CMS Physics Analysis Summary

Contact: cms-pag-conveners-higgs@cern.ch 2012/03/08

Search for a light pseudoscalar Higgs boson in the dimuon channel

The CMS Collaboration

Abstract

The dimuon invariant mass spectrum is investigated between 5.5 GeV and 14 GeV using 1.3 fb^{-1} of pp collision data at $\sqrt{s}$ = 7 TeV collected by the CMS detector at the LHC to search for a light pseudoscalar Higgs boson. Such a light Higgs boson is predicted in many models, including the Next-to-Minimal Supersymmetric Standard Model (NMSSM) which avoids the μ-term problem of the MSSM. A limit on the cross section times branching ratio $\sigma(pp \to a) \times BR(a \to \mu^+\mu^-)$ is set.

1 Introduction

Low energy supersymmetry (SUSY) is an elegant solution to the hierarchy problem that arises in the Standard Model (SM), while providing a candidate for dark matter and allowing in addition the unification of gauge couplings at the GUT scale. The simplest SUSY model is the Minimal Supersymmetric Model (MSSM). However, the ad hoc nature of the μ term in the MSSM superpotential and the very large stop masses needed if the LHC hints of a SM-like Higgs boson with mass near 125 GeV survive [1, 2] argue strongly for an extension. By far the simplest is the Next-to MSSM (NMSSM) (see [3] for a review) which extends the MSSM by introducing a complex singlet superfield. The latter has a scalar component field and associated superpotential term that automatically generates an effective μ term of order the electroweak scale as a result of spontaneous symmetry breaking. In addition, scalar interactions associated with the new superpotential term can easily raise the mass of the light Higgs boson into the mass range above 125 GeV. The added scalar field expands the Higgs sector to three CP-even scalars (h_1, h_2, h_3), two CP-odd scalars (a_1, a_2) and two charged scalars (H^+, H^-). The NMSSM has two very natural symmetries that emerge when one of two pairs of parameters are set to zero. If either is imposed (e.g. at the GUT scale), it is quite natural for the a_1 to be very light. Indeed, $m_{a_1} < 2m_B$, where m_B is the B meson mass, is a very natural possibility. The light a_1 is a superposition of the MSSM CP-odd doublet scalar and the additional CP-odd singlet scalar of the NMSSM: $a_1 = \cos\theta_A a_{MSSM} + \sin\theta_A a_S$. Small mixing, $|\cos\theta_A| \ll 1$, corresponding to the a_1 being mainly singlet, occurs when one of the above $U(1)$ symmetries is imposed at the GUT scale. However, the coupling of the a_1 to $\mu^+\mu^-$, $\tau^+\tau^-$ and $b\bar{b}$ is proportional to $C_{a_1 b\bar{b}} = \tan\beta\cos\theta_A$ and is therefore appreciable for large values of $\tan\beta$, even if $\cos\theta_A$ is small. In the following, we refer to the light pseudoscalar as a except when specifically probing the NMSSM parameter space (a_1).

More generally, superstring modeling suggests the possibility of many light a's, at least some of which couple to $\mu^+\mu^-$, $\tau^+\tau^-$ and $b\bar{b}$. In the SUSY context, detection of a light CP-odd boson a would be compelling evidence for a model with more than two Higgs doublets and the a would be most naturally accommodated in the NMSSM or an extended NMSSM structure. A light a could also alter Higgs phenomenology significantly by virtue of $h \to aa$, with $a \to 2\tau, 2g, 2c, 2s$ decays being dominant [4].

Searches for a light a are mainly sensitive to $C_{ab\bar{b}}$. For $m_a < M_{\Upsilon_{1S}}$ the strongest constraints on $C_{ab\bar{b}}$ are those from BaBar in the $\Upsilon_{3S} \to \gamma a$ channel [5, 6]. However, for m_a above $M_{\Upsilon_{1S}}$ only the Tevatron and LHC have sensitivity. The a can be produced at the Tevatron and LHC via $gg \to a$, where the coupling C_{agg} derives from quark (especially bottom and top) triangle loops, as shown in Figure 1. This process, plus higher corrections thereto, leads to a large cross section due to the large gg "luminosity" at small gluon momentum fractions, provided the $C_{aq\bar{q}}$ ($q = t, b$ in particular) couplings deriving from the doublet component of the a are not too suppressed. This large cross section will typically lead to a significant number of $gg \to a \to \mu^+\mu^-$ events even though $BR(a \to \mu^+\mu^-)$ is small.

The LEP experiments considered various Higgs scenarios in the MSSM [7] and placed limits on certain Higgs-like event topologies that can be relevant for the NMSSM Higgs searches. Also, there are constraints on $gg \to a \to \mu^+\mu^-$ from the CDF collaboration [8, 9]. While these can be reinterpreted in terms of limits on $C_{ab\bar{b}}$ [10], the analysis was only performed in the mass range 6.3 GeV $\leq M_{\mu\mu} \leq$ 9 GeV; additionally, the limits on $C_{ab\bar{b}}$ are weaker than those set by BaBar. In the NMSSM context, where $C_{a_1 b\bar{b}} = \tan\beta\cos\theta_A$, the weakness of the pre-LHC limits translates to rather modest limits on $|\cos\theta_A|$ that do not strongly constrain the NMSSM models of interest. As estimated in [10] and explicitly demonstrated here, CMS can definitely improve

the BaBar and CDF limits, the latter due to the higher production yield [$\sigma_{LHC}(pp \to a) \sim 4.5\ \sigma_{TeV}(p\bar{p} \to a)$] and by virtue of better performance in the muon reconstruction. Further, the CMS analysis can extend the limits into the $m_a > M_{\Upsilon_{3S}}$ mass range.

The branching ratio for $a \to \mu^+\mu^-$ depends on the mass and on $\tan\beta$ but not on $\cos\theta_A$ at tree level [10]. It is nearly constant for $m_a > 5\,\text{GeV}$ and ranges from 10^{-3} to 4×10^{-3} for $\tan\beta = 1$ to $\tan\beta = 20$. Figure 1 also shows the cross section for $gg \to a$ at 7 TeV for several values of $\tan\beta$ and $\cos\theta_A = 1$. While $BR(a \to \mu\mu)$ changes very little with increasing values of $\tan\beta$ once $\tan\beta > 2$, in contrast, the production $\sigma(gg \to a)$ increases rapidly with $\tan\beta$ due to the fact that $C_{ab\bar{b}} \propto \tan\beta$. However, because of subtleties related to top-quark loop contributions and higher-order corrections this increase is slower than $\tan^2\beta$. In the context of the NMSSM, all $q\bar{q}$ couplings of the a_1 are proportional to $\cos\theta_A$, implying that $\sigma(gg \to a_1) \propto \cos^2\theta_A$.

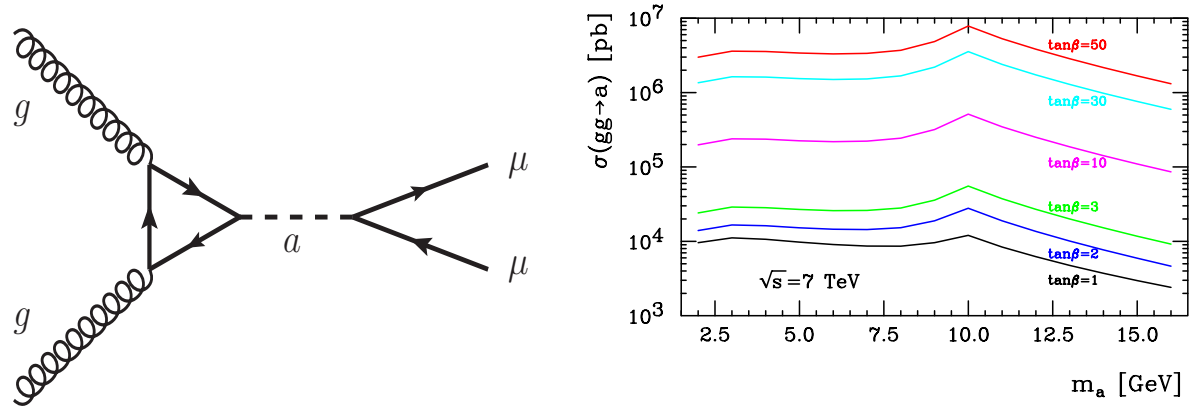

Figure 1: Left: Feynman diagram for a production and decay. Right: cross section for $gg \to a$ production at $\sqrt{s} = 7$ TeV for $\tan\beta$ = 1, 2, 3, 10, 30, 50. We plot the cross section including all available higher order corrections as per HIGLU [11], especially those from resolvable parton final state contributions ($gg \to ag$).

This paper presents the results of a search for a light a decaying into two oppositely charged muons in the invariant mass range around the Upsilon resonances in proton-proton collisions. The data were recorded at the LHC at $\sqrt{s}$ = 7 TeV, by the Compact Muon Solenoid (CMS) experiment in 2011, corresponding to an integrated luminosity of 1.3 fb^{-1}.

2 CMS Detector

The central feature of the CMS detector is a superconducting solenoid, of 6 m internal diameter, providing a field of 3.8 T. Within the field volume are the silicon pixel and strip tracker, the crystal electromagnetic calorimeter (ECAL) and the brass/scintillator hadron calorimeter (HCAL). Muons are measured in gas-ionization detectors embedded in the steel return yoke. In addition to the barrel and endcap detectors, CMS has extensive forward calorimetry. The muons are measured in the pseudorapidity range $|\eta| < 2.4$, ($\eta \equiv -\ln[\tan(\theta/2)]$, where θ is the polar angle of the trajectory of a particle with respect to the direction of the counterclockwise proton beam) with detection planes made using three technologies: Drift Tubes (DT) (for the range $|\eta| < 1.2$), Resistive Plate Chambers (RPC) (for $|\eta| < 1.6$) and Cathode Strip Chambers (CSC) (for $|\eta| < 2.4$). The DT and RPC are indicated as the central "barrel" while the CSC detector comprises the "endcaps". Matching the muons to the tracks measured in the silicon tracker results in a transverse momentum resolution between 1 and 5 %, for p_T values up to

1 TeV/c. The detector systems are aligned and calibrated using LHC collision data and cosmic-ray muons [12]. The first level (L1) of the CMS trigger system, composed of custom hardware processors, uses information from the calorimeters and muon detectors to select, in less than 1 μs, the most interesting events. The High Level Trigger (HLT) processor farm further decreases the event rate from around 100 kHz to around 300 Hz, before data storage. A more detailed description can be found in [13].

3 Analysis Strategy

We search for a narrow resonance a in the dimuon invariant mass distribution between 5.5 and 8.8 GeV (defined as "mass range 1") and between 11.5 and 14 GeV ("mass range 2"). We avoid the range between 9 and 11 GeV because the abundant contributions of the bottomonium resonances to the mass spectrum would make this search unfeasible. After the selection criteria are applied, we perform a mass scan dividing the spectrum into 110 (100) bins of 30 MeV each for mass range 1 (2). We fit each bin to determine a potential contribution from an a signal. Given the better mass resolution in the barrel part of the detector than in the endcaps, we also separate the two acceptance regions in the mass scan in order to improve the sensitivity.

4 Data and Monte Carlo samples

The data sample used for this study was recorded by the CMS detector in pp collisions at a center-of-mass energy of 7 TeV, between August and November 2011. The sample corresponds to a total integrated luminosity of 1.3 fb^{-1}. Data are included in the analysis for all runs where the LHC was providing stable beams, where the silicon tracker, the muon detectors, the calorimeters and the trigger were performing well and the luminosity measurement was available.

We analyze events collected with a muon trigger that requires the detection of two opposite sign muons with $p_{\mathrm{T}} > 3.5\,\mathrm{GeV}$ at the hardware level (L1), with additional cuts at the High Level Trigger (HLT). All three muon systems, DT, CSC and RPC, take part in the trigger decision. A good primary vertex is also required, as defined in [14]. The additional HLT requirements include $p_{\mathrm{T}}(\mu^+\mu^-) > 6\,\mathrm{GeV}$, $5.5 < m_{\mu^+\mu^-} < 14\,\mathrm{GeV}$, and a distance of closest approach to the beam axis, $DCA < 0.5$. To maintain the trigger rate at a reasonable level in the 2011 LHC run for low p_{T} dimuons with this selection, a prescale factor of 2 is imposed.

Our main sources of background arise from QCD processes and, in the lower invariant mass range, from a residual tail of the Y(1S) resonance. We derive our knowledge of the background shape in the invariant mass directly from data, and use simulated events as a crosscheck. The signal samples, QCD, and Y resonances are simulated with PYTHIA 6 with the D6T tune. As the NMSSM is not fully implemented in PYTHIA, we instead generate the MSSM pseudoscalar A boson in the mass range of 5.5 to 14 GeV. We then force the decay $a \to \mu^+\mu^-$. It is not necessary to simulate the cross section accurately for the purpose of this search, as we determine the upper limit normalized to unit cross section. We use $\tan\beta = 10$ in the MC generation. These samples also contain a valid simulation of the pile-up (overlapping pp interactions in the same bunch crossing) effects on the number of primary vertices. In the analysis, the events are reweighted to account for the differences between the pile-up simulation and the real pile-up in data.

5 Muon Reconstruction

There are two kinds of muon objects reconstructed by two of the different CMS algorithms [15] used in this analysis, global muons and tracker muons.

- Global muon reconstruction starts from a stand-alone muon track reconstructed in the muon system, and finds matching tracks in the silicon tracker. A global fit is performed using the hits of both tracks. The global track with the best normalized χ^2 is identified as a global muon.
- Tracker muons are tracks in the silicon tracker that are identified as muon tracks if the extrapolation to the muon system matches a muon segment.

6 Event and Dimuon Candidate Selection

In order to select the best dimuon candidate in each event, a set of quality criteria is applied to the tracks to reject misidentified muons and muons from kaon and pion decays. First, the muons are required to be in the geometrical acceptance ($|\eta| \leq 2.4$) and to be well in the plateau of the trigger efficiency, with $p_{\mathrm{T}} > 5.5$ GeV. The tracks are required to have at least 11 hits in the silicon tracker, at least one of which must be in the pixel detector, and a track fit $\chi^2/dof < 1.8$. This value is chosen to maximize the signal significance with respect to the QCD background, which is extracted directly from data.

Next, we impose isolation requirements to suppress misidentified (fake) leptons from jets and non-prompt muons from hadron decays. Muons are required to be isolated within a narrow cone of radius $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.3$ around the muon direction. This cone radius allows for excellent signal retention, given the two signal muons in each event. The muon relative isolation is defined as the sum of the transverse momentum (as measured in the silicon tracker) and transverse energy (as measured in the calorimeters) of all objects within this cone (excluding the muon itself), divided by the muon p_{T} as measured by the tracker. We require this quantity to be less than 0.2:

$$I_\mu = \frac{\sum_{R<0.3} p_{\mathrm{T}} + \sum_{R<0.3} E_{T_{ECAL}} + \sum_{R<0.3} E_{T_{HCAL}}}{p_{\mathrm{T}}} < 0.2. \quad (1)$$

This cut value is optimized by comparing the a signal with opposite-sign dimuons from data, and we verify that this value is appropriate for both the central (barrel) and forward (endcap) dimuon pairs. By requiring the muons to be isolated with this selection, we reject a large fraction of the background arising from QCD production of jets.

At this analysis stage, dimuon candidates are required to consist of two opposite-sign global muons with an invariant mass between 5.5 and 14 GeV. If such a candidate does not exist, possible dimuon candidates comprised of one global muon and one tracker muon are analyzed. If more than one dimuon candidate passes all these selections, the one with the highest χ^2 probability associated to the kinematic fit of the dimuon vertex is retained.

7 Fit to the invariant mass spectrum

The invariant mass spectrum in the range of our search presents mainly two contributions: the QCD continuum and the bottomonium resonances. To characterize these shapes for use in the mass scan, we perform a binned maximum likelihood fit to the total invariant mass

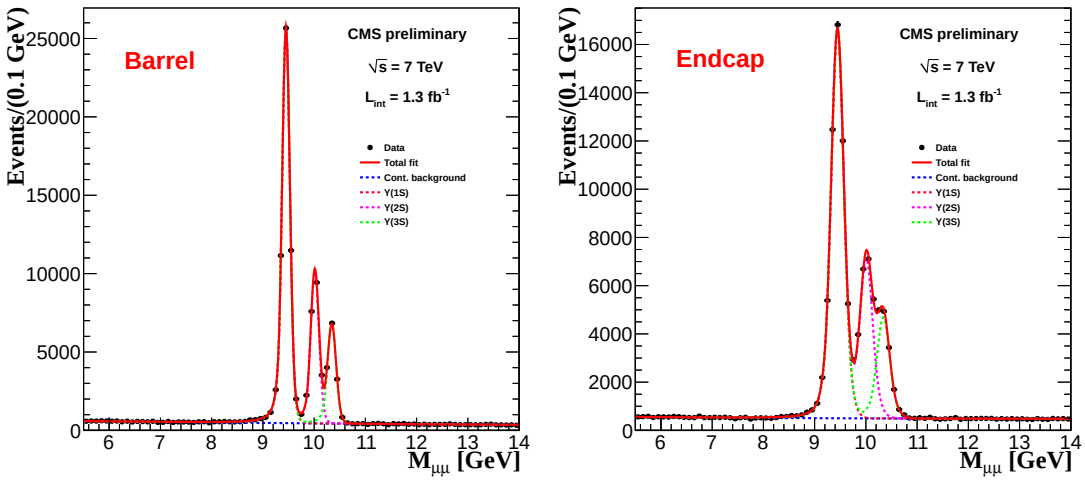

Figure 2: Dimuon invariant mass distribution for the central region (left) and for the endcaps (right) after the optimization of the cuts for the background reduction. The invariant mass distributions are fitted (red curves) accounting for the three bottomonium resonances and QCD background.

distribution. For the QCD background, we use a first-order polynomial probability density function (PDF). Each Υ is parametrized via a double Crystal Ball (CB) function. A CB function is formed convoluting a core Gaussian resolution with a power law side tail describing final state radiation. The resolution of one CB is left free in the fit but is constrained to be the same for all the three resonances. The resolution of the other CB function is determined from the fit on the Υ(1S) peak, and forced to scale with the mass of the other two resonances. As the resonances overlap, we fit for the presence of all three Υ states simultaneously. Therefore the PDF consists of three double CB functions. The mean of the CB of the Υ(1S) is left free in the fit, to accommodate a possible bias in the momentum scale calibration. The number of free parameters is reduced by fixing the Υ(2S) and Υ(3S) mass difference, relative to Υ(1S), to their world average values.

The fit to the Υ shape and continuum background is performed in the two acceptance regions (barrel and endcaps) separately, as shown in Figure 2. The number of events of Υ and continuum determined from the fit are given in Table 1.

Table 1: Summary of the number of events of Υ and continuum background from the invariant mass fit. The Υ contribution is summed over the three resonances.

Contribution	Number of events (barrel)	Number of events (endcap)
Υ	93753 ± 396	95876 ± 454
Continuum background	41210 ± 320	45792 ± 385

8 Efficiency

The efficiency of the selection of the a decaying into dimuon pairs can be factorized into three main contributions:

$$\epsilon = \epsilon_{acc} \times \epsilon_{trig} \times \epsilon_{sel} \tag{2}$$

where ϵ_{acc} is the kinematic acceptance for the a, ϵ_{trig} is the efficiency of the muon trigger, and ϵ_{sel} is the efficiency of the selection cuts applied to the dimuon candidates.

We do not treat ϵ_{acc} in the analysis. The main reason is that there are currently no Monte Carlo generators that reliably describe the a production mechanism (see also Section 10). The trigger and selection efficiencies (ϵ_{trig} and ϵ_{sel}) are measured with J/ψ events in data using the tag-and-probe technique. We perform this study in bins of η and p_{T} of the probe muon. The efficiency values extracted from data are compared with those obtained from the simulation of the $J/\psi \to \mu^+\mu^-$. The difference between the efficiency in data and MC is evaluated in bins of p_{T} and η and used as correction factor to weight the MC events in order to accommodate for possible discrepancies. These correction factors are typically of the order of few percent. For each dimuon candidate, the weight is the product of the correction factors for the two muons in the candidate.

The isolation cut efficiency, that contributes in ϵ_{sel}, cannot be measured using the J/ψ dataset as one of the main production mechanisms for J/ψ is through B-meson decays, resulting in non-isolated muons. This is not well accounted for in simulation, and would result in biased data/MC efficiency corrections. In order to estimate this correction, we use $Z \to \mu\mu$ events and consider the lower p_{T} spectrum of the probe muon. We assign a 5% systematic uncertainty for the isolation requirement efficiency to account for this, as described below.

For each a mass sample, the efficiency is defined as the weighted fraction of generated candidates which survive all the optimized cuts. This ratio contains in itself all the contributions to the efficiency and ranges from about 1% to 3.5% for the a mass range of 6 to 13 GeV. We fit these distributions with second (third) order polynomial functions in the barrel (endcaps) to use in the mass scan. The increase in the efficiency as a function of the invariant mass is mainly due to the p_{T} requirements on the muons at the HLT level. The fit is used in the determination of the upper limit on the cross section, taking into account the systematic uncertainty on the efficiency as described below.

9 Signal invariant mass resolution

As the light pseudoscalar is a narrow resonance ($\Gamma_\mu(a \to \mu^+\mu^-) \sim$ MeV for low $\tan\beta$ values), the a width is mostly driven by the experimental resolution of the CMS detector. Because of the different integrated magnetic field as a function of η, the muon p_{T} resolution (and therefore the a mass resolution) is better in the barrel than in the endcaps. For this reason, we split the MC and data samples in two regions based on η_a.

For each signal mass point, we determine the resolution by fitting the a invariant mass spectrum with two Crystal Ball (CB) functions (similar to the fit on the Y resonances) and the mass resolution is calculated as the weighted average of the widths of the two functions.

The resulting experimental resolution on the dimuon invariant mass ranges from 50 to 120 MeV (90 to 190 MeV) in the barrel (endcaps) for the mass range 6 to 13 GeV. These agree well with the resolution obtained from the Y resonances in data and MC. The resolution systematic uncertainty is discussed below. We perform a fit of the resolution as a function of the mass using the different signal samples. This fit is then used to extract the values of the dimuon mass resolution for each mass bin needed in the scan to determine the upper limit.

10 Systematic uncertainties

Several sources of systematic uncertainty affect these results, including a 4.5% uncertainty on the integrated luminosity. The efficiency corrections are determined using the tag-and-probe results described above. We determine event-by-event the uncertainty on the total efficiency corrections by propagating the uncertainties on the single muon corrections. This total event uncertainty on the efficiency is largely independent of mass, with a maximum value of 12%. We conservatively apply this value as a systematic uncertainty for every bin in the scan.

For the isolation cut, we assign as a systematic uncertainty the largest discrepancy between data and MC in the entire relevant p_{T} range, which is 5%. To take into account a possible systematic effect due to the simulation of the pseudoscalar experimental width, we evaluate the systematic uncertainty on the resolution of the a as the quadrature sum of the difference between the mass resolution of the a with a mass of 10 GeV and the resolution of the Y(2S) (which has the same mass) in MC simulation, and the difference between the latter and the mass resolution obtained for the Y(2S) from data. Additionally, the finite statistics for the determination of the mass resolution as a function of the dimuon mass contributes a source of uncertainty. We consider the mass ranges separately and include this systematic uncertainty in the calculation of the upper limit on the cross section times branching ratio. Overall, this adds a 11% (4%) effect for the barrel (endcaps).

We also consider systematic uncertainties on the background description (see Section 11). To estimate the shape uncertainty of the first-order polynomial background PDF, we fit the background with alternative shape hypotheses (a second order polynomial and an exponential function). We then generate toy-MC experiments using the alternative background hypotheses and fit the experiments using the first order polynomial. We then determine the systematic error on the first order polynomial from the distribution of the fitted parameters in the toy-MC experiments.

Finally, we note that the a modeling lacks a complete MC generator that accurately represents the kinematic of the scalar production and decay. Since no generator is currently considered to be appropriate for this model, at present we do not assign any systematic uncertainty to account for this effect.

11 Determination of the upper limit

In Figure 3 we show the invariant mass spectrum of data for the barrel and endcap regions. We perform the mass scan on this spectrum, dividing mass range 1 into 110 bins and mass range 2 in 100 bins of 30 MeV each. For each bin, we build a signal Gaussian PDF with a mean fixed to the center of the bin and a width determined from the fits described above. We use a first order polynomial to characterize the background. In mass range 1, we take into account the radiative tail of the Y(1S) by including its shape determined from the full invariant mass spectrum fit.

No significant discrepancy with SM background predictions is observed. Thus, we determine the 95% Confidence Level (CL) upper limit of the $pp \to a \to \mu\mu$ cross section times branching ratio as a function of the dimuon mass using the CLs approach [16]. Figure 4 shows the upper limit results for mass ranges 1 and 2 including the systematic uncertainties discussed above. These limits are significant in the context of the NMSSM. Constraints within the NMSSM can be phrased in terms of upper limits on $|\cos\theta_A|$. The larger the value of $\tan\beta$, the stronger the constraint. In Figure 5 we give upper limits, $|\cos\theta_A|^{\max}$, as a function of m_{a_1} for $\tan\beta = 1, 2, 3, 10, 30, 50$ using the $\sigma(pp \to a)BR(a \to \mu^+\mu^-)$ limits of Figure 4 and the fact that $\sigma(gg \to$

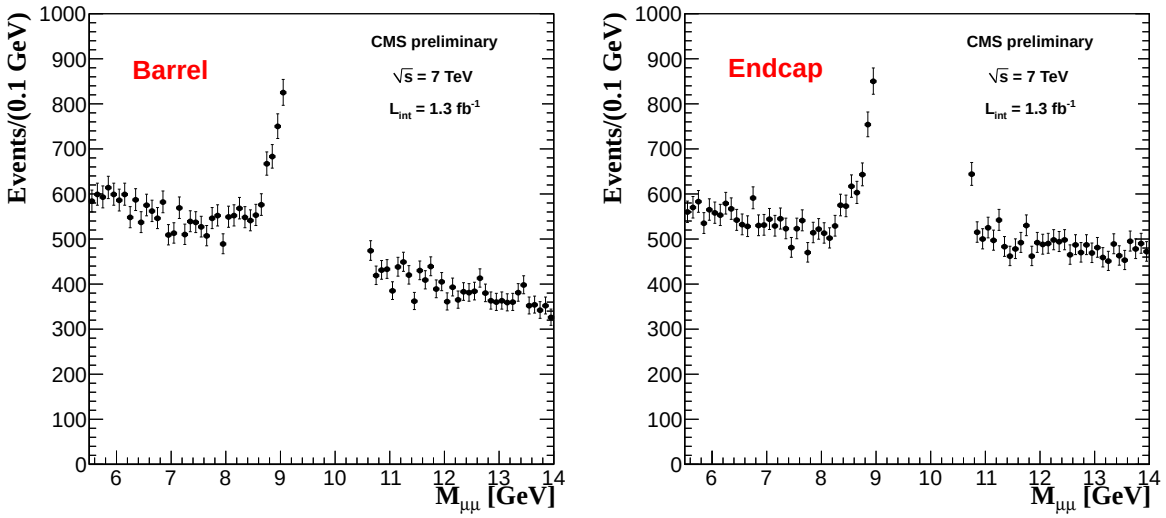

Figure 3: Invariant mass spectrum of data for barrel (left) and endcap (right) in the interval used for the scan.

$a) \leq \sigma(pp \to a)$. The upper limits are compared to an earlier analysis of the BaBar Y 1S and 3S data [17]. The CMS limits are superior for $m_{a_1} \geq 7.5\,\text{GeV}$ for $\tan\beta = 50$ decreasing to $m_{a_1} \geq 6\,\text{GeV}$ for $\tan\beta = 2$ and are superior for all masses at $\tan\beta = 1$.

12 Summary

We perform a search for a narrow, low mass, pseudoscalar a which is produced by $gg \to a$ and decays via $a \to \mu\mu$ in the mass ranges $5.5 - 8.8$ GeV and $11.5 - 14$ GeV, using a data sample corresponding to 1.3 fb^{-1} collected with the CMS detector. Given the good agreement with SM background predictions, we set upper limits on the cross section times branching ratio of $pp \to a \to \mu\mu$ in the mass ranges considered. These upper limits are applied in the context of the light pseudoscalar a_1 of the NMSSM to yield upper limits on the NMSSM parameter $|\cos\theta_A|$, which specifies the fraction by which the a_1 overlaps the MSSM doublet Higgs fields at the amplitude level. These limits are superior to those from BaBar for a significant portion of the $m_{a_1} < M_{\Upsilon_{1S}}$ mass range, and are the only limits available in the $m_{a_1} > M_{\Upsilon_{3S}}$ mass range.

References

[1] CMS Collaboration, "Combined results of searches for the standard model Higgs boson in pp collisions at sqrt(s) = 7 TeV", arXiv:1202.1488.

[2] ATLAS Collaboration, "Combined search for the Standard Model Higgs boson using up to 4.9 fb-1 of pp collision data at sqrt(s) = 7 TeV with the ATLAS detector at the LHC", arXiv:1202.1408.

[3] M. Maniatis, "The Next-to-Minimal Supersymmetric extension of the Standard Model reviewed", *Int.J.Mod.Phys.* **A25** (2010) 3505–3602, arXiv:0906.0777. doi:10.1142/S0217751X10049827.

[4] R. Dermisek and J. F. Gunion, "Escaping the large fine tuning and little hierarchy problems in the next to minimal supersymmetric model and $h \to aa$ decays",

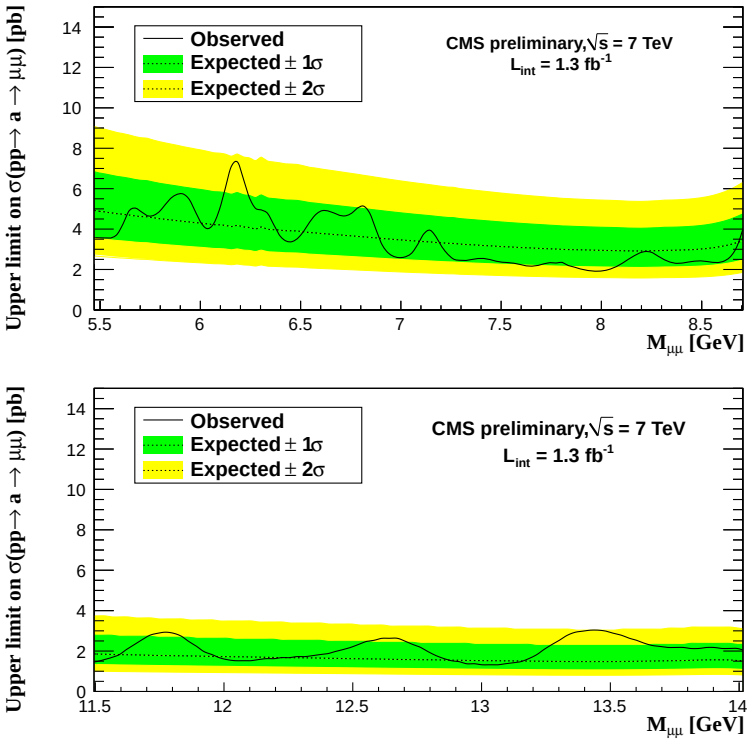

Figure 4: Upper limits at 95% C.L. on $\sigma(pp \rightarrow a \rightarrow \mu\mu)$ in mass range 1 and 2 including systematic uncertainties. The dotted lines correspond to the expected limits, and the green and yellow bands correspond to 1- and 2-sigma level uncertainties on the expected limits.

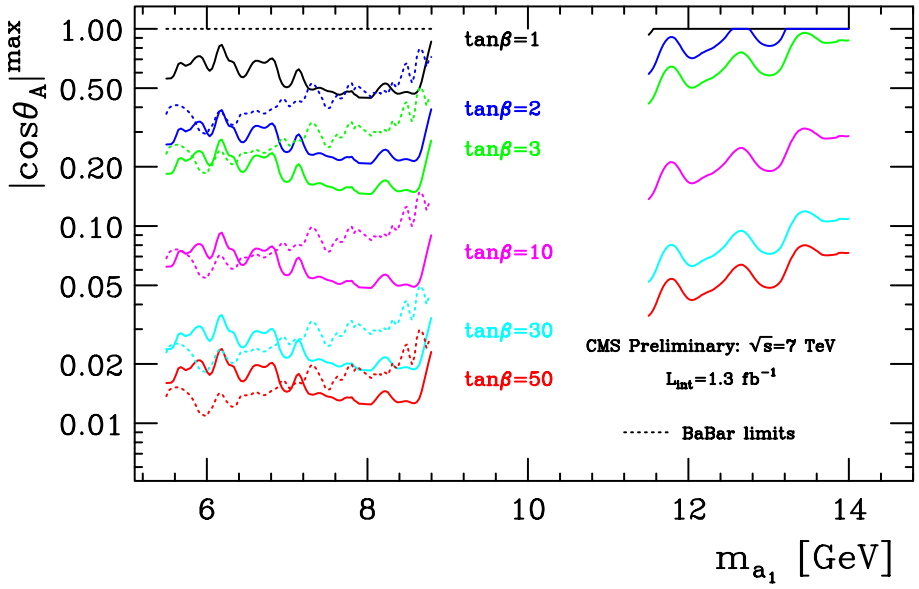

Figure 5: Upper limit on the NMSSM parameter $|\cos\theta_A|$ in mass ranges 1 and 2. The different solid curves correspond to different $\tan\beta$ values: $\tan\beta = 1$ (black), $\tan\beta = 2$ (blue), $\tan\beta = 3$ (green), $\tan\beta = 10$ (magenta), $\tan\beta = 30$ (cyan) and $\tan\beta = 50$ (red). For each $\tan\beta$ value in mass range 1 there is a second dotted curve that shows the limits from the BaBar Y analyses. There are no BaBar limits for $\tan\beta = 1$ in mass range 1, or for any $\tan\beta$ in mass range 2. The line at $|\cos\theta_A|^{max} = 1$ is equivalent to no limit.

Phys.Rev.Lett. **95** (2005) 041801, arXiv:hep-ph/0502105. doi:10.1103/PhysRevLett.95.041801.

[5] BaBar Collaboration, "Search for Dimuon Decays of a Light Scalar Boson in Radiative Transitions $\Upsilon \to \gamma A^0$", *Phys.Rev.Lett.* **103** (2009) 081803, arXiv:0905.4539. doi:10.1103/PhysRevLett.103.081803.

[6] BaBar Collaboration, "Search for a low-mass Higgs boson in $\Upsilon(3S) \to \gamma A^0$, $A^0 \to \tau^+\tau^-$ at BABAR", *Phys.Rev.Lett.* **103** (2009) 181801, arXiv:0906.2219. doi:10.1103/PhysRevLett.103.181801.

[7] ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaborations, LEP Working Group for Higgs Boson Searches Collaboration, "Search for neutral MSSM Higgs bosons at LEP", *Eur.Phys.J.* **C47** (2006) 547–587, arXiv:hep-ex/0602042. doi:10.1140/epjc/s2006-02569-7.

[8] CDF Collaboration, "Search for narrow resonances below the Upsilon mesons", *Phys. Rev.* **D72** (2005) 092003, arXiv:hep-ex/0507044. doi:10.1103/PhysRevD.72.092003.

[9] CDF Collaboration, "Search for narrow resonances lighter than Y mesons", *Eur. Phys. J.* **C62** (2009) 319–326, arXiv:0903.2060. doi:10.1140/epjc/s10052-009-1057-4.

[10] R. Dermisek and J. F. Gunion, "Direct production of a light CP-odd Higgs boson at the Tevatron and LHC", *Phys.Rev.* **D81** (2010) 055001, arXiv:0911.2460. doi:10.1103/PhysRevD.81.055001.

[11] M. Spira, "HIGLU: A program for the calculation of the total Higgs production cross-section at hadron colliders via gluon fusion including QCD corrections", arXiv:hep-ph/9510347.

[12] CMS Collaboration, "Alignment of the CMS Silicon Tracker during Commissioning with Cosmic Rays", *JINST* **5:T03009** (2010).

[13] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **0803** (2008) S08004.

[14] CMS Collaboration, "Tracking and Primary Vertex Results in First 7 TeV Collisions", *CMS Physics Analysis Summary* **TRK-10-005** (2010).

[15] CMS Collaboration, "Muon Reconstructon", *PAS* **MUO-10-002** (2010).

[16] T. Junk, "Confidence Level Computation for Combining Searches with Small Statistics", arXiv:hep-ex/9902006. doi:10.1016/S0168-9002(99)00498-2.

[17] R. Dermisek and J. F. Gunion, "New constraints on a light CP-odd Higgs boson and related NMSSM Ideal Higgs Scenarios", *Phys.Rev.* **D81** (2010) 075003, arXiv:1002.1971. doi:10.1103/PhysRevD.81.075003.