''_{ijk} U_R^i D_R^j D_R^k - \kappa_i L_L^i \cdot H_U, \tag{1.2}$$

that are also renormalisable and consistent with Lorentz and gauge invariance, in the superpotential. In the above expression, L_L^i, Q_L^i, E_R^i, D_R^i and U_R^i are the left-handed lepton and quark superfield doublets and the right-handed electron, down-type and up-type quark superfield singlets, respectively, and H_U is the scalar superfield doublet coupling to up-type fermions, all containing the standard model fields and their superpartners. Standard model fields have an R-parity quantum number of 1 and superpartner fields -1. This has two important consequences: superpartners are always produced in pairs, a feature exploited by most direct searches, and the lightest superpartner is stable, which is necessary for it to be a good candidate dark matter particle. The latter also provides an experimental signature in the form of missing momentum.

1.2 Long-lived particles in theories of New Physics

Hidden valleys and hidden sectors

Long-lived particles emerge naturally in models for physics beyond the standard model that feature a hidden sector: a set of matter fields and interactions that is only weakly coupled to the standard model, due to *e.g.* the high mass of the messengers that couple to both sectors, a small mixing of visible and hidden gauge bosons, or suppressed interactions of the hidden sector with the Brout-Englert-Higgs scalar field *etc.*

The distinguishing feature of hidden valley theories is the decay of one or more hidden states to standard model particles [16]. Such a phenomenology may be realized in supersymmetric extensions of the standard model [17, 18], and could explain dark matter. A simple hidden valley theory can be constructed by adding a confining gauge interaction $SU(N)$ — similar to QCD — that does not couple to the standard model fields, and a number of standard model singlet fermion fields that carry "valley" or *v*-charge: the *v*-quarks. A *v*-hadron spectrum is then present, with details depending on the *v*-quark content, but in most simple cases one or more light states are found *e.g.* a *v*-pion triplet in the case of two light *v*-quarks. The valley fields can be coupled to the standard model by a $U(1)$ symmetry, spontaneously broken by a scalar ϕ_v field, such that the mediating Z' is heavy. As decays of the lightest *v*-hadrons to standard model particles can only proceed through Z' exchange, they are heavily suppressed in case the Z' is much heavier than the *v*-hadrons. Decays

of the lightest v-hadron, which is scalar for all spectra studied in [16], into heavy fermions are preferred because of the smaller helicity suppression.

The most stringent constraints on such a model come from the measurements of the invisible decay width of the Z boson, which forbids new particles lighter than half the Z mass, unless they have very small couplings. Big-bang nucleosynthesis requires the lifetime of at least one v-hadron to be well below 1 s, or the scale of the hidden valley to be larger than 1 GeV [16].

Depending on the details of the model, v-hadrons may be produced in varying multiplicities, in the decay of the Z′ or in decays of the Brout-Englert-Higgs boson, through mixing with the ϕ_v field [19], or even as v-jets if the mediator is much heavier than the valley degrees of freedom, which might give distinctive signatures if the v-hadrons are also long-lived [20, 21].

A benchmark model adopted by all experimental searches is the decay of a scalar resonance, *e.g.* the standard model Brout-Englert-Higgs boson, to two identical v-hadrons $H^0 \to \pi_v\pi_v$, where each v-hadron decays to heavy quarks. Many searches specifically exploit the fact that two π_v decays are produced, and should be found approximately back-to-back in the detector, to a certain extent. An overview of the experimental searches targeting this model is given in section 1.4.

For the specific case where the Brout-Englert-Higgs boson decays to neutral long-lived particles, also the constraints on its invisible decays obtained by the CMS [22] and ATLAS [23] collaborations are relevant, both excluding a branching fractions of larger than about 25–30 % to undetectable particles with the full LHC run 1 data set.

Light hidden sectors — at the GeV-scale or below — have received a lot of attention in view of a number of possible dark matter detection signals. The PAMELA [25], ATIC [24], FERMI-LAT [25] and AMS [25] experiments found excesses of electron and positron yields at high energy. An X-ray line at 511 keV from the galactic centre was found in the INTEGRAL data [26] and more recently an X-ray line at 3.55 GeV from galaxy clusters was found [27, 28]. All of these might be due to dark matter annihilation, or be explained by additional astrophysical sources. Possible signals of direct dark matter detection have been reported by the DAMA/LIBRA [29], CoGeNT [30], CRESST [31] and CDMS [32] experiments, most of which are however incompatible with other results.

In the case of secluded dark matter, where the dark matter particle is coupled to the standard model through a $U(1)$ gauge interaction with a vector boson that is lighter than the WIMP particle itself, the mediator may be long-lived [33]. [34] and [35] studied the phenomenology of this model at B-factories and in fixed-target experiments, respectively. The case of a GeV-scale mediator and particle degrees of freedom at the same scale was also studied, and found to give rise to long-lived particles [36], as well as general hidden $U(1)$ bosons that mix with the photon [37] or

also couple to the Brout-Englert-Higgs boson [38]. Various supersymmetric models with a gauge interaction at the GeV scale have also been studied, see *e.g.* [39–43], and often lead to detectable long-lived decays. A possible spectacular signature of such models are lepton jets [44, 45], collimated clusters of leptons, some of which may be produced at displaced positions.

Alternative models for dark matter production, *e.g.* asymmetric dark matter [46], where the baryon asymmetry and the dark matter component are due to the same mechanism, or freeze-in dark matter [47, 48], where dark matter particles are produced by the decay of a heavier particle, may also include long-lived particles that produce a displaced vertex signature at a collider experiment [49, 50].

Long-lived superpartner decays

Supersymmetric extensions of the standard model can include long-lived particles in several ways. One possibility is the breaking of the R-parity symmetry. In many cases, the LSP is then no longer a dark matter candidate, but if R-parity violation is sufficiently small, it may be long-lived. R-parity can be broken in three ways: spontaneously, through an R-parity odd field that acquires a vacuum expectation value, explicitly through non-zero values for the bilinear couplings in the last term of eq. (1.2), or explicitly through non-vanishing bilinear and trilinear couplings[1] [51]. The values of R-parity violating couplings however are severely constrained, *e.g.* from their contribution to the decay of the proton, electric and magnetic dipole moments, flavour changing neutral currents in hadron decays, neutrino properties *etc.* [51, 52].

In case trilinear couplings are present, a sfermion LSP can decay to a pair of standard model fermions and a gaugino or higgsino LSP can decay through exchange of a virtual sfermion to a three-body final state of three leptons, a lepton and two quarks or three quarks, depending on the structure of the R-parity violating couplings. The last case is not possible in the minimal scenario of bilinear R-parity violation, which is discussed below, and is particularly interesting because there is no large invisible transverse momentum from the LSP, which reduces the sensitivity of standard superpartner searches. The possibility of light long-lived neutralinos decaying to three quarks [53], possibly pair-produced in the decay of a scalar boson [54], and the potential sensitivity of LHCb to such decays [55] has been studied in the literature. Bounds on the λ'' couplings from baryogenesis require LSP lifetimes that lead to displaced vertices or undetectable particles at the LHC [56].

[1] non-zero values for trilinear couplings only are not possible, as they induce contributions to the bilinear couplings through loop diagrams

A special case that has been studied in detail is bilinear R-parity violation, where only the bilinear couplings in the last term of eq. (1.2) are nonzero. These can be rotated away by a redefinition of the H_D doublet, due to the presence of the $\mu H_U \cdot H_D$ term in the MSSM scalar potential, but this generates lepton-number violating trilinear terms, of the same form as the first two trilinear terms in eq. (1.2), from the down-type Yukawa terms. The possibility of a long-lived LSP has been studied in the context of minimal supergravity [57], models where supersymmetry breaking is dominated by anomaly contributions [58, 59] and the μ-from-ν supersymmetric standard model [60, 61] in respectively [62, 63], [64] and [65], and in a more signature-based way in [54, 66].

In supergravity theories where mediation through gauge interactions dominates supersymmetry breaking, and where the supersymmetry breaking scale is thus well below the Planck scale, the effect of additional flavour violation at the Planck scale on the soft supersymmetry-breaking terms is suppressed by the ratio of the supersymmetry breaking scale and the Planck scale. As the contribution from gravity mediation is quantified by the gravitino mass, this also means that the gravitino should be the lightest superpartner. The gravitino is then a good candidate for dark matter, and decays of the next-to-lightest superpartner to it may be suppressed, such that the NLSP is long-lived [67]. These features were shown to be independent of the details of the gauge mediation mechanism [68]. The compatibility of observed light element yields with the model of big-bang nucleosynthesis gives constraints on the properties of a stable or long-lived gravitino and a long-lived NLSP [69] — a very long-lived NLSP could however be avoided by allowing R-parity violating decays, which would still keep the gravitino sufficiently long-lived [70]. The collider phenomenology in this class of scenarios has been studied for different fields dominating the NLSP, and found to give displaced vertices for the (in minimal models of gauge mediation typically bino-dominated) neutralino [65, 71, 72], stau [73–75], wino- and higgsino-dominated neutralino [76, 77] and stop [78] cases.

Supersymmetric extensions of the standard model with a split spectrum, *i.e.* with sfermion masses at a much higher scale than the gaugino masses, may also give rise to heavily suppressed, and thus long-lived, gaugino decays, possibly in the lifetime range that gives a displaced vertex [79, 80].

Overviews of the bounds on long-lived superpartner decays with hadronic final states can be found in [81], and with a single lepton from displaced decay in [82]. [83] gives an overview of the constraints on R-parity violating superpartner decays.

1.3 Production of long-lived particles in hadron collisions

Depending on the details of the model they are embedded in, exotic long-lived particles may be produced through a variety of mechanisms, leading to different kinematic distributions. This can have a large impact on the sensitivity of an analysis with the LHCb detector, because the forward geometry requires the long-lived particles to be produced in the forward direction and with a sufficiently large forward boost in order to be able to detect all the decay products.

Many features of high-energy hadron collisions can more easily be understood by exploiting the fact that the collision process between the proton constituents happens at such a high energy (or short distance) scale, that the partons are essentially free, and the cross-section can be written as

$$\sigma(pp \to \mathrm{X}) = \sum_{i,j} \int \mathrm{d}x_1 \, \mathrm{d}x_2 \, f_i(x_1) \, f_j(x_2) \, \mathrm{d}\hat{\sigma}_{i,j}, \tag{1.3}$$

where i, j indicate the different partons, x_1, x_2 the fraction of the hadron momentum each of them carries, and $f_i(x)$ and $\hat{\sigma}_{i,j}$ the corresponding parton distribution functions and parton-level cross-sections, respectively. If the cross-section for a process has a qualitatively simple behaviour, *e.g.* it is dominated by a single resonance, the main features of the kinematic distributions can be inferred from the hadron structure. The parton distribution functions are extracted from measurements of benchmark processes, where the final state gives sufficient information about the parton-level kinematic variables and the parton-level cross-sections are precisely known, *e.g.* inelastic scattering with a lepton beam, or Drell-Yan production of heavy electroweak vector bosons in hadron collisions. Figure 1.4 shows the parton distribution functions for the proton at two different energy scales.

If long-lived particles are produced in the cascade decay of heavy resonances, *e.g.* TeV-scale superpartners, the di-superpartner system would be approximately at rest in the laboratory reference frame, and the angular distribution of the long-lived particles approximately isotropic. If, on the other hand, long-lived particles are light and no heavy scales are involved in the production process, the initial state will be dominated by low-x partons, typically asymmetric due to the sharp rise of parton distribution functions, and they will be mostly found in the forward and backward direction.

The benchmark production model adopted by most experimental searches is the production of long-lived particles in two-body decays of a Brout-Englert-Higgs boson. At the LHC, such a particle is sufficiently light for the production to be dominated by gluon fusion processes. The perturbative calculation of the cross-section

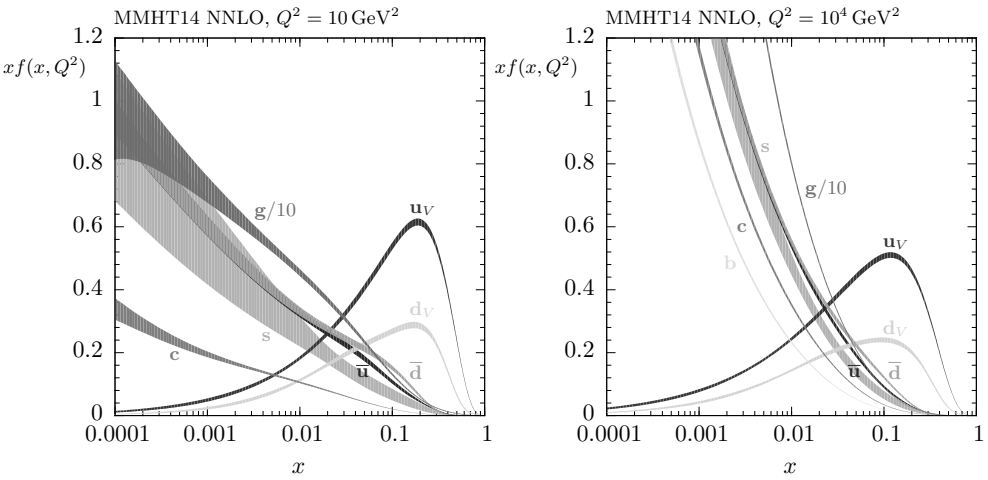

Figure 1.4: Illustration of the behaviour of the parton distributions at two different energy scales, $Q^2 = 10\,\mathrm{GeV}^2$ (left) and $Q^2 = 10^4\,\mathrm{GeV}^2$ (right). u_v and d_v represent the valence quark contributions, while the other curves show the sea quark and gluon contributions. Figure from MMHT14 [84].

and kinematic distributions is complicated due to large corrections from higher-order diagrams and soft gluon emissions, which also affect the kinematic distributions [85], *cf.* fig. 1.5. Currently, cross-section calculations with three-loop perturbative corrections are available [86]. Figure 1.6 shows the transverse and longitudinal momentum spectra obtained using the standard LHCb simulation, with the PYTHIA generator and the CTEQ6 leading order parton distribution function fit.

The distribution of the π_v kinematic variables in fig. 1.7 shows mainly the effect of the longitudinal boost of the resonance and the phase space available for its two-body decay: the p_T of the π_v particles is larger and their longitudinal boost follows the resonance's less closely for lower π_v mass. Only for longitudinal π_v boosts larger than about 3 one expects the decay products to be all found inside the LHCb acceptance of $\theta < 0.3$ rad.

The sensitivity of a search for dijet decays in the forward region with a minimum and maximum transverse displacement requirement is affected by the p_T distribution in two ways: higher (transverse) momentum of the long-lived particles increases the efficiency for short lifetimes, following

$$\mathrm{FD} = \gamma\beta ct = \frac{p}{m}t \qquad \text{and} \qquad R_{xy} = \gamma\beta ct \sin\theta = \frac{p_T}{m}t, \tag{1.4}$$

where t is the decay time of a long-lived particle, FD the total flight distance between the production and decay point, R_{xy} the corresponding transverse distance and β

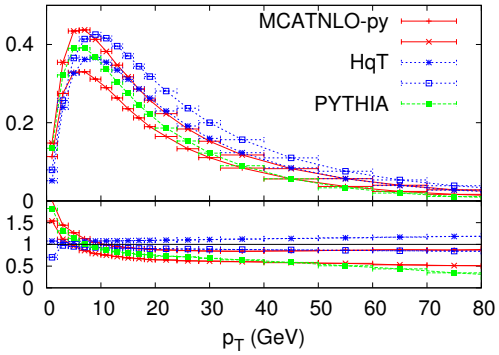

Figure 1.5: p_T spectrum of a Brout-Englert-Higgs boson with mass 120 GeV/c^2 produced in proton-proton collisions at 7 TeV, as calculated with NNLO+NNLL precision (HqT), using a NLO+PS approach (MCATNLO-py) and using a LO+PS approach (PYTHIA), including uncertainty bands and the ratio with respect to HqT, figure from [85]

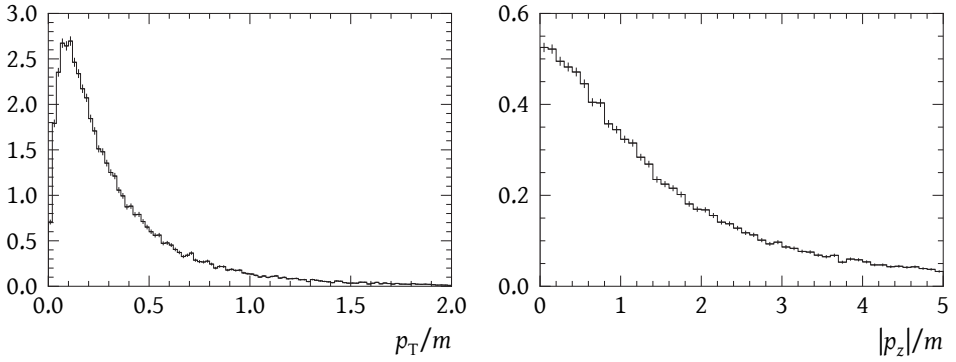

Figure 1.6: Basic kinematic distributions for the production of a standard model-like Brout-Englert-Higgs boson with a mass m of 125 GeV/c^2 through the gluon fusion process in 8 TeV pp collisions, generated using PYTHIA8 as described above

and γ, defined using the particle's speed as $\beta = v/c$ and $\gamma = 1/\sqrt{1-\beta^2}$, describe the Lorentz boost with respect to the laboratory frame. For large boosts — or equivalently, light long-lived particles, if produced in decay of the same resonance — however, the jets can be too closely together to be separately reconstructed, and a specialised analysis, possibly using the reconstruction of a single large jet and jet substructure techniques, is in order to target this signature.

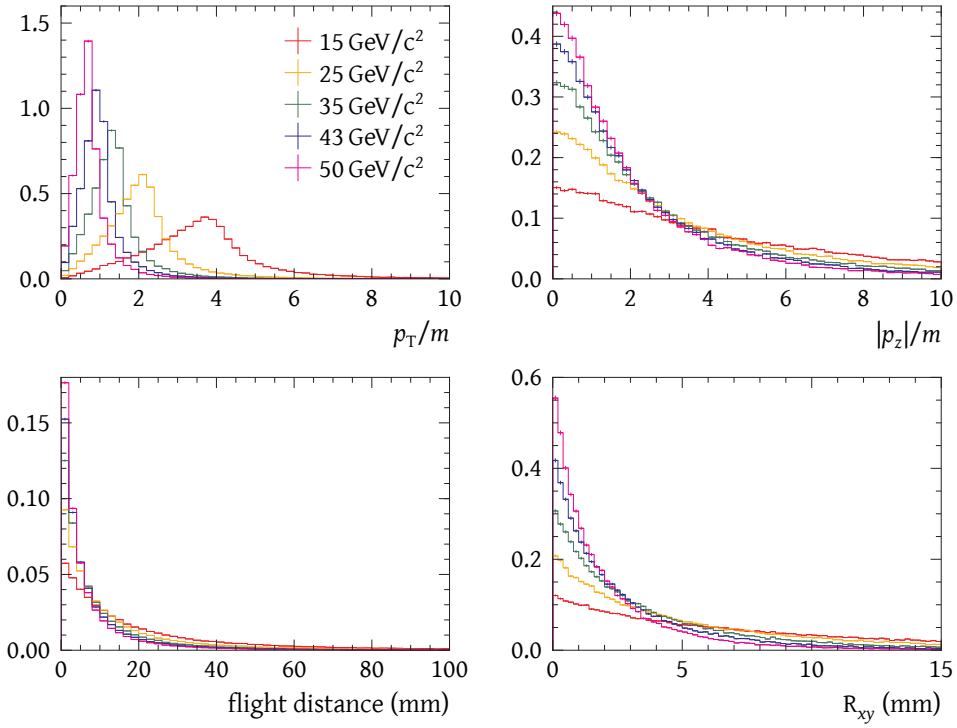

Figure 1.7: Kinematic and lifetime-related distributions for π_v particles with a lifetime of 10 ps produced in the decay of a standard model-like Brout-Englert-Higgs boson, for various π_v masses

1.4 Searches for exotic long-lived particles

Due to the diversity of the possible production mechanisms and the large range of lifetimes, long-lived particles can be searched for in a many ways: the recoil of dark matter particles scattering off nuclei, missing momentum when produced in high-energy particle collisions, annihilation in astrophysical systems *etc.* For short

lifetimes, where only the decay products can be observed, production in a fixed-target or colliding beam setup can be studied, where the long-lived particles can be produced directly or in the decay cascade of known or unknown, short-lived or long-lived, particles. The discussion in this section focuses on searches for particles produced in hadron collisions decaying to hadronic final states, and within the detector volume of a collider experiment, typically up to 20 m, as these can be directly compared to the analysis described in this thesis.

Long-lived particle searches at the Tevatron

The first such search was carried out by the D0 collaboration [87], using 3.6 fb^{-1} of $\mathrm{p\bar{p}}$ collisions at a centre-of-mass energy of 1.96 TeV at the Tevatron collider. It assumed a scalar boson with a mass of 90–200 GeV/c^2 produced through the gluon fusion mechanism decaying to two π_v particles with a mass in the range of 15–40 GeV/c^2, which in turn decay into $\mathrm{b\bar{b}}$ quark pairs. The trigger relied on a single high-p_T muon from the decay of one of the b quarks, which should also be contained inside one of the two jets required in the offline selection. Displaced vertices were reconstructed using the tracking detectors, in a radius between 1.6 cm and 20 cm from the beam line, and required to have at least four associated tracks. In addition, either the total momentum of the displaced vertex reconstructed from the associated tracks should be compatible with its flight direction, or these track momenta should constitute a large invariant mass, depending on whether the assumed π_v is above or below 20 GeV/c^2, respectively. Upper limits of a few pb, about 1–10 times the total standard model Brout-Englert-Higgs boson production cross-section when assuming 100 % branching ratios to the final state under study, on $\sigma(\mathrm{p\bar{p}} \to \mathrm{HX}) \times \mathcal{B}(\mathrm{H} \to \pi_v\pi_v) \times \mathcal{B}(\pi_v \to b\bar{b})^2$, were obtained for proper decay lengths in the cm range. The strongest exclusion was obtained for the lowest π_v masses, likely because the π_v particles have a larger transverse momentum in that case, which reduces the efficiency loss due to the $R_{xy} > 1.6$ cm requirement.

The CDF collaboration performed a search [88] using a similar data set as D0, and for a similar set of signal model parameters: scalar resonance masses of 130 GeV/c^2 and 170 GeV/c^2, π_v masses of 20–65 GeV/c^2 and lifetimes of about 1 cm. At the trigger level, CDF had the advantage of a displaced track and vertex trigger, which improved the efficiency to detect π_v decays to other final states. The complete trigger selection relied on displaced tracks, calorimeter deposits in the central region, and a veto on jets in the forward plug calorimeter. At least three jets in the central detector were required for the signal region, two of which with a secondary vertex, reconstructed using a jet flavour tagging algorithm. Final selections were made on the displacement of the jet axes from the collision point, the decay lengths, the opening

angle between the jets, and the separation of the secondary vertices or their combined track invariant mass. The upper limits on $\sigma(p\bar{p} \to HX) \times \mathcal{B}(H \to \pi_v\pi_v) \times \mathcal{B}(\pi_v \to b\bar{b})^2$ obtained vary between about five and a few tens of pb.

Long-lived particle searches by the general-purpose LHC experiments

At the LHC, searches for exotic long-lived particles decaying in the detector volume were also carried out by the ATLAS and the CMS collaboration. The ATLAS searches can be divided in two groups: searches for long-lived superpartner decays using standard trigger selections based on high-p_T objects and missing transverse momentum, and more signature-driven searches using dedicated triggers. In the first category, ATLAS has searched for superpartner decays with a muon in the final state [89, 90] in 7 TeV pp collisions, triggered by a segment reconstructed in the muon spectrometer, where the decay takes place inside the barrel pixel detector fiducial region, with $|z| < 300$ mm and $40 < R_{xy} < 180$ mm. The displaced vertex was required to have an invariant mass above 10 GeV/c^2 and at least four associated tracks. In the analysis of 2010 data [89], the high-p_T muon was not required to come from the displaced vertex, in contrast to the analysis of 2011 data [90], where also the muon trigger criteria were tighter and the muon required to be well reconstructed in the inner detector and compatible with coming from the displaced vertex. The vertex reconstruction at large displacements was improved by a so-called retracking procedure, which increased the track reconstruction efficiency at large displacement. The results were interpreted in terms of the R-parity violating decay of the lightest superpartner, a neutralino from the decay of a pair-produced squark, through the λ'_{211} coupling $\tilde{\chi}^0 \to \mu ud$. Upper limits down to 0.5 pb and 0.002 pb were obtained, using the 2010 and 2011 data, respectively, with the best sensitivity for a proper lifetime of $\mathcal{O}(1\,\text{cm})$.

The related analysis of 2012 data [91] used a broader range of trigger and offline signatures to also study long-lived particle decays with a high-p_T electron and decays in events with multiple jets or a large amount of missing transverse momentum. Neutralino decays through the R-parity violating couplings λ_{ijk} with a lepton pair in the final state, long-lived NLSP decays to a massive gauge boson and a light LSP in general gauge-mediated supersymmetry breaking scenarios, and R-hadron decays into the LSP and two quarks suppressed by the high squark masses in split supersymmetry, were added to the searched-for signals. The R_{xy} acceptance was extended to 300 mm. The obtained upper limits go down to about 1 fb for all channels.

The second group of ATLAS searches employs dedicated trigger selections for long-lived particle decays. One of these triggers is based on the presence of a cluster

of tracks in the muon spectrometer, and is sensitive to long-lived particle decays at the end of the hadronic calorimeter and inside the muon spectrometer, from about 4 to 7 m of transverse displacement from the interaction region. In the search using 2011 data [92] the Higgs to Hidden Valley scenario was assumed, and thus two such vertices required to be reconstructed, with $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} > 2.2$, which makes this search sensitive to proper lifetimes of the order of 1 m. Upper limits on the production down to about 1.7 pb were obtained, for scalar boson masses in the 120–140 GeV/c^2 and π_v masses in the 20–40 GeV/c^2 range.

The related search using 2012 data [93] added the possibility of reconstructing one of the long-lived particles in the inner detector for the Higgs to Hidden Valley model. For the other Hidden Valley production mechanism with a Z′ particle decaying into a large number of π_v particles, which is used as a benchmark model triggered by a signature with jets and missing transverse momentum rather than a cluster of tracks in the muon spectrometer, two decay vertices can be reconstructed in the muon spectrometer, one in the muon spectrometer and one in the inner detector, or both in the inner detector. Cross section upper limits down to about 0.1 pb for the Higgs to Hidden Valley phenomenology were obtained, and similar values for the other models, mostly sensitive for proper lifetimes of $\mathcal{O}(1\,\mathrm{m})$, without strong dependence on the mass of the resonance and the long-lived particles.

The ATLAS collaboration added more dedicated triggers for long-lived particle decays for 2012 data taking [94]. One of these is based on the ratio of corresponding energy deposits in the electromagnetic and the hadronic calorimeter and was used to search for pair-produced π_v particles decaying between about 2 m and 3.5 m from the beam line [95]. Upper limits down to 0.5 pb, for proper lifetimes of 0.5–2 m depending on the π_v mass, were obtained.

The CMS collaboration performed a search using the signature of a single displaced vertex with two associated jets [96], using a dedicated trigger selection, which also allowed to search for decays with short lifetimes, down to the cm range. Due to the stringent p_T requirements on the jets employed there, and the requirement on the sum of jet transverse momenta in the first trigger stage, the best sensitivity, leading to cross section upper limits down to 0.001 pb, was achieved for long-lived particles produced in the decay of a heavier resonance, 200–1000 GeV/c^2, and, due to the requirement of two separated jets, with higher mass. For the kinematic range probed by the analysis described in this thesis, the sensitivity is limited, as was also observed in [97].

Searches using different final states and related signatures

Searches specifically for exotic long-lived particles decaying to lepton pairs were also performed by the D0 [98] and CMS [99] collaborations.

The search by the D0 collaboration for displaced dimuon pairs used an R-parity violating neutralino decay as benchmark signature, and obtained upper limits on the production cross-section down to a few 10^{-1} pb also for low long-lived particle masses, in the range 3–40 GeV/c^2, with best sensitivity for decay times of the order 20 ps. The displaced di-electron search performed by the D0 collaboration relied mostly on the electromagnetic calorimeter, and is thus mainly sensitive to proper lifetimes in the 10 cm range, with best exclusion limits of a few pb.

The CMS searches assumed that the long-lived particles are pair-produced in the decay of a scalar resonance, but only one was required to be reconstructed. The most stringent exclusion was obtained for proper lifetimes in the cm range, down to a few 10^{-4} pb for heavy (400–1000 GeV/c^2) resonances, and for light resonances down to about 10^{-3} pb for muon final states and 10^{-2}–10^{-3} pb for electron final states, with long-lived particle masses of at least 20 GeV/c^2.

Searches for the phenomenologically related but experimentally different signatures of a long-lived particle decaying to an invisible particle and a photon, and of a charged particle that does not decay inside the detector volume, have also been performed by several collider experiments, see *e.g.* [100] and [101–103] for recent results.

Long-lived particle searches at LHCb

As illustrated in fig. 1.8, the searches for long-lived particles decaying to hadronic final states by the ATLAS and CMS experiments require either a relatively high p_{T} (CMS) or displacement (ATLAS), mostly due to trigger constraints. The region of a lifetimes in the 10 ps and masses of a few tens of GeV/c^2, which was also probed by the Tevatron experiments, is much less constrained. The LHCb detector has a unique ability to probe this region due to the precise vertex detector and the high-level trigger system built for selecting heavy flavour decays based on mostly tracking information.

The LHCb collaboration has searched for pairs of long-lived particles produced in the decay of a resonance with the signature of two decay vertices on opposite sides of the vertex detector [1, 104]. An update of this analysis with a larger data set is in preparation [105]. Final states with quarks and a muon have also been studied, using the 2010 [106], 2011 [107] and combined 2011 and 2012 data sets [108]. Long-lived particles decaying to opposite-sign muon pairs have been searched for

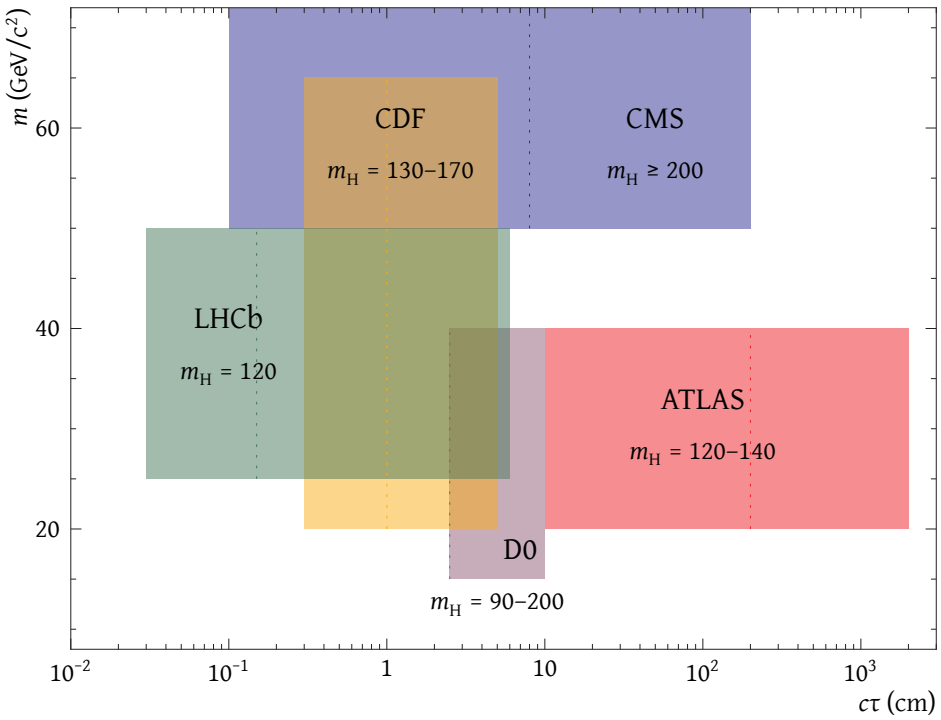

Figure 1.8: Illustration of the parameter space regions probe by different experiments, assuming pair-production of π_v particles in the decay of a Brout-Englert-Higgs boson, whose assumed mass range (in GeV/c^2) is indicated in each region. The vertical dotted lines indicate the proper lifetime where each search is the most sensitive. For a full comparison, also the sensitivity of each search and the experimental signature, one or two vertices, should be taken into account.

in B-decays [109, 110]. Dimuon decays of heavier, promptly produced, long-lived particles have also been studied [111].

For long-lived particles decaying to quarks, the reconstruction of the final-state jets can be used to improve the suppression of instrumental backgrounds. This has the advantage that the requirement of reconstructing two vertices can be dropped, which increases the acceptance for the signal model and makes the search sensitive to other production mechanisms. The lifetime reach is also enhanced. Such an analysis was developed and performed on the 2011 data set [2, 112].

For 2012 data taking, the selection efficiency was improved through the additional of a trigger selection targeted at very displaced decays — beyond 1 cm from the beam line — in the first stage of the software trigger. The dedicated selections in the second stage of the software trigger were also reoptimised. The analysis described in this thesis includes both the 2011 and 2012 data set, and also takes advantage of modifications in the offline reconstruction strategy that lead to a higher efficiency and reduced systematic uncertainties.

2

Experimental setup

The LHCb detector [113, 114] at the LHC accelerator was built for the precise study of the decays of beauty and charm hadrons produced in proton-proton collisions. Since the main features used to discriminate these heavy hadron decays from the large background of light hadron production are a higher transverse momentum and the macroscopic lifetimes of the weakly decaying hadrons, which translate into flight distances of the order of a centimetre, it is an interesting place to search for yet unknown massive long-lived particles.

2.1 The LHC accelerator

The LHC accelerator was designed to accelerate bunches of protons and nuclei up to an energy of 7 TeV per proton. Two beams circulate in opposite directions and are kept on their approximately 27 km long orbits by superconducting magnets. The protons are accelerated through rapidly oscillating electromagnetic fields inside so-called radio frequency cavities, which create buckets interspaced by 25 ns where particles are kept together. Focusing and steering magnets make bunches of protons from the two beams collide in four interaction points, where the particle detectors study the particles emerging from the collision.

As illustrated in fig. 2.1, the ATLAS and CMS experiments are located at interaction points opposite from each other in the LHC ring. They take full advantage of the beam intensity and focusing capabilities of the accelerator, collecting as many collisions as possible to search for new physics phenomena at the high-energy frontier. ALICE, a dedicated experiment for the study of the quark-gluon plasma, and LHCb, mainly focused on quark flavour physics, are both located at interaction points with different beam optics, as they do not require the maximum possible instantaneous luminosity. For LHCb, this is because with more proton-proton interactions per bunch crossing, *pileup*, the increasing particle multiplicity in the forward region makes it difficult to efficiently reconstruct heavy flavour decays. For ALICE, this is because the detector makes use of a time projection chamber with large drift volume, which integrates over a relatively long drift time, and would suffer from the same problems with pileup if inelastic collisions are produced too frequently.

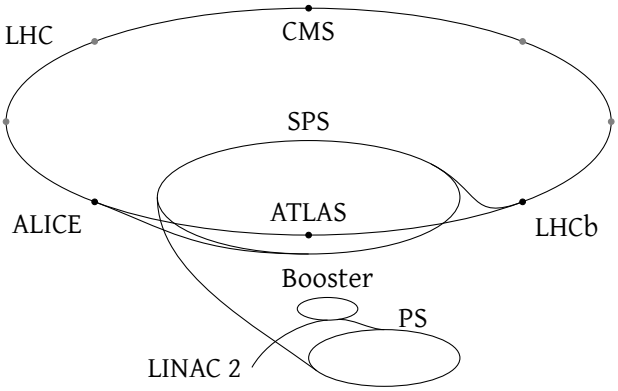

Figure 2.1: Sketch of the LHC proton injection chain and the location of the main experiments

During most of 2010, 2011 and 2012 data taking, the accelerator was operated with a 50 ns bunch filling scheme and beam energy of 3.5–4 TeV instead of the nominal values, for operational reasons. In order to compensate for the lower bunch-crossing rate and to maximise the delivered luminosity, bunch intensities were increased to higher values than originally planned [115], which lead to more pileup for the high-luminosity experiments.

The LHCb collision settings also changed with respect to the design values, from an average number μ of visible interactions per bunch crossing of 0.4, in order to maximise the number of crossings with a single collision, to about 1.5.[1] In contrast to the other interaction points, the instantaneous luminosity was kept constant, at a value of $4 \times 10^{32}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$, as is illustrated for a specific fill in fig. 2.2. This was achieved by a luminosity levelling procedure which slightly offsets the position of the beams in the transverse plane, moving them closer together as they lose intensity, until after about ten hours of circulating, when they maximally overlap in the transverse plane. The instantaneous luminosity at that point is still much smaller than for the high-luminosity interaction points due to the different focusing and number of colliding bunches. As this levelled phase covers most of the data taking, the conditions were essentially constant for readout and further processing, and the resulting data set very homogeneous.

[1] All numbers on the operational conditions and detector performance in this and the following sections are taken from [114] unless specified otherwise.

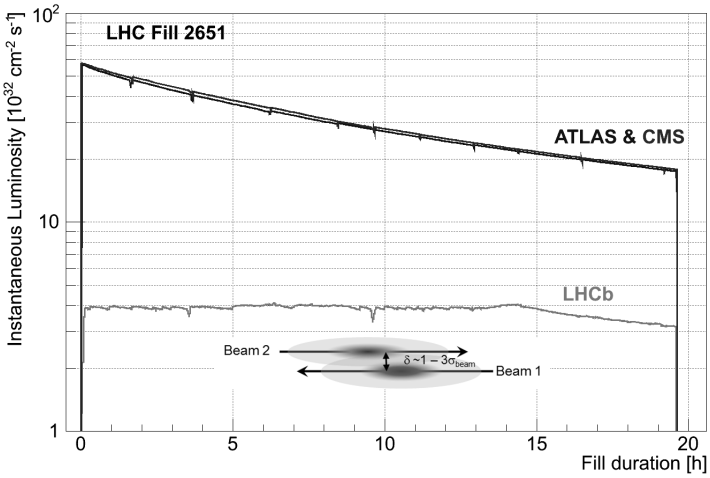

Figure 2.2: Instantaneous luminosity at the ATLAS, CMS and LHCb interaction points during LHC fill 2651, figure from [114]

2.2 The LHCb detector

The most striking characteristic of the LHCb setup is the geometry of the detector: all subdetectors cover approximately the same, relatively small, solid angle around the beam line. This is motivated by the typical kinematics of heavy flavour production at the LHC: the dominant process is gluon fusion, where a gluon from each of the protons combines to produce a $b\bar{b}$ pair through a $2 \rightarrow 2$ scattering process. Since the threshold for the production of a $b\bar{b}$ pair, $10\,\mathrm{GeV}/c^2$, is only a tiny fraction of the collision energy 7–14 TeV, the gluons only need to carry a small fraction x of the proton momentum for this process to be kinematically allowed. Due to the rapid rise of the gluon distribution function for decreasing x, the vast majority of the gluon-gluon interactions producing a $b\bar{b}$ pair have highly asymmetric incoming momenta. The $b\bar{b}$ system is then highly boosted in the forward or backward direction, and both B hadrons are produced under a small angle from one of the beam directions. It is important to have a good acceptance for both b hadrons from the same collision, because the non-signal b hadron provides information about the flavour of the signal b quark at production, which is crucial for the precise study of lifetime-dependent asymmetries and oscillation phenomena.

For the study of exotic long-lived particles, however, many different production mechanisms are possible, some of which yielding a much smaller fraction of particles inside the LHCb acceptance. The analyses described in the subsequent chapters

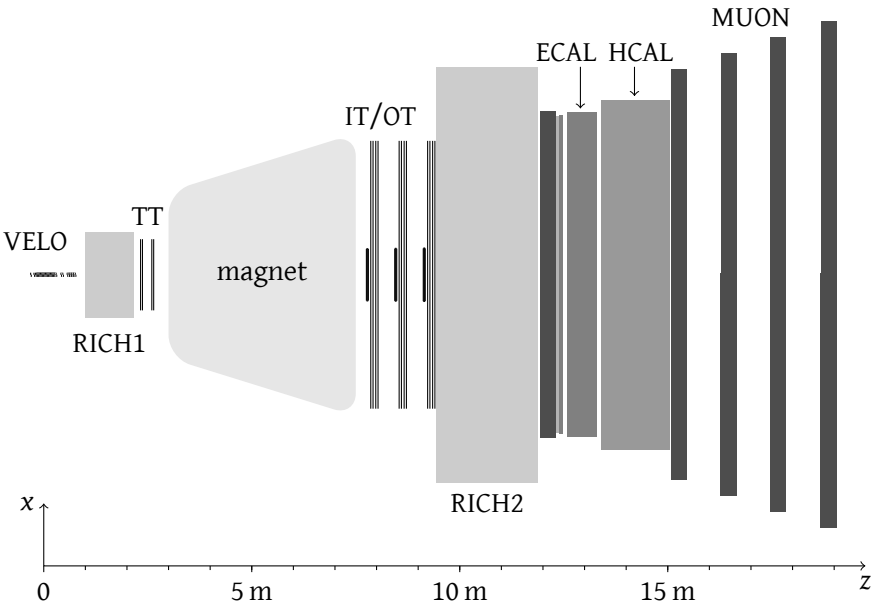

Figure 2.3: Schematic top view of the LHCb detector in the bending plane

search for the decay of one such particle, even though the benchmark models used predict them to be produced in pairs.

Overview and coordinate system

The LHCb detector can roughly be subdivided into two parts. The first is the tracking system which provides a momentum measurement for charged particles and precise vertex information. The second is the particle identification system which allows to separate the different kinds of detectable particles. It includes two ring-imaging Cherenkov detectors, a calorimeter system and muon chambers. In addition, the calorimeter system measures the energy of electrons, photons and hadrons.

At the center of the tracking system, a large dipole magnet generates a magnetic field along the vertical direction, such that charged particles are bent in the horizontal plane depending on their charge and momentum. The integrated magnetic field amounts to about 4 T m, which means that charged particles produced in the primary interaction need a minimal momentum of about 2 GeV/c to reach the detectors beyond the magnet.

The part of the detector further away from the interaction point than the mag-

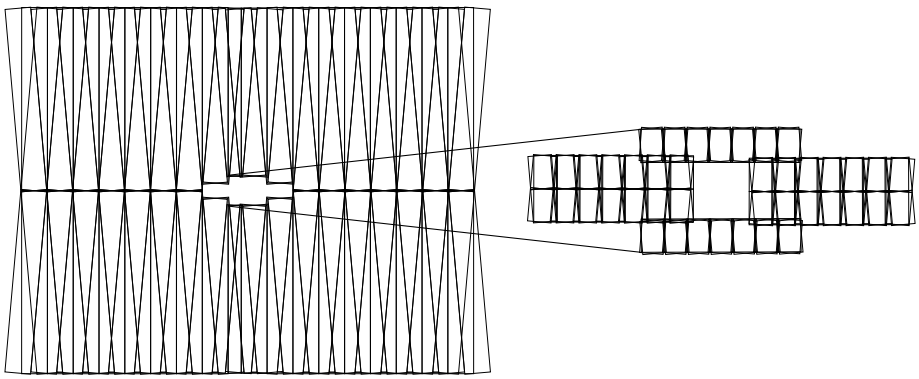

Figure 2.4: Layout of a downstream tracking station, visualized by the OT module boundaries on the left and the IT sensors in the zoomed area on the right

net is conventionally called the downstream region, as opposed to the upstream region. The direction along the beam line, pointing from the interaction point towards the magnet, provides the z-axis of the LHCb reference coordinate system. The vertical direction gives the y-axis pointing upwards. The x-axis pointing away from the center of the storage ring in the horizontal plane completes this right-handed coordinate system.

Tracking detectors

The three downstream tracking stations each consist of a fine-grained inner region, instrumented with silicon strip detectors, called the Inner Tracker (IT), and an outer region, the Outer Tracker (OT), consisting of straw drift tubes. Both measure the x-position of a particle passing through the detection layer very precisely: the OT straws, arranged in two staggered monolayers per detection plane, have a diameter of 5 mm, but the drift time measurement allows to improve the hit resolution for high-momentum tracks to about 200 µm. The IT has a strip pitch of 198 µm and reaches a hit resolution of about 50 µm. The y-segmentation, on the other hand, of these detectors is rather coarse. The OT straws are split in an upper and a lower half, each grouped in 18 modules per detection plane. The IT consists of four rectangular boxes, those left and right of the beam pipe with two rows of seven sensor modules, one above and one below the horizontal plane, while those on top of and below the beam pipe contain only one row of vertical strips. This results in a cross-shaped coverage of the highest occupancy part of the detection plane, as illustrated in fig. 2.4.

The necessary sensitivity to the y-coordinate to disentangle the sets of measure-

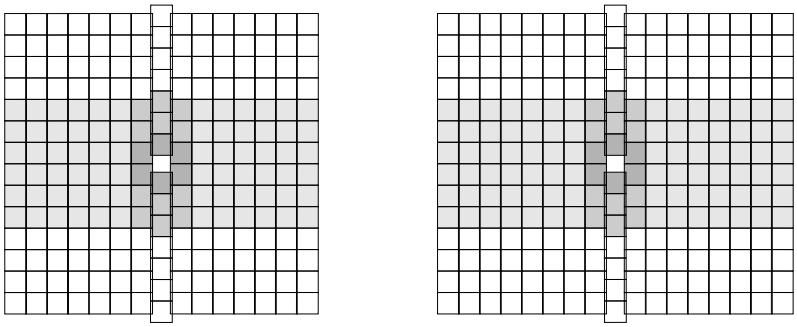

Figure 2.5: Layout of the TT tracking layers, visualized by the TT sensors boundaries in the TTa (left) and TTb (right) x-layers. The shaded areas indicate the different readout sectors.

ments corresponding to the trajectories of different particles is provided by having, for both IT and OT, four z-layers per detection station: while the outer layers have vertical straws or strips, the modules of the inner two layers are rotated in opposite direction over a 5° angle around the z-axis.

The tracking station upstream of the magnet, called Tracker Turicensis (TT), has the same layout with four stereo layers, grouped two by two (TTa and TTb) with a 30 cm gap in z between the pairs. The same type of silicon sensors as for the IT is used, but now covering the whole fiducial area. In TTa (TTb) each layer has seven (eight) full vertical modules with fourteen sensors on either side of the beam pipe and two half modules with seven sensors each in between, one above and one below the beam pipe, as illustrated in fig. 2.5. The strip pitch and single hit resolution are comparable to those for the IT.

The vertex detector (VELO) is instrumental for the precise reconstruction of displaced vertices by providing the first track segment with very precise position information. It encloses the interaction region and is made out of two halves, both consisting of a series of stations that each provide measurements of the radial and azimuthal coordinate of particle hits by means of a sensor with the appropriate strip layout for each coordinate, which are called R and Φ sensors, respectively. The R strips are further split into four 45° sectors. This setup minimises the amount of detector material that particles have to traverse, because the readout electronics and cooling elements can be kept outside the acceptance. The need to correctly combine R and Φ hit clusters when reconstructing tracks, on the other hand, comes with a computational cost and increases the ghost rate at high occupancy, when not all ambiguities can be resolved. Each sensor has the shape of a half disc, with inner radius 8 mm and outer radius 42 mm. The strip pitches vary from about 40 µm in the innermost region to about 100 µm, which leads to hit resolutions between 5 µm and

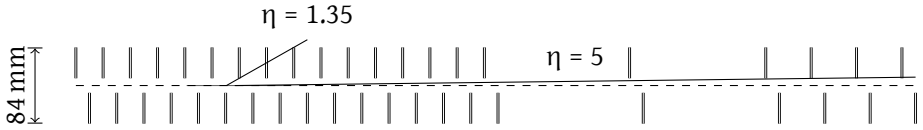

Figure 2.6: Cross-section of the VELO sensors in the horizontal plane, or almost any other plane containing the beam axis. In the region around the vertical plane where the two detector halves overlap, this picture should be overlayed with its mirror image around the beam line.

25 µm depending on the angle between the traversing track and the sensor.

The stations of the two halves are placed in staggered positions along the z-axis, more densely around the interaction region and more sparsely downstream. Each halve is enclosed by a thin metal box to separate the detector and the accelerator vacuum, as well as to protect the silicon sensors from the strong electromagnetic fields induced by the circulating proton beams. This so-called RF foil has a complicated corrugated shape in order to minimise the median amount of material crossed by a particle inside the full detector acceptance.

Particle identification detectors

Two ring-imaging Cherenkov (RICH) detectors are situated in between the VELO and TT and behind the downstream tracking stations. These exploit the Cherenkov effect where a charged particle propagating through a medium at a speed larger than the speed of light in that medium, in this case a gas mixture engineered to cover the desired velocity range, emits light in a cone with an angle depending on its velocity. The mirror system inside the detector maps this cone onto a ring whose radius is related to the cone opening angle. By combining this velocity measurement with the momentum measurement from the tracking system, the mass of a charged particle can be determined and pions disentangled from protons and kaons. RICH1 uses aerogel and C_4F_{10} as radiators to cover the low momentum range from about 2 to 60 GeV/c, while RICH2 uses a CF_4 radiator and is sensitive in the momentum range from about 15 to 100 GeV/c.

Located downstream of the RICH2 detector are the muon and calorimeter systems. The muon system consists of five stations, the first of which is placed upstream of the calorimeter system and the other four downstream, interleaved with layers of absorbing material. Each station is divided into four regions, segmented into pads with increasing granularity towards the beam line, using MWPC detector technology except for the inner region of M1, the most upstream station, where

GEM detectors are used, because these are less sensitive to radiation damage. The muon system is used in the hardware trigger as well as for the identification of muons further on in the reconstruction.

The calorimeter identifies photons, electrons and hadrons, and provides a measurement of their energy and position. A detection layer of scintillator pads (SPD) is placed in front of the electromagnetic calorimeter to identify charged particles. The electromagnetic calorimeter (ECAL) consists of two parts: first a lead converter layer and a scintillator pad detection layer, the preshower detector (PS), followed by the main part of the ECAL. The addition of the converter layer and PS allows to gather more information about the longitudinal development of showers and thus improves the separation between hadronic and electromagnetic showers.

The ECAL detector itself, which is focused on reaching a good energy resolution for electrons and photons from heavy flavour decays, uses the шашлык (*shashlik*) technology and consists of alternating layers of iron and scintillating tiles, with wavelength-shifting fibres going through the stack to collect light and guide it to photodetectors. The energy resolution was determined in a test beam and, using the parameterization $\sigma_E/E = a/\sqrt{E} \oplus b \oplus c/E$, where a is the stochastic, b the systematic and c the noise contribution, with E in GeV, measured to be $8.5\,\% < a < 9\,\%$, $b \sim 0.8\,\%$ and c at the level of a few per mille.

The hadronic calorimeter (HCAL) is also a sampling device made out of iron and scintillating tiles, with the particularity that these are oriented parallel to the beam axis: each module consists of a repetition of layers where the position of scintillator tiles and steel are staggered between consecutive layers. An energy resolution of $\sigma_E/E = (69 \pm 5)\,\%/\sqrt{E} \oplus (9 \pm 2)\,\%$, with E in GeV, is achieved [116].

The SPD, PS and ECAL detectors are segmented in three regions with different granularity, amounting to about 6000 cells per subdetector. The HCAL detector has a much coarser granularity, with about 1500 cells in two regions, as illustrated in fig. 2.7.

Readout and trigger

The LHC produces collisions with a bunch crossing rate up to 40 MHz — during the first run typically 10 MHz–15 MHz at the LHCb interaction point. The complete detector, however, can only be read out at a rate of about 1 MHz. Therefore the first stage of trigger system, L0, implemented using custom hardware, accepts collision events at a maximum rate of 1 MHz, based on the information in the calorimeter and muon systems. In both the muon and the calorimeter system, the recorded signals are combined into a set of candidates that are passed on to the L0 decision unit, together with global event information, such as the number of pads hit in the SPD and

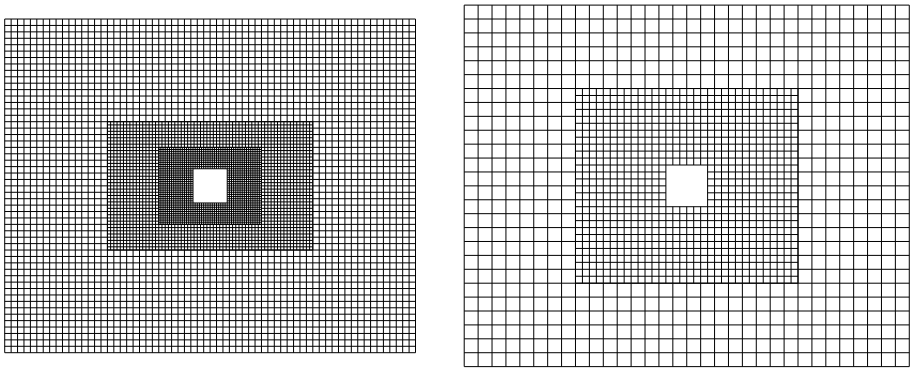

Figure 2.7: Transverse segmentation of the ECAL (left) and HCAL (right)

the summed transverse energy recorded in the calorimeter system in the current and the two preceding and subsequent 25 ns intervals.

More precisely, for each calorimeter readout board the 2 × 2 cluster of cells with the highest E_T corresponding to a photon, electron or hadron candidate is retained. In the muon system, candidates are created from coincidences in pads of the five stations.

The bulk of the events recorded in 2011 and 2012 were triggered by either the L0Hadron selection, which requires a single high-E_T hadron candidate, similar L0Electron and L0Photon selections, or a L0Muon selection requiring a high-p_T muon candidate. The L0DiMuon selection requires two muon candidates, with lower p_T thresholds than L0Muon. All of these selections also apply a requirement on the charged particle multiplicity inside the acceptance, measured by the number of SPD hits. This is because busy events require considerably more computing time in the subsequent software trigger reconstruction, so a lower average number of tracks per event allows to increase the total trigger efficiency, especially for high-multiplicity beauty and charm decays which also have low-p_T tracks. The L0DiMuon selection applies only a loose criterion on the number of SPD hits, 900 instead of the default 600, so the $Z \to \mu^+\mu^-$ sample it selects can be used to quantify the inefficiency of the other selections. This needs to be done in a data-driven way, because the SPD multiplicity distribution is not reliably modelled in simulation.

In case of a positive L0 trigger decision, the detector is fully read out and the event sent to a computing farm, where a software trigger applies further selections, until the remaining data volume can be stored and further processed. As most of the employed selections rely on partial event reconstruction, similar to the full offline reconstruction, this high-level trigger is described in section 2.4.

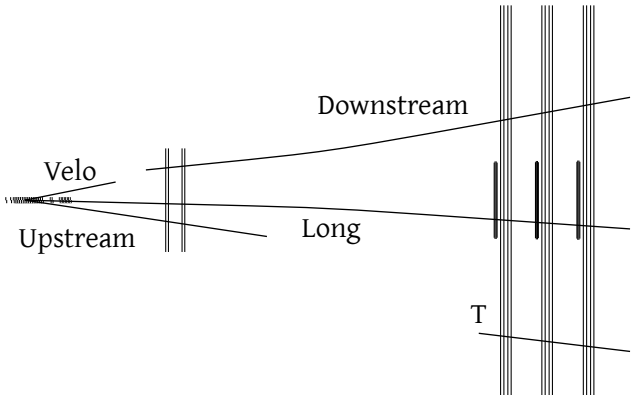

Figure 2.8: Schematic view of the reconstructed track types crossing the VELO, TT and T-station detection planes

2.3 Event reconstruction

The reconstruction of particles produced in an LHCb collision is commonly divided in three stages. Firstly, tracks are reconstructed, using mostly the information from the VELO, TT, IT and OT detectors. Secondly, the reconstructed tracks are classified as muons, electrons or charged hadrons using information from the RICH, muon and calorimeter detectors. Thirdly, neutral particles are reconstructed from calorimeter deposits. These three steps are described in sections 2.3.1 to 2.3.3, respectively.

The quality requirements on these reconstructed particles employed for the reconstruction of beauty and charm decays are usually quite minimal. The selections on kinematic variables, vertex quality, the alignment of the reconstructed candidate momentum and flight direction used to separate the signal from background are hard to satisfy with tracks or neutrals that are poorly reconstructed or wrongly combined. In addition, small inefficiencies on the final state particles quickly add up to a sizeable effect on the candidate selection efficiency. When reconstructing more inclusive signatures, however, *e.g.* jets compatible with a displaced vertex as described in chapter 3, it is beneficial to trade some single particle efficiency for a more accurate global view of the collision. This is done by the particle flow algorithm described in section 2.3.4, which combines all reconstructed objects, recovering hadron calorimeter deposits that are otherwise not included, and applies common selections to create a particle list suitable for jet clustering.

2.3.1 Track reconstruction

Different algorithms are used to maximise the efficiency for reconstructing the trajectories of charged particles from the hits they leave in the tracking detectors. Depending on their path through the LHCb detector, tracks are classified into different types, as defined in fig. 2.8. First the standalone VELO tracking finds all VELO track segments, that need to cross at least three stations. Then the forward tracking algorithm tries to extend these tracks to the TT and T stations, thus obtaining long tracks. As a subsequent step, the seeding algorithm tries to find track segments among the yet unused T station hits, after which another attempt is made to match these to VELO segments in order to maximise the efficiency for long tracks. The remaining T segments are matched to TT hits, creating downstream tracks, or kept as T tracks if this matching fails. Downstream and T tracks can be due to V^0 decays like $K^0_S \to \pi^+\pi^-$ and $\Lambda \to p\pi^-$, or secondary particles produced by interactions with the detector material outside the VELO acceptance. Low-momentum particles that are bent out of the acceptance by the magnet, *e.g.* charged hadrons from excited hadron decays like $D^{*+} \to D^0\pi^+$, are reconstructed as VELO-only or upstream tracks.

The track finding efficiency for charged particles produced close to the interaction region and inside the forward acceptance is above 96 % in the momentum range between 5 and 200 GeV/c.

All tracks are fitted using a Kalman filter method to determine the track parameters and their covariance matrix. The precision of the combined tracking system can be quantified by the momentum resolution and the impact parameter resolution of the tracks. The impact parameter is the minimal distance between the extrapolation of a track and a reference point, *e.g.* the assumed production point. Three factors, all related to the vertex detector design, influence the precision on this quantity: the resolution of the detector hits, multiple scattering of the particle when crossing detector material and the extrapolation distance to the reference point from the closest hit. A resolution of $15\,\mu\text{m} + (29\,\mu\text{m}\,\text{GeV/c})/p_\text{T}$ is achieved on the impact parameter along the *x* and *y* direction, as illustrated in fig. 2.9a.

The momentum is measured from the track curvature in the magnetic field, so the precision depends on the measurement of the track slopes before and after the magnet, and the integrated field along the particle trajectory. Besides the aforementioned hit resolutions, the precise alignment of the tracking detectors plays an important role. The momentum resolution varies from $\Delta p/p = 5 \times 10^{-3}$ for low momentum tracks to 8×10^{-3} at 100 GeV/c, as is shown in fig. 2.9b. This translates into a mass resolution of about 5×10^{-3} for two-body decays like $J/\psi \to \mu^+\mu^-$ and $\Upsilon(1S) \to \mu^+\mu^-$, *cf.* fig. 2.10.

Especially important for the sensitivity of long-lived particle searches to decays

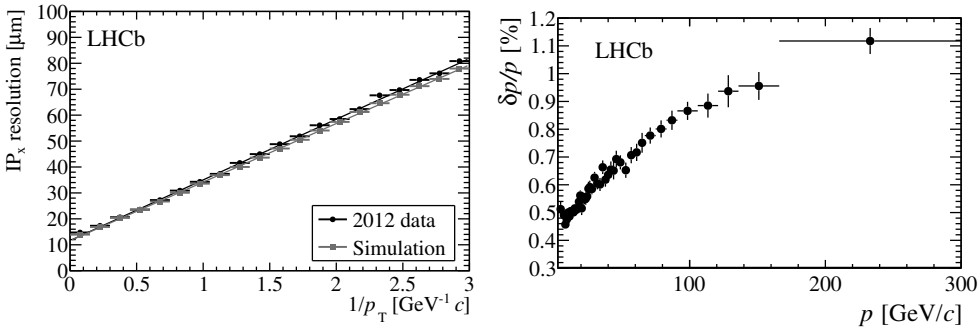

(a) (b)

Figure 2.9: (a) impact parameter resolution in data and simulation and (b) momentum resolution for long tracks measured in data, figures from [114]

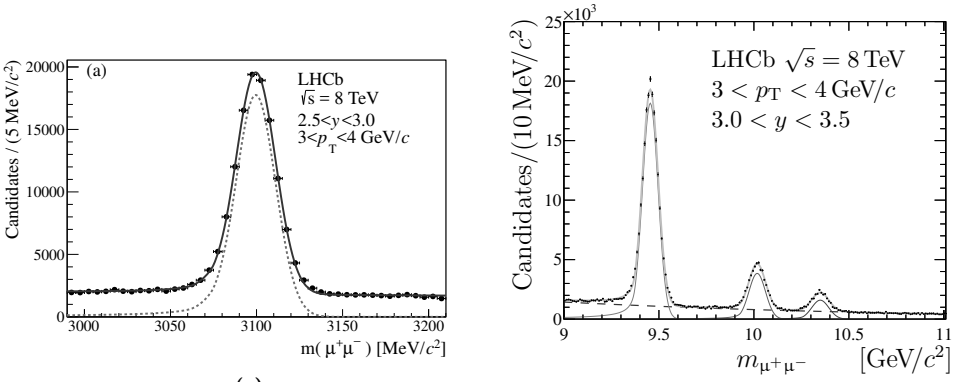

(a) (b)

Figure 2.10: Reconstructed $\mu^+\mu^-$ invariant mass spectrum in the region of the (a) J/ψ and (b) Y resonances, figures from [117] and [118], respectively

with large lifetimes is the acceptance of the VELO detector and of the associated track reconstruction. Figure 2.6 shows the typical cross-section of the sensitive elements by a plane containing the beam axis. A track needs to traverse the sensitive area of at least three stations to be reconstructible, but the efficiency is higher for tracks crossing more stations. In addition, tracks that originate from the beam line have a higher efficiency than tracks that do not point back to it. The latter is due to the specific configuration of the track finding algorithm used in HLT1 [119]: it proceeds through a search for seeds of at least four R clusters in the same 45° sector of five consecutive R sensors, first backward starting from the high-z side of the detector, then forward from the other side. Those are then extended within the same sector and upgraded to three-dimensional space tracks by adding matching Φ hits. In the HLT2 and offline configurations, shorter and non-pointing tracks are recovered in two steps: first all seeds of three R clusters in four consecutive sensors are extended with hits in neighbouring 45° sectors. Only clusters that were not used by the first procedure are considered for this recovery. Second, tracks are searched for using the remaining Φ clusters. The effective angular acceptance goes maximally from about $\eta = 1.35$ to $\eta = 5$ for particles from the beam axis. For particles produced away from the beam line, the acceptance changes depending on the production position and loses its approximate cylindrical symmetry: some tracks at high rapidity fall inside the acceptance, while some tracks at low rapidity fall outside.

2.3.2 Charged particle identification

Information from the RICH detectors, calorimeters and the muon system is added to each track to allow identification of the different particles. A compatible muon segment and small energy deposits in the calorimeter are characteristic of muons. Electrons are identified mostly by the shape of the associated shower in the calorimeter: they deposit most of their energy in the PS and ECAL, only a small fraction in the HCAL. The charged hadrons — pions, kaons and protons — are distinguished by the radius of the associated RICH ring. As this is a measurement of the particle's speed, the discriminating power decreases towards higher momenta.

The information from the different subdetectors is combined to obtain maximal separation between the different particle identification hypotheses. Two approaches are used for this. The first defines a likelihood in each particle identification detector for the observed signals under the different hypotheses, such that combined likelihood ratios can be used as selection variables. The second uses a multivariate classifier for each hypothesis, the output of which forms a good classifier. The latter method has the advantage that it is easier to incorporate other, *e.g.* track quality, inputs and to re-optimise for a specific purpose. The optimal balance between signal

efficiency and background suppression strongly depends on the specific analysis application.

Even if accurate particle identification such as kaon-pion separation is not as crucial for inclusive jet final states as it is for separating hadronic heavy flavour decays, it offers interesting additional information. A by-product of the multivariate classifiers mentioned before is an improved ghost track rejection: a charged particle with sufficient momentum should have associated deposits in the calorimeters or the muon chambers, and a corresponding RICH ring. Another example is the kaon identification that may help in distinguishing beauty and charm from light quark and gluon jets.

2.3.3 Neutral particle reconstruction

Photons and neutral pions are reconstructed using the electromagnetic calorimeter. Energy deposits are grouped into 3×3 clusters with centres separated by at least one cell, which are, if not compatible with any of the tracks in the event, assumed to be due to neutral particles. Isolated clusters are classified as photons or neutral pions based on the compatibility of their cluster shape with being caused by a single photon or by two nearby photons coming from a π^0 decay [120]. Neutral pions with smaller momentum are reconstructed by making pairs of sufficiently separated photons compatible with originating from a $\pi^0 \to \gamma\gamma$ decay. Neutral pions reconstructed using these two strategies are called merged and resolved π^0, respectively.

2.3.4 Particle Flow

The particle flow algorithm takes all standard reconstructed charged and neutral particles and the remaining calorimeter clusters, and combines them into a list of unique inputs for jet clustering [121].

For charged particles, a relatively loose track preselection, employing the multivariate classifiers mentioned in section 2.3.2, is used to reject ghost tracks. Tracks with measurements before and after the magnet are also removed in case they do not have a significant measured curvature, because the charge of these tracks is not unambiguously known and their reconstructed momentum can become very large. Most of these are ghost tracks, typically wrongly combined upstream and downstream segments. VELO and upstream tracks are retained, but with zero energy, so they do not influence the reconstructed jet momentum and energy. Tracks with a momentum estimate are retained as charged hadron, electron or muon depending on their particle identification information.

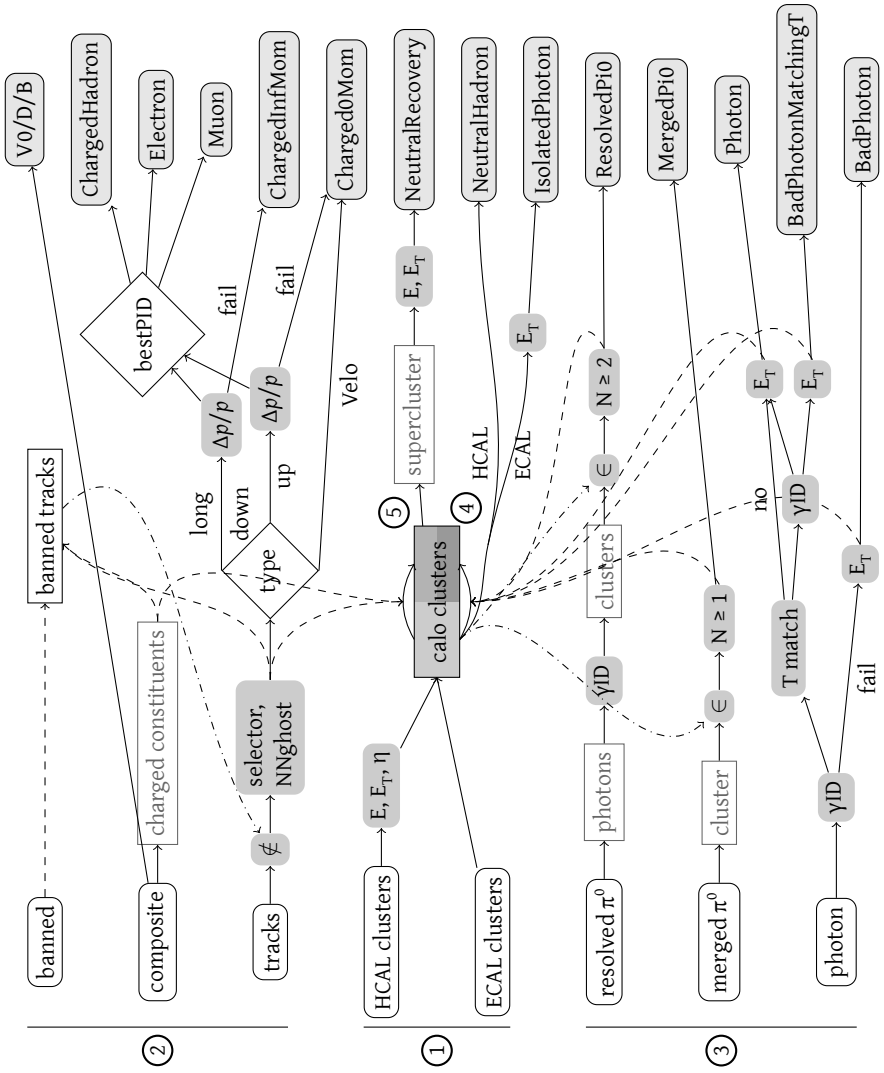

Figure 2.11: Technical overview of the particle flow algorithm. The leftmost nodes, at the bottom of the page, represent the different inputs, while the rightmost nodes indicate the output categories. All output types are used for jet clustering, except for BadPhotonMatchingT. The numbers in steps correspond to the different steps in the procedure described in section 2.3.4. The rectangular nodes in the middle, a list of tracks and calorimeter clusters, ensure that no reconstructed object is used twice.

Reconstructed photons as well as resolved and merged π^0 particles, as described in the previous section, are added after applying quality criteria. Most notably, photons that are near a T-track segment are subject to tighter quality requirements in order to improve the rejection of charged particles produced at high pseudorapidity. These could be too close to the beam line to be reconstructible in the VELO, but bent into the calorimeter acceptance by the magnet.

Calorimeter deposits from neutral particles are recovered in two steps. First the clusters that have not been associated to any neutral or charged particle up to this point are interpreted as neutral hadrons and photons, depending on their cluster type, and retained as inputs for the jet clustering. Second, the remaining clusters that have been associated to charged particles are used to recover nearby neutral energy deposits: after merging nearby clusters in both calorimeters, the total expected energy from all associated charged particles is subtracted from the recorded energy and the remainder added as a neutral particle.

A graphical overview of how the different inputs are combined is given by fig. 2.11.

2.4 High-level trigger

The high-level trigger is a real-time software application running on a large computing farm, with the task to reduce the million of events read out every second to an amount that can be stored: about 3 kHz in 2011 and 5 kHz in 2012. Spreading the events over about 30 000 processes results in about 30 ms of available processing time per event. During 2012, the effective throughput of the high-level trigger was increased by buffering events to disk storage after L0: this allowed to over-commit the computing farm by about 10–20 % and process the buffered events while the LHC was being refilled.

The event rate reduction is achieved in two stages: HLT1 searches for a limited number of signatures that inclusively identify the presence of a heavy flavour decay or another physics process of interest, while HLT2 employs a larger set of more specific selections. In both cases, the different selections are organised in so-called *lines*, which are executed independently: either one of them can retain the event, but the application will always run all other lines as well.

HLT1

HLT1 starts by reconstructing VELO tracks with the FastVeloTracking algorithm described in section 2.3.1. These tracks are then used to search for primary interaction vertices. In order to obtain a momentum estimate for the tracks, they must be ex-

tended to the downstream detector region. The size of the search window in the downstream region depends on the minimal momentum where a high track finding efficiency is desired: the lower the targeted momenta are, the larger the search windows must be. Large search windows, however, are computationally expensive, because they quickly increase the number of hit combinations to be considered. Various strategies are used by the different trigger lines to keep the CPU consumption within the available budget. The muon selections require a L0 decision based on the muon system and first match VELO tracks to track segments in the muon stations.

The main selection for hadrons [122] runs the forward tracking only for VELO tracks that are displaced from all primary vertices by at least 100 µm and have at least ten hits, while no more than two hits expected from the trajectory through the detector are not found. In addition, the forward tracking can be performed with a smaller search window because the next selection step in the procedure is a high transverse momentum requirement. For the most inclusive line, which runs on all L0-triggered events, the track p_T needs to exceed 1.7 GeV/c; for lines that only accept L0-electron and L0-muon triggered events the selection requirements are considerably looser. Finally the remaining tracks are fitted and a good fit χ^2 required.

Final states containing leptons are also selected by a high-p_T single muon line, targeting massive electroweak gauge boson decays, and single electron and muon lines with looser p_T requirements, no displacement requirement, but tighter track quality criteria. The dimuon lines select muon pairs with invariant mass above 2.7 GeV/c^2 without displacement requirement and down to 1 GeV/c^2 with an impact parameter selection for each of the muons [123].

HLT2

HLT1 brings the event rate down to about 60 kHz and leaves considerably more time per event for reconstruction and selection in the second stage, HLT2. In HLT2 a larger set of reconstruction algorithms is executed: a more complete version of the VELO track reconstruction, the forward tracking for all tracks above a certain momentum threshold and subsequently the Kalman filter track fit. There is a much broader set of different selection lines for this final online selection, all exploiting more specific characteristics of certain signals of interest, in particular fully reconstructed beauty and charm decays.

A few rather inclusive selections remain, *e.g.* a single muon line with a p_T threshold at 10 GeV/c for electroweak physics, a displaced single muon line for semi-leptonic B-decays *etc.* Most of the fully hadronic B-decays are selected by a family of so-called topological trigger selections, which is discussed in more detail in section 3.4.2.

Other selections aim at more specific signatures, *e.g.* a line selecting opposite-sign non-prompt dimuon pairs for B-decays with charmonium resonances, a line for dimuon combinations with mass close to the B hadron masses, lines selecting exclusive or partial charm final states *etc.* The trigger strategy used in this thesis for inclusively selecting high-multiplicity high-mass displaced vertices, which uses dedicated reconstruction and selection steps, is described in section 3.4.2.

The trigger system during run 1 is described in more detail in [124]. A more extensive overview of the selections can be found in [125]. [126] shows the performance for a set of reference signatures.

2.5 Offline data processing

All events selected by HLT2 are stored in a raw data format, which contains the full event as read out from the detector, supplemented with information about the trigger selections that retained it. For each positive HLT1 or HLT2 decision the list of detector channels corresponding to each selected candidate is kept, together with a limited set of calculated quantities.

This raw data is not directly suitable for analysis. First a full event reconstruction step, as described in section 2.3 is applied[2], whose outputs are saved in files of the DST format. These go through another preselection step which is called stripping. Both are run centrally by the LHCb computing team, the reconstruction typically a few times per data set to take advantage of improvements in the algorithms, the stripping more often to allow for new and improved selections to be developed.

As many copies of the stripping output are available to allow for quick analysis turnaround times, the disk space used by it is relatively costly. Therefore, a typical exclusive selection line should not accept more than a fraction of the order 5×10^{-4} out of the total triggered data set. For cleanly reconstructed signatures, it is often sufficient to only save the selected candidates rather than the full event — the output files are then called μDST — in which case the accept rate can be about an order of magnitude larger.

[2] Logically the particle flow algorithm also belongs there, but in order to take advantage of more recent developments and not to increase the event size, it was executed on demand inside the stripping step.

3

Search for long-lived particles decaying to jet pairs

The search for long-lived particles decaying to quarks described in this chapter extends the analysis [2, 112] with a data set that is about twice as large as before: 1.38 fb^{-1} recorded in 2012 at a centre-of-mass energy of 8 TeV in addition to the 0.62 fb^{-1} data set collected in 2011 at 7 TeV. Furthermore, the analysis benefits from an improved reconstruction and selection efficiency, both in the 2012 trigger and offline.

3.1 Introduction

In this analysis, long-lived particles are studied that decay into a hadronic final state that can be reconstructed as a pair of jets produced at the displaced decay vertex. As a benchmark production model, the decay of a scalar resonance, such as the standard model Brout-Englert-Higgs massive scalar boson, into two long-lived particles is used. Only one long-lived particle is required to be found in the analysis, so the results do not strongly depend on the details of this production model.

No long-lived particles with masses above 10 GeV/c^2 are known, so all backgrounds are due to misreconstruction or interactions outside the proton-proton collision region. Therefore, the analysis focuses on correctly reconstructing and identifying the signal signature of two jets being produced in the hadronic decay of a long-lived particle. A single candidate that can unambiguously be identified as the decay of an exotic long-lived particle would immediately imply a discovery.

Backgrounds to this search are discussed in section 3.2, together with the tools developed to suppress them. Then, after a short description in section 3.3 of the data samples considered, section 3.4 gives an overview of how the reconstruction and first selection steps are integrated in the data processing sequence. The final selection is outlined in section 3.5, and its efficiency for the signal model is given in section 3.6.

3.2 Backgrounds

The reconstruction strategy that was found to be the most effective first finds decay vertices using tracks in the VELO and then adds jets from particles compatible with these vertices [2]. A detailed description will be given in section 3.4. Backgrounds from material interactions and high-occupancy events mostly play a role in the first stage, while other combinatorial backgrounds also remain important after jet reconstruction.

3.2.1 Material interactions

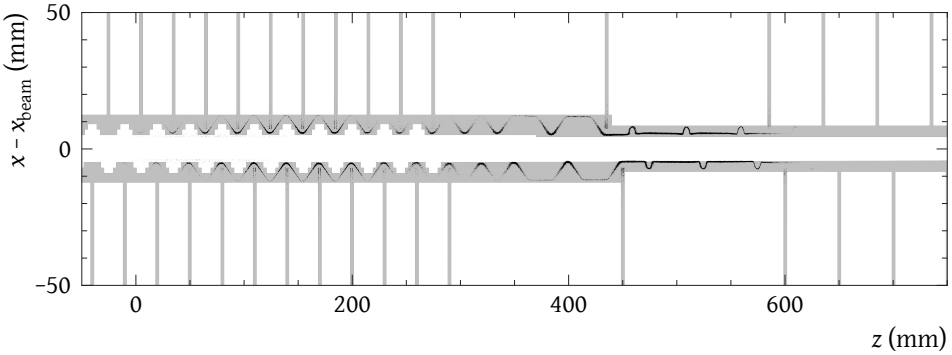

Figure 3.1: Geometry of the veto region around the detector material in the horizontal plane, overlayed with the spatial distribution of vertices with $R_{xy} > 4$ mm

Hadronic interactions of energetic particles produced in the collisions with detector material constitute a large background to this analysis. In the region where vertices can be reconstructed from VELO tracks, the most important detector parts are the VELO sensors and the RF foil. An algorithm has been developed to veto an envelope region around these, which was validated using preselected data in 2011 and has been used to reject background in the trigger selection since.

As can be seen in fig. 3.1, the vetoed region is located mostly between a radial distance of 5 mm and 13.5 mm from the VELO half centre, except for the 4 mm thick region around each station which also extends up to larger displacements. The inner radius of the veto is decreased to 4.3 mm in the region with the so-called gap sensor, which is located at $z \approx 450$ mm. In the central part of the detector where stations are equidistant, the variation of the RF foil radius along the z axis is taken into account by removing two cylinders from the vetoed region: vertices with R_{xy} up to 9 mm are accepted if they are further than 11 mm away in z from any detector station, and

those with R_{xy} up to 7 mm are accepted if they are more than 8 mm away from the detectors. In the downstream part of the VELO, envelopes are made larger, while in the overlap region close to the *yz*-plane, all vertices with radius larger than a few mm are vetoed.

This veto is effective in removing material interaction backgrounds, but significantly reduces the efficiency for R_{xy} in the range 5–13.5 mm. Some of this inefficiency could be recovered by following the shape of the RF foil more closely or by using the fitted vertex uncertainties to reduce the average envelope thickness. Both however require more precise knowledge of the shape and location of the RF foil, and a correct modelling of the vertex uncertainties also for vertices made of only few and soft tracks. Due to the importance of the material veto for controlling the accept rate of the dedicated trigger lines, the simplicity of the current approach was preferred.

3.2.2 High-occupancy events

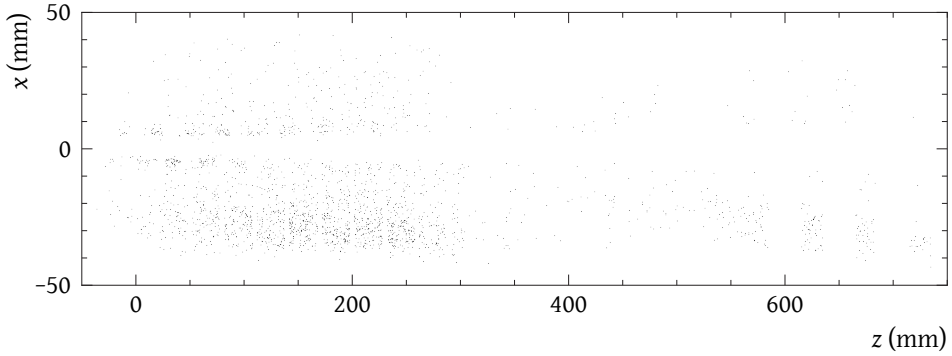

Figure 3.2: Spatial distribution of loosely selected vertices (using a small subset of the data sample) outside the material veto region, and with $R_{xy} > 5$ mm

A different class of fake candidates is due to events where tracks from the proton-proton interaction cross with particles produced in interactions of the beam with parts of the accelerator. The latter cross the VELO in the horizontal plane and create sufficient confusion for the VELO track reconstruction to find many nearby tracks. Since these tracks are by construction close together, they give rise to, often high-multiplicity, vertices that are displaced from the beam line. Figure 3.2 shows the spatial distribution of these vertices, which is highly asymmetric between the two sides of the detector. This supports the hypothesis that interactions of the beams with collimators play a role — the fact that the vertices are found on the inside of

the accelerator ring disagrees with the most naive expectation, but the beam optics around the interaction is rather complicated.

Two selections were designed to veto these events. The first relies on the fact that the VELO track reconstruction creates many more tracks than can be expected from the number of hits, so events where there is on average more than one track for every ten hit clusters are vetoed. The second criterion uses the characteristic that the machine-related background particles are almost parallel to the beam line, and thus create an imbalance in the ϕ coordinate of the hit distribution. This can be quantified by taking the vector sum of a unit vector with the ϕ value of the cluster in the transverse plane, for all ϕ hit clusters. Events are removed where the length of this sum vector, $D_\phi = \left|\sum_j e^{i\phi_j}\right|$, whose distribution for a preselected data sample is shown in fig. 3.3, exceeds 250. Comparing the distribution in fig. 3.3 — which is for data events with at least one reconstructed displaced vertex — to simulation, the peak at low values is well described, but the distribution tends to zero above 250 in simulation.

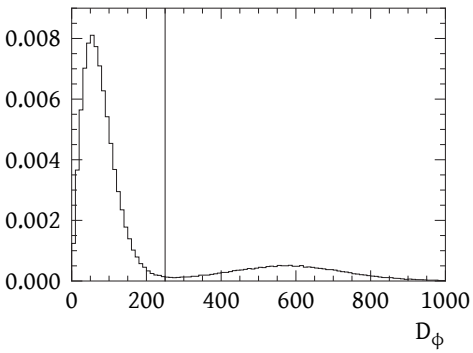

Figure 3.3: Distribution of the D_ϕ variable defined in section 3.2.2 for a preselected data sample (normalised to unity)

3.2.3 Heavy flavour decays

The remaining background is found to be due to combinations of heavy flavour decay product tracks with other displaced tracks, *e.g.* from material interactions or V^0 decays, or imprecisely reconstructed primary vertex tracks. Attempts were made to simulate sufficiently large background samples, but this turned out to be not feasible. The number of $b\bar{b}$ pairs created per fb^{-1} is of the order 10^{11}, the charm production cross-section is about an order of magnitude larger, and V^0 decays are found

in almost every event passing the trigger. Even if biasing simulated events towards larger B hadron lifetimes can gain about two orders of magnitude, the number of events required to sufficiently sample all relevant features of this small background is clearly out of reach with the computing time and storage space currently available[1].

Most of the selection is aimed at rejecting this background to a negligible level. At the level of the vertex selection, *i.e.* before the jet reconstruction, the most discriminating features are the typically smaller track multiplicity than the signal, the smaller p_{T} carried by the tracks and the smaller invariant mass of the vertex. Furthermore, the average beauty lifetimes are about 1.5 ps, leading to flight distances of the order 1 cm, and B hadrons are on average produced under smaller angles with the beam line than the signal. Together with the larger background track density around the interaction region, this leads to a rapid decrease of this background with the displacement of the vertex from the beam line, which is why the vertex requirements are relaxed for large R_{xy}.

The jet reconstruction provides additional information to reject more of these backgrounds: in most cases they do not contain two real jets, but rather part of one jet and a few random tracks, or two parts of jets. By combining jet quality with decay vertex selections, and requiring the final candidate to be compatible with the decay of a particle produced in the primary interaction, this background can be further suppressed. In addition, background candidates mostly have a much smaller reconstructed dijet invariant mass, so by combining the knowledge on the shape of invariant mass distribution and the yield at low invariant mass, their final number under the signal peak can be constrained.

3.3 Data samples

3.3.1 LHCb pp collision data

As in a previous analysis [2] the data sample recorded in 2011 at a centre-of-mass-energy of 7 TeV with the HLT configurations 0x0076 and 0x0079[2] — the largest data sample with the same configuration of the trigger selections used — is considered,

[1] for reference, the size of large simulated samples is expressed in millions of events, the largest single productions contain about 5×10^8 events

[2] Trigger configurations are uniquely identified by a hexadecimal number with eight figures, which is called a TCK (trigger configuration key), the first four related to the HLT configuration and the latter four to the L0 settings. Simulation TCKs have the same number as the corresponding data TCK, with the first number replaced by 4 and differ only by the prescales and postscales that are disabled.

which amounts to an integrated luminosity of about 0.62 fb^{-1}. The overview of configurations and their corresponding integrated luminosity is given in table 3.1.

Out of the data taken in 2012 at a centre-of-mass energy of 8 TeV, 1.38 fb^{-1} of integrated luminosity is used. This is not the complete available data set due to changing conditions and HLT settings during the first months of data taking. An overview can be found in table 3.1: the HLT configurations from `0x0099` up to `0x00ad` use an improved and retuned version of the dedicated lines, thus providing a uniform data sample for analysis, which contains the bulk of the available data. Almost half of the missing third of the luminosity was recorded with the `0x0097003d` TCK, where the dedicated selections had a different tuning. There were also changes made to the track reconstruction between the releases v13r3 and v14r2 of the high-level trigger software application. It is possible to recover the data taken with these settings, but given that this adds only a small fraction to the total available data set at the expense of more complicated studies of the systematic uncertainty, it was not found worthwhile.

The legacy run 1 version of the offline reconstruction and stripping [128, 129] are used. The dedicated lines that construct the dijet candidates based on displaced vertices reconstructed in HLT2 or offline, `DisplVerticesLinesJetHltSingleHighMass` and `DisplVerticesLinesJetSingleHighMass`, respectively, are described below in section 3.4.

3.3.2 Simulated signal samples

Hidden Valley π_v from Higgs decays

The simulated signal samples used for most of the trigger and stripping reconstruction optimisation and the final selection are based on the Hidden Valley model introduced in chapter 1. This model adds another gauge group to the standard model, inspired by QCD. The lightest of a new set of heavy bound states, called π_v in analogy with the strong interaction, can only decay to standard model particles, but since this decay needs to proceed through a very heavy mediator, it is suppressed and the π_v has a macroscopic lifetime. One of the possible production models, through the decay of the standard model Brout-Englert-Higgs massive scalar boson into a pair of π_v particles, is assumed. The mass of the scalar boson is set to 125 GeV/c^2, and the $gg \to H$ process selected, which dominates the production in 7–8 TeV pp collisions, *cf.* section 1.3.

Samples are generated for π_v masses of 25 GeV/c^2, 35 GeV/c^2, 43 GeV/c^2 and 50 GeV/c^2 and lifetimes of 10 ps and 100 ps; this gives enough events in the reconstructible range to reweight to π_v lifetimes up to a few times 100 ps. All events were produced

Table 3.1: List of trigger configurations used during 2011 and 2012, with corresponding integrated luminosity. The configurations with the same configuration of the dedicated lines are grouped, and both for 2011 and 2012 only the last of these groups is used. For entries with a dash in the last column, the dedicated HLT2 lines were turned off. Their contribution is therefore excluded from the sums. The luminosity numbers are only indicative, as they do not include the precise calibration [127]

2011

Release	TCK	$\int \mathcal{L}\,dt$ (pb^{-1})
Early data		373.90
v12r3	0x00360032	3.31
v12r4	0x00480032	1.94
	0x004a0033	2.08
v12r5	0x005a0032	64.61
	0x005b0032	2.01
	0x005d0033	2.25
v12r6	0x006d0032	98.19
	0x00700034	1.10
	0x00710035	0.90
	0x00730035	192.49
	0x00740036	5.02
June technical stop		694.16
v12r8	0x00760037	291.93
v12r9	0x00790037	38.47
	0x00790038	363.76

2012

Release	TCK	$\int \mathcal{L}\,dt$ (pb^{-1})	
First version		0.34	
v13r2	0x007e003a	0.16	
	0x007f0040	0.18	
	0x00860040	5.07	—
Retuned version		330.79	
v13r3	0x008c0040	59.68	
	0x008e0040	0.43	
v14r2	0x0094003d	271.69	—
	0x0095003d	0.81	—
	0x0097003d	270.68	
New implementation		1375.94	
v14r6	0x00990042	553.89	
	0x00990043	2.66	
	0x00990044	140.68	
	0x009a0042	0.02	
v14r8	0x009f0045	26.29	
	0x00a10044	25.62	
	0x00a10045	70.41	
	0x00a20044	1.64	
v14r9	0x00a30044	303.34	
	0x00a30046	30.17	
v14r11	0x00a90046	61.56	
	0x00ab0046	50.14	
	0x00ac0046	109.40	
	0x00ad0046	0.12	

with the PYTHIA8 generator [130, 131], with a specific LHCb configuration [132] and using the CTEQ6 leading order parton distribution function fit [133]. The interaction of the generated particles with the detector, and its response, are implemented using the GEANT4 toolkit [134, 135] as described in [136]. The collision settings, trigger (with the TCK `0x409f0045` for 2012 conditions) and offline reconstruction are consistent with those used to record and process the data sample.

In order to save computing time, a selection is applied after the generator stage: only events with at least four reconstructible tracks ($p > 2\,\text{GeV}/c$ and $\theta < 400$ mrad) from one π_v decay are kept — this is the minimum requirement for finding a displaced vertex. The fraction of events accepted by this generator-level requirement is about 25–28 %, increasing with the mass of the π_v and the mass of its decay products. It takes slightly higher values for 2012 conditions due to the higher collision energy, which leads to a larger longitudinal boost of the scalar boson.

3.4 Preselection and candidate definition

This section gives a description of the different trigger and offline reconstruction and selection steps leading to the candidate sample prior to the final selection optimisation. More details on the optimisation of the reconstruction can be found in [1, 104] and [2, 112].

3.4.1 L0 and HLT1

In the L0 trigger, the searched-for signal is, due to the typical transverse momentum of the final state, efficiently selected by the standard high-bandwidth trigger lines for flavour physics: L0Hadron and L0Muon. Events selected by the L0Photon, L0Electron and L0DiMuon lines are also included, despite their minor contribution to the efficiency. The final candidates are required to be compatible with one of the candidates satisfying the trigger selections, based on either hit-level overlap criteria — as commonly used in the TISTOS method [137] — or based on the presence of a L0 candidate that passes all trigger criteria inside the cone of one of the final state jets.

Also in HLT1 the inclusive flavour physics selections based on a non-prompt track with high transverse momentum are mostly relied upon. An important limitation, however, is that all of these require a track reconstructed in the VELO. This effectively limits the acceptance to decays within a few cm from the beam line and a few tens of cm from the interaction point, *cf.* section 2.3.1.

A sizeable inefficiency for highly displaced (*i.e.* beyond the RF foil, from about 10 mm from the beam line on) long-lived particle decays was observed, as can be

seen in fig. 3.4. This was traced back to the stringent VELO track quality requirements made as the first step of the most inclusive line, Hlt1TrackAllL0, which is targeted at hadronic heavy flavour decays: 10 VELO hits need to be found on the track, and there can be no more than two hits that are expected based on the trajectory through the detector but not found. In order to recover some of this efficiency, a dedicated HLT1 line was added in 2012, trading this track quality requirement for the presence of another nearby displaced track.

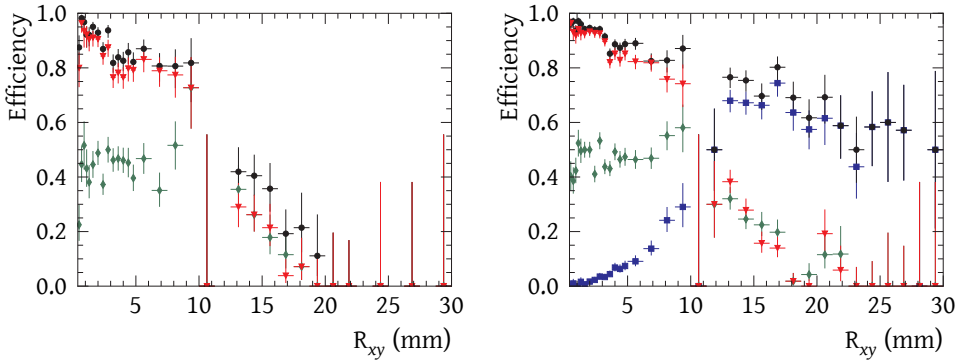

(a) 2011 (b) 2012

Figure 3.4: HLT1 efficiency as a function of transverse displacement for 35 GeV/c^2 π_v particles with both B hadrons inside the acceptance ($\eta > 2$), that pass HLT2 (see section 3.4.2) and either of the two-jet stripping lines (see section 3.4.3), using only Hlt1TrackAllL0 (red triangles), Hlt1TrackPhoton (green diamonds), Hlt1VertexDisplVertex (blue squares) or the combination of these selections (black circles) in (a) 2011 and (b) 2012

More specifically, the track quality requirement is loosened to finding at least three so-called spacepoints — coincidences of R and Φ hits in the same station — two of which should be in consecutive stations crossed by the track. Then all two-track combinations are made out of the tracks with a closest distance to the beam line of at least 2 mm. Only if the closest approach of these tracks is smaller than 0.3 mm and the vertex is displaced from the beam line by more than 12 mm, the next reconstruction steps and selections are performed, with the same settings as in the Hlt1TrackAllL0 line and requiring at least one of the two tracks to pass all reconstruction and selection stages. The first of these — the most time consuming part of HLT1 — is the forward tracking algorithm that searches for downstream hits compatible with the VELO segment. The transverse momentum of one of the tracks should be above 1.7 GeV/c and its total momentum should exceed 10 GeV/c. The track fit χ^2 of this track is also required to be below 2.5 per degree of freedom.

Figure 3.4 shows the efficiency of the different trigger lines as a function of transverse displacement: the Hlt1VertexDisplVertex line improves the HLT1 efficiency for decays beyond the RF foil to about 60 %.

All other selections targeted at heavy flavour decays, as well as the inclusive single- and dilepton trigger lines, are included, but only contribute a small additional efficiency.

3.4.2 HLT2

In HLT2, a combination of dedicated lines and the topological B-lines is used. The topological lines make 2-, 3- and 4-track combinations resembling the typical hadronic B decay topologies with possible charm intermediate states. Selections are made based on the transverse momentum, both of the individual constituents and of the combination, the compatibility of the tracks coming from a common vertex or decay cascade, and the displacement of the tracks and vertices from the primary interaction vertex. The final discrimination, mostly against charm decays, is provided by a corrected mass variable. That is obtained as the minimal mass compatible with the measured flight direction from the primary to the secondary vertex and the measured momentum,

$$m_{\text{corr}} = \sqrt{m^2 + (p \sin\theta)^2} + p \sin\theta, \tag{3.1}$$

where p and m are the measured momentum and mass, respectively, and θ is the angle between the displacement vector and the momentum [138]. This formula can be obtained as follows: for the candidate to point back, one should add a momentum four-vector that aligns the candidate momentum with the direction of flight given by the vertices; this constraint fixes the two components transverse to the flight direction; for the combined mass to be minimal, the added four-momentum vector should have zero invariant mass; the corrected candidate mass can then be written as a function of measured quantities and the added momentum along the direction of flight. The corrected mass is minimal when in the centre of mass system no momentum difference along the flight direction is present, or, equivalently, the opening angle between the observed and added momentum is maximal, which results in eq. (3.1). The selection variables are combined in a multivariate classifier [139] for the final selection [124].

The dedicated inclusive displaced vertex lines focus on very displaced vertex signatures with high track multiplicities and invariant mass, with a very low retention on background. The basic approach was similar in 2011 and 2012: a vertex finding step followed by a set of selections on track multiplicity, mass and displacement to optimise efficiency over as wide a range as possible in mass and displacement.

The vertex finding algorithm used, a specially tuned version of the primary vertex reconstruction, was the same for all data (details can be found in [140]): an iterative approach of seed finding and an adaptive fit to collect all tracks belonging to each of the seeds, which are then removed from the list of inputs for the next seed finding step. An important difference, however, is that in 2012 only forward VELO tracks not pointing back to any primary interaction vertex within 0.1 mm are considered for the vertex finding[3]. The main advantage of this is the reduced competition for tracks between the displaced and primary vertex seeds, where the latter take precedence because of their typically higher track multiplicity. A more detailed discussion of the performance will be given in section 4.1.

All vertex candidates need to have at least four tracks, be displaced from the beam line by more than 0.4 mm in the transverse plane — in order to remove primary vertices and a large fraction of the heavy flavour secondary vertices — and be located outside the vetoed region around the detector material, *cf.* section 3.2.1. In addition, in the 2012 version fewer than half of the tracks can have hits upstream of the vertex position and a single track should not carry more than 90 % of the energy — two mild ghost and low-multiplicity background rejection criteria. Furthermore, there should be at least one primary vertex that is located upstream of the candidate and inside the luminous region. This primary vertex should have at least ten tracks out of which at least one is forward and one backward, to ensure it is due to a highly inelastic proton-proton interaction.

The specific selections applied afterwards make different trade-offs between displacement and invariant mass, covering as different signatures as possible while keeping the overall retention low. The thresholds are listed in table 3.2. Due to the tighter input track selection, these numbers have different meanings for 2011 and 2012 data: in general the 2012 cuts are tighter, except for large displacements. For evaluating kinematic quantities of the vertex, the VELO track segments are searched for in the common list of long tracks also used by exclusive lines. VELO track segments that are not found there, are assigned a default transverse momentum of 400 MeV/c (in 2011) or are ignored (in 2012).

3.4.3 Stripping

The main changes in candidate definition with respect to a previous analysis [2] are related to the addition of jet reconstruction to the stripping. Even though the

[3]The distinction between forward and backward VELO tracks is made by the VELO tracking algorithm: the tracks found in the first pass, searching from the high-z end of the detector towards the interaction point, are classified as forward, while those found in the second pass, starting from the other end of the detector, are classified as backward.

Table 3.2: Specific selections on the vertex candidates used in the dedicated HLT2 and stripping lines. The columns give the minimal number of tracks N^{vertex}_{track}, the minimal radial displacement from the beam line R_{xy}, the minimal invariant mass m and the minimal sum of the track transverse momenta $\sum p_T$ (where applicable)

Line	N^{vertex}_{track}	R_{xy} mm	m GeV/c^2	$\sum p_T$ GeV/c
2011				
SingleHighFDPostScaled	5	2.0	4.5	3.0
SingleHighMassPostScaled	4	0.4	10.0	3.0
2012				
SingleVeryHighFD	4	5.0	2.0	—
SingleHighFD	5	3.0	2.8	—
Single	4	1.5	5.0	—
SingleHighMass	5	0.5	8.0	—
Stripping21				
JetSingleHighMass	5	0.4	5.0	7.0

distinction between the stripping and the offline analysis step is generally organisational rather than conceptual, this gives a clear advantage in view of the event retention constraint: whereas the selection of vertices before could not retain more than a fraction of 0.0005 of all events in order to keep the output size under control, it can now accept a fraction of up to 0.01, after which the 2-jet signature can be used instead to obtain the remaining retention. This allows to increase the signal efficiency while increasing the total retention factor. The new bottleneck is the CPU time consumption of the particle flow procedure, which includes the on-demand decoding and reconstruction of calorimeter information and amounts to about 100 ms per event[4], while the average should not exceed 1 ms per event.

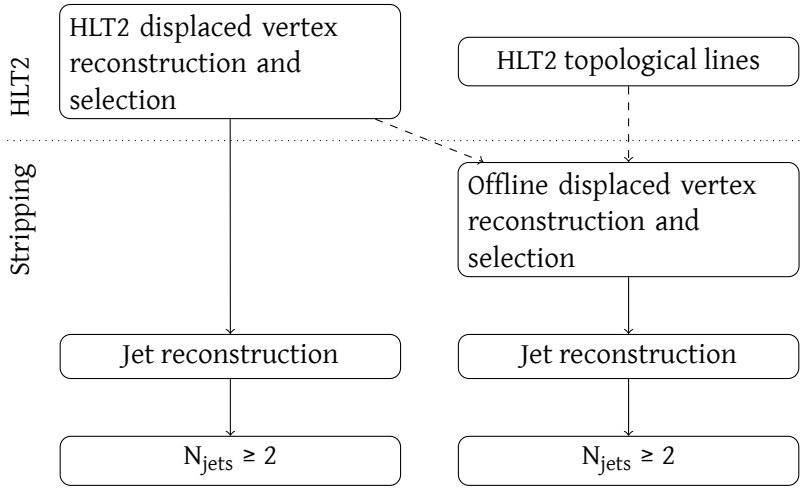

Figure 3.5: Illustration of the relation between the different HLT2 and stripping lines

As illustrated in fig. 3.5, two sets of vertices are used: the candidates from the dedicated HLT2 lines, and the candidates from an offline vertex finding step. The former have already passed the tight trigger vertex selections, so they can be used directly. The offline vertex finding step went through the same changes as HLT2 in 2012, *i.e.* mainly the filtered input tracks list and updated selection cuts, which are listed in table 3.2. In addition, fewer than half of the tracks can have hits upstream of the vertex and no signal track can carry more than 80 % of the candidate energy. Figure 3.6 shows the vertex reconstruction and selection performance as a function of radial displacement for both cases. The effect of the tight selections close to the

[4]This may be improved in the future, by further tuning the procedure, by improving the implementation of the algorithms, or by making the particle flow part of the standard event reconstruction executed for every event.

background-dominated beam line region, the loss of efficiency due to the detector acceptance at high R_{xy} and the lower efficiency in the region of the RF foil are clearly visible, as well as an increased efficiency at large displacements in 2012.

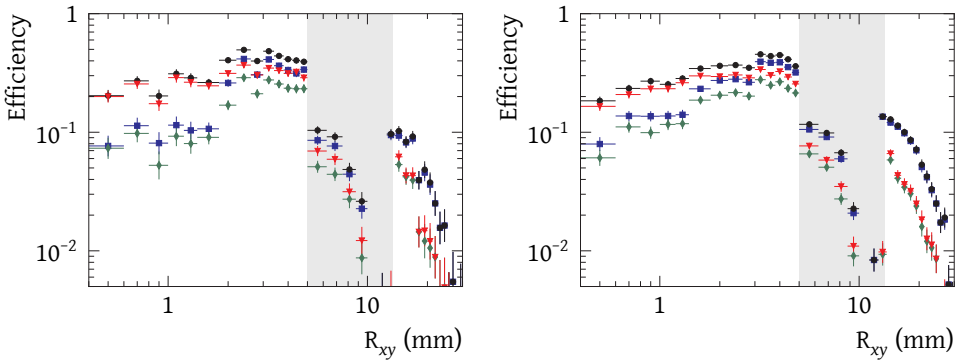

Figure 3.6: Combined HLT2 and stripping efficiency for 35 GeV/c^2 π_v particles with both B hadrons inside the acceptance ($\eta > 2$) for candidates based on a vertex reconstructed in HLT2 (blue squares), offline (red triangles) or their combination (black circles), for the 2011 (left) and 2012 (right) trigger configuration. To illustrate the overlap between the two strategies, the efficiency to select a π_v using both approaches is also drawn (green diamonds). No L0 or HLT1 decision is required, so in a sequential breakdown of the efficiency as in section 3.6, this figure fits before fig. 3.4.

The jet reconstruction used in Stripping21 uses the particle flow inputs described in section 2.3.4 with the final run 1 calorimeter calibration. Particle flow inputs with VELO information that are not compatible with the vertex are removed, after which the anti-k_T algorithm is executed with a cone size parameter of $R = 0.7$. All jets with $p_T > 5$ GeV/c are kept if they satisfy the default jet identification criteria, which will be described in detail in section 3.5.1 together with the tighter selections applied for this analysis.

3.4.4 Further offline reconstruction

As a first step of the offline analysis, the jets of the stripping candidates are updated with geometric information: a jet reference point is determined using an adaptive vertex fit of all constituents and the direction is defined by the sum of the four-momenta of the constituents. Two requirements are then applied to ensure that the jets are well reconstructed and compatible with the vertex used as a reference point to reconstruct them: jets that point more closely to a primary vertex are removed, as well as those that do not point back to their vertex within less than 2 mm. Only

candidates with at least two jets passing these criteria are retained. This selection has a typical signal efficiency above 90 %.

A technical problem appears with candidates reconstructed using both stripping lines. Due to the relatively loose selections on the compatibility of the particle flow inputs with the vertices, these will have almost identical jets if the reconstructed vertices are close. Such pairs of candidates are found by comparing the vertex position of different candidates in the event, and only the candidate based on an HLT2 vertex is retained.

3.5 Selection

The final selection of dijet candidates consists of three parts: firstly the identification of each of the jets as the result of hadronisation and decay of an energetic quark, secondly the further selection of the decay vertex to suppress B-decay and other combinatorial backgrounds, and finally the criteria applied on the dijet candidate to ensure that it is due to the two-body decay of a particle produced in the proton-proton collision.

3.5.1 Jet selection

The particle flow based jet reconstruction was designed and optimised for efficiency and energy resolution, in order to provide the most inclusive view of the event possible: the selections applied on the inputs to remove ghost and other misreconstructed particles are only as tight as is possible without compromising the efficiency. For a search, however, especially if low-p_{T} jets are included, it is beneficial to apply an additional set of jet identification criteria to further suppress fake jet backgrounds.

Properties of the reconstructed jets that are typically used for this include the ratio of the p_{T} contribution from charged particles over the total jet p_{T} (charged p_{T} fraction, CPF), the maximal p_{T} fraction carried by a single charged or neutral particle (maximum track and neutral p_{T} fraction, MTF and MNF, respectively), the largest p_{T} of a single constituent (maximum p_{T}, MPT) and the minimal number of constituents needed to collect 90 % of the jet p_{T} (N90). Most of these exploit the effect that fake jets often have a single or very few constituents carrying all energy, or that they are almost completely charged or neutral. The matching of the jet to a primary or candidate decay vertex in the reconstruction relies on the presence of at least one VELO track segment among the jet constituents. For a jet based on a vertex that has been required to have a high track multiplicity $\mathrm{N}_{\mathrm{track}}^{\mathrm{vertex}}$, background can be further reduced by requiring a larger number of jet constituents with vertex

information (number with primary vertex information, NPV). The distributions of the energy contribution from each of the particle flow categories from fig. 2.11 and from each allowed track type were also studied. Figure 3.7 collects the distribution of the jet identification variables for jets reconstructed in events with a $Z \to \mu^+\mu^-$ candidate, and shows in general excellent agreement between data and simulation, also for the energy contribution from neutral objects that are not used anywhere else than in the particle flow jet reconstruction, *cf.* figs. 3.7f to 3.7h.

Eventually, the following selections are applied: CPF > 0.1, MPT > 0.9 GeV/c, N90 ≥ 3, MTF < 0.7, MNF < 0.5. In addition, each jet is required to contain at least two constituents with a VELO track segment.

Jet energy response

The jet energy response can be studied in simulation by constructing truth jets. Truth jets are created by collecting all visible detector-stable products[5] of the long-lived particle decay and clustering those using the same jet algorithm as is used for reconstructed objects. These jets exclude detector resolution, pile-up and underlying event contributions, but include fragmentation and other strong interaction effects that are inherent to the hadronic decay of the long-lived particle. Truth jets with $p_T > 5\,\text{GeV}/c$ and $2 < \eta < 4.5$ are considered reconstructible. Figure 3.8 shows the $\Delta\phi$ and $\Delta\eta$ distance between true and reconstructed jets, for all truth jets where a reconstructed jet with at least one truth-matched constituent was found, and the correlation between truth jet energy and reconstructed jet energy for reconstructible jets.

Figures 3.9a and 3.9b illustrate the dependence on true jet p_T and the primary vertex multiplicity in the event. In both cases a clear dependence is observed. Knowing that the low-p_T jet sample also has the largest contamination from fake jets made out of particles from the underlying event and pileup interactions, and suffers from a poorer direction resolution, no p_T-dependent jet energy correction is applied, as it would only increase the combinatorial background in the signal region.

3.5.2 Vertex selection and vertex-jet overlap

Close to the beam line, the candidate sample is largely dominated by B decay vertices combined with random tracks, especially the subsample where the vertex was reconstructed in the stripping and subject to minimal selections. Therefore, a set of vertex multiplicity and invariant mass selections is applied to align it more with the

[5] all particle types with an average lifetime of at least 1 m/c, removing their decay products and all neutrinos

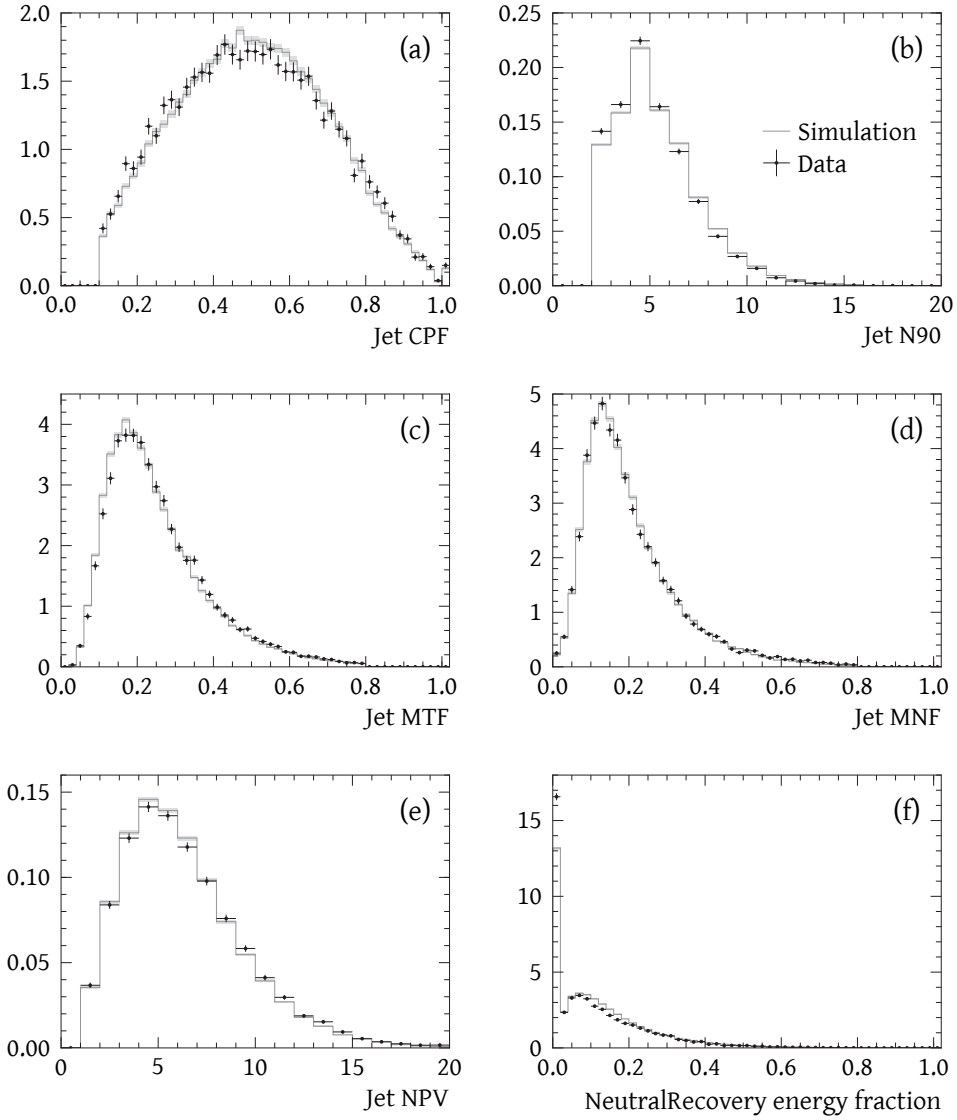

Figure 3.7: Distribution of jet identification variables for jets that are kinematically selected to be back-to-back with a $Z \to \mu^+\mu^-$ candidate in 2012 data (black) and simulation (gray)

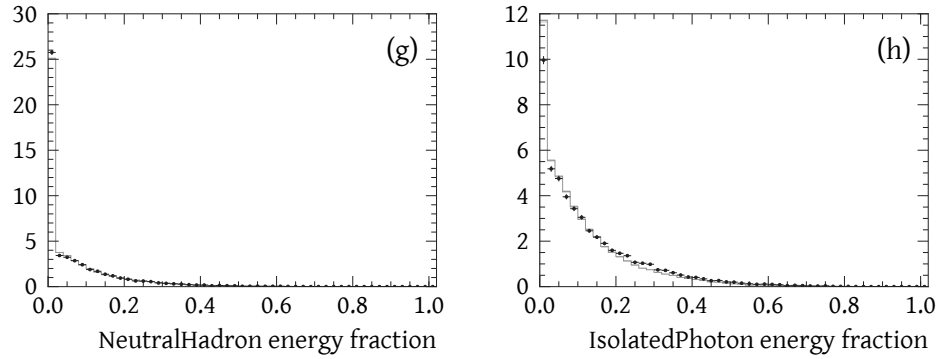

Figure 3.7: Distribution of jet identification variables for jets that are kinematically selected to be back-to-back with a $Z \to \mu^+\mu^-$ candidate in 2012 data (black) and simulation (gray) (continued)

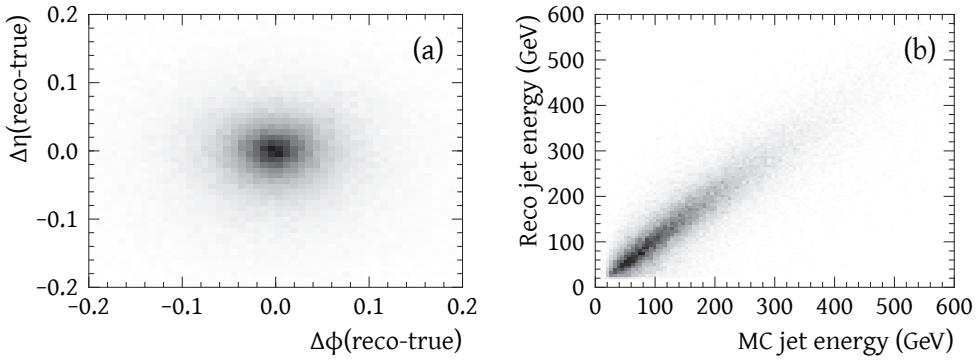

Figure 3.8: (a) kinematic distance Δη and Δϕ between truth and reconstructed jets and (b) energy correlation between reconstructed and truth jets

Table 3.3: Additional offline requirements on the minimal vertex track multiplicity N_{track}^{vertex} and minimal vertex invariant mass. The selection on the minimal number of tracks tighter than 2 shown in the last column is only applied for the 2012 data set

R_{xy}	N_{track}^{vertex}	m (GeV/c^2)	NPV
0.4–1 mm	8	10	4
1–2 mm	6		3
2–5 mm	6		2
> 5 mm	5		2

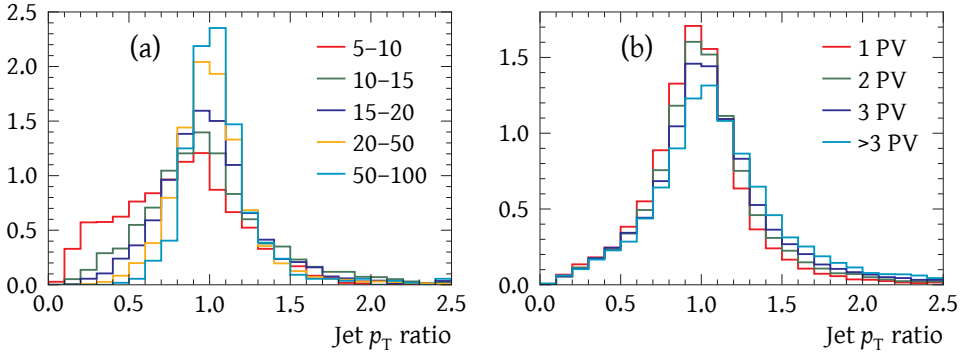

Figure 3.9: Distribution of the ratio of the reconstructed and reconstructible jet p_T (a) for several values of reconstructible jet p_T (in GeV/c), for jets in events with a single reconstructed primary vertex, and (b) for different numbers of primary vertices (distributions normalised to unity)

HLT2 vertices, which have had tight selections on these quantities applied already, but loose enough to keep the efficiency loss minimal: candidates with $R_{xy} < 1$ mm should have at least eight vertex tracks and an invariant mass of more than 10 GeV/c^2, while those up to $R_{xy} = 5$ mm need to have at least six tracks and those beyond $R_{xy} = 5$ mm at least five tracks. Due to the larger observed background in 2012 data, the selection on the number of VELO segments inside each jet has been tightened accordingly. Table 3.3 lists the selections for each R_{xy} category. Figure 3.10 shows the distribution of each variable in the R_{xy} bin where it removes most of the signal and illustrates the large difference between these distributions in the different R_{xy} bins, caused partly by trigger selections and partly by the larger contamination from tracks produced near the primary interactions. The larger background in 2012 may be due to the increased cross-section for heavy flavour production, a larger boost of the produced beauty and charm hadrons, and the beam optics that was slightly changed to improve the luminosity levelling procedure, but this should be a small effect.

3.5.3 Dijet candidate selection

Two final selections are applied on the dijet candidate: the first requires the momentum vector to align with the flight direction from a primary vertex to the reconstructed decay position. This is implemented by a requirement on the dijet invariant mass divided by the corrected mass $m_{\text{corr}} = \sqrt{m^2 + (p \sin\theta)^2} + p \sin\theta$, *cf.* section 3.4.2,

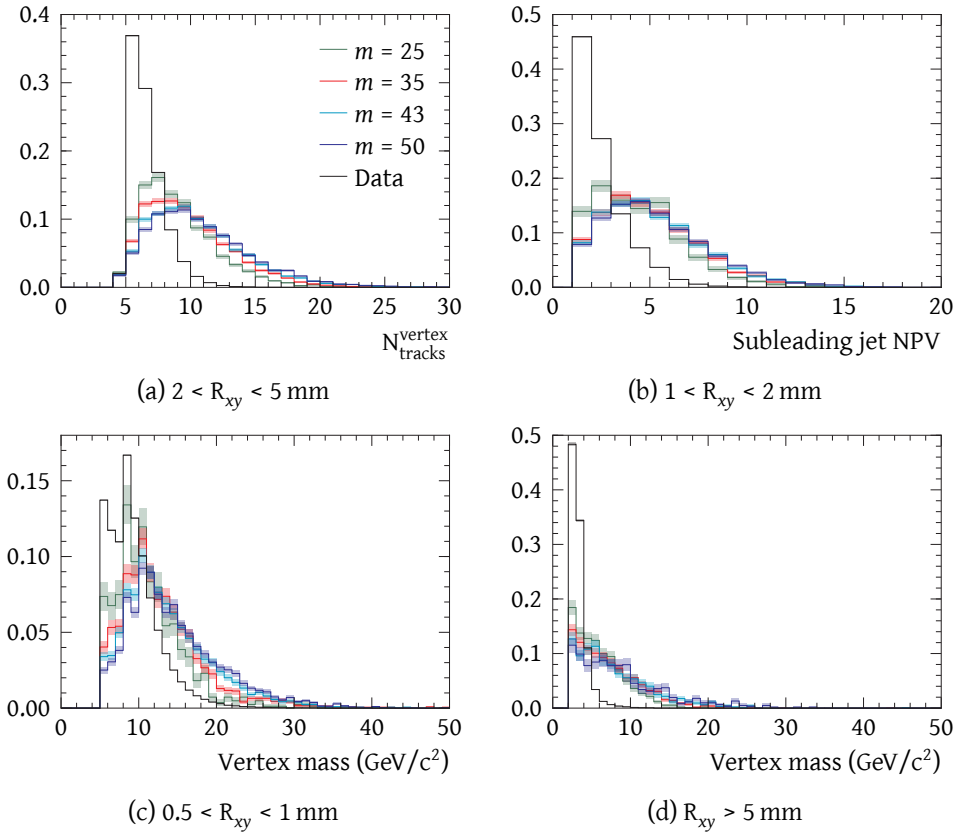

(a) $2 < R_{xy} < 5$ mm

(b) $1 < R_{xy} < 2$ mm

(c) $0.5 < R_{xy} < 1$ mm

(d) $R_{xy} > 5$ mm

Figure 3.10: Distribution (normalised to unity) of vertex related quantities in different bins of R_{xy} for 2012 data (black) and different π_v signal samples with a lifetime of 10 ps (green, red, cyan and blue for a π_v mass of 25 GeV/c^2, 35 GeV/c^2, 43 GeV/c^2 and 50 GeV/c^2, respectively) after preselection (all stripping candidates passing L0, HLT1 and HLT2)

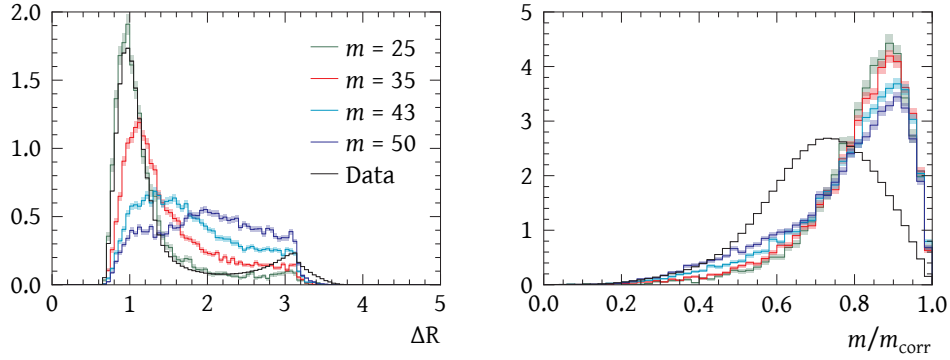

(a) (b)

Figure 3.11: Distribution (normalised to unity) of (a) the ΔR between the two jets and (b) the invariant mass divided by the corrected mass for the dijet candidate in different bins of R_{xy} for 2012 data (black) and different π_v signal samples with a lifetime of 10 ps (same colors as in fig. 3.10) after preselection

calculated using the primary vertex for which the latter quantity is minimal. The ratio m/m_{corr}, of which distributions for signal and a background-dominated data sample are shown in fig. 3.11b, should be at least 0.7.

The second selection is made on the opening angle between the jets, $\Delta R < 2.2$, in order to veto the remaining back-to-back dijet background, which can be seen around $\Delta R = \pi$ in fig. 3.11a.

3.6 Selection efficiency

This section summarizes the efficiency of the reconstruction and selection described above for the $H^0 \to \pi_v\pi_v$ signal samples. The efficiency is defined as the number of retained candidates divided by the number of events with a H^0 particle decaying to two π_v particles. In data, the latter is given by the product of the integrated luminosity, inclusive cross-section and branching fraction

$$N_{H^0\to\pi_v\pi_v} = \left(\int \mathcal{L}\,dt\right) \times \sigma(pp \to H^0X) \times \mathcal{B}(H^0 \to \pi_v\pi_v), \tag{3.2}$$

whereas in simulation, this corresponds to the number of signal events before applying any generator-level acceptance requirements.

The total efficiency can be split into parts as

$$\epsilon_{\text{total}} = A \times \epsilon_{\text{vertex}} \times \epsilon_{2\,\text{jets}|\text{vertex}} \times \epsilon_{\text{HLT1}|2\,\text{jets}} \times \epsilon_{\text{L0}|\text{HLT1}} \times \epsilon_{\text{selected}|\text{L0}}, \tag{3.3}$$

where ϵ_{vertex} is the probability for a signal decay in acceptance to give rise to a reconstructed vertex that passes the vertex selection criteria in the trigger and stripping, $\epsilon_{2\,\text{jets}|\text{vertex}}$ the efficiency for finding two associated jets, and $\epsilon_{i|j}$ the efficiency of the consecutive selection criteria applied on the candidate. The acceptance factor A represents the ratio of the number of events that leave sufficient decay products visible to the detector with respect to the total number of signal events. Its precise definition leaves the final result invariant, but it is helpful to quantify the inefficiency of the reconstruction and selection sequence compared to an ideal case — with the same detector geometry. For the studies in this section, the acceptance will be defined as the number of π_v particles with two weakly decaying B hadron decay products in the forward region $\eta > 2$ (A_{forward}). An additional acceptance requirement is added that specifies that the π_v particle decays inside the vertex detector, $R_{xy} < 50$ mm and $z < 610$ mm (A_{VELO}).

Tables 3.4 and 3.5 collect the signal efficiency for the different signal samples for 2011 and 2012 data taking conditions and trigger selections, respectively.

Table 3.4: Number of candidates per event (efficiency) in percent of the selection for different $H^0 \to \pi_v\pi_v$, $\pi_v \to b\bar{b}$ models in 2011 conditions. For each mass, the acceptance of the spectrometer ($A_{forward}$) and of the vertex detector (A_{VELO}) are given for a lifetime of 10 ps and 100 ps, together with the efficiency of the vertex reconstruction in the dedicated HLT2 lines or the stripping (in the latter case, the HLT2 efficiency is included), the probability to find two associated jets, and the efficiency of L0, HLT1 and the analysis selection on this candidate. The two subsamples (indicated by the "Hlt" and "Off" columns) are made independent by only considering candidates based on an offline vertex if no corresponding HLT2-based candidate is found. The statistical uncertainty on the efficiency due to the limited size of the simulated sample is at most at the level of a few percent

	25 GeV/c^2		35 GeV/c^2		43 GeV/c^2		50 GeV/c^2	
$A_{forward}$	12.260		9.469		7.579		6.242	
10 ps	Hlt	Off	Hlt	Off	Hlt	Off	Hlt	Off
A_{VELO}	12.235		9.465		7.578		6.242	
Vertex	1.088	0.756	1.470	0.901	1.548	0.758	1.461	0.620
2 jets	0.425	0.356	0.853	0.582	1.034	0.550	1.075	0.498
L0	0.358	0.330	0.765	0.552	0.947	0.527	1.003	0.479
HLT1	0.329	0.330	0.710	0.551	0.882	0.525	0.926	0.476
Final	0.223	0.226	0.445	0.357	0.458	0.268	0.327	0.167
100 ps	Hlt	Off	Hlt	Off	Hlt	Off	Hlt	Off
A_{VELO}	7.472		6.878		6.131		5.304	
Vertex	0.357	0.137	0.559	0.176	0.671	0.171	0.697	0.158
2 jets	0.132	0.067	0.288	0.113	0.416	0.125	0.490	0.128
L0	0.105	0.062	0.246	0.105	0.369	0.118	0.441	0.121
HLT1	0.080	0.062	0.196	0.105	0.304	0.117	0.366	0.121
Final	0.053	0.044	0.110	0.066	0.132	0.058	0.101	0.040

Table 3.5: Number of candidates per event (efficiency) in percent of the selection for different $H^0 \to \pi_v\pi_v$, $\pi_v \to b\bar{b}$ models in 2012 conditions, *cf.* table 3.4

	25 GeV/c^2		35 GeV/c^2		43 GeV/c^2		50 GeV/c^2	
$A_{forward}$	13.431		10.594		8.719		7.281	
10 ps	Hlt	Off	Hlt	Off	Hlt	Off	Hlt	Off
A_{VELO}	13.398		10.590		8.717		7.280	
Vertex	1.429	0.613	1.860	0.738	1.887	0.693	1.761	0.605
2 jets	0.596	0.279	1.126	0.471	1.317	0.505	1.339	0.479
L0	0.539	0.255	1.039	0.440	1.218	0.477	1.248	0.456
HLT1	0.503	0.255	0.971	0.439	1.135	0.475	1.159	0.451
Final	0.313	0.143	0.591	0.243	0.561	0.210	0.400	0.141
100 ps	Hlt	Off	Hlt	Off	Hlt	Off	Hlt	Off
A_{VELO}	8.061		7.560		6.838		6.159	
Vertex	0.476	0.105	0.710	0.146	0.811	0.137	0.894	0.132
2 jets	0.178	0.047	0.388	0.094	0.512	0.097	0.610	0.103
L0	0.155	0.042	0.349	0.086	0.467	0.090	0.557	0.099
HLT1	0.137	0.042	0.301	0.086	0.400	0.090	0.472	0.098
Final	0.084	0.027	0.172	0.052	0.179	0.043	0.140	0.031

4

Systematic uncertainties

Systematic uncertainties on the signal efficiency as determined from simulation, are obtained from studies of data-simulation differences using control samples. The contributions are evaluated in the same order as eq. (3.3). This order does not strictly follow the data flow because L0 and HLT1 efficiencies can be studied more accurately for a well-defined candidate.

4.1 Vertex reconstruction efficiency

The vertex reconstruction efficiency can be split into two parts: the track reconstruction efficiency and the vertex finding efficiency given that a sufficient number of tracks has been reconstructed. The former is dominated by the placement of the VELO tracking stations and the reconstruction algorithm, whose main features have been described in section 2.3.1, and determines the lifetime acceptance at large lifetimes. This is one of the main limitations of this analysis, but it is known to be accurately described in simulation [141] to within a few percent, also for very displaced and low-momentum tracks [142, 143]. A standard 2 % inefficiency is added for each track in a signal sample, after which all selections on the vertex are applied again. The candidate efficiency difference, ranging up to 5 % depending on the sample and trigger selections, is applied as a systematic uncertainty.

The second effect can be studied using offline selected B decays. The track parameters, which could in a fraction of the cases lead to *e.g.* a pair of tracks not passing a selection if poorly reconstructed, are known to be correctly modelled in MC. In addition, none of the selections on impact parameter and its associated χ^2 is tight, so no impact on the vertex reconstruction efficiency from the small differences is expected. More important is the influence of other tracks in the event: the vertex finding algorithm described in section 3.4.2 reduces computing time by excluding tracks from the vertex search as soon as they have been found to be sufficiently close to another track or vertex, instead of considering all possible combinations. This deteriorates the performance for low-multiplicity vertices near higher-multiplicity vertices. The loose selections that are used to make sure that heavy flavour secondaries are also included in the vertex, make this effect extend up to several mm

from the beam line. The primary vertex reconstruction is much less affected, because primary vertices with a highly inelastic collision have a much higher track multiplicity. In addition, the larger of the two vertices will still be found, so this will translate in a long tail of the primary vertex resolution, which does not a priori lead to the B-decay candidates not being selected.

The modelling of the strong interaction processes at low energy, where nonperturbative effects are important, is known to have inaccuracies: simulated samples on average contain fewer low-p_{T} particles than real data. Due to multiple scattering, the position information of the tracks reconstructed from these particles also has a large uncertainty, which increases the potential influence of these particles on the displaced vertex finding efficiency. The impact on the performance may be different between the HLT version and the offline version of the vertexing algorithm, because in the former case unfitted tracks are used, whereas the latter has access to fitted tracks, with accurate momentum information if hits are found over a sufficiently large integrated magnetic field. This changes the estimated covariance matrix of the track parameters.

A method has been developed to study the systematic uncertainty due to the vertex performance differences in data and simulation, exploiting the large available samples of B hadron decays: a B decay with a sufficient number of charged particles in the final state and long decay time should also be found as a vertex, and only afterwards rejected based on the R_{xy}, track multiplicity and kinematic criteria. A very clean case is the $\mathrm{B}^0 \to \mathrm{J}/\psi\,\mathrm{K}^{*0}$ decay with four tracks produced at the same point. Using the default stripping selection augmented with a loose additional requirement on the total decay chain fit χ^2 [144], about 150 000 candidates are selected from the 2012 data sample, and about twice as much in simulation. The vertex finding algorithm is rerun on events with a selected decay, without applying the invariant mass and track multiplicity selections used in the corresponding trigger or stripping line. This approach can be used both for the offline vertex finding, where the input tracks are stored, and the HLT vertex finding, where the appropriate version of the VELO tracking needs to be executed again because not all tracks from this reconstruction stage are stored in the raw file.

Both for data and simulation, the background is subtracted using the *sPlot* method [145] with a fit to the invariant mass distribution and the B hadron production kinematic distributions are subsequently reweighted for the simulated samples to match the data. Figure 4.1 shows the vertex finding efficiency in data and simulation for the relevant cases. For the 2011 HLT configuration, very little difference is observed, while for the 2012 HLT and Stripping21 cases, the efficiency in data is lower at small radius, but the relative difference decreases rapidly with displacement. The hypothesis that the differences are due to the simulation of the underlying event

seems to be confirmed by fig. 4.2: the primary vertex track multiplicity distribution is much higher in data than in simulation, and binned as a function of this variable, the vertex finding efficiency differences are compatible with being due to statistical fluctuations. Nevertheless, the difference observed as for the $B^0 \to J/\psi K^{*0}$ sample is assigned as a systematic, weighted with the R_{xy} distribution of the signal sample. Beyond R_{xy} = 5 mm, the same difference as between 4.5 mm is assumed. Due to the lower efficiency at low R_{xy}, this leads only to a moderate systematic uncertainty of about 4 %.

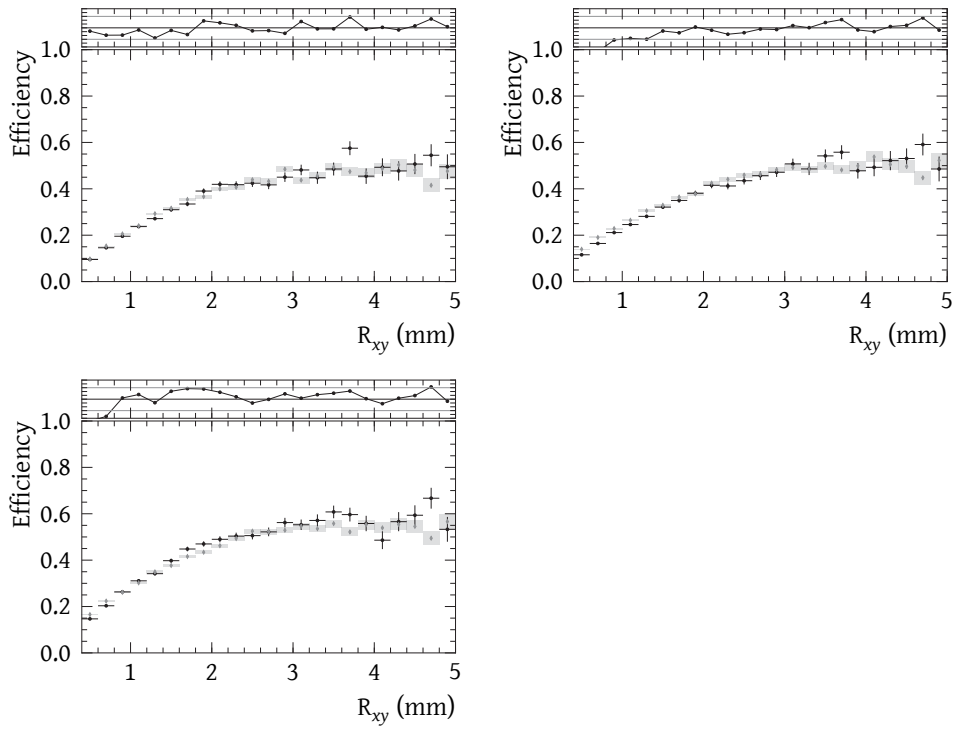

Figure 4.1: Vertex reconstruction efficiency for $B^0 \to J/\psi K^{*0}$ decays using the 2011 HLT (top left), the 2012 HLT (top right) and the Stripping21 (bottom left) configuration of the displaced vertex finding for data (black circles) and simulation (gray diamonds). The inset shows the efficiency difference, divided by its uncertainty.

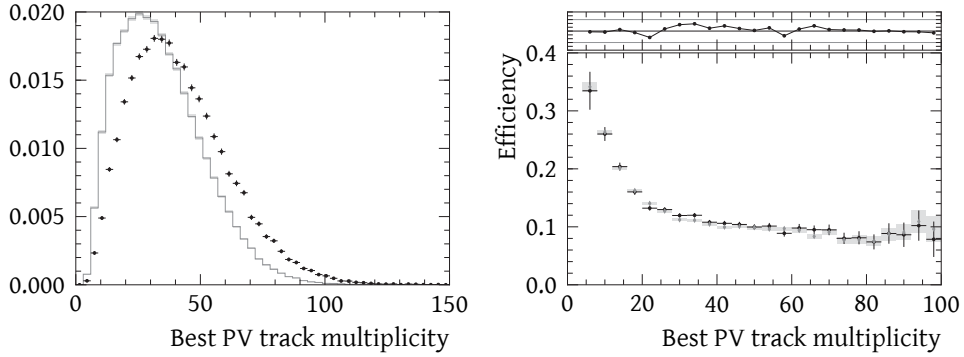

Figure 4.2: Distribution of the associated primary vertex track multiplicity for $B^0 \to J/\psi K^{*0}$ decays using the 2012 HLT configuration of the VELO track reconstruction (left) for data (black, circles) and simulation (gray, diamonds), and the efficiency of the displaced vertex finding as a function of this variable (right). The inset shows the efficiency difference, divided by its uncertainty.

4.2 Jet reconstruction efficiency and resolution

Systematic uncertainties related to the jet reconstruction can be introduced in two ways: through differences between data and simulation in the jet reconstruction efficiency and through differences in the resolution on the jet energy and direction, which enter the dijet candidate kinematic and pointing selection.

4.2.1 Reconstruction efficiency

The jet reconstruction efficiency is known from the Z + jet cross-section measurement [121] to be high and in excellent agreement between data and simulation. The only reconstruction difference with respect to that analysis is the selection of the input particles that contain a VELO track segment, but these are relatively loose and known to be well described in simulation. Only at low p_T there is some discrepancy, but VELO and upstream tracks do not contribute to the jet energy, so this effect can be neglected.

For all studies of systematic uncertainties presented here, the same reconstruction is used on the $Z \to \mu^+\mu^-$ events as on the sample of interest[1]. About 60 000 and 150 000 $Z \to \mu^+\mu^-$ candidates in respectively the 2011 and 2012 data sample pass the

[1]Technically this is done by removing the primary vertex where the Z is produced from the primary vertex list and passing it as an displaced vertex candidate input to the subsequent algorithms, such that these can be executed unmodified. The muons are banned from the particle flow.

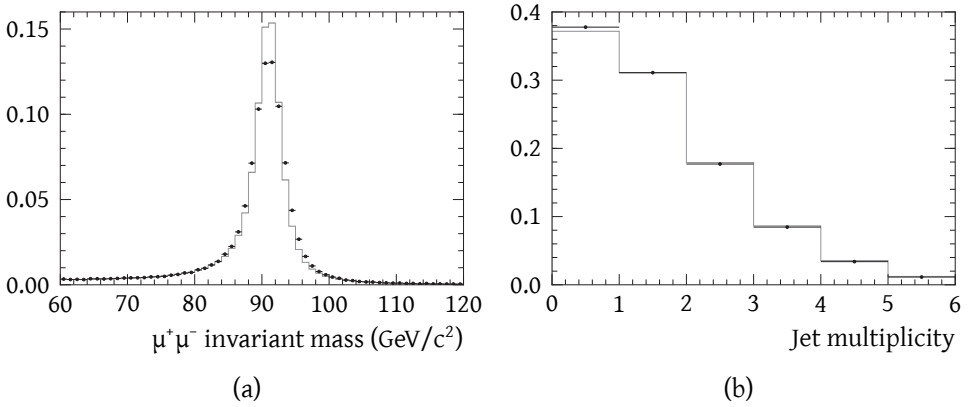

Figure 4.3: Distribution (normalised to unity) of (a) the $Z \to \mu^+\mu^-$ candidate invariant mass and (b) the number of jets associated to the corresponding primary interaction, for 2012 data (black) and simulation (gray)

preselection, similar to the one used by the cross-section measurements [146, 147] but with slightly looser trigger requirements. The simulated samples contain about five times more candidates than the corresponding data samples. Depending on the systematic effect under study, further selections are applied on the jets. Figure 4.3 shows the $Z \to \mu^+\mu^-$ invariant mass and the number of jets associated to the same primary interaction, which shows an excellent agreement between data and simulation. The difference in $\mu^+\mu^-$ mass resolution is known, *cf.* the measurement of the forward-backward asymmetry in $Z \to \mu^+\mu^-$ production [148], but has no effect for these and related studies, as the full range 60–120 GeV/c^2 is used. The data sample has a very high signal purity across this range.

The difference between data and simulation with the largest impact on the jet reconstruction efficiency is the response to very low p_T jets, close to the threshold of 5 GeV/c. Figure 4.4 shows a reasonable agreement between data and simulation of the distribution of the p_T ratio between the highest-p_T jet and the Z candidate in $Z \to \mu^+\mu^-$ events, down to low p_T. Naively one expects this distribution to tend to one — with a tail on the lower side due to QCD radiation — for sufficiently high-p_T Z candidates due to the contribution from the process where the Z recoils against a single parton that forms the hardest jet. This is only observed for the case where also the jet has rather high p_T, due to the smaller contamination from events where the recoiling jet is outside the acceptance of the LHCb detector. The sensitivity to a different response in data and simulation is evaluated by increasing the minimum jet p_T for candidates passing the full offline selection by 10 %, to 5.5 GeV/c, which

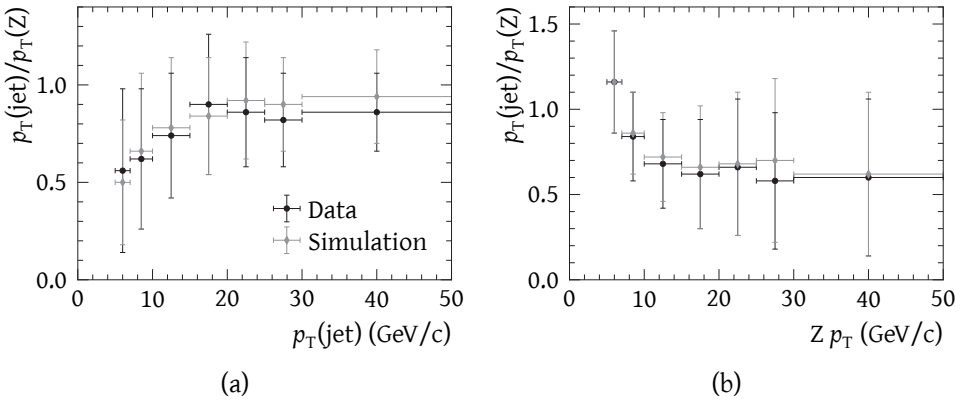

Figure 4.4: Full width at half maximum of the distribution of the p_T ratio of the leading jet and the Z candidate in $Z \rightarrow \mu^+\mu^-$ events, as a function of (a) jet p_T and (b) Z p_T. The jet is required to be back-to-back with the Z ($\Delta\phi > 3\pi/4$) and carry at least four times as much p_T than any other reconstructed jet from the same primary vertex.

affects the selection efficiency by about 2 %.

4.2.2 Jet identification efficiency

By replacing the jet identification selection with a selection on a p_T balance of the leading jet and the Z, the $Z \rightarrow \mu^+\mu^-$ sample can also be used to study the difference in jet identification efficiency between data and simulation. Figure 4.5 shows the efficiency versus jet p_T of the three different selections, when only one of them is applied and when all three are applied simultaneously. After applying the selection on the N90 variable, those on MTF and MNF have virtually no effect on real jets, which is not surprising as they all quantify that the jet energy should be distributed over several constituents — the definitions of the jet identification variables can be found in section 3.5.1. No relative difference larger than 3 % is seen, which is assigned as a systematic uncertainty.

4.2.3 Jet direction resolution

The difference in jet direction resolution is difficult to study due to the presence of other effects in the Z + jet — or any other — control sample, such as the presence of undetectable particles and additional QCD radiation. Figure 4.6b shows the width of the peak around π in the distribution of the azimuthal angle difference between the Z candidate and the leading jet in the event, obtained from fits of the sum of

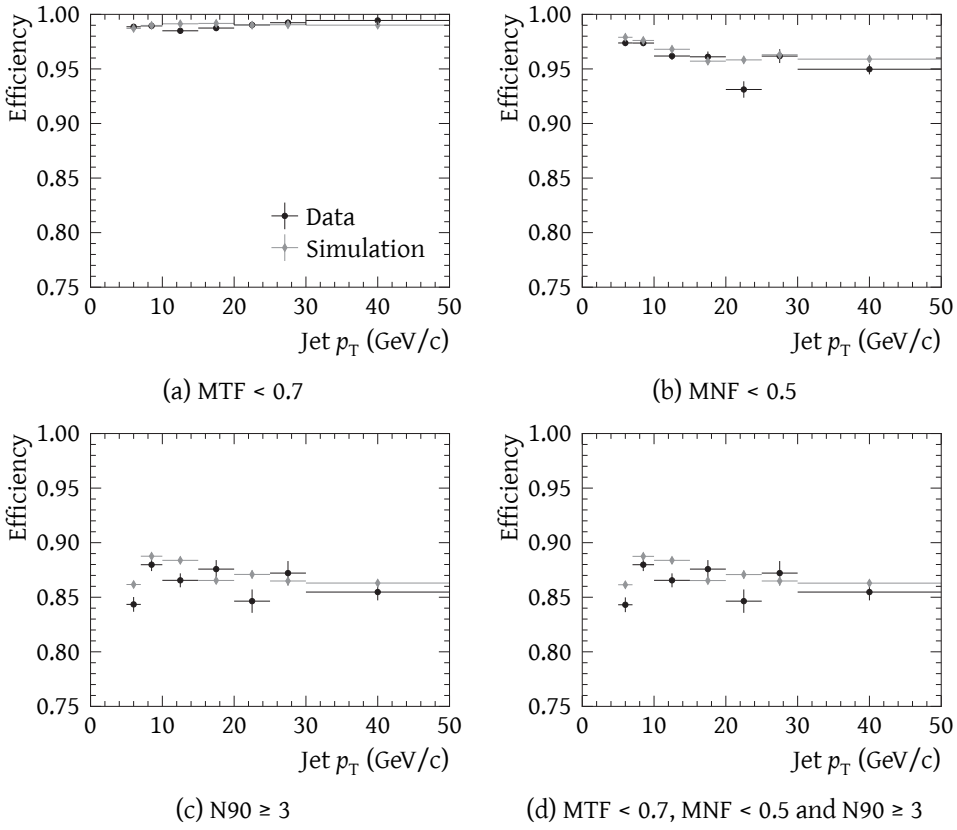

(a) MTF < 0.7

(b) MNF < 0.5

(c) N90 ≥ 3

(d) MTF < 0.7, MNF < 0.5 and N90 ≥ 3

Figure 4.5: Efficiency of the jet identification selection in data and simulation, using jets produced back-to-back with a $Z \rightarrow \mu^+\mu^-$ candidate

two Gaussians and a flat background component like fig. 4.6a. Only the jet identification criteria are applied, and a loose selection requiring that the Z and jet carry similar p_T. The Z + jet control sample mostly shows that there is no evidence for a large discrepancy between data and simulation. The direction resolution due to detector and reconstruction effects, which can be obtained from a comparison with truth jets, see fig. 4.7, is only a minor contribution to the width of the $\Delta\phi$ peak. The sensitivity of the efficiency to a difference in this resolution is studied by smearing the jet slopes independently with a factor 1.4 GeV/E inspired by fig. 4.6b, which is about one third of the resolution measured by comparing truth with reconstructed jets at low p_T, where the Z + jet sample shows the largest discrepancy. This has a 5–6 % effect on the efficiency, which is assigned as a systematic uncertainty.

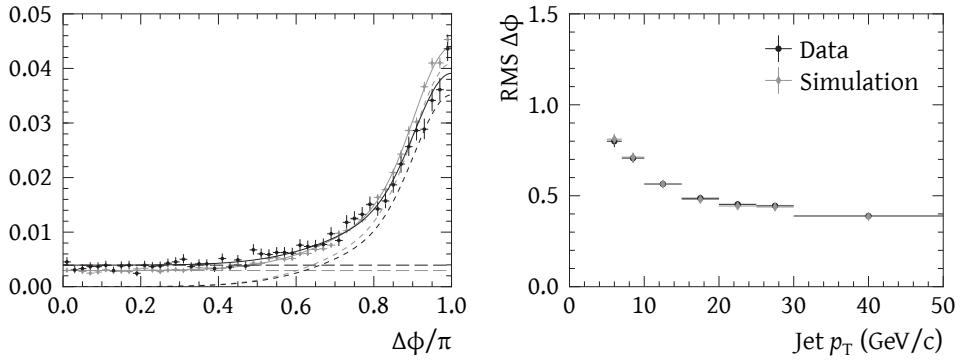

(a) $10 < p_T < 15$ GeV/c (b)

Figure 4.6: (a) $\Delta\phi$ distribution between $Z \to \mu^+\mu^-$ candidates and the leading jet from their primary vertex, with an example of the fit used to extract the root mean square (RMS) in (a). (b) shows the RMS of the $\Delta\phi$ distribution found in data and simulation as a function of the jet transverse momentum.

4.3 Trigger efficiency

4.3.1 L0 trigger

The systematic uncertainty related to the L0 trigger selections can be split into two parts: first the difference in calorimeter candidates between data and simulation, and second the difference in the distribution of the SPD hit multiplicity.

The former difference is quite small, as can be seen from fig. 4.8a, which shows the efficiency of the different calorimeter-based L0 selections for jets in $Z \to \mu^+\mu^-$

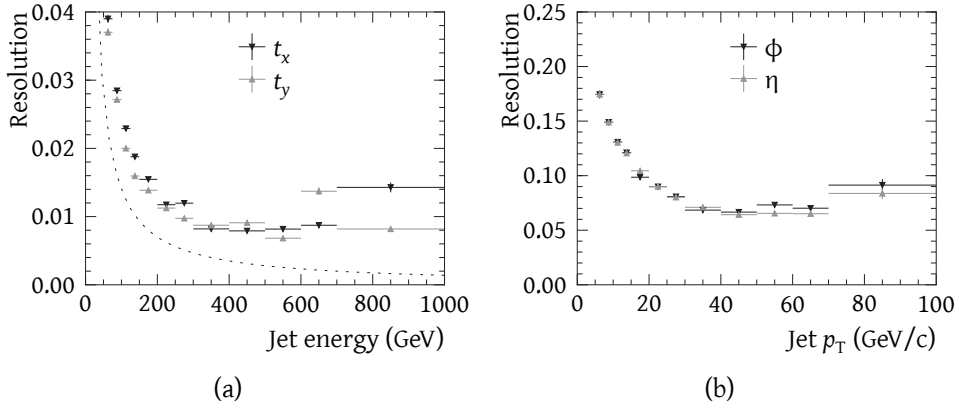

Figure 4.7: RMS of (a) the jet slope residuals and (b) the pseudorapidity and azimuthal angle residuals between truth and reconstructed jets, as a function of jet energy and transverse momentum, respectively. The dotted line in (a) shows the empirical smearing factor used to determine the resulting systematic uncertainty.

events. The events are required to be triggered based on the $Z \to \mu^+\mu^-$ candidates, such that the jet sample is unbiased. In addition, the requirement on the SPD hit multiplicity is applied beforehand because its associated systematic effect is evaluated separately. The agreement of the L0Hadron and combined efficiency in data and simulation for high-p_T jets is very good, with the largest relative difference at low p_T. The L0Photon and L0Electron efficiency, on the other hand, are significantly different in data and simulation. This is probably due to inaccuracies in the modelling of the material between the interaction point and the calorimeter system. The first muon station, for instance, which is located just in front of the calorimeter system, is known to contain more material than is assumed in simulation. Interactions of neutral hadrons can then produce charged particles, which produce a signal in the SPD detector such that the combined calorimeter object looks more similar to an electron than to a photon. The systematic uncertainty is evaluated by taking the maximal efficiency difference when folding the combined L0 efficiency with the two jet p_T spectra, where one selected jet is sufficient to retain the event. Systematic uncertainties of 2–4 % are obtained, depending on the π_v mass through the jet p_T distributions.

The SPD hit multiplicity requirement requires more attention: this variable is, due to the imperfect description of detector material and modelling of forward physics, not very well described in simulation. Most of the physics lines do not accept events with more than 600 SPD hits, which removes a non-negligible fraction

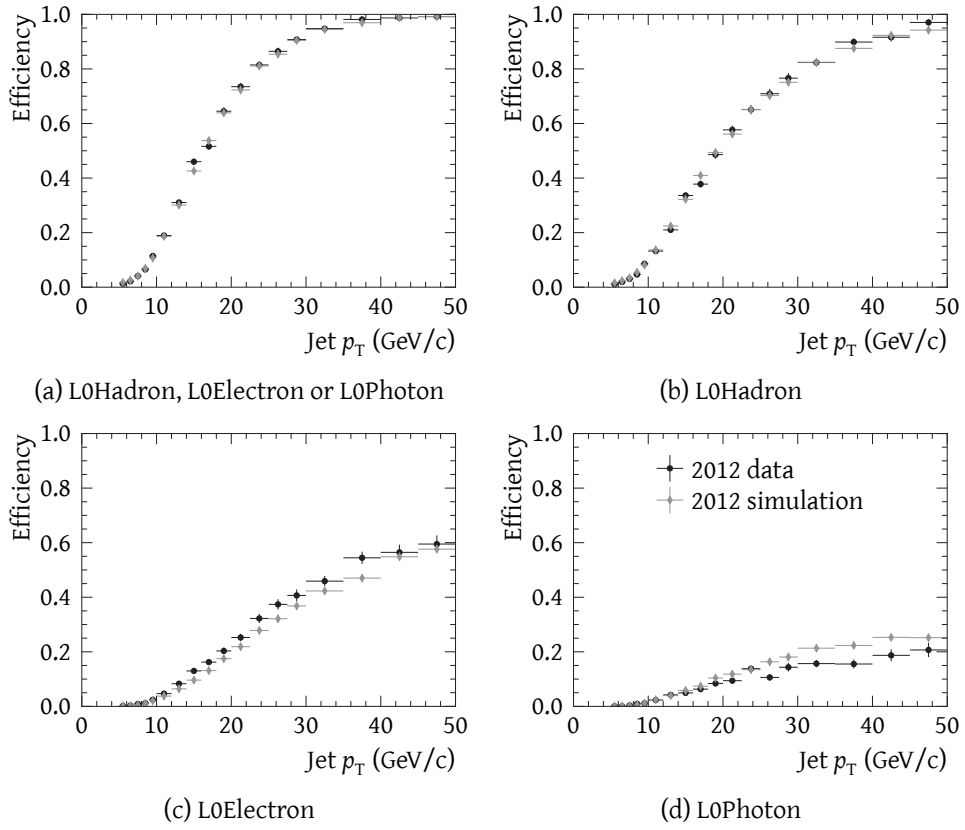

(a) L0Hadron, L0Electron or L0Photon

(b) L0Hadron

(c) L0Electron

(d) L0Photon

Figure 4.8: L0 trigger efficiency as a function of jet p_T for jets in events with a $Z \to \mu^+\mu^-$ candidate and fewer than 600 SPD hits in 2012 data and simulation

of the events. Fortunately the L0DiMuon accepts events with up to 900 SPD hits, so the sample triggered by it can be used to quantify the efficiency loss. By scaling the distribution for the SPD multiplicity in simulation, a scaling factor of about 1.5 is found to give a good agreement on the right-hand side of the distribution, for different numbers of primary vertices and for the overall distribution, *cf.* figs. 4.9 and 4.10. Applying this to the signal samples gives an efficiency correction of about 90 % for 2011 and 85 % for 2012, with a systematic uncertainty of 2–3 %, depending on the sample. The signal efficiency is corrected, and the uncertainty on this number is added as a systematic uncertainty.

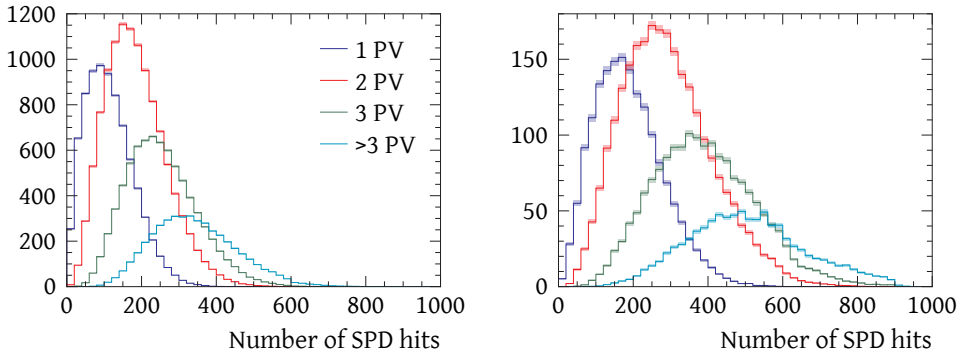

(a) Simulation (b) Data

Figure 4.9: SPD hit multiplicity for events with one, two, three or more than three reconstructed primary vertices, for (a) a simulation sample and (b) a data sample of events with a $Z \to \mu^+\mu^-$ candidate triggered by the L0DiMuon selection

4.3.2 High-level trigger

The largest systematic uncertainty related to the high-level trigger is included in the vertex reconstruction efficiency for the candidates based on a vertex reconstructed in HLT2. The other set, based on an offline reconstructed vertex, is mostly selected by the topological HLT2 selections. The performance of these is known to be accurately described in simulation. As a cross-check, their efficiency on the dimuon-triggered $B^0 \to J/\psi K^{*0}$ samples used in section 4.1 is compared, as a function of displacement, in fig. 4.11b. A maximal difference of 2–3 % is observed, which is assigned as a systematic uncertainty.

The additional selections employed in HLT1 that are the most sensitive to differences between data and simulation are the requirements on the number of observed

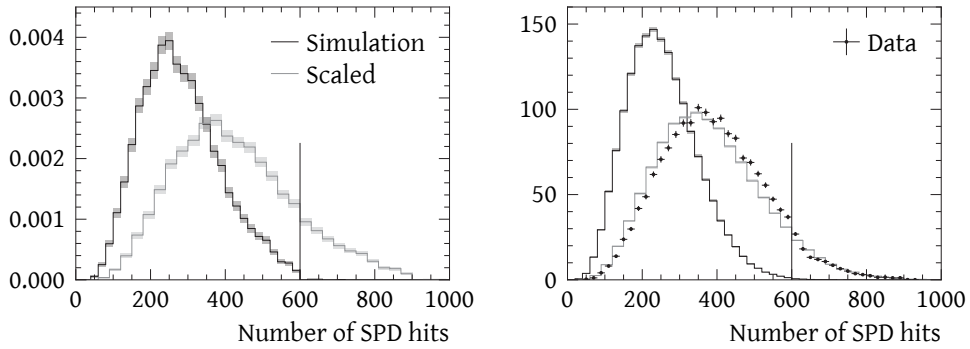

Figure 4.10: SPD hit multiplicity distribution in 2012 conditions for (a) a simulated signal sample before and after correction (scaled to unit area) and (b) for data and simulation events with a $Z \to \mu^+\mu^-$ candidate and three reconstructed primary vertices (scaled to the number in data)

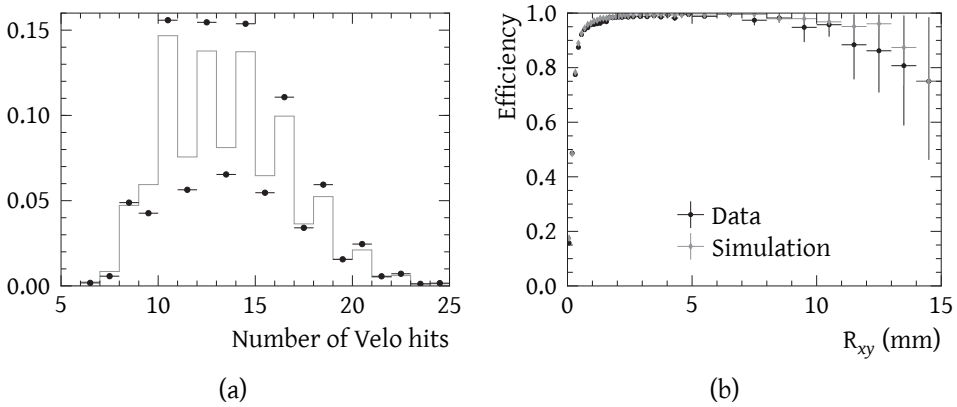

Figure 4.11: Comparison between data and simulation of (a) the number of hits per VELO track, where the markers show the distribution in data, and (b) the topological trigger efficiency versus transverse displacement, both for the dimuon-triggered $B^0 \to J/\psi K^{*0}$ sample from section 4.1

hits and on the number of hits expected from the trajectory but absent from the VELO track. Figure 4.11a shows the distribution of the former for the $B^0 \to J/\psi K^{*0}$ samples. The difference between the fractions of tracks with even and odd numbers of hits shows that the VELO hit efficiency in simulation is pessimistic compared to data. The relative efficiency difference for the requirement of at least 10 hits is 2 % for this sample. For π_v decays the final-state track multiplicity is much higher, which dilutes any effect due to a mismodelling of the hit efficiency: a decay at any position that gives one sufficiently long reconstructible VELO track will typically give several such tracks.

4.4 Total systematic uncertainties

Tables 4.1 and 4.2 show the total systematic uncertainty for each sample in 2011 and 2012 conditions, respectively. The largest contributions are related to the resolution on the jet direction, the vertex reconstruction and selection efficiency and the L0 trigger response. In all three cases, the differences are related to the description of low-energy strong interaction effects and of the description of the detector material and response, so the uncertainty could be reduced by more detailed studies of the differences between data and simulation, and by a further tuning of the description in simulation to be in agreement with data.

Table 4.1: Overview of the contributions to the systematic uncertainty on the signal efficiency and yield (in percent) for different signal samples in 2011 conditions

π_v mass (GeV/c^2)	25		35		43		50	
π_v lifetime (ps)	10	100	10	100	10	100	10	100
Vertex finding	3.8	4.2	3.3	3.9	2.8	3.7	3.7	2.6
Vertex selection	5.5	5.1	4.8	4.6	4.1	4.4	4.2	4.6
Jet reconstruction	3.1	3.1	1.6	1.6	0.7	0.7	0.5	0.5
Jet identification	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0
Jet resolution	7.0	7.0	6.0	6.0	7.4	7.4	8.5	8.5
HLT2	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0
HLT1	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0
L0	4.0	4.0	3.0	3.0	3.0	3.0	2.0	2.0
N_{SPD}	1.7	1.7	2.0	2.0	1.6	1.6	2.3	2.3
Efficiency	12.0	12.0	10.4	10.5	10.7	11.1	11.7	11.5
Luminosity	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.7
Yield	12.1	12.0	10.6	10.7	10.8	11.2	11.8	11.7

Table 4.2: Overview of the contributions to the systematic uncertainty on the signal efficiency and yield (in percent) for different signal samples in 2012 conditions

π_v mass (GeV/c^2)	25		35		43		50	
π_v lifetime (ps)	10	100	10	100	10	100	10	100
Vertex finding	4.2	4.5	3.8	4.4	3.4	4.1	3.1	3.9
Vertex selection	2.7	2.2	2.0	1.9	1.8	1.6	1.4	1.4
Jet reconstruction	2.7	2.7	1.1	1.1	0.7	0.7	0.3	0.3
Jet identification	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0
Jet resolution	5.8	5.8	5.3	5.3	6.1	6.1	7.9	7.9
HLT2	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0
HLT1	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0
L0	4.0	4.0	2.5	2.5	2.0	2.0	2.0	2.0
N_{SPD}	2.2	2.2	2.5	2.5	2.5	2.5	2.5	2.5
Efficiency	10.5	10.5	9.1	9.4	9.3	9.5	10.4	10.6
Luminosity	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
Yield	10.5	10.5	9.2	9.4	9.3	9.6	10.4	10.7

5

Results

In this chapter the statistical procedure used to quantify the sensitivity of the search and to extract results is outlined, signal efficiencies are given as a function of the long-lived particle lifetime and mass and the extracted results are presented.

5.1 Statistical procedure

The compatibility of the observed data with expectations can be quantified using a classical statistical hypothesis test: in order to establish the existence of a new particle, it should be shown that the probability to obtain a data set that is at least as incompatible with the null hypothesis of no new long-lived particles, is very small. This probability is also known as the p-value. It is evaluated by comparing the value of a function of the observed data set, the test statistic, with its expected distribution in case no signal is present, as illustrated in fig. 5.1a. By convention $p \leq 2.7 \times 10^{-3}$ and $p \leq 5.7 \times 10^{-7}$, corresponding to three and five standard deviations for a Gaussian distribution, are required to claim "evidence" or "discovery" of a new phenomenon, respectively. A confidence interval, in case of a non-observation most appropriately an upper limit, for a cross-section or signal strength parameter can also be constructed, provided that the p-value of the experiment can be determined for any value of this parameter of interest: the confidence interval at a certain confidence level CL contains all values of the parameter of interest for which a p-value larger than $1 - \mathrm{CL}$ is found.

It is often useful to combine the observed upper limit with the expected upper limit, and an estimate of the statistical fluctuations on the latter, to get a graphical representation of the upper limit and significance at the same time, for an interval or discrete set of values of a signal model parameter, *e.g.* the lifetime or mass of a new resonance. This can be done by evaluating the upper limit for the quantiles of the test statistic under the hypothesis of no signal, *e.g.* the median and the endpoints of the central 68 % and 95 % intervals that correspond to one and two standard deviations for a Gaussian distribution, in the same way as for the observed test statistic value, as is illustrated in fig. 5.1b.

A choice then needs to be made for a test statistic, and for a computationally

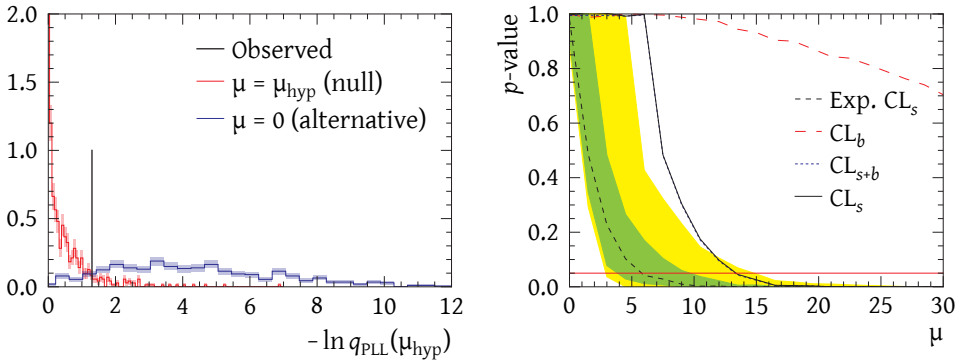

Figure 5.1: Illustration of the limit setting procedure: (a) shows the distribution of the test statistic under the signal and the no-signal hypothesis for a value of the hypothesized signal strength of the order of the obtained upper limit (the y axis range is reduced to below the height of the peak at zero for visibility); (b) illustrates how an upper limit is extracted out of a series of such distributions (the dashed black line and the green and yellow band show the median and the one and two σ quantiles of the expected CL_s distribution, respectively).

feasible method to obtain its distribution for different values of the parameter of interest. For this analysis, the profile likelihood ratio test statistic

$$q^{\mu}_{\mathrm{PLL}} = \frac{L(\mu, \hat{\theta}(\mu))}{L(\hat{\mu}, \hat{\theta})} \tag{5.1}$$

is chosen, where $L(\mu, \theta)$ denotes the likelihood as a function of the parameter of interest, the signal strength $\mu = (\sigma/\sigma^{\mathrm{SM}}_{gg\to H^0}) \cdot \mathcal{B}(H^0 \to \pi_v\pi_v)$[1], and other parameters θ that also need to be obtained from the measurement, known as nuisance parameters, $L(\mu, \hat{\theta}(\mu))$ is the maximum likelihood for a hypothesized value of μ and $L(\hat{\mu}, \hat{\theta})$ is the global maximum likelihood. The likelihood $L(\mu, \theta) = \prod_{i \in \text{data set}} P(x_i; \mu, \theta)$ requires a model of the probability density function for the observables x_i, which is in this case the sum of a background and a signal component as functions of the dijet invariant mass, for a number of R_{xy} categories. It will be described in detail in section 5.2.

This test statistic has the advantage that also systematic uncertainties can straightforwardly be taken into account. The signal distributions and efficiencies are additionally parameterized as a function of the associated unknown parameters that is

[1]The gluon fusion cross-sections at the measured standard model Brout-Englert-Higgs boson mass from [149] are used.

left floating in the likelihood minimisation, but can additionally be constrained using auxiliary measurements. In case these parameters correlate with the extracted signal yield, they will increase the constrained maximum likelihood for values of the parameter of interest close to the maximum, make the maximum less steep and decrease the significance.

The test statistic distribution can be obtained by generating pseudo-experiment data sets from the likelihood function. This is a robust method that takes the systematic uncertainties into account in a frequentist approach. It is rather time consuming, however, as a large number of pseudo-experiments need to be generated for each value of the parameter of interest, and for each of them two numerical minimisations of the likelihood need to be done. Alternatively, the asymptotic expressions for the test statistic distribution from [150] can be used to obtain approximate results. These generalise the result by Wilks [151] that in the large sample limit, for a single parameter of interest and $\mu_{\text{hyp}} = 0$, the distribution of $-2 \ln q_{\text{PLL}}$ follows a χ^2 distribution with one degree of freedom. By employing the expression obtained by Wald [152]

$$-2 \ln q_{\text{PLL}} = \frac{(\mu - \mu_{\text{hyp}})^2}{\sigma^2} + \mathcal{O}(1/\sqrt{N}), \tag{5.2}$$

generalised to more than one parameter, it is shown in [150] that $-2 \ln q_{\text{PLL}}$ follows a noncentral χ^2 distribution[2] for one degree of freedom, which reduces the problem to obtaining the noncentrality parameter $\Lambda = \frac{(\mu-\mu_{\text{hyp}})^2}{\sigma^2}$. This can be done using the second derivatives of the likelihood function around the minimum, or by generating a so-called Asimov data set, a single representative pseudo-experiment data set where the estimator of each parameter gives the assumed value, *i.e.* a pseudo-experiment with no statistical fluctuations, which is obtained in practice by exploiting the law of large numbers and generating a very large pseudo-experiment data set. The approximate formulas were used for studies of the fit model and the impact of the systematic uncertainties, and found to be in reasonable agreement with the frequentist method, which was used to produce the final results.

[2]The noncentral χ^2 distribution is a generalisation of the χ^2 distribution, and describes the distribution of the sum of squares of k independent random variables, each normally distributed with unit variance around an expectation value μ_i, which results in an additional "non-centrality" parameter $\lambda = \sum_i \mu_i^2$. The case $\lambda = 0$ corresponds to the χ^2 distribution. For one degree of freedom, the probability density function is given by [150]

$$f(x;\lambda) = \frac{1}{2\sqrt{x}\sqrt{2\pi}} \left(\exp\left(-\frac{1}{2} \left(\sqrt{x} + \sqrt{\lambda} \right)^2 \right) + \exp\left(-\frac{1}{2} \left(\sqrt{x} - \sqrt{\lambda} \right)^2 \right) \right).$$

One additional correction is applied to avoid rejecting the null hypothesis if the observation is compatible with neither the null nor the alternative hypothesis — *i.e.* to avoid obtaining an upper limit that is too low due to a downward fluctuation of the background: the CL_s method [153] prescribes the division of CL_{s+b} by CL_b (p_{null} by $p_{\mathrm{alternative}}$) to arrive at the effective p-values to extract the confidence interval from.

When generating the test statistic distributions using pseudo-experiments in the fully frequentist procedure, all p-values are re-evaluated after removing a small fraction of pseudo-experiments with failed fits, where $-\ln q_{\mathrm{PLL}} < 0$.

5.2 Likelihood model

The dijet invariant mass distribution shown in figs. 5.3 and 5.4 is modelled by the sum of a background and a signal component in six bins of R_{xy} to take advantage of the different R_{xy} distribution of background and signal. Each bin is further split into categories for each data-taking year to take the different trigger efficiencies into account. The R_{xy} bin boundaries include the values used by the trigger selections in 2011 and 2012. The largest remaining intervals are split further to avoid large differences in data set size between consecutive bins, and to make sure that the background dijet invariant mass distribution is uncorrelated with displacement in each bin.

In each of the R_{xy} bins, the signal is modelled using a Gaussian distribution whose parameters are obtained from fully simulated signal candidates. A template histogram has also been considered, but as the influence of the specific signal model was found to be small, the simpler option of an analytical shape with only a few parameters was chosen[3]. This is also numerically more stable, especially when adding a nuisance parameter that shifts the signal shape to incorporate the uncertainty on the jet energy scale.

For the background, an empirical model must be adopted, as no sufficiently large simulated sample could be generated to accurately describe it, *cf.* section 3.2.3. The profile likelihood test statistic allows to take the uncertainty on its shape into account in the same way as other systematic effects. As long as the model, including signal component, can describe the data, the statistical procedure is valid: a more flexible model leads, due to the minimisation of the two likelihoods for each pseudo-

[3] The statistical uncertainty can in a template be taken into account by adding a "true event count" nuisance parameter for every bin in the histogram, constrained by the event count in the simulated sample, as in [154], with the Poissonian constraint replaced by a Gaussian with variance given by the sum of squares of the weights of the simulated events found in that bin.

experiment data set separately, to smaller values of $-\ln q_{\mathrm{PLL}}$ and more conservative results.

All background candidates are due to combinations of tracks from heavy flavour decays and poorly reconstructed soft particle tracks, from different primary interactions or from material interactions. Two main categories can be identified: one contains candidates with a small opening angle and thus a steeply decreasing invariant mass spectrum, *e.g.* due to the combination of a heavy flavour decay vertex with tracks from another primary or material interaction. As these candidates are mostly not due to two jets, this category is strongly suppressed by the vertex and jet quality requirements. The second category is due to dijet events, and gives candidates with jets that are approximately back-to-back in the transverse plane, with naturally much higher invariant masses. This contribution is the main motivation for applying the selections on the opening angle and pointing of the candidate. After the final selection, most of the remaining candidates are due to the first component. Its invariant mass spectrum is modelled by the convolution of a steeply falling exponential distribution with a bifurcated Gaussian, such that low mass cutoff, which is not of interest, can be fitted: all parameters are left floating in the fit. The convolution is done analytically, *cf.* appendix A. The remaining contribution from the second component is modelled by a similar shape, but with a slope $\exp(-m/\alpha)$ fixed to $\alpha = 8.8\,\mathrm{GeV}/c^2$, and a fixed resolution model at the lower end only for numerical stability. The value for the slope is obtained from a fit to the high-mass part of the distribution observed at the preselection stage. The extracted results are not very sensitive to this value because of the small number of candidates at high invariant mass in each R_{xy} bin: the statistical uncertainty on α is only $0.2\,\mathrm{GeV}/c^2$, and variations over the range 6.8–$9.8\,\mathrm{GeV}/c^2$ were checked using the asymptotic approximation to give variations in the observed and expected limits of about 1 %, which is negligible compared to other sources of systematic uncertainties. The evolution of the R_{xy}-integrated dijet invariant mass distribution is illustrated in fig. 5.2, where a slope $\alpha = 8.8\,\mathrm{GeV}/c^2$ is shown by the dashed line.

All parameters of the fitted invariant mass distribution are allowed to float independently in each bin, except for the mass scale nuisance parameter and the nuisance parameter on the total signal efficiency. All relevant systematic uncertainties have been incorporated in the fit model: the total uncertainty on the efficiency obtained in chapter 4, the uncertainty on the dijet invariant mass scale and the uncertainty on the shape parameters and relative normalisation — due to the finite size of the simulated samples — of the signal model in the R_{xy} bins. Gaussian constraints representing the external knowledge on these parameters are added to the likelihood.

The fit model is implemented using RooFit [155] and the RooStats [156] tools

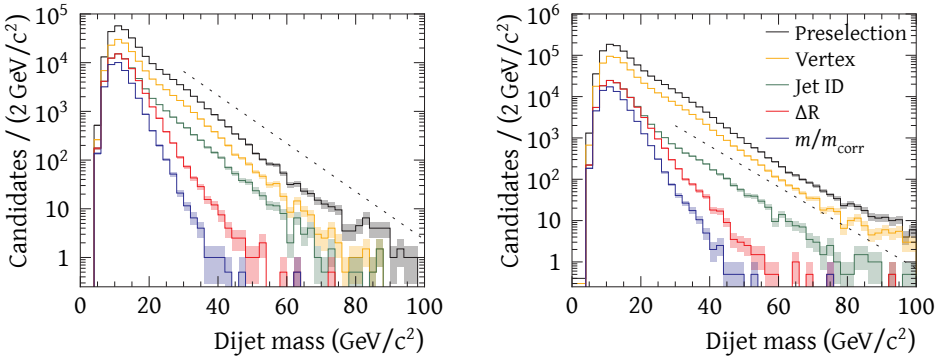

(a) 2011 (b) 2012

Figure 5.2: Invariant mass distribution observed in 2011 and 2012 data after different stages of the selection, integrated over all R_{xy} bins

are used for the limit setting procedure.

5.3 Lifetime dependence of the signal efficiency

The signal efficiency for each R_{xy} bin is obtained from the simulated samples, *cf.* section 3.6. The π_v samples with lifetimes of 10 ps and 100 ps are reweighted to mimic different lifetime values.

If there were only one sample of N_i events, with a lifetime-independent generator efficiency ϵ_{gen} and average long-lived particle lifetime τ_i available, the weight as a function of the decay time t_{gen} to obtain an average lifetime τ_{hyp} would be given by

$$w_i(t_{gen}) = \frac{\frac{1}{\tau_{hyp}} e^{-t_{gen}/\tau_{hyp}}}{\frac{N_i}{\epsilon_{gen}} \frac{1}{\tau_i} e^{-t_{gen}/\tau_i}}.$$

After weighting the candidates from both samples according to this formula, there is a residual choice to be made for the relative weight of one sample with respect to the other. The optimal choice is to make this ratio proportional to the expected number of events with a given lifetime in each sample,

$$w(t_{gen}) = w_i(t_{gen}) \cdot \frac{\frac{N_i}{\tau_i} e^{-t_{gen}/\tau_i} \tau_i}{\sum_j \frac{N_j}{\tau_j} e^{-t_{gen}/\tau_j}} = \frac{\frac{1}{\tau_{hyp}} e^{-t_{gen}/\tau_{hyp}}}{\sum_j \frac{N_j}{\epsilon_{gen}} \frac{1}{\tau_j} e^{-t_{gen}/\tau_j}},$$

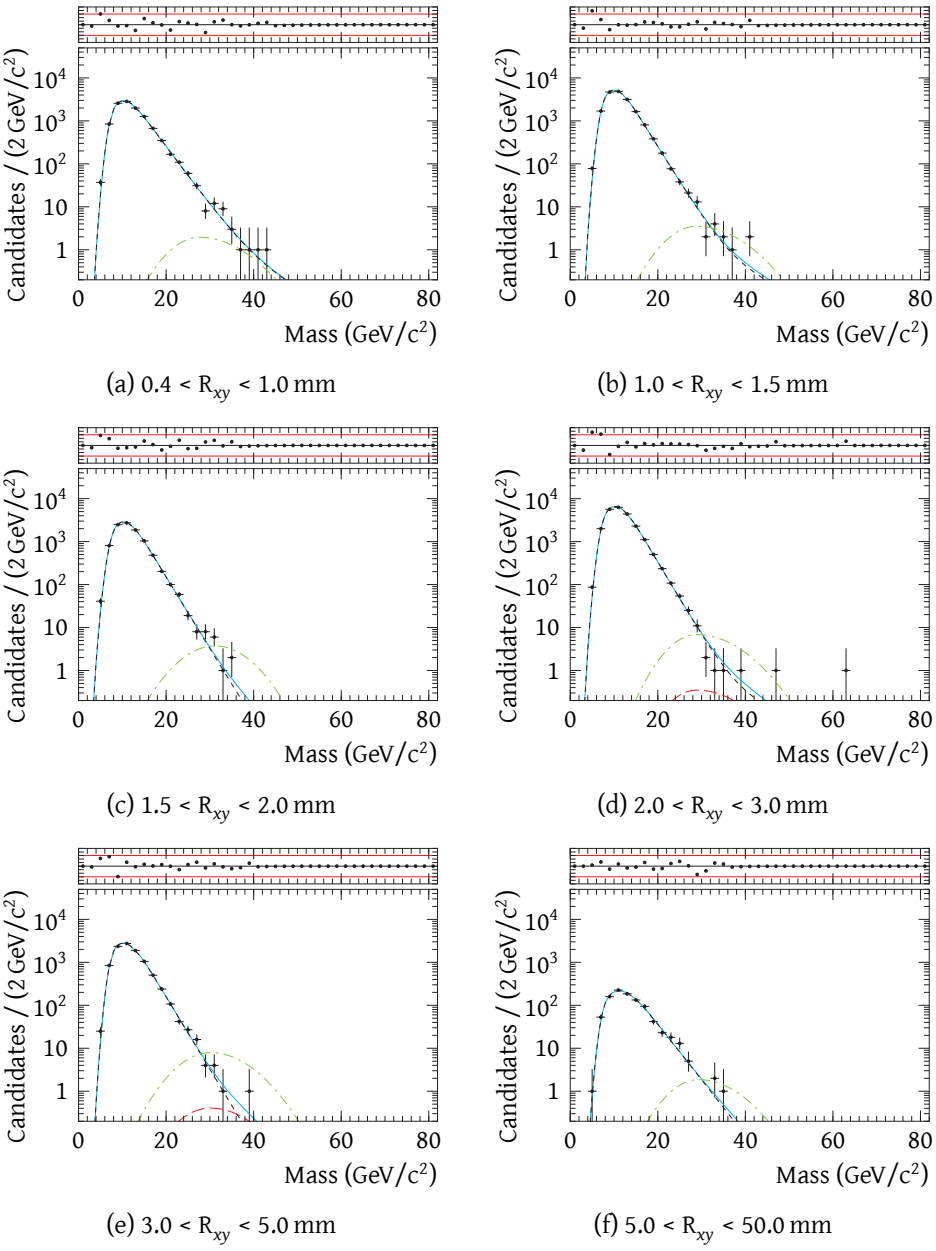

Figure 5.3: Fit to the invariant mass distribution in six bins of R_{xy} for 2011 data (black points) and a signal wit π_v mass of 35 GeV/c^2 and lifetime of 10 ps. The solid blue line indicates the total fit model, whereas the short-dashed black and long-dashed red line show the background and signal components. The long-short dashed green line shows the signal yield expected for a cross-section of 15.11 pb [149].

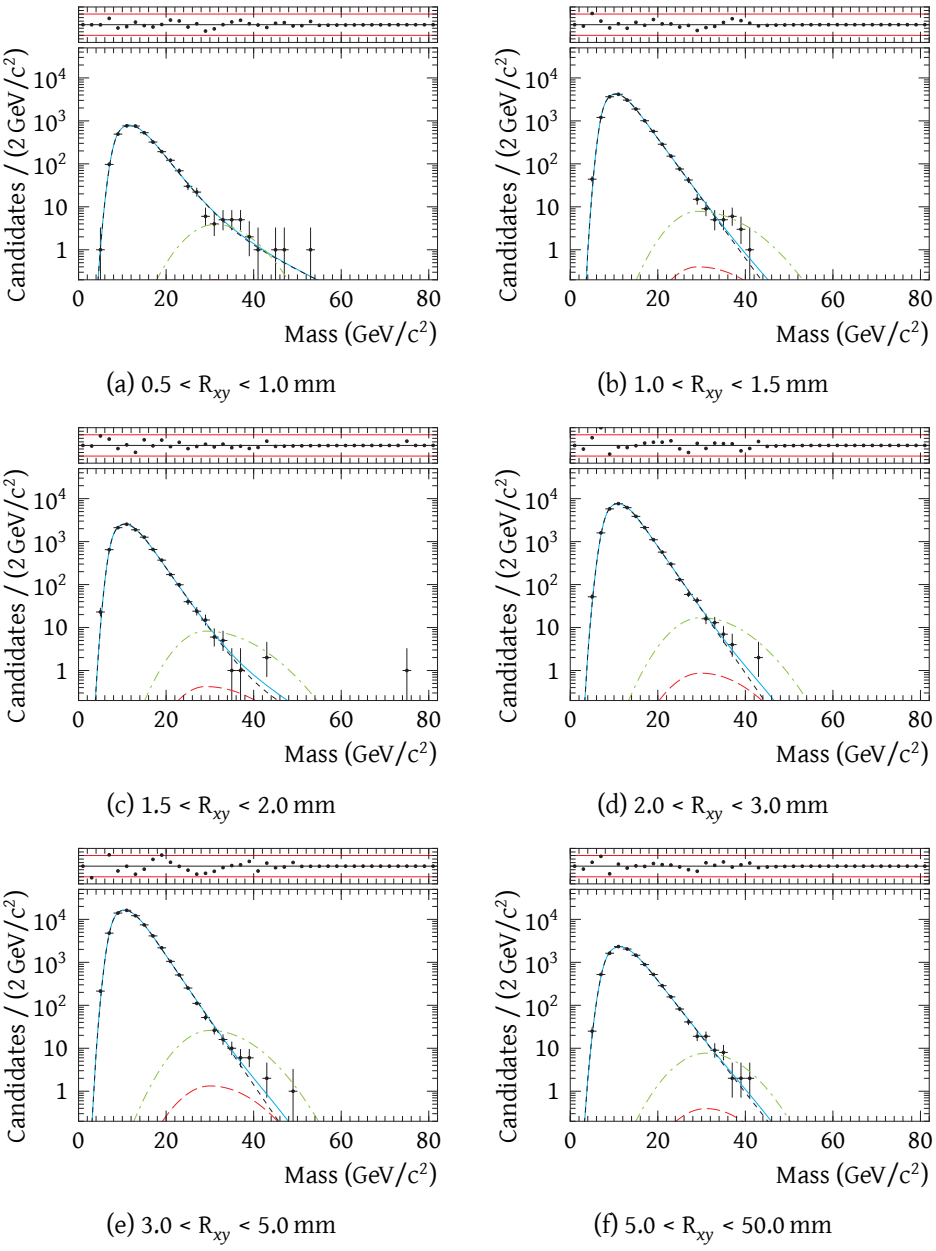

Figure 5.4: Fit to the invariant mass distribution in six bins of R_{xy} for 2012 data (black points) and a signal wit π_v mass of 35 GeV/c^2 and lifetime of 10 ps. The solid blue line indicates the total fit model, whereas the short-dashed black and long-dashed red line show the background and signal components. The long-short dashed green line shows the signal yield expected for a cross-section of 19.24 pb [149].

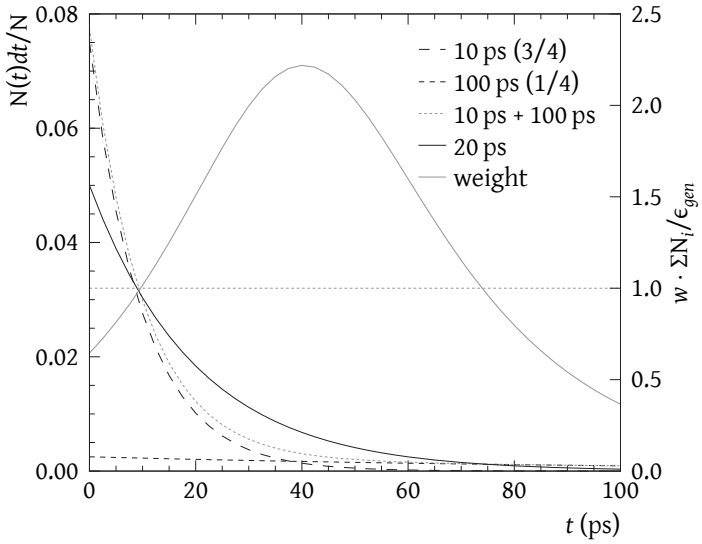

Figure 5.5: Lifetime distributions for a sample with average lifetime 10 ps (long-dashed black line), with three times more statistics than a sample with average lifetime of 100 ps (medium-dashed black line), their sum (short-dashed gray line), a lifetime distribution with an average of 20 ps (solid black line) and the combined weight for going from the first two to the latter (solid gray line). The scale for the weight is given by the *y*-axis on the right of the figure.

which gives a final expression that is independent of the sample the candidate originates from, such that the relative uncertainty on the number of events in a narrow lifetime range reduces to the optimal $\sqrt{N}/N$. The typical behaviour of the weight as a function of lifetime is illustrated in fig. 5.5.

Tables 5.1 and 5.2 show the number of events after final selection and the resulting corrected efficiency in each R_{xy} bin. The effect of the lifetime-dependent weights is illustrated by the efficiency for a number of intermediate lifetimes in table 5.3.

5.4 Upper limits

Finally, 95 % CL upper limits as a function of lifetime for Hidden Valley models with different π_v lifetimes are shown in fig. 5.6 and summarized in table 5.4.

The main differences with respect to the previously published results [2] are the generally better sensitivity through the increase of the data sample by a factor of three, and an increased sensitivity on top of that for larger lifetimes, through the

Table 5.1: 2011 data yields and contributions to the signal efficiency ($\times 10^{-4}$) for signal samples with π_v lifetime 10 ps and different masses

R_{xy} (mm)	Data	25 GeV/c^2	35 GeV/c^2	43 GeV/c^2	50 GeV/c^2
0.4–1.0	10 975	2.06 ± 0.21	6.12 ± 0.37	8.42 ± 0.43	7.02 ± 0.38
1.0–1.5	17 629	4.44 ± 0.31	11.02 ± 0.49	11.03 ± 0.49	7.73 ± 0.40
1.5–2.0	9854	4.83 ± 0.33	11.02 ± 0.49	9.57 ± 0.46	7.36 ± 0.39
2.0–3.0	22 764	10.42 ± 0.48	21.05 ± 0.68	20.38 ± 0.67	13.34 ± 0.53
3.0–5.0	9842	16.48 ± 0.60	24.60 ± 0.73	21.48 ± 0.69	13.75 ± 0.53
5.0–50.0	955	4.74 ± 0.32	5.60 ± 0.35	3.77 ± 0.29	1.75 ± 0.19

Table 5.2: 2012 data yields and contributions to the signal efficiency ($\times 10^{-4}$) for signal samples with π_v lifetime 10 ps and different masses

R_{xy} (mm)	Data	25 GeV/c^2	35 GeV/c^2	43 GeV/c^2	50 GeV/c^2
0.4–1.0	3431	1.02 ± 0.11	3.85 ± 0.20	5.22 ± 0.23	4.61 ± 0.21
1.0–1.5	16 154	3.37 ± 0.19	8.65 ± 0.30	10.06 ± 0.32	7.83 ± 0.28
1.5–2.0	9862	4.39 ± 0.22	9.28 ± 0.31	10.11 ± 0.32	7.83 ± 0.28
2.0–3.0	29 484	9.19 ± 0.32	18.51 ± 0.43	17.80 ± 0.42	13.30 ± 0.36
3.0–5.0	63 195	16.94 ± 0.43	28.00 ± 0.53	22.88 ± 0.47	14.40 ± 0.37
5.0–50.0	10 038	6.43 ± 0.27	7.64 ± 0.28	4.81 ± 0.22	2.22 ± 0.15

Table 5.3: 2012 signal efficiency contributions for the 35 GeV π_v sample, reweighted from the 10 ps and the 100 ps sample to different lifetime hypotheses

R_{xy} (mm)	10 ps ϵ ($\times 10^{-4}$)	25 ps ϵ ($\times 10^{-4}$)	50 ps ϵ ($\times 10^{-5}$)	100 ps ϵ ($\times 10^{-5}$)
0.4–1.0	3.86 ± 0.19	1.771 ± 0.085	9.28 ± 0.45	4.75 ± 0.23
1.0–1.5	8.81 ± 0.28	4.40 ± 0.14	23.76 ± 0.76	12.34 ± 0.40
1.5–2.0	9.59 ± 0.29	5.19 ± 0.16	28.77 ± 0.87	15.16 ± 0.46
2.0–3.0	18.70 ± 0.40	11.43 ± 0.25	66.2 ± 1.5	35.70 ± 0.78
3.0–5.0	28.13 ± 0.48	21.25 ± 0.36	133.1 ± 2.3	74.7 ± 1.4
5.0–50.0	8.06 ± 0.24	10.46 ± 0.29	85.3 ± 2.5	56.1 ± 1.8

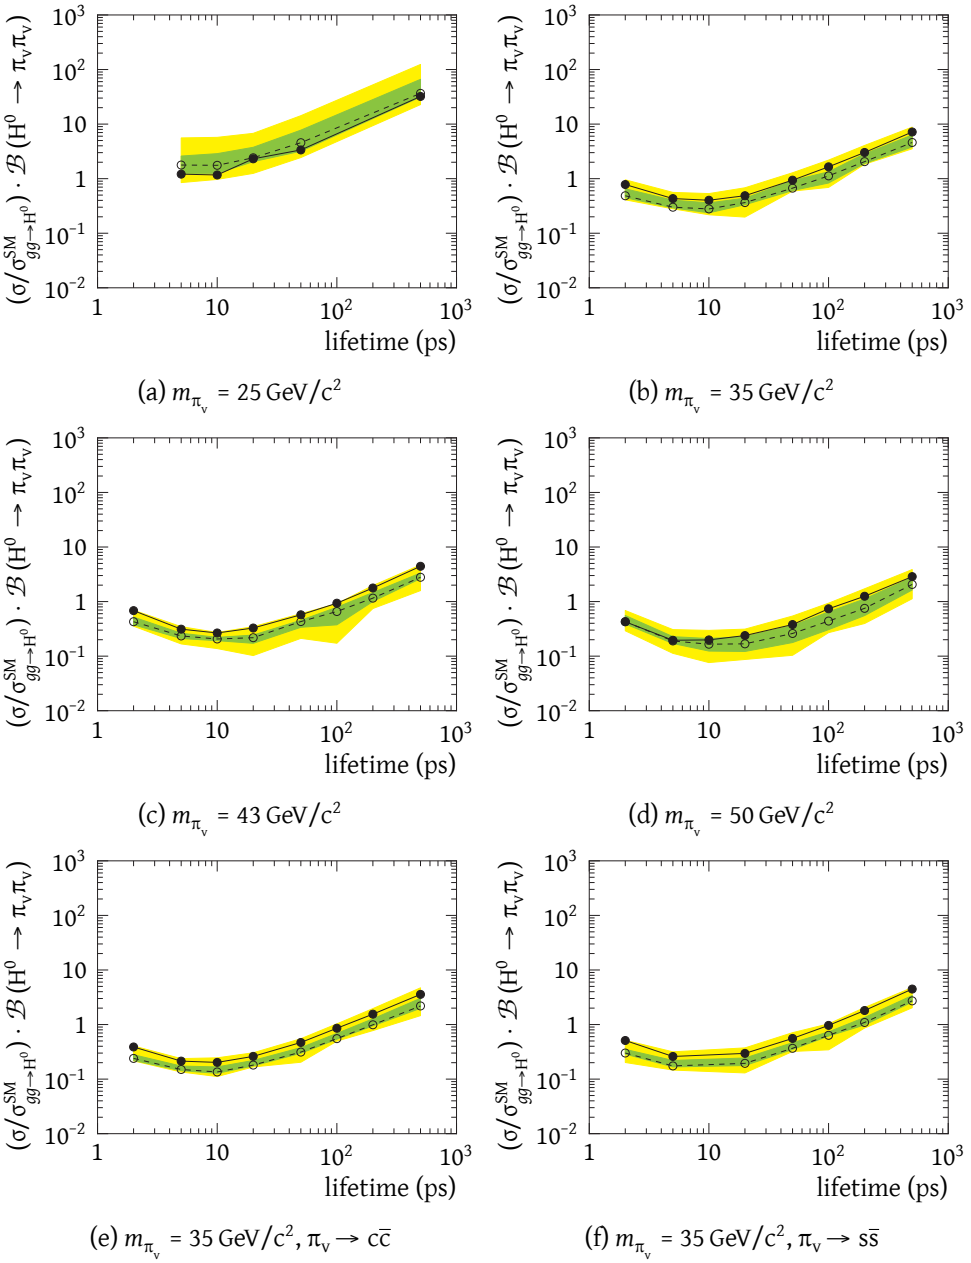

Figure 5.6: Expected (dashed line and open circles — the quantiles corresponding to ±1σ and ±2σ for a Gaussian distribution are indicated by the green and yellow band, respectively) and observed (solid line and filled circles) upper limit versus lifetime for different π_v masses and decay modes

Table 5.4: Observed 95 % CL signal strength (μ) upper limits on different π_v models

mass / lifetime (ps)	2	5	10	20	50	100	200	500
25 GeV/c^2	—	1.21	1.17	2.31	3.36	—	—	32.13
35 GeV/c^2	0.78	0.43	0.40	0.49	0.94	1.65	3.02	7.17
43 GeV/c^2	0.68	0.31	0.27	0.33	0.57	0.93	1.77	4.44
50 GeV/c^2	0.43	0.20	0.20	0.24	0.38	0.74	1.25	2.87
35 GeV/c^2, $\pi_v \to c\bar{c}$	0.39	0.21	0.20	0.26	0.47	0.86	1.55	3.57
35 GeV/c^2, $\pi_v \to s\bar{s}$	0.51	0.26	—	0.30	0.56	0.96	1.81	4.47

addition of the R_{xy} bin above 5 mm and the improved trigger selections in 2012.

5.5 Discussion and outlook

Compared to the result published as [2] that used only the 2011 data sample, and employed looser requirements on the vertices at small displacement, the efficiency is slightly decreased for small lifetimes, and increased for large lifetimes, where the obtained sensitivity is significantly better than can be expected from the larger data sample alone. The addition of a second component to the background model is, due to its anticorrelation with the signal, expected to decrease the best fit signal strength, and thus to reduce the discovery potential. The expected upper limit is expected to increase, but the impact on the observed upper limit should be small. This is confirmed by fig. 5.7, which shows the expected and observed upper limits without this alternative background model. In this case a discrepancy between expected and observed upper limit is observed, which can be traced back mostly to the low R_{xy} bin in 2012, *cf.* tables 5.5 and 5.6.

The main part of the search that could be improved is the vertex reconstruction: adding decays outside the VELO, using vertices made out of downstream tracks, could boost the sensitivity for lifetimes in the 0.1–1 ns range. Also for the region where vertices can be reconstructed, even when excluding the region around the RF foil and the collision point, the efficiency is not higher than about 50 % at a few mm, and much smaller close to the interaction region and beyond the RF foil. A reconstruction strategy that attempts a complete classsification of all tracks in an event may help to exclude the background from heavy flavour decays and material interactions. A better exploitation of the momentum information from the vertex seed finding step onwards may also be beneficial. The new vertex detector with pixel sensors that will be install as a part of the LHCb detector upgrade may also

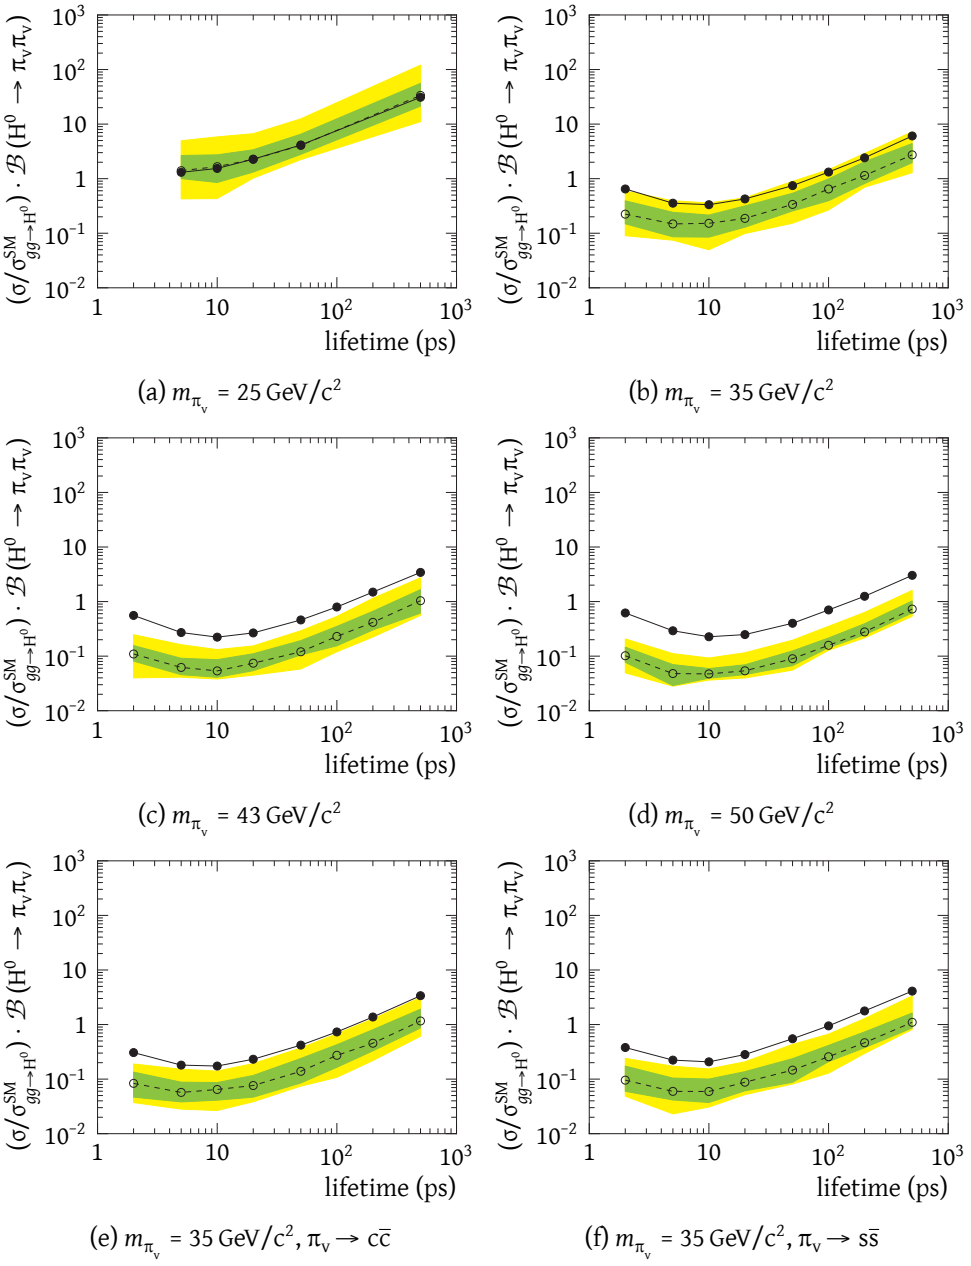

Figure 5.7: Expected (dashed line and open circles — the quantiles corresponding to ±1σ and ±2σ for a Gaussian distribution are indicated by the green and yellow band, respectively) and observed (solid line and filled circles) upper limit versus lifetime for different π_v masses and decay modes, with the one-component background model employed in [2]

Table 5.5: Best fit signal yields per bin in 2011 data, using a 10 ps π_v signal model with different masses, from a fit where the yields in every bin are floated independently and with asingle-component background model. The quoted uncertainties are approximate, and were obtained using the second-order derivatives of the likelihood.

R_{xy} (mm)	25 GeV/c^2	35 GeV/c^2	43 GeV/c^2	50 GeV/c^2
0.4–1.0	31 ± 21	18 ± 11	4.7 ± 4.9	1.5 ± 3.0
1.0–1.5	40 ± 22	11.0 ± 6.4	4.8 ± 3.7	2.5 ± 2.6
1.5–2.0	34 ± 18	6.0 ± 5.4	−1.9 ± 3.3	−11.1 ± 2.7
2.0–3.0	−5 ± 18	6.6 ± 4.4	4.2 ± 2.7	3.6 ± 2.4
3.0–5.0	19 ± 17	2.3 ± 3.5	0.9 ± 1.9	0.3 ± 1.5
5.0–50.0	5 ± 14	−2.1 ± 3.1	−4.2 ± 1.9	−45.0 ± 1.5

Table 5.6: Best fit signal yields per bin in 2012 data, using a 10 ps π_v signal model with different masses, from a fit where the yields in every bin are floated independently and with a single-component background model. The quoted uncertainties are approximate, and were obtained using the second-order derivatives of the likelihood.

R_{xy} (mm)	25 GeV/c^2	35 GeV/c^2	43 GeV/c^2	50 GeV/c^2
0.4–1.0	−153 ± 37	22.4 ± 7.9	12.9 ± 5.4	7.3 ± 4.0
1.0–1.5	−22 ± 35	10.5 ± 8.0	3.9 ± 5.0	−2.4 ± 2.9
1.5–2.0	−5 ± 20	8.0 ± 5.7	4.8 ± 3.1	3.4 ± 2.4
2.0–3.0	13 ± 39	6.4 ± 8.7	1.2 ± 3.8	−0.8 ± 2.4
3.0–5.0	−70 ± 46	15 ± 11	7.4 ± 5.5	3.3 ± 3.4
5.0–50.0	35 ± 40	11.1 ± 10.0	1.5 ± 5.3	−5.1 ± 3.2

improve the displaced vertex reconstruction capabilities, by allowing the fast reconstruction of more tracks than can currently be found within the available time budget and reducing the amount of ghost tracks.

5.6 Interpretation

The obtained results constrain the production of long-lived particles with a mass in the range 25–50 GeV/c^2, where the upper bound is given by the assumption of production in the decay of a scalar resonance with a mass of 125 GeV/c^2. In case the long-lived particles were to be produced by a different mechanism, this analysis would also be sensitive to higher long-lived particle masses, but with a different efficiency due to the acceptance. The lifetime reach is determined by the lifetime of heavy flavour hadrons, which demands $R_{xy} > 0.4$ mm in order to keep the vertex reconstruction in the trigger and stripping tractable, and the size of the region in the VELO where tracks can be reconstructed, as a VELO track (HLT1) and a vertex made out of VELO tracks (HLT2) are required in the run 1 trigger setup. The highest sensitivity is reached for a lifetime around 10 ps, but the whole interval between the B hadron lifetime of about 1.5 ps and 1 ns can be probed. This covers the parameter space region that is difficult to probe by the ATLAS and CMS experiments, which can currently record the largest data sets. The same region was probed by the Tevatron experiments CDF and D0, but with smaller data sets than LHCb.

When interpreted in terms of an exotic decay of the standard model Brout-Englert-Higgs boson to long-lived particles, the obtained results bound the branching fraction to be below about 40 %, 30 % and 20 % for long-lived particle lifetimes in the 5–10 ps range and masses of 35 GeV/c^2, 43 GeV/c^2 and 50 GeV/c^2, respectively, in the same range as the current constraints on the branching ratio to undetectable final states. For smaller and larger lifetimes, the bounds are less stringent. The experimental search was designed to only exploit the decay of a single long-lived particle itself, and not possible other features of the same collision present in specific models, and to cover as large a region in parameter space as possible. Therefore, the constraints also apply for any other model that predicts the production of long-lived particles with hadronic decays in proton-proton collisions, but the possible differences in production spectra that influence the acceptance and efficiency are taken into account. The largest effect is likely to be the limited geometrical acceptance of the LHCb detector, which favours models with a reasonably light portal over production mechanisms that involve very heavy resonances.

A

RooBifurGaussModel implementation

In this appendix, analytic expressions are derived for the convolution

$$(\mathrm{D} * \mathrm{R})(x) = \int_{-\infty}^{+\infty} \mathrm{D}(t)\mathrm{R}(x-t)\,\mathrm{d}t = \int_{-\infty}^{+\infty} \mathrm{D}(x-t)\mathrm{R}(t)\,\mathrm{d}t, \tag{A.1}$$

with D an exponential distribution

$$\mathrm{D}(x) = \begin{cases} 0 & \text{if } x < 0 \\ \frac{1}{\tau}\exp\left(-\frac{x}{\tau}\right) & \text{if } x > 0 \end{cases} \tag{A.2}$$

and R a bifurcated Gaussian distribution

$$\mathrm{R}(x) = \begin{cases} \mathrm{R_L}(x) = \sqrt{\frac{2}{\pi}}\frac{1}{\sigma_\mathrm{L}+\sigma_\mathrm{R}}\exp\left(-\frac{(x-\mu)^2}{2\sigma_\mathrm{L}^2}\right) & \text{if } x \le \mu \\ \mathrm{R_R}(x) = \sqrt{\frac{2}{\pi}}\frac{1}{\sigma_\mathrm{L}+\sigma_\mathrm{R}}\exp\left(-\frac{(x-\mu)^2}{2\sigma_\mathrm{R}^2}\right) & \text{if } x \ge \mu \end{cases} \tag{A.3}$$

and its integrals.

The expression in eq. (A.3) for the bifurcated Gaussian distribution is obtained by requiring that the shape follows a Gaussian, with different width, on both sides of $x = \mu$, that it should be continuous in $x = \mu$, and normalised to unity.

Using the definition of the convolution from eq. (A.1), its commutativity, and the two pieces of R (eq. (A.3)), we obtain

$$(\mathrm{D} * \mathrm{R})(x) = \int_{-\infty}^{+\infty} \mathrm{D}(x-t)\mathrm{R}(t)\,\mathrm{d}t \tag{A.4}$$

$$= \int_{-\infty}^{\mu} \mathrm{D}(x-t)\mathrm{R_L}(t)\,\mathrm{d}t + \int_{\mu}^{+\infty} \mathrm{D}(x-t)\mathrm{R_R}(t)\,\mathrm{d}t \tag{A.5}$$

$$= \int_{x-\mu}^{+\infty} \mathrm{D}(u)\mathrm{R_L}(x-u)\,\mathrm{d}u + \int_{-\infty}^{x-\mu} \mathrm{D}(u)\mathrm{R_R}(x-u)\,\mathrm{d}u, \tag{A.6}$$

which, using $\mathrm{D}(u) = 0$ for $u < 0$, reduces to

$$(\mathrm{D} * \mathrm{R})(x) = \int_{0}^{+\infty} \mathrm{D}(u)\mathrm{R_L}(x-u)\,\mathrm{d}u \tag{A.7}$$

for $x \leq \mu$ and to

$$(D * R)(x) = \int_0^{x-\mu} D(u) R_R(x-u)\,du + \int_{x-\mu}^{+\infty} D(u) R_L(x-u)\,du \tag{A.8}$$

for $x \geq \mu$.

Both parts will require the evaluation of an integral of the form

$$I(a,b) = \int_a^b \exp(-\alpha x) \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx, \tag{A.9}$$

which can be rewritten as

$$\begin{aligned} I(a,b) &= \int_a^b \exp\left[-\frac{1}{2\sigma^2}\left(x^2 - 2\mu x + \mu^2 + 2\sigma^2\alpha x\right)\right] dx \\ &= \int_a^b \exp\left[-\frac{1}{2\sigma^2}\left((x-(\mu-\sigma^2\alpha))^2 + \mu^2 - (\mu^2 - 2\mu\sigma^2\alpha + \sigma^4\alpha^2)\right)\right] dx \\ &= \exp\left[-\alpha\left(\mu - \frac{\sigma^2\alpha}{2}\right)\right] \int_a^b \exp\left(-\frac{(x-(\mu-\sigma^2\alpha))^2}{2\sigma^2}\right) dx \end{aligned}$$

$$I(a,b) = \exp\left[-\alpha\left(\mu - \frac{\sigma^2\alpha}{2}\right)\right] \frac{\sigma\sqrt{2\pi}}{2} \left[\operatorname{erf}\left(\frac{x-(\mu-\alpha\sigma^2)}{\sqrt{2}\sigma}\right)\right]_a^b, \tag{A.10}$$

where in the last step the definition of the error function

$$\operatorname{erf} z = \frac{2}{\sqrt{\pi}} \int_0^z \exp(-t^2)\,dt \tag{A.11}$$

has been used.

Now the convolution integral for $x \leq \mu$, eq. (A.7) can be solved as

$$\begin{aligned} (D * R)(x) &= \int_0^{+\infty} \frac{1}{\tau} \exp\left(-\frac{u}{\tau}\right) \sqrt{\frac{2}{\pi}} \frac{1}{\sigma_L + \sigma_R} \exp\left(-\frac{(x-u-\mu)^2}{2\sigma_L^2}\right) du \\ &= \frac{1}{\tau}\sqrt{\frac{2}{\pi}} \frac{1}{\sigma_L + \sigma_R} \int_0^{+\infty} \exp\left(-\frac{u}{\tau}\right) \exp\left(-\frac{(x-u-\mu)^2}{2\sigma_L^2}\right) du \end{aligned}$$

$$(D * R)(x) = \frac{1}{\tau} \frac{\sigma_L}{\sigma_L + \sigma_R} \exp\left[-\frac{1}{\tau}\left(x - \mu - \frac{\sigma_L^2}{2\tau}\right)\right] \left[1 + \operatorname{erf}\left(\frac{x-\mu-\sigma_L^2/\tau}{\sqrt{2}\sigma_L}\right)\right], \tag{A.12}$$

and, analogously, for $x \geq \mu$, eq. (A.8) as

$$(\mathrm{D} * \mathrm{R})(x) = \frac{1}{\tau}\frac{1}{\sigma_\mathrm{L} + \sigma_\mathrm{R}} \exp\left(-\frac{x-\mu}{\tau}\right)$$
$$\left\{\sigma_\mathrm{R} \exp\left(\frac{\sigma_\mathrm{R}^2}{2\tau^2}\right)\left[\operatorname{erf}\left(\frac{\sigma_\mathrm{R}}{\sqrt{2}\tau}\right) + \operatorname{erf}\left(\frac{x-\mu-\sigma_\mathrm{R}^2/\tau}{\sqrt{2}\sigma_\mathrm{R}}\right)\right]\right.$$
$$\left. + \sigma_\mathrm{L} \exp\left(\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right)\left[1 - \operatorname{erf}\left(\frac{\sigma_\mathrm{L}}{\sqrt{2}\tau}\right)\right]\right\}. \quad \text{(A.13)}$$

From eqs. (A.12) and (A.13), it is clear that an analytic expression for the integral of the convolution will require solving integrals of the form

$$J(a,b) = \int_a^b \exp(-\alpha x) \operatorname{erf}\left(\frac{x-\beta}{\sqrt{2}\sigma}\right) \mathrm{d}x, \quad \text{(A.14)}$$

which can be done using partial integration, eq. (A.11) and eq. (A.10):

$$J(a,b) = \left[-\frac{\exp(-\alpha x)}{\alpha} \operatorname{erf}\left(\frac{x-\beta}{\sqrt{2}\sigma}\right)\right]_a^b$$
$$+ \int_a^b \frac{\exp(-\alpha x)}{\alpha} \frac{2}{\sqrt{\pi}} \exp\left(-\frac{(x-\beta)^2}{2\sigma^2}\right) \frac{1}{\sqrt{2}\sigma} \mathrm{d}x$$

$$J(a,b) = \left[-\frac{\exp(-\alpha x)}{\alpha} \operatorname{erf}\left(\frac{x-\beta}{\sqrt{2}\sigma}\right)\right]_a^b$$
$$+ \frac{1}{\alpha} \exp\left[-\alpha\left(\beta - \frac{\sigma^2\alpha}{2}\right)\right]\left[\operatorname{erf}\left(\frac{x-(\beta-\alpha\sigma^2)}{\sqrt{2}\sigma}\right)\right]_a^b. \quad \text{(A.15)}$$

We now calculate a primitive function of the convolution by taking $\int_\mu^y (\mathrm{D} * \mathrm{R})(x)\,\mathrm{d}x$.

For $y \leq \mu$, this yields

$$\begin{aligned}
\int_y^{\mu} (\mathrm{D} * \mathrm{R})(x)\,\mathrm{d}x &= \frac{\sigma_\mathrm{L} \exp\left(\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right)}{\tau\,(\sigma_\mathrm{L}+\sigma_\mathrm{R})} \int_{y-\mu}^{0} \exp\left(-\frac{z}{\tau}\right) \left[1 + \operatorname{erf}\left(\frac{z - \sigma_\mathrm{L}^2/\tau}{\sqrt{2}\sigma_\mathrm{L}}\right)\right] \mathrm{d}z \\
&= \frac{\sigma_\mathrm{L} \exp\left(\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right)}{\tau\,(\sigma_\mathrm{L}+\sigma_\mathrm{R})} \left\{ \left[-\tau \exp\left(-\frac{z}{\tau}\right)\right]_{y-\mu}^{0} \right. \\
&\quad + \left[-\tau \exp\left(-\frac{z}{\tau}\right) \operatorname{erf}\left(\frac{z - \sigma_\mathrm{L}^2/\tau}{\sqrt{2}\sigma_\mathrm{L}}\right)\right]_{y-\mu}^{0} \\
&\quad \left. + \tau \exp\left[-\frac{1}{\tau}\left(\frac{\sigma_\mathrm{L}^2}{\tau} - \frac{\sigma_\mathrm{L}^2}{2\tau}\right)\right] \left[\operatorname{erf}\left(\frac{z}{\sqrt{2}\sigma_\mathrm{L}}\right)\right]_{y-\mu}^{0} \right\} \\
&= \frac{\sigma_\mathrm{L} \exp\left(\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right)}{\sigma_\mathrm{L}+\sigma_\mathrm{R}} \left\{ \exp\left(-\frac{y-\mu}{\tau}\right) - 1 \right. \\
&\quad + \operatorname{erf}\left(\frac{\sigma_\mathrm{L}}{\sqrt{2}\tau}\right) + \exp\left(-\frac{y-\mu}{\tau}\right) \operatorname{erf}\left(\frac{y - \mu - \sigma_\mathrm{L}^2/\tau}{\sqrt{2}\sigma_\mathrm{L}}\right) \\
&\quad \left. - \exp\left(-\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right) \operatorname{erf}\left(\frac{y-\mu}{\sqrt{2}\sigma_\mathrm{L}}\right) \right\}
\end{aligned}$$

$$\begin{aligned}
\int_y^{\mu} (\mathrm{D} * \mathrm{R})(x)\,\mathrm{d}x = \frac{\sigma_\mathrm{L} \exp\left(\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right)}{\sigma_\mathrm{L}+\sigma_\mathrm{R}} &\left\{ \exp\left(-\frac{y-\mu}{\tau}\right) \operatorname{erfc}\left(\frac{\sigma_\mathrm{L}^2/\tau + (\mu - y)}{\sqrt{2}\sigma_\mathrm{L}}\right) \right. \\
&\left. - \operatorname{erfc}\left(\frac{\sigma_\mathrm{L}}{\sqrt{2}\tau}\right) + \exp\left(-\frac{\sigma_\mathrm{L}^2}{2\tau^2}\right) \operatorname{erf}\left(\frac{\mu - y}{\sqrt{2}\sigma_\mathrm{L}}\right) \right\}, \quad \text{(A.16)}
\end{aligned}$$

where in the last step the complementary error function,

$$\operatorname{erfc}(z) = 1 - \operatorname{erf}(z), \quad \text{(A.17)}$$

has been introduced for notational simplicity.

For $y \geq \mu$ we find

$$\int_{\mu}^{y} (D * R)(x)\,dx = \frac{1}{\tau}\frac{1}{\sigma_L + \sigma_R}\left\{\sigma_R \exp\left(\frac{\sigma_R^2}{2\tau^2}\right) \int_0^{y-\mu} \exp\left(-\frac{z}{\tau}\right) \operatorname{erf}\left(\frac{z - \sigma_R^2/\tau}{\sqrt{2}\sigma_R}\right) dz \right.$$
$$+ \left[\sigma_R \exp\left(\frac{\sigma_R^2}{2\tau^2}\right) \operatorname{erf}\left(\frac{\sigma_R}{\sqrt{2}\tau}\right) + \sigma_L \exp\left(\frac{\sigma_L^2}{2\tau^2}\right) \operatorname{erfc}\left(\frac{\sigma_L}{\sqrt{2}\tau}\right)\right] \cdot$$
$$\left. \int_0^{y-\mu} \exp\left(-\frac{z}{\tau}\right) dz \right\}$$

$$\int_{\mu}^{y} (D * R)(x)\,dx = \frac{1}{\sigma_L + \sigma_R}\left\{\sigma_R \exp\left(\frac{\sigma_R^2}{2\tau^2}\right) \left[\exp\left(-\frac{\sigma_R^2}{2\tau^2}\right) \operatorname{erf}\left(\frac{y-\mu}{\sqrt{2}\sigma_R}\right)\right.\right.$$
$$\left. - \exp\left(-\frac{y-\mu}{\tau}\right) \operatorname{erf}\left(\frac{y - \mu - \sigma_R^2/\tau}{\sqrt{2}\sigma_R}\right) - \operatorname{erf}\left(\frac{\sigma_R}{\sqrt{2}\tau}\right)\right]$$
$$+ \left[\sigma_R \exp\left(\frac{\sigma_R^2}{2\tau^2}\right) \operatorname{erf}\left(\frac{\sigma_R}{\sqrt{2}\tau}\right) + \sigma_L \exp\left(\frac{\sigma_L^2}{2\tau^2}\right) \operatorname{erfc}\left(\frac{\sigma_L}{\sqrt{2}\tau}\right)\right] \cdot$$
$$\left.\left(1 - \exp\left(-\frac{y-\mu}{\tau}\right)\right)\right\} \quad \text{(A.18)}$$

Bibliography

[1] LHCb collaboration, *Search for (Higgs-like) bosons decaying into long-lived exotic particles*, LHCb-CONF-2012-014.

[2] LHCb collaboration, *Search for long-lived particles decaying to jet pairs*, Eur. Phys. J. C75 (2015) 152, arχiv:1412.3021.

[3] Particle Data Group, *Review of particle physics*, Chin. Phys. C38 (2014) 090001, and 2015 update.

[4] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B716 (2012) 1, arχiv:1207.7214; CMS collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, ibid. B716 (2012) 30, arχiv:1207.7235.

[5] CMS collaboration, *Precise determination of the mass of the Higgs boson and tests of compatibility of its couplings with the standard model predictions using proton collisions at 7 and 8* TeV, Eur. Phys. J. C75 (2015) 212, arχiv:1412.8662; ATLAS collaboration, *Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at* $\sqrt{s}$ *= 7 and* 8 TeV *in the ATLAS experiment*, arχiv:1507.04548, submitted to Eur.Phys.J.; ATLAS and CMS collaborations, *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC* pp *collision data at* $\sqrt{s}$ *= 7 and* 8 TeV, ATLAS-CONF-2015-044, CMS-PAS-HIG-15-002, preliminary.

[6] Muon g-2 collaboration, *Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL*, Phys. Rev. D73 (2006) 072003, arχiv:hep-ex/0602035.

[7] CMS collaboration, *Search for Lepton-Flavour-Violating Decays of the Higgs Boson*, Phys. Lett. B749 (2015) 337, arχiv:1502.07400.

[8] LHCb collaboration, *Test of lepton universality using* $B^+ \to K^+\ell^+\ell^-$ *decays*, Phys. Rev. Lett. 113 (2014) 151601, arχiv:1406.6482.

[9] BaBar collaboration, *Evidence for an excess of* $\bar{B} \to D^{(*)}\tau^-\bar{\nu}_\tau$ *decays*, Phys. Rev. Lett. 109 (2012) 101802, arχiv:1205.5442; BaBar collaboration, *Measurement of an Excess of* $\overline{B} \to D^{(*)}\tau^-\overline{\nu}_\tau$ *Decays and Implications for Charged Higgs Bosons*, Phys. Rev. D88 (2013) 072012, arχiv:1303.0571; LHCb collaboration, *Measurement of the ratio of branching fractions* $\mathcal{B}(\overline{B}^0 \to D^{*+}\tau^-\overline{\nu}_\tau)/\mathcal{B}(\overline{B}^0 \to D^{*+}\mu^-\overline{\nu}_\mu)$, Phys. Rev. Lett. 115 (2015) 111803, arχiv:1506.08614; Belle collaboration, *Measurement of the branching ratio of* $\overline{B} \to D^{(*)}\tau^-\overline{\nu}_\tau$ *relative to* $\overline{B} \to D^{(*)}\ell^-\overline{\nu}_\ell$ *decays with hadronic tagging at Belle*, Phys. Rev. D92 (2015) 072014, arχiv:1507.03233.

[10] LHCb collaboration, *Angular analysis of the* $B^0 \to K^{*0}\mu^+\mu^-$ *decay*, JHEP 02 (2016) 104, arχiv:1512.04442.

[11] ATLAS collaboration, *Search for resonances decaying to photon pairs in* 3.2 fb^{-1} *of* pp *collisions at* $\sqrt{s}$ = 13 TeV *with the ATLAS detector*, ATLAS-CONF-2015-081, preliminary.

[12] CMS collaboration, *Search for new physics in high mass diphoton events in proton-proton collisions at* 13 TeV, CMS-PAS-EXO-15-004, preliminary.

[13] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory* (1995).

[14] S. P. Martin, *A Supersymmetry primer*, FERMILAB-PUB-97-425-T, arχiv:hep-ph/9709356, [Adv. Ser. Direct. High Energy Phys.18,1(1998)].

[15] D. E. Morrissey and M. J. Ramsey-Musolf, *Electroweak baryogenesis*, New J. Phys. 14 (2012) 125003, arχiv:1206.2942; L. Canetti, M. Drewes, and M. Shaposhnikov, *Matter and Antimatter in the Universe*, ibid. 14 (2012) 095012, arχiv: 1204.4186.

[16] M. J. Strassler and K. M. Zurek, *Echoes of a hidden valley at hadron colliders*, Phys. Lett. B651 (2007) 374, arχiv:hep-ph/0604261.

[17] M. J. Strassler, *Possible effects of a hidden valley on supersymmetric phenomenology*, arχiv:hep-ph/0607160.

[18] Y. Nakai, M. Reece, and R. Sato, *SUSY Higgs Mass and Collider Signals with a Hidden Valley*, arχiv:1511.00691.

[19] M. J. Strassler and K. M. Zurek, *Discovering the Higgs through highly-displaced vertices*, Phys. Lett. B661 (2008) 263, arχiv:hep-ph/0605193.

[20] P. Schwaller, D. Stolarski, and A. Weiler, *Emerging Jets*, JHEP 05 (2015) 059, arχiv:1502.05409.

[21] T. Cohen, M. Lisanti, and H. K. Lou, *Semivisible Jets: Dark Matter Undercover at the LHC*, Phys. Rev. Lett. 115 (2015) 171804, arχiv:1503.00009.

[22] CMS collaboration, *Search for invisible decays of Higgs bosons in the vector boson fusion and associated ZH production modes*, Eur. Phys. J. C74 (2014) 2980, arχiv:1404.1344; CMS collaboration, *Search for invisible decays of Higgs bosons in the vector boson fusion production mode*, CMS-PAS-HIG-14-038, preliminary; CMS collaboration, *A combination of searches for the invisible decays of the Higgs boson using the CMS detector*, CMS-PAS-HIG-15-012, preliminary.

[23] ATLAS collaboration, *Search for Invisible Decays of a Higgs Boson Produced in Association with a Z Boson in ATLAS*, Phys. Rev. Lett. 112 (2014) 201802, arχiv: 1402.3244; ATLAS collaboration, *Search for invisible decays of the Higgs boson produced in association with a hadronically decaying vector boson in pp collisions at* $\sqrt{s} = 8$ TeV *with the ATLAS detector*, Eur. Phys. J. C75 (2015) 337, arχiv: 1504.04324; ATLAS collaboration, *Search for invisible decays of a Higgs boson using vector-boson fusion in* pp *collisions at* $\sqrt{s} = 8$ TeV *with the ATLAS detector*, arχiv:1508.07869, submitted to JHEP; ATLAS, *Constraints on new phenomena via Higgs boson couplings and invisible decays with the ATLAS detector*, JHEP 11 (2015) 206, arχiv:1509.00672.

[24] J. Chang et al., *An excess of cosmic ray electrons at energies of 300-800 GeV*, Nature 456 (2008) 362.

[25] PAMELA collaboration, *An anomalous positron abundance in cosmic rays with energies 1.5-100 GeV*, Nature 458 (2009) 607, arχiv:0810.4995; AMS collaboration, *High Statistics Measurement of the Positron Fraction in Primary Cosmic Rays of* 0.5–500 GeV *with the Alpha Magnetic Spectrometer on the International Space Station*, Phys. Rev. Lett. 113 (2014) 121101.

[26] G. Weidenspointner et al., *The sky distribution of positronium annihilation continuum emission measured with spi/integral*, Astron. Astrophys. 450 (2006) 1012, arχiv:astro-ph/0601673.

[27] E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall, *Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters*, Astrophys. J. 789 (2014) 13, arχiv:1402.2301.

[28] A. Boyarsky, O. Ruchayskiy, D. Iakubovskyi, and J. Franse, *Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster*, Phys. Rev. Lett. 113 (2014) 251301, arχiv:1402.4119.

[29] DAMA and LIBRA collaborations, *New results from DAMA/LIBRA*, Eur. Phys. J. C67 (2010) 39, arχiv:1002.1028.

[30] C. E. Aalseth et al., *Search for an Annual Modulation in a P-type Point Contact Germanium Dark Matter Detector*, Phys. Rev. Lett. 107 (2011) 141301, arχiv: 1106.0650.

[31] G. Angloher et al., *Results from 730 kg days of the CRESST-II Dark Matter Search*, Eur. Phys. J. C72 (2012) 1971, arχiv:1109.0702.

[32] CDMS collaboration, *Silicon Detector Dark Matter Results from the Final Exposure of CDMS II*, Phys. Rev. Lett. 111 (2013) 251301, arχiv:1304.4279.

[33] M. Pospelov, A. Ritz, and M. B. Voloshin, *Secluded WIMP Dark Matter*, Phys. Lett. B662 (2008) 53, arχiv:0711.4866.

[34] B. Batell, M. Pospelov, and A. Ritz, *Probing a Secluded U(1) at B-factories*, Phys. Rev. D79 (2009) 115008, arχiv:0903.0363.

[35] B. Batell, M. Pospelov, and A. Ritz, *Exploring Portals to a Hidden Sector Through Fixed Targets*, Phys. Rev. D80 (2009) 095024, arχiv:0906.5614.

[36] R. Essig, P. Schuster, and N. Toro, *Probing Dark Forces and Light Hidden Sectors at Low-Energy e+e- Colliders*, Phys. Rev. D80 (2009) 015003, arχiv:0903.3941.

[37] M. Reece and L.-T. Wang, *Searching for the light dark gauge boson in GeV-scale experiments*, JHEP 07 (2009) 051, arχiv:0904.1743.

[38] D. Curtin, R. Essig, S. Gori, and J. Shelton, *Illuminating Dark Photons with High-Energy Colliders*, JHEP 02 (2015) 157, arχiv:1412.0018.

[39] D. E. Morrissey, D. Poland, and K. M. Zurek, *Abelian Hidden Sectors at a GeV*, JHEP 0907 (2009) 050, arχiv:0904.2567.

[40] C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, *Kinetic Mixing as the Origin of Light Dark Scales*, Phys. Rev. D80 (2009) 035008, arχiv:0902.3246.

[41] A. Katz and R. Sundrum, *Breaking the Dark Force*, JHEP 06 (2009) 003, arχiv: 0902.3271.

[42] Y. F. Chan, M. Low, D. E. Morrissey, and A. P. Spray, *LHC Signatures of a Minimal Supersymmetric Hidden Valley*, JHEP 1205 (2012) 155, arχiv:1112.2705.

[43] D. E. Morrissey and A. P. Spray, *New Limits on Light Hidden Sectors from Fixed-Target Experiments*, JHEP 06 (2014) 083, arχiv:1402.4817.

[44] A. Falkowski, J. T. Ruderman, T. Volansky, and J. Zupan, *Hidden Higgs Decaying to Lepton Jets*, JHEP 05 (2010) 077, arχiv:1002.2952.

[45] A. Falkowski, J. T. Ruderman, T. Volansky, and J. Zupan, *Discovering Higgs Decays to Lepton Jets at Hadron Colliders*, Phys. Rev. Lett. 105 (2010) 241801, arχiv: 1007.3496.

[46] K. Petraki and R. R. Volkas, *Review of asymmetric dark matter*, Int. J. Mod. Phys. A28 (2013) 1330028, arχiv:1305.4939.

[47] J. McDonald, *Thermally generated gauge singlet scalars as selfinteracting dark matter*, Phys. Rev. Lett. 88 (2002) 091304, arχiv:hep-ph/0106249.

[48] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, *Freeze-In Production of FIMP Dark Matter*, JHEP 03 (2010) 080, arχiv:0911.1120.

[49] D. E. Kaplan, M. A. Luty, and K. M. Zurek, *Asymmetric Dark Matter*, Phys. Rev. D79 (2009) 115016, arχiv:0901.4117.

[50] R. T. Co, F. D'Eramo, L. J. Hall, and D. Pappadopulo, *Freeze-In Dark Matter with Displaced Signatures at Colliders*, JCAP 1512 (2015) 024, arχiv:1506.07532.

[51] R. Barbier, C. Berat, M. Besancon, M. Chemtob, A. Deandrea, E. Dudas, P. Fayet, S. Lavignac, G. Moreau, E. Perez, and Y. Sirois, *R-parity violating supersymmetry*, Phys. Rept. 420 (2005) 1, arχiv:hep-ph/0406039.

[52] H. Dreiner, M. Kramer, and B. O'Leary, *Bounds on R-parity violating supersymmetric couplings from leptonic and semi-leptonic meson decays*, Phys. Rev. D75 (2007) 114016, arχiv:hep-ph/0612278.

[53] L. M. Carpenter, D. E. Kaplan, and E.-J. Rhee, *Reduced fine-tuning in supersymmetry with R-parity violation*, Phys. Rev. Lett. 99 (2007) 211801, arχiv:hep-ph/0607204.

[54] L. M. Carpenter, D. E. Kaplan, and E. J. Rhee, *New Light Windows for Sparticle Masses and Higgs Decays in the R Parity Violating MSSM*, arχiv:0804.1581.

[55] D. E. Kaplan and K. Rehermann, *Proposal for Higgs and superpartner searches at the LHCb experiment*, JHEP 0710 (2007) 056, arχiv:0705.3426.

[56] K. Barry, P. W. Graham, and S. Rajendran, *Displaced vertices from* R*-parity violation and baryogenesis*, Phys. Rev. D89 (2014) 054003, arχiv:1310.3853.

[57] M. A. Diaz, J. C. Romao, and J. W. F. Valle, *Minimal supergravity with R-parity breaking*, Nucl. Phys. B524 (1998) 23, arχiv:hep-ph/9706315; J. W. F. Valle, *Supergravity unification with bilinear R-parity violation*, in Proceedings, 6th International Symposium on Particles, strings and cosmology (1998) pp. 502–512, arχiv:hep-ph/9808292.

[58] F. De Campos, M. A. Diaz, O. J. P. Eboli, M. B. Magro, and P. G. Mercadante, *Anomaly mediated supersymmetry breaking without R parity*, Nucl. Phys. B623 (2002) 47, arχiv:hep-ph/0110049.

[59] F. de Campos, M. A. Diaz, O. J. P. Eboli, R. A. Lineros, M. B. Magro, and P. G. Mercadante, *Neutrinos in anomaly mediated supersymmetry breaking with R-parity violation*, Phys. Rev. D71 (2005) 055008, arχiv:hep-ph/0409043.

[60] D. E. Lopez-Fogliani and C. Munoz, *Proposal for a Supersymmetric Standard Model*, Phys. Rev. Lett. 97 (2006) 041801, arχiv:hep-ph/0508297.

[61] N. Escudero, D. E. Lopez-Fogliani, C. Munoz, and R. R. de Austri, *Analysis of the parameter space and spectrum of the* μν*SSM*, JHEP 12 (2008) 099, arχiv:0810.1507.

[62] F. de Campos, O. J. P. Eboli, M. B. Magro, W. Porod, D. Restrepo, M. Hirsch, and J. W. F. Valle, *Probing bilinear R-parity violating supergravity at the LHC*, JHEP 05 (2008) 048, arχiv:0712.2156.

[63] F. de Campos, O. J. P. Eboli, M. B. Magro, and D. Restrepo, *Searching supersymmetry at the LHCb with displaced vertices*, Phys. Rev. D79 (2009) 055008, arχiv:0809.0007.

[64] F. de Campos, M. A. Diaz, O. J. P. Eboli, M. B. Magro, W. Porod, and S. Skadhauge, *CERN LHC signals for neutrino mass model in bilinear R-parity violating mAMSB*, Phys. Rev. D77 (2008) 115025, arχiv:0803.4405.

[65] P. Ghosh, D. E. Lopez-Fogliani, V. A. Mitsou, C. Munoz, and R. Ruiz de Austri, *Probing the* μ*-from-*ν *supersymmetric standard model with displaced multileptons from the decay of a Higgs boson at the LHC*, Phys. Rev. D88 (2013) 015009, arχiv:1211.3177.

[66] P. W. Graham, D. E. Kaplan, S. Rajendran, and P. Saraswat, *Displaced Supersymmetry*, JHEP 07 (2012) 149, arχiv:1204.6038.

[67] G. F. Giudice and R. Rattazzi, *Theories with gauge mediated supersymmetry breaking*, Phys. Rept. 322 (1999) 419, arχiv:hep-ph/9801271.

[68] M. Buican, P. Meade, N. Seiberg, and D. Shih, *Exploring General Gauge Mediation*, JHEP 03 (2009) 016, arχiv:0812.3668; P. Meade, N. Seiberg, and D. Shih, *General Gauge Mediation*, Prog. Theor. Phys. Suppl. 177 (2009) 143, arχiv:0801.3278.

[69] M. Kawasaki, K. Kohri, T. Moroi, and A. Yotsuyanagi, *Big-Bang Nucleosynthesis and Gravitino*, Phys. Rev. D78 (2008) 065011, arχiv:0804.3745.

[70] F. Takayama and M. Yamaguchi, *Gravitino dark matter without R-parity*, Phys. Lett. B485 (2000) 388, arχiv:hep-ph/0005214.

[71] M. Hirsch, W. Porod, and D. Restrepo, *Collider signals of gravitino dark matter in bilinearly broken R-parity*, JHEP 03 (2005) 062, arχiv:hep-ph/0503059.

[72] S. Bobrovskyi, W. Buchmuller, J. Hajer, and J. Schmidt, *Quasi-stable neutralinos at the LHC*, JHEP 09 (2011) 119, arχiv:1107.0926.

[73] W. Buchmuller, K. Hamaguchi, M. Ratz, and T. Yanagida, *Supergravity at colliders*, Phys. Lett. B588 (2004) 90, arχiv:hep-ph/0402179.

[74] W. Buchmuller, L. Covi, K. Hamaguchi, A. Ibarra, and T. Yanagida, *Gravitino Dark Matter in R-Parity Breaking Vacua*, JHEP 03 (2007) 037, arχiv:hep-ph/0702184.

[75] S. Bobrovskyi, W. Buchmuller, J. Hajer, and J. Schmidt, *Broken R-Parity in the Sky and at the LHC*, JHEP 10 (2010) 061, arχiv:1007.5007.

[76] P. Meade, M. Reece, and D. Shih, *Long-Lived Neutralino NLSPs*, JHEP 10 (2010) 067, arχiv:1006.4575.

[77] S. Bobrovskyi, J. Hajer, and S. Rydbeck, *Long-lived higgsinos as probes of gravitino dark matter at the LHC*, JHEP 02 (2013) 133, arχiv:1211.5584.

[78] L. Covi and F. Dradi, *Long-Lived stop at the LHC with or without R-parity*, JCAP 1410 (2014) 039, arχiv:1403.4923.

[79] J. L. Hewett, B. Lillie, M. Masip, and T. G. Rizzo, *Signatures of long-lived gluinos in split supersymmetry*, JHEP 09 (2004) 070, arχiv:hep-ph/0408248.

[80] Y. Cui, *Natural Baryogenesis from Unnatural Supersymmetry*, JHEP 12 (2013) 067, arχiv:1309.2952.

[81] Z. Liu and B. Tweedie, *The Fate of Long-Lived Superparticles with Hadronic Decays after LHC Run 1*, JHEP 06 (2015) 042, arχiv:1503.05923.

[82] J. A. Evans and J. Shelton, *Long-Lived Staus and Displaced Leptons at the LHC*, arχiv:1601.01326.

[83] C. Csaki, E. Kuflik, S. Lombardo, O. Slone, and T. Volansky, *Phenomenology of a Long-Lived LSP with R-Parity Violation*, JHEP 08 (2015) 016, arχiv:1505.00784.

[84] L. A. Harland-Lang, A. D. Martin, P. Motylinski, and R. S. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, Eur. Phys. J. C75 (2015) 204, arχiv:1412.3989.

[85] LHC Higgs Cross Section Working Group, *Handbook of LHC Higgs Cross Sections: 2. Differential Distributions*, CERN-2012-002, arχiv:1201.3084.

[86] C. Anastasiou, C. Duhr, F. Dulat, F. Herzog, and B. Mistlberger, *Higgs Boson Gluon-Fusion Production in QCD at Three Loops*, Phys. Rev. Lett. 114 (2015) 212001, arχiv:1503.06056.

[87] D0 collaboration, *Search for Resonant Pair Production of long-lived particles decaying to $b\bar{b}$ in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, Phys. Rev. Lett. 103 (2009) 071801, arχiv:0906.1787.

[88] CDF collaboration, *Search for heavy metastable particles decaying to jet pairs in* $p\bar{p}$ *collisions at* $\sqrt{s}$ = 1.96 TeV, Phys. Rev. D85 (2012) 012007, arχiv:1109.3136.

[89] ATLAS collaboration, *Search for displaced vertices arising from decays of new heavy particles in* 7 TeV pp *collisions at ATLAS*, Phys. Lett. B707 (2012) 478, arχiv: 1109.2242.

[90] ATLAS collaboration, *Search for long-lived, heavy particles in final states with a muon and multi-track displaced vertex in proton-proton collisions at* $\sqrt{s}$ = 7 TeV *with the ATLAS detector*, Phys. Lett. B719 (2013) 280, arχiv:1210.7451.

[91] ATLAS collaboration, *Search for massive, long-lived particles using multitrack displaced vertices or displaced lepton pairs in* pp *collisions at* $\sqrt{s}$ = 8 TeV *with the ATLAS detector*, Phys. Rev. D92 (2015) 072004, arχiv:1504.05162.

[92] ATLAS collaboration, *Search for a light Higgs boson decaying to long-lived weakly-interacting particles in proton-proton collisions at* $\sqrt{s}$ = 7 TeV *with the ATLAS detector*, Phys. Rev. Lett. 108 (2012) 251801, arχiv:1203.1303.

[93] ATLAS collaboration, *Search for long-lived, weakly interacting particles that decay to displaced hadronic jets in proton-proton collisions at* $\sqrt{s}$ = 8 TeV *with the ATLAS detector*, Phys. Rev. D92 (2015) 012010, arχiv:1504.03634.

[94] ATLAS collaboration, *Triggers for displaced decays of long-lived neutral particles in the ATLAS detector*, JINST 8 (2013) P07015, arχiv:1305.2284.

[95] ATLAS collaboration, *Search for pair-produced long-lived neutral particles decaying in the ATLAS hadronic calorimeter in* pp *collisions at* $\sqrt{s}$ = 8 TeV, Phys. Lett. B743 (2015) 15, arχiv:1501.04020.

[96] CMS collaboration, *Search for Long-Lived Neutral Particles Decaying to Quark-Antiquark Pairs in Proton-Proton Collisions at* $\sqrt{s}$ = 8 TeV, Phys. Rev. D91 (2015) 012007, arχiv:1411.6530.

[97] C. Csaki, E. Kuflik, S. Lombardo, and O. Slone, *Searching for displaced Higgs boson decays*, Phys. Rev. D92 (2015) 073008, arχiv:1508.01522.

[98] D0 collaboration, *Search for neutral, long-lived particles decaying into two muons in* $p\bar{p}$ *collisions at* $\sqrt{s}$ = 1.96 TeV, Phys. Rev. Lett. 97 (2006) 161802, arχiv: hep-ex/0607028; D0 collaboration, *Search for long-lived particles decaying into electron or photon pairs with the D0 detector*, ibid. 101 (2008) 111802, arχiv: 0806.2223.

[99] CMS collaboration, *Search in leptonic channels for heavy resonances decaying to long-lived neutral particles*, JHEP 02 (2013) 085, arχiv:1211.2472; CMS collaboration, *Search for long-lived particles that decay into final states containing two electrons or two muons in proton-proton collisions at* $\sqrt{s}$ = 8 TeV, Phys. Rev. D91 (2015) 052012, arχiv:1411.6977.

[100] CMS collaboration, *Search for long-lived particles decaying to photons and missing energy in proton-proton collisions at* $\sqrt{s}$ = 7 TeV, Phys. Lett. B722 (2013) 273, arχiv:1212.1838.

[101] ATLAS collaboration, *Searches for heavy long-lived charged particles with the ATLAS detector in proton-proton collisions at* $\sqrt{s}$ = 8 TeV, JHEP 01 (2015) 068, arχiv:1411.6795.

[102] CMS collaboration, *Searches for long-lived charged particles in* pp *collisions at* $\sqrt{s}$ = 7 TeV*and*8 TeV, JHEP 07 (2013) 122, arχiv:1305.0491.

[103] LHCb collaboration, *Search for long-lived heavy charged particles using a ring imaging Cherenkov technique at LHCb*, arχiv:1506.09173, to appear in JHEP.

[104] N. Gauvin, *Search for Higgs Bosons Decaying into Long-Lived Exotic Particles in the LHCb Experiment*, PhD thesis (Ecole Polytechnique, Lausanne), presented 25 March 2011, EPFL-THESIS-5025.

[105] J. Rouvinet, *Search for (Higgs-like) bosons decaying into a pair of long-lived exotic particles in the LHCb experiment*, PhD thesis (Ecole Polytechnique, Lausanne), presented 27 June 2014, EPFL-THESIS–6178.

[106] O. A. de Aguiar Francisco, *Procura de SUSY com violação da Paridade-R no LHCb através do canal* $\tilde{\chi}_0^1 \to \mu q' \bar{q}$, MA thesis (Rio de Janeiro Federal U., 2011), CERN-THESIS-2011-306.

[107] S. Klaver, *A Search for Long-Lived Neutralinos in LHCb*, MA thesis (University of Amsterdam, 2013).

[108] B. L. Muster, *Search for exotic massive long-lived particles decaying semileptonically in the LHCb detector*, PhD thesis (Ecole Polytechnique, Lausanne, 2015), presented 3 July 2015, EPFL-THESIS-6618.

[109] LHCb collaboration, *Search for Majorana neutrinos in* $B^- \to \pi^+\mu^-\mu^-$ *decays*, Phys. Rev. Lett. 112 (2014) 131802, arχiv:1401.5361.

[110] LHCb collaboration, *Search for hidden-sector bosons in* $B^0 \to K^{*0}\mu^+\mu^-$ *decays*, Phys. Rev. Lett. 115 (2015) 161802, arχiv:1508.04094.

[111] E. M. Löbker, *Search for long-lived exotic particles decaying to muon pairs : a study performed with* 3 fb^{-1} *data collected at the LHCb experiment in 2011 and 2012*, MA thesis (University of Twente, 2015).

[112] V. Heijne, *Search for long-lived exotic particles at LHCb*, PhD thesis (Vrije Universiteit, Amsterdam), presented 26 March 2015, CERN-THESIS-2014-294.

[113] LHCb collaboration, *The LHCb detector at the LHC*, JINST 3 (2008) S08005.

[114] LHCb collaboration, *LHCb detector performance*, Int. J. Mod. Phys. A30 (2015) 1530022, arχiv:1412.6352.

[115] M. Lamont, *Status of the LHC*, J. Phys. Conf. Ser. 455 (2013) 012001.

[116] LHCb collaboration, *First Years of Running for the LHCb Calorimeter System*, PoS TIPP2014 (2014) 030, arχiv:1407.4289.

[117] LHCb collaboration, *Production of* J/ψ *and* Υ *mesons in pp collisions at* $\sqrt{s}$ = 8 TeV, JHEP 06 (2013) 064, arχiv:1304.6977.

[118] LHCb collaboration, *Forward production of* Υ *mesons in pp collisions at* $\sqrt{s}$ = 7 *and* 8 TeV, JHEP 11 (2015) 103, arχiv:1509.02372.

[119] O. Callot, *FastVelo, a fast and efficient pattern recognition package for the Velo*, LHCb-PUB-2011-001. CERN-LHCb-PUB-2011-001, LHCb, Geneva: CERN.

[120] O. Deschamps, F. P. Machefert, M. H. Schune, G. Pakhlova, and I. Belyaev, *Photon and neutral pion reconstruction*, LHCb-2003-091.

[121] LHCb collaboration, *Study of forward* Z + *jet production in pp collisions at* $\sqrt{s}$ = 7 TeV, JHEP 01 (2014) 033, arχiv:1310.8197.

[122] V. V. Gligorov, *A single track HLT1 trigger*, LHCb-PUB-2011-003.

[123] R. Aaij and J. Albrecht, *Muon triggers in the High Level Trigger of LHCb*.

[124] R. Aaij et al., *The LHCb trigger and its performance in 2011*, JINST 8 (2013) P04022, arχiv:1211.3055.

[125] A. Puig, *The LHCb trigger in 2011 and 2012*, LHCb-PUB-2014-046.

[126] J. Albrecht, V. Gligorov, G. Raven, and S. Tolk, *Performance of the LHCb High Level Trigger in 2012*, J. Phys. Conf. Ser. 513 (2014) 012001, arχiv:1310.8544.

[127] LHCb collaboration, *Precision luminosity measurements at LHCb*, JINST 9 (2014) P12005, arχiv:1410.0149.

[128] LHCb collaboration, *Official productions of LHCb real data*, (2015)

[129] LHCb collaboration, *Stripping21 selections*, (2014)

[130] T. Sjöstrand, S. Mrenna, and P. Skands, *PYTHIA 6.4 physics and manual*, JHEP 05 (2006) 026, arχiv:hep-ph/0603175.

[131] T. Sjöstrand, S. Mrenna, and P. Skands, *A brief introduction to PYTHIA 8.1*, Comput. Phys. Commun. 178 (2008) 852, arχiv:0710.3820.

[132] I. Belyaev et al., *Handling of the generation of primary events in Gauss, the LHCb simulation framework*, J. Phys. Conf. Ser. 331 (2011) 032047.

[133] J. Pumplin, D. Stump, J. Huston, H. Lai, P. M. Nadolsky, et al., *New generation of parton distributions with uncertainties from global QCD analysis*, JHEP 0207 (2002) 012, arχiv:hep-ph/0201195.

[134] Geant4 collaboration, *Geant4 developments and applications*, IEEE Trans. Nucl. Sci. 53 (2006) 270.

[135] Geant4 collaboration, *Geant4: A simulation toolkit*, Nucl. Instrum. Meth. A506 (2003) 250.

[136] M. Clemencic et al., *The LHCb simulation application, Gauss: Design, evolution and experience*, J. Phys. Conf. Ser. 331 (2011) 032023.

[137] H. Dijkstra, N. Tuning, and N. Brook, *Some Remarks on Systematic Effects of the Trigger and Event Generator Studies*, LHCb-2003-157; S. Tolk, J. Albrecht, F. Dettori, and A. Pellegrino, *Data driven trigger efficiency determination at LHCb*, LHCb-PUB-2014-039.

[138] SLD collaboration, *A Measurement of R(b) using a vertex mass tag*, Phys. Rev. Lett. 80 (1998) 660, arχiv:hep-ex/9708015.

[139] V. V. Gligorov and M. Williams, *Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree*, JINST 8 (2013) P02013, arχiv:1210.6861.

[140] A. Bay, V. Coco, P. N. Y. David, V. Heijne, W. Hulsbergen, M. Martinelli, B. Muster, J. Rouvinet, S. Klaver, and S. Tourneur, *Exotic massive long-lived particles search in* 0.62 fb^{-1} *at LHCb*, LHCb-ANA-2014-006, Linked to LHCb-PAPER-2014-062.

[141] LHCb collaboration, *Measurement of the track reconstruction efficiency at LHCb*, JINST 10 (2015) P02007, arχiv:1408.1251.

[142] LHCb collaboration, *Prompt* K^0_S *production in pp collisions at* $\sqrt{s} = 0.9$ TeV, Phys. Lett. B693 (2010) 69, arχiv:1008.3105.

[143] LHCb collaboration, *Measurement of* V^0 *production ratios in pp collisions at* $\sqrt{s} = 0.9$ *and* 7 TeV, JHEP 08 (2011) 034, arχiv:1107.0882.

[144] W. D. Hulsbergen, *Decay chain fitting with a Kalman filter*, Nucl. Instrum. Meth. A552 (2005) 566, arχiv:physics/0503191.

[145] M. Pivk and F. R. Le Diberder, *sPlot: A statistical tool to unfold data distributions*, Nucl. Instrum. Meth. A555 (2005) 356, arχiv:physics/0402083.

[146] LHCb collaboration, *Measurement of the forward* Z *boson cross-section in pp collisions at* $\sqrt{s}$ = 7 TeV, JHEP 08 (2015) 039, arχiv:1505.07024.

[147] LHCb collaboration, *Measurement of forward* W *and* Z *boson production in pp collisions at* $\sqrt{s}$ = 8 *TeV*, JHEP 01 (2015) 155, arχiv:1511.08039.

[148] LHCb collaboration, *Measurement of the forward-backward asymmetry in* $Z/\gamma^* \to \mu^+\mu^-$ *decays and determination of the effective weak mixing angle*, JHEP 11 (2015) 190, arχiv:1509.07645.

[149] LHC Higgs Cross Section Working Group, *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties*, CERN-2013-004, arχiv:1307.1347.

[150] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C71 (2011) 1554, arχiv:1007.1727.

[151] S. S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, The Annals of Mathematical Statistics 9 (1938) 60.

[152] A. Wald, *Tests of statistical hypotheses concerning several parameters when the number of observations is large*, Transactions of the American Mathematical Society 54 (1943) 426.

[153] A. L. Read, *Presentation of search results: The CL(s) technique*, J. Phys. G28 (2002) 2693.

[154] R. J. Barlow and C. Beeston, *Fitting using finite Monte Carlo samples*, Comput. Phys. Commun. 77 (1993) 219.

[155] W. Verkerke and D. P. Kirkby, *The RooFit toolkit for data modeling*, eConf C0303241 (2003) MOLT007, arχiv:physics/0306116.

[156] L. Moneta, K. Belasco, K. S. Cranmer, S. Kreiss, A. Lazzaro, et al., *The RooStats Project*, PoS ACAT2010 (2010) 057, arχiv:1009.1003.

Summary

Are the currently known elementary particles and their interactions the most fundamental picture of nature? If yes, how did the universe as we observe it come about?

The standard model of particle physics describes the electromagnetic interaction and the strong and weak nuclear force as relativistic quantum field theories with interactions derived from gauge symmetries between twelve elementary fermions. It extends quantum electrodynamics, which very successfully describes the interactions between electrons, nuclei and light or other types of electromagnetic radiation that are relevant for processes at atomic and molecular scales. The strong nuclear force, which holds nuclei together, is described by a theory that confines the colored quarks in color-neutral bound states, *e.g.* the proton and the neutron. The weak interaction that is responsible for nuclear decays, is modelled by the exchange of charged and neutral vector bosons. The large mass of the messenger particles causes this interaction to be weak compared to the strong and electromagnetic interaction, and nuclear decays to proceed slowly.

In general the theory is in excellent agreement with experimental tests. It does, however, not explain the structure observed in the mass spectrum of the fermions and in their properties, and why the scale of electroweak symmetry breaking, which determines the the masses of the gauge bosons associated with the weak interaction, is so small: the natural scale of a more fundamental theory would be the Planck scale, where gravity needs a quantum description. Therefore, the standard model is considered incomplete, and a natural solution of the aforementioned problems prefers new effects to become visible at energy scales in the TeV range.

Another argument why the standard model is incomplete comes from cosmology. Current models can describe the large-scale evolution of the universe as well as the structure observed at small scales under the assumption that the universe contains a non-baryonic gravitationally interacting matter-like component, dark matter, and a cosmological constant or dark energy component, besides the matter that we observe. Both are well-motivated — dark matter is also needed to describe the observed dynamics of galaxies; the accelerated expansion of the universe is an experimental fact — but their precise nature and origin are not known. New inter-

actions are also required to arrive at a correct prediction of the amount of matter produced in the early universe.

These questions motivate the search for a more complete description of nature, and for experimental evidence to guide its development. Over the last decades, many models for physics beyond the standard model have been proposed, but so far experimental results have not given conclusive evidence for any of those. On the contrary, the possible visible effects of new particles and interactions have been severely constrained.

The large collision data samples collected by the LHC experiments since 2009 have allowed to perform experimental tests of the standard model and searches for new particles at unprecedented energies and precision. Since most inelastic proton-proton collisions at TeV-scale energies are due to the strong interaction and most of the rare processes are accurately described by the standard model, collisions that include new states or other deviations are obscured by large backgrounds. To separate the processes of interest from background, a variety of experimental signatures is used. These include collisions with a large amount of undetected outgoing momentum in the transverse plane, inferred from the recoiling visible particles, fully reconstructible decays of resonances to known particles *etc.*

The focus of this work is the detection of new particles through their decay at a position displaced from the primary proton-proton collision point. Decays as nearby as a few mm to the collision point can be reconstructed by using the precise particle tracking capabilities close to the interaction region provided by silicon vertex detectors. These are nowadays commonly used by collider experiments to identify and study beauty and charm hadron decays, with lifetimes of the order of 1 ps, which translates to typical displacements of the order of 1 cm at the LHC.

The LHCb detector, which was specifically designed for studying heavy flavour decays, can naturally also be used to search for exotic long-lived particle decays, because its vertex detector can reconstruct tracks from decays displaced up to 20 cm along the beam line, and up to about 4 cm in the transverse direction, with high precision and efficiency. Furthermore, it is fully instrumented in the forward region, for angles with the beam line up to 250–300 mrad, with charged particle tracking and identification detectors and calorimeter detectors sensitive to all electromagnetically and strongly interacting particles.

The experimental challenge of discovering new phenomena in LHC detector datasets consists of removing $\mathcal{O}(10^{14})$ background collisions while retaining the capability to identify the searched-for signal process based on as few occurences as possible. This is done in several stages, most with tight constraints on the computing resources used. The first and least flexible step is a hardware trigger system that steers the readout of the detector: only for 1×10^6 out of the about 12×10^6 colli-

sions produced during one second of LHC operation can the LHCb detector fully be read out. This decision is based on a rough estimate of the particles with the highest transverse momenta detected by the calorimeter and muon systems.

The fully read out events are then sent to a trigger software application that runs in about 30×10^3 processes on a computing cluster, where the next two stages of triggering are performed. The first is based on full track reconstruction using the vertex detector information. Interesting tracks, with a displacement typical for heavy flavour decay products, or compatible with a muon detector segment, are further reconstructed beyond the magnetic field, such that their transverse momentum, typically larger for processes of interest than for background, can be used for further discrimination. For the signal studied in this work, a displaced multi-track vertex, an additional selection was added that keeps pairs of tracks, such that the track quality criteria that penalise tracks from decays at more than about 5 mm from the interaction region could be relaxed with respect to the nominal selection requirement. The second software trigger stage consists of a large number of different selections, all reconstructing candidates that are similar to those used for the final measurements. Exotic long-lived particle decays are reconstructed here by an algorithm that considers all tracks in the vertex detector to find their production vertices. High-multiplicity vertices outside the interaction region are candidate decays. The reconstructed invariant mass and the track momenta are used to limit the accept rate. All events that contain at least one candidate passing these selections are stored.

The first analysis step performed on the stored datasets usually consists of the complete re-reconstruction of the candidates, from the level of the read out detector signals up, using the full-precision version of the reconstruction that is executed on the worldwide LHC computing grid, without the tight trigger computing time constraints. For this search, the complete re-reconstruction of the vertex is supplemented by a parallel approach based on the trigger candidates, in order to maximise the efficiency and overlap with the triggered sample. The vertices are used as reference points to attempt the full reconstruction of the two quarks in the final state: a pair of jets of hadrons. This is done with a clustering algorithm that combines reconstructed particles with similar momentum directions, charged as well as neutral. For the charged particles, the information of the vertex detector and the tracking system is maximally used: tracks incompatible with the reconstructed decay position, the vertex, are not considered, and the momentum estimate helps to improve the precision on the jet momentum. The requirement of two jets compatible with a trigger candidate provides sufficient reduction at this stage. The resulting sample is augmented with a set of candidates built from vertices found in the larger track collection that is reconstructed offline.

The data sample is then split into categories based on the transverse displacement from the interaction region. The background yield decreases much faster as a function of this variable than typical signals, because of their different lifetime. Additional vertex, jet quality, and jet-vertex matching requirements are applied in order to further reduce the background. Two classes of instrumental backgrounds can completely be removed: when a particle from a proton-protoncollision interacts with detector material, more particles may be produced,whose tracks give rise to a pattern similar to a signal vertex. Such vertices are removed by vetoing candidates in a region around the detector parts. A second class of background was found to be related to near-parallel tracks, possibly from interactions of the beam with collimators, and collision tracks. If many silicon strips are hit in the same detector region, the reconstruction algorithm can find a large number of tracks that are very close together, such that high-multiplicity vertices can be made. This background is vetoed by exploiting some distinctive properties of the track and hit distributions observed in such events.

The remaining background vertices are due to random combinations of tracks from heavy flavour decays, from material interactions and poorly reconstructed low-momentum particles from the same or different primary collisions. Two categories can be distinguished: candidates dominated by a single heavy flavour jet or material interaction, where the particles are clustered in two nearby jets, and candidates where final-state particles from either side of a collision that produces a heavy quark pair contribute, which gives two jets that are nearly back-to-back in the transverse plane. The invariant mass distribution of the former category, with small opening angle, peaks at low mass and falls rapidly, while the second category leads to a less steeply falling spectrum. The high-mass part of the latter contribution is strongly reduced by two selection requirements applied on the candidates: the candidate should point back to a primary proton-proton collision point and the opening angle should not be too large.

The dijet invariant mass spectrum is then fitted, separately in each of the transverse displacement categories, using an empirical model that contains a contribution for the searched-for signal and for each of these background types: the dominating small-angle background contribution is described by an analytical model that tends to an exponential at high values, and is fully determined from data, in every category separately. For the second contribution, a fixed shape is added, also with an exponential high-mass tail, with the slope fixed from the high-mass distribution of the whole sample before the final selections. Statistical hypothesis tests are performed, and upper limits on the production of long-lived particles with the assumed decay properties are obtained.

A sensitivity down to the level of a few pb, or a branching fraction of 20 % of

the standard model Brout-Englert-Higgs boson is obtained, for long-lived particle lifetimes in the range 2–500 ps and masses in the range 25–50 GeV/c^2. This covers the low-mass low-lifetime region that is difficult to probe by the other LHC experiments, and that has previously been studied at the Tevatron collider with a lower centre-of-mass energy and smaller datasets. For the studied part of parameter space, the LHCb results are the most stringent to date.

Due to the minimal assumptions made to perform this search — only a single long-lived particle decay is assumed, and the selection only exploits the decay of such a particle to quarks, and no other features of the same collision — the obtained results apply to any other model that predicts the production of long-lived particles that decay to quarks, if corrected for the possibly different fraction of decays contained in the forward region visible to the detector. The developed methods can be applied to the large data sets that will be collected at a higher collision energy in the coming years, and to other final states, in order to discover or further constrain the production of exotic long-lived particles.

Samenvatting

Vormen de tot nu toe gekende elementaire deeltjes en interacties de fundamenteelste beschrijving van de natuur? En als dat zo is, hoe is het universum zoals we het waarnemen dan tot stand gekomen?

Het standaardmodel van de deeltjesfysica beschrijft de elektromagnetische interactie en de sterke en zwakke kernkracht als relativistische kwantumveldentheorieën, met interacties afgeleid van ijksymmetrieën tussen twaalf elementaire fermionen. Het is een uitbreiding van de kwantumelektrodynamica, die de interacties tussen elektronen, atoomkernen en licht en andere vormen van elektromagnetische straling die een rol spelen in processen op atomaire en moleculaire schaal met veel succes beschrijft. De sterke kernkracht, die atoomkernen samenhoudt, wordt beschreven door een theorie die gekleurde quarks vastzet in kleur-neutrale gebonden toestanden, *e.g.* protonen en neutronen. De zwakke kernkracht die verantwoordelijk is voor nucleaire vervalprocessen, wordt gemodelleerd door de uitwisseling van geladen of neutrale vectorbosonen. De grote massa van deze boodschapperdeeltjes zorgt ervoor dat de interactie zwak is in vergelijking met de sterke en elektromagnetische interactie, en maakt de vervalprocessen traag.

Over het algemeen komen de voorspellingen van de theorie zeer goed overeen met experimentele waarnemingen. Ze biedt evenwel geen verklaring voor de waargenomen structuur in het massaspectrum en de eigenschappen van de fermionen, of waarom de schaal van de elektrozwakke symmetriebreking zo laag ligt: de natuurlijke schaal voor een fundamentelere theorie is de Planckschaal, waar kwantumeffecten en de zwaartekracht verenigd moeten worden. Daarom wordt het standaardmodel verondersteld onvolledig te zijn, en bij een natuurlijke oplossing voor de zopas vermelde problemen worden nieuwe waarneembare effecten verwacht bij energieën in de buurt van de TeV-schaal.

Een andere reden waarom het standaardmodel onvolledig is, komt uit de kosmologie. De huidige modellen kunnen de evolutie van het heelal, op grote zowel als op kleine schaal, beschrijven, in de veronderstelling dat het universum, naast de baryonische materie die we zien, ook een niet-baryonische materie-achtige component bevat, donkere materie, en een kosmologische constante of donkere energiecomponent. Beide aannames zijn goed gemotiveerd — donkere materie is ook nodig om

de dynamica van melkwegstelsels te beschrijven; de versnelde uitdeining van het universum is een experimenteel gegeven — maar de juiste aard en oorsprong zijn een raadsel. Ook voor een juiste voorspelling van de hoeveelheid materie die geproduceerd is in het jonge universum, zijn nieuwe interacties nodig.

Deze vragen geven een reden om op zoek te gaan naar een volledigere beschrijving van de natuur, en naar experimentele resultaten die bij de ontwikkeling daarvan kunnen helpen. In de voobije decennia zijn allerlei modellen voor uitbreidingen van het standaardmodel voorgesteld, maar tot dusver hebben experimenten voor geen daarvan duidelijk bewijs geleverd. Integendeel, de mogelijke zichtbare effecten van nieuwe deeltjes en interacties zijn sterk beperkt.

De grote hoeveelheid botsingen die de LHC-experimenten sinds 2009 hebben geregistreerd, maken het mogelijk om het standaardmodel experimenteel te testen en nieuwe deeltjes te zoeken bij een tot nog toe onbereikbare energie en precisie. Aangezien de meeste inelastische proton-protonbotsingen bij een energie van enkele TeV veroorzaakt worden door de sterke kernkracht, en zelfs de meeste zeldzame processen goed beschreven worden door het standaardmodel, worden de botsingen met nieuwe deeltjes of andere afwijkingen aan het oog onttrokken door een grote hoeveelheid achtergrond. De interessante processen worden daarom van de achtergrond gescheiden door gebruik te maken van karakteristieke eigenschappen, bijvoorbeeld een eindtoestand waarvan de zichtbare onderdelen aangeven dat een groot deel niet gedetecteerd is, volledig gereconstrueerde vervallen van gekende zwaardere deeltjes *etc.*

Dit werk richt zich op de waarneming van nieuwe deeltjes door middel van hun verval op een andere plaats dan het proton-proton botsingspunt waar ze geproduceerd zijn. Nauwkeurige sporenreconstructie laat toe om vervallen te onderscheiden tot op enkele mm van het productiepunt, dankzij de precisie van silicium vertexdetectoren. Deze detectoren worden door de meeste collider-experimenten gebruikt om vervallen van beauty- en charmhadronen te herkennen en te bestuderen. Zulke hadronen hebben een gemiddelde levensduur van de grootte-orde 1 ps, wat aan de LHC overeenkomt met een vertex die ongeveer een centimeter verplaatst is.

De LHCb-detector is specifiek ontworpen om zware quarkvervallen te bestuderen, en kan dus ook gebruikt worden om te zoeken naar de vervallen van langlevende exotische deeltjes. De vertexdetector kan sporen reconstrueren van vervallen tot 20 cm van het interactiepunt langs de bundellijn en tot ongeveer 4 cm in het vlak loodrecht daarop. Bovendien is het voorwaartse gebied, met hoeken tot 250–300 mrad met de bundellijn, volledig uitgerust met detectoren: sporendetectoren en identificatiedetectoren voor geladen deeltjes, en calorimeters die gevoelig zijn aan alle elektromagnetisch en sterk interagerende deeltjes.

De experimentele uitdaging om nieuwe verschijnselen te ontdekken in de dataset

van een LHC-detector bestaat eruit $\mathcal{O}(10^{14})$ achtergrondbotsingen te verwijderen, en toch nog het bestaan van het signaalproces te kunnen aantonen op basis van zo weinig mogelijk botsingen waarin dit plaatsvindt. Deze reductie gebeurt in verschillende stappen, met veelal strikte beperkingen op de gebruikte rekenkracht en harde schrijfruimte. De eerste en minst flexibele stap is een hardwaretrigger die het uitlezen van de detector controleert: de LHCb detector kan ten hoogste voor 1×10^6 van de ongeveer 12×10^6 botsingen die de LHC per seconde produceert, volledig uitgelezen worden. De beslissing wordt genomen op basis van een ruwe schatting van de deeltjes met de hoogste loodrechte impulsen in de calorimeters en de muondetector.

De volledig uitgelezen *events* worden naar een rekencentrum gestuurd, waar ongeveer 30×10^3 processen van een softwaretrigger de twee volgende reductiestappen uitvoeren. De eerste stap begint met de volledige sporenreconstructie in de vertexdetector. Interessante sporen, die niet recht naar een botsingspunt terugwijzen, zoals de vervalproducten van zware quarks, of die overeenkomen met een spoor in de muondetector, worden verder gereconstrueerd aan de andere kant van het magnetisch veld, zodat hun loodrechte impuls, die typisch groter is voor de bestudeerde processen dan voor achtergrond, voor de verdere selectie gebruikt kan worden. Voor het in dit werk bestudeerde signaal, een verplaatste vertex met een groot aantal sporen, werd een selectie toegevoegd voor paren van sporen, zodat de spoorkwaliteitsvereisten, die sporen van vervallen op meer dan 5 mm benadelen, losser konden worden gemaakt dan in de standaard één-spoor selectie. De tweede softwaretriggerstap bestaat uit een groot aantal verschillende selecties die allemaal kandidaten reconstrueren die zeer gelijkaardig zijn aan die gebruikt voor de uiteindelijke metingen. De vervallen van exotisch langlevende deeltjes worden hier gereconstrueerd door een algoritme dat alle sporen in de vertexdetector gebruikt om hun productievertices te zoeken. Vertices buiten het interactiegebied en met een groot aantal sporen zijn signaalkandidaten. De gereconstrueerde invariante massa en de impuls van de sporen worden gebruikt om het aantal aanvaarde events te beperken. Alle events met ten minste één kandidaat die aan de vereisten voldoet, worden opgeslaan.

De eerste stap van een analyse die uitgevoerd wordt op een opgeslagen dataset is meestal een her-reconstructie van de kandidaten, vanaf de detectorsignalen, met de meest nauwkeurige versie van de reconstructie, op het wereldwijde LHC computing grid, zonder de strikte rekentijdbeperking van de trigger. In dit specifieke geval werden, naast de volledig hergereconstrueerde vertices, ook de triggerkandidaten hergebruikt, om de efficiëntie en overlap met de triggerselectie te maximaliseren. De vertices worden gebruikt als referentiepunten voor de volledige reconstructie van de uitgaande quarks van het verval als twee hadronjets. Daarvoor wordt een

clusteralgoritme gebruikt dat gereconstrueerde deeltjes, geladen en neutraal, met gelijkaardige impulsrichtingen groepeert. Voor de geladen deeltjes wordt de informatie van de vertexdetector en de sporendetectoren zo veel mogelijk gebruikt: sporen die niet compatibel zijn met het veronderstelde vervalpunt, de vertex, komen niet in aanmerking, en de impulsbepaling helpt de precisie op de totale jetimpuls te verbeteren. De vereiste van twee jets die passen bij de triggerkandidaat zorgt voor een voldoende reductie. Deze worden aangevuld met de kandidaten gevonden bij de vertices die gereconstrueerd zijn met behulp van de uitgebreidere offline sporenverzameling.

De dataset wordt vervolgens onderverdeeld in categorieën aan de hand van de loodrechte verplaatsing van de interactieregio. Het aantal gevonden achtergrondkandidaten neemt sneller af als functie van die variabele dan voor een typisch signaal, door de verschillende levensduur. De achtergrond wordt verder onderdrukt met selecties voor de vertex, de kwaliteit van de jets en de compatibiliteit van de jets met de vertex. Twee types instrumentele achtergrond kunnen volledig geëlimineerd worden: een deeltje afkomstig van de proton-protonbotsing kan botsen met het detectormateriaal, waarbij andere deeltjes kunnen geproduceerd worden, met sporen die op een signaalvertex lijken. Dergelijke vertices worden verwijderd door kandidaten uit het gebied rondom de detectoronderdelen te verwijderen. Een tweede groep achtergrondvertices werd geïdentificeerd als een combinatie van bijna evenwijdige sporen, mogelijk afkomstig van interacties van de bundel met collimators, en sporen van de proton-protonbotsingen. Als veel siliciumstrips in hetzelfde stuk van de detector geraakt worden, kan het reconstructie-algoritme een groot aantal tracks dicht bij elkaar vinden, waaruit vertices gemaakt kunnen worden. Deze achtergrond wordt verwijderd op basis van de karakteristieke spoor- en hitdistributies in dergelijke events.

De overblijvende achtergrondvertices komen voort uit willekeurige combinaties van sporen van zware quarkvervallen, materiaalinteracties en slecht gereconstrueerde deeltjes met een laag momentum, afkomstig van dezelfde of verschillende primaire botsingen. Twee categorieën kunnen onderscheiden worden: kandidaten die gedomineerd worden door een enkele zware quarkjet of materiaalinteractie, waar de deeltjes geclusterd worden in twee jets dicht bij elkaar, en kandidaten waar de eindproducten van beide kanten van een botsing met een paar zware quarks bijdragen, wat jets oplevert met ongeveer tegenovergestelde richtingen in het loodvlak op de bundellijn. De invariante massadistributie van de eerste categorie, met kleine openingshoek, heeft een piek bij lage massa en neemt snel af, terwijl de tweede categorie een minder snel afnemend spectrum oplevert. De bijdrage bij hoge massa van de laatste component wordt sterk verminderd door de twee laatste selectievereisten: de kandidaat moet terugwijzen naar een proton-protonbotsingspunt en de

openingshoek mag niet te groot zijn.

De invariante massadistributie van de dijetkandidaten wordt dan gefit, in iedere categorie van verplaatsing afzonderlijk, met een empirisch model dat een bijdrage voor het veronderstelde signaal en voor elk van de achtergrondtypes bevat: de grootste achtergrondbijdrage met kleine openingshoek wordt beschreven door een analytisch model dat naar een exponentiële distributie neigt bij hoge invariante massa, en wordt helemaal bepaald door de data, in iedere categorie afzonderlijk. Voor de tweede bijdrage wordt een vaste vorm toegevoegd, ook met een exponentiële staart bij hoge massa, met een helling die vastgezet wordt op de waarde die de distributie voor de hele dataset samen het beste beschrijft bij hoge invariante massa, voordat de laatste twee selecties toegepast worden. Uit statistische hypothesetests worden bovenlimieten op de productie van langlevende deeltjes met de veronderstelde eigenschappen bekomen.

Een gevoeligheid tot het niveau van enkele pb, of een vertakkingsverhouding van 20 % van het Brout-Englert-Higgsboson in het standaardmodel, werd bereikt, voor langlevende deeltjes met gemiddelde levensduur in het interval 2–500 ps en massa 25–50 GeV/c^2. Dit beslaat het gebied met lage massa en levensduur dat moeilijk bereikbaar is voor de andere LHC-experimenten, en dat voorheen bestudeerd is aan de Tevatronversneller, met een lagere botsingsenergie en kleinere datasts. Voor het bestudeerde gedeelte van de parameterruimte zijn de LHCb-resultaten de meest gevoelige tot nu toe.

Door de minimale aannames in deze zoektocht — het verval van slechts één langlevend deeltje wordt verondersteld, en de selectie berust enkel op het verval van zo'n deeltje naar quarks, en geen andere eigenschappen van de botsing die het voortbrengt — zijn de resultaten ook toepasbaar op andere modellen die de productie van langlevende deeltjes met quarkvervallen voorspellen, mits een correctie wordt toegepast voor de mogelijk verschillende fractie van de vervallen dat zich in het geïnstrumenteerde gedeelte van de detector bevindt. De ontwikkelde methodes kunnen ook toegepast worden bij de analyse van de grotere datasets die in de komende jaren bij een hogere botsingsenergie verzameld zullen worden, en bij verschillende eindtoestanden, om de productie van exotische langlevende deeltjes te ontdekken of verder te begrenzen.

Acknowledgements

I would like to use the last pages of my thesis to specifically thank some of the many people without whom this thesis would not exist, and the work that lead to it would not have been as fun and interesting as it has been.

First of all Wouter, my supervisor, who was always available to answer questions and to give advice, on analysis technicalities, physics in general, the first versions of my thesis chapters *etc.* He also kept an eye on the bigger picture and schedule, but still gave me the freedom to get sidetracked a bit from time to time — both were important. I also want to thank Marcel, my advisor, for his support throughout my PhD, for the interesting discussions we had, for the fast but thorough reading of the thesis manuscript, and for the helpful comments.

Carrying out the analysis described in this thesis in the time that was left for it, would also not have been possible without all the insights and techniques developed as part of the other long-lived particle searches, so I would like to thank the exotica and jet reconstruction working groups and the many other helpful LHCb colleagues. Special thanks to Victor, who helped me get started with all simulation, trigger, stripping and jet reconstruction code, and also with some crazy analysis ideas that made the work atmosphere at CERN even nicer. Veerle not only developed many parts of the predecessor analysis that also ended up in this work, but has also been a nice office mate, and was always available for questions and moral support, even after finishing her PhD.

I also very much enjoyed and learned a lot from the many discussions during coffee, tea and lunch breaks, many other activities with the b-physics group at Nikhef, and the weekly PhD student discussions that we organised for some time. Thanks Suvayu, Jeroen, Rose, Siim, Panos, Vasilis, Jacco, Roel, Kristof, Serena, Alex, Barbara, Chiara, Elena, Laurent, Maarten, Mick, Katya, Maurizio, Francesco, Manuel, Rob, Greg, Jeroen, Hella, Niels, Patrick, Eddy, Tjeerd, Gerhard and Antonio for all the shared insights and opinions, and the always stimulating group atmosphere.

One only learns how well one understands something by trying to explain to someone else — I noticed once again how true that is thanks to Suzanne and Elise, who did their master's projects with Wouter, Veerle and me, with great enthusiasm, lots of interesting questions, and a firm belief in the existence of exotic long-lived

particles.

That brings me to the long list of friends and colleagues that I met at CERN and at Nikhef, at summer schools and conferences, with whom I had lots of nice lunches, coffee breaks, Friday afternoon beers and dinners, with discussions about physics, but much more often other topics — thanks for all the nice moments. I find it very difficult to make a list here, so I will apologise for that, and just try to give as many of you as possible a copy of my thesis, and an invitation for the defence.

Ik wil ook mijn Belgische vrienden bedanken, de ondertussen iets meer dan zeven van de middelbare school, met wie het jaarlijkse weekend en enkele andere evenementen intussen echte tradities geworden zijn, en de vrienden van de fysica uit Gent, in het bijzonder Karen en Nadja. Het was altijd fijn om jullie nog een keer terug te zien tijdens een vakantie of verlengd weekend.

Ten slotte wil ik ook mijn ouders, mijn zussen Sara en Lies, en de rest van de familie bedanken voor hun steun, de bezoekjes aan Genève en Amsterdam, de vele weekendpauzes thuis, en hun voortdurende interesse en medeleven, zelfs met het schrijfproces.

Thanks,

Pieter