Published for SISSA by Springer

Received: June 30, 2022
Revised: November 11, 2022
Accepted: January 13, 2023
Published: February 3, 2023

Optimizing trigger-level track reconstruction for sensitivity to exotic signatures

K.F. Di Petrillo,[a,b] **J.N. Farr,**[c] **C. Guo,**[d] **T.R. Holmes,**[c] **J. Nelson**[e] **and K. Pachal**[f]

[a] *Particle Physics Division, Fermi National Accelerator Laboratory, Kirk and Pine Street, Batavia, IL, U.S.A.*
[b] *Department of Physics and Enrico Fermi Institute, The University of Chicago, 5720 South Ellis Avenue, Chicago, IL, U.S.A.*
[c] *Department of Physics & Astronomy, University of Tennessee, 1408 Circle Drive, Knoxville, TN, U.S.A.*
[d] *Department of Physics, Illinois Institute of Technology, 10 West 35th Street, Chicago, IL, U.S.A.*
[e] *Department of Physics, Brown University, 182 Hope Street, Providence, RI, U.S.A.*
[f] *TRIUMF, 4004 Wesbrook Mall, Vancouver, BC, Canada*

E-mail: karri@uchicago.edu, jfarr7@vols.utk.edu, cguo19@hawk.iit.edu, tholmes@utk.edu, Jessica_n_Nelson@brown.edu, kpachal@triumf.ca

Abstract: Many compelling beyond the Standard Model scenarios predict signals that result in unconventional charged particle trajectories. Signatures for which unusual tracks are the most conspicuous feature of the event pose significant challenges for experiments at the Large Hadron Collider (LHC), particularly for the trigger. This article presents a study of track-based triggers for a representative set of long-lived and unconventional signatures at the upcoming High Luminosity LHC, as well as resulting recommendations for the target parameters of a hardware-based tracking system. Scenarios studied include large multiplicities of low-p_{T} tracks produced in a soft-unclustered-energy-pattern model, displaced leptons and anomalous prompt tracks predicted in a Supersymmetry model with long-lived staus, and displaced hadrons predicted in a Higgs portal scenario with long-lived scalars.

Keywords: Beyond Standard Model, Exotics, Hadron-Hadron Scattering, Lifetime

ArXiv ePrint: 2211.05720

https://doi.org/10.1007/JHEP02(2023)034

JHEP02(2023)034

Contents

JHEP02(2023)034

1 Introduction

Decades of experimental tests have demonstrated the predictive power of the Standard Model (SM). However, there are a myriad of phenomena it does not explain, including dark matter, baryon asymmetry, naturalness, and more. Many beyond the Standard Model (BSM) scenarios that address these inefficiencies predict unconventional track signatures that could be observed at high energy colliders such as the Large Hadron Collider (LHC). For example, long-lived BSM particles can travel a measurable distance before decaying, resulting in displaced track or anomalous prompt track signatures [1]. Alternatively, models with a strongly coupled hidden valley can result in large multiplicities of soft charged particles [2]. These unconventional track signatures pose tremendous challenges for general purpose experiments at the LHC, particularly for the trigger.

Due to bandwidth and storage constraints, the ATLAS and CMS experiments both rely on real-time reconstruction to select interesting collisions before detector data can be read out and stored. The system that performs this task is called the trigger. A particular challenge for trigger systems at the LHC is differentiating signal processes from a large background of low energy proton-proton interactions, including the multiple interactions that occur per event, referred to as pile-up. The trigger is primarily designed to select events with high-p_{T} jets, photons, or leptons that originate directly from a pp collision of interest, or events characterized by large missing transverse momentum. These prompt signatures are easily identified in the first, hardware-based, stage of the trigger, which makes use of calorimeter and muon system information. Limited tracking information is only available at later stages of the trigger decision.

Unfortunately, standard triggers often result in low trigger efficiency for unconventional BSM signatures. If a BSM particle decays at a distance and produces jets or leptons, the decay products will not point back to the pp collision. Decay products can be misidentified as pile-up or fail standard quality criteria. Charged BSM particles which travel a long distance before decaying can result in anomalous signals in the detector that can be confused with noise. Models with many low energy charged particles can be difficult to differentiate from pile-up.

When searching for these unconventional BSM signatures, the most distinctive features of signal events are located in the tracker. Dedicated track-based triggers are necessary to access these challenging scenarios, and would require track information is available at an early stage of the trigger decision. A long history of hardware-level or intermediate stage tracking in various experiments provides a guideline for what rates and performance can be realistically anticipated at ATLAS and CMS. While possible implementations and track reconstruction algorithms are discussed elsewhere [3], including unconventional tracking, there has been no explicit study of how best to allocate the limited resources available to optimize sensitivity across a range of signatures thus far. This paper addresses that need.

1.1 History and current state of track-triggers

High energy physics experiments typically use multi-step trigger systems to select events which are saved for analysis. The first stage, Level 1 (L1), performs an initial reduction using simple decisions often implemented in custom hardware. The following stages (collectively, the High Level Trigger or HLT) have a longer latency budget and use standard software routines to further reduce the event rate.

The usage of track information in the trigger at previous experiments has prioritized the identification of Standard Model processes. Searches for long-lived or displaced signatures have typically used triggers based on calorimeter or muon detector information [4, 5]. The CDF and D0 experiments, which can be considered the predecessors to ATLAS and CMS, operated in conditions with a 2.5 MHz event rate and matched tracking detector hits to predefined patterns or equations for fast track reconstruction at L1 [6, 7]. Experiments specializing in b physics have long prioritized hardware-level track-triggers. BABAR [8], HERA-B [9], and Belle-II [10] all included track reconstruction at L1, though in some cases restricted to regions of interest identified by other event activity. The L1 track reconstruction

JHEP02(2023)034

in these three cases constrains tracks to originate from the z axis, diminishing sensitivity to displaced decays. Track reconstruction at HLT can be more sophisticated. LHCb contends with an L1 event rate of 40 MHz, like ATLAS and CMS but with much lower instantaneous luminosity. In Run 2 LHCb did not have any tracking at L1, but has moved to an HLT-only design for Run 3 that will permit full event tracking with high quality on each event [11, 12].

ATLAS and CMS currently use similar two-step trigger systems. In both experiments, the L1 stage reduces the event rate from 40 MHz to 100 kHz within a latency of $\sim 3\,\mu s$ while the second (HLT) stage reduces the event rate from 100 kHz to ~ 1 kHz, with ~ 200 ms on average to make a decision [13–15]. Due to bandwidth and latency constraints, tracking information is only produced at the HLT, where it is the most CPU-intensive reconstruction process. There, limited track reconstruction routines are run in software, balancing the amount of information necessary to make a decision with the amount of time needed to reconstruct the event. For example in lepton triggers, track reconstruction is often only performed in small regions of interest that were identified by the calorimeter or muon system at L1. Full detector tracking is possible with fast algorithms and a reduced event rate. These fast algorithms result in reduced efficiency across wide ranges of track transverse momentum, p_T, and transverse impact parameter, d_0 with respect to offline reconstruction.

The upcoming High Luminosity LHC (HL-LHC) will see significant upgrades to the track-based triggering capacity of both experiments. ATLAS and CMS both plan to install new tracking detectors and completely overhaul their trigger designs. In CMS, tracker p_T-modules will enable sufficient bandwidth reduction such that tracks with $p_T > 2\,\text{GeV}$ can be reconstructed at L1 [16, 17]. While the ATLAS tracker will not participate in the hardware-level trigger decision at the HL-LHC, it will participate at the HLT, where tracking will be run on a selected subset of events within the full HLT input rate of 1 MHz. Under the baseline plan, track reconstruction will be performed at the HLT with a CPU based processing farm potentially complemented by commodity accelerator hardware [18].

With additional effort, it may be possible to extend these hardware track-triggers to include a range of long-lived or unconventional signatures [19, 20] by allocating some fraction of the available tracking resources to unconventional track reconstruction and a fraction of the rate to triggers based on these tracks. Tracking performance is limited by detector readout rates, the hardware-based tracking technology used, output rate, and latency constraints. In general, increasing the maximum curvature, displacement, or target efficiency of reconstructed tracks increases complexity, rate, latency, or a combination thereof. In order to maintain a constant latency but extend its d_0 range, a track-trigger could increase its p_T threshold or decrease its target efficiency. Depending on the track-trigger design, it may be possible to alter these requirements independently for prompt and displaced tracks, resulting in a highly efficient prompt track-trigger with a low p_T threshold, while also exploring a large range of displacements.

1.2 Expanding track-triggers for future experiments

There has been a recent proliferation of studies on the efficacy of track-based triggers for long-lived particles and other exotic signatures, including rare Higgs decays [19], soft-unclustered-energy-patterns (SUEPs) [2], and generic displaced signatures [21, 22]. ATLAS demonstrated

JHEP02(2023)034

that a hardware-based track finder algorithm could extend its impact parameter range from $|d_0| < 2$ mm to $|d_0| < 10$ mm without significantly degrading efficiency for prompt tracks [20]. Displaced hardware-based tracking has been extensively studied for the HL-LHC on CMS [19].

However, because the signatures of these models are so varied, it is challenging to translate individual studies into an optimal design for a future trigger. In particular, for a realistic rate- and latency-limited trigger, it is unclear which efficiency trade-offs are most beneficial across a range of signatures: is it better to expand d_0 coverage, reduce p_{T} thresholds, or maximize overall tracking efficiency? The answer to this question depends on the design of the trigger being considered, but it also depends on what kind of exotic signatures are being targeted. For hadronic decays, low p_{T} thresholds are important, but high efficiency may not be needed. For leptonic decays the reverse is true.

This study aims to provide a recommendation on how to prioritize a track-trigger's p_{T} and d_0 coverage in order to maximize sensitivity to a range of unconventional signatures. This aim is based on the idea that most tracking algorithms are configured to reconstruct particles within a restricted range of tracking parameter space to meet bandwidth and latency constraints, and that increasing coverage in one tracking parameter will require decreasing coverage elsewhere. This study will not discuss specific tracking algorithms, as these vary widely between experiments and are discussed elsewhere. The overall efficiency to reconstruct the unconventional tracks expected in a variety of BSM models will be presented as a function of track parameters and simple trigger concepts, which can act as a guide for future trigger design. Because the practicality of such triggers depends on the rates at which they can operate, this study will also show expected Standard Model backgrounds for possible triggers based on these tracking parameters, demonstrating their feasibility.

2 Benchmark model performance

This work studies a representative set of unconventional signatures motivated by three beyond the Standard Mode scenarios. Figure 1 shows Feynman diagrams corresponding to the benchmark models considered. Events with a large multiplicity of prompt low p_{T} tracks are motivated by a soft-unclustered-energy-pattern (SUEP) model. Displaced decays to hadrons are motivated by a model with long-lived scalars coupling to the Higgs boson. A gauge mediated Supersymmetry breaking scenario with direct production of long-lived staus results in displaced leptons as well as heavy (meta)stable charged particles (HSCPs) which can be detected directly.

Together, these four signatures span a wide range of possible track momenta and displacement, as well as event-level track multiplicity. In order to directly compare these vastly different signatures, a common strategy is used to evaluate the acceptance and efficiency of possible track-trigger configurations. The following sections outline this overall analysis strategy as well as the specific studies performed for each model.

JHEP02(2023)034

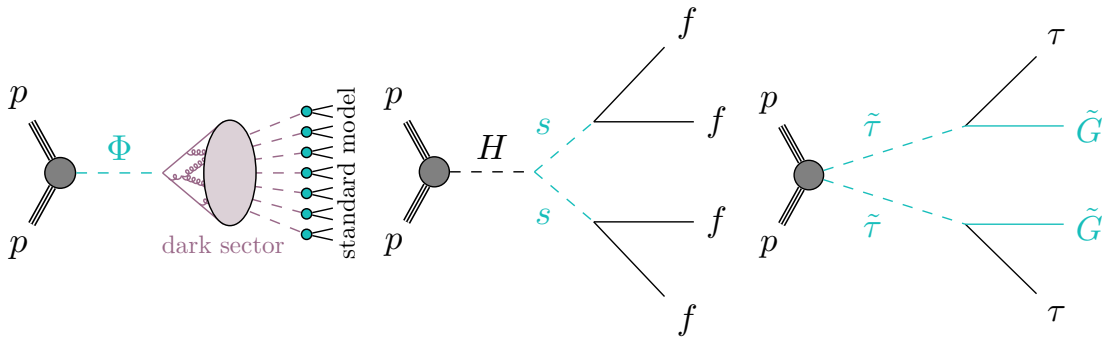

Figure 1. Feynman diagrams for benchmark processes. From left to right: SUEPs, Higgs portal, and long-lived staus.

2.1 Analysis strategy

All efficiency studies are performed assuming HL-LHC conditions, using truth-level simulation of pp collisions at a center of mass energy of $\sqrt{s} = 14\,\text{TeV}$. Charged particles considered for analysis are stable (or metastable) decay products of short-lived BSM particles, except for the HSCP case in which direct detection of the metastable particle is considered. For simplicity, there is no simulation of detector or material interactions, and pile-up collisions are not included in signal simulation. Discussion of how signal efficiencies and background rates are influenced by pile-up is included in section 3.

The acceptance and efficiency of possible trigger configurations are evaluated for each signature. First, baseline criteria are defined based on potential detector specifications (e.g. η coverage, tracker radius). Only stable charged particles satisfying these criteria are considered to be in acceptance. This geometric acceptance simply asks whether particles could be reconstructed by an offline tracking algorithm with the specified detector layout, and is independent of what could then be defined by a track-trigger's configuration. Certain effects of adjusting the detector layout are explored for the HSCP signature. For other models the detector layout is kept fixed. The acceptance for a signal point is then the fraction of events with at least n particles in the final state that meet the tracking criteria, where n is also model dependent.

For displaced signatures, only tracks with $|d_0| > 1$ mm are counted in the acceptance criteria. This displacement requirement separates tracks which would likely be picked up by a prompt track reconstruction algorithm from those which require a displaced tracking algorithm, and also defines a region in which SM backgrounds are small.

Efficiency is then defined to capture the effects of track reconstruction criteria which could be set at the track-trigger level rather than enforced by the physical detector. The two most important criteria studied are the minimum p_T and maximum $|d_0|$ of tracks which would be reconstructed by the track-trigger. These factors have different effects in different BSM models, depending on whether the final state tracks are prompt or displaced, what the average decay angles of long-lived particles may be, and how much momentum the final state particles carry. To study the interplay of these effects, for a fixed acceptance, the p_T and d_0 thresholds are varied, and the fraction of surviving tracks is measured. An

JHEP02(2023)034

event-level efficiency is then defined, as for the acceptance, based on whether at least n particles in the final state pass these thresholds.

In realistic tracking detectors, although the reconstruction efficiency for prompt isolated tracks is very close to 100% independent of track p_{T}, the efficiency tends to decrease roughly linearly with increasing $|d_0|$ [23]. To emulate this effect in the following studies, the likelihood of reconstructing a given track is set not as a binary by the d_0 threshold (e.g. all tracks up to $|d_0| = 5$ mm are reconstructed; all above 5 mm are not reconstructed), but instead as a linearly decreasing function. The likelihood of reconstructing the track begins at unity for $d_0 = 0$ and decreases linearly with increasing $|d_0|$ until it reaches zero at the specified $|d_0|$ threshold. This linearly falling assumption is a good match for the behavior observed in offline track reconstruction in ATLAS, though the maximum $|d_0|$ values explored here for trigger-level reconstruction are more conservative.

The efficiency definition considered in this section is designed to test if a hypothetical track-trigger can reconstruct a sufficient number of unconventional tracks such that a dedicated trigger could be defined. For all BSM scenarios considered, background rates in HL-LHC conditions for dedicated triggers are studied in section 3. For the heavy stable charged particle scenario and displaced lepton scenarios, simple track counting requirements will already result in negligible backgrounds. For the SUEP case, a large QCD background has the potential to mimic signal if the track-trigger is unable to differentiate between pile-up vertices. For the lower p_{T} Higgs portal scenario, counting displaced tracks at large pile-up is not sufficient to reduce background from Standard Model long-lived particles and randomly crossing displaced tracks produced in high pile-up conditions, and more complex selections are studied.

2.2 SUEPs

Motivated in certain hidden valley models, SUEPs represent a worst case scenario for triggers at hadron colliders [2, 24]. In SUEP events, low-p_{T} final state particles are not collimated into Standard Model-like jets. The most conspicuous feature is a large multiplicity of diffuse low p_{T} tracks. This scenario is very difficult for the trigger to distinguish from pile-up, unless the many soft tracks can be reconstructed and associated to a pp collision of interest. As a result, there are currently no constraints on SUEP scenarios from the LHC.

In SUEP models, a dark SU(N) sector is connected to the Standard Model by a mediator with a mass much greater than the mass splittings amongst the dark sector particles. Energy from the production of the mediator is distributed amongst a large multiplicity of light dark hadrons, which decay back to SM particles. The 't Hooft coupling of the dark sector, $\lambda_D \equiv \alpha_D N_{C_D}$, where α_D is the dark coupling constant and N_{C_D} is the number of colors, determines the angular emission of partons during showering. In SUEP scenarios this coupling is assumed to be large, and dark mesons are isotropically distributed in the mediator's rest frame.

SUEP samples are generated in Pythia8 [25] using the custom plugin described in ref. [2]. Events are generated assuming a scalar mediator, Φ, with a range of masses: 125 GeV (matching the SM Higgs boson) and then 200 GeV to 1000 GeV in 200 GeV intervals.

JHEP02(2023)034

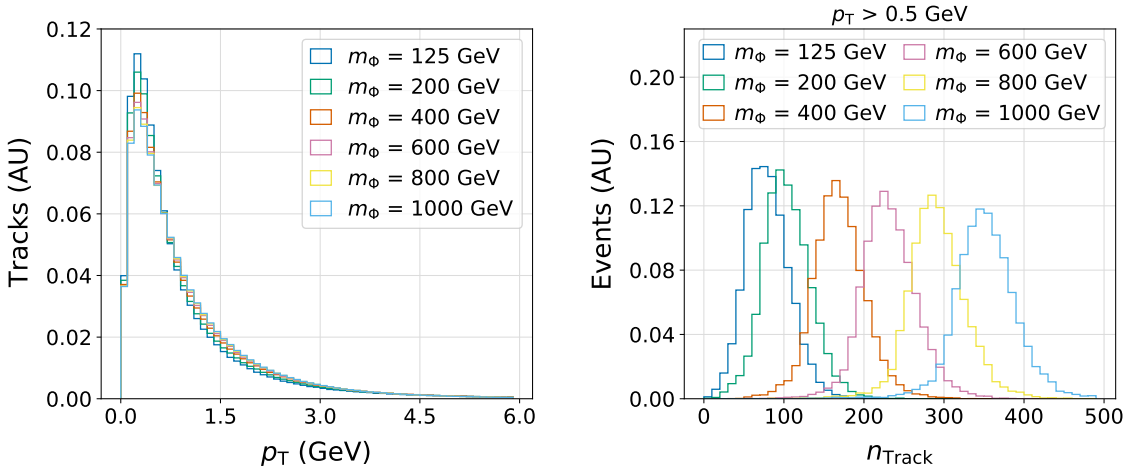

(a) Charged particle p_T

(b) Charged particle multiplicity

Figure 2. Distribution of all charged SM particle transverse momenta (a) and multiplicity of stable charged particles with $p_T > 0.5\,\text{GeV}$. (b) in SUEP events for a range of mediator masses.

The SUEP shower consists of a large multiplicity of a single flavor of dark mesons with mass $m_{\pi_d} = 1\,\text{GeV}$ and temperature $T = 1\,\text{GeV}$.

In principle, dark mesons may decay to a variety of final state SM particles. Possible decay modes include pairs of leptons, hadrons, photons, and light or heavy flavor quarks. Decay products may be also be displaced, or the final state may also include undetected dark matter particles. For simplicity, this study focuses on a scenario in which each dark meson decays promptly to a pair of light quarks, $d\bar{d}$.

Figure 2(a) shows charged particle transverse momenta for the SUEP benchmark models under consideration. Unlike most BSM scenarios, the p_T spectra of final state particles do not depend on the mediator mass, but on the dark meson mass and temperature. As a result, all generated momenta distributions are similar in shape, and peak sharply in the lowest p_T bins.

The number of final state particles depends on mass of the scalar mediator and dark meson mass, $n \sim m_\Phi / m_{\pi_d}$, as shown in figure 2(b). For a fixed mediator and dark meson mass, increasing the temperature reduces the number of final state particles and increases their mean momenta, as described in ref. [26].

Because all tracks from the model considered here are prompt, the selection strategy is relatively straightforward. The number of stable charged particles within tracker acceptance and above a given p_T threshold are counted per event. It is assumed that the track-trigger would have sufficient longitudinal impact parameter resolution to associate tracks to the correct primary vertex, and neutral final state particles are neglected.

Table 1 summarizes SUEP object-level acceptance and efficiency definitions. Tracks within $|\eta| < 2.5$ are considered to be within the detector acceptance, which is consistent with the η range of the CMS L1 track finder [27]. The event-level acceptance is 100% in all cases because of the large multiplicity of charged particles. The efficiency is defined as

JHEP02(2023)034

Variable	Requirement
Acceptance	
$\lvert\eta\rvert$	< 2.5
Efficiency	
p_{T}	$> 0.5, 1, 2\,\mathrm{GeV}$
n_{Track}	$> 100, 150, 200$

Table 1. Acceptance and efficiency requirements for a track-trigger targeting events with a high multiplicity of soft tracks.

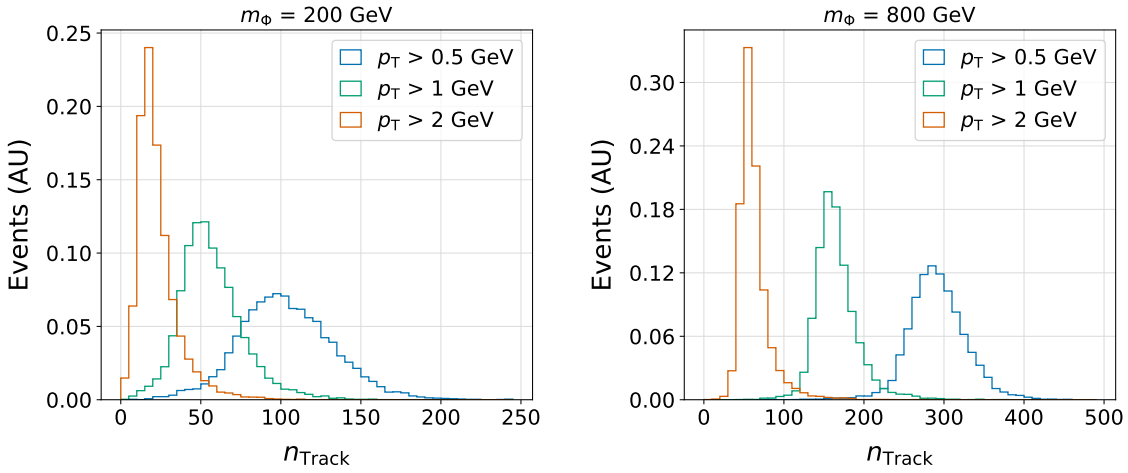

(a) (b)

Figure 3. Number of charged particles per event which pass the specified track p_{T} threshold, for mediators of mass 200 GeV (a) and 800 GeV (b).

the fraction of events where the number of tracks, n_{Track}, that pass the specified p_{T} and η requirements is above a certain threshold.

The most important parameter to consider is the p_{T} threshold at which track reconstruction begins. For SUEP signatures, signal to background discrimination decreases with increasing track p_{T} requirements, because track multiplicity is the most distinctive feature of these events. This study investigates a range of p_{T} thresholds, from $0.5 < p_{\mathrm{T}} < 2\,\mathrm{GeV}$. The lowest threshold represents the minimum charged particle p_{T} for most standard offline track reconstruction algorithms, aimed at charged particles that traverse a sufficient number of tracking detector layers [28, 29]. The highest p_{T} threshold examined, $p_{\mathrm{T}} > 2\,\mathrm{GeV}$, is based on the minimum L1 tracking threshold of 2 to 3 GeV estimated by CMS for their Phase 2 trigger upgrade [17]. For the temperature considered, a threshold above $p_{\mathrm{T}} = 2\,\mathrm{GeV}$, counting the n_{Track} is no longer expected to be a powerful discriminator. Figure 3 shows the number of reconstructed tracks in SUEP events satisfying different track p_{T} thresholds.

Efficiencies for all simulated mediator masses are shown as a function of the minimum n_{Track} requirement for three different p_{T} thresholds in figure 4. The efficiency increases

JHEP02(2023)034

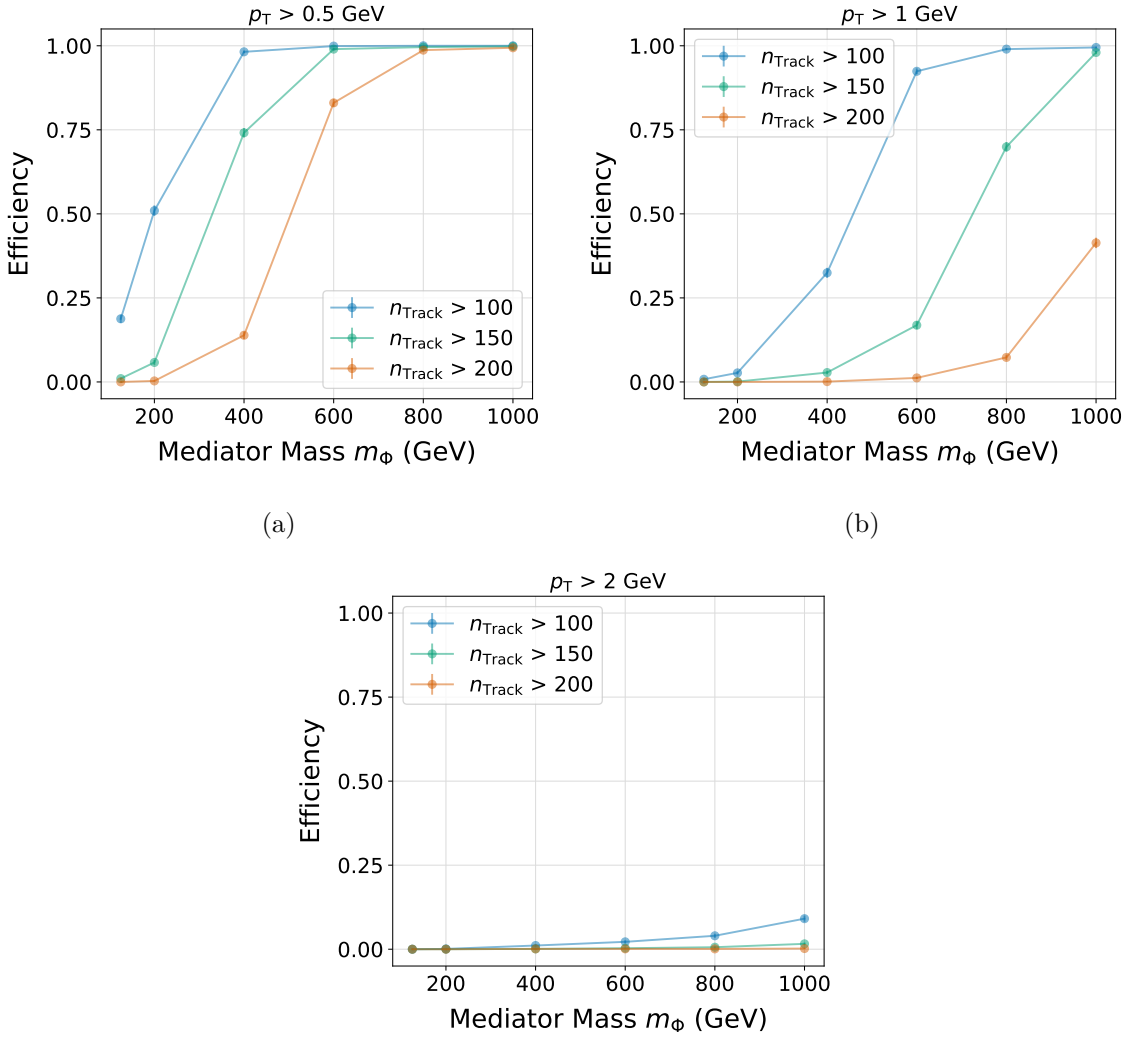

Figure 4. Efficiency, defined as number of events for which at least n_{Track} within $\eta < 2.5$ pass the track reconstruction p_T threshold, as a function of mediator mass for minimum track p_T of 0.5 GeV (a), 1 GeV (b), and 2 GeV (c).

with mediator mass, due to the larger number of tracks produced. The track p_T spectra do not change with mediator mass, and this effect is simply because a sufficient number of particles are reconstructed more frequently for heavy mediators. With a track reconstruction threshold of $p_T > 0.5$ GeV, efficiencies of 50% are obtained for mediators with masses above 200 GeV with a $n_{\text{Track}} > 100$ trigger, or for mediators with masses above 400 GeV with a $n_{\text{Track}} > 150$ trigger. If the track reconstruction threshold increases to $p_T > 1$ GeV, 50% efficiency is only possible for mediator masses above ~ 600 GeV. With a reconstruction threshold of $p_T > 2$ GeV, there is negligible efficiency for all signals with any n_{Track} threshold considered.

JHEP02(2023)034

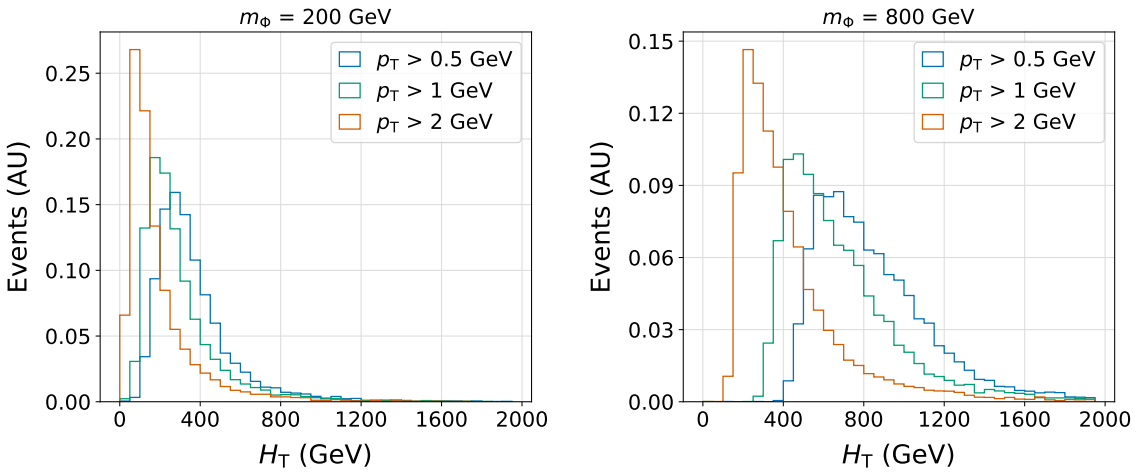

(a) (b)

Figure 5. H_T of charged particles passing the specified track p_T threshold, for mediators of mass 200 GeV (a) and 800 GeV (b).

For a track-trigger similar to the CMS HL-LHC upgrade, where $p_T > 2$ GeV is the baseline threshold, a trigger requiring a large track multiplicity is inefficient for most SUEP signal models. Other possible metrics for a SUEP trigger should be investigated, such as the scalar sum of jet energy, the scalar sum of track p_T (H_T, shown in figure 5), or event shapes such as isotropy or sphericity [30, 31]. At the LHC, SUEP final state particles are expected to be isotropically distributed in ϕ, but localized in η, resulting in a "belt of fire". This characteristic shape has been shown to be an efficient discriminator at the HLT [2]. The most optimal trigger available to the LHC experiments will likely use H_T or an event shape in combination with the number of reconstructed tracks.

2.3 Higgs portal to long-lived scalars

In this model, referred to here as the "Higgs portal", a SM Higgs boson decays to a pair of long-lived scalar particles. Each scalar subsequently decays to a pair of SM particles. Since the long-lived particle is not electromagnetically charged, no direct detection trigger strategies are available. The identifiable characteristic of this scenario is several relatively low-p_T displaced jets. The low-p_T nature of these jets presents a challenge for current L1 triggers because it is difficult for the calorimeter to distinguish these displaced jets from a large multi-jet background. Instead, this scenario can be targeted with a trigger that selects events with a number of displaced tracks.

The benchmark model considered in this paper explores the well-motivated possibility that the SM Higgs boson acts as a mediator to a hidden sector via mixing with a new neutral scalar [32–34]. A dark U(1) gauge group with a dark Higgs potential is introduced, leading to a total of two new particles, of which only the scalar has a mass low enough to be accessible at the LHC. This model serves as a good benchmark for a wide variety of Higgs portal scenarios, which will be a high priority target for the HL-LHC. If the only available

JHEP02(2023)034

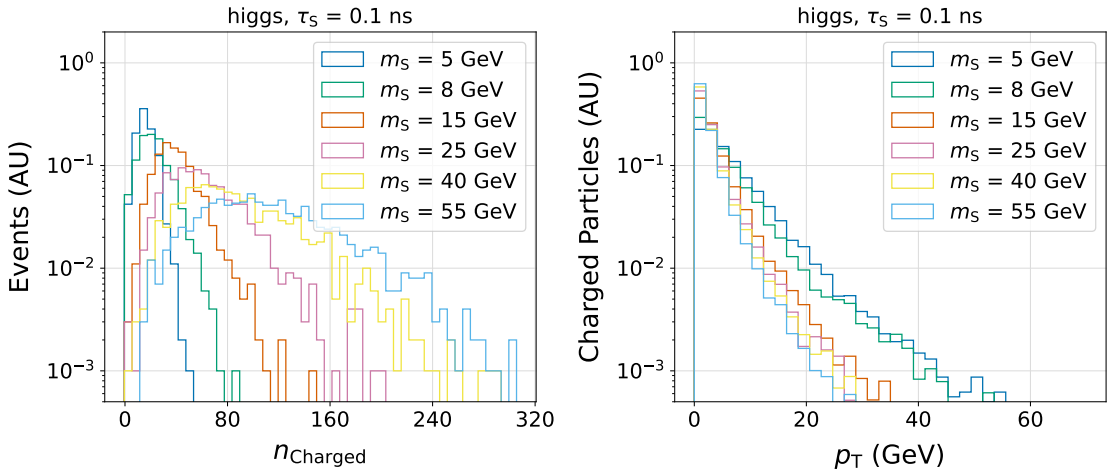

Figure 6. (Left) The number of displaced charged particles per event. Particles are required to be decay products of the long-lived scalar and have $p_T > 0.5$ GeV. (Right) Displaced charged particle p_T for each scalar mass.

decay channel for the dark scalar is via its (small) mixing with the Higgs, this process will be suppressed, and the scalar will travel a macroscopic distance before decaying in to SM particles. The resulting signature consists of low p_T displaced tracks arising from the point at which the scalar decayed. The light mass of the scalar also means these displaced tracks tend to have modest impact parameters, even when the scalar decays at a high radius. The branching ratio for the decay of the dark scalar follows that of the Higgs, restricted by its mass.

Higgs portal events are generated using MadGraph5_aMC@NLO 2.9.3 [35] and the SM + Dark Vector + Dark Higgs model described in refs. [36, 37]. The Higgs mixing parameter is taken to be 10^{-4} and the kinetic mixing parameter to be 10^{-10}. The dark Z mass is set to 1 TeV, and effectively decoupled. In the generated scenario the dark scalar, s, is the only relevant new particle and it decays only into Standard Model particles. The properties of the dark scalar are varied across the mass range m_s = 5, 8, 15, 25, 40, and 55 GeV and proper lifetime range 0.01, 0.1, and 1 ns. The scalar's branching ratio follows those of the Higgs, such that the scalars decay predominantly to $c\bar{c}$ and $\tau\bar{\tau}$ for $m_s = 5$ and 8 GeV and predominantly to $b\bar{b}$ for heavier m_s.

The number of displaced charged particles per event and their p_T is shown in figure 6. The low parent masses and fully hadronic final state result in a moderate number of low-p_T particles. However, there are a sufficient number of displaced tracks to make a track-based trigger feasible.

A baseline track-level acceptance for this scenario is defined using requirements described in table 2. In addition to minimum p_T and η requirements, displaced charged particles are required to be produced within a minimum radius and pass a minimum length requirement. These requirements ensure that displaced particles will traverse a sufficient number of tracker layers to be considered for reconstruction.

JHEP02(2023)034

Variable	Requirement
Acceptance	
L_{xy}	< 300 mm
L_{track}	> 200 mm
$\lvert d_0 \rvert$	≥ 1 mm
p_{T}	$> 0.5\,\mathrm{GeV}$
$\lvert \eta \rvert$	< 2.5
Efficiency	
p_{T}	$> 0.5, 1, 2, 5, 10\,\mathrm{GeV}$
$\lvert d_0 \rvert$	$< 10, 20, 50, 100$ mm

Table 2. Details of the displaced track acceptance and efficiency definitions for the Higgs portal model. The L_{xy} parameter refers to the maximum allowed production radius of the track, while the L_{track} variable refers the minimum track length.

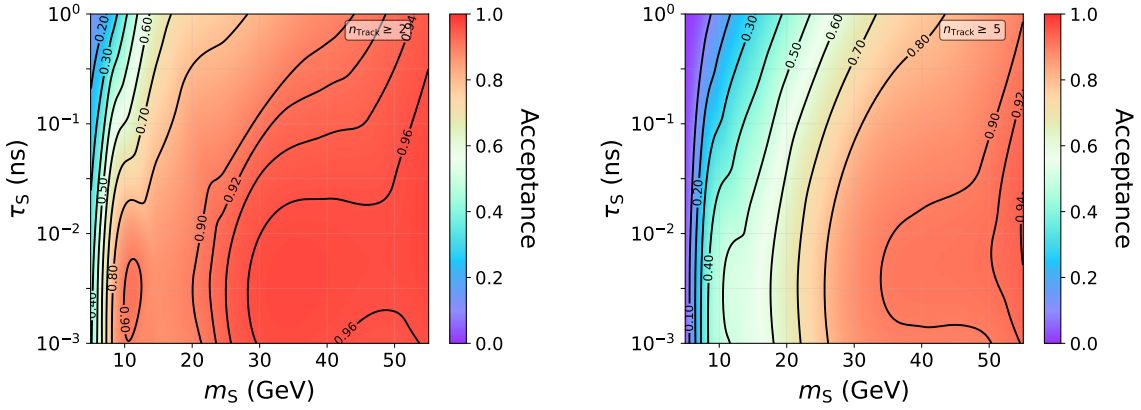

Figure 7. Acceptances for the Higgs portal scenario requiring at least two tracks (left) and least five tracks (right) passing the requirements listed in table 2.

For this model, two event-level acceptance definitions are considered due to the low momenta of the displaced tracks. The first requires a minimum of five tracks per event to pass the acceptance criteria defined in table 2, and the second requires least two tracks per event. Event-level acceptances for these two definitions are shown in figure 7.

Discontinuities in the event-level acceptances and efficiencies around scalar masses of $m_s = 10\,\mathrm{GeV}$ are a result of the sudden change in dominant decay modes, from pairs of charm quarks and tau leptons to $b\bar{b}$. The acceptance requirement has the largest impact at low masses, where there are fewer tracks likely to pass the minimum p_{T} threshold.

Track-level efficiency requirements corresponding to potential online tracking configurations are defined in table 2. The track-trigger parameters considered include a minimum p_{T} and maximum $\lvert d_0 \rvert$ requirement, where the exact values are varied to understand the impact each selection has on the efficiency. Event-level efficiencies are defined as the fraction of events in acceptance which also have at least n_{Track} passing efficiency requirements.

JHEP02(2023)034

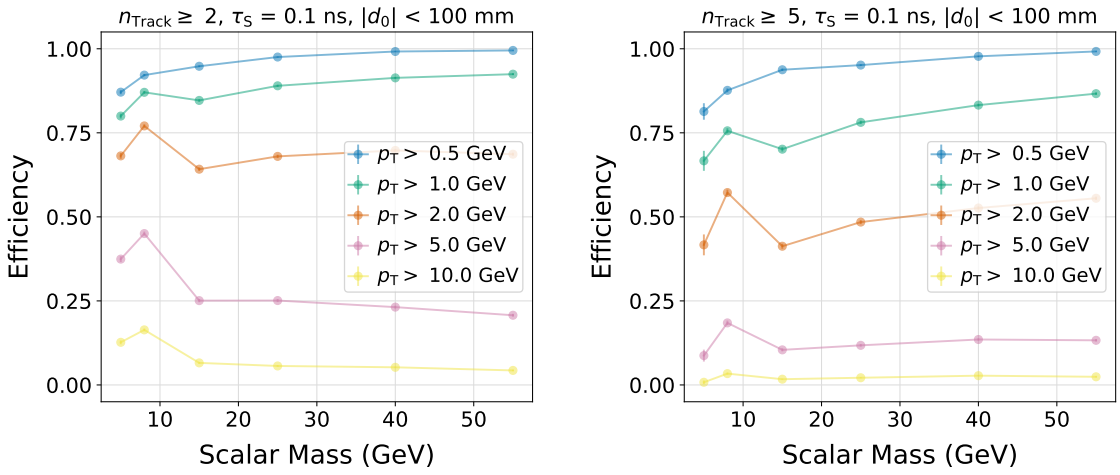

Figure 8. Event-level efficiency as a function of scalar mass for a range of minimum track p_T values. Efficiencies are shown for the two (left) or five (right) scenarios. In all cases, the track reconstruction efficiency is assumed to decrease linearly with $|d_0|$ up until $|d_0| = 100$ mm.

As with the acceptance definitions, two different values of n_{Track} are studied for the Higgs portal efficiencies. Efficiencies are defined on top of the corresponding n_{Track} acceptance. Figures in this section largely correspond to $n_{\text{Track}} \geq 5$, a likely realistic baseline for a trigger targeting hadronic long-lived particle decays. The additional study performed with $n_{\text{Track}} \geq 2$ is to allow for direct comparison to the stau model. Results for this two-track trigger are used in section 4 and directly compared in figure 8.

As expected, the Higgs portal efficiencies are very sensitive to the minimum p_T requirement, as shown in figure 8. The overall trend in efficiency is the same for each n_{Track} considered. For a minimum p_T of 2 GeV, efficiencies for $n_{\text{Track}} \geq 5$ drop to around 50% of their maximum value, and p_T thresholds any higher make this strategy ineffectual.

The Higgs portal efficiency's dependence on the $|d_0|$ range, shown in figure 9 is more modest than the dependence on p_T. Endpoints of $|d_0| < 1$ cm are at worst about 50% as efficient as endpoints of $|d_0| < 10$ cm for the lifetimes considered. Given the high track multiplicity per event and long-lived particle's boost, it is easier to find tracks with small $|d_0|$ for this model, even for very displaced decays.

Figure 10 shows the event-level efficiency as a function of scalar mass and lifetime for a variety of minimum p_T thresholds. The nominal threshold of $p_T > 2$ GeV for the CMS Phase 2 tracker results in a peak efficiency of around 50% for higher mass scalars.

JHEP02(2023)034

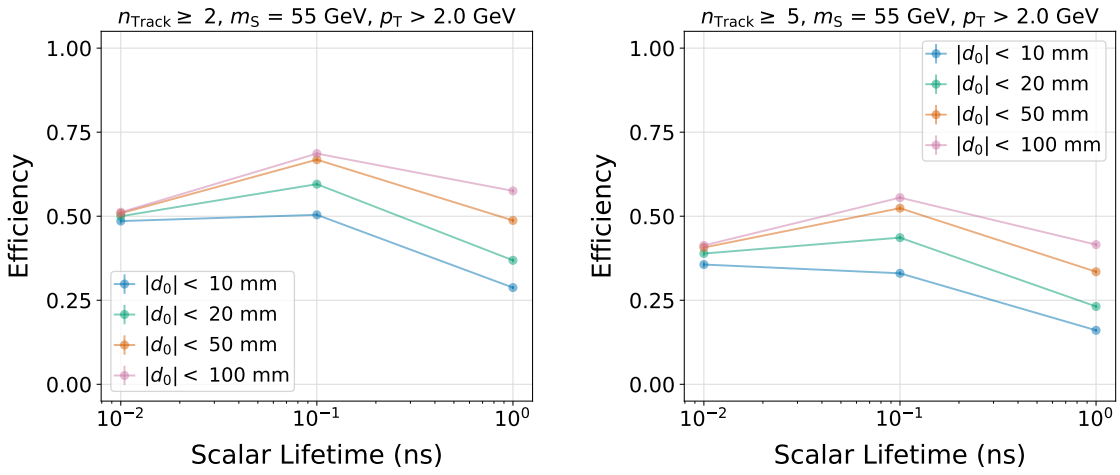

Figure 9. Event-level efficiency as a function of long-lived scalar lifetime for different $|d_0|$ endpoints and p_T larger than 2 GeV. Efficiencies are shown for the two (left) or five (right) n_{Track} scenarios. In all cases, the track reconstruction efficiency is assumed to decrease linearly with $|d_0|$ up until the endpoint specified.

JHEP02(2023)034

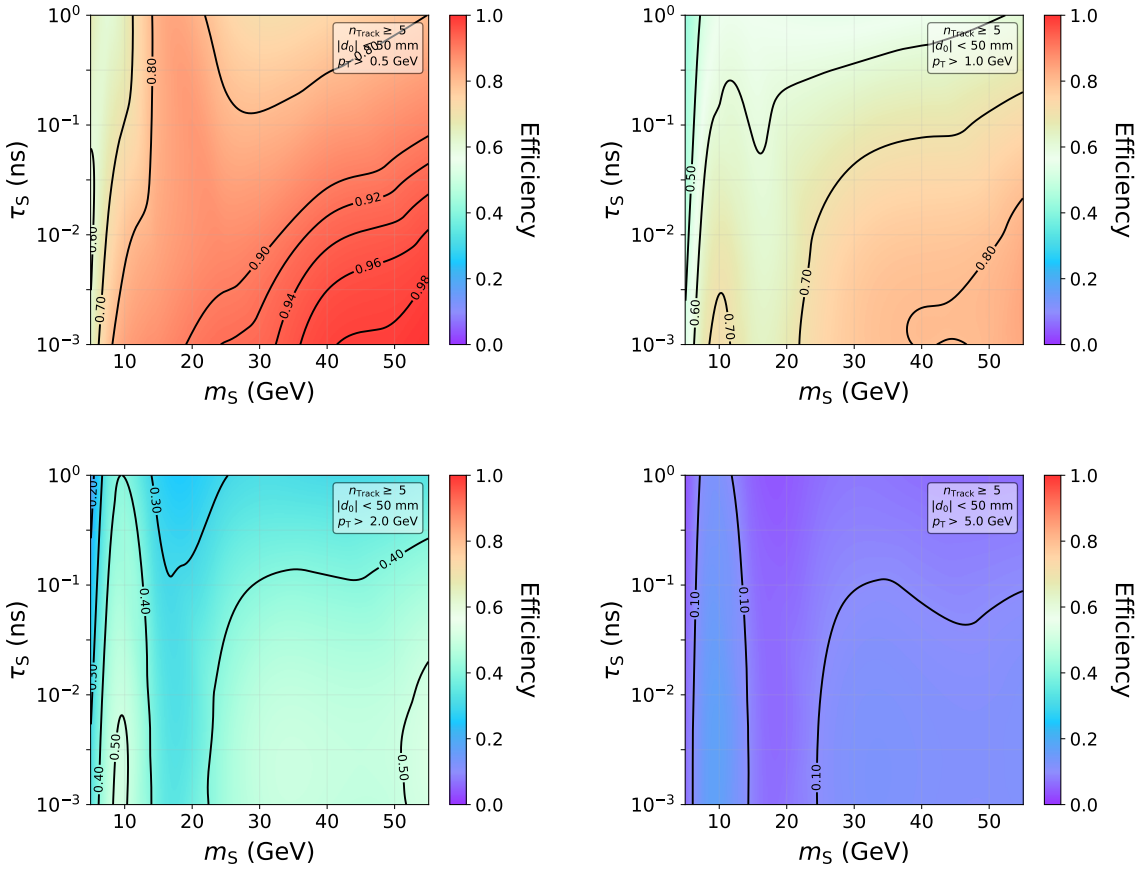

Figure 10. Efficiencies for the Higgs to long-lived scalar model for a variety of track p_T thresholds: 0.5 GeV (top left), 1 GeV (top right), 2 GeV (bottom left), 5 GeV (bottom right). Across a wide range of masses and lifetimes, efficiency falls off steeply as the p_T threshold increases above 2 GeV. In all cases, a linearly parameterized $|d_0|$ efficiency is used with an endpoint at 50 mm, and the five-track selection is used.

JHEP02(2023)034

2.4 Long-lived staus

Long-lived staus serve as a useful benchmark for two long-lived particle detection techniques. Staus can be either be directly detected as HSCPs or indirectly identified via displaced decay products. The current most comprehensive limits are from the OPAL Experiment, which exclude staus with mass $m_{\tilde{\tau}} < 90\,\text{GeV}$ for all possible lifetimes [38]. LHC searches for highly ionizing particles, disappearing tracks, and displaced leptons extend exclusions to larger stau masses [39–43]. However, these searches do not cover the full lifetime space. Searches which aim to indirectly detect displaced decay products are limited by trigger p_{T} thresholds and efficiency for displaced particles. Searches which aim to directly detect charged particles such as the stau itself are currently trigger-limited when a track is not reconstructed in the muon system.

The benchmark model considered here assumes direct production of a pair of staus in a gauge mediated Supersymmetry breaking (GMSB) scenario [44–46]. While some GMSB slepton models assume the three sleptons are mass degenerate, here the stau alone is taken to be the next-to-lightest supersymmetric particle (NLSP). The gravitino is the lightest supersymmetric particle in GMSB models. The stau decays to a tau and a gravitino via a small gravitational coupling, thereby gaining a significant lifetime.

Stau samples are generated and decayed with MadGraph5_aMC@NLO 2.9.3. A simplified model is used wherein the left-handed and right-handed staus are assumed to be mass degenerate and have a mixing angle $\sin\theta = 0.95$ [42]. The gravitino is given a negligible mass of 1 GeV. Stau samples are generated for masses $m_{\tilde{\tau}} = 100$, 200, 300, 400, 500, and 600 GeV and lifetimes 0.001, 0.01, 0.1, 1 ns. For these lifetimes, displaced tracks from the decay of the tau (pions and light leptons) can be reconstructed. Additional samples are generated with 10 ns and infinite lifetimes. For these lifetimes, the stau is sufficiently long-lived that it can be directly identified as an anomalous prompt track with slow velocity or high ionization.

2.4.1 Displaced leptons

For lifetimes below 1 ns, long-lived staus are most easily identified by their displaced decay products. In the GMSB model considered, where the light gravitino does not interact with the detector, the only visible signatures of the stau decays are the displaced taus. In the case of pair produced staus, two taus are present in the final state, each the only visible product of its own displaced vertex. The taus then decay as usual, leading to a mix of final state scenarios ranging from fully leptonic when both taus decay leptonically to fully hadronic.

While relatively rare (12% of di-tau events), the fully leptonic scenario produces the distinctive signature of a pair of displaced, high-quality leptons, and is the easiest signature to target. The fully leptonic decay mode is the only scenario explored at the LHC thus far, with Run 2 results from both ATLAS [42] and CMS [43]. Triggering on leptonic tau decays is also the most straightforward. Standard photon triggers can be used to recover displaced electrons. Triggers which only require a track reconstructed in the muon system, and are agnostic to tracker activity or impact parameter requirements, are highly efficient for displaced muons up to $|d_0| \lesssim 10$ cm [47]. In Run 2, displaced electrons and muons were

JHEP02(2023)034

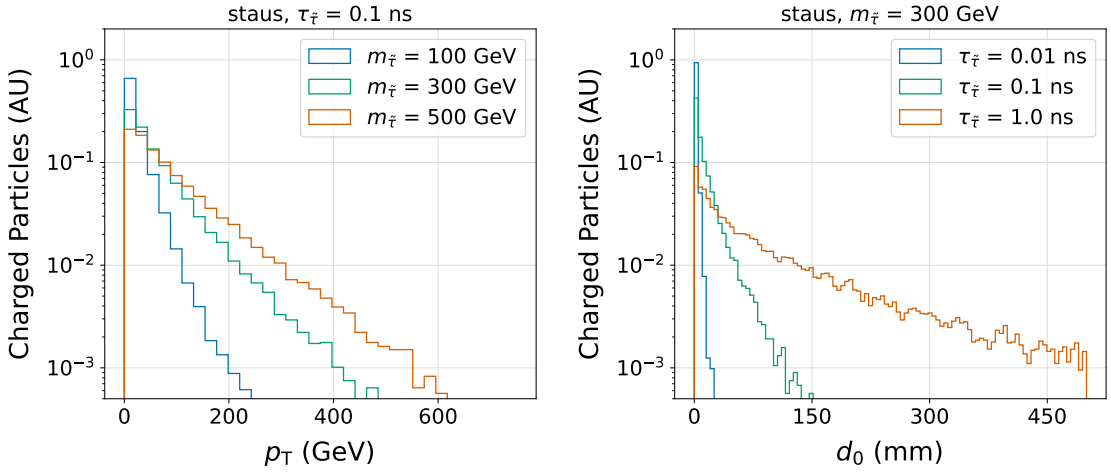

Figure 11. (Left) Displaced charged particle p_T. Particles are required to be decay products of the long-lived stau. (Right) Displaced charged particle $|d_0|$ for varied lifetimes.

both required to pass a relatively high p_T threshold, upwards of 40 GeV. Because each lepton carries roughly 1/6 of the parent stau's energy, these thresholds are prohibitively high for the identification of low mass staus.

Hadronic decays of displaced taus are harder to identify. The complex algorithms that are used to identify prompt taus must be adapted to recover efficiency for taus with larger impact parameters. Approximately 15% of taus decay to three charged particles. In this unique case, a visible displaced vertex can be reconstructed. The remaining 50% of taus decay hadronically to a single charged particle. In both cases, charged particles may be accompanied by one or more neutral hadrons, and are always accompanied by a single tau neutrino. The fraction of the tau's momentum carried by the charged particles depends on the decay mode.

In all scenarios, at least two displaced tracks are produced per signal event. Figure 11 shows the p_T and $|d_0|$ of displaced charged particles produced in the stau decay. For staus with masses above 100 GeV, displaced tracks typically have p_T values well above the minimum thresholds for potential track-triggers. Displaced tracks tend to have larger impact parameters than in the Higgs portal model as a result of the larger parent mass.

This section follows a similar strategy to the Higgs portal model, as described in section 2.3. The baseline track acceptance is described in table 2. Events are considered to pass acceptance if they contain at least 2 tracks passing all acceptance requirements. The resulting event-level acceptance, for a range of masses and lifetimes, can be seen in figure 12.

Efficiency requirements for individual tracks are also described in table 2. The event-level efficiency requires events have at least two tracks passing the per-track efficiency, with respect to events passing acceptance. A range of minimum p_T thresholds are explored, as shown in figure 13. In this scenario p_T thresholds have a minimal impact on overall efficiency, especially for higher stau masses. For the CMS track-trigger baseline p_T threshold of 2 GeV, there is no significant reduction of efficiency compared to lower thresholds.

JHEP02(2023)034

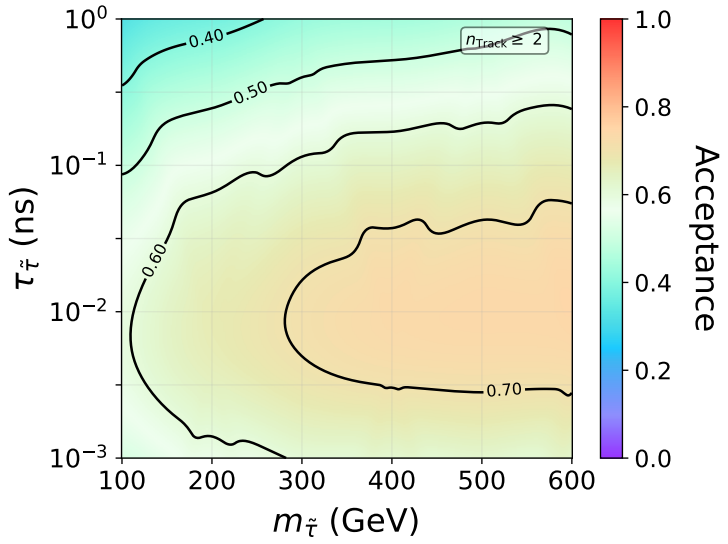

Figure 12. Acceptance for displaced lepton signature requiring at least two tracks and passing the requirements listed in table 2.

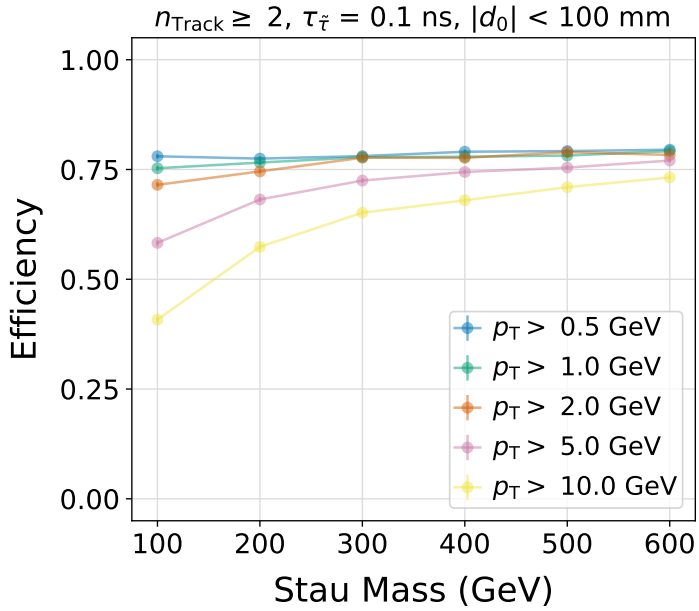

Figure 13. Efficiency as a function of stau mass for a range of minimum track p_{T} values. Events are defined to have passed if they contain two tracks with a p_{T} value larger than the minimum value, as well as a $|d_0|$ larger than 1 mm and smaller than 100 mm. The proper lifetime is $\tau = 0.1$ ns. Track reconstruction efficiency decreases linearly from 100% for prompt tracks to 0% at $|d_0| = 100$ mm.

JHEP02(2023)034

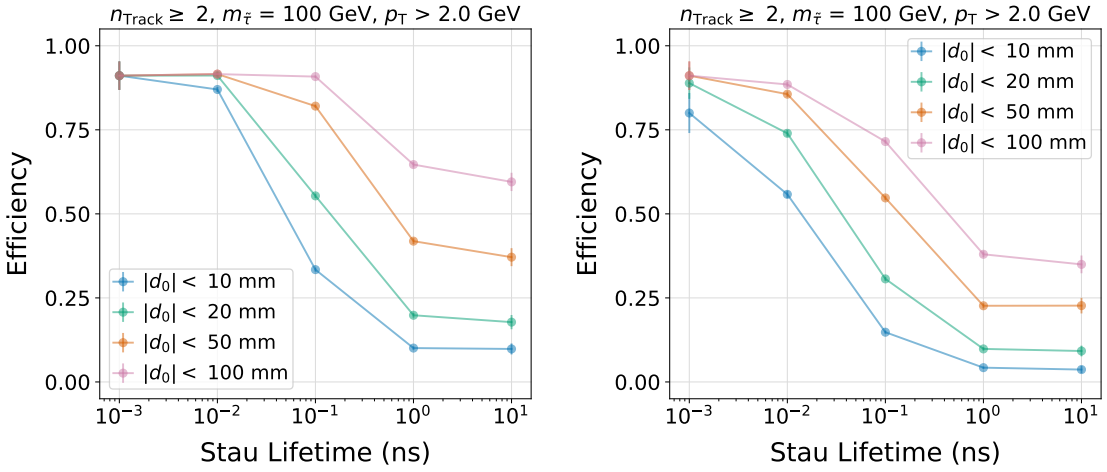

Figure 14. Event-level efficiency as a function of stau lifetime for a $m_{\tilde{\tau}} = 100\,\text{GeV}$ stau. (Left) Assumes 100% efficiency up until a $|d_0|$ endpoint. (Right) Efficiency decreases linearly from 100% for prompt tracks to 0% at the $|d_0|$ endpoint. In both cases, a minimum track p_{T} of 2 GeV and a minimum track $|d_0|$ of 1 mm are required.

The event-level efficiency for this model is much more sensitive to changes in the track-trigger efficiency as a function of $|d_0|$. In this study, two parameterizations of efficiency are compared. The first parameterization is a binary model in which all tracks with $|d_0|$ less than a given value are assumed to be reconstructed. The second scenario follows the more realistic model described in section 2.1. The efficiency for prompt tracks is assumed to be 100%, decreasing linearly to 0% at the $|d_0|$ endpoint. The two scenarios are compared in figure 14, which shows event-level efficiencies for a range of stau lifetimes. Though the binary model results in a higher efficiency, by definition, the difference is modest ($< 30\%$) for most lifetimes.

Figure 15 shows the event-level efficiency as a function of stau mass and lifetime for a variety of $|d_0|$ endpoints, assuming a linearly decreasing efficiency as a function of $|d_0|$. For endpoints below 2 cm, efficiencies are small ($\leq 30\%$) across the plane, but for endpoints at 10 cm, closer to typical offline efficiencies, efficiencies are high ($30 - 80\%$) for the full range of lifetimes. Given the low multiplicity of tracks expected per event with impact parameters roughly consistent with the stau decay radius, this sensitivity to $|d_0|$ endpoint is expected.

JHEP02(2023)034

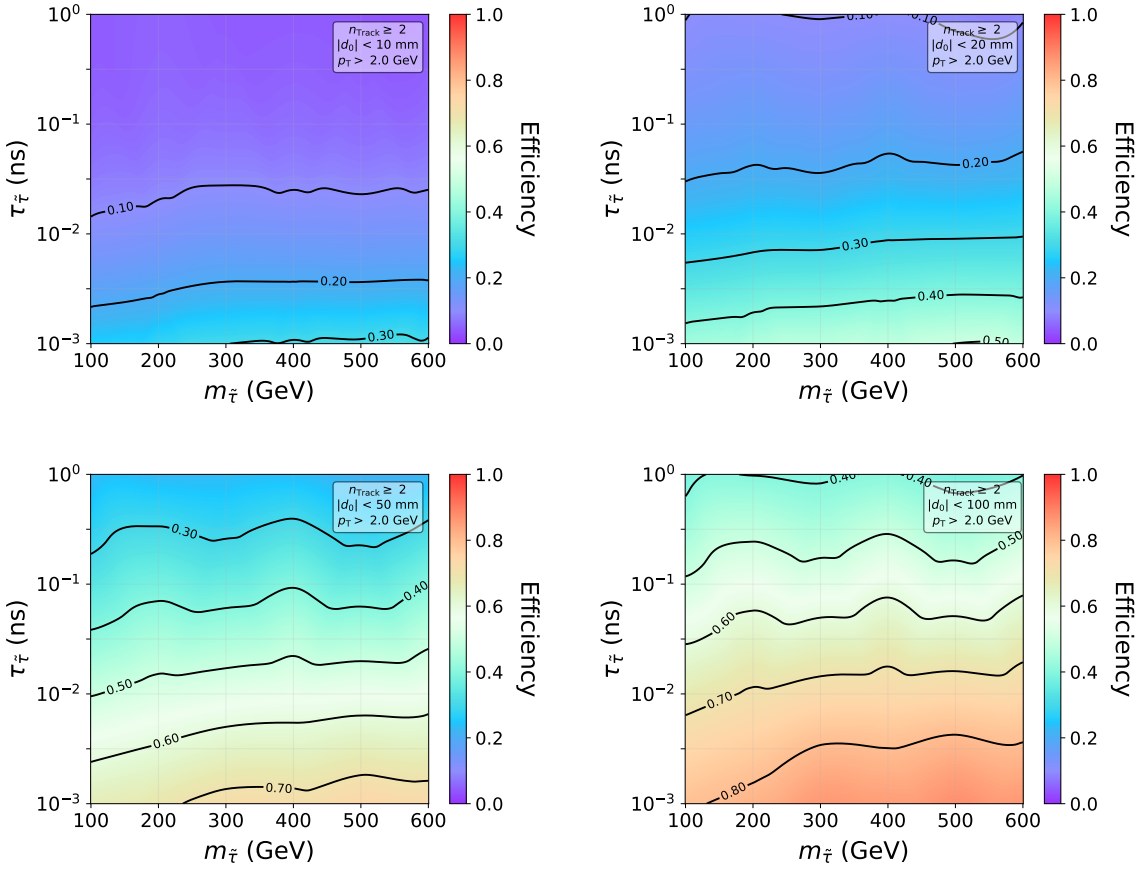

Figure 15. Event-level efficiencies for the stau model for a variety of track $|d_0|$ endpoints: 10 mm (top left), 20 mm (top right), 50 mm (bottom left), 100 mm (bottom right). In all cases, a linearly parameterized $|d_0|$ efficiency is used. Tracks are required to have $p_{\text{T}} > 2\,\text{GeV}$ and events are required to have $n_{\text{Track}} \geq 2$.

2.4.2 Direct detection

For long-lived staus with proper lifetimes of $\tau \gtrsim 1$ ns, the charged stau will traverse a sufficient number of tracker layers such that its trajectory can be directly reconstructed as a prompt, high-p_{T}, isolated track. These staus will be slowly moving and highly ionizing with respect to Standard Model charged particles produced at the LHC. Most SM charged particles have mass $m < 1\,\text{GeV}$, $\beta \sim 1$, and can be approximated as minimum ionizing particles. In contrast, staus with masses $m \geq 100\,\text{GeV}$ will most often be produced with $\beta\gamma \sim 1$ and be slowly moving, $\beta < 1$.

This study considers long-lived staus with masses between $100 \leq m_{\tilde{\tau}} \leq 1000\,\text{GeV}$ and proper lifetimes between $0.01 \leq \tau_{\tilde{\tau}} \leq 10$ ns. Two possible trigger strategies are studied. The first scenario looks for at least one high-p_{T} isolated track, and investigates the impact of geometric acceptance on trigger efficiency. This scenario is meant to be as inclusive as possible to direct long-lived particle detection, and can serve as a baseline for triggers which additionally incorporate delay or anomalous ionization measurements to reduce backgrounds.

JHEP02(2023)034

Variable	Requirement
Acceptance	
$\lvert\eta\rvert$	$< 1.0, 2.5, 4.0$
L_{xy}	$> 600, 800, 1000, 1200$ mm
OR $\lvert z\rvert$	> 3000 mm
Inclusive Efficiency	
p_{T}	$> 10, 20, 50, 100$ GeV
Time-of-flight Efficiency	
delay	$> 0.25, 0.33, 0.50$ ns
OR β_{TOF}	$< 0.96, 0.95, 0.9$
OR m_{TOF}	$> 15, 30, 60$ GeV

Table 3. Track acceptance and efficiency requirements considered for the inclusive high-p_{T} charged particle selection and for the time-of-flight selection. Events are required to have at least one stau passing these selections.

The second scenario assumes the outermost layer of the tracker also serves as a time of flight layer, and incorporates a timing measurement to further distinguish staus from SM tracks.

The selections considered for the inclusive trigger scenario are summarized in table 3, and require at least one stau to decay within detector acceptance. The first geometric effect considered is the pseudorapidity range of the tracking detector or track-trigger. Typically the tracker is designed to emphasize efficiency and performance in the barrel. In comparison the endcap and further forward regions have degraded performance and higher background rates. The maximum pseudorapidity is varied from $|\eta| < 1.0,\ 2.5,\ 4.0$ to represent a barrel only scenario (if endcap rates are prohibitively high), a nominal scenario, and an extended scenario, respectively.

Another geometric constraint is that the stau must traverse a sufficient number of tracker layers to be reconstructed. To investigate this effect, the minimum L_{xy} beyond which the stau must decay is varied from $600 < L_{\mathrm{xy}} < 1200$ mm. The minimum L_{xy} corresponds to particles which would traverse the pixel detector and at least three p_{T}-layers of the upgraded CMS tracker, or at least seven silicon layers in the upgraded ATLAS tracker. The maximum L_{xy} roughly corresponds to the full tracker for both experiments. In all cases the minimum $|z|$ of the endcap is kept fixed at 3000 mm.

Signatures of disappearing tracks with shorter L_{xy} could be considered, and have been used for offline discrimination of signal processes in both ATLAS and CMS [40, 41]. However, this method is prone to high fake rates due to the reduced number of track hits, and is therefore outside of the scope of this study.

At the HL-LHC, the inclusive trigger studied here would likely include a track-based isolation requirement in order to reduce backgrounds. This study assumes sufficient tracker longitudinal impact parameter resolution such that any track-based stau isolation requirement would be fully efficient.

Figure 16 shows event-level acceptance as a function of varying the track-trigger maximum $|\eta|$ for a fixed $L_{\mathrm{xy}} = 1200$ mm requirement. Varying the track-trigger pseudorapidity

JHEP02(2023)034

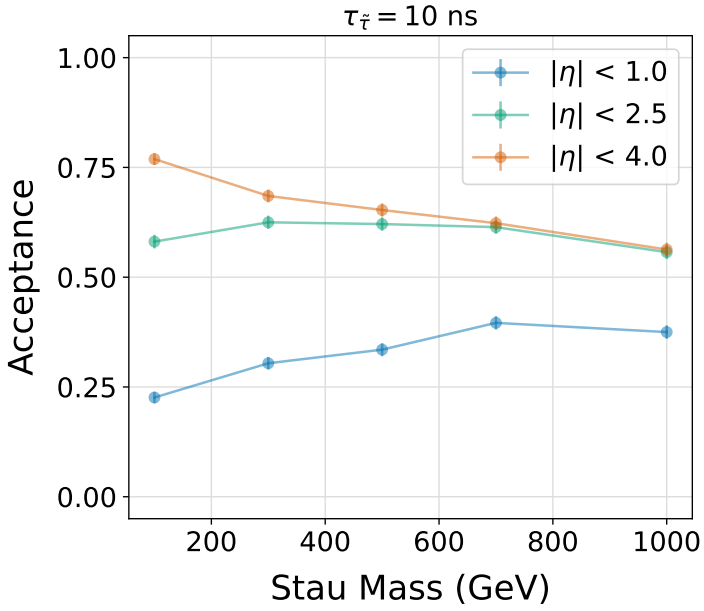

Figure 16. Stau model event-level acceptance for a variety of minimum tracker $|\eta|$ values. Acceptance is shown as a function of stau mass for proper lifetime $\tau = 10$ ns. The stau is required decay beyond $L_{xy} = 1200$ mm.

has a larger effect on the event-level acceptance for lower mass staus. Extending the track-trigger to $|\eta| < 4.0$ would improve the acceptance for $m_{\tilde{\tau}} = 100\,\text{GeV}$ by 30% but provides negligible benefit for $m_{\tilde{\tau}} = 1\,\text{TeV}$. Extending the track-trigger acceptance to the far forward region would be extremely challenging from a technical perspective, and offer little benefit in terms of physics reach for metastable charged particles because lower mass BSM particles also benefit from larger cross sections. In contrast, reducing the pseudorapidity coverage from $|\eta| < 2.5$ to the barrel only scenario, $|\eta| < 1.0$, would reduce the overall acceptance by approximately 50% or more, and provides strong motivation for including the endcaps in any track-trigger.

The impact of varying the minimum number of layers to reconstruct a track is shown in figure 17. This study assumes a fixed pseudorapidity requirement of $|\eta| < 2.5$. Reducing the minimum L_{xy} from 1200 mm to 600 mm, or requiring fewer layers per track, improves the overall acceptance for stau lifetimes of $\tau_{\tilde{\tau}} = 1$ ns by roughly a factor of four. However, the event-level efficiency for this lifetime never exceeds that of a displaced track trigger, as shown in figure 14. Varying the number of layers required per track does not significantly improve the acceptance for lifetimes $\tau_{\tilde{\tau}} > 1$ ns, because most staus decay well beyond the tracker, $L_{xy} \sim 1.2$ m.

Increasing the minimum p_T required per track is a useful handle to reduce background rates while retaining nearly full signal efficiency. The impact of varying the minimum track p_T threshold on the event-level efficiency is shown in figure 33 of appendix A.3. All p_T thresholds considered are nearly fully efficient except for the lowest stau mass considered.

JHEP02(2023)034

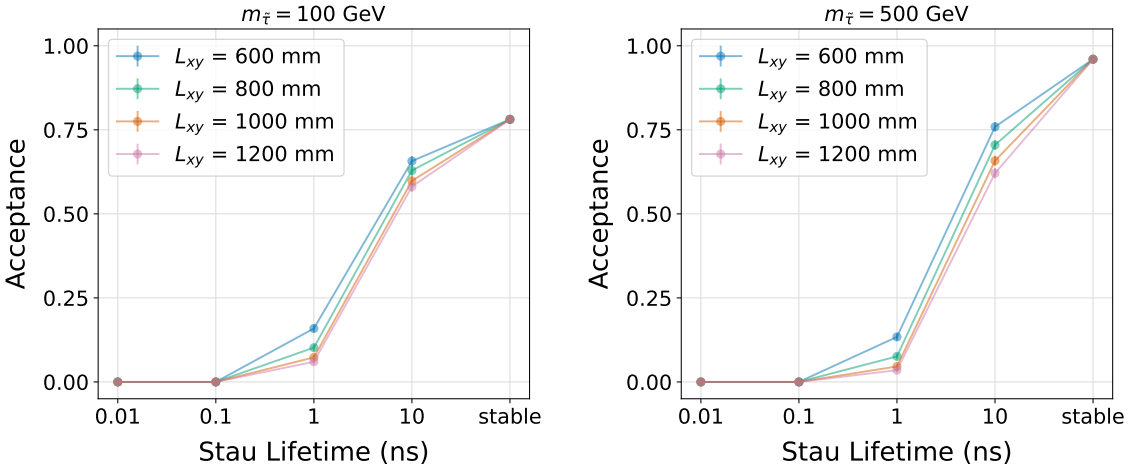

(a) (b)

Figure 17. Stau event-level acceptance for a variety of minimum tracker L_{xy} values. Acceptance is shown as a function of stau lifetime for staus with mass $m_{\tilde{\tau}} = 100\,\mathrm{GeV}$ (a) and $m_{\tilde{\tau}} = 500\,\mathrm{GeV}$ (b). The track-trigger covers $|\eta| < 2.5$.

The second track-trigger scenario investigates using a time-of-flight detector as an additional handle to improve signal to background discrimination. A timing layer similar to the planned CMS MIP timing detector is considered [48]. The timing layer is assumed to be located at roughly $L_{\mathrm{xy}} = 1150$ mm, $z = 3000$ mm, covering $|\eta| < 2.5$, and provide a 50 picosecond resolution timestamp for all charged particles passing through the detector. ATLAS also plans to incorporate precision timing at the HL-LHC, but only at larger pseudorapidity [49]. For ATLAS, it may be more optimal to consider calorimeter or muon spectrometer timing measurements to target heavy metastable charged particles.

The time-of-flight detector's measurement for a given particle is computed according to

$$t_{\mathrm{hit}} = \frac{L(p_{\mathrm{T}}, \eta)}{c \cdot p}\sqrt{p^2 + m^2}$$

An increase in path length, L, due to the magnetic field is negligible for charged particles with $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ in a 4 Tesla solenoidal magnetic field, and therefore neglected by this study. The detector resolution is assumed to have a Gaussian uncertainty with a width of $\sigma = 50$ ps. The spread of collisions in z and the spread in collision time are assumed to follow a Gaussian distribution with widths $\sigma = 50$ mm and $\sigma = 200$ ps, respectively [50]. This study assumes the hardware track-trigger is able to measure the track's longitudinal impact parameter, or origin in z, but not the primary vertex time. It is likely that computing the primary vertex time will only be possible in the HLT or offline, and that the uncertainty in the L1's time-of-flight measurement will be dominated by the beamspot's spread in time. For simplicity, this study focuses on stable lifetimes, and the selections considered are summarized in table 3.

Time-of-flight and delay distributions are shown in figure 18 for different stau masses. For comparison, a "background" sample is simulated by using SM particles with $p_{\mathrm{T}} > 10\,\mathrm{GeV}$

JHEP02(2023)034

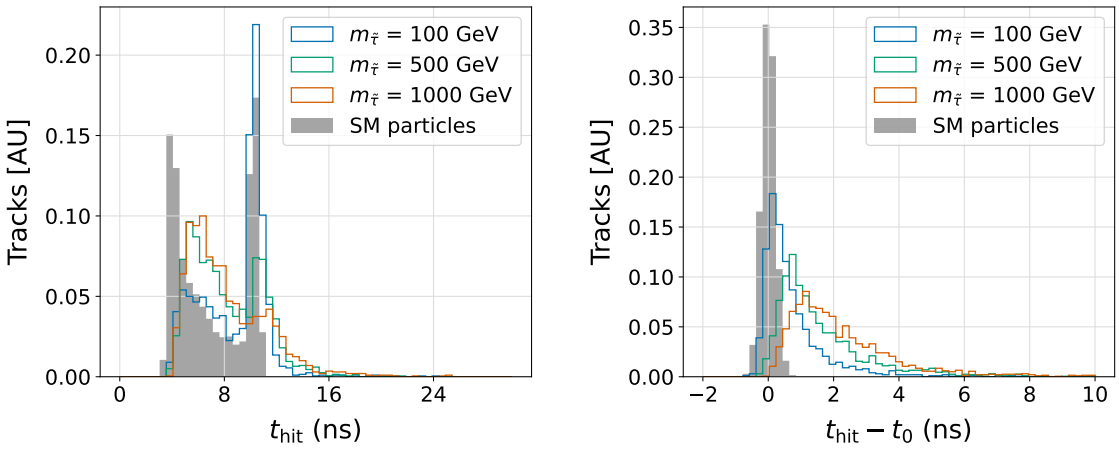

(a) (b)

Figure 18. Timestamp as measured by the time-of-flight detector (a) and delay in time of arrival with respect to particles travelling at the speed of light (b). These distributions incorporate a realistic spread in collision time and z position, as well as the timing layer's temporal resolution. SM particles with $p_{\text{T}} > 10\,\text{GeV}$ from the signal samples are shown for comparison.

in the signal samples as a proxy. The time-of-flight measurements contain two peaks, at $t_{\text{hit}} \sim 3$ and $t_{\text{hit}} \sim 10$ ns, corresponding to the difference in path length for tracks which traverse the barrel or the endcaps, respectively. The delay is defined to be the time-of-flight measurement subtracted by the time it would take a particle traveling at the speed of light to reach the same region of the detector. Staus arrive noticeably later than background particles, and the measured delay in arrival increases with stau mass.

With the time-of-flight measurement it is also possible to extract the stau velocity, β, and mass as shown in figure 19. Computing β_{TOF} accounts for differences in path length, due to the stau η and the longitudinal position of the primary vertex z. As expected, background tracks peak at $\beta_{\text{TOF}} = 1$, while the measured velocity decreases with increasing mass.

Computing a charged particle's mass requires a track p_{T} measurement, and the mass resolution can be described as

$$(\Delta m_{\text{TOF}})^2 = m^2 \left[\left(\frac{\Delta p_{\text{T}}}{p_{\text{T}}} \right)^2 + \left(\frac{1}{1-\beta^2} \right)^2 \left(\frac{\sigma_{t_{\text{hit}}}}{t_{\text{hit}}} \right)^2 \right]$$

For simplicity, this study assumes a p_{T} resolution of 1% which increases up to 10% for particles with p_{T} of 1 TeV. This assumption is somewhat optimistic for the endcaps, but reasonable for the barrel [16, 17]. Measured masses peak at the actual stau mass, while background is peaked near zero, with long tails due to resolution effects.

The event-level efficiency for staus to pass a variety of β_{TOF} and m_{TOF} requirements are shown in figure 34 of appendix A.3. Efficiencies are shown for stable staus with respect to having at least one track in geometric acceptance. Selections on $\beta_{\text{TOF}} < 0.96, 0.95, 0.90$ or $m_{\text{TOF}} > 15, 30, 60\,\text{GeV}$ were chosen to provide comparable background efficiencies of

JHEP02(2023)034

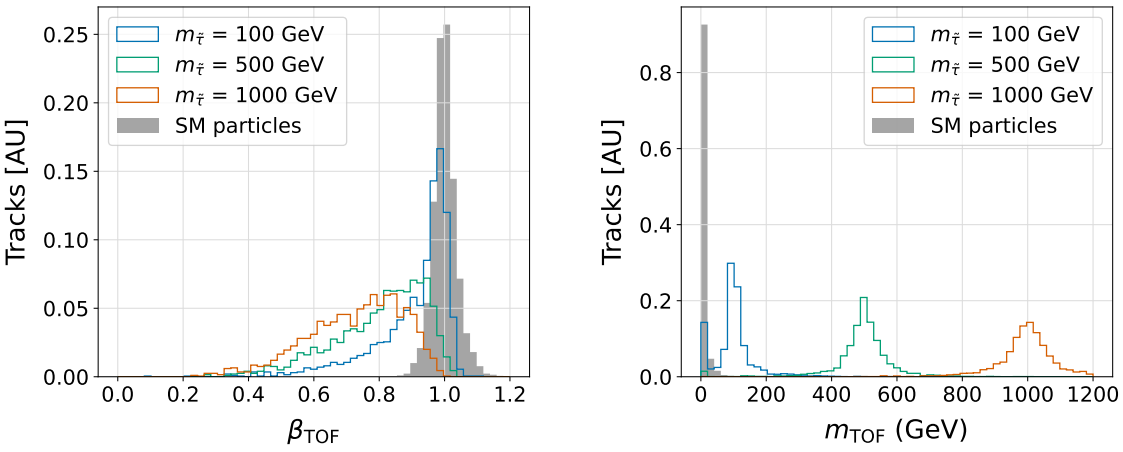

(a) (b)

Figure 19. Measured β (a) and mass (b) for different stau masses. These distributions incorporate a realistic spread in collision time and z position, as well as the timing layer's temporal resolution. SM particles with $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ from the signal samples are shown for comparison.

10%, 5%, and 1% per track, respectively. In both cases, the event-level efficiency increases with stau mass. Requirements on β_{TOF} are less efficient than the measured mass.

3 Background rates and trigger feasibility

Future track-triggers must operate within strict latency constraints and reject HL-LHC backgrounds to reduce data flow to a manageable rate. The required latency depends on the overall design of the trigger system, and the ability to meet latency requirements depends on the exact implementation of the hardware track-trigger. Any latency related studies are therefore beyond the scope of this publication, which intends to remain independent of implementation details. However, it is possible to investigate whether HL-LHC backgrounds would overwhelm the BSM signals given the track-trigger reconstruction parameters considered. While the previous section focused on the efficiency of a potential track-trigger to reconstruct anomalous tracks in BSM signal events, this section investigates signal efficiency and background rates for specific trigger selections. *The goal of this section is not to identify an optimal selection, but to demonstrate trigger feasibility.*

The vast majority of background events input to any track-trigger at the HL-LHC will consist of low-p_{T} QCD processes with ~ 200 overlapping pile-up collisions. Cross sections for electroweak processes are several orders of magnitude lower than QCD and can be safely neglected.

Displaced tracks additionally have backgrounds from long-lived SM particles, material interactions inside the detector, cosmic ray muons, and tracks which are reconstructed from an incorrect combination of hits. Material interactions and cosmic ray backgrounds can be rejected with fairly simple selections [47]. Fake tracks due to the incorrect combination

JHEP02(2023)034

Variable	Requirement
p_{T}	> 1.0, 2.0, 10
$\lvert\eta\rvert$	< 2.5
production vertex	$L_{\mathrm{xy}} < 300$ mm AND $\lvert z\rvert < 300$ mm
decay vertex	$L_{\mathrm{xy}} > 1$ m OR $\lvert z\rvert > 3$ m

Table 4. Track criteria for minimum bias background studies.

of hits are highly dependent on tracker geometry and the exact implementation of the track-trigger. In CMS's L1 track finder, fake rates are expected to be around 5% for prompt tracks after applying quality cuts [51]. Prompt signatures are therefore dominated by SM backgrounds and fake contributions can be safely ignored.

Due to reduced SM backgrounds, fake tracks are a more significant background for displaced track reconstruction. To minimize this effect, this study considers only displaced tracks with relatively small production radii and long path-lengths, such that a sufficient number of hits and quality of fit requirements could be applied per track to keep combinatoric backgrounds low. While these quality requirements are not explicitly addressed in this study, the shape of the efficiency vs. $|d_0|$ parameterization is taken from existing algorithms that employ these cuts to keep fake rates low. Additional requirements on vertexing or more complex topological requirements could be added to further reduce the contribution of fake tracks for a given trigger selection.

In order to estimate SM background rates, a minimum bias sample of 100,000 events is generated using Pythia8. Half of the sample consists of events with $p_{\mathrm{T}} < 20\,\mathrm{GeV}$, which are produced via Pythia's non-diffractive soft QCD class. The other half of events are produced with Pythia's hard QCD class, and required to have $p_{\mathrm{T}} > 20\,\mathrm{GeV}$. Because higher p_{T} events are more likely to pass a given trigger, the rate of hard QCD events produced is biased by a ${p_{\mathrm{T}}}^4$ selection factor. Event weights are applied to correct for this bias.

Tracks studied in this section are required to pass the criteria outlined in table 4. A threshold of $p_{\mathrm{T}} > 0.5\,\mathrm{GeV}$ was found to result in an overwhelming rate of QCD background tracks. Therefore thresholds of $p_{\mathrm{T}} > 1\,\mathrm{GeV}$ or higher are considered. Requirements are designed to be consistent with previous studies in this paper. Particles are required to be produced within 300 mm in L_{xy} and $|z|$, and decay outside the tracker barrel $L_{\mathrm{xy}} > 1$ m or endcap $|z| > 3$ m. In this section, the linearly falling efficiency in $|d_0|$ with a 100 mm endpoint presented in previous sections is used, which is pessimistic with respect to offline displaced tracking efficiency shown in ref. [23]. The resulting track p_{T} and $|d_0|$ for the minimum bias sample are compared to representative signal samples in figure 20. As expected, QCD tracks tend to be much lower in p_{T} than signal tracks, especially stau samples, and QCD $|d_0|$ drops off rapidly compared to long-lived signals.

HL-LHC pile-up conditions are mimicked by overlaying minimum bias collisions on each generated event. For each event of interest, 199 pile-up collisions are randomly chosen from the minimum bias sample with a probability consistent with the event weight. The z position of each primary and pile-up collision is sampled from a Gaussian distribution with

JHEP02(2023)034

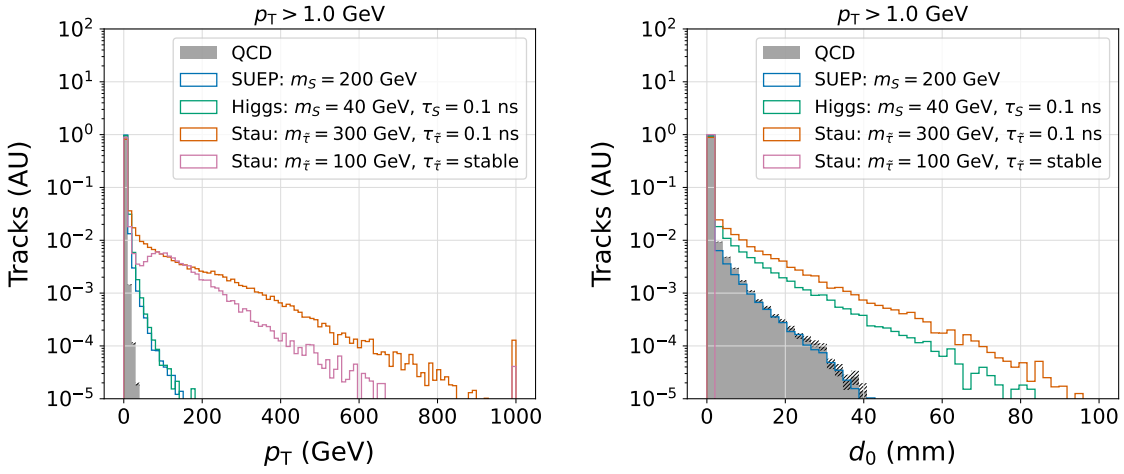

(a) (b)

Figure 20. Track p_T (a) and $|d_0|$ (b) for QCD and representative signal samples.

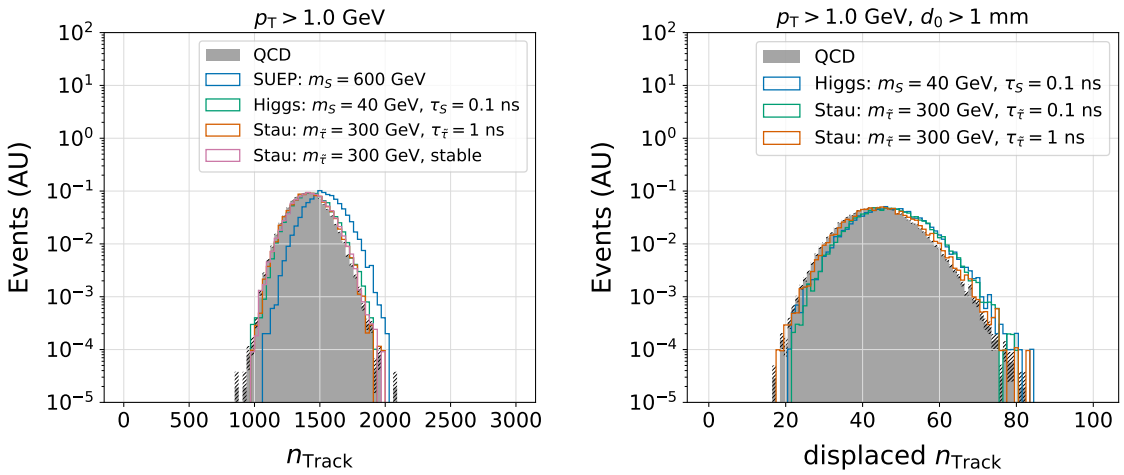

(a) (b)

Figure 21. The total number of tracks per event(a) and total number of displaced tracks per event (b) for minimum bias QCD and representative signal samples at pile-up=200. Slashed bands on the QCD distribution represent the statistical uncertainty.

a width $\sigma_z = 50$ mm in order to account for the longitudinal spread of the beamspot [50]. The production and decay vertex positions of each truth particle in the collision are adjusted accordingly.

The total number of tracks passing the criteria described in table 4 with $p_T > 1$ GeV are shown in figure 21. Tracks which additionally pass a displacement requirement of $|d_0| > 1$ mm are shown for comparison. On average ~ 1400 tracks are expected per event, and 50 of those tracks are displaced.

JHEP02(2023)034

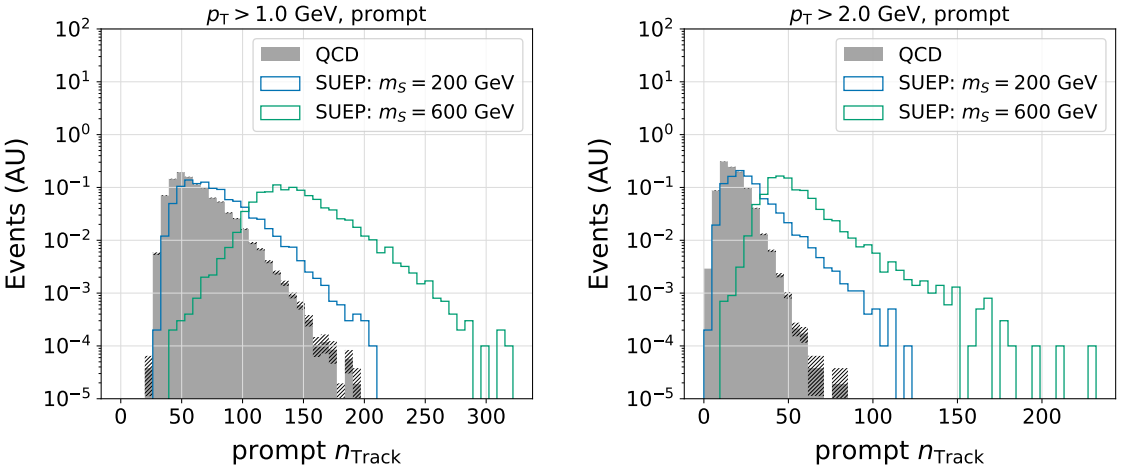

(a) (b)

Figure 22. The number of prompt tracks with $p_T > 1$ GeV (a) and $p_T > 2$ GeV (b) for minimum bias and representative SUEP signal samples at pile-up 200. Slashed bands on the QCD distribution represent the statistical uncertainty.

For the purpose of this study, an acceptable background rate for an individual trigger is assumed to be at most a few percent of the total L1 output. When correcting for empty bunches, the LHC collides protons bunches at a rate of ~ 30 MHz. As currently planned, ATLAS and CMS L1 triggers will have output rates of 1 MHz and 750 kHz, respectively [18, 52]. To be considered acceptable for this study, background rates must be below 10 kHz, corresponding to QCD selection factors smaller than $\sim 1/3000$.

Figure 21 demonstrates that a primary vertex selection is essential to separate prompt signals from QCD in high pile-up conditions. The track-trigger is assumed to have a longitudinal vertex resolution of 1 mm, and merge truth-level pile-up vertices with $|\Delta z| <$ 1 mm. Once the primary vertex associated with the signal process is identified, any tracks with $|\Delta z| < 1$ mm and $|d_0| < 1$ mm are defined to pass the prompt selection.

The primary vertex is typically identified as the vertex with the largest sum of track ${p_T}^2$, in order to prioritize higher p_T tracks. This definition is highly efficient for the HSCP scenario, and correctly selects the primary vertex $95 - 100\%$ of the time, for lower and higher mass staus respectively. However, for signals with lower p_T tracks, such as the SUEP scenario, this criteria selects the correct primary vertex $30 - 70\%$ of the time, for the lowest and highest mediator masses considered. For SUEPs, the vertex with largest number of tracks is taken to be the primary vertex, which is the correct choice for $70 - 100\%$ of events.

Figure 22 shows the number of prompt tracks for the minimum bias and SUEP samples. It is possible to reject QCD to acceptable levels by requiring $n_{\text{Track}} > 150$ or $n_{\text{Track}} > 50$ for $p_T > 1$ and 2 GeV tracks, respectively. Resulting efficiency would be between $1 - 10\%$ depending on the mediator mass.

Alternative handles to retain SUEP signal efficiency include the number of track jets per event and event H_T, which are shown in figure 23. To compute these observables,

JHEP02(2023)034

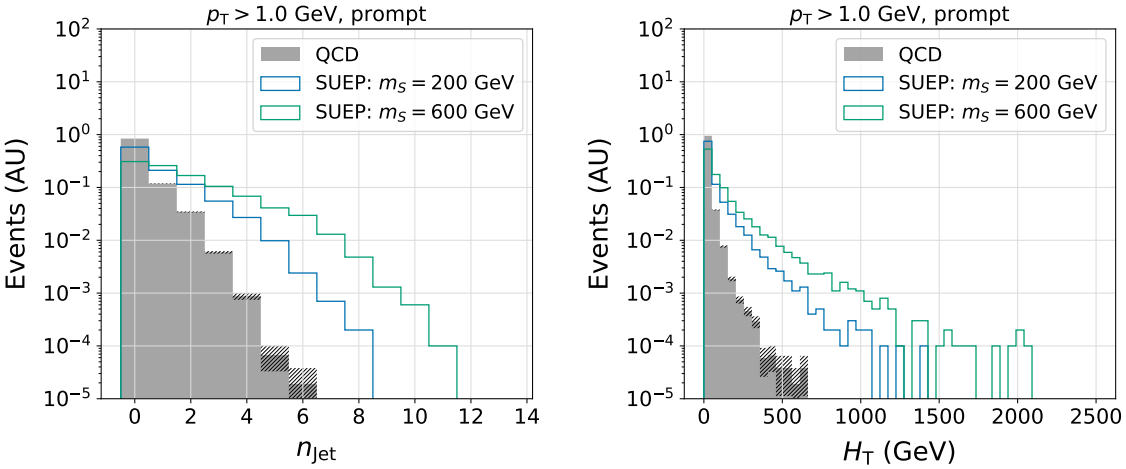

(a) (b)

Figure 23. The number of prompt jets (a) and H_T (b) for minimum bias QCD and representative SUEP signal samples at 200 pile-up. Jets are reconstructed from prompt tracks with $p_T > 1$ GeV, and jets are required to have $p_T > 25$ GeV. Slashed bands on the QCD distribution represent the statistical uncertainty.

Anti-k_T $R = 0.4$ jets are reconstructed from prompt tracks with $p_T > 1$ GeV, and jets are required to have $p_T > 25$ GeV. No requirements on n_{Track} are applied. A requirement of $n_{jet} > 4$ or $H_T > 300$ would enable SUEP efficiencies of a few percent. More complex event-shape variables such as isotropy or sphericity could be considered for additional discrimination [30, 31].

For the HSCP scenario counting the number of high-p_T, isolated, or slowly moving tracks per event is sufficient to enable a feasible trigger. To compute isolation, prompt tracks within $\Delta R < 0.4$ of the HSCP candidate are considered. Tracks are considered isolated if the p_T-sum of other tracks in that cone divided by the HSCP candidate p_T is less than 0.05. The time of flight, β_{TOF}, and m_{TOF}, are computed using the same approach as described in section 2.4.2. Figure 24 shows efficiencies of above 80% can easily be achieved.

For signals with displaced tracks, a primary vertex requirement cannot be applied, and high-$|d_0|$ tracks from all pile-up vertices must be considered. For higher mass scenarios, such as long-lived staus, a low rate, efficient trigger can be defined by increasing the p_T threshold to 10 GeV and requiring $n_{Track} \geq 2$, as shown in figure 25. For the lower mass Higgs portal scenario, this strategy would result in an efficiency less than one percent, so a different approach is needed.

The Higgs portal signal is the most challenging scenario considered, but there are several possible strategies to reduce background rates. In Run 2, CMS operated with a dedicated trigger that selected events with $H_T > 400$ GeV and trackless or displaced jets [53]. ATLAS uses a similar concept for their displaced decay triggers based on calorimeter and muon spectrometer information [54]. For the HL-LHC, there are several strategies to identify displaced jets or displaced vertices [19, 55].

JHEP02(2023)034

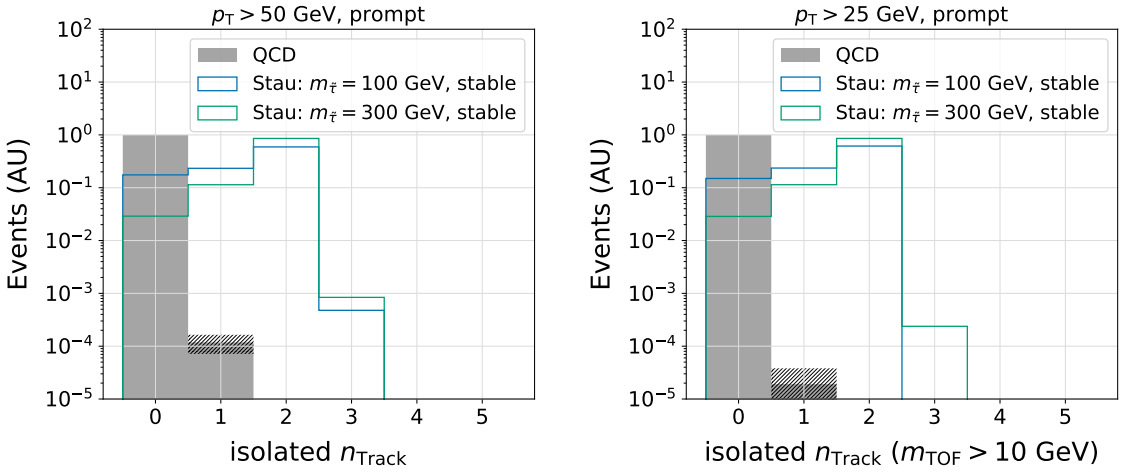

(a) (b)

Figure 24. The number of $p_T > 50\,\text{GeV}$ isolated tracks (a) and $m_{\text{TOF}} > 25\,\text{GeV}$, isolated $p_T > 25\,\text{GeV}$ tracks (b) per event for minimum bias QCD and representative HSCP signal samples with pile-up 200. Slashed bands on the QCD distribution represent the statistical uncertainty.

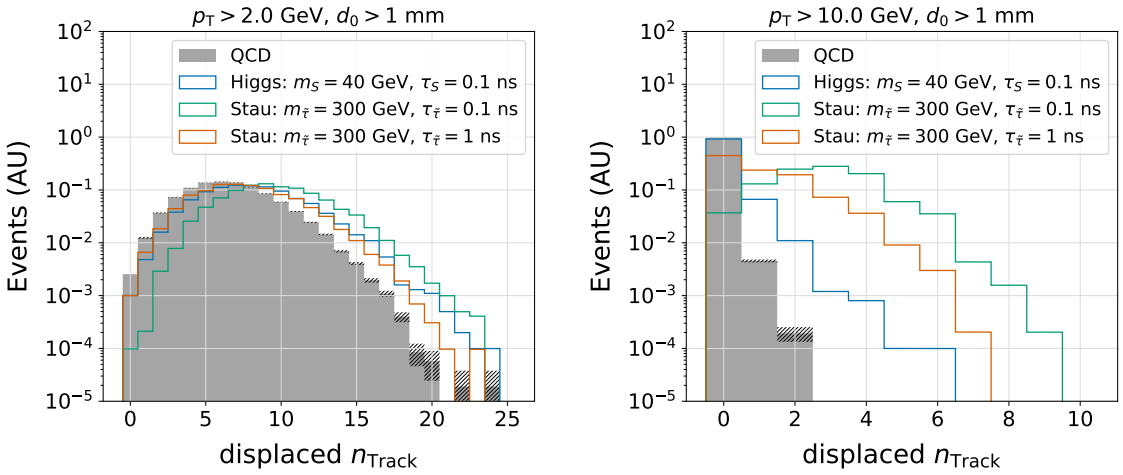

(a) (b)

Figure 25. The number of displaced tracks with $p_T > 2\,\text{GeV}$ (a) and $p_T > 10\,\text{GeV}$ (b) per event for minimum bias QCD and representative long-lived signals with pile-up 200. Slashed bands on the QCD distribution represent the statistical uncertainty.

JHEP02(2023)034

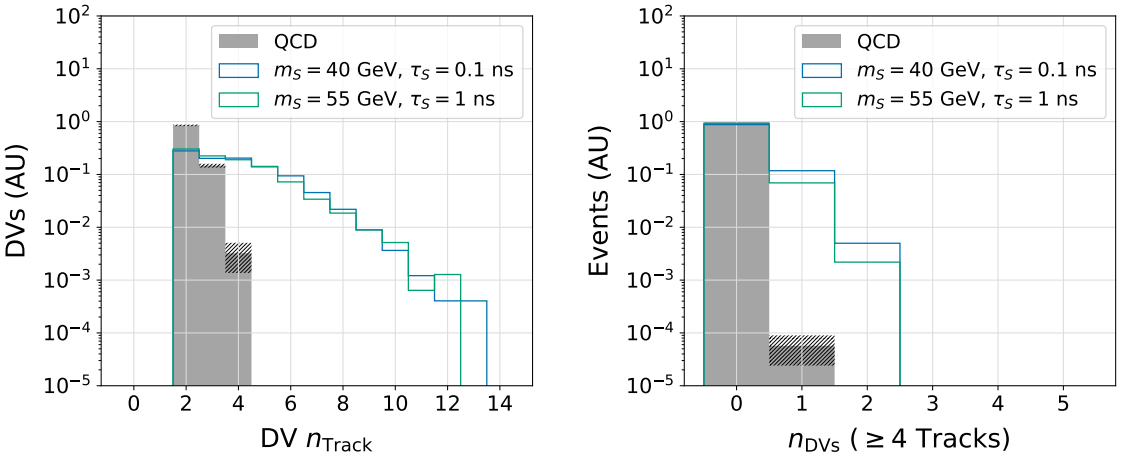

(a) (b)

Figure 26. The number of tracks per displaced vertex (a) and the number of vertices with at least four tracks (b) per event for minimum bias QCD and representative Higgs portal signals with pile-up 200. Slashed bands on the QCD distribution represent the statistical uncertainty.

For this study, a crude displaced vertex reconstruction algorithm is considered. First, all possible pairs of two-track vertices are formed from tracks with $p_{\text{T}} > 1\,\text{GeV}$ and $|d_0| > 1$ mm. The distance of closest approach between the two tracks is required to be < 1.5 mm. Next, any two-track vertices < 1 mm apart are merged. Remaining tracks with $p_{\text{T}} > 1\,\text{GeV}$ and $|d_0| > 0.5$ mm are added to the highest n_{Track} vertex within < 1.5 mm, where one exists. Finally, any ambiguity (tracks associated to multiple vertices) is resolved by uniquely assigning them to the possible vertex with the highest n_{Track}.

Displaced vertices are required to be located within $L_{\text{xy}} < 300$ mm, and $|z| < 300$ mm. To reject real SM long-lived particles, vertices must have a sum of track ${p_{\text{T}}}^2 > 10\,\text{GeV}^2$, $m_{\text{DV}} > 3\,\text{GeV}$, and $L_{\text{xy}} > 4$ mm. Figure 26 shows requiring displaced vertices with at least four tracks results in a feasible trigger with $\sim 10\%$ efficiency. For comparison, triggering on a prompt lepton from associated Higgs production would be less than one tenth as efficient.

An efficient trigger strategy that reduces SM backgrounds to a reasonable rate is possible for all BSM scenarios and tracking parameters considered. Incorporating realistic detector effects and a proper optimization is left for future work.

4 Comparison of tracking parameters

The effect of varying primary tracking parameters is compared across BSM models and signatures in the following sections. The minimum p_{T} threshold is relevant to some extent for all models, primarily affecting SUEPs and the Higgs portal model. Maximum reconstructed $|d_0|$ affects only models with displaced tracks. The final factor considered is the detector layout, as studied through the HSCP model.

JHEP02(2023)034

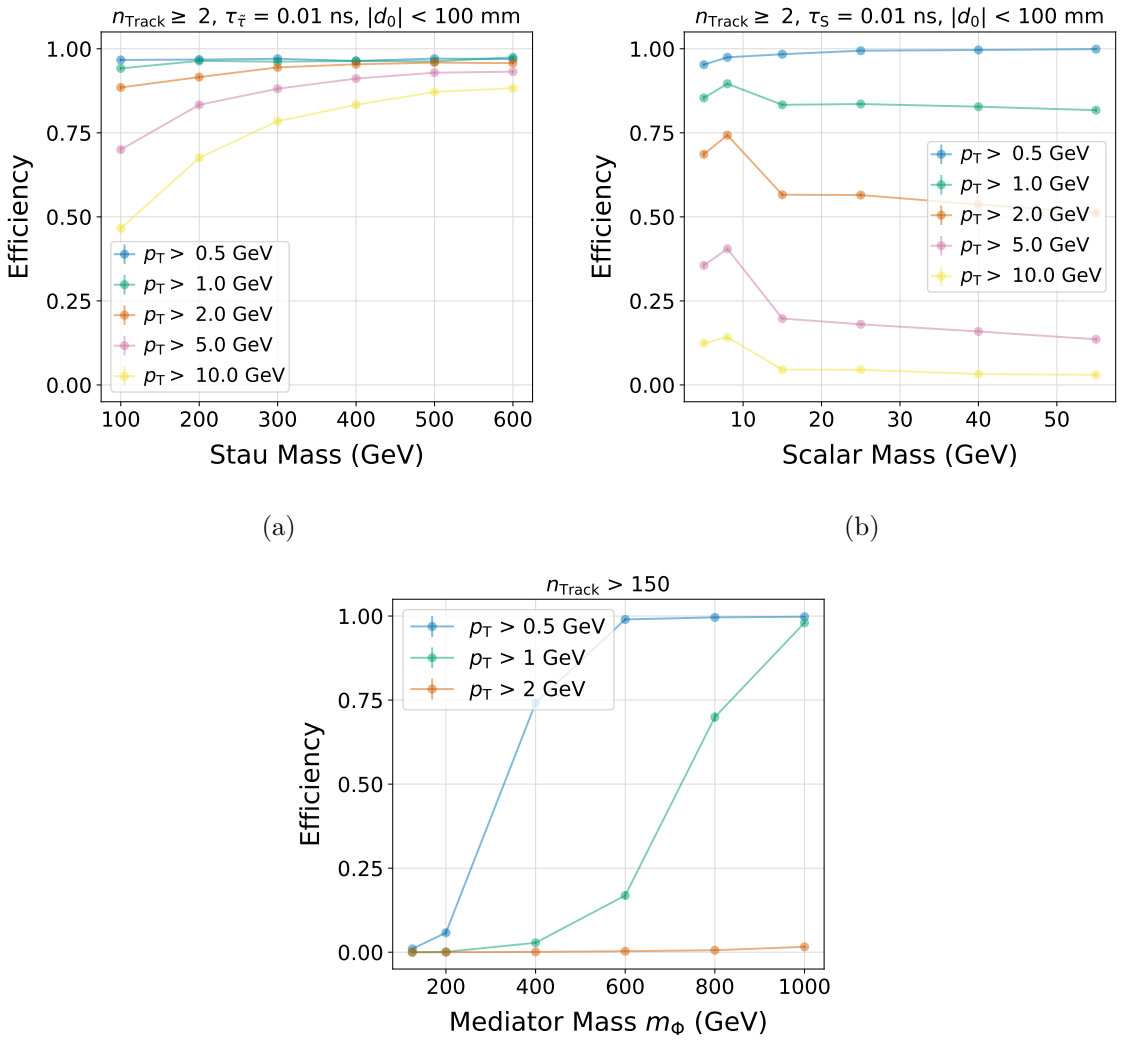

Figure 27. Trigger efficiency as a function of BSM particle mass for the stau (a), Higgs portal (b), and SUEPs (c) models for a range of minimum track p_T thresholds. The stau and Higgs portal samples both correspond to a lifetime of $\tau = 0.01$ ns in order to target displaced decay products. A linearly decreasing efficiency is assumed up to a $|d_0|$ endpoint of 100 mm, and both long-lived scenarios require $n_{\text{Track}} \geq 2$. In the SUEP model all tracks are prompt.

4.1 Minimum transverse momentum

Event-level efficiencies as a function of the minimum track p_T threshold and BSM parent mass are shown for several different models in figure 27. The p_T threshold has a negligible effect on the direct detection of most stable charged particles, and is not considered here. For this comparison, the Higgs portal trigger selection is loosened to consider an event accepted if at least two, rather than five, tracks are successfully reconstructed. This adjustment makes the two selections identical and the results directly comparable.

JHEP02(2023)034

Displaced leptons, like heavy stable charged particles, are relatively robust to the p_{T} threshold, though efficiencies would drop below 50% for the lightest staus with a $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ requirement. The p_{T} threshold has a larger effect on the Higgs portal scenario, where the hadronic decay products are usually softer. Thresholds above 2 GeV would reduce the efficiency of even this two-track trigger to below 20% for all Higgs portal scalar masses. The largest impact of the p_{T} threshold is on the SUEP signature, where all masses have negligible efficiency for $p_{\mathrm{T}} \geq 2\,\mathrm{GeV}$. The lowest reconstruction threshold considered, $p_{\mathrm{T}} \geq 0.5\,\mathrm{GeV}$, would result in a high rate of QCD background. A practical minimum p_{T} for both models would be 1 GeV.

In most systems the reconstruction for prompt and displaced tracks is likely to be controlled and parameterized separately. In these cases, it would be advisable to set a lower p_{T} threshold for prompt track reconstruction than for displaced tracks [56]. This would enable triggering on prompt SUEP scenarios, where a low p_{T} requirement is the most essential, without imposing the same low p_{T} on displaced tracks. However, the p_{T} threshold for displaced tracks would still need to be on the order of 2 to 3 GeV to achieve significant efficiency for the Higgs portal model.

While the minimum track reconstruction p_{T} threshold should be kept as low as possible within the balance of rate and other factors, individual triggers tuned for particular BSM scenarios may benefit from using a much higher minimum p_{T} requirement to suppress background. For example, the efficiency for detection of HSCPs was found to be nearly independent of minimum track p_{T} for cut values up to 100 GeV.

4.2 Impact parameter dependence

Trigger efficiency dependence on the maximum reconstructed $|d_0|$ is shown as a function of long-lived particle lifetime for the stau and Higgs portal models in figure 28. As is the case in the p_{T} comparison, the Higgs portal scenario requires $n_{\mathrm{Track}} \geq 2$, rather than five. The requirements for an event to pass are therefore identical between these two models. Two mass points for each scenario are chosen to illustrate the full range of displaced track kinematics. For the stau model, one low- and one high-mass point are chosen. For the Higgs portal, mass points are chosen such that scalar decays are either primarily to b quarks or to c quarks and taus.

The trends in figure 28 demonstrate that the largest change in efficiency as a result of reduced $|d_0|$ endpoints occurs for intermediate lifetimes ($10^{-2} \lesssim \tau \lesssim 1$ ns) for both benchmark models. For proper lifetimes of $\tau = 0.1$ ns, figure 29 shows the impact of varying the $|d_0|$ endpoint for a variety of stau and scalar masses. The difference in efficiency for different endpoints is consistent for the range of long-lived particle masses considered.

Figure 29 also shows that the impact of $|d_0|$ endpoints is significantly more important for the stau model than for the Higgs portal model. This effect is due to a difference in boost. The long-lived scalar's decay products are more collimated and often point back to the pp collision. As a result, displaced tracks in the Higgs portal model tend to have smaller $|d_0|$ for a given long-lived particle decay position, with respect to the stau model. It is therefore possible to achieve comparable efficiencies between tight and extended $|d_0|$ endpoints in the Higgs portal model. In contrast, the less boosted stau model suffers major efficiency losses as the $|d_0|$ endpoint is reduced from 10 cm to 10 mm.

JHEP02(2023)034

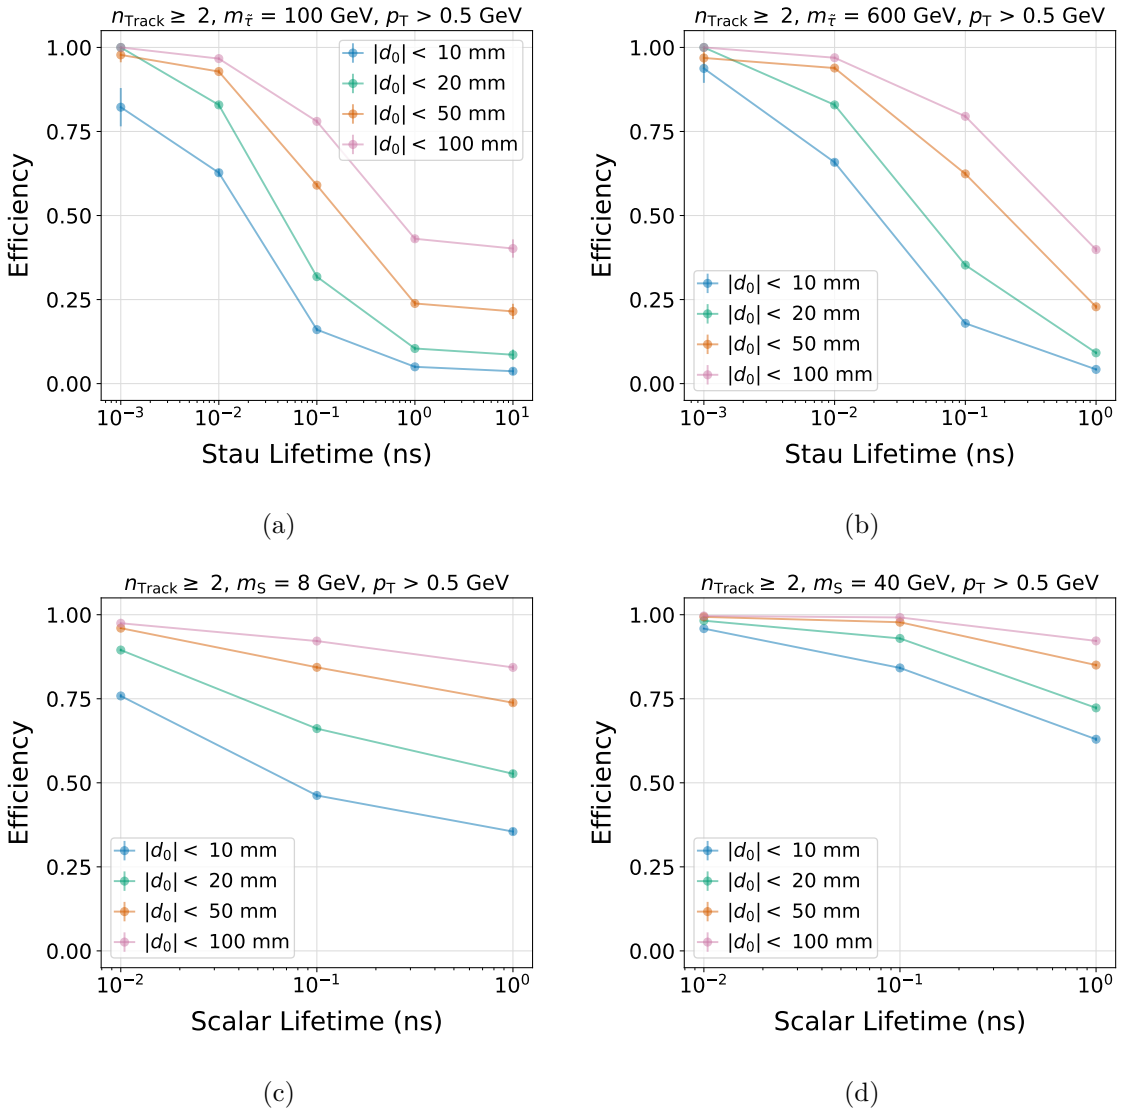

Figure 28. Trigger efficiency as a function of lifetime at a range of maximum track $|d_0|$ values for staus with masses 100 GeV (c) and 600 GeV (b) and a Higgs portal with scalar masses 8 GeV (b) and 40 GeV (d). Only displaced tracks produced in the long-lived particle decay are considered. Effects of the p_{T} threshold are minimized by requiring $p_{\text{T}} > 0.5$ GeV, and both long-lived scenarios require $n_{\text{Track}} \geq 2$.

JHEP02(2023)034

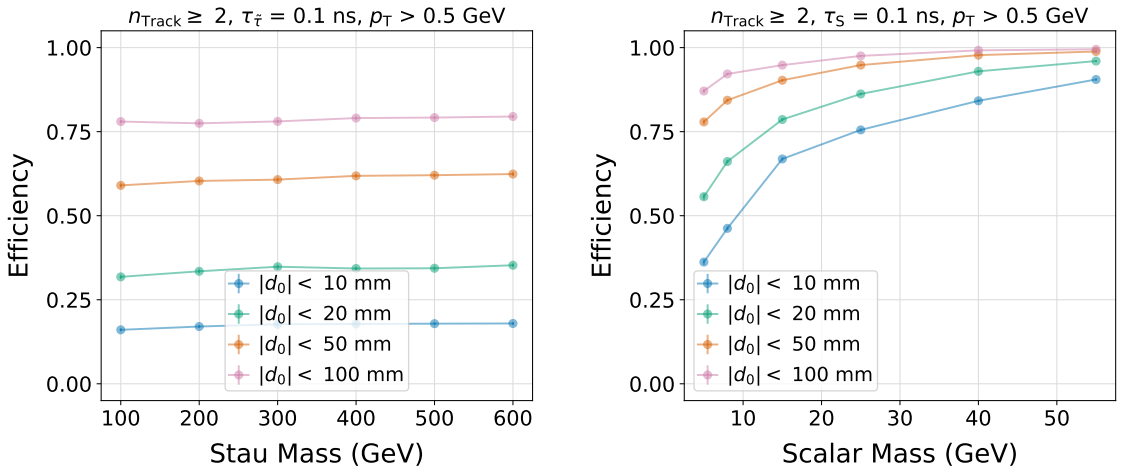

(a) (b)

Figure 29. Event-level efficiency for different $|d_0|$ endpoints as a function of long-lived particle mass. Both the stau (a) and Higgs portal model (b) have lifetimes of $\tau = 0.1$ ns. Only displaced tracks produced in the long-lived particle decay are considered. Effects of the p_{T} threshold are minimized by requiring $p_{\text{T}} > 0.5$ GeV, and both long-lived scenarios require $n_{\text{Track}} \geq 2$.

JHEP02(2023)034

4.3 Number of tracking detector layers

The minimum number of tracking detector layers required per track was studied for the heavy metastable charged particle scenario by requiring the stau to decay beyond a varying minimum L_{xy}. Reducing the minimum L_{xy} from 1.2 to 0.6 m was not found to have a significant effect on the event-level acceptance. For staus with a lifetime of $\tau_{\tilde{\tau}} = 1$ ns, the difference between the most stringent and least stringent requirements on the effective number of tracker layers led to a factor four improvement in event-level acceptance. However, the acceptance remained low in all cases. For even longer lifetimes, most staus decay well beyond $L_{xy} \sim 1$ m. Reducing the minimum number of layers required per track resulted in minimal acceptance improvements.

The effects of the number of tracking layers will still be significant for displaced signatures. The radius at which a long-lived particle decays will directly affect how many tracking layers are available to reconstruct the displaced track. For example, if the particle decays after the first tracking layer, only $n - 1$ hits will be associated to the reconstructed track. The long-lived particle's decay distance also correlates to $|d_0|$, with the mean decay distance longer than the mean $|d_0|$. If a long-lived particle decays beyond multiple tracker layers, there may not be enough hits to reconstruct displaced tracks. The efficiency for reconstructing tracks beyond the equivalent $|d_0|$ would drop to zero.

As seen in the section 4.2, access to high-$|d_0|$ tracks significantly increases the efficiency for both stau and Higgs portal models. For staus, reconstructing high-$|d_0|$ tracks enables the displaced track trigger to outperform the stable charged particle trigger at a lifetime of 1 ns. For the stable charged particle approach, staus with masses between 100 GeV and 500 GeV at 1 ns lifetimes all show an acceptance times efficiency below 15% regardless of the number of layers required per track. For the displaced track approach, the acceptance times efficiency for a stau with a 1 ns lifetime increases from $\sim 0\%$ to $\sim 20\%$ if the $|d_0|$ endpoint is extended from 10 mm to 100 mm, assuming a linearly decreasing efficiency (see appendix A.2). Extending the efficiency to a larger $|d_0|$ endpoint relies on having a sufficient number of tracking layers at larger radii.

5 Conclusions and recommendations

Current ATLAS and CMS triggers are primarily designed to target prompt decays of heavy particles to high-p_T Standard Model particles. However, many compelling BSM scenarios lead to decays which are significantly displaced from the interaction point or include low p_T decay products. For these unconventional scenarios, the anomalous charged particle trajectories are the most distinctive feature of the event. This study explored the impact of varying tracking parameter requirements on trigger efficiency for a collection of representative models.

In general, for these unconventional signatures, perfect track reconstruction efficiency is rarely achievable. However, even introducing a partial track reconstruction efficiency in a p_T or d_0 range can make a substantial difference in sensitivity, and high efficiencies across these parameter spaces are not required. Triggering on 20% of events is infinitely better than triggering on none.

JHEP02(2023)034

HSCP identification is relatively insensitive to choice of tracker parameters, due to their distinctive, prompt, and high-p_{T} tracks. If available, a time of flight estimate is the most discriminating strategy. For HSCPs, the most impactful variation studied was η coverage, with typical efficiencies from a barrel-only scenario at $\sim 30\%$, but those from an extended endcap scenario as high as $\sim 90\%$.

To tackle highly diffuse models such as SUEPs, a track-trigger with a low p_{T} threshold is required. For the SUEP models studied here, the optimal threshold for prompt track reconstruction is $p_{\mathrm{T}} > 1\,\mathrm{GeV}$. Higher thresholds result in low signal efficiency while lower thresholds lead to high QCD rates. For experiments where a low p_{T} threshold is not feasible at L1, alternatives to track counting should be considered, and could be complemented with lower track p_{T} thresholds at HLT. For the SUEP scenario considered here, the overall efficiency is still likely to be low, because high energy initial state radiation is required for any alternate trigger strategy. A p_{T} threshold above 1 GeV will be acceptable for SUEP models with higher temperatures or dark meson masses. In all cases, a primary vertex requirement is essential to separate any SUEP signal from QCD backgrounds at the HL-LHC. The vertex reconstruction method, selection, and longitudinal resolution are intimately connected to a SUEP trigger's performance.

For the other models considered here, a tracking threshold of $p_{\mathrm{T}} > 2\,\mathrm{GeV}$ is sufficient to provide useful sensitivity. To gain sensitivity to the decay products of more massive particles, such as long-lived staus, a large $|d_0|$ range is more essential for overall trigger efficiency than a low p_{T} threshold. Simply counting the number of high-p_{T} displaced tracks is sufficient to reject QCD backgrounds. To address models with light mediators, lower mass long-lived particles, and hadronic decays, the most useful design strategy is a track-trigger that pushes the p_{T} threshold as low as possible with modest $|d_0|$ coverage. In the case of the Higgs portal model, substantial improvement could be gained from reduced p_{T} thresholds. QCD backgrounds pose an additional challenge for these lower p_{T} scenarios, and simple track counting is not sufficient for separating signal from background. Displaced vertex reconstruction or displaced jet tagging will be essential to improve sensitivity to Higgs portal scenarios.

To address all of these cases simultaneously while keeping resource usage to a minimum, the best design scenario is one that allows for increasingly large $|d_0|$ coverage as particle p_{T} increases. In track-trigger designs using pattern banks, this strategy is relatively easy to employ [56]. An approximately constant number of patterns is required to cover a given area in the 2 dimensional space defined by maximum $|d_0|$ and $1/p_{\mathrm{T}}$ values. An optimal pattern bank could be defined by sampling from a triangle rather than a rectangle in that 2D space. A more generic strategy is to create a displaced tracking algorithm that is applied only to tracks passing a minimum p_{T} threshold that is higher than the nominal prompt one. For example, a track-trigger that reconstructs prompt tracks with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$ and displaced tracks with $p_{\mathrm{T}} > 2\,\mathrm{GeV}$ would provide new sensitivity to all models considered here. If the algorithm allows, the p_{T} threshold could be increased along with the maximum $|d_0|$. This strategy is depicted in figure 30.

Track-triggers at the High Luminosity LHC will provide a unique opportunity to directly trigger on these challenging exotic scenarios for the first time. Acting on the recommendations provided here will enable track-trigger designs for the HL-LHC to strike a

JHEP02(2023)034

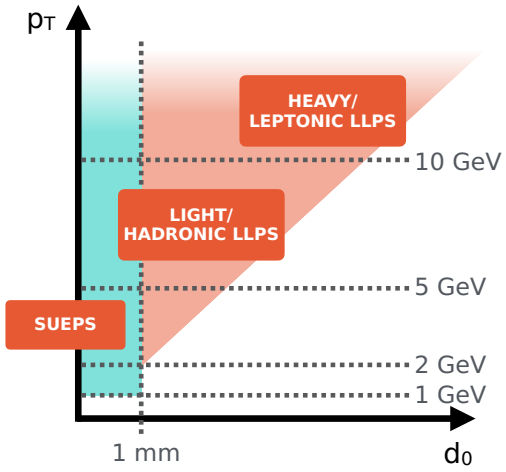

Figure 30. Recommended strategy for optimizing sensitivity to a large range of unconventional signals while minimizing resource usage. Teal represents prompt tracking, while orange represents displaced tracking. HSCP sensitivity is not depicted, but relies on very accessible high-p_{T} prompt tracks.

balance that can best serve a full range of unconventional signatures. The same principles will apply to tracker and trigger designs for future detectors.

Acknowledgments

Tova Holmes and Jesse Farr's research is supported by U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences Energy Frontier Research Centers program under Award Number DE-SC0020267.

Karri Folan DiPetrillo and Chris Guo's work utilizes resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. DiPetrillo's research is also supported in part by the Neubauer Family Foundation Program for Assistant Professors.

Jess Nelson's research is supported by the Duke Research Experiences for Undergraduates program in particle physics, under National Science Foundation Award Number 1757783.

Katherine Pachal's research is supported by TRIUMF, which receives federal funding via a contribution agreement with the National Research Council (NRC) of Canada.

JHEP02(2023)034

A Additional distributions

A.1 Higgs portal

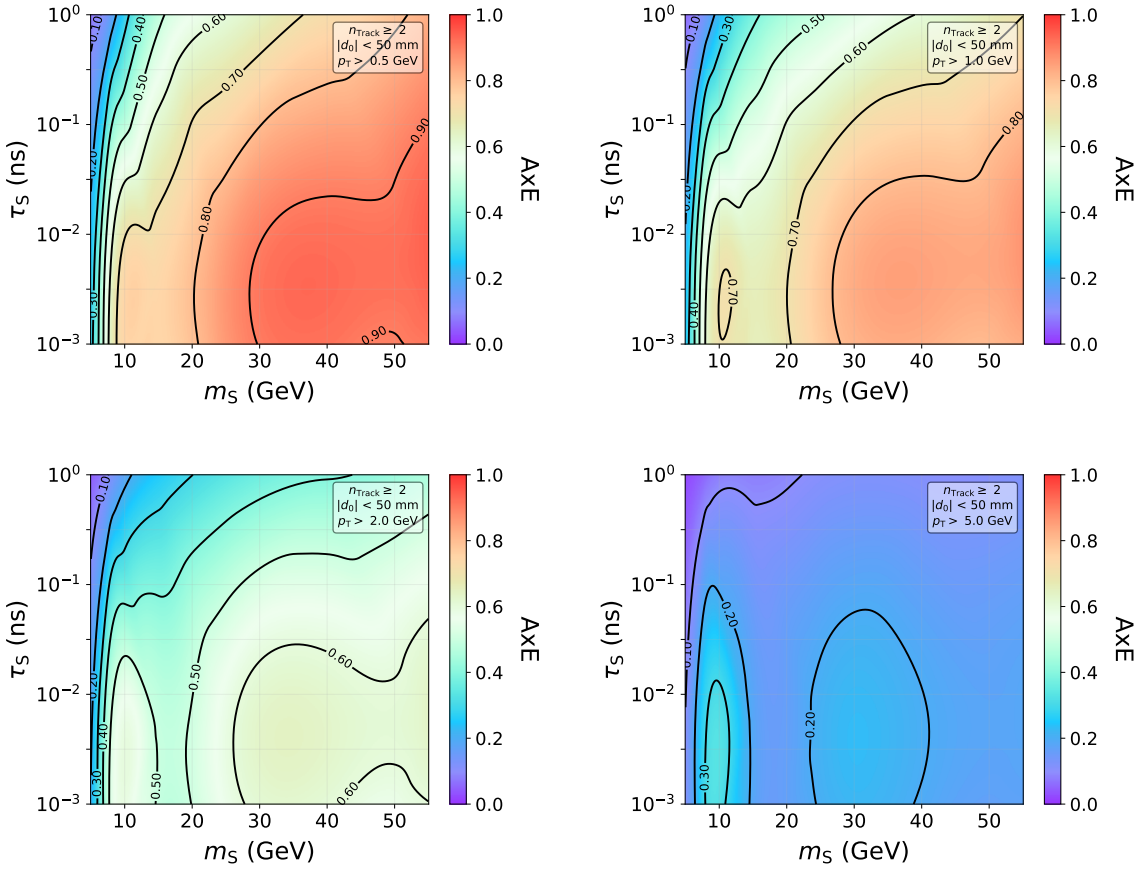

Figure 31. Event-level acceptance times efficiency for Higgs portal model with a constant $|d_0|$ endpoint and p_T thresholds ranging (clockwise from top left to bottom left) from 0.5 Gev, 1.0 GeV, 2.0 Gev, and 5.0 GeV. In all cases, the two-track selection and a linearly decreasing efficiency is used.

JHEP02(2023)034

A.2 Displaced leptons

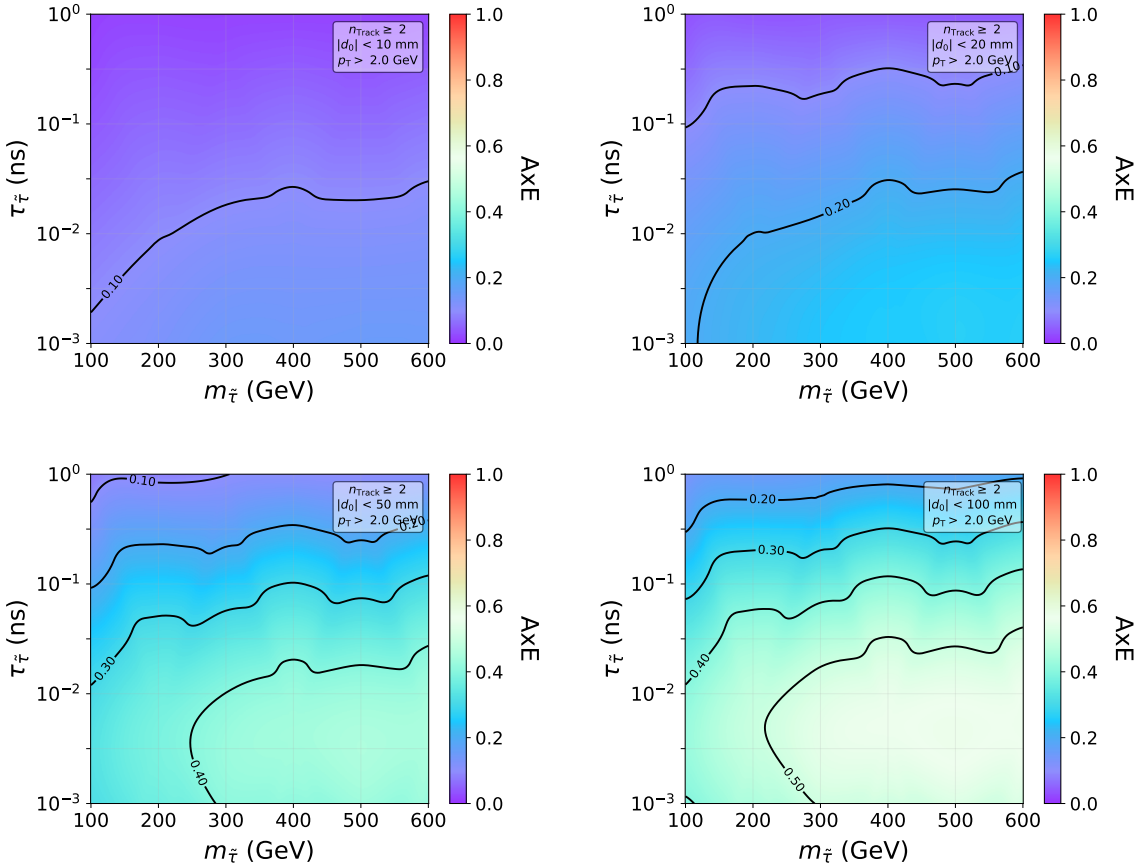

Figure 32. Event-level acceptance times efficiency for a displaced track trigger targeting the stau scenario. The constant p_{T} threshold is kept constant and $|d_0|$ endpoint is varied from (clockwise from top left to bottom left) from 10 mm, 20 mm, 50 mm, and 100 mm. In all cases, the two-track selection and a linearly decreasing efficiency is used.

JHEP02(2023)034

A.3 HSCP

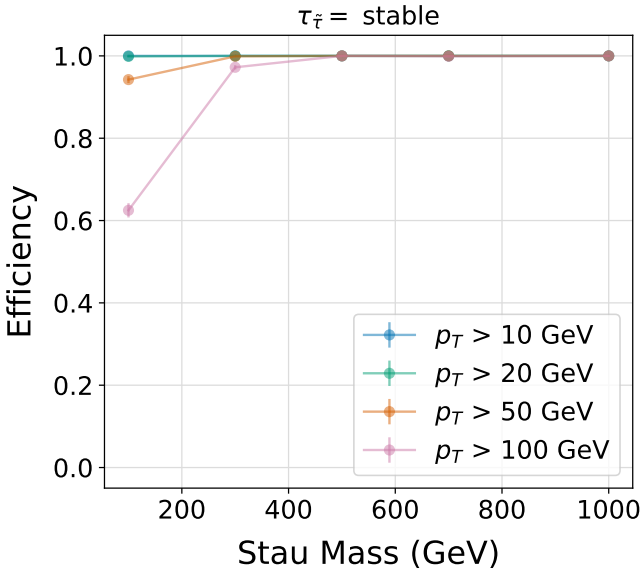

Figure 33. Stau efficiency for a variety of minimum tracker p_{T} values. Efficiency is shown as a function of stau mass for a stable lifetime and minimum $L_{\mathrm{xy}} = 1200$ mm.

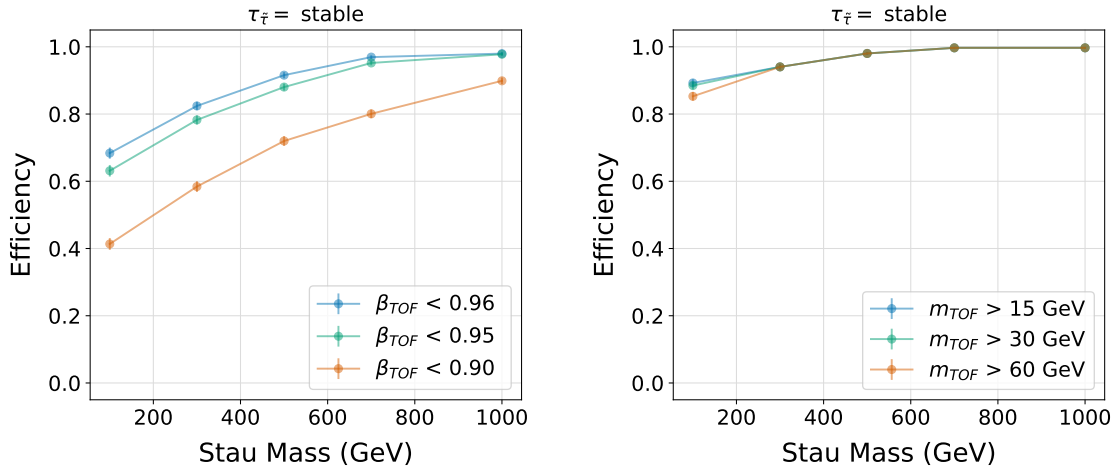

(a) (b)

Figure 34. Efficiency with respect to various requirements on measured track β (a) and measured mass (b).

JHEP02(2023)034

References

[1] L. Lee, C. Ohm, A. Soffer and T.-T. Yu, *Collider Searches for Long-Lived Particles Beyond the Standard Model*, *Prog. Part. Nucl. Phys.* **106** (2019) 210 [*Erratum ibid.* **122** (2022) 103912] [arXiv:1810.12602] [INSPIRE].

[2] S. Knapen, S. Pagan Griso, M. Papucci and D.J. Robinson, *Triggering Soft Bombs at the LHC*, *JHEP* **08** (2017) 076 [arXiv:1612.00850] [INSPIRE].

[3] A. Ryd and L. Skinnari, *Tracking Triggers for the HL-LHC*, *Ann. Rev. Nucl. Part. Sci.* **70** (2020) 171 [arXiv:2010.13557] [INSPIRE].

[4] LHCb collaboration, *Search for Dark Photons Produced in 13 TeV pp Collisions*, *Phys. Rev. Lett.* **120** (2018) 061801 [arXiv:1710.02867] [INSPIRE].

[5] T. Ferber, C. Garcia-Cely and K. Schmidt-Hoberg, *BelleII sensitivity to long-lived dark photons*, *Phys. Lett. B* **833** (2022) 137373 [arXiv:2202.03452] [INSPIRE].

[6] A. Bardi et al., *SVT: An online silicon vertex tracker for the CDF upgrade*, *Nucl. Instrum. Meth. A* **409** (1998) 658 [INSPIRE].

[7] D0CTT Group collaboration, *The D0 central track trigger*, *IEEE Trans. Nucl. Sci.* **51** (2004) 345 [INSPIRE].

[8] S. Bailey et al., *Rapid 3D track reconstruction with the BABAR trigger upgrade*, *Nucl. Instrum. Meth. A* **518** (2004) 544 [INSPIRE].

[9] HERA-B collaboration, *First level trigger for HERA-B*, *Nucl. Instrum. Meth. A* **384** (1996) 131 [INSPIRE].

[10] Y. Iwasaki, B. Cheon, E. Won and G. Varner, *Level 1 trigger system for the Belle II experiment*, in *17th IEEE NPSS Real Time Conference*, (2010) DOI [INSPIRE].

[11] LHCb collaboration, *Design and performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC*, 2019 *JINST* **14** P04013 [arXiv:1812.10790] [INSPIRE].

[12] R. Aaij et al., *Allen: A high level trigger on GPUs for LHCb*, *Comput. Softw. Big Sci.* **4** (2020) 7 [arXiv:1912.09161] [INSPIRE].

[13] CMS collaboration, *Performance of the CMS Level-1 trigger in proton-proton collisions at* $\sqrt{s} = 13$ *TeV*, 2020 *JINST* **15** P10017 [arXiv:2006.10165] [INSPIRE].

[14] CMS collaboration, *The CMS trigger system*, 2017 *JINST* **12** P01020 [arXiv:1609.02366] [INSPIRE].

[15] ATLAS collaboration, *Operation of the ATLAS trigger system in Run 2*, 2020 *JINST* **15** P10004 [arXiv:2007.12539] [INSPIRE].

[16] CMS collaboration, *The Phase-2 Upgrade of the CMS Tracker*, Tech. Rep. CERN-LHCC-2017-009 (2017), DOI [INSPIRE].

[17] CMS collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, Tech. Rep. CERN-LHCC-2020-004 (2020) [INSPIRE].

[18] ATLAS collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*, Tech. Rep. CERN-LHCC-2017-020 DOI [INSPIRE].

JHEP02(2023)034

[19] Y. Gershtein, *CMS Hardware Track Trigger: New Opportunities for Long-Lived Particle Searches at the HL-LHC*, *Phys. Rev. D* **96** (2017) 035027 [arXiv:1705.04321] [INSPIRE].

[20] ATLAS collaboration, *FTK public results*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/FTKPublicResults (2019).

[21] M. Mårtensson, M. Isacson, H. Hahne, R. Gonzalez Suarez and R. Brenner, *To catch a long-lived particle: hit selection towards a regional hardware track trigger implementation*, 2019 *JINST* **14** P11009 [arXiv:1907.09846] [INSPIRE].

[22] B. Bhattacherjee, S. Mukherjee, R. Sengupta and P. Solanki, *Triggering long-lived particles in HL-LHC and the challenges in the rst stage of the trigger system*, *JHEP* **08** (2020) 141 [arXiv:2003.03943] [INSPIRE].

[23] ATLAS collaboration, *Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector*, Tech. Rep. ATL-PHYS-PUB-2017-014 (2017) [INSPIRE].

[24] M.J. Strassler, *Why Unparticle Models with Mass Gaps are Examples of Hidden Valleys*, arXiv:0801.0629 [INSPIRE].

[25] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [arXiv:1410.3012] [INSPIRE].

[26] C. Cesarotti and J. Thaler, *A Robust Measure of Event Isotropy at Colliders*, *JHEP* **08** (2020) 084 [arXiv:2004.06125] [INSPIRE].

[27] I. Tomalin et al., *An FPGA based track finder for the L1 trigger of the CMS experiment at the High Luminosity LHC*, 2017 *JINST* **12** P12019 [INSPIRE].

[28] ATLAS collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673 [arXiv:1704.07983] [INSPIRE].

[29] CMS collaboration, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, 2014 *JINST* **9** P10009 [arXiv:1405.6569] [INSPIRE].

[30] C. Cesarotti, M. Reece and M.J. Strassler, *The efficacy of event isotropy as an event shape observable*, *JHEP* **07** (2021) 215 [arXiv:2011.06599] [INSPIRE].

[31] J.F. Donoghue, F.E. Low and S.-Y. Pi, *Tensor Analysis of Hadronic Jets in Quantum Chromodynamics*, *Phys. Rev. D* **20** (1979) 2759 [INSPIRE].

[32] ATLAS collaboration, *Search for long-lived neutral particles produced in pp collisions at* $\sqrt{s} = 13$ *TeV decaying into displaced hadronic jets in the ATLAS inner detector and muon spectrometer*, *Phys. Rev. D* **101** (2020) 052013 [arXiv:1911.12575] [INSPIRE].

[33] J.D. Wells, *How to Find a Hidden World at the Large Hadron Collider*, Tech. Rep. MCTP-07-51 (2008) [INSPIRE].

[34] S. Gopalakrishna, S. Jung and J.D. Wells, *Higgs boson decays to four fermions through an abelian hidden sector*, *Phys. Rev. D* **78** (2008) 055002 [arXiv:0801.3456] [INSPIRE].

[35] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [arXiv:1405.0301] [INSPIRE].

[36] D. Curtin et al., *Exotic decays of the 125 GeV Higgs boson*, *Phys. Rev. D* **90** (2014) 075004 [arXiv:1312.4992] [INSPIRE].

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

JHEP02(2023)034

[38] OPAL collaboration, *Searches for gauge-mediated supersymmetry breaking topologies in e^+e^- collisions at LEP2*, *Eur. Phys. J. C* **46** (2006) 307 [hep-ex/0507048] [INSPIRE].

[39] ATLAS collaboration, *Search for heavy, long-lived, charged particles with large ionisation energy loss in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS experiment and the full Run 2 dataset*, *arXiv e-prints* (2022) arXiv:2205.06013 [arXiv:2205.06013].

[40] CMS collaboration, *Search for disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Lett. B* **806** (2020) 135502 [arXiv:2004.05153] [INSPIRE].

[41] ATLAS collaboration, *Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **82** (2022) 606 [arXiv:2201.02472] [INSPIRE].

[42] ATLAS collaboration, *Search for Displaced Leptons in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector*, *Phys. Rev. Lett.* **127** (2021) 051802 [arXiv:2011.07812] [INSPIRE].

[43] CMS collaboration, *Search for long-lived particles decaying to leptons with large impact parameter in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **82** (2022) 153 [arXiv:2110.04809] [INSPIRE].

[44] J.A. Evans and J. Shelton, *Long-Lived Staus and Displaced Leptons at the LHC*, *JHEP* **04** (2016) 056 [arXiv:1601.01326] [INSPIRE].

[45] J. Alwall, P. Schuster and N. Toro, *Simplified Models for a First Characterization of New Physics at the LHC*, *Phys. Rev. D* **79** (2009) 075020 [arXiv:0810.3921] [INSPIRE].

[46] LHC New Physics Working Group collaboration, *Simplified Models for LHC New Physics Searches*, *J. Phys. G* **39** (2012) 105005 [arXiv:1105.2838] [INSPIRE].

[47] ATLAS collaboration, *Search for long-lived, massive particles in events with a displaced vertex and a muon with large impact parameter in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **102** (2020) 032006 [arXiv:2003.11956] [INSPIRE].

[48] CMS collaboration, *A MIP Timing Detector for the CMS Phase-2 Upgrade*, Tech. Rep. CERN-LHCC-2019-003 (2019) [INSPIRE].

[49] ATLAS collaboration, *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade: Technical Design Report*, Tech. Rep. CERN-LHCC-2020-007, ATLAS-TDR-031 [INSPIRE].

[50] ATLAS collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton-proton collisions at the LHC*, *Eur. Phys. J. C* **77** (2017) 332 [arXiv:1611.10235] [INSPIRE].

[51] E. Bartz et al., *FPGA-based tracking for the CMS Level-1 trigger using the tracklet algorithm*, 2020 *JINST* **15** P06024 [arXiv:1910.09970] [INSPIRE].

[52] CMS collaboration, *The Phase-2 Upgrade of the CMS L1 Trigger Interim Technical Design Report*, Tech. Rep. CERN-LHCC-2017-013, CERN, Geneva (2017), DOI.

[53] CMS collaboration, *Search for long-lived particles using displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **104** (2021) 012015 [arXiv:2012.01581] [INSPIRE].

[54] ATLAS collaboration, *Triggers for displaced decays of long-lived neutral particles in the ATLAS detector*, 2013 *JINST* **8** P07015 [arXiv:1305.2284] [INSPIRE].

[55] Y. Gershtein, S. Knapen and D. Redigolo, *Probing naturally light singlets with a displaced vertex trigger*, *Phys. Lett. B* **823** (2021) 136758 [arXiv:2012.07864] [INSPIRE].

[56] ATLAS collaboration, *The ATLAS Fast TracKer system*, 2021 *JINST* **16** P07006 [arXiv:2101.05078] [INSPIRE].

JHEP02(2023)034