Available on the CERN CDS information server **CMS PAS EXO-14-012**

CMS Physics Analysis Summary

Contact: cms-pag-conveners-exotica@cern.ch 2015/04/01

Search for long-lived particles that decay into final states containing two muons reconstructed using only the CMS muon chambers

The CMS Collaboration

Abstract

A search is performed for long-lived neutral particles that decay into final states that include a pair of muons. These are reconstructed using only the CMS muon chambers. The experimental signature is a distinctive topology consisting of a pair of muons originating from a displaced secondary vertex. Events corresponding to an integrated luminosity of 20.5 fb^{-1} were collected with the CMS detector at the CERN LHC in pp collisions at $\sqrt{s} = 8$ TeV. No significant excess is observed above standard model expectations. Upper limits on the product of the cross section and branching fraction of such a signal are presented as a function of the long-lived particle's mean proper decay length. The limits are provided for two specific models. In the first model, a Higgs boson decays into a pair of long-lived, neutral bosons, each of which can decay into a pair of muons. In the second model, squarks are pair produced and the decay of each produces a long-lived neutralino that subsequently decays to two muons and a neutrino. The limits are also presented in an approximately model-independent way, allowing them to be applied to a wide class of models yielding the above topology. A combination of these limits with those from a related search that used muons reconstructed in the CMS silicon tracker is also presented.

1 Introduction

The existence of massive, long-lived particles, which might be produced at the LHC, is conjectured by several extensions to the standard model, such as "split SUSY" [1], SUSY with very weak R-parity violation [2], "hidden valley" models [3] and Z′ models that contain long-lived neutrinos [4]. In models where the long-lived massive particles decay into final states that include a pair of charged leptons, these leptons can be differentiated from those produced by standard model (SM) processes by virtue of their significant displacement from the LHC beamline.

This study presents the results of a search for long-lived neutral particles decaying to muon pairs reconstructed using only the muon chambers of the Compact Muon Solenoid (CMS) detector. The analysis uses data taken during 2012 in proton-proton (pp) collisions at $\sqrt{s} = 8\,\text{TeV}$, corresponding to a total integrated luminosity of $20.5\,\text{fb}^{-1}$. As benchmarks for this physics signature, two models are considered to quantify the sensitivity of the analysis. The first model postulates pair production of long-lived, spinless X bosons from the decay of a non-SM Higgs boson, $\text{H}^0 \to \text{XX}$, where H^0 is produced by gluon-gluon fusion and each X boson can decay to lepton pairs, $\text{X} \to \mu^+\mu^-$ [5]. In the second model, a pair of squarks each decays as $\tilde{\text{q}} \to q\tilde{\chi}^0$, $\tilde{\chi}^0 \to \mu^+\mu^-\nu$, where the long-lived neutralino, χ^0, can produce a lepton pair and a neutrino as a result of R-parity violation [2]. Even though the results are given for these specific models, they are presented in a way that is approximately model independent, and can be used to derive concrete limits on a wide array of models producing decays of long-lived particles to muons.

This study is closely related to another CMS analysis [6], which employed similar analysis techniques to search for the same signal topology, but differed in that it relied on leptons (in that case, electrons or muons) whose trajectories were reconstructed in the CMS silicon tracker. Since the muon chambers are able to reconstruct muons produced an order of magnitude farther away from the beam-line than what the silicon tracker can, the new search improves sensitivity to particles that are especially long-lived. This search is fully complementary to the tracker-based one in that it explicitly excludes any muon whose trajectory is reconstructed in the silicon tracker. Combined results from the two searches are presented.

This analysis complements two recent CMS publications: one searching for events that contain one electron and one muon from long-lived (LL) particle decays [7], and another that searches for LL particles decaying to dijets [8]. The D0 Collaboration has published the results of a search for leptons from non-prompt decays in its tracker volume [9, 10], performed using $\text{p}\bar{\text{p}}$ collisions at $\sqrt{s} = 1.96\,\text{TeV}$ at the Fermilab Tevatron. The ATLAS Collaboration has also performed related searches for long-lived particles using different decay channels [11, 12], or lower-mass LL particles [13], compared to those considered in this paper.

2 The CMS detector

The central feature of the CMS apparatus 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, strip tracker, the lead-tungstate crystal electromagnetic calorimeter (ECAL) and the brass/scintillator hadron calorimeter. Muons are identified in gas-ionisation detectors embedded in the steel magnetic-flux return yoke of the solenoid. A detailed description of the complete CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in Ref. [14].

The silicon tracker can reconstruct the trajectories of charged particles such as muons, electrons and hadrons as well as their momentum with high precision. It is composed of pixel detectors (three barrel layers and two forward disks on each 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.

In order to stop electrons and photons, the electromagnetic calorimeter is placed on the periphery of the silicon tracker. The ECAL consists of nearly 76 000 lead tungstate crystals in a barrel and two endcap sections, which provide coverage in pseudorapidity $|\eta| < 3$. The hadron calorimeter is between the muon chambers and the electromagnetic calorimeter. It measures the energy of hadrons and it is made up of barrel, endcap and forward sections.

Muons are measured in the pseudorapidity range $|\eta| < 2.4$ with detection planes based on one of three technologies: drift tubes (DT) in the barrel region (radius between 4 and 7.5 m), cathode strip chambers (CSC) in the endcaps, and resistive plate chambers (RPC) in the barrel and endcaps. The muon system has three main functions: triggering on muons, muon identification, and the improvement of muon momentum measurement. The drift tube subsystem is responsible of measuring the position and direction of muons in the $|\eta| < 1.2$ region. Cathode strip chambers consist of anode wires and cathode strips positioned perpendicular to each other in the region $0.9 < |\eta| < 2.4$. The resistive plate chambers located in the pseudorapidity range of $|\eta| < 1.6$ provide additional fast muon trigger capability. Tracks can be reconstructed independently in the tracker or in the muon system. The precision of the measured track parameters can be improved by performing a single fit combining information from both systems. Muon reconstruction performance has been studied in great detail with data [15].

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

The analysis uses data taken from pp collisions at a centre-of-mass energy of 8 TeV corresponding to an integrated luminosity of 20.5 ± 0.5 fb^{-1}. Events are selected by a trigger requiring two muons in an event, each reconstructed in the muon detectors without imposing any beam spot constraint and having $p_T > 23\,\text{GeV}/c$. Both muons are also required to have at least two reconstructed hits per chamber in at least two CSC or DT chambers. To prevent cosmic ray muons from passing these criteria, the opening angle between the two muons must be less than 2.5 radians. This trigger selection is independent of the silicon tracker activity.

For the $H^0 \to XX$ model, simulated signal samples are generated using PYTHIA V6.426 [16] to simulate H^0 production through gluon fusion ($gg \to H^0$). Subsequently the H^0 is forced to decay to XX, with the X bosons each decaying to lepton pairs ($X \to \mu^+\mu^-$). The analysis focuses on the final states with at least one muon pair. The generated samples have M_{H^0} = 125, 200, 400, 1000 GeV/c^2 and M_X = 20, 50, 150, 350 GeV/c^2. Each sample is produced with three X boson lifetimes, which correspond to mean transverse decay lengths in the laboratory frame of approximately 20 cm, 200 cm and 2000 cm.

For the $\tilde{\chi}^0 \to \mu^+\mu^-\nu$ model, squark pair production is simulated by PYTHIA as well as the decay to $\tilde{\chi}^0$. The following four pairs of squark and neutralino masses are used : $(M_{\tilde{q}}, M_{\tilde{\chi}^0})$ =

(1500, 494), (1000, 148), (350, 148) and (120, 48) GeV/c^2. The R-parity violating couplings λ_{122} and λ_{121} are set to non-zero values so that the decay of the $\widetilde{\chi}^0$ into two charged leptons and a neutrino is permitted. The values of λ_{122} and λ_{121} are chosen to give a mean transverse decay length of approximately 200 cm.

Several simulated background samples are generated with PYTHIA. Among those, the dominant background is Drell–Yan production of dimuons: prompt $\mu^+\mu^-$ pairs can be misidentified as displaced from the primary vertex due to detector resolution effects, and the production and decay of $\tau^+\tau^-$ pairs can produce genuinely displaced leptons, although the probability that both τ-leptons decay leptonically is small. Other simulated backgrounds are $t\bar{t}$, W/Z boson pair production with leptonic decays, and QCD multijet events. All these backgrounds produce negligible contributions. In all the samples, the response of the detector is simulated in detail using GEANT4 [17]. The samples are then processed through the trigger emulation and event reconstruction chain of the CMS experiment. The random cosmic ray background is not simulated, so must be estimated from data.

4 Event reconstruction and selection

To select pp collisions, events are required to contain a primary vertex within a distance from the nominal interaction point of less than 2 cm in the direction transverse to the beam, and of less than 24 cm in the direction along the beam. Furthermore, for events with at least 10 tracks, the fraction of tracks classified as "high purity", as defined in [18], must exceed 25%, to reject events produced by the interaction of beam-related protons with the LHC collimators. Where more than one primary vertex is reconstructed in an event, the one with the largest sum of the p_{T}^2 of the tracks associated to it is selected.

By design, since it is complementary to the search presented in [6], this analysis does not use the silicon tracker information in muon track reconstruction. Muon candidates are reconstructed utilizing only the hits in the muon chambers and a final refit step is applied to minimize possible biases from a loose beamspot constraint in the seeding step. We will refer to them as refitted stand-alone (RSA) muons.

We require the RSA muons to satisfy $p_{\mathrm{T}} > 26\,\mathrm{GeV}/c$ and $|\eta| < 2$. The momentum threshold is slightly higher than the corresponding trigger requirement, to ensure that the trigger has a good efficiency and its systematic uncertainty is minimal.

A distinctive track rejection step is applied to make the analysis fully complementary to the tracker-based analysis described in [6] and to exclude prompt muons in the most effective way. All muons reconstructed in the muon chambers are rejected if they can be matched to a track reconstructed in the silicon tracker with $p_{\mathrm{T}} > 10\,\mathrm{GeV}/c$. The matching is done by extrapolating the track from the silicon tracker to the muon's innermost hit in the muon system. The track and the muon are considered matched if ΔR (where $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$) between the innermost hit of the muon and the extrapolated position of the tracker track in the muon chambers is less than 0.1. The p_{T} requirement on the tracker tracks is relaxed compared to [6] to allow for the poor p_{T} resolution of RSA muons. No overlap is found in the events selected by the two analyses when applying their full selection to all simulated signal samples.

To select muons of good quality, the muon track fitted to the hits in the muon chambers must satisfy the condition $\chi^2/\mathrm{dof} < 2$. Each muon must be assigned hits in at least three of the muon chambers. Given non-negligible cosmic muon contamination, each muon is also required to have at least 17 hits (in DT+CSC+RPC), since in most cases out-of-time muons of cosmic origin

tend to have a lower number of hits than in-time muons, that is, muons arising from pp collisions. Finally, muons should have a transverse impact parameter significance, $|d_0|/\sigma_d > 4$, where $|d_0|/\sigma_d$ is the ratio of the magnitude of the transverse impact parameter to its uncertainty, evaluated with respect to the primary vertex. For muons passing the above selection, omitting the track rejection step, the typical d_0 resolution is 2 cm and the typical relative p_T resolution is 40%.

Non-isolated muons, such as those produced from semileptonic b hadron decay, can have a significant impact parameter and represent a possible source of background. However, the track rejection step suppresses this source of background and no additional isolation requirement is necessary. The absence of an isolation requirement makes this analysis sensitive to models producing highly displaced b quarks, such as long-lived particles decaying to $b\bar{b}$ where the two b quarks each produce a muon. However, due to the softness of the muon p_T spectrum and to the fact that the muons do not necessarily point back to the b quark production vertex the sensitivity to models as $X \to b\bar{b}$ is not competitive with other direct searches for displaced jets, such as [8].

The LL particle candidates are formed by pairing all possible combinations of selected muons in the event. No opposite charge requirement is imposed when building the dimuon candidates, since muons in simulated signal events that are reconstructed only in the muon chambers have a probability of being assigned the correct charge that can be as low as 80% for muons displaced by $|d_0| > 50$ cm. This is due to poorer momentum resolution when compared to tracks reconstructed in the tracker, caused by the lower magnetic field in the muon chambers and the larger traversed material. We discard dimuons consistent with coming from J/ψ and Υ decays and γ conversions by requiring an invariant mass greater than $15\,\text{GeV}/c^2$, although this background is actually negligible as a result of the track rejection step. The two muon tracks are required to form a secondary vertex whose fit has $\chi^2/\text{dof} < 4$. When one muon is assigned to more than one LL candidate, we keep only the LL candidate with the smallest χ^2/dof of the secondary vertex. This procedure avoids the double counting of muons.

A significant amount of background arises from cosmic rays, which may be reconstructed as back-to-back muons that are often displaced from the primary vertex. Such events should, in principle, be removed at trigger level. However, the trigger requirement on the opening angle between the two muons (α) is tightened to $\cos(\alpha) > -0.75$. Furthermore, a dimuon candidate can also be reconstructed from the combination of a cosmic muon and another (fake or real) muon produced in the pp collision. To remove these combinations, candidates are rejected when one of the two muons is back-to-back ($\cos(\alpha) \leq -0.75$) to a third muon in the same event that is not included in another dimuon candidate.

The efficiency of the double muon trigger becomes difficult to model if the two muons are very close to each other, since the trigger is inefficient in this case. It is therefore required that the two muons are separated by $\Delta R > 0.2$. Finally, LL candidates should have a transverse decay length significance of $L_{xy}/\sigma_{L_{xy}} > 12$, where L_{xy} is defined as the distance between the primary and the secondary vertices in the transverse plane. The L_{xy} resolution is approximately 3 cm. This and the d_0/σ_d cuts are selected to give an expected background significantly smaller than one event, which produces the best signal sensitivity for most of the lifetimes considered in this paper.

The angular difference in the azimuthal plane, $\Delta\Phi$, between the dimuon momentum vector and the vector from the primary vertex to the dilepton vertex should satisfy $|\Delta\Phi| < \pi/2$, where $\Delta\Phi$ is measured in the range $-\pi < \Delta\Phi < \pi$. The region, $|\Delta\Phi| < \pi/2$, is called the signal region and the one with $|\Delta\Phi| > \pi/2$ is defined as the control region. The control region

should be signal-free, whereas the background should be symmetrically distributed in both regions. Figure 1 shows the $|\Delta\Phi|$ distribution in data and for two signal Monte Carlo samples for a relaxed selection where the d_0/σ_d and $L_{xy}/\sigma_{L_{xy}}$ cuts are not applied.

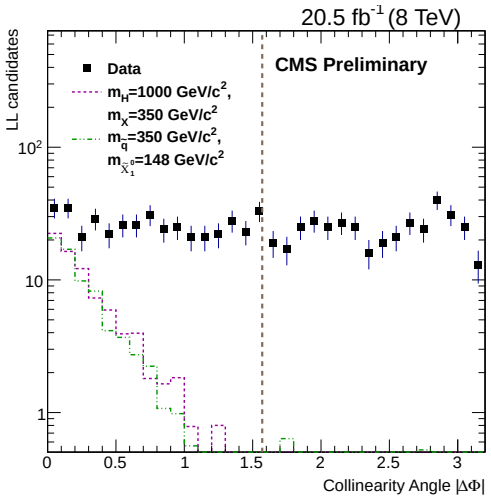

Figure 1: Distribution of the collinearity angle $|\Delta\Phi|$ in data and signal Monte Carlo simulation. The full selection is applied except the d_0/σ_d and $L_{xy}/\sigma_{L_{xy}}$ cuts. For both signal samples the product of the cross section and branching ratio is taken to be 1 pb. The $H^0 \to XX$ sample has a lifetime of 350 cm while the $\tilde{q} \to q\tilde{\chi}^0$ sample has a lifetime of 173 cm. The vertical line separates the signal region (left) from the control region (right).

We define the signal to be within the detector acceptance if it satisfies the following conditions.

- The generated transverse decay length L_{xy} of the LL particle must be $< 500\,\text{cm}$.
- The generated muon pseudorapidity must be $|\eta| < 2$.
- The generated muon momentum must satisfy $p_\text{T} > 26\,\text{GeV}/c$.

This defines a region where it is possible to reconstructed long-lived particle decays in CMS. We call acceptance (A) the fraction of long-lived particle decays that fall in the acceptance region.

Figure 2 demonstrates that our analysis is mostly sensitive to LL particles with long lifetimes while being completely insensitive to prompt events. Although RSA muons have non-null reconstruction efficiency up to 5 meters away from the beam spot in the transverse plane, the effective range of the analysis is restricted to 2.5 meters since the dimuon trigger efficiency vanishes around that distance. The inefficiency of the trigger is due to a bias in the reconstructed muon p_T as a function of the transverse impact parameter. This bias causes the p_T requirement of the trigger to be inefficient for longer decay lengths. The modest selection efficiency at low L_{xy} is caused mainly by the track-veto and the $L_{xy}/\sigma_{L_{xy}}$ cuts.

The selection efficiency and the limits are determined in terms of the number of events passing our selection, rather than the number of dimuon candidates. The signal efficiency is simply the ratio of the number of simulated events in which at least one LL candidate passes the selection requirements to the total number of simulated events. It is computed separately for two cases: firstly using events that have only one generated LL particle (X or $\tilde{\chi}^0$) decaying to muons, giving the efficiency ϵ_1; and secondly using events in which two LL particles decay to muons, giving efficiency ϵ_2. The efficiencies can be estimated for any LL particle lifetime by reweighting the generated events.

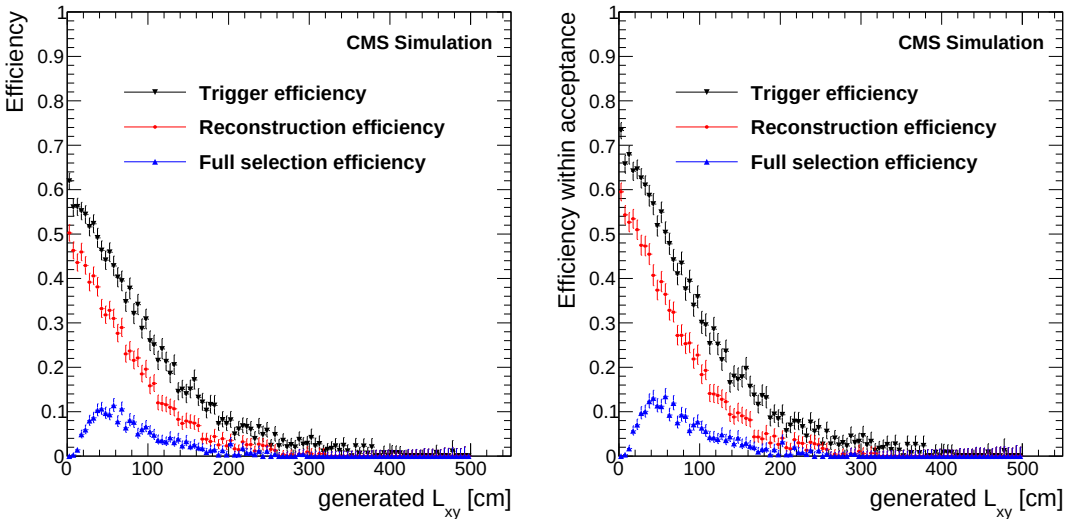

Figure 2: (Left) The efficiency to find muons from long-lived particles decays as a function of the generated L_{xy}. (Right) The efficiency computed only for particles decaying within the detector acceptance. Although the RSA muon reconstruction efficiency is non-null up to ≈ 5 meters in the transverse plane, the effective range of the analysis is up to 2.5 meters in the transverse plane due to the trigger efficiency. The signal sample shown on the diagram has the following mass points: $M_{\mathrm{H}^0} = 1000\,\mathrm{GeV}/c^2$ and $M_{\chi} = 350\,\mathrm{GeV}/c^2$ with $c\tau = 350$ cm.

5 Estimated background and associated systematic uncertainties

The signal and the background have different distributions in $|\Delta\Phi|$. The signal is expected to have small $|\Delta\Phi|$ values, assuming that the dilepton system produced when the LL particle decays is usually boosted with respect to its flight direction. In contrast, for background events, the vector from the primary to secondary vertices does not correspond to the flight direction of any long-lived particle, so its angular distribution with respect to the dilepton momentum vector should not show any forward-backward asymmetry. The background distribution is therefore expected to be symmetric in $|\Delta\Phi|$ around $\pi/2$ due to the absence of a genuine displaced secondary vertex, i.e. equal numbers of background events are expected in the signal and control regions.

No data events are observed after the full selection is applied in the control region, and this determines the expected background in the signal region to be zero events. The systematic uncertainty on this estimate is computed in Section 5.1.

The background symmetry is confirmed for data by comparing the $L_{xy}/\sigma_{L_{xy}}$ tail-cumulative distribution in the signal region with that in the control region at modest $L_{xy}/\sigma_{L_{xy}}$ and $|d_0|/\sigma_d$ values, where the data is background-dominated. The study is also repeated using simulated background events. For both studies, the full selection except the $L_{xy}/\sigma_{L_{xy}}$ requirement is implemented, although to reject any signal events, the requirement on the muon $|d_0|/\sigma_d$ is reversed, requiring that at least one muon should have $|d_0|/\sigma_d < 4$. The $L_{xy}/\sigma_{L_{xy}}$ tail-cumulative plot for data excludes the region $L_{xy}/\sigma_{L_{xy}} > 6$, where potential signal events might appear.

Figure 3 shows the tail-cumulative distributions of $L_{xy}/\sigma_{L_{xy}}$ in the signal and control regions for events in data. The difference between the signal and control regions is seen to be negligible, as expected under the symmetric background hypothesis.

The full selection is also implemented on a cosmic-enriched sample to test how effectively the

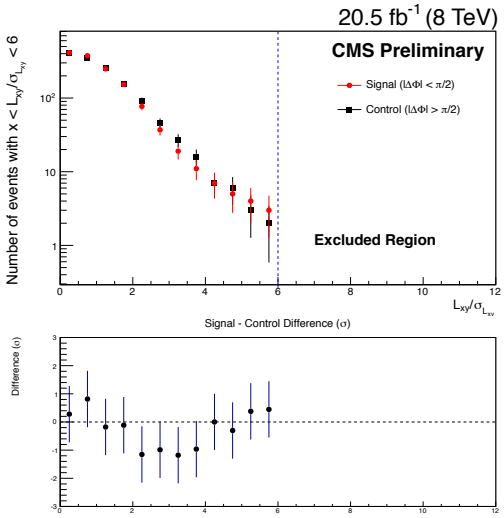

Figure 3: Comparisons of the $L_{xy}/\sigma_{L_{xy}}$ tail-cumulative distributions between the signal ,$|\Delta\Phi| < \pi/2$, and control, $|\Delta\Phi| > \pi/2$, regions for data (top). The full selection is applied with the exception of the $L_{xy}/\sigma_{L_{xy}}$ cut. The $|d_0|/\sigma_d$ requirement is reversed to $|d_0|/\sigma_d < 4$. The plot on the bottom shows the statistical significance of the difference between the two regions. The $L_{xy}/\sigma_{L_{xy}} > 6$ region is excluded to avoid possible signal contamination.

analysis rejects events of cosmic origin. The analysis is run over data collected by CMS in 2012 in cosmic runs when no beam activity was present. These data contain 13 million events triggered by a dedicated cosmic muon trigger and we reconstruct about 160,000 events with dimuon candidates. Only one of these events passes the analysis selection, corresponding to an efficiency about 200 times smaller than the lowest signal efficiency of the analysis. Therefore, dimuon candidates originating purely from cosmics are negligible after the analysis selection.

5.1 Systematic uncertainties on the estimated background

There are three main classes of systematic uncertainty in this analysis. These are the uncertainty in the integrated luminosity, the uncertainty in the signal selection efficiency, which will be discussed in Section 6, and the uncertainty in the estimated background. Since we estimate the background from the control region, a possible source of systematic uncertainty is any difference between the background distributions in the signal and control regions.

To quantify this systematic uncertainty, we compare the cumulative tail distribution in data for signal and control region for the $L_{xy}/\sigma_{L_{xy}}$ range $[x, 9]$ where x varies from zero to 8. There are not enough statistics to perform a meaningful comparison after the full selection. Therefore, we utilise a relaxed selection where one of the two RSA muons is allowed to be assigned hits in as few as two muon chambers (instead of the three required by the full selection). Since the muon chambers lie beyond the region where we are able to reconstruct secondary vertices, the decrease in the number of valid stations is not associated to muons being produced inside the muon chambers. Instead, we are selecting a set of slightly lower quality muons. The maximum deviation between the background estimated in the control region and the one from the signal region is 67%. We repeat this study further relaxing the selection to allow both RSA muons to have at least two muon chambers with valid hits. The biggest discrepancy in this case is found to be 58%. We take the larger of these two results (67%) as the systematic uncertainty on the background estimate. An explanation of how the expected background and the related systematic uncertainty are utilized in the limit calculation is given in Section 7.

6 Signal systematic uncertainties

The dominant systematic uncertainties in the analysis are associated with the signal efficiency and are caused by uncertainties in the trigger efficiency and in the reconstruction efficiencies of the displaced RSA muons. A summary of all significant sources of systematic uncertainty affecting the signal efficiency is presented in Table 1. In addition, we investigate the effect of pileup on the cosmic muon and tracker track rejection cuts, and also the effect of the poor RSA muon p_{T} resolution. These effects appear to be negligible and we do not assign any additional systematic uncertainty for them, as explained later. Additional systematics from the parton distribution function sets, the renormalisation and factorisation scales used in generating simulated events, and the effect of higher order QCD corrections are described in [6].

Table 1: Systematic uncertainties related to the signal selection. The uncertainty specified is a relative uncertainty on the signal efficiency. The relative uncertainty in the luminosity is 2.6%.

Source	Uncertainty
Pileup modelling	2%
Tracking efficiency from cosmics	3%
Trigger efficiency	15%
Parton distribution functions	$< 1\%$
Renormalisation and factorisation scales	$< 0.5\%$
NLO effects	$5 - 7\%$

6.1 Luminosity

For the running period corresponding to this analysis, CMS estimates the relative uncertainty on the luminosity to be 2.6% [19]. This uncertainty is used when calculating the final cross section estimates.

6.2 Effect of pileup

In order for the simulation to describe the pileup events in data realistically, the simulated background events are reweighted to match the pileup in data. The systematic uncertainty on the pileup modelling is estimated varying the average number of reconstructed primary vertices in the background MC events by $\pm 5\%$. The systematic uncertainty on the signal efficiency due to the pileup modelling is found to be less than 2% relative for all signal MC samples used.

The more pp collisions that occur during each LHC bunch crossing, the more tracks would be reconstructed in the silicon tracker. This would increase the probability that we reject LL particle candidates due to accidental matching of RSA muons to unrelated tracker tracks. Therefore, the possible dependence of the track rejection and cosmic rejection cuts on pileup is investigated on simulated $\mathrm{H}^0 \to \mathrm{XX}$ events with $M_{\mathrm{H}^0} = 1000\,\mathrm{GeV}/c^2$, $M_{\mathrm{X}} = 350\,\mathrm{GeV}/c^2$ and $c\tau = 350\,\mathrm{cm}$. We find no significant dependence after the full selection. To increase the available statistics, we also study the pileup dependence of the two cuts when the $|d_0|/\sigma_d$ and $L_{xy}/\sigma_{L_{xy}}$ cuts are removed from the full selection and the minimum valid muon station requirement is loosened from 3 to 2. Again, we do not observe any significant dependence. Therefore, we do not assign an additional systematic uncertainty.

6.3 Track finding and selection efficiency

To assess if the efficiency to reconstruct displaced muons in the muon chambers is correctly modeled by the simulation, a direct measurement is performed using 'tag and probe' tech-

niques [20] applied to cosmic ray muons. Events are selected from dedicated runs with no beam activity and the cosmic ray muons are reconstructed as two separate RSA muons in opposite halves of the CMS detector. To select a consistent sample between data and simulation and to ensure that the bottom half of the cosmic muon has a timing consistent with those coming form a pp collision event, we require that the top half of the cosmic muon has a time delay from a pp collision event as measured with respect to the interaction point in the range [-40, -20]ns.

We perform two measurements: the first one uses a 'tag' that requires the presence of a track in the silicon tracker, while the second uses a 'tag' that requires an RSA muon in the muon chambers in the upper half of CMS. In both cases, we evaluate the fraction of these selected events in which an RSA muon (the 'probe') is then reconstructed in the lower half of CMS, which also passes the analysis muon selection requirements. The two measurements provide estimates of the efficiency of an RSA muon to be reconstructed and pass these selection requirements.

The results are shown in Figure 4 as a function of the transverse impact parameter of the muon. The results of these two measurements do not yield identical efficiency values as they are obtained over samples with different timing distributions. We can select the timing for the muons in the top half of CMS, while for the silicon tracker tracks the timing is constrained by the charge integration time of the tracker, for which no direct measurement of the track arrival time is available. The aim of these measurements is to provide a comparison between data and simulation, so the conditions need only to be consistent within a given method.

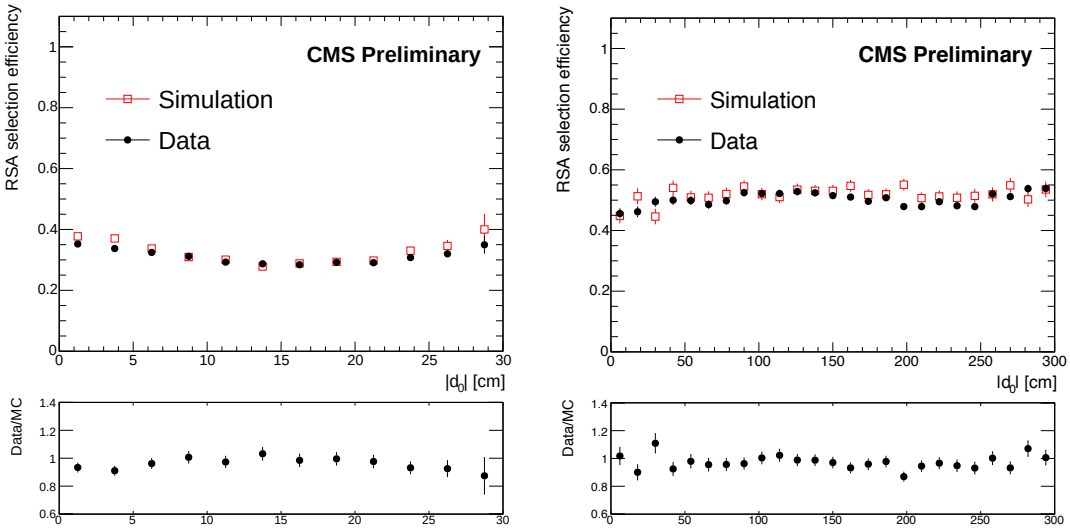

Figure 4: RSA muon reconstruction and selection efficiency measured using cosmic muons in data and simulation by requiring the presence of a reconstructed track in the silicon tracker (left) and using only the muon chambers (right) as a function of $|d_0|$.

To estimate the systematic uncertainty associated to the simulation of the track reconstruction and selection efficiency for the dimuon candidates, we evaluate the weighted average of the relative difference between data and simulation using the $|d_0|$ distributions of each muon in simulated signal samples. We find that a 3% systematic uncertainty per candidate covers the variations for all considered signals and this is taken as the systematic uncertainty.

6.4 Trigger efficiency measurement

The efficiency for a muon to pass the trigger selection is measured with the tag and probe method utilizing events from Z boson candidates decaying to muon pairs. The regions where the trigger efficiencies of the two muons are correlated are excluded from the analysis, namely when they are close to each other ($\Delta R < 0.2$) or when they are close to back-to-back ($\cos(\alpha) < -0.75$). The discrepancy between data and simulation is no larger than 10%, which is taken as the systematic uncertainty on the trigger efficiency measured for Z decays.

To estimate a systematic uncertainty on the dependence of the trigger efficiency on the vertex displacement we utilize cosmic muons data and simulation. We use the same approach as described in the previous section for measuring the tracking and selection efficiency. However, we perform the measurement on muons reconstructed offline rather than muons reconstructed by the trigger (L2Muons) to minimize the dependence of the results on the timing of cosmic muons. The muon reconstruction algorithms used in the trigger are similar to those used offline, but the final refit of the RSA muon algorithm is not applied to L2Muons. We therefore utilize offline RSA muons without the final refit for this measurement. We find that a systematic uncertainty of 15% per candidate would cover all considered signals. This method tests the dependence of the trigger efficiency on the decay length in data and simulation and it applies to both prompt and displaced muons. Therefore, this systematic is expected to cover also the one derived from the tag and probe method on prompt L2Muons (10%), and we take this (15%) as the overall systematic uncertainty on the trigger efficiency.

6.5 Effect of the p_{T} resolution

To study the effect of the modest RSA muon p_{T} resolution on this analysis we scale up and down the reconstructed p_{T} of the RSA muons in simulated signal events by 10%. In both cases we find no significant change in the signal efficiency for all signal samples and we conclude that no additional systematic uncertainty needs to be assigned.

7 Results

No events are observed in the signal region after the full selection. The upper limits on the two signal models for various particle masses and lifetime values are set with 95% confidence level (CL). The computation is performed using the Bayesian method described in Ref. [21]. The limits are derived by comparing the number of events N_S expected in the signal region with the number of events that the signal plus background hypothesis predicts.

The systematic uncertainties on the signal selection efficiency given in Section 6 are introduced in the limit calculation as nuisance parameters with log-normal prior distributions. The expected number of background events in the signal region, μ_B, is taken as an additional nuisance parameter, which depends on the number of observed events in the control region, N_C. The probability distribution of μ_B, $p(\mu_B)$, is given by $p(\mu_B) = \frac{\mu_B^{N_C}}{N_C!}\exp(-\mu_B)$, as can be shown using Bayesian method assuming a flat prior in μ_B [21]. An additional multiplicative nuisance parameter allows for the expected background in the signal region to differ from that in the control region within the systematic uncertainty on the background estimated earlier (67%). This is implemented as a multiplicative log-normal term with $\sigma = 1.67$.

The expected number of signal events, μ_S, takes the following form:

$$\begin{aligned}\mu_S &= \mathcal{L}\sigma\left[2\mathrm{B}(1-\mathrm{B})\epsilon_1+\epsilon_2\mathrm{B}^2\right](1-f)\\ &= 2\mathcal{L}\sigma\mathrm{B}\epsilon_1\left[1-\mathrm{B}\left(1-\frac{\epsilon_2}{2\epsilon_1}\right)\right](1-f)\end{aligned} \quad (1)$$

where $\mathcal{L}$ is the integrated luminosity, $\epsilon_{(1,2)}$ are the signal efficiencies defined in Section 4, σ is the production cross section of $\mathrm{H}^0 \to \mathrm{XX}$ (or $\tilde{\mathrm{q}}\bar{\tilde{\mathrm{q}}}+\tilde{\mathrm{q}}\tilde{\mathrm{q}}$) and B is the branching fraction for the decay $\mathrm{X} \to \mu^+\mu^-$ (or $\tilde{\mathrm{q}} \to q\tilde{\chi}^0$, $\tilde{\chi}^0 \to \mu^+\mu^-\nu$). The parameter f represents the ratio of the number of signal events falling into the control region as fake background to the number of signal events in the signal region. Although the effect is negligible for all signal samples used, the conservative value, 0.5% (8.1%) for $\mathrm{H}^0 \to \mathrm{XX}$ ($\tilde{\mathrm{q}}\bar{\tilde{\mathrm{q}}}+\tilde{\mathrm{q}}\tilde{\mathrm{q}}$), is set for this parameter.

In signal events where two LL particles each decay to a final state containing dimuons, if the efficiency to reconstruct one dimuon pair is independent of the efficiency to reconstruct the second one in the same event, then $\epsilon_2 = 1-(1-\epsilon_1)^2$. However, because the dimuon trigger can sometimes fire by using one muon from each LL particle, one actually finds that $\epsilon_2 \geq 1-(1-\epsilon_1)^2$. It is conservative to calculate the limits assuming that $\epsilon_2 = 1-(1-\epsilon_1)^2$ in equation (1), since this minimises the predicted number of signal events.

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

In equation (2), the upper bounds on σB depend on the branching fraction. They scale approximately with the expression $[1-\frac{1}{2}\mathrm{B}\epsilon_1]$ and are thus best for low values of B, though the dependence of the limits on B is weak, especially if ϵ_1 is small. We compute limits setting B in the square brackets expression to unity so as to obtain conservative limits that are valid for any value of B.

The 95% CL upper limits are calculated for all mass points of $\mathrm{H}^0 \to \mathrm{XX}$ signal samples as a function of X boson lifetime. The expected limits are illustrated in the plots in Figure 5. The analysis is least sensitive to the $M_{\mathrm{H}^0} = 125\,\mathrm{GeV}/c^2$ case due to the low signal selection efficiencies in particular when $M_{\mathrm{X}} = 20\,\mathrm{GeV}/c^2$. The corresponding limits on $\sigma(\tilde{\mathrm{q}}\bar{\tilde{\mathrm{q}}}+\tilde{\mathrm{q}}\tilde{\mathrm{q}})\mathrm{B}(\tilde{\mathrm{q}} \to q\tilde{\chi}^0, \tilde{\chi}^0 \to \mu^+\mu^-\nu)$ are shown in Fig. 6. The most stringent limits are set for lifetimes between 10 cm and 200 cm, as expected, since the analysis has a negligible sensitivity for transverse decay lengths less than 40 cm and since for longer lifetimes a high fraction of the LL particles decay outside the detector. The shaded bands in these limit plots represent the $\pm 1\sigma$ range of variation of the expected 95% CL limits, shown for one choice of masses. All observed limits are consistent with the corresponding expected ones.

We also compute limits on the product of cross section, branching ratio and acceptance, thus accounting for the effects of the acceptance A, as defined in Section 4. Figure 7 shows the resulting limits on $\sigma(\mathrm{H}^0 \to \mathrm{XX})\mathrm{B}(\mathrm{X} \to \mu^+\mu^-)A(\mathrm{X} \to \mu^+\mu^-)$ and Fig. 8 shows the corresponding limits on $\sigma(\tilde{\mathrm{q}}\bar{\tilde{\mathrm{q}}}+\tilde{\mathrm{q}}\tilde{\mathrm{q}})\mathrm{B}(\tilde{\mathrm{q}} \to q\tilde{\chi}^0, \tilde{\chi}^0 \to \mu^+\mu^-\nu)A(\tilde{\mathrm{q}} \to q\tilde{\chi}^0, \tilde{\chi}^0 \to \mu^+\mu^-\nu)$. These limits restricted to the acceptance region are less model dependent, as can be seen by the fact that the limits on $\sigma\mathrm{B}A$ are similar for $\mathrm{X} \to \mu^+\mu^-$ and $\tilde{\chi}^0 \to \mu^+\mu^-\nu$. They also show substantially less dependence on the Higgs boson and X boson masses and on the mean proper decay length $c\tau$ of the X boson. The residual dependence of the limits on $c\tau$ is due to the $|d_0|/\sigma_d > 4$ and $L_{xy}/\sigma_{L_{xy}} > 12$ requirements at small values of $c\tau$; whereas at larger values of $c\tau$, it is caused by the fact that the track finding efficiency falls for leptons produced far from the beamline with very large impact parameters.

The limits described above are determined in the context of two specific models. However, the analysis is sensitive to any process in which a LL particle is produced and subsequently decays to a final state that includes dimuons. The limits within the acceptance region (i.e. on $\sigma \mathrm{B}A$), which are less model dependent, can be utilized to place approximate limits on this more general class of models. The limits on $\sigma \mathrm{B}A$ shown in Fig. 7 and Fig. 8 should remain approximately valid in most signal models in which each event contains two identical LL particles that decay in this way. (The variation amongst the limit curves shown in these plots for different signal models and particle masses gives an indication of the accuracy of this statement.) Exceptions could arise for models that give poor efficiency within the acceptance criteria, e.g. for models in which the leptons are not isolated; have impact parameters with significance below $|d_0|/\sigma_d < 4$, corresponding to $|d_0| \lesssim 8$ cm; are almost collinear with each other (with the dilepton mass below $15\,\mathrm{GeV}/c^2$, or $\Delta R < 0.2$); or do not usually satisfy the $|\Delta\Phi| < \pi/2$ criterion, such that the parameter f becomes large (e.g. if the LL particle is slow-moving and decays to many particles).

In models where only one LL particle that can decay inclusively to dimuons is produced in each event, the expected number of selected signal events for given $\sigma \mathrm{B}$ will be up to a factor of two lower, and the limits on $\sigma \mathrm{B}A$ will be up to a factor of two worse than shown in Fig. 7 and Fig. 8.

7.1 Combined Limits

The analysis described in this document and the one based on muons reconstructed in the silicon tracker in [6] are orthogonal in the sense that, by design, there is no overlap in the events selected by the two analyses. The limits from the two analyses can therefore be combined by treating their results as independent measurements of the same physics signal.

The principal systematic uncertainties of the two channels are uncorrelated except for those discussed here. The systematic uncertainty on the integrated luminosity is fully correlated. That on the trigger efficiency is partially correlated, so we take it to be fully correlated, which is the most conservative assumption. The systematic uncertainty in the tracker tracking efficiency is also partially anticorrelated because of the tracker track rejection cut applied in the muon chambers-based analysis. It is therefore conservative to assume no correlation in it.

The results of the combination are shown in Fig. 9 for the $\mathrm{H}^0 \to \mathrm{XX}$ and in Fig. 10 for the $\widetilde{\mathrm{q}} \to q\widetilde{\chi}^0$, $\widetilde{\chi}^0 \to \mu^+\mu^-\nu$. The same combination is performed within the acceptance region as defined in Section 4 and the results are presented in Fig. 11 for the $\mathrm{H}^0 \to \mathrm{XX}$ and in Fig. 12 for the $\widetilde{\mathrm{q}} \to q\widetilde{\chi}^0$, $\widetilde{\chi}^0 \to \mu^+\mu^-\nu$. The limits coincide with the ones from the tracker-based analysis for lower lifetime values, where the tracker-based analysis has much higher efficiency. For higher lifetime values the limits are improved by a factor of up to two by the combination.

8 Summary

A search for long-lived particles decaying to final states that include dimuons, using muons reconstructed only in the muon chambers, has been performed on pp collision data taken by the CMS detector at $\sqrt{s} = 8\,\mathrm{TeV}$ in 2012. It extends the scope of a similar search for displaced dimuon signatures based on the silicon tracker. No events pass the selection criteria of the analysis. Upper limits are computed for two specific models. The first model predicts a Higgs boson with mass in the range 125–1000 GeV/c^2, decaying to pairs of long-lived neutral particles, with masses in the range 20–350 GeV/c^2, that can decay to dimuon pairs. The limits are typically in the range 1–50 fb, and can weaken to a few pb for the lowest masses and longest lifetimes,

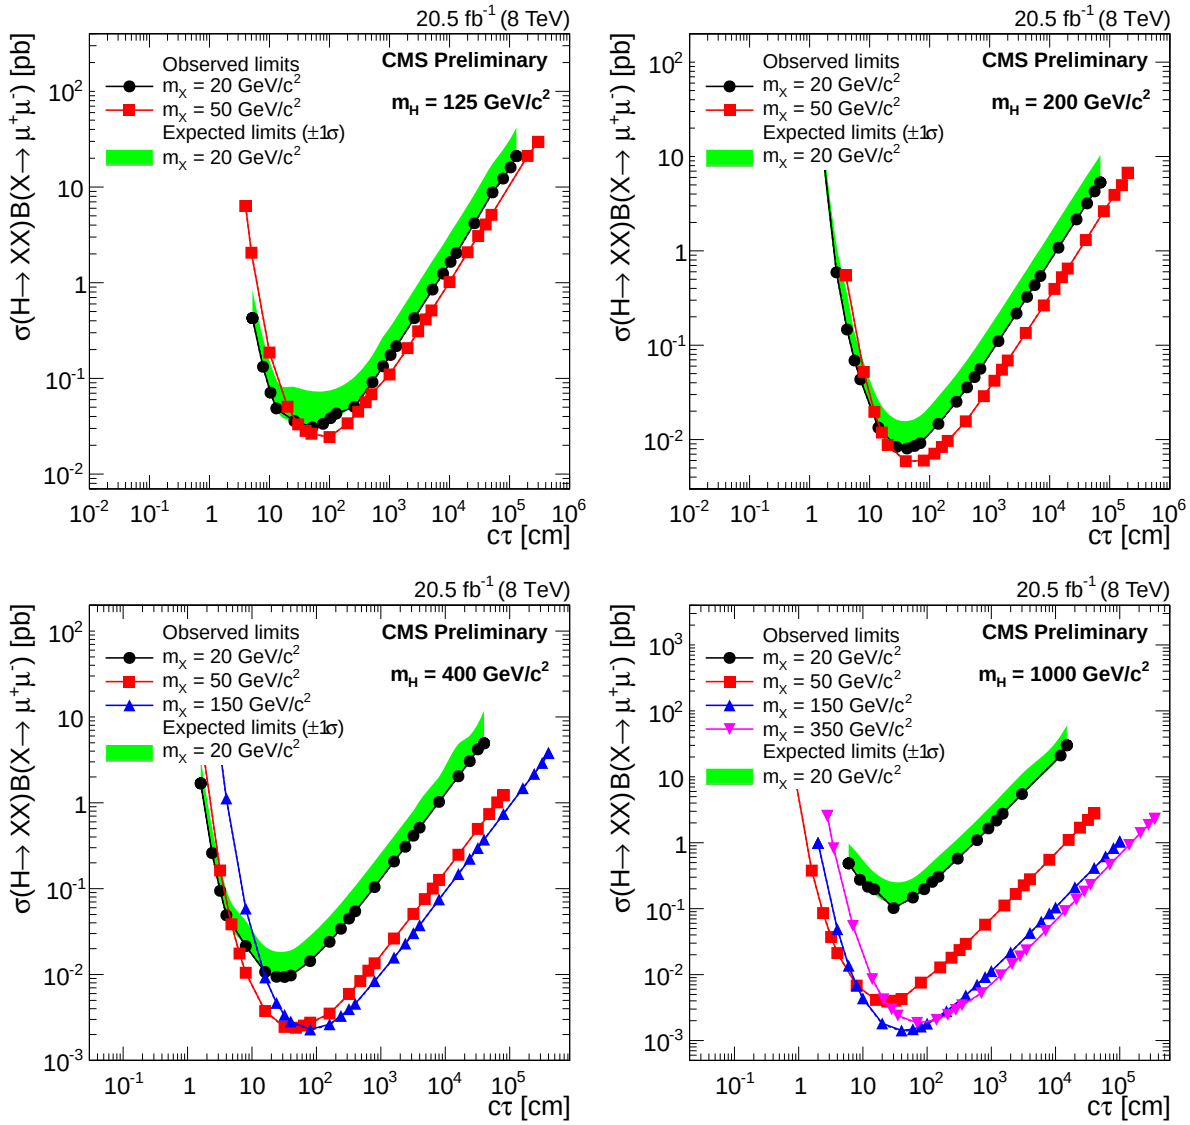

Figure 5: 95% CL upper limits on $\sigma(\mathrm{H}^0 \to \mathrm{XX})\mathrm{B}(\mathrm{X} \to \mu^+\mu^-)$ for $M_{H^0} = 125\,\mathrm{GeV}/c^2$, $200\,\mathrm{GeV}/c^2$, $400\,\mathrm{GeV}/c^2$ and $1000\,\mathrm{GeV}/c^2$ with various X mass points. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a $20\,\mathrm{GeV}/c^2$ X boson mass. Corresponding bands for the other X boson masses show a similar level of agreement and are omitted for clarity of presentation.

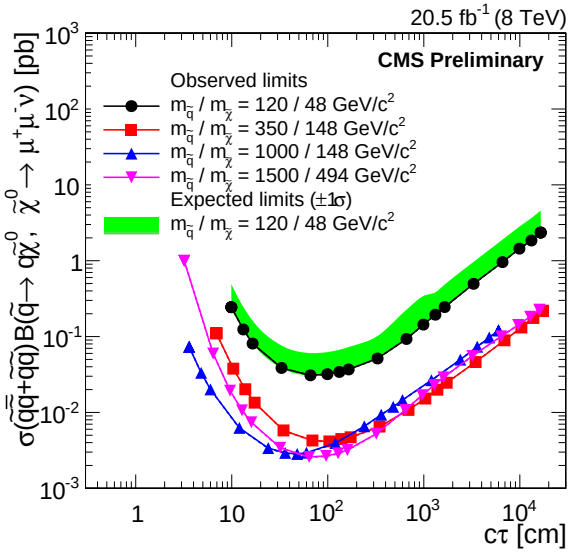

Figure 6: The 95% CL upper limits on $\sigma(\tilde{q}\bar{\tilde{q}} + \tilde{q}\tilde{q})B(\tilde{q} \to q\tilde{\chi}^0,\ \tilde{\chi}^0 \to \mu^+\mu^-\nu)$ as a function of the neutralino lifetime. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 120 GeV/c^2 squark and a 48 GeV/c^2 neutralino mass. Corresponding bands for the other squark and neutralino masses show a similar level of agreement and are omitted for clarity of presentation.

and are given for lifetimes in the range $1 < c\tau < 10000$ cm. The second model produces squark pairs which each decay to long-lived neutralinos. The neutralino can decay to a pair of muons and a neutrino. No events are observed in the signal region after the full selection. The limits are combined with those from a previous search [6] for the same signal models using leptons reconstructed in the silicon tracker. In the muon channel, for particles with very long lifetimes, the combined limits improve on the limits obtained from the tracker-based search, and provide the most stringent limits obtained so far.

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: BMWF 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

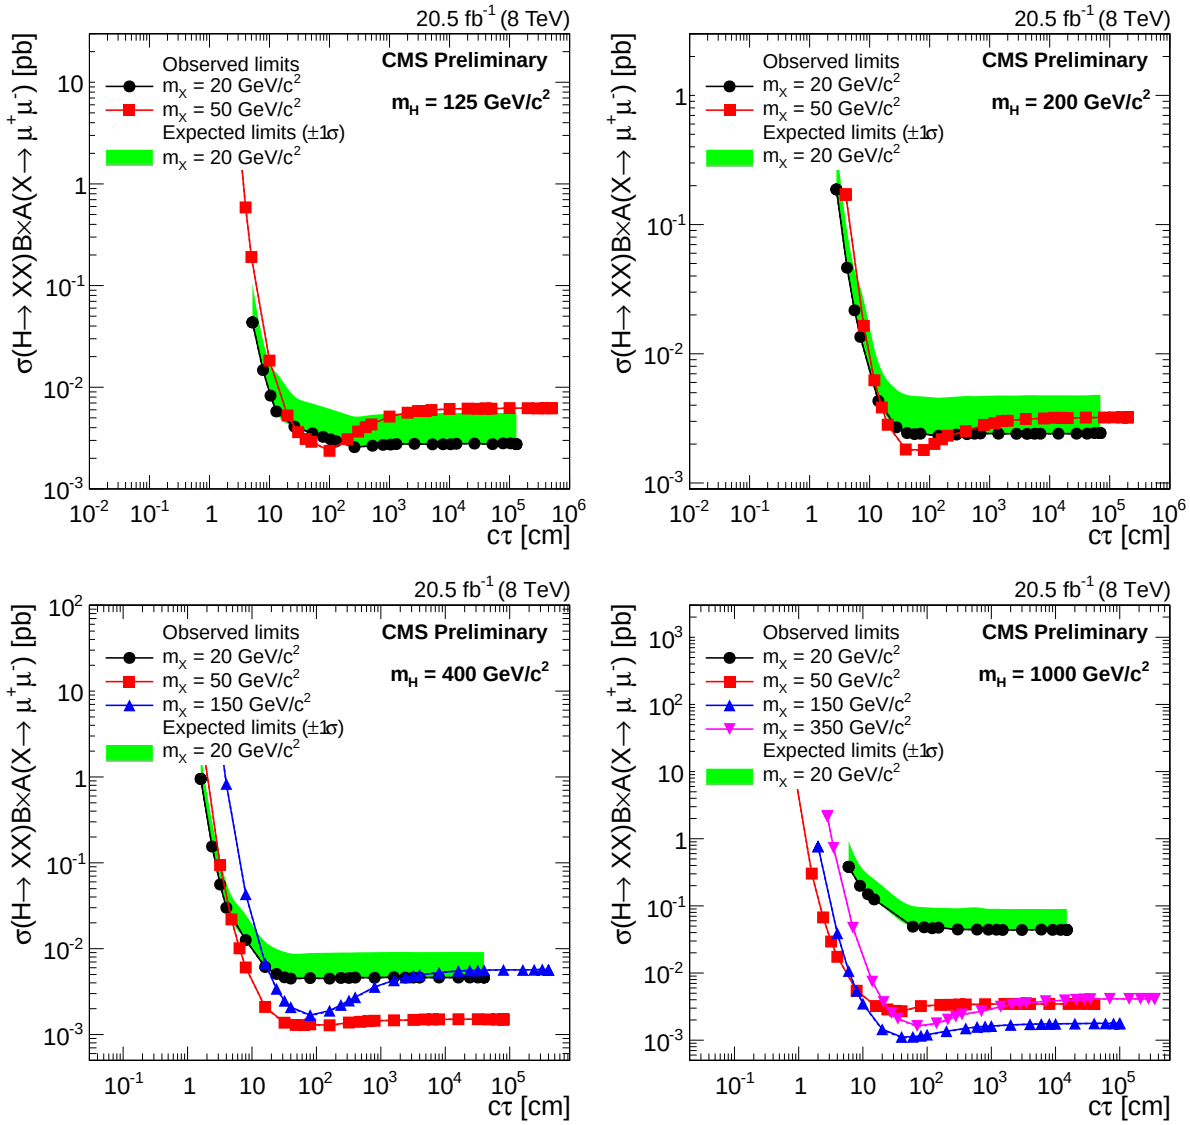

Figure 7: 95% CL upper limits on $\sigma(\mathrm{H}^0 \to \mathrm{XX})\mathrm{B}(\mathrm{X} \to \mu^+\mu^-)A(\mathrm{X} \to \mu^+\mu^-)$ for M_{H^0} = 125 GeV/c^2, 200 GeV/c^2, 400 GeV/c^2 and 1000 GeV/c^2 with various X mass points. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 20 GeV/c^2 X boson mass. Corresponding bands for the other X boson masses show a similar level of agreement and are omitted for clarity of presentation.

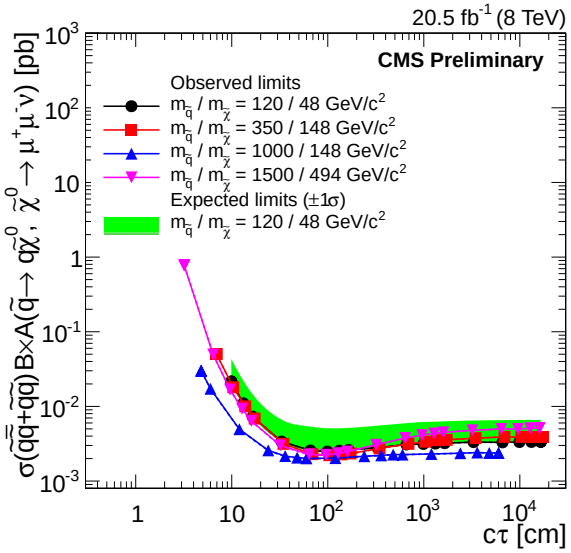

Figure 8: The 95% CL upper limits on $\sigma(\tilde{q}\bar{\tilde{q}} + \tilde{q}\tilde{q})B(\tilde{q} \to q\tilde{\chi}^0,\ \tilde{\chi}^0 \to \mu^+\mu^-\nu)A(\tilde{q} \to q\tilde{\chi}^0,\ \tilde{\chi}^0 \to \mu^+\mu^-\nu)$ as a function of the neutralino lifetime. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 120 GeV/c^2 squark and a 48 GeV/c^2 neutralino mass. Corresponding bands for the other squark and neutralino masses show a similar level of agreement and are omitted for clarity of presentation.

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 for long-lived particles that decay into final states that include two electrons or two muons in proton-proton collisions at $\sqrt{s} = 8$ TeV", arXiv:1411.6977.

[7] CMS Collaboration, "Search for "Displaced Supersymmetry" in events with an electron and a muon with large impact parameters", (2014). arXiv:1409.4789. Submitted to PRL.

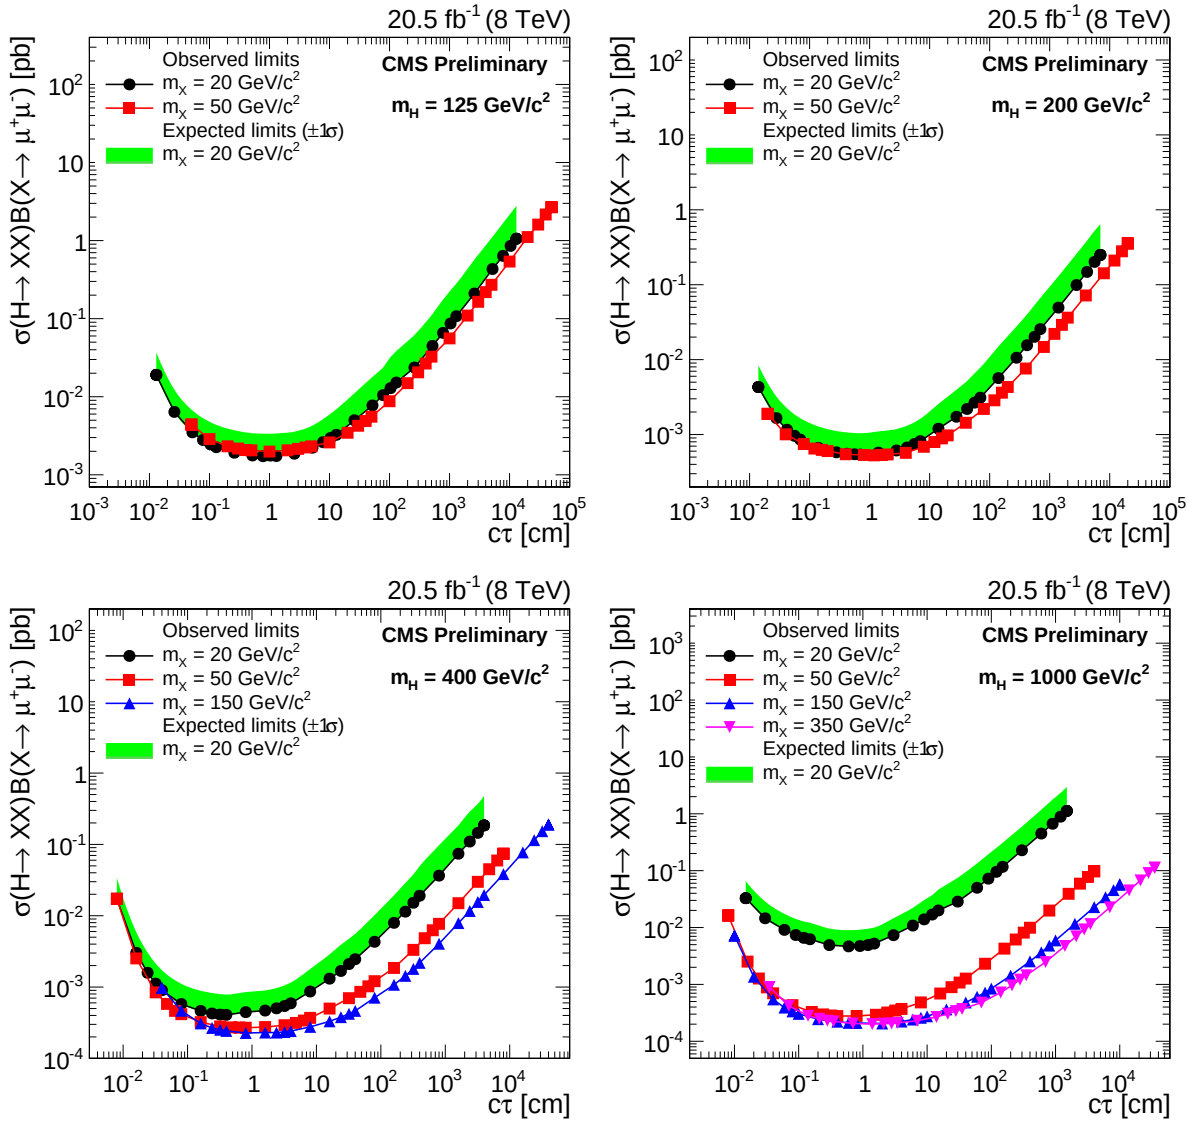

Figure 9: Combined 95% CL upper limits on $\sigma(\mathrm{H}^0 \to \mathrm{XX})\mathrm{B}(\mathrm{X} \to \mu^+\mu^-)$ for all H^0 mass values of 125 GeV/c^2, 200 GeV/c^2, 400 GeV/c^2 and 1000 GeV/c^2, with various X mass points. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 20 GeV/c^2 X boson mass. Corresponding bands for the other X boson masses show a similar level of agreement and are omitted for clarity of presentation.

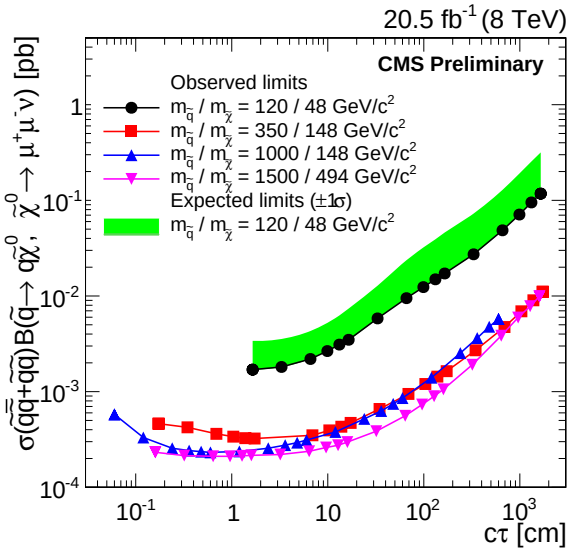

Figure 10: Combined 95% CL upper limits on $\sigma(\tilde{q}\bar{\tilde{q}} + \tilde{q}\tilde{q})B(\tilde{q} \to q\tilde{\chi}^0,\ \tilde{\chi}^0 \to \mu^+\mu^-\nu)$ as a function of the neutralino lifetime. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 120 GeV/c^2 squark and a 48 GeV/c^2 neutralino mass. Corresponding bands for the other squark and neutralino masses show a similar level of agreement and are omitted for clarity of presentation.

[8] CMS Collaboration, "Search for long-lived neutral particles decaying to dijets", CMS Physics Analysis Summary CMS-PAS-EXO-12-038 (in preparation for publication), 2012.

[9] 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, doi:10.1103/PhysRevLett.97.161802, arXiv:hep-ex/0607028.

[10] 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.

[11] 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, doi:10.1103/PhysRevLett.108.251801, arXiv:1203.1303.

[12] 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, doi:10.1016/j.physletb.2013.01.042, arXiv:1210.7451.

[13] ATLAS Collaboration, "Search for long-lived neutral particles decaying into lepton jets in proton–proton collisions at $\sqrt{s}$ = 8 TeV with the ATLAS detector", arXiv:1409.0746.

[14] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.

[15] CMS Collaboration, "Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV", *Journal of Instrumentation* **7** (2012), no. 10, P10002.

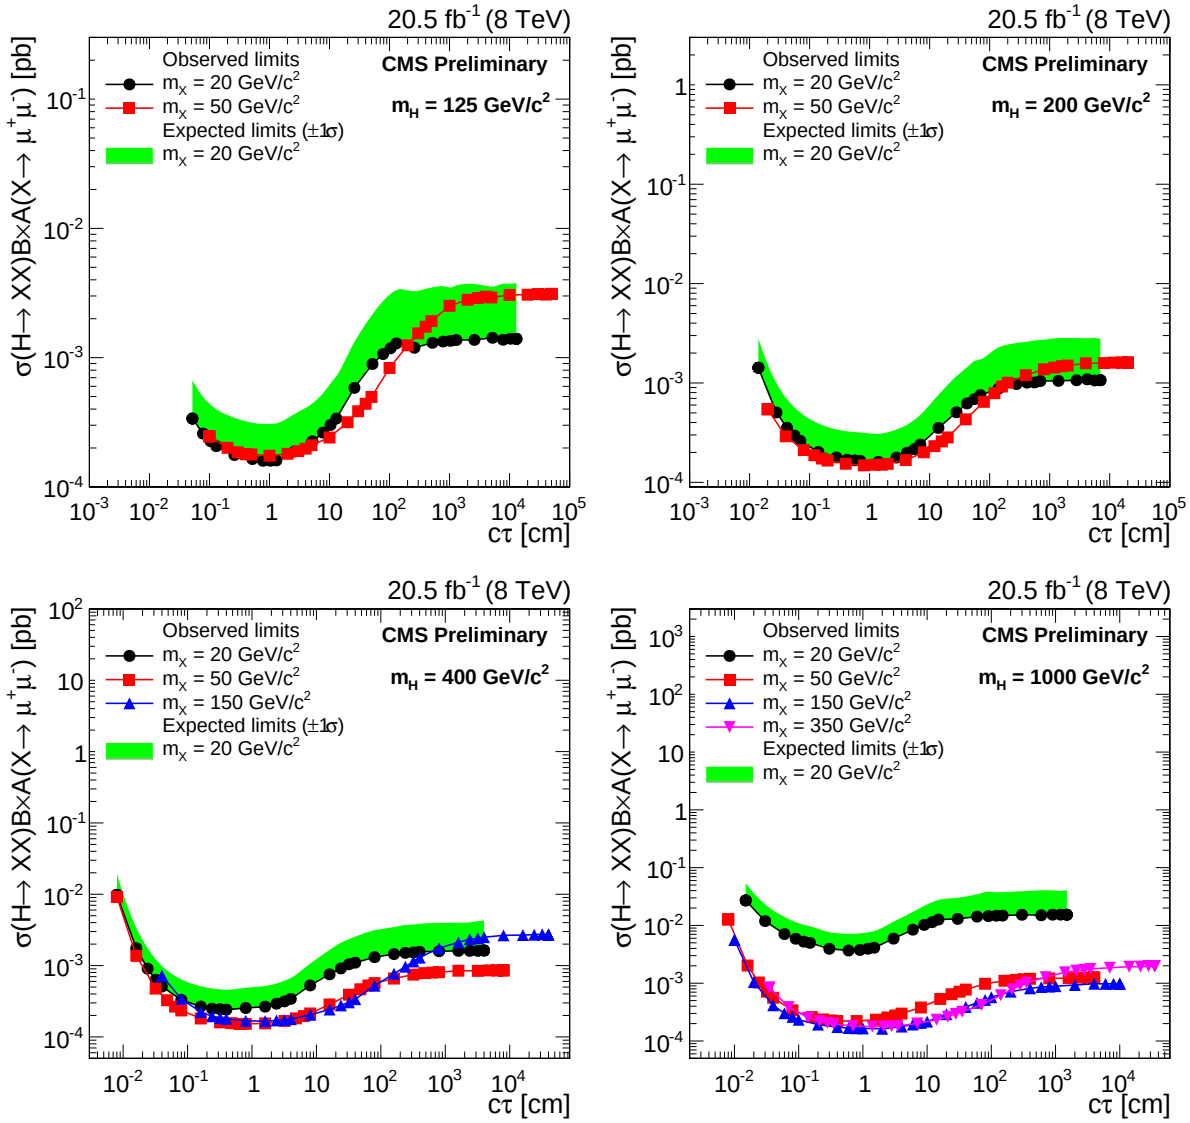

Figure 11: Combined 95% CL upper limits on $\sigma(\mathrm{H}^0 \to \mathrm{XX})\mathrm{B}(\mathrm{X} \to \mu^+\mu^-)A(\mathrm{X} \to \mu^+\mu^-)$ for all H^0 mass values of 125 GeV/c^2, 200 GeV/c^2, 400 GeV/c^2 and 1000 GeV/c^2, with various X mass points. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 20 GeV/c^2 X boson mass. Corresponding bands for the other X boson masses show a similar level of agreement and are omitted for clarity of presentation.

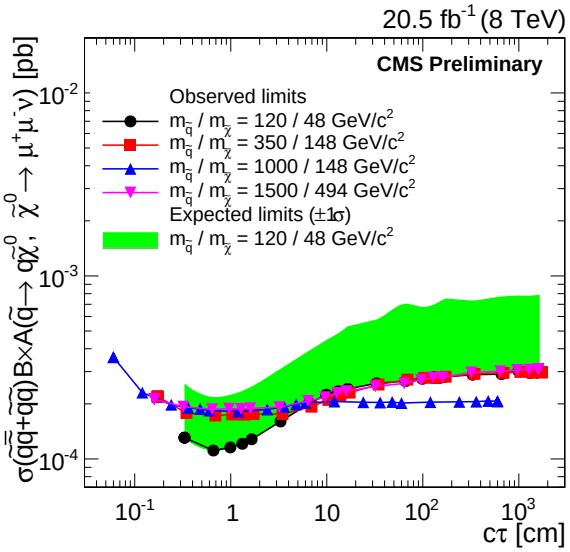

Figure 12: Combined 95% CL upper limits on $\sigma(\tilde{q}\bar{\tilde{q}} + \tilde{q}\tilde{q})B(\tilde{q} \to q\tilde{\chi}^0,\ \tilde{\chi}^0 \to \mu^+\mu^-\nu)A(\tilde{q} \to q\tilde{\chi}^0,\ \tilde{\chi}^0 \to \mu^+\mu^-\nu)$ as a function of the neutralino lifetime. Shaded bands show the $\pm 1\sigma$ range of variation of the expected 95% CL limits for the case of a 120 GeV/c^2 squark and a 48 GeV/c^2 neutralino mass. Corresponding bands for the other squark and neutralino masses show a similar level of agreement and are omitted for clarity of presentation.

[16] S. M. T. Sjöstrand and P. Z. Skands, "PYTHIA 6.4 Physics and Manual", *JHEP* **05** (2006) 576, doi:10.1088/1126-6708/2006/05/026, arXiv:hep-ph/0603175.

[17] GEANT4 Collaboration, "GEANT4: A simulation toolkit", *Nucl. Instrum. Meth. A* **506** (2003) 250, doi:10.1016/S0168-9002(03)01368-8.

[18] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker", *JINST* **9** (2014) P10009, doi:DOI:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569.

[19] CMS Collaboration, "CMS luminosity based on pixel cluster counting - Summer 2013 update", CMS Physics Analysis Summary CMS-PAS-LUM-13-001, 2013.

[20] CMS Collaboration, "Measurement of the inclusive W and Z production cross sections in pp collisions at $\sqrt{s} = 7$ TeV with the CMS experiment", *Journal of High Energy Physics* **2011** (2011), no. 10, doi:10.1007/JHEP10(2011)132.

[21] 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.