Modern Physics Letters A
 © World Scientific Publishing Company

Recent results on $B \rightarrow \mu^+\mu^-$ decays with the CMS experiment*

URS LANGENEGGER

*Paul Scherrer Institute
 CH-5232 Villigen PSI
 Switzerland*

*urs.langenegger@psi.ch
 (on behalf of the CMS collaboration)*

Received (Day Month Year)

Revised (Day Month Year)

Results on $B \rightarrow \mu^+\mu^-$ decays with the CMS experiment are reported, using 61 fb^{-1} of data recorded during LHC Run 1 and 2016. With an improved muon identification algorithm and refined unbinned maximum likelihood fitting methods, the decay $B_s^0 \rightarrow \mu^+\mu^-$ is observed with a significance of 5.6 standard deviations. Its branching fraction is measured to be $\mathcal{B}(B_s^0 \rightarrow \mu^+\mu^-) = [2.9 \pm 0.7 (\text{exp}) \pm 0.2 (\text{frag})] \times 10^{-9}$, where the first error is the combined statistical and systematic uncertainty and the second error quantifies the uncertainty of the B_s^0 and B^+ fragmentation probability ratio. The $B_s^0 \rightarrow \mu^+\mu^-$ effective lifetime is $\tau_{\mu^+\mu^-} = 1.70^{+0.61}_{-0.44}$ ps. No evidence for the decay $B^0 \rightarrow \mu^+\mu^-$ is found and an upper limit of $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-) < 3.6 \times 10^{-10}$ (at 95% confidence level) is determined. All results are consistent with the standard model of particle physics.

Keywords: B mesons; leptonic decays; CMS; LHC

PACS Nos.: 13.20.He

1. Introduction

The leptonic B meson^a decays $B^0 \rightarrow \mu^+\mu^-$ and $B_s^0 \rightarrow \mu^+\mu^-$ allow precision tests of the standard model (SM) of particle physics because their branching fractions can be calculated with small theoretical uncertainties. They are forbidden at tree level in the SM and are mediated via effective flavor-changing neutral-current Z -penguin and box processes, as illustrated in Fig. 1. The branching fractions^{1–5} in the SM are $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-) = (1.03 \pm 0.05) \times 10^{-10}$ and $\mathcal{B}(B_s^0 \rightarrow \mu^+\mu^-) = (3.66 \pm 0.14) \times 10^{-9}$, integrated over the B_s^0 meson decay time. The helicity suppression of these decays in the SM provides sensitivity to hypothetical (pseudo-)scalar interactions beyond the SM (BSM). The hierarchical nature of the Cabibbo-Kobayashi-Maskawa^{6,7} (CKM)

*Based on a seminar given at CERN on Sep. 21, 2019, and published in JHEP, 04, 188 (2020).

^aThe symbol B is used to denote B^0 , B_s^0 , and B^+ mesons and/or Λ_b baryons. Charge conjugation is implied throughout, except as noted.

2 Urs Langenegger

matrix implies that the decay $B^0 \rightarrow \mu^+ \mu^-$ is CKM-suppressed compared to $B_s^0 \rightarrow \mu^+ \mu^-$, since $|V_{td}| < |V_{ts}|$.

Fig. 1. The decays $B_s^0 \rightarrow \mu^+ \mu^-$ and $B^0 \rightarrow \mu^+ \mu^-$ are mediated by (left) Z -penguin and (right) box diagrams. The latter is suppressed by $(m_W/m_t)^2$ with respect to the former.

The heavy and light mass eigenstates of the B_s^0 meson, $|B_{s,L(H)}^0\rangle = p|B_s^0\rangle \pm q|\bar{B}_s^0\rangle$ (with $|p|^2 + |q|^2 = 1$), have different lifetimes. In the absence of CP violation, only the CP -odd heavy B_s^0 state, with a lifetime of $\tau_{B_{s,H}^0} = 1.615 \pm 0.009$ ps,⁸ can decay into the dimuon final state via the SM interactions. Because this fact is independent of the predicted numerical value of the branching fraction, it is of high interest to measure also the $B_s^0 \rightarrow \mu^+ \mu^-$ effective lifetime, defined⁹ as the time expectation value of the untagged rate by

$$\tau_{\mu^+ \mu^-} \equiv \frac{\int_0^\infty t [\Gamma(B_s^0(t) \rightarrow \mu^+ \mu^-) + \Gamma(\bar{B}_s^0(t) \rightarrow \mu^+ \mu^-)] dt}{\int_0^\infty [\Gamma(B_s^0(t) \rightarrow \mu^+ \mu^-) + \Gamma(\bar{B}_s^0(t) \rightarrow \mu^+ \mu^-)] dt}, \quad (1)$$

where t is the proper decay time of the B_s^0 meson. Experimentally, $\tau_{\mu^+ \mu^-}$ is determined by fitting a single^b exponential function, corrected for experimental artifacts like efficiency and resolution, to the decay time distribution of $B_s^0 \rightarrow \mu^+ \mu^-$ decays.

In the SM the branching fractions for these decays have been calculated beyond leading order in quantum chromodynamics (QCD) since more than a quarter century (cf. Fig. 2 and references therein). Most recently, computations of three-loop QCD corrections,² electroweak effects at next-to-leading order,³ and enhanced electromagnetic corrections^{4,5} have been completed. The (relative) errors are estimated to be smaller than 5% in the most recent calculation. By now, the theoretical error budget is dominated by external parametric uncertainties (either from CKM matrix element magnitudes or the B meson decay constant, depending on the number of dynamical quark flavors in the lattice QCD calculations¹⁰).

The experimental effort has been pursued both at e^+e^- machines for $B^0 \rightarrow \mu^+ \mu^-$ and at hadron colliders for $B_s^0 \rightarrow \mu^+ \mu^-$ (and $B^0 \rightarrow \mu^+ \mu^-$). The impressive sensitivity progression in the past four decades is illustrated in Fig. 2 (right). At the Large Hadron Collider (LHC), the decay $B_s^0 \rightarrow \mu^+ \mu^-$ has been measured with at least four standard deviation (σ) significances by the ATLAS,¹¹ CMS,¹² and LHCb¹³ collaborations. More recently, LHCb¹⁴ and CMS¹⁵ (with the result discussed here) each have reached more than 5σ individually. The $B_s^0 \rightarrow \mu^+ \mu^-$

^bIn general, the untagged decay rate of B_s^0 mesons, with heavy and light states, should be described by two exponential functions, not a single one.

Fig. 2. (left) Time evolution of the SM prediction. Note that the top quark was not yet discovered at the time of the first prediction beyond leading order and that many parameters were significantly off from today's values. (right) The experimental results for the branching fraction $\bar{B}(B_s^0 \rightarrow \mu^+ \mu^-)$, shown with blue symbols for upper limits and black symbols for measurements, and upper limits for $B(B^0 \rightarrow \mu^+ \mu^-)$, shown in red symbols.

effective lifetime has been measured first by the LHCb collaboration¹⁴ and now by the CMS collaboration.¹⁵ All confirmed results to date are in agreement with the SM predictions.

2. Experimental Strategy

The experimental approach starts with reconstructing dimuon candidates in a wide invariant mass region. The number of background candidates is reduced with advanced muon identification algorithms and multivariate analysis techniques in the selection. Finally, the number of signal decays is determined with an unbinned maximum likelihood fit to the mass distribution and other variables. Signal $B \rightarrow \mu^+ \mu^-$ decays are characterized by two muons, with an invariant mass $m_{\mu\mu}$ around the B^0 or B_s^0 mass, originating from a common point in space where the B meson decayed. The background has several components with characteristics that allow its reduction with respect to the signal:

- Combinatorial background from two semileptonic B decays or from one semileptonic B decay together with a hadron misidentified as a muon (fake muon). The two muons do not originate from the same point in space. This component is the limiting factor for the measurement of $B_s^0 \rightarrow \mu^+ \mu^-$.
- Rare decays of a single B hadron, illustrated in Fig. 3 (left). They consist of decays with (1) two muons (*e.g.*, $B^0 \rightarrow \pi^0 \mu^+ \mu^-$, where the pion is not considered in the final state reconstruction), (2) one muon combined with a fake muon (*e.g.*, $B_s^0 \rightarrow K^- \mu^+ \nu$), or (3) with two fake muons (*e.g.*, $B_s^0 \rightarrow K^+ K^-$). These decays constitute a dangerous background affecting in particular $B^0 \rightarrow \mu^+ \mu^-$. In the first two cases, the mass distribution is

leaking into the B^0 mass region, but the missing particle provides a handle to reduce this contribution. The last case constitutes a peaking background near the B^0 mass region. The wrong mass hypothesis, muon instead of kaon or pion, shifts the mass distribution from the B_s^0 mass to lower values.

In the discussion above, a fake muon is a hadron misidentified in the detector as a muon, either because of its decay-in-flight or punch-through (hits in the muon system associated to a charged track).

Fig. 3. (left) Background invariant mass distribution from rare B decays with 0, 1, or 2 fake muons (hadrons misidentified as muons). It can be deduced from the plot that these decays are much more dangerous for $B^0 \rightarrow \mu^+ \mu^-$ than for $B_s^0 \rightarrow \mu^+ \mu^-$ because of (1) their mass distribution covering the B^0 mass region ($m_{B^0} = 5.280 \text{ GeV}^8$) more strongly than the B_s^0 mass region ($m_{B_s^0} = 5.367 \text{ GeV}^8$) and (2) the suppressed rate of $B^0 \rightarrow \mu^+ \mu^-$ compared to $B_s^0 \rightarrow \mu^+ \mu^-$. (right) Illustration of measurements of f_s/f_u , the ratio of the B_s^0 and B^+ fragmentation probabilities for different mean B meson transverse momentum $p_\perp$ and at different center-of-mass energies $\sqrt{s}$. The bands indicate the parametrizations of the LHCb experiment. The value adopted by CMS corresponds to the PDG⁸ value, with an ad-hoc enlargement of the error (see text for more details). It is interesting to note that the most significant slope determination, shown in green, results in the smallest slope.

The branching fraction $\mathcal{B}(B_s^0 \rightarrow \mu^+ \mu^-)$ is determined relative to a normalization sample with a well-known branching fraction. Starting from Ref. 16, the decay $B^+ \rightarrow J/\psi K^+$, with $J/\psi \rightarrow \mu^+ \mu^-$, has been used as normalization schematically as follows:

$$\overline{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-) = \frac{N_{B_s^0}}{N_{B^+}} \frac{f_u}{f_s} \frac{\varepsilon_{B^+}^{\text{tot}}}{\varepsilon_{B_s^0}^{\text{tot}}} \mathcal{B}(B^+ \rightarrow J/\psi K^+) \mathcal{B}(J/\psi \rightarrow \mu^+ \mu^-), \quad (2)$$

where $N_{B_s^0}(N_{B^+})$ is the number of signal $B_s^0 \rightarrow \mu^+ \mu^-$ ($B^+ \rightarrow J/\psi K^+$) decays, $\varepsilon_{B_s^0}^{\text{tot}}$ ($\varepsilon_{B^+}^{\text{tot}}$) is the total signal (B^+) efficiency, $\mathcal{B}(B^+ \rightarrow J/\psi K^+) = (1.01 \pm 0.03) \times 10^{-3}$ and $\mathcal{B}(J/\psi \rightarrow \mu^+ \mu^-) = (5.96 \pm 0.03) \times 10^{-2}$,⁸ and f_u/f_s is the ratio of the B^+ and B_s^0 fragmentation functions. A similar approach is used for $\mathcal{B}(B^0 \rightarrow \mu^+ \mu^-)$, using $f_d/f_u = 1$.⁸

In addition to serving as a ‘normalization sample’, the $B^+ \rightarrow J/\psi K^+$ candidates also constitute a large sample of B^+ mesons where the Monte Carlo (MC) simulation can be validated against data. To compare B_s^0 mesons in data and MC simulation, a ‘control sample’ of $B_s^0 \rightarrow J/\psi \phi$ (with $J/\psi \rightarrow \mu^+ \mu^-$ and $\phi \rightarrow K^+ K^-$) candidates is used. Because the analysis relies on MC simulation for the efficiency determination, the validation of the MC simulation is essential. Furthermore, these $B \rightarrow J/\psi X$ samples can be used to study differences between data and MC simulation regarding b quark production processes (*e.g.*, by combining such decays with another muon from a semileptonic $\bar{B}$ decay) and to study in detail the b quark hadronization into B_s^0 or B^+ mesons.

The usage of $B^+ \rightarrow J/\psi K^+$ as a normalization sample is motivated by the minimal difference in the final state with respect to the signal decay (one additional charged particle). This leads to a substantial reduction of the systematic uncertainties. However, it implies a direct dependence of $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-)$ on f_s/f_u [though not for the effective $B_s^0 \rightarrow \mu^+ \mu^-$ lifetime or $\mathcal{B}(B^0 \rightarrow \mu^+ \mu^-)$]. In Fig. 3 (right) the experimental situation of f_s/f_u measurements and combinations is illustrated. The LHCb collaboration¹⁷ obtains a significant slope vs. the transverse momentum ($p_\perp$) of the B meson, while the ATLAS collaboration¹⁸ and the CMS collaboration (internal study performed for the result discussed here) see no such effect. Nevertheless the CMS experiment decided to account for a hypothetical $p_\perp$ dependence and possible center-of-mass energy ($\sqrt{s}$) dependence by adding an ad-hoc error to f_s/f_u . An uncertainty of 0.008 is derived from the difference between the value of f_s/f_u in Ref. 8, obtained at $\sqrt{s} = 7$ TeV, and that in Ref. 19, obtained at $\sqrt{s} = 13$ TeV. In addition, with the parametrization of the $p_\perp$ dependence in Ref. 19, a difference of 0.013 is determined between the f_s/f_u values at the average $p_\perp$ of Ref. 19 and the average $p_\perp$ of the $B_s^0 \rightarrow \mu^+ \mu^-$ candidates in this analysis. Ref. 17 would imply a much smaller $p_\perp$ dependence, but was published too late to be included. In summary, the CMS experiment uses $f_s/f_u = 0.252 \pm 0.012(\text{exp}) \pm 0.015(\text{CMS})$ where the first error is from the PDG⁸ and the second error is the ad-hoc error of CMS.

In the future, a normalization to other decay modes may provide a less contentious solution. While $\mathcal{B}(B_s^0 \rightarrow J/\psi \phi)$ normally has a dependence on f_s/f_d (or f_s/f_u) when determined at hadron colliders, results obtained at e^+e^- machines at the $\Upsilon(5S)$ provide additional input and may lead to a smaller overall error.

To avoid a possible bias, the analyses searching for and measuring the $B \rightarrow \mu^+ \mu^-$ decays have been pursued as ‘blind’ analyses since a long time. This implies that a signal region, often defined in terms of the invariant mass, is hidden during the development and optimization of the analysis methodology. To take full advantage of

6 *Urs Langenegger*

analysis improvements (*e.g.*, improved muon identification, re-processed data with better tracking resolutions, etc.) it is advantageous to re-analyze the old, previously published, datasets. This implies a ‘re-blinding’ of the data, which is, however, not a problem because the improvements normally change the set of candidates noticeably. An alternative would be to combine new results only statistically with the old results, albeit at a loss of sensitivity. In this analysis, the first approach is used and the dimuon mass range $5.2 < m_{\mu\mu} < 5.45$ GeV was kept (re-)blinded until the entire selection and fitting procedure was finalized.

3. Detector and Data

The data for this analysis was collected by the CMS experiment²⁰ in LHC Run 1 (25 fb^{-1}) and in 2016 (36 fb^{-1}), as summarized in Table 1. The CMS experiment is very well suited for $B \rightarrow \mu^+ \mu^-$ measurements because its silicon tracker, composed of a pixel detector with 66 million pixels of $100 \mu\text{m} \times 150 \mu\text{m}$ and a micro-strip detector with 10 million strips with pitches between 80 and $180 \mu\text{m}$, provides outstanding three-dimensional (3D) vertexing and tracking capabilities in a very homogeneous solenoidal magnetic field of 3.8 T. The tracker is divided into barrel and endcap parts. In Run 2, the micro-strip detector was subject to operational instabilities and the data are therefore divided into two separate data-taking periods, 2016A and 2016B, of 16 fb^{-1} and 20 fb^{-1} , respectively. The systematic error of the tracking efficiency is estimated^{21,22} to be 4% (2.3%) in Run 1 (2016).

Table 1. Summary of the data, together with center-of-mass energy $\sqrt{s}$, the integrated luminosity $\mathcal{L}$, the pileup quantified as the average number of pp collision vertices reconstructed as primary vertices $\langle N_{\text{PV}} \rangle$, and the channel definition based on the pseudorapidity of the most-forward muon $|\eta_{\mu}^{\text{f}}|$. Note that in Run 2 (2016A and 2016B) the region $|\eta_{\mu}^{\text{f}}| > 1.4$ (the forward channel of Run 1) is no longer present because of trigger rate constraints.

Data-taking period	$\sqrt{s}$ [TeV]	$\mathcal{L}$ [fb^{-1}]	Pileup $\langle N_{\text{PV}} \rangle$	Channels	
				central	forward
2011	7	5	8	$0 < \eta_{\mu}^{\text{f}} < 1.4$	$1.4 < \eta_{\mu}^{\text{f}} < 2.1$
2012	8	20	15	$0 < \eta_{\mu}^{\text{f}} < 1.4$	$1.4 < \eta_{\mu}^{\text{f}} < 2.1$
2016A	13	16	18	$0 < \eta_{\mu}^{\text{f}} < 0.7$	$0.7 < \eta_{\mu}^{\text{f}} < 1.4$
2016B	13	20	18	$0 < \eta_{\mu}^{\text{f}} < 0.7$	$0.7 < \eta_{\mu}^{\text{f}} < 1.4$

Muons are detected in four muon stations, using three complementary detector types, interspersed among the steel flux-return plates. Standalone muons are formed from hits in the muon stations and combined with silicon tracker tracks to form so-called global muons.^{23,24} A dedicated boosted decision tree (BDT) was trained separately for Run 1 and Run 2 data to obtain the best possible hadron-to-muon misidentification probability. The starting point for this BDT are global muons. The variables used in the BDT are based on measurements from the (1) silicon tracker, (2) the muon system, and (3) the combined global muon reconstruc-

tion. This new muon BDT achieves an average muon misidentification probability of 6×10^{-4} and 10^{-3} for pions and kaons, respectively, with a muon identification efficiency of about 75%. Compared to the previous analysis,¹² the muon BDT is operated at a significantly lower muon misidentification rate and roughly the same muon identification efficiency. The muon BDT is extensively validated with kinematically identified samples of muons, pions, kaons, and protons (from the decays $J/\psi \rightarrow \mu^+ \mu^-$, $K_s^0 \rightarrow \pi^+ \pi^-$, $\phi \rightarrow K^+ K^-$, and $\Lambda \rightarrow p \pi^-$, respectively) in data and MC simulation. All distributions of variables used in the muon BDT, the BDT discriminator distributions, and the absolute muon misidentification probability are found to be consistent between data and simulation. The systematic error for the muon efficiency is determined from the difference of the efficiency ratio of the muon BDT discriminator requirement for $B^+ \rightarrow J/\psi K^+$ and $B_s^0 \rightarrow J/\psi \phi$ between data and simulation, which agrees to better than 3%. The systematic error on the muon misidentification is derived from a direct comparison of the absolute muon misidentification probability in data and simulation (10% relative uncertainty for pions and kaons). For protons, the very small sample size of $\Lambda \rightarrow p \pi^-$, with a misidentified proton, does not allow this approach and the error of the average muon misidentification probability, in data and simulation, is used instead (60% relative uncertainty).

The trigger²⁵ of the CMS experiment has two stages: the first stage is based on custom hardware processors and selects two muons with either no or minimal $p_\perp^\mu$ threshold (because of the strong magnetic field, there is an implicit $p_\perp^\mu$ threshold of about 3.5 GeV in the central region). The second stage, the high-level trigger (HLT) consists of a processor farm running the full event reconstruction software with reduced sets of calibration constants. The normalization sample is triggered with a setup that is very similar to the signal setup with the exception that the two muons must be consistent with originating from a J/ψ meson from a B decay (‘displaced J/ψ ’). The signal (normalization) trigger efficiency varies over the data-taking periods from 65–75% (50–75%). The systematic uncertainty on the ratio of the trigger efficiency is estimated to be 3%.

The tracking detectors of the CMS experiment induce a strong pseudorapidity (η) dependence of the mass resolution. Therefore the analysis sensitivity benefits from subdividing the data into ‘channels’, according to the η of the most forward muon. Because of trigger changes over the years, the boundary between the ‘central’ and ‘forward’ channels is different between Run 1 and Run 2. In total, there are eight channels in the analysis (central and forward in four data-taking periods) as summarized in Table 1.

At the instantaneous luminosity of the LHC, multiple pp interactions (pileup) occur in each bunch crossing. Tracks from other pp interaction vertices (primary vertices or PVs) increase the combinatorial background and complicate the determination of key variables in the selection, as discussed below. Table 1 lists the average number of PVs reconstructed in the different data-taking periods.

4. Candidate Selection

The $B \rightarrow \mu^+ \mu^-$ candidate reconstruction starts with two global muons with $p_{\perp}^{\mu} > 4 \text{ GeV}$ and a small distance of closest approach $d_{\text{ca}} < 0.8 \text{ cm}$ between their trajectories. The two muons are constrained to originate from a common point (secondary vertex or SV) and to have an invariant mass $4.9 < m_{\mu\mu} < 5.9 \text{ GeV}$, after refitting their momenta to include the SV as an additional hit. For each reconstructed B candidate one specific PV is chosen as the B -meson origin (denoted below as B -PV), based on the longitudinal impact parameter ℓ_z along the beam axis of the extrapolated B -meson trajectory. The PVs are refitted by excluding tracks from the B candidate. The variables involved in the selection are sketched in Fig. 4 (left) for a signal event. The $B \rightarrow \mu^+ \mu^-$ selection exploits the differences between signal and background. Signal $B \rightarrow \mu^+ \mu^-$ decays are characterized by a SV with a good fit χ^2/dof , separated from the B -PV by a large flight length ℓ_{3D} and significance $\ell_{3D}/\sigma(\ell_{3D})$, where $\sigma(\ell_{3D})$ is the error of ℓ_{3D} . The B -meson proper decay time is measured as $t = m_{\mu\mu} \ell_{3D} / p^{\mu^+ \mu^-}$ in 3D space. The B -meson momentum is well aligned with the flight direction (the direction from the B -PV to the SV), implying a small opening angle α_{3D} and a small B impact parameter δ_{3D} and significance $\delta_{3D}/\sigma(\delta_{3D})$.

Fig. 4. (left) Sketch of a signal decay with the variables involved (see text for a description). (right) Sketch of a background decay showing aspects of the candidate isolation. Tracks from pileup PVs that would fall within the isolation cone (shown in pink) are excluded from the isolation calculation by requiring that all tracks considered are associated with the B -PV or are close to the SV (see text for more details).

For a signal decay, the two muons are the only final-state particles of the B meson decay, and therefore the B meson and the muons are isolated, *i.e.*, not many other charged tracks are nearby. This isolation is quantified by two sets of variables. The first set determines a macroscopic isolation I using

$$I = \frac{p_{\perp}^B}{p_{\perp}^B + \sum_{\text{trk}} p_{\perp}^{\text{trk}}}, \quad (3)$$

where the sum includes all tracks, not part of the B candidate, inside a cone with

radius $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\varphi)^2} < 0.7$, where $\Delta\eta$ ($\Delta\varphi$) is the difference between the track and the B candidate η (azimuthal angle φ). Because of tracks from other PVs not related to the B -PV, it is essential to only use tracks that are either associated to the B -PV or are close to the SV. Without these requirements, the tracks from other pileup PVs would bias I , as indicated in Fig. 4 (right) for a background event. The tracks in the sum must fulfill $p_\perp > 0.9$ GeV and have $d_{\text{ca}}^{\text{SV}} < 0.05$ cm with respect to the SV. These requirements were optimized with $B \rightarrow J/\psi X$ decays for strongest background rejection and best agreement between data and simulation. Similar isolation variables, I_{μ_1} and I_{μ_2} , are determined for each of the two muons, although with different track requirements in the sum, $p_\perp > 0.5$ GeV and $d_{\text{ca}}^{\text{SV}} < 0.1$ cm, and a smaller cone radius of $\Delta R = 0.5$.

The second set of isolation variables is based on microscopic observables determined with tracks of $p_\perp > 0.5$ GeV that are not part of the B candidate and are not associated to any non- B -PV (*i.e.*, if they are associated to a PV it must be the B -PV): the minimum distance of closest approach, d_{ca}^0 , of any qualifying track to the SV and $N_{\text{trk}}^{\text{close}}$, the number of tracks with $d_{\text{ca}}^{\text{SV}} < 0.03$ cm to the SV.

Many of these variables are correlated with each other. For instance, the flight length and its significance are correlated with $N_{\text{trk}}^{\text{close}}$, and the isolation variables I , I_{μ_1} , and I_{μ_2} are correlated with each other. A robust approach in such a situation is to train a multivariate analysis technique²⁶ in the form of a BDT. This selection BDT was trained on $B_s^0 \rightarrow \mu^+ \mu^-$ and $B^0 \rightarrow \mu^+ \mu^-$ signal decays from MC simulation and combinatorial background from the data sideband $5.45 < m_{\mu\mu} < 5.9$ GeV. Since data events are used in the training, the MC and data samples were randomly split into three subsets to ensure that the *training* and *validation* of a BDT are performed on subsets completely independent of its *application*. This procedure implies that three BDTs are required for the analysis of each channel. A preselection removes candidates with extreme outlier values in the variables and requires the SV to be well separated from the B -PV, $\ell_{3D}/\sigma(\ell_{3D}) > 4$. With the subdivision of the data into three subsets per channel, the number of events available for training is somewhat limited after the preselection (at least 6000 candidates in any subset). For the per-channel optimization of the BDT configuration, a set of core variables [$\ell_{3D}/\sigma(\ell_{3D})$, α_{3D} , $\delta_{3D}/\sigma(\delta_{3D})$, d_{ca}^0 , χ^2/dof , $N_{\text{trk}}^{\text{close}}$, I , I_{μ_1} , I_{μ_2}] is iteratively combined with a subset of other B -candidate variables [d_{ca} , δ_{3D} , ℓ_{3D} , $\ell_{xy}/\sigma(\ell_{xy})$, $p_\perp^B$, η_B]. The best BDT configurations are chosen based on (1) the maximum of $S/\sqrt{S+B}$, where S (B) is the expected $B_s^0 \rightarrow \mu^+ \mu^-$ signal (combinatorial background, extrapolated from the sideband) yield in the mass region $5.3 < m_{\mu\mu} < 5.45$ GeV and (2) the visual assessment of the agreement between data and MC simulation using large samples of exclusive $B \rightarrow J/\psi X$ candidates.

The decays $B^+ \rightarrow J/\psi K^+$ and $B_s^0 \rightarrow J/\psi \phi$, both with $J/\psi \rightarrow \mu^+ \mu^-$, allow the validation of the selection BDT and serve as normalization and control samples, respectively. Their reconstruction starts with two oppositely charged muons with $p_\perp^\mu > 4$ GeV, $p_\perp^{\mu^+ \mu^-} > 7$ GeV, and $2.9 < m_{\mu\mu} < 3.2$ GeV. They are combined with

one or two tracks, respectively, assumed to be kaons and with $p_{\perp} > 0.6$ GeV. To reduce combinatorial background, the maximum distance of closest approach between any pair of tracks is required to fulfill $d_{\text{ca}}^{\text{max}} < 0.08$ cm. For $B_s^0 \rightarrow J/\psi \phi$, the two tracks must fulfill $1.01 < m_{K^+K^-} < 1.03$ GeV. To allow the selection of the $B \rightarrow J/\psi X$ candidates with the same selection BDT as for the signal $B \rightarrow \mu^+ \mu^-$ decays, their variable distributions should mirror the corresponding ones from signal decays. This implies that the SV χ^2/dof is not based on the full SV fit with three or four tracks, but only on the dimuon vertex fit. In addition, for all isolation variables, the kaon track(s) are not part of the track sums of Eq. 3. Example mass distributions, together with fits to the data using signal (double Gaussian functions with a common mean) and background components, are shown in Fig. 5. The background components for $B^+ \rightarrow J/\psi K^+$ contain an exponential function for the combinatorial background, an error function for partially reconstructed $B \rightarrow J/\psi K^+ X$ decays, and a MC simulation based shape for the peaking background from $B^+ \rightarrow J/\psi \pi^+$ (fixed to 4%⁸ of the total signal yield). For $B_s^0 \rightarrow J/\psi \phi$ the background is parametrized by an exponential function for the combinatorial component and a MC simulation based shape for the peaking background from $B^0 \rightarrow J/\psi K^{*0}$, with $K^{*0} \rightarrow K^+ \pi^-$ (where the pion is treated as a kaon). The total $B^+ \rightarrow J/\psi K^+$ normalization yield used for the determination of $\mathcal{B}(B_s^0 \rightarrow \mu^+ \mu^-)$ is $N_{B^+} = (1.43 \pm 0.06) \times 10^6$, where the error is dominated by the systematic component of (relative) 4%, obtained from the yield comparison between fits without and with J/ψ -mass constraints.

Fig. 5. Invariant mass distribution for (left) the normalization sample $B^+ \rightarrow J/\psi K^+$ and (right) the control sample $B_s^0 \rightarrow J/\psi \phi$. The data are shown by solid black circles, the result of the fit is overlaid with the thick solid black line, and the different components are shown with hatched regions.

Background-subtracted variable distributions in data are compared to the corresponding distributions in the MC simulation. In Fig. 6 the distributions for the $p_\perp$ of the subleading muon (the muon with the lower $p_\perp$) and the $\ell_{3D}/\sigma(\ell_{3D})$ of the B^+ meson are shown as examples. The MC simulation provides a reasonable description of the data; the remaining discrepancies are fully accounted for in the systematic uncertainty. The pileup dependence of the HLT tracking in Run 2 affects the normalization sample stronger than the signal sample because of the displacement requirement in the $B^+ \rightarrow J/\psi K^+$ HLT path. This is corrected for with an offline reweighting depending on the number of reconstructed PV and $\ell_{xy}/\sigma(\ell_{xy})$. The remaining systematic error from this corrections is estimated to be 6% for 2016A and 5% for 2016B for the branching fraction measurement. For the $B_s^0 \rightarrow \mu^+ \mu^-$ effective lifetime measurement, a systematic error of 0.07 ps is estimated. This is the second-largest systematic uncertainty for both results.

Fig. 6. Comparison of the measured and simulated distributions of (left) the subleading muon $p_\perp$ and (right) the flight length significance $\ell_{3D}/\sigma(\ell_{3D})$ for $B^+ \rightarrow J/\psi K^+$ decays in data and MC simulation. Background-subtracted data are shown by solid black circles, the MC simulation by hatched histograms. In the lower panels the ratio between data and MC simulation is shown. The band at $\pm 20\%$ is just to guide the eye.

In Fig. 7 the BDT discriminator response is shown for dimuon candidates, illustrating the background rejection, and for $B^+ \rightarrow J/\psi K^+$ candidates, illustrating the agreement between data and MC simulation. The systematic error of the selection efficiency is estimated from the double ratio

$$D = \frac{\left[\frac{\varepsilon(B^+ \rightarrow J/\psi K^+)}{\varepsilon(B_s^0 \rightarrow J/\psi \phi)} \right]_{\text{data}}}{\left[\frac{\varepsilon(B^+ \rightarrow J/\psi K^+)}{\varepsilon(B_s^0 \rightarrow J/\psi \phi)} \right]_{\text{MC}}}, \quad (4)$$

where the control sample $B_s^0 \rightarrow J/\psi \phi$ is used as a placeholder for the signal sample. Depending on the channel, this systematic error varies between 5% and 10%. It constitutes the largest contribution to the overall systematic error for the branching fraction. The selection efficiency depends on the unknown true effective lifetime because of the displacement and isolation criteria. This uncertainty of 1–3%, depending on data-taking period and analysis channel, is estimated with $\Delta \equiv [\varepsilon_{\text{tot}}(\tau_{B_{s,H}^0}) - \varepsilon_{\text{tot}}(\tau_{B_{s,L}^0})]/\sqrt{12}$ using simulated samples with different effective lifetimes.

Fig. 7. Comparison of the BDT discriminator response for (left) dimuons in data (background from the invariant mass sideband $5.45 < m_{\mu\mu} < 5.9$ GeV) and MC simulation ($B_s^0 \rightarrow \mu^+\mu^-$ signal) and (right) $B^+ \rightarrow J/\psi K^+$ candidates in data and MC simulation. The left plot illustrates the rejection power of the BDT against (combinatorial) background, the right plot shows the level of agreement between data and MC simulation, an ingredient in the determination of the systematic error of the analysis efficiency. The arrow in the left plot indicates the category boundary of the BDT discriminator range used in the $\mathcal{B}(B_s^0 \rightarrow \mu^+\mu^-)$ result determination (cf. Table 2 below).

The mixture of $b\bar{b}$ quark production processes in the MC simulation is not necessarily the same as in data. At leading order three processes contribute to heavy quark production in pp collisions: gluon splitting, gluon-gluon fusion, and flavor excitation. The b quarks from gluon splitting are closer together in phase space than for the other two processes, where the two b quarks tend to be back-to-back in the transverse plane. Therefore a $B_s^0 \rightarrow \mu^+\mu^-$ decay from a B_s^0 meson produced in gluon splitting will be, on average, less isolated than from a B_s^0 meson from gluon-gluon fusion. Therefore a mismatch of the production process mixture in data and simulation will imply a systematic uncertainty on the selection efficiency. The mix-

ture of production processes is studied by combining a $B^+ \rightarrow J/\psi K^+$ ($B_s^0 \rightarrow J/\psi \phi$) candidate with another muon μ_3 (assumed to originate from the semileptonic decay of the other b hadron in the event) and studying the $\Delta R(B, \mu_3)$ distribution. Fitting templates from MC simulation for gluon-splitting and the sum of gluon-gluon fusion plus flavor excitation to the data allows an estimate of the systematic uncertainty for the efficiency ratio at the 3% level.

As a cross check, the effective^c branching fraction $\mathcal{B}(B_s^0 \rightarrow J/\psi \phi)$ was determined with an approach equivalent to Eq. 2, where the number of $B_s^0 \rightarrow J/\psi \phi$ decays is used for $N_{B_s^0}$, for all channels in all data-taking periods. The standard deviation of these eight measurements is 4%, smaller than the combination of the systematic uncertainties due to analysis efficiency, tracking efficiency (the decay $B_s^0 \rightarrow J/\psi \phi$ has one additional kaon track compared to $B^+ \rightarrow J/\psi K^+$), hadronization uncertainties (B_s^0 vs. B^+), and yield determinations. Therefore we conclude that the systematic error is not underestimated.

Efficiency-corrected yield ratios for $B^+ \rightarrow J/\psi K^+$, $B_s^0 \rightarrow J/\psi \phi$, and $B^0 \rightarrow J/\psi K^{*0}$ (with $K^{*0} \rightarrow K^+ \pi^-$) decays are studied in the range $10 < p_{\perp}^B < 100$ GeV and $0 < |\eta_B| < 2.2$ to estimate a possible dependence on these kinematic variables. No significant slope is observed vs. $p_{\perp}^B$ or η_B for f_s/f_u , f_s/f_d , or f_d/f_u (as a control measurement).

5. Results

For the determination of the results, the per-channel event samples are further subdivided into mutually exclusive categories of different signal-to-background ratios by introducing high- and low-BDT categories in the BDT discriminator distributions. This is illustrated in Fig. 7 (left), where the arrow shows the boundary between these categories. This categorization was optimized separately for the branching fraction measurement and the effective lifetime measurements. Table 2 provides the category boundaries. They vary over the channels because the BDT configuration has been optimized independently for each channel and results in different BDT discriminator distributions.

The branching fractions $\overline{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-)$ and $\mathcal{B}(B^0 \rightarrow \mu^+ \mu^-)$ are determined with a 3D unbinned extended maximum likelihood fit to the dimuon invariant mass distribution, the relative mass resolution, and the binary distribution for the dimuon pairing configuration $\mathcal{C}$ ($\mathcal{C} = \pm 1$ for the two muons bending towards or away from each other, respectively). While the first two variables have been used already in the past,¹² the last variable was added as a protection against a possible underestimation of the $B \rightarrow h^+ h^-$ background (critical for the $B^0 \rightarrow \mu^+ \mu^-$ measurement). Such an underestimate is possible because the $B \rightarrow h^+ h^-$ contribution is determined under the assumption that the dimuon fake rate is the product of the single

^cThis is not an absolute measurement of $\mathcal{B}(B_s^0 \rightarrow J/\psi \phi)$ because of f_s/f_u in Eq. 2. However, the point of this study is not the absolute measurement but rather the study of the stability of this quantity in different channels, data-taking periods, and detector configurations.

Table 2. BDT discriminator category boundaries for the branching fraction measurement (left part) and the effective lifetime measurements (right part). These boundaries are illustrated in Fig. 7 (left). In the 2011 data-taking period, there is no low-BDT category because of the limited number of events.

	branching fraction measurement		effective lifetime measurement	
	central	forward	central	forward
2011	{0.28, 1}	{0.21, 1}	{0.22, 1}	{0.19, 1}
2012	{0.27, 0.35, 1}	{0.23, 0.32, 1}	{0.32, 1}	{0.32, 1}
2016A	{0.19, 0.30, 1}	{0.19, 0.30, 1}	{0.22, 1}	{0.30, 1}
2016B	{0.18, 0.31, 1}	{0.23, 0.38, 1}	{0.22, 1}	{0.29, 1}

muon fake rates. While there is no evidence for such an underestimation for $\mathcal{C} = -1$ with the muons bending away from each other, the situation is less clear for $\mathcal{C} = +1$ where the muon tracks can be close together in the muon system. Since the effect is different for the two configurations, the $\mathcal{C}$ distribution is introduced. Its shape is taken from MC simulation for the signal and all background components. In the fit, a scale factor is used to correct the expected background component yields for the case $\mathcal{C} = +1$. A second, independent approach to control such an underestimate is to strongly reduce the muon misidentification probability (in this analysis, the misidentification probability was reduced by about 50% compared to the previous analysis¹²).

The signal probability density functions (PDFs) are based on a Crystal Ball function²⁷ for the invariant mass and a nonparametric kernel estimator²⁸ based on Gaussian kernels for the relative mass resolution. Table 3 provides a summary of all PDFs, for the branching fraction fit and the effective lifetime fit. The width of the Crystal Ball function is a conditional parameter with linear dependence on the dimuon mass resolution. All parameters, except for the signal yields, are fixed to values obtained from the MC simulation. Differences in the mass scale, studied with $J/\psi \rightarrow \mu^+\mu^-$ and $\Upsilon(1S) \rightarrow \mu^+\mu^-$ and interpolated to $m_{B_s^0}$, are taken into account by shifting the MC mass distributions. The difference in the mass resolution between data and MC simulation has an effect of less than 0.2% on the final results and is neglected.

The combinatorial background is modeled with a nonnegative Bernstein polynomial of the first degree (basis polynomials $p_{0,1}(x) = 1 - x$ and $p_{1,1}(x) = x$ for $x \in [0, 1]$) with floating parameters in the fit. Using an exponential function instead changes the result by 2.3% (0.6%) for $B_s^0 \rightarrow \mu^+\mu^-$ ($B^0 \rightarrow \mu^+\mu^-$), which is included in the systematic uncertainties. The relative mass resolution is modeled with a kernel function, determined from the data invariant mass sideband.

The rare background components are grouped together according to the number of muons in the final state (zero, one, or two muons). Each group is the weighted sum of various components. The peaking background $B \rightarrow h^+h^-$ combines all modes with $h \in \{\pi, K, p\}$, while the semileptonic group $B \rightarrow h\mu\nu$ consists of $B^0 \rightarrow \pi^-\mu^+\nu$, $B_s^0 \rightarrow K^-\mu^+\nu$, and $\Lambda_b \rightarrow p\mu^-\bar{\nu}$. Finally, the group $B \rightarrow h\mu^+\mu^-$ includes $B^- \rightarrow \pi^-\mu^+\mu^-$ and $B^0 \rightarrow \mu^+\mu^-\pi^0$. The weights in the sum are the prod-

Table 3. Summary of the PDFs used for signal and background components in the unbinned maximum likelihood fit for (top part of the table) $\mathcal{B}(B_s^0 \rightarrow \mu^+ \mu^-)$ and $\mathcal{B}(B^0 \rightarrow \mu^+ \mu^-)$ as well as for (bottom part of the table) $\tau_{\mu^+ \mu^-}$. Parameters that are floated in the fit are explicitly indicated (normalizations N , $\tau_{\mu^+ \mu^-}$, and background polynomial parameters p_0 and p_1 and lifetime t_0). The other parameters, fixed to the values obtained in MC simulation, are not shown explicitly. Functions whose parameters are determined (and fixed) in the sideband are indicated with /SB. The function abbreviations are as follows: Crystall-Ball (CB), Gaussian (G), Gaussian kernel estimator (KEYS), Bernstein polynomial of first degree (BE), binary distribution (BD), exponential including resolution and efficiency modeling (Exp), and exponential including resolution (Exp'). The arrow $\rightarrow$ indicates that a component is absorbed into the entry to the right.

Variable	$B_s^0 \rightarrow \mu\mu$	$B^0 \rightarrow \mu\mu$	$B \rightarrow hh$	$B \rightarrow h\mu\nu$	$B \rightarrow h\mu\mu$	Combinatorial Background
$m_{\mu\mu}$	CB(N)	CB(N)	CB+G	KEYS	KEYS	BE(p_0, p_1)
$\sigma(m_{\mu\mu})/m_{\mu\mu}$	KEYS	KEYS	KEYS	KEYS	KEYS	KEYS/SB
$\mathcal{C}$	BD	BD	BD	BD	BD	BD/SB
$m_{\mu\mu}$	CB(N)	$\rightarrow$	CB+G	$\rightarrow$	G	BE(p_0, p_1)
t	Exp($N, \tau_{\mu^+ \mu^-}$)	$\rightarrow$	Exp	$\rightarrow$	Exp	Exp'(N, t_0)

uct of the misidentification probabilities of each hadron (depending on charge, $p_\perp$, and η), the decay mode branching fraction, the analysis efficiency, and the trigger efficiency. Using Eq. 2 with the known branching fractions and the measured $B^+ \rightarrow J/\psi K^+$ normalization yield, it is possible to predict absolutely the expected yield per decay mode. To account for the missing components in the two groups with one or two muons, a common scaling factor is applied such that their sum plus the combinatorial background, extrapolated from the sideband $5.45 < m_{\mu\mu} < 5.9$ GeV, matches the event yield in data in the mass region $4.9 < m_{\mu\mu} < 5.2$ GeV. The trigger efficiency for these modes cannot be determined easily because of possible correlations between the two (fake) muons and the very limited sample size of the MC simulation where 1–2 hadrons are misidentified as fake muons. Studies based on samples with alternative muon identification algorithms indicate that the trigger efficiency for $B \rightarrow h^+ h^-$ is $\approx 50\%$ of the signal trigger efficiency while for the other groups (with at least one muon) it is at the same level as for the signal. A relative systematic error of 100% is assigned to the $B \rightarrow h^+ h^-$ trigger efficiency and this dominates the overall systematic uncertainty of the $B \rightarrow h^+ h^-$ yield. For $B \rightarrow h\mu\nu$ and $B \rightarrow h\mu^+ \mu^-$ the systematic error on the yield is about 15%. In the fit, the rare background yields are constrained to the expectations within these uncertainties.

The only parameters of interest in the fit are $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-)$ and $\mathcal{B}(B^0 \rightarrow \mu^+ \mu^-)$. All other parameters are nuisance parameters and are subject to Gaussian constraints, except for the rare background yields where log-normal priors are used as constraints. In Fig. 8 (left) the mass distribution of the high BDT categories is shown. The $B_s^0 \rightarrow \mu^+ \mu^-$ signal is clearly visible. As a consequence of the substantially improved muon identification algorithm, the peaking background is virtually invisible—a significant improvement compared to Ref. 12. The result of the fit to the data in the 14 categories, as defined in Table 2 (left part), is

$$\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-) = [2.9 \pm 0.7 (\text{exp}) \pm 0.2 (\text{frag})] \times 10^{-9}, \quad (5)$$

where the large statistical ($\pm 0.6 \times 10^{-9}$) and small systematical ($\pm 0.3 \times 10^{-9}$) errors are combined into one experimental uncertainty, and the second error is due to the f_s/f_u uncertainty. Using Wilks' theorem,²⁹ the observed (expected) significance amounts to 5.6σ (6.5σ). The fit likelihood contours are shown in Fig. 8 (right); the correlation between the two branching fractions is -0.181 . Summing over all BDT categories, a total $B_s^0 \rightarrow \mu^+\mu^-$ signal yield of 61^{+15}_{-13} events is observed, with an average $p_\perp^B$ of 17.2 GeV . For $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+\mu^-)$ the systematic error is dominated by the efficiency difference between data and MC simulation and by the pileup dependent effects of the HLT tracking in Run 2. All other components of the systematic error are much less important.

The fit also determines $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-) = (0.8^{+1.4}_{-1.3}) \times 10^{-10}$ with an observed (expected) significance of 0.6σ (0.8σ). Because no significant result was expected in this case, the primary result here is $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-) < 3.6 \times 10^{-10}$ (3.1×10^{-10}) at 95(90)% confidence level, using the CL_s method^{30,31} with the standard LHC-type profiled likelihood. The corresponding expected upper limit is $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-) < 3.0 \times 10^{-10}$ (2.4×10^{-10}), assuming no signal. The systematic error for $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-)$ is similar as for $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+\mu^-)$, with the exception of the rare background yields. However, given the precision of the upper limits quoted, this difference has no numerical impact.

Fig. 8. (left) Invariant dimuon mass distribution for the high BDT categories (cf. Table 2), with the fit projections overlaid. The peaking background is virtually invisible. (right) Likelihood contours of the branching fraction fit for $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+\mu^-)$ and $\mathcal{B}(B^0 \rightarrow \mu^+\mu^-)$, with the best-fit value (black cross) and the SM expectation (red solid square).

For the determination of the $B_s^0 \rightarrow \mu^+\mu^-$ effective lifetime $\tau_{\mu^+\mu^-}$, two independent fit frameworks were established, a 2D unbinned maximum likelihood fit to the invariant mass and decay time distributions and a 1D binned maximum likelihood fit to the decay time distribution where the background is subtracted with the *sPlot* method.³² Prior to unblinding the data, the former was chosen as the

primary method based on its better median expected performance. The fits are performed for the decay time range $1 < t < 11$ ps in eight BDT categories provided in Table 2. The decay time restriction is motivated by the very low selection efficiency at smaller decay times t , because of the $\ell_{3D}/\sigma(\ell_{3D})$ and isolation requirements, while the efficiency at large t is strongly decreasing in Run 2, due to HLT requirements on the muon impact parameter. Concerns regarding the fit stability, given the small expected number of signal events, motivated using eight instead of 14 BDT categories.

The PDFs for the invariant mass of the 2D unbinned maximum likelihood fit are very similar to those in the branching fraction fits with the notable exception that $B^0 \rightarrow \mu^+ \mu^-$ is included in the peaking background. In the decay time PDFs, the exponential functions are convolved with Gaussian functions to account for detector resolution. The decay time efficiency is included in all fit components except for the combinatorial background, because its PDF is modeled from data directly. In the fit, the $B_s^0 \rightarrow \mu^+ \mu^-$ effective lifetime $\tau_{\mu^+ \mu^-}$, the signal yield, and the parameters of the combinatorial background are floated (cf. Table 3 for a summary). All other parameters are constrained or fixed to the MC simulation values. There is no common $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-)$ constraint for the $B_s^0 \rightarrow \mu^+ \mu^-$ signal yields in the eight BDT categories. The result of the fit is

$$\tau_{\mu^+ \mu^-}(B_s^0 \rightarrow \mu^+ \mu^-) = 1.70_{-0.44}^{+0.61} \text{ ps}, \quad (6)$$

where the error combines the large statistical ($_{-0.43}^{+0.60}$ ps) and small systematic (± 0.09 ps) uncertainty. In Fig. 9 (left) the decay time distribution is shown, together with the fit results overlaid. The observed errors are about one root-mean-square deviation larger than expected ($_{-0.30}^{+0.39}$ ps). This is attributed to fluctuations in the small sample size.

The second determination of $\tau_{\mu^+ \mu^-}$ uses the complete model of the branching fraction fit to determine the *sPlot* weights. An exponential function, modified to include the channel-dependent resolution and efficiency effects, is fit to the *sPlot* distribution. Special care is applied to determine asymmetric uncertainties and to reduce the bias due to the large bin widths. The fit yields $\tau_{\mu^+ \mu^-} = 1.55_{-0.33}^{+0.52}$ ps, where a fit bias of $+0.09$ ps has been corrected for. This bias, together with the dependence on the Run 2 data-taking periods, is the largest systematic error in this approach. The two determinations of the effective lifetime are consistent with each other.

6. Conclusions

The CMS experiment has analyzed rare leptonic $B \rightarrow \mu^+ \mu^-$ decays with 61 fb^{-1} of data collected in LHC Run 1 and 2016. The decay $B_s^0 \rightarrow \mu^+ \mu^-$ is observed with a significance of 5.6σ (6.5σ expected). Its branching fraction is measured to be $\bar{\mathcal{B}}(B_s^0 \rightarrow \mu^+ \mu^-) = [2.9 \pm 0.7 (\text{exp}) \pm 0.2 (\text{frag})] \times 10^{-9}$ and the $B_s^0 \rightarrow \mu^+ \mu^-$ effective lifetime is determined to be $\tau_{\mu^+ \mu^-} = 1.70_{-0.44}^{+0.61}$ ps. Both results are limited by the

Fig. 9. Decay time distributions for (left) the 2D unbinned maximum likelihood fit and (right) the *sPlot* fit approach. The data combine all events in the eight BDT categories of Table 2 (right part). In the left plot, the total fit is shown by the solid line, the signal by the single-hatched distribution, and the different background components by the broken lines and the cross-hatched distributions. In the right plot, only the signal component is shown as the background has been subtracted beforehand with the *sPlot* method (see text for details).

small signal sample. No evidence is found for $B^0 \rightarrow \mu^+ \mu^-$ and an upper limit of $\mathcal{B}(B^0 \rightarrow \mu^+ \mu^-) < 3.6 \times 10^{-10}$ (at 95% confidence level) is determined. These results are consistent with the SM. The branching fraction results are also consistent with the previous CMS analysis¹² when restricting this analysis to the Run 1 dataset.

Acknowledgments

It is a pleasure to acknowledge the common effort and discussions on $B \rightarrow \mu^+ \mu^-$ decays over the past 15 years with many colleagues of the CMS collaboration. Countless stimulating and instructive discussions with Andreas Crivellin are very much appreciated.

References

1. C. Bobeth, M. Gorbahn, T. Hermann, M. Misiak, E. Stamou and M. Steinhauser, *Phys. Rev. Lett.* **112**, 101801 (2014), [arXiv:1311.0903 \[hep-ph\]](#).
2. T. Hermann, M. Misiak and M. Steinhauser, *JHEP* **12**, 097 (2013), [arXiv:1311.1347 \[hep-ph\]](#).
3. C. Bobeth, M. Gorbahn and E. Stamou, *Phys. Rev. D* **89**, 034023 (2014), [arXiv:1311.1348 \[hep-ph\]](#).
4. M. Beneke, C. Bobeth and R. Szafron, *Phys. Rev. Lett.* **120**, 011801 (2018), [arXiv:1708.09152 \[hep-ph\]](#).
5. M. Beneke, C. Bobeth and R. Szafron, *JHEP* **10**, 232 (2019), [arXiv:1908.07011 \[hep-ph\]](#).
6. N. Cabibbo, *Phys. Rev. Lett.* **10**, 531 (1963).
7. M. Kobayashi and T. Maskawa, *Prog. Theor. Phys.* **49**, 652 (1973).

8. Particle Data Group, M. Tanabashi *et al.*, *Phys. Rev. D* **98**, 030001 (2018).
9. K. De Bruyn, R. Fleischer, R. Knegjens, P. Koppenburg, M. Merk and N. Tuning, *Phys. Rev. D* **86**, 014027 (2012), [arXiv:1204.1735 \[hep-ph\]](#).
10. Flavour Lattice Averaging Group Collaboration, S. Aoki *et al.*, *Eur. Phys. J. C* **80**, 113 (2020), [arXiv:1902.08191 \[hep-lat\]](#).
11. ATLAS Collaboration, M. Aaboud *et al.*, *JHEP* **04**, 098 (2019), [arXiv:1812.03017 \[hep-ex\]](#).
12. CMS Collaboration, S. Chatrchyan *et al.*, *Phys. Rev. Lett.* **111**, 101804 (2013), [arXiv:1307.5025 \[hep-ex\]](#).
13. LHCb Collaboration, R. Aaij *et al.*, *Phys. Rev. Lett.* **111**, 101805 (2013), [arXiv:1307.5024 \[hep-ex\]](#).
14. LHCb Collaboration, R. Aaij *et al.*, *Phys. Rev. Lett.* **118**, 191801 (2017), [arXiv:1703.05747 \[hep-ex\]](#).
15. CMS Collaboration, A. M. Sirunyan *et al.*, *JHEP* **04**, 188 (2020), [arXiv:1910.12127 \[hep-ex\]](#).
16. D0 Collaboration, V. Abazov *et al.*, *Phys. Rev. Lett.* **94**, 071802 (2005), [arXiv:hep-ex/0410039](#).
17. LHCb Collaboration, R. Aaij *et al.*, *Phys. Rev. Lett.* **124**, 122002 (2020), [arXiv:1910.09934 \[hep-ex\]](#).
18. ATLAS Collaboration, G. Aad *et al.*, *Phys. Rev. Lett.* **115**, 262001 (2015), [arXiv:1507.08925 \[hep-ex\]](#).
19. LHCb Collaboration, R. Aaij *et al.*, *Phys. Rev. D* **100**, 031102 (2019), [arXiv:1902.06794 \[hep-ex\]](#).
20. CMS Collaboration, S. Chatrchyan *et al.*, *JINST* **3**, S08004 (2008).
21. CMS Collaboration, V. Khachatryan *et al.*, *Eur. Phys. J. C* **70**, 1165 (2010), [arXiv:1007.1988 \[physics.ins-det\]](#).
22. CMS Collaboration, CMS Collaboration, *Tracking POG results for pion efficiency with the D^{*+} meson using data from 2016 and 2017*, CMS Detector Performance Note CMS-DP-2018-050 (2018).
23. CMS Collaboration, S. Chatrchyan *et al.*, *JINST* **7**, P10002 (2012), [arXiv:1206.4071 \[physics.ins-det\]](#).
24. CMS Collaboration, A. M. Sirunyan *et al.*, *JINST* **13**, P06015 (2018), [arXiv:1804.04528 \[physics.ins-det\]](#).
25. CMS Collaboration, V. Khachatryan *et al.*, *JINST* **12**, P01020 (2017), [arXiv:1609.02366 \[physics.ins-det\]](#).
26. H. Voss, A. Höcker, J. Stelzer and F. Tegenfeldt, *TMVA, the toolkit for multivariate data analysis with ROOT*, in *XIth International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT)*, (2007). p. 40. [arXiv:physics/0703039 \[PoS\(ACAT\)040\]](#).
27. M. J. Oreglia, *A study of the reactions $\psi' \rightarrow \gamma \gamma \psi$* , PhD thesis, Stanford University, (1980). SLAC-R-236, UMI-81-08973. See Appendix D.
28. K. S. Cranmer, *Comput. Phys. Commun.* **136**, 198 (2001), [arXiv:hep-ex/0011057 \[hep-ex\]](#).
29. S. S. Wilks, *Annals Math. Statist.* **9**, 60 (1938).
30. A. L. Read, *J. Phys. G* **28**, 2693 (2002).
31. T. Junk, *Nucl. Instrum. Meth. A* **434**, 435 (1999), [arXiv:hep-ex/9902006](#).
32. M. Pivk and F. R. Le Diberder, *Nucl. Instrum. Meth. A* **555**, 356 (2005), [arXiv:physics/0402083 \[physics.data-an\]](#).