
CMS Physics Analysis Summary

Contact: cms-pag-conveners-exotica@cern.ch

2014/03/18

Search for long-lived particles decaying to final states that include dileptons

The CMS Collaboration

Abstract

A search has been performed for long-lived particles, which decay to a final state that includes a pair of leptons (e^+e^- or $\mu^+\mu^-$). The experimental signature is a distinctive topology consisting of a pair of charged leptons originating from a displaced secondary vertex. Events were collected by the CMS detector at the LHC during pp collisions at $\sqrt{s} = 8$ TeV, and selected from data samples corresponding to 19.6 (20.5) fb^{-1} of integrated luminosity in the electron (muon) channel. No significant excess is observed above standard model expectations. Upper limits are set on such a signal, as a function of the long-lived particle lifetime, in the context of two specific models. In the first, Higgs bosons decay to a pair of long-lived neutral particles that each subsequently decay to dileptons. In the second, events contain a pair of squarks that each decay to a long-lived neutralinos, which then each decay to dileptons and a neutrino.

1 Introduction

Several models of new physics predict the existence of massive, long-lived particles which could manifest themselves through their delayed decays to leptons. Such scenarios arise, for example, in various supersymmetric (SUSY) scenarios such as “split SUSY” [1] or SUSY with very weak R-parity violation [2], “hidden valley” models [3], and Z' models that contain long-lived neutrinos [4].

This Letter presents a search, using data from the Compact Muon Solenoid (CMS) collected during 2012 in proton-proton (pp) collisions at $\sqrt{s} = 8$ TeV, for massive, long-lived exotic particles that decay to a pair of charged leptons. We search for events containing a pair of charged electrons or muons (dileptons) originating from a common secondary vertex within the volume of the CMS tracker, where the leptons have a significant transverse displacement from the event primary vertex. This topological signature has the potential to provide clear evidence for physics beyond the standard model (SM). It is also very powerful in suppressing backgrounds from SM processes.

This signature can probe a wide class of models. However, we will present search results in the context of two specific models, so as to give a quantitative indication of the typical sensitivity. In the first model, the long-lived particle is a spinless boson X , which has a non-zero branching fraction to dileptons. The X is pair-produced in the decay of a (non-SM) Higgs boson, i.e. $H^0 \rightarrow XX, X \rightarrow \ell^+\ell^-$ [5], where the Higgs boson is produced through gluon-gluon fusion. In the second model, the long-lived particle is a neutralino $\tilde{\chi}^0$ which has an R-parity violating decay into a neutrino and two leptons. The neutralino is produced in events containing a pair of squarks, where a squark can decay via the process $\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu$. Both models predict up to two displaced dilepton vertices in the tracking volume per event. In this paper, we will use *LL particle* to refer to any long-lived particle, such as the X or $\tilde{\chi}^0$ particle considered in our signal models.

The search presented here is an update of a previous CMS analysis that used a smaller data sample collected at $\sqrt{s} = 7$ TeV during 2011 [6]. The higher integrated luminosity collected in 2012 makes the new analysis sensitive to signal cross sections about a factor of five smaller than previously. In addition, an improved analysis strategy used in 2012 substantially broadens the range of signal models to which the analysis is sensitive. The analysis presented here complements a similar CMS analysis that searches for LL particles decaying to dijets [7].

The D0 Collaboration has performed searches for leptons from delayed decays in its tracker volume [8, 9], but these searches are sensitive to a much smaller kinematic phase space region than CMS. The ATLAS Collaboration has performed searches that are sensitive to decay lengths up to about 20 m by exploiting the ATLAS muon spectrometer [10, 11], using different decay channels from those considered in this Letter.

2 The CMS detector

The central feature of the CMS apparatus [12] is a superconducting solenoid of 6 m internal diameter providing an axial field of 3.8 T. Within the field volume are the silicon pixel and strip tracker, the lead-tungstate crystal electromagnetic calorimeter (ECAL), and the brass/scintillator hadron calorimeter (HCAL). Muons are identified in gas-ionisation detectors embedded in the steel magnetic-flux return yoke of the solenoid.

The silicon tracker is composed of pixel detectors (three barrel layers and two forward disks on either end of the detector) surrounded by strip detectors (ten barrel layers plus three inner

disks and nine forward disks at each end of the detector). The tracker covers the pseudorapidity range $|\eta| < 2.5$, where $\eta = -\ln[\tan(\theta/2)]$ and θ is the polar angle with respect to the anticlockwise-beam direction. The pixel tracker and a subset of the strip tracker layers provide three-dimensional hit position measurements. The other strip tracker layers measure hit position only in $r\phi$, where ϕ is the azimuthal angle about the anticlockwise-beam direction in the barrel, or $z\phi$ in the endcap. Owing to the strong magnetic field and the high granularity of the silicon tracker, promptly produced charged particles with transverse momentum $p_T = 100 \text{ GeV}/c$ are reconstructed with a resolution in p_T of $\approx 1.5\%$ and in transverse impact parameter d_0 of $\approx 15 \mu\text{m}$. The track reconstruction algorithms are able to reconstruct displaced tracks with transverse impact parameters up to $\approx 25 \text{ cm}$ from particles decaying up to $\approx 50 \text{ cm}$ from the beam line. The performance of the track reconstruction algorithms has been studied with data [13]. The silicon tracker is also used to reconstruct the primary vertex position with a precision of $\approx 20 \mu\text{m}$ in each dimension.

The ECAL consists of nearly 76 000 lead tungstate crystals, which provide coverage in pseudorapidity $|\eta| < 3$. Muons are measured in the pseudorapidity range $|\eta| < 2.4$ with detection planes based on one of three technologies: drift tubes in the barrel region, cathode strip chambers in the endcaps, and resistive plate chambers in the barrel and endcaps.

The first level of the CMS trigger system, composed of custom hardware processors, selects events of interest using information from the calorimeters and the muon detectors. A high-level trigger processor farm then employs the full event information to further decrease the event rate.

3 Data and Monte Carlo simulation samples

For this analysis, pp collision data at a centre-of-mass energy of 8 TeV corresponding to an integrated luminosity of 19.6 ± 0.5 (20.5 ± 0.5) fb^{-1} are used in electron (muon) channels. The lower luminosity in the electron channel is due to different data quality requirements for various sub-detectors to those in the muon channel.

For the electron channel, these data are collected with a high-level trigger that requires two clustered energy deposits in the ECAL. The leading (sub-leading) energy deposit is required to have transverse energy $E_T > 36$ (22) GeV, and both clusters are required to pass loose requirements on their compatibility with a photon/electron hypothesis. For the muon channel, the trigger requires two muons, each reconstructed in the muon detectors without using any primary vertex constraint and having $p_T > 23 \text{ GeV}/c$. To prevent cosmic ray muons from passing these criteria, the opening angle between the two muons must be less than 2.5 radians. Tracker information is not used in either trigger.

For the $H^0 \rightarrow XX$ model, simulated signal samples are generated using PYTHIA V6.426 [14] to simulate H^0 production through gluon fusion ($gg \rightarrow H^0$). Subsequently the H^0 is forced to decay into XX , with the X bosons each decaying to dileptons ($X \rightarrow \ell^+\ell^-$), where ℓ represents either a muon or an electron. Several samples with different combinations of H^0 masses ($M_{H^0} = 125, 200, 400, 1000 \text{ GeV}/c^2$) and X boson masses ($M_X = 20, 50, 150, 350 \text{ GeV}/c^2$) are generated. The Higgs boson resonance is assumed to be narrow for the purposes of simulation, but the analysis has negligible dependence on this. Each sample is furthermore produced with three different lifetimes of the X bosons, corresponding to mean transverse decay lengths of approximately 2 cm, 20 cm and 200 cm in the laboratory frame. For the $\tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu$ model, PYTHIA is used to simulate squark pair production and decay to $\tilde{\chi}^0$, with four combinations of squark mass and neutralino mass ($M_{\tilde{q}}, M_{\tilde{\chi}^0} = (1500, 494), (1000, 148), (350, 148)$, and $(120, 48) \text{ GeV}/c^2$)

used. The R-parity violating parameter λ'_{211} is set to a non-zero value to enable the decay of the $\tilde{\chi}^0$ into two charged leptons and a neutrino. The value of λ'_{211} is chosen to give a mean transverse decay length of approximately 20 cm.

Several simulated background samples generated with PYTHIA are used. The dominant background is Drell-Yan production of dileptons, with contributions from e^+e^- , $\mu^+\mu^-$ and $\tau^+\tau^-$ all being significant. Although only a small fraction of the taus decay to electrons or muons, when they do so, the electrons and muons are displaced from the primary vertex. Other simulated backgrounds are $t\bar{t}$, W/Z boson pair production with leptonic decays, and QCD multijet events. In all the samples, the response of the detector is simulated in detail using GEANT4 [15]. The samples are then processed through the trigger emulation and event reconstruction chain of the CMS experiment.

4 Event reconstruction and selection

Events are required to contain a primary vertex with at least four associated tracks whose position is displaced from the nominal interaction point by no more than 2 cm in the direction transverse to the beam and no more than 24 cm in the direction along the beam. Furthermore, to reject events produced by the interaction of beam-related protons with the LHC collimators, the fraction of tracks classified as ‘high purity’, as defined in Ref. [13], must exceed 25% in any event with at least 10 tracks.

Standard CMS offline lepton identification algorithms are not applied, since they are inefficient for leptons from highly displaced vertices. However, these algorithms are not needed to suppress the very low backgrounds present in this analysis. Leptons are reconstructed using tracks reconstructed in the tracker, which must be classified as high-purity, and have pseudorapidity $|\eta| < 2$, since the efficiency for finding tracks from displaced secondary vertices falls off at large $|\eta|$.

A track is considered to be identified as a muon if it matches a muon trigger object within $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 0.1$. Here, $\Delta\eta$ and $\Delta\phi$ are the differences between the track and the muon trigger object in η and ϕ .

A track is identified as an electron if its direction is consistent within a cone of size $\Delta R < 0.1$ with an energy deposit in the ECAL identified as a photon. The energy of the electron is taken from this ECAL deposit, since it is less affected by bremsstrahlung than the track. Some quality requirements are placed on the ECAL deposit: it must not lie in the narrow region between the ECAL barrel and endcap; the ratio of hadronic energy in the HCAL tower behind the seed crystal of the ECAL energy deposit to the energy of the ECAL energy deposit (H/E) must be less than 0.05; and a measure of the spread in η of the ECAL deposit, known as $\sigma_{i\eta i\eta}$, must satisfy $\sigma_{i\eta i\eta} < 0.012$ (0.034) if the ECAL deposit is in the barrel (endcap) region of the ECAL.

The LL particle candidates are formed from pairs of charged lepton candidates. In the dimuon channel, the two tracks must each have $p_T > 26\text{ GeV}/c$ and be oppositely charged. In the dielectron channel, the higher (lower) E_T electron must satisfy $E_T > 40\text{ GeV}$ (25 GeV), where the energy is taken from the ECAL measurement. These thresholds are slightly higher than the corresponding trigger requirements, to ensure that the trigger has good efficiency. In the dielectron channel, the two tracks must also satisfy $p_T > 36\text{ GeV}/c$ (21 GeV/c) if associated to the higher (lower) E_T electron. This p_T requirement, which is slightly lower than the corresponding E_T requirement placed on the ECAL deposit, suppresses electrons that emit large amounts of bremsstrahlung, and which thus tend to have poor impact parameter resolution. No charge

requirement is applied in the electron case, as bremsstrahlung can result in incorrect charge reconstruction of high p_T electrons.

To reject promptly produced particles, the tracks must have a transverse impact parameter significance with respect to the primary vertex of $|d_0|/\sigma_d > 12$. (The value of 12 was chosen so as to give an expected background significantly below one event, which gives the best signal sensitivity for the vast majority of the LL particle lifetimes considered in this paper). Both lepton candidates are required to be isolated, to reject background from jets. Specifically, a hollow isolation cone is constructed around each candidate, with a radius $0.04 < \Delta R < 0.3$ for electrons and $0.03 < \Delta R < 0.3$ for muons. Within this isolation cone, the $\sum p_T$ of all tracks with $p_T > 1 \text{ GeV}/c$, excluding the other lepton candidate, divided by the p_T of the lepton must be less than 0.1. This requirement has very little effect on the signal efficiency, which is relatively insensitive to the number of primary vertices in each event.

The two tracks are fitted to a common vertex, which is required to have a chi-squared per degree of freedom $\chi^2/\text{dof} < 10$ (5) in the electron (muon) channel. The candidate is rejected if the two tracks together have more than one hit assigned to them, which lies between the centre of CMS and their common vertex position. The candidate is also rejected if, at distances from the centre of CMS lying beyond the vertex position but smaller than the position of the innermost hit on each track, the tracks together have a number of missing hits exceeding 3 (4) in the electron (muon) channel. A missing hit is defined as occurring when a track passes through an active sensor without being assigned a hit. To eliminate background from J/ψ and Υ decays and from γ conversions, LL particle candidates are required to have a dilepton invariant mass larger than $15 \text{ GeV}/c^2$. Finally, the difference in azimuthal angle, $\Delta\Phi$, between the dilepton momentum vector and the vector from the primary vertex to the dilepton vertex, is required to satisfy $|\Delta\Phi| < \pi/2$, where $\Delta\Phi$ is measured in the range $-\pi < \Delta\Phi < \pi$. (Events satisfying the other cuts but with $|\Delta\Phi| > \pi/2$ are used to define a control region, as detailed in Section 5.)

Cosmic ray muons may be reconstructed as back-to-back tracks. To reject them, a requirement of $\cos(\alpha) > -0.79$ is applied, where α is the opening angle between the two tracks. Again, this requirement is slightly tighter than the requirement by the trigger. Background from misidentified leptons is reduced by requiring that the two lepton candidates are not both matched to the same trigger object or offline photon. Owing to the difficulty of modelling the very poor trigger efficiency for closely spaced muon pairs, the two tracks in muon channel candidates must be separated by $\Delta R > 0.2$.

The selection efficiency and the limits are determined in terms of the number of events passing our selection, rather than the number of candidates. An event is determined to pass if it contains at least one candidate passing the selection.

4.1 Selection efficiency and acceptance

The overall signal efficiency is defined as the fraction of events in which at least one dilepton candidate passes all selection criteria. It is determined from the simulation, separately for the electron and muon channels, and independently for two different classes of events: first for events in which only one LL exotic particle (X or $\tilde{\chi}^0$) decays to the chosen lepton species, defining efficiency ϵ_1 , and second for events in which both LL exotic particles decay to chosen lepton species, defining efficiency ϵ_2 . The efficiencies are estimated for a range of LL particle lifetimes, corresponding to mean transverse decay lengths of $\approx 0.7 - 600 \text{ cm}$, by reweighting the simulated signal events. The maximum efficiency (for $H^0 \rightarrow XX$ with $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 150 \text{ GeV}/c^2$, $c\tau = 1 \text{ cm}$) is approximately 36% (46%) in the electron (muon) channel, but becomes significantly smaller at lower H^0 masses or longer lifetimes. For example, if $c\tau$ is in-

creased to 20 cm for this set of masses, the efficiency drops to 14% (20%) in the electron (muon) channel.

For simulated events, it is useful to define a set of *acceptance criteria* defining events that could possibly be reconstructed at CMS. Specifically, the generated transverse decay length of the LL particle should be no more than 50 cm, and the generated leptons should pass the p_T and η requirements listed in Section 4; that is, both must have $|\eta| < 2$, and in the muon channel, both must have $p_T > 26 \text{ GeV}/c$, whilst in the electron channel, one must have $E_T > 40 \text{ GeV}$ and the other $E_T > 25 \text{ GeV}$. The *acceptance* is defined as the fraction of LL particle decays that pass the acceptance criteria. If the efficiency is measured using only LL particle decays within the acceptance, it is higher; e.g., for $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 150 \text{ GeV}/c^2$, $c\tau = 1 \text{ cm}$, it is approximately 44% (58%) in the electron (muon) channel. More importantly, it shows much less dependence on the exact choice of signal model; e.g., for this same choice of masses, but with $c\tau = 20 \text{ cm}$, the efficiency defined in this way falls to 28% (40%) in the electron (muon) channel.

5 Estimated background and associated systematic uncertainties

The difference in azimuthal angle, $\Delta\Phi$, between the dilepton momentum vector and the vector from the primary vertex to the dilepton vertex, should be small for signal events, as the dilepton system is likely to be thrown forward in the direction in which the LL particle was moving. Indeed, this is certain for LL particles that decay *only* to dileptons. For background events, the collinearity angle will be more randomly distributed, as in this case, the displacement of the dilepton vertex relative to the primary vertex does not correspond to the flight direction of any LL particle. (The requirement that the dilepton invariant mass should exceed $15 \text{ GeV}/c^2$, mentioned in Section 4, rules out the possibility that both leptons come from the decay of the same long-lived SM particle, such as a B hadron, or from a γ conversion. And cosmic rays, which would each be reconstructed as two muons tracks, and so could also be a potential background, are eliminated by the requirement that $\cos(\alpha) > -0.79$.) In particular, defining a ‘signal region’ and ‘control region’ corresponding to $|\Delta\Phi| < \pi/2$ and $|\Delta\Phi| > \pi/2$, respectively, signal should inhabit almost exclusively the signal region, whilst background should populate the two regions equally. For example, in the case of Drell-Yan production of dimuons, any displacement of the reconstructed dimuon vertex from the primary vertex is due to detector resolution. As each muon constrains the position of the dimuon vertex in a way that is symmetric under reflection, the covariance matrix of the vertex is also symmetric under reflection. To take another example, in the case of Drell-Yan production of ditaus, where each tau decays to a muon, the muons may have significant $|d_0|/\sigma_d$ due to the non-zero lifetime of the τ lepton. However, the dimuon vertex is not a genuine vertex in this case, so its position is not correlated with the dimuon momentum vector. The expected background level in the signal region with $|d_0|/\sigma_d > 12$ can therefore be derived from the number of candidates seen in the control region with the same requirement. We observe zero events in data with $|d_0|/\sigma_d > 12$ in the control region, and this determines the probability distribution of the expected background level, as explained in Section 7. The systematic uncertainty on this estimate is derived below.

We validate this method by comparing the $|d_0|/\sigma_d$ distribution in the signal region with that in the control region, using data at modest $|d_0|/\sigma_d$ values, where the data is background-dominated. The same comparison can be made using simulated background events. In both cases, we observe good agreement in both shape and normalization between the signal and control regions, supporting the validity of the method. Figure 1 shows the tail-cumulative distributions (i.e., the integral from the plotted value to infinity) of $|d_0|/\sigma_d$ in the signal and

control regions, excluding the region with $|d_0|/\sigma_d > 6$ (4.5) in the electron (muon) channel in case signal is present. No statistically significant difference is seen. Figure 2 shows the comparison between the data and the simulated background for the signal region; we can see that the expected background is predominantly Drell-Yan dilepton production.

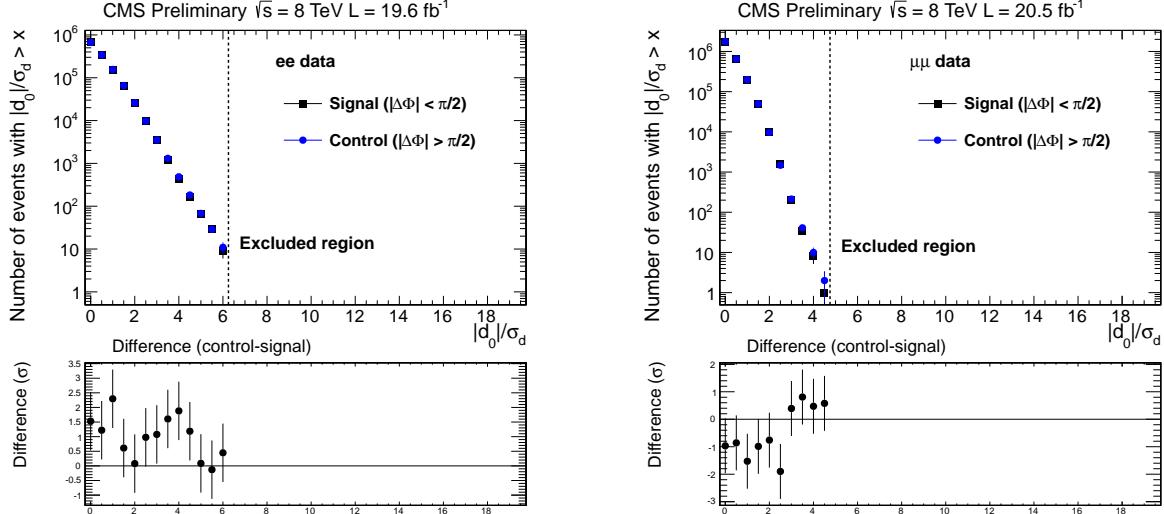


Figure 1: Comparison of the tail-cumulative distributions of $|d_0|/\sigma_d$ for events in the signal region ($|\Delta\Phi| < \pi/2$) and the control region ($|\Delta\Phi| > \pi/2$) for the dielectron channel (left) and the dimuon channel (right). The points for the signal and control regions lie almost on top of each other. The bottom plots show the statistical significance of the difference between the two.

Residual misalignment of the tracker is the only known effect that can cause the expected background to differ in the signal and control regions. This effect can be corrected for using *impact-parameter corrections* that are derived as follows: The mean offset from zero of the signed d_0 and z_0 of prompt (i.e., $|d_0|$ and $|z_0|$ below 500 μm) muon tracks is measured, as a function of the track η and ϕ , and also as a function of run period. This bias, which arises from misalignment and is always less than 5 μm , can then be subtracted from the measured impact parameters of individual tracks. The validity of this method has been checked by applying it to promptly reconstructed data, which is known to have a less precise alignment to the reprocessed data used for the rest of the analysis. If plots such as Fig. 1 are made for the promptly reconstructed data, we observe a significant asymmetry between the control and signal regions arising from the misalignment. This asymmetry is largely removed by applying the impact parameter corrections derived from data.

To define the difference between two distributions, and thus the systematic uncertainty arising from this difference, we employ the following procedure: first, the tail-cumulative distribution of the $|d_0|/\sigma_d$ distribution is calculated (excluding the high- $|d_0|/\sigma_d$ region for the signal region). Then, the ratio of the two tail cumulative distributions is computed. Points where the relative statistical uncertainty on this ratio is greater than 30% are excluded, to avoid problems arising from regions where the number of events is very small. Finally, the maximum difference of the ratio from unity for all values of $|d_0|/\sigma_d$ is taken to be the systematic uncertainty.

We consider two sources of the systematic uncertainty on the background. The first accounts for the effect of the impact-parameter corrections applied, and is obtained by comparing the distributions in the signal and control regions (taking the larger of the differences in the two regions) before and after the impact-parameter corrections are applied. This results in a systematic uncertainty of 7% and 12% in the electron and muon channels, respectively. The second

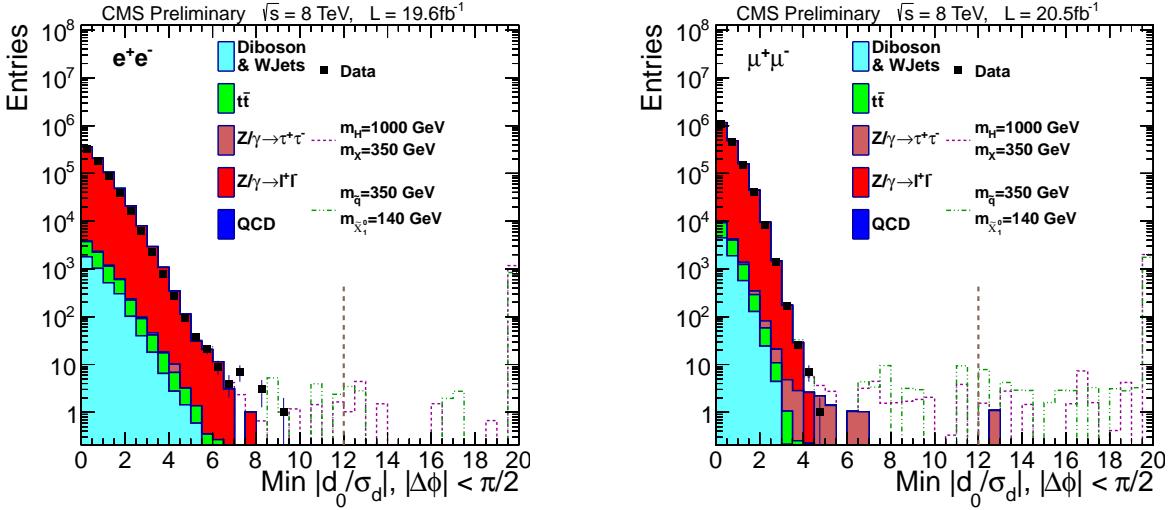


Figure 2: The minimum transverse impact parameter significance $|d_0|/\sigma_d$ for events in the signal region ($|\Delta\Phi| < \pi/2$) for the dielectron (left) and dielectron (right) channels. The solid points indicate the data, the coloured histograms are the simulated background, and the dashed lines show the simulated signal. The vertical dashed line shows the selection requirement of $|d_0|/\sigma_d = 12$ used.

accounts for any residual misalignment present after the application of the impact-parameter corrections, and is obtained by comparing the two regions with $\Delta\Phi < 0$ and $\Delta\Phi > 0$ after the impact-parameter corrections are applied. Both signal and background are expected to be equally divided between these two regions, so any significant asymmetry between them can only arise through systematic effects. In this case, we observe a maximum variation in the ratio of 11% and 21% in the electron and muon channels, respectively. To be conservative, we take the larger value of the two variations (before and after impact-parameter corrections, and with $\Delta\Phi < 0$ and > 0), yielding a final systematic of 11% and 21% in the electron and muon channels, respectively, on the estimated amount of background. The statistical significance of the asymmetry due to misalignment decreases with increasing $|d_0|/\sigma_d$ requirement, and is insignificant for $|d_0|/\sigma_d > 5$, suggesting that with the more stringent $|d_0|/\sigma_d$ requirement used in the analysis, misalignment will actually have a negligible effect on the results.

One potential issue with the above method is that since the systematic uncertainty due to misalignment is determined at values of $|d_0|/\sigma_d$ lower than our selection requirement, the possibility exists that the systematic due to misalignment issues could grow for larger $|d_0|/\sigma_d$ values. In order to account for this possibility, when computing our final limits, we do so twice, once with the impact-parameter corrections applied, and once without them, and then take the worse limits as our final result.

6 Systematic uncertainties affecting the signal

Aside from the systematic uncertainty on the estimated background, which was discussed in Section 5, one must also consider the systematic uncertainty on the efficiency to select signal events and the uncertainty on the integrated luminosity. These are addressed here. The systematic effects influencing the signal efficiency relate to: uncertainties in the efficiency of reconstructing tracks from displaced vertices; the trigger efficiency; the modelling of pileup in the simulation; the parton distribution function sets; the renormalisation and factorisation

scales used in generating simulated events; and the effect of higher order QCD corrections.

Table 1 summarises the non-negligible sources of systematic uncertainty affecting the signal efficiency. The most important of these is that related to the tracking efficiency. The relative uncertainty in the luminosity is 2.6% [16].

Table 1: Systematic uncertainties affecting the signal efficiency over the range of M_{H^0} and M_X values considered. In all cases, the uncertainty specified is a relative uncertainty. Note that the NLO uncertainty (*) shown is for the $M_{H^0} = 125 \text{ GeV}/c^2$ case, and is negligible for the other cases. The relative uncertainty in the luminosity is 2.6%.

Source	Uncertainty
Pileup modelling	2%
Parton distribution functions	< 1%
Renormalisation and factorisation scales	< 0.5%
Tracking efficiency from cosmics	6.1%
Tracking efficiency in high hit occupancy environment	3.4%
Tracking efficiency loss from bremsstrahlung (e only)	5.8%
Trigger efficiency	1.7% (e), 6.2% (μ)
NLO effects (*)	5 – 7%

Varying the modelling of the pileup within its estimated uncertainties yields a relative change in the signal selection efficiency of less than 2%, irrespective of the mass point chosen. The relative uncertainty due to parton distribution functions (PDF) is studied using the PDF4LHC procedure [17] and is less than 1% for all mass points. The dependence of the acceptance on the choice of the renormalisation and factorisation scales, which are chosen to be equal and are varied by factors of 0.5 and 2, is found to be well below 0.5%. These uncertainties are applied in the cross-section limit calculation.

6.1 Track-finding efficiency

Understanding the efficiency to reconstruct a track as a function of its impact parameter is a crucial aspect of the analysis. Three methods are used to assess if the efficiency to reconstruct displaced tracks is correctly modelled by the simulation.

The first method consists of a direct measurement of the efficiency to reconstruct isolated tracks, using cosmic ray muons. Events are selected from dedicated runs with no beam activity and the cosmic ray muons are reconstructed by combining the hits in the muon chambers from opposite halves of CMS. The efficiency to reconstruct a tracker track associated with a cosmic ray muon, as a function of the transverse and longitudinal impact parameters, is shown in Fig. 3. We focus principally on the region $|d_0| < 20 \text{ cm}$, since in simulated signal events, the reconstructed tracks from displaced vertices lie predominantly in this region. To estimate the systematic uncertainty on the dilepton efficiency, we use the measured tracking efficiency as a function of impact parameter, in data and simulation, to estimate the efficiency to reconstruct a pair of leptons of given impact parameters. We then weight this efficiency according to the simulated impact parameter distributions of the dileptons in signal Monte Carlo samples. The ratio of the estimated efficiency in data to simulation differs from unity by no more than 6.1% for any of the considered LL particle masses, so is taken as the systematic uncertainty.

A second method is used to study how the presence of a high density of tracker hits around displaced leptons degrades the track reconstruction. An upper bound on the effect of any inaccuracy in the simulation of this, is obtained by comparing the efficiency to reconstruct the tracks

of simulated single muons, firstly in a background of hits obtained from LHC collision events and secondly from simulated collision events. The relative disagreement is 1.7% (resulting in 3.4% for a dilepton candidates with two tracks).

A third method [7] uses charged pions from K_s^0 decay to establish that the tracking efficiency is simulated with a relative systematic uncertainty of 5%. This verifies that the p_T spectrum of the K_s^0 is correctly simulated using K_s^0 with small decay lengths, for which the tracking efficiency is close to 100%. Then it compares in data and simulation, distributions such as that of the reconstructed K_s^0 decay length or that of the charged pion impact parameter distribution. Any differences in these can be attributed to imperfect simulation of the tracking efficiency for non-prompt particles. Since this method is mainly sensitive to the tracking efficiency for soft hadrons in jets, it is only used to provide additional reassurance that the displaced tracking efficiency is well modelled.

These methods do not explicitly measure tracking efficiency for dielectron candidates. However, simulation studies indicate that the electron tracking efficiency is only about 22% relative smaller than the muon efficiency, where the difference can be attributed to bremsstrahlung in the electron case. The material budget of the tracker is modelled in simulation with an accuracy better than 10% [18]. Since the amount of bremsstrahlung should be proportional to the amount of material in the tracker, this implies a corresponding relative uncertainty in the difference between the tracking efficiencies for electrons and muons. The corresponding systematic uncertainty for the dielectron candidates, which have two tracks, is double this, namely 5.8%.

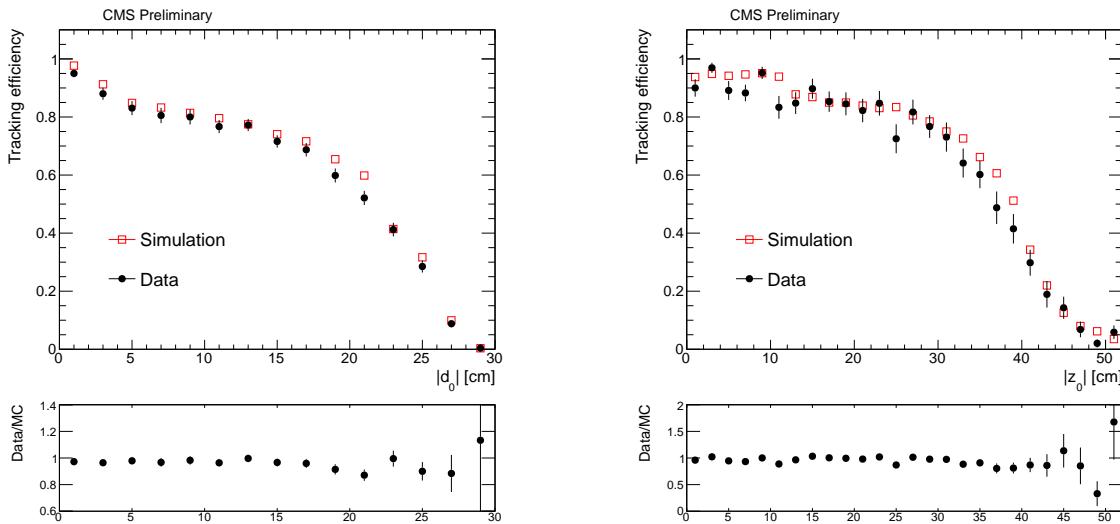


Figure 3: Efficiency of the tracker to find a track, given a cosmic ray muon reconstructed in the muon chambers, as a function of the transverse (left) and longitudinal (right) impact parameters (with respect to the nominal interaction point of CMS). The efficiency is plotted in bins of 2 cm width. For the left plot, the longitudinal impact parameter $|z_0|$ is required to be less than 10 cm, and for the right plot, the transverse impact parameter $|d_0|$ must be less than 4 cm. The bottom plots show the ratio of the efficiencies in data over simulation.

6.2 Trigger efficiency measurement

The trigger efficiency is measured using the ‘tag-and-probe’ method [19]. In the muon channel, the decays of Z bosons to dimuons are reconstructed in data collected with single-muon triggers. They are used to measure the efficiency for a lepton to pass one leg of the dimuon triggers used in this analysis. The dimuon trigger efficiency is then obtained as the square of this

single-lepton efficiency, which assumes that there is no correlation in efficiency between the two leptons. This is generally a good assumption except in the muon channel, where dimuons separated by $\Delta R < 0.2$ must be excluded because the trigger is inefficient for closely spaced dimuons. In the electron channel, the analysis is similar, although since the two legs of the trigger have different thresholds, each leg is measured separately. The systematic uncertainty associated with the trigger efficiency is evaluated by taking the difference between the estimates of the efficiency from data and simulation, yielding a total relative uncertainty of 1.7% for the electron channel and 6.2% for the muon channel. To ensure that the efficiencies obtained from the Z sample, which contains prompt electrons and muons, are valid for our signal sample as well, we also examine the trigger efficiency in simulated signal events as a function of the mean lifetime of the X boson. We observe that the electron efficiency changes by less than $\pm 1.5\%$ over the range of $c\tau$ values considered in our analysis, which is smaller than the systematic uncertainty on the trigger efficiency assigned in this channel; the change in the muon channel is approximately $\pm 3\%$, significantly smaller than the assigned systematic, so we can conclude that the existing systematic is adequate to cover the uncertainty from this source.

6.3 Effect of higher-order QCD corrections

For $M_{H^0} = 125 \text{ GeV}/c^2$, the leptons from the X boson decay have a combined efficiency of only a few percent for passing the lepton p_T requirements. For this reason the signal efficiency at this mass is sensitive to the modelling of the Higgs p_T spectrum, which may in turn be influenced by higher order QCD corrections. To study this effect, we reweight the LO $H^0 p_T$ spectrum from our signal sample to match the corresponding Higgs p_T spectrum evaluated at NLO. For $M_{H^0} = 125 \text{ GeV}/c^2$ and $M_X = 20$ (50) GeV/c^2 the signal efficiency changes by 5% (7%). This change is taken as an additional systematic uncertainty in the efficiency for the $M_{H^0} = 125 \text{ GeV}/c^2$ case. For larger H^0 masses, the corresponding systematic uncertainty is below 0.5%, and hence neglected.

7 Results

Events from background are equally likely to populate the signal and control regions, whereas any events arising from LL particles will populate almost exclusively the signal region. In consequence, the presence of a signal in the data would reveal itself as a statistically significant excess of events in the signal region compared to the control region. After all selection requirements are applied, no events survive in the signal or control regions of either the muon or electron channels. There is thus no statistically significant excess.

We set 95% confidence level (CL) upper limits on the signal processes using the Bayesian method described in Ref. [20]. The limits are determined from a comparison of the number of events N_S seen in the signal region with the number expected in the signal plus background hypothesis.

The limit calculation takes into account the systematic uncertainties on the signal described in Section 6 by introducing a nuisance parameter for each uncertainty, marginalized by a log-normal prior distribution. The expected number of background events μ_B in the control region, and hence also in the signal region, is an additional nuisance parameter. It is constrained by the observed number of events N_C in the control region. Its probability distribution $p(\mu_B)$ is given by $p(\mu_B) = \frac{\mu_B^{N_C}}{N_C!} \exp(-\mu_B)$, as can be shown using Bayesian methodology assuming a flat prior in μ_B [20]. The expected background in the signal region may differ from that in the control region, as a result of tracker misalignment. This is taken into account by allowing for

the systematic uncertainty on it described in Section 5, and in addition, by evaluating the limits twice, once with and once without correcting the track impact parameters for misalignment, as described in Section 5. The worse of these two sets of limits is taken as the result.

If a genuine signal were present, it would give rise to an excess of events in the signal region with an expected number of:

$$\mu_S = \mathcal{L}\sigma [2B(1-B)\epsilon_1 + \epsilon_2 B^2] (1-f) \quad (1)$$

where \mathcal{L} is the integrated luminosity, $\epsilon_{(1,2)}$ are the signal efficiencies defined in Section 4.1, σ is the production cross section of $H^0 \rightarrow XX$ (or $\tilde{q}\tilde{q} + \tilde{q}\tilde{q}$) and B is the branching fraction for the decay $X \rightarrow \ell^+\ell^-$ (or $\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu$). The parameter f is the mean number of genuine signal events expected to fall in the control region for each genuine signal event in the signal region. It is very small, being less than 3% (5%) for all the $X \rightarrow \ell^+\ell^-$ ($\tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu$) samples considered here. Its effect is to reduce slightly the effective signal efficiency, by causing some of the signal to be misinterpreted as background. One expects $\epsilon_2 \geq 1 - (1 - \epsilon_1)^2$, where the two terms will be equal if the efficiency to select each of the two LL particles in an event is independent of the other, or the first term will be larger if the presence of one LL particle increases the efficiency to select the other (as can happen if one lepton from each causes the event to trigger). Assuming $\epsilon_2 = 1 - (1 - \epsilon_1)^2$, which is conservative since it minimises the value of μ_S , transforms Eq. (1) into:

$$\mu_S = 2\mathcal{L}\sigma B\epsilon_1 [1 - B\epsilon_1/2] (1-f) \quad (2)$$

Since μ_S in Eq. (2) depends not only on σB , but also on B , the upper limits on σB depend on the assumed value of B , scaling approximately as $1/[1 - 0.5B\epsilon_1]$. They are thus best for low values of B , though the dependence of the limits on B is fairly small, particularly if ϵ_1 is small.

For each combination of the H^0 and X boson masses that are modelled, and for a range of X boson lifetimes, the 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)$ are calculated. The observed limits are shown in Figs. 4–7. (The worse limits for the muon channel in the $M_{H^0} = 1000 \text{ GeV}/c^2, M_X = 20 \text{ GeV}/c^2$ case are caused by the poor trigger efficiency, and associated ΔR requirement, for nearby muons.) The corresponding limits on $\sigma(\tilde{q}\tilde{q} + \tilde{q}\tilde{q})B(\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu)$ are shown in Fig. 8. The solid curves in these plots show the limits which are valid for $B\epsilon_1 \ll 1$, whilst the dotted curve shows limits valid for $B = 1$. There is negligible difference between them, except for the highest Higgs or squark masses, when the efficiencies are largest. For intermediate values of B , the limit will lie between the solid and dotted curves. The bands in all these plots show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

At $\sqrt{s} = 8 \text{ TeV}$, the theoretical cross section for SM Higgs production through the dominant gluon-gluon fusion mechanism is 19.3, 7.1, 2.9 and 0.03 pb for H^0 boson masses of 125, 200, 400 and $1000 \text{ GeV}/c^2$, respectively [21]. The theoretical cross section for $\tilde{q}\tilde{q} + \tilde{q}\tilde{q}$ production is 2590, 10, 0.014, 0.00067 pb for \tilde{q} masses of 120, 350, 1000 and $1500 \text{ GeV}/c^2$, as evaluated with the PROSPINO generator [22] assuming a gluino mass of $5 \text{ TeV}/c^2$. The observed limits on σB are usually well below these theoretical cross sections, implying that non-trivial bounds are being placed on the decay modes involving LL particles.

Figs. 9–12 show the limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)A(X \rightarrow \ell^+\ell^-)$, where A is the acceptance defined in Sect. 4.1. Fig. 13 shows the corresponding limits on $\sigma(\tilde{q}\tilde{q} + \tilde{q}\tilde{q})B(\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu)A(\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu)$. These non-acceptance corrected limits show substantially less dependence on the Higgs and X boson masses and the lifetime. They should also be less model dependent. Indeed, it can be seen that the limits on $\sigma B A$ are similar for $X \rightarrow \ell^+\ell^-$ and $\tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu$.

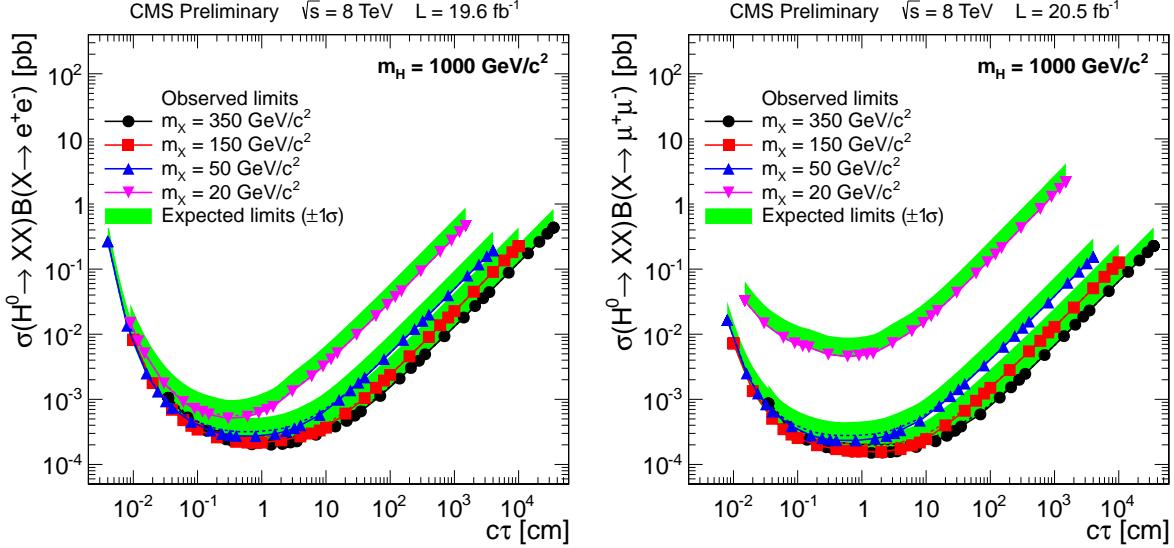


Figure 4: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $1000 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $Be_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

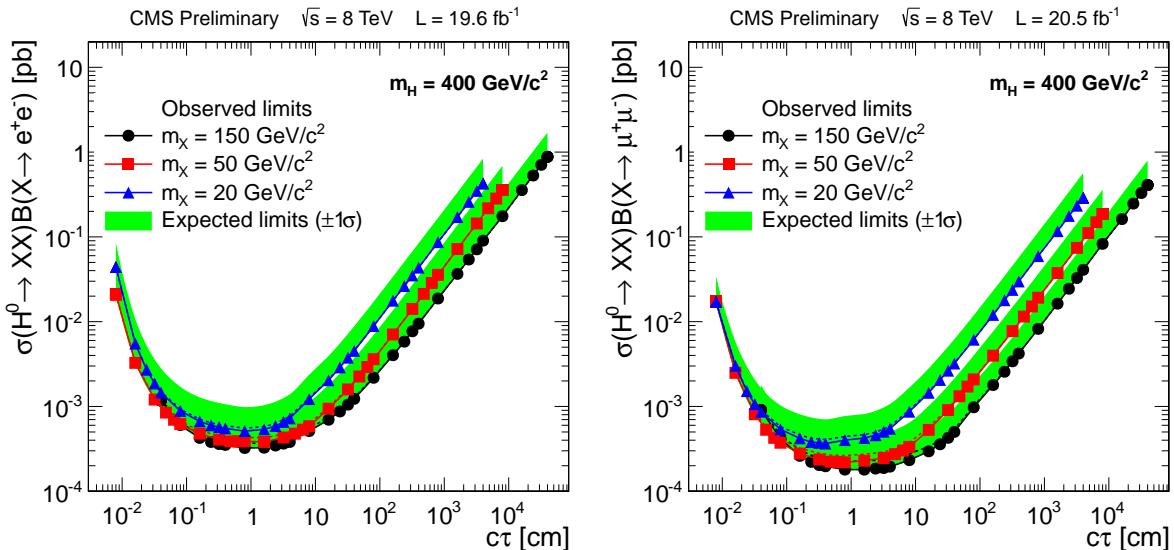


Figure 5: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $400 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $Be_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

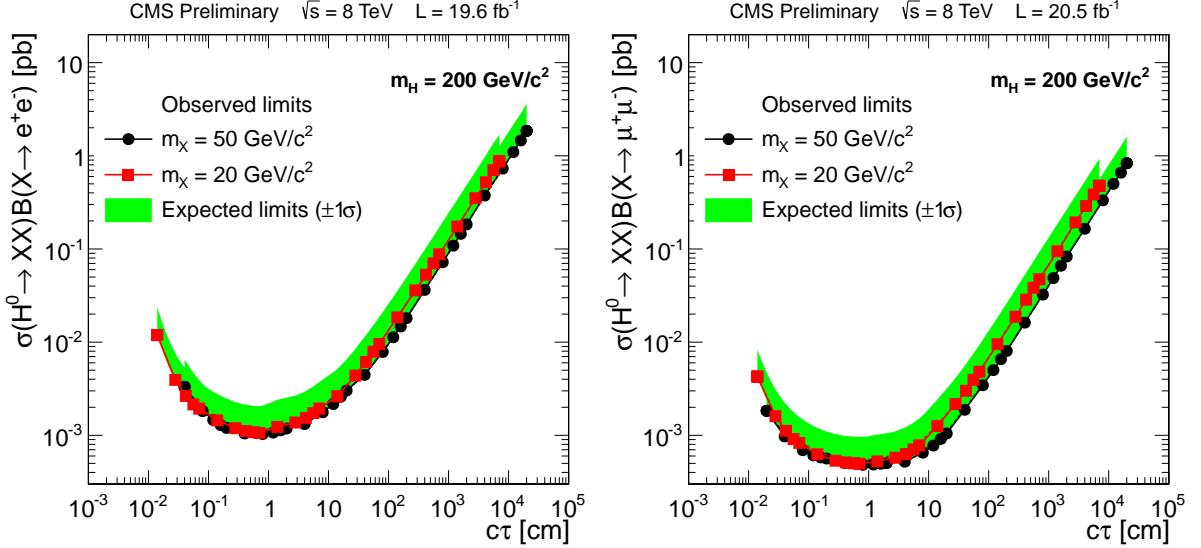


Figure 6: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $200 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

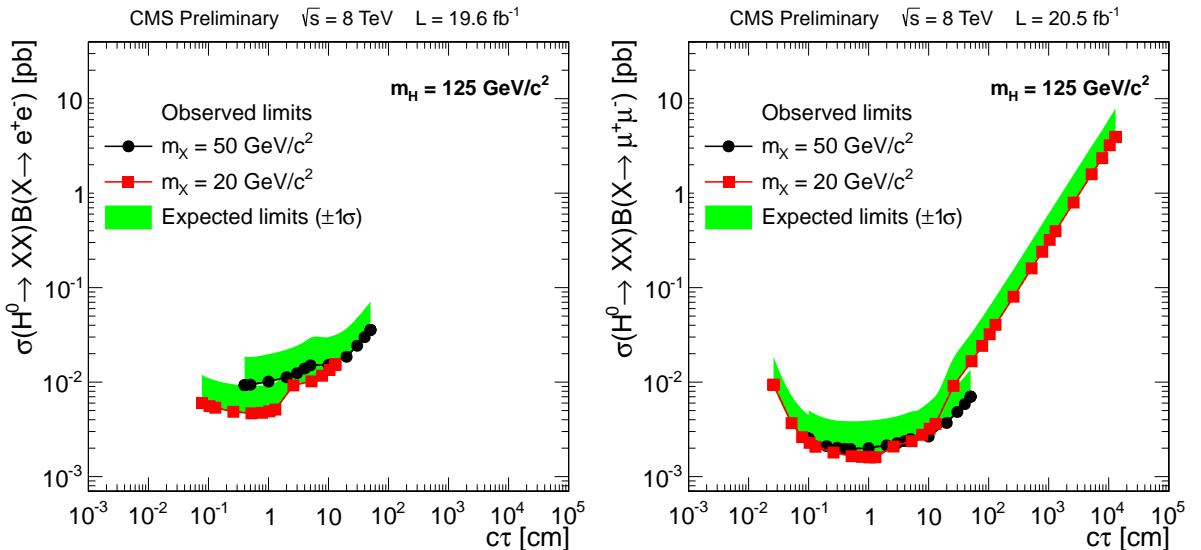


Figure 7: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $125 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

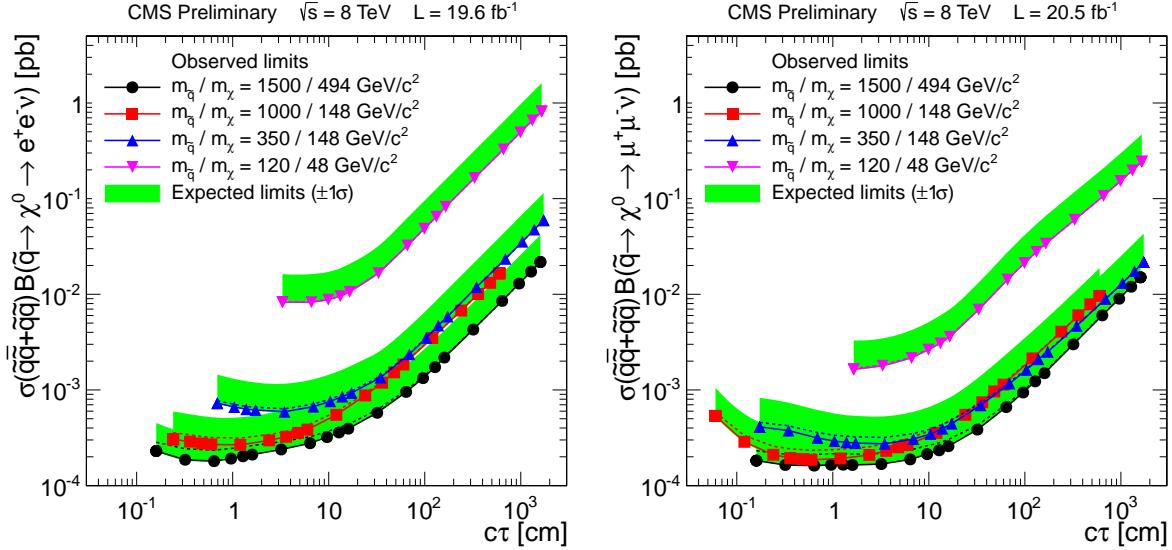


Figure 8: The 95% CL upper limits on $\sigma(\tilde{q}\tilde{q} + \tilde{q}\tilde{q})B(\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu)$ for the electron (left) and muon channel (right), as a function of the neutralino lifetime. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

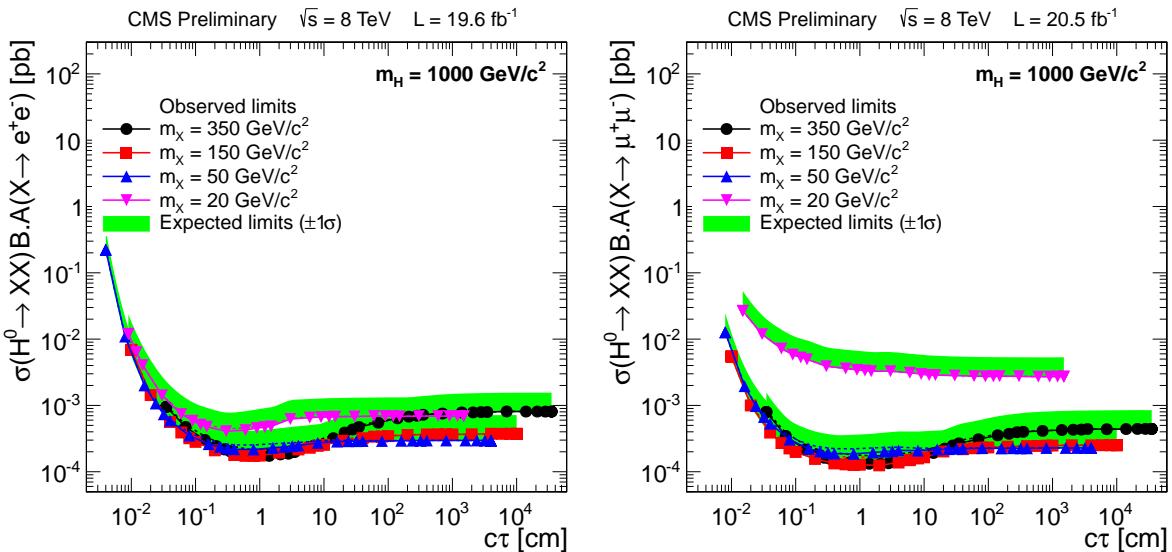


Figure 9: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)A(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $1000 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

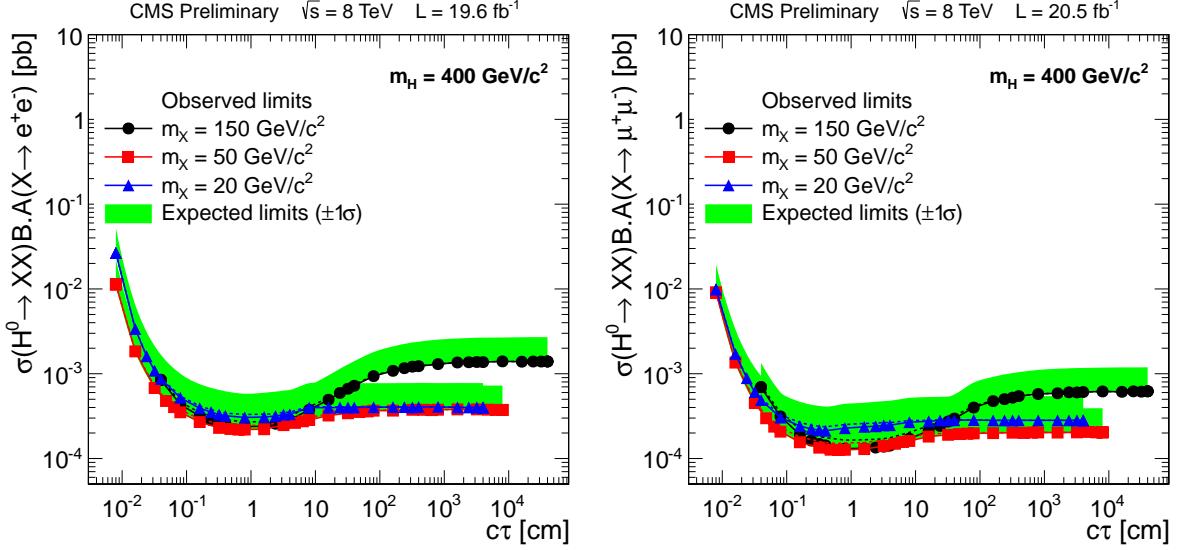


Figure 10: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)A(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $400 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

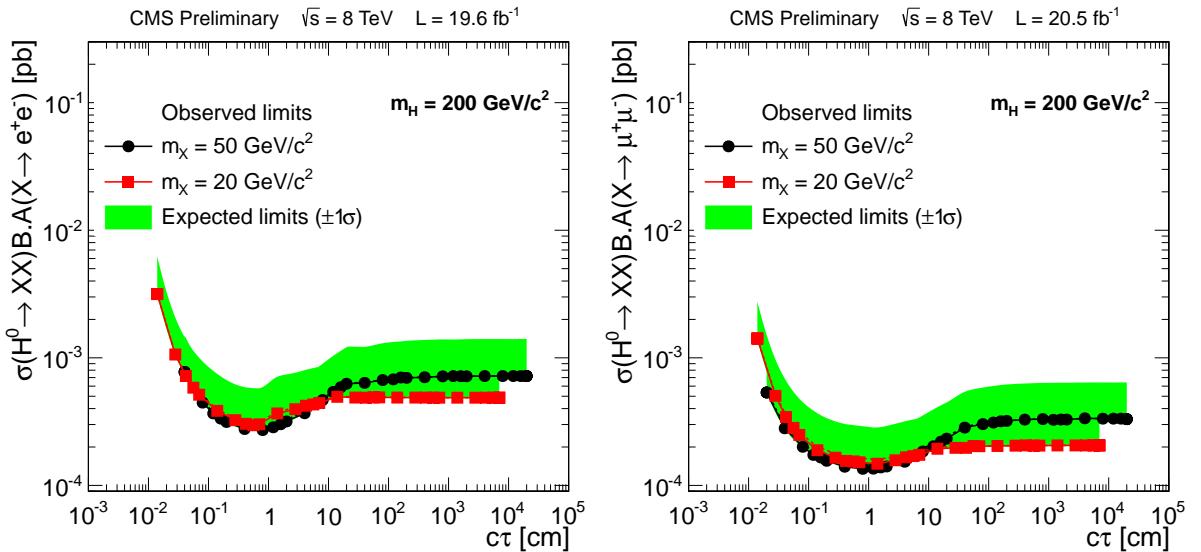


Figure 11: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)A(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $200 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

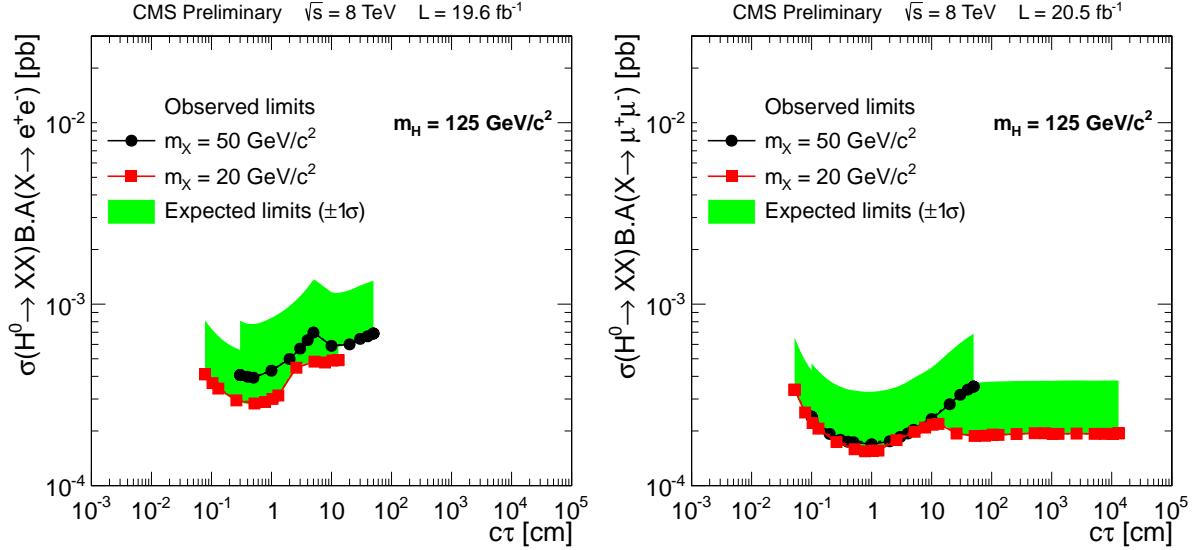


Figure 12: The 95% CL upper limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)A(X \rightarrow \ell^+\ell^-)$ for the electron (left) and muon channel (right) for a H^0 mass of $125 \text{ GeV}/c^2$. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

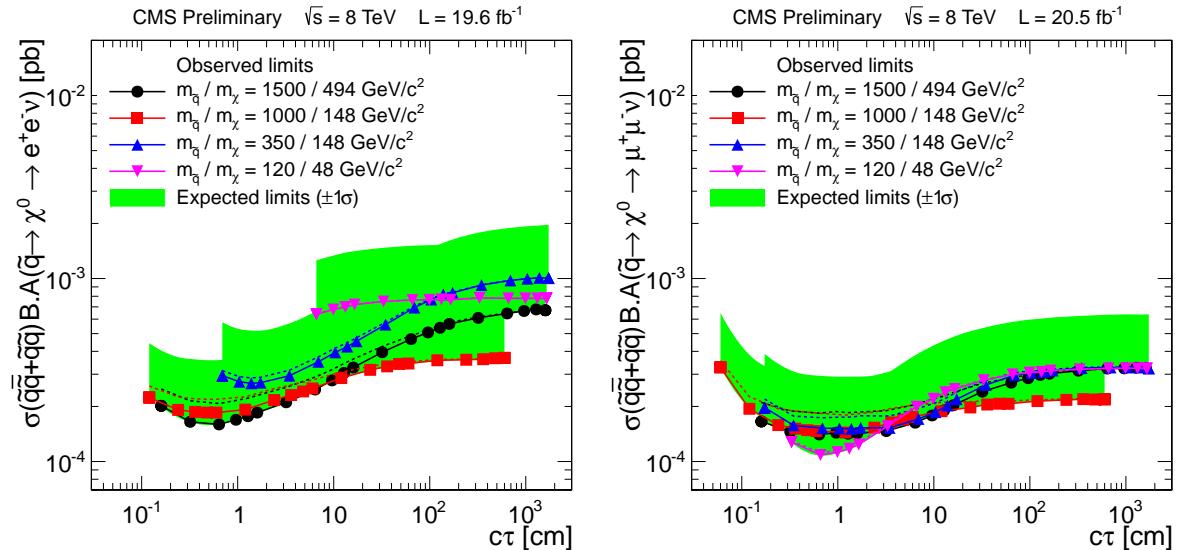


Figure 13: The 95% CL upper limits on $\sigma(\tilde{q}\tilde{q} + \tilde{q}\tilde{q})B(\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu)A(\tilde{q} \rightarrow q\tilde{\chi}^0, \tilde{\chi}^0 \rightarrow \ell^+\ell^-\nu)$ for the electron (left) and muon channel (right), as a function of the neutralino lifetime. The solid curves in these plots show the observed limits evaluated for $B\epsilon_1 \ll 1$, whilst nearby dotted curves (often hidden under a solid curve) show them evaluated for $B = 1$. Green shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits.

For H^0 or X boson (or \tilde{q} or $\tilde{\chi}^0$) masses other than those plotted in Figs. 4–13, the exact limit is not known, since no simulated signal samples are available with which to determine the signal selection efficiency. However, one can infer approximate limits for other masses, provided the latter lie within the range of the simulated samples. For example, for $M_{H^0} = 1000 \text{ GeV}/c^2$, it should be safe to assume that the limits for $M_X = 170 \text{ GeV}/c^2$ would be at least as good as the weaker of the limits for $M_X = 150 \text{ GeV}/c^2$ and $M_X = 350 \text{ GeV}/c^2$. The limits on σBA show very little mass dependence, allowing such interpolations to be made with confidence.

Although the limits shown above are determined in the context of two specific models, the analysis is sensitive to any process in which a LL particle is produced and subsequently decays to dileptons (plus possibly additional particles). To place approximate limits on this more general class of models, it is best to use the non-acceptance corrected limits (on σBA) shown in Figs. 9–13, because of their smaller model dependence.

In any model where each event contains two identical LL particles that can decay to dileptons (plus possibly additional particles), the above limits on σBA should remain approximately valid. However, four effects can introduce a model dependence into the efficiency to select signal inside the acceptance criteria, and so render these limits inaccurate: (i) Even within the acceptance requirement $L_{xy} < 50 \text{ cm}$, the lepton reconstruction efficiency depends on both the impact parameters of the lepton and the radial and z -coordinates of its production point relative to the primary vertex. It decreases for larger values of these parameters. The lepton selection efficiency also decreases for small transverse impact parameters, due to the selection requirement $|d_0|/\sigma_d > 12$, which since $\sigma_d \approx 15 \mu\text{m}$ for high p_T leptons, makes the analysis insensitive to leptons with $|d_0| \lesssim 180 \mu\text{m}$. (ii) The requirement that the dilepton mass should exceed $15 \text{ GeV}/c^2$, and the additional requirement in the dimuon channel that the two muons should have $\Delta R > 0.2$, makes the analysis insensitive to models in which the two leptons are close together. (iii) Models in which the parameter f is large, for example because the LL particle is slow moving and has a high decay multiplicity, will have worse limits. If this parameter is estimated from simulation, the corresponding worsening of the limit can be inferred from Eq. (2). (iv) The leptons are assumed to be isolated.

In models where each event contains only one LL particle that can decay to dileptons (plus possibly additional particles), the expected number of selected signal events for given σB will be up a factor of two lower, and so the limits on σBA will be up to a factor of two worse. The same considerations about model-dependence discussed in the previous paragraph still apply.

The acceptance A for any given model can be determined with a generator-level simulation, allowing limits on σBA to be converted to limits on σB . The following example illustrates this. The limits on $\sigma(H^0 \rightarrow XX)B(X \rightarrow \ell^+\ell^-)$ quoted above are for H^0 bosons produced through gluon-gluon fusion. If the H^0 bosons were instead produced by the sum of all SM production mechanisms, their momentum spectra would be slightly harder. For $M_{H^0} = 125 \text{ GeV}/c^2$, the acceptance would then be larger by a factor of approximately 1.18 (1.12) for $M_X = 20$ (50) GeV/c^2 , with a corresponding improvement in the limits on σB . The change is smaller for larger H^0 boson masses.

8 Summary

A search has been performed, using pp collision data collected at $\sqrt{s} = 8 \text{ TeV}$, for long-lived particles that decay to a final state, which includes a pair of leptons (e^+e^- or $\mu^+\mu^-$). No such events have been seen. Quantitative limits have been placed on such a signal in the context of two specific models. In the first model, a Higgs boson, in the mass range $125 - 1000 \text{ GeV}/c^2$,

decays into a pair of X bosons, in the mass range $20 - 350 \text{ GeV}/c^2$, which each decay to dileptons. The upper limits obtained are typically in the range $0.1 - 5 \text{ fb}$ for X bosons with lifetimes in the range $0.01 < c\tau < 100 \text{ cm}$, but they weaken to $5 - 50 \text{ fb}$ for the lowest considered Higgs mass of $125 \text{ GeV}/c^2$. In the second model, events contain a pair of squarks, which each decay to a long-lived neutralino $\tilde{\chi}^0$, which subsequently decays to $e^+e^-\nu$ or $\mu^+\mu^-\nu$. In this case, the upper limits are typically in the range $2 - 5 \text{ fb}$ for $\tilde{\chi}^0$ lifetimes in the range $0.1 < c\tau < 100 \text{ cm}$ and squark masses above $350 \text{ GeV}/c^2$. For a lower squark mass of $120 \text{ GeV}/c^2$, the limits are typically a factor ten weaker. These limits are the most stringent in these channels to date. To allow the results to be reinterpreted in the context of other models, limits are also presented within the detector acceptance, rendering them less model dependent.

Acknowledgements

We congratulate our colleagues in the CERN accelerator departments for the excellent performance of the LHC and thank the technical and administrative staffs at CERN and at other CMS institutes for their contributions to the success of the CMS effort. In addition, we gratefully acknowledge the computing centres and personnel of the Worldwide LHC Computing Grid for delivering so effectively the computing infrastructure essential to our analyses. Finally, we acknowledge the enduring support for the construction and operation of the LHC and the CMS detector provided by the following funding agencies: BMBF and FWF (Austria); FNRS and FWO (Belgium); CNPq, CAPES, FAPERJ, and FAPESP (Brazil); MEYS (Bulgaria); CERN; CAS, MoST, and NSFC (China); COLCIENCIAS (Colombia); MSES (Croatia); RPF (Cyprus); MoER, SF0690030s09 and ERDF (Estonia); Academy of Finland, MEC, and HIP (Finland); CEA and CNRS/IN2P3 (France); BMBF, DFG, and HGF (Germany); GSRT (Greece); OTKA and NKTH (Hungary); DAE and DST (India); IPM (Iran); SFI (Ireland); INFN (Italy); NRF and WCU (Korea); LAS (Lithuania); CINVESTAV, CONACYT, SEP, and UASLP-FAI (Mexico); MSI (New Zealand); PAEC (Pakistan); MSHE and NSC (Poland); FCT (Portugal); JINR (Armenia, Belarus, Georgia, Ukraine, Uzbekistan); MON, RosAtom, RAS and RFBR (Russia); MSTD (Serbia); SEIDI and CPAN (Spain); Swiss Funding Agencies (Switzerland); NSC (Taipei); ThEP, IPST and NECTEC (Thailand); TUBITAK and TAEK (Turkey); NASU (Ukraine); STFC (United Kingdom); DOE and NSF (USA).

References

- [1] J. L. Hewett, B. Lillie, M. Masip, and T. G. Rizzo, “Signatures of long-lived gluinos in split supersymmetry”, *JHEP* **09** (2004) 070, doi:10.1088/1126-6708/2004/09/070, arXiv:hep-ph/0408248.
- [2] R. Barbier et al., “R-parity violating supersymmetry”, *Phys. Rept.* **420** (2005) 1, doi:10.1016/j.physrep.2005.08.006, arXiv:hep-ph/0406039.
- [3] T. Han, Z. Si, K. M. Zurek, and M. J. Strassler, “Phenomenology of hidden valleys at hadron colliders”, *JHEP* **07** (2008) 008, doi:10.1088/1126-6708/2008/07/008, arXiv:0712.2041.
- [4] L. Basso, A. Belyaev, S. Moretti, and C. H. Shepherd-Themistocleous, “Phenomenology of the minimal B-L extension of the Standard model: Z' and neutrinos”, *Phys. Rev. D* **80** (2009) 055030, doi:10.1103/PhysRevD.80.055030, arXiv:0812.4313.

- [5] M. J. Strassler and K. M. Zurek, “Discovering the Higgs through highly-displaced vertices”, *Phys. Lett. B* **661** (2008) 263, doi:10.1016/j.physletb.2008.02.008, arXiv:hep-ph/0605193.
- [6] CMS Collaboration, “Search in leptonic channels for heavy resonances decaying to long-lived neutral particles”, *JHEP* **1302** (2013) 085, doi:10.1007/JHEP02(2013)085, arXiv:1211.2472.
- [7] CMS Collaboration, “Search for long-lived neutral particles decaying to dijets”, CMS Physics Analysis Summary CMS-PAS-EXO-12-038, 2012.
- [8] D0 Collaboration, “Search for neutral, long-lived particles decaying into two muons in $p\bar{p}$ collisions at $\sqrt{s} = 1.96\text{-TeV}$ ”, *Phys. Rev. Lett.* **97** (2006) 161802, doi:10.1103/PhysRevLett.97.161802, arXiv:hep-ex/0607028.
- [9] D0 Collaboration, “Search for long-lived particles decaying into electron or photon pairs with the D0 detector”, *Phys. Rev. Lett.* **101** (2008) 111802, doi:10.1103/PhysRevLett.101.111802, arXiv:0806.2223.
- [10] ATLAS Collaboration, “Search for a light Higgs boson decaying to long-lived weakly-interacting particles in proton-proton collisions at $\sqrt{s} = 7\text{ TeV}$ with the ATLAS detector”, *Phys. Rev. Lett.* **108** (2012) 251801, arXiv:1203.1303.
- [11] ATLAS Collaboration, “Search for displaced vertices arising from decays of new heavy particles in 7 TeV pp collisions at ATLAS”, *Phys. Lett. B* **707** (2012) 478, doi:10.1016/j.physletb.2011.12.057, arXiv:1109.2242.
- [12] CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **03** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.
- [13] CMS Collaboration, “CMS Tracking Performance Results from Early LHC Operation”, *Eur. Phys. J. C* **70** (2010) 1165, doi:10.1140/epjc/s10052-010-1491-3, arXiv:1007.1988.
- [14] S. M. T. Sjöstrand and P. Z. Skands, “PYTHIA 6.4 Physics and Manual”, *JHEP* **05** (2006) 576, arXiv:0603175.
- [15] GEANT4 Collaboration, “GEANT4: A simulation toolkit”, *Nucl. Instrum. Meth. A* **506** (2003) 250, doi:10.1016/S0168-9002(03)01368-8.
- [16] CMS Collaboration, “CMS luminosity based on pixel cluster counting - Summer 2013 update”, CMS Physics Analysis Summary CMS-PAS-LUM-13-001, 2013.
- [17] D. Bourilkov, R. C. Group, and M. R. Whalley, “LHAPDF: PDF use from the Tevatron to the LHC”, (2006). arXiv:hep-ph/0605240.
- [18] CMS Collaboration, “Studies of tracker material”, CMS Physics Analysis Summary CMS-PAS-TRK-10-003, 2010.
- [19] CMS Collaboration, “Measurements of Inclusive W and Z Cross Sections in pp Collisions at $\sqrt{s} = 7\text{ TeV}$ ”, *JHEP* **01** (2011) 080, doi:10.1007/JHEP01(2011)080, arXiv:1012.2466.
- [20] ATLAS Collaboration, CMS Collaboration, LHC Higgs Combination Group, “Procedure for the LHC Higgs boson search combination in Summer 2011”, Public Note ATL-PHYS-PUB-2011-011, ATL-COM-PHYS-2011-818, CMS-NOTE-2011-005, 2011.

- [21] LHC Higgs Cross Section Working Group Collaboration, “Handbook of LHC Higgs Cross Sections: 3. Higgs Properties”, [doi:10.5170/CERN-2013-004](https://doi.org/10.5170/CERN-2013-004), [arXiv:1307.1347](https://arxiv.org/abs/1307.1347).
- [22] W. Beenakker, R. Hopker, M. Spira, and P. Zerwas, “Squark and gluino production at hadron colliders”, *Nucl.Phys.* **B492** (1997) 51–103, [doi:10.1016/S0550-3213\(97\)80027-2](https://doi.org/10.1016/S0550-3213(97)80027-2), [arXiv:hep-ph/9610490](https://arxiv.org/abs/hep-ph/9610490).