



On the LHC sensitivity for non-thermalised hidden sectors [☆]

Felix Kahlhoefer

Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, D-52056 Aachen, Germany



ARTICLE INFO

Article history:

Received 6 February 2018

Accepted 14 February 2018

Available online 21 February 2018

Editor: A. Ringwald

Keywords:

Cosmology of theories beyond the Standard Model

LHC searches for hidden sectors

ABSTRACT

We show under rather general assumptions that hidden sectors that never reach thermal equilibrium in the early Universe are also inaccessible for the LHC. In other words, any particle that can be produced at the LHC must either have been in thermal equilibrium with the Standard Model at some point or must be produced via the decays of another hidden sector particle that has been in thermal equilibrium. To reach this conclusion, we parametrise the cross section connecting the Standard Model to the hidden sector in a very general way and use methods from linear programming to calculate the largest possible number of LHC events compatible with the requirement of non-thermalisation. We find that even the HL-LHC cannot possibly produce more than a few events with energy above 10 GeV involving states from a non-thermalised hidden sector.

© 2018 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>). Funded by SCOAP³.

1. Introduction

One of the central motivations for the LHC has been to probe models of dark matter (DM) in which the DM candidate is a weakly interacting massive particle (WIMP), meaning that it has mass and interaction strength comparable to other known particles (i.e. comparable to the electroweak scale). In such models the DM particle is predicted to be in thermal equilibrium with the bath of Standard Model (SM) particles, so that its relic abundance can be calculated via the freeze-out mechanism [1].

The non-observation of new physics at the LHC has however cast doubt on the idea of WIMP DM and has led to increased interest in alternative ideas. A simple modification of the WIMP framework is to assume that even though DM was in thermal equilibrium in the early Universe, the annihilation processes that are responsible for setting its relic abundance are *secluded* from the SM and therefore much more difficult to observe at the LHC [2]. Within this framework (recently coined “WIMP next door” [3]) it is easily possible to reconcile LHC constraints with the idea of thermal freeze-out [4].

Nevertheless, LHC results may also be taken to indicate that DM was not produced via thermal freeze-out and that in fact the DM particle was never in thermal equilibrium with the SM thermal bath. Among the many possible alternatives (collectively referred to as FIMPs for feebly-interacting massive particles [5]) are models

in which DM is produced via out-of-equilibrium decays (so-called superWIMPs [6]) or via the freeze-in mechanism [7]. A question of general interest is therefore whether the LHC may be able to also shed light on models in which the dark sector is so weakly coupled to the SM that it was never in thermal equilibrium.

In this letter we address this question by comparing interaction rates at the LHC and in the early Universe. This comparison is based on the following simple observations:

1. If the same physics can be used to describe the early Universe and the LHC, any process that can be induced at the LHC also took place in the early Universe (provided the process happens on short enough timescales). Given a temperature of the SM thermal bath, it is straight-forward to calculate the rate at which this process must have occurred.
2. Assuming that SM states dominate the energy density of the early Universe, we can directly compare the rate of any such process (at a given temperature) to the corresponding Hubble rate. If it exceeds the Hubble rate, the process happened sufficiently frequently that the final state particles were produced with large abundance. This typically implies that the inverse process must also have occurred frequently, leading to thermal equilibrium between initial and final states.

It is straight-forward to perform such a comparison in the context of specific models. Here we demonstrate that it is in fact also possible to compare rates at the LHC and in the early Universe in a model-independent way. To do so, we evaluate the thermalisation

[☆] This article is registered under preprint number: TTK-18-03.

E-mail address: kahlhoefer@physik.rwth-aachen.de.

condition between the SM and a hidden sector¹ for a very general functional form of the cross section of the process linking the two sectors. Requiring that the two sectors do not thermalise then leads to constraints on this cross section, which can be translated into an upper bound on the rate of the same process at the LHC. This approach allows us to show that any process that does not reach thermal equilibrium in the early Universe can only induce a negligibly small number of events with sufficiently high energy at the LHC.

We conclude that, provided the Universe once was at a temperature comparable to LHC energies, any hidden sector that we could hope to observe at the LHC must have been at least partially in thermal equilibrium in the early Universe. Conversely, any hidden sector that was not in thermal equilibrium in the early Universe is unobservable for the LHC.

This letter is structured as follows. In section 2 we present the derivation of the upper bound on the number of observable LHC events, first for a rather specific assumption on the functional form of the cross section and then using a general parametrisation. A detailed discussion of the underlying assumptions is provided in section 3. Appendix A takes a closer look at the case of s -channel resonances.

2. Thermal equilibrium and LHC predictions

We consider a process that converts two SM particles into one or more hidden sector states. Rather than specifying the details of the underlying interactions, we simply assume that the total cross section σ of this process depends in some way on the centre-of-mass energy \sqrt{s} of the colliding particles. The fundamental requirement we impose is that the cross section is sufficiently small that the hidden sector states never reach thermal equilibrium (in short, we require *non-thermalisation*). To make this statement quantitative, we define the reaction rate Γ as the product of the thermally averaged cross section times relative velocity and the equilibrium number density of the SM particles in the initial state [8]:

$$\Gamma \equiv \langle \sigma v \rangle n^{\text{eq}} = \int \frac{N_c s^2 K_1(\sqrt{s}/T)}{4\pi^2 T^2} \sigma(\sqrt{s}) d\sqrt{s}, \quad (1)$$

where N_c denotes the number of colour degrees of freedom in the initial state and we have assumed the masses of the SM particles to be negligible.² The requirement that no thermal equilibrium is achieved can then be rephrased as

$$\Gamma(T) < H(T) = \sqrt{\frac{4\pi^3 g_*}{45}} \frac{T^2}{M_{\text{pl}}} \quad (2)$$

for all temperatures T , where g_* counts the number of effective relativistic degrees of freedom at temperature T and M_{pl} denotes the Planck mass.

It will be convenient to rewrite this requirement by introducing the dimensionless reaction rate

$$\begin{aligned} \gamma(T) &\equiv \frac{\Gamma(T)}{H(T)} \\ &= \int d\sqrt{s} \sqrt{\frac{45}{\pi g_*}} \frac{N_c M_{\text{pl}} s^2 K_1(\sqrt{s}/T)}{8\pi^3 T^4} \sigma(\sqrt{s}) \end{aligned} \quad (3)$$

The non-thermalisation constraint then simply becomes $\gamma(T) < 1$ for all T .

Let us now consider the same process at the LHC. Clearly, the process will only be relevant if both of the particles in the initial state can be supplied via proton-proton collisions.³ In this case, the leading-order production cross section at the LHC is given by

$$\sigma_{\text{LHC}} = \int dx_1 dx_2 f_1(x_1) f_2(x_2) \sigma(\sqrt{s_{\text{tot}} x_1 x_2}), \quad (4)$$

where $\sqrt{s_{\text{tot}}}$ is the total centre-of-mass energy of the LHC, x_i denote the fraction of the proton momentum carried by the particles in the initial state and f_i denote the parton distribution functions (pdfs). Here we use the MSTW 2008 NNLO pdfs [10], setting the factorisation scale to the partonic centre-of-mass energy $\mu_F = E_{\text{cm}} \equiv \sqrt{s_{\text{tot}} x_1 x_2}$.

The total number of production processes of hidden sector states that will occur at the LHC is then given by $N_{\text{LHC}} = \sigma_{\text{LHC}} \mathcal{L}$, where \mathcal{L} denotes the integrated luminosity. Again, it will be convenient to rewrite this expression slightly:

$$N_{\text{LHC}} = \int d\sqrt{s} \frac{dx}{x} f_1(x) f_2\left(\frac{s}{s_{\text{tot}} x}\right) \frac{2\mathcal{L}\sqrt{s}}{s_{\text{tot}}} \sigma(\sqrt{s}). \quad (5)$$

2.1. A first example

Our aim is to find the functional form $\sigma(\sqrt{s})$ that maximises N_{LHC} while satisfying the constraint $\gamma(T) < 1$. To gain some intuition on the nature of these requirements, let us first make a very simple ansatz and write

$$\sigma(\sqrt{s}) = \sigma_0 \sqrt{s_0} \delta(\sqrt{s} - \sqrt{s_0}), \quad (6)$$

where $\delta(x)$ is the Dirac δ -function. Such a cross section could arise for example from the exchange of an on-shell s -channel mediator connecting the SM initial state to the hidden sector (see Appendix A for further details). With this ansatz we find

$$\begin{aligned} \gamma(T) &= \sqrt{\frac{45}{\pi g_*}} \frac{N_c M_{\text{pl}} \sqrt{s_0}^5 K_1(\sqrt{s_0}/T)}{8\pi^3 T^4} \sigma_0 \\ &= \sqrt{\frac{45}{\pi g_*}} \frac{N_c \sqrt{s_0} M_{\text{pl}} x_0^4 K_1(x_0)}{8\pi^3} \sigma_0, \end{aligned} \quad (7)$$

where in the second line we have defined $x_0 = \sqrt{s_0}/T$. It is straight-forward to see that if we vary T (or equivalently x_0), the expression is maximised for $x_0 \approx 3.414$. Thus, the constraint $\gamma(T) < 1$ is satisfied for all T exactly if

$$\sigma_0 < 6.45 \cdot 10^{-6} \text{ fb} \left(\frac{1 \text{ GeV}}{\sqrt{s_0}} \right) \frac{1}{N_c} \left(\frac{g_*(T_0)}{106.75} \right)^{1/2}, \quad (8)$$

where $T_0 = \sqrt{s_0}/3.414$.

Clearly, the non-thermalisation constraint requires the cross section connecting SM states to the hidden sector to be extremely small. We observe in particular that the constraint becomes more

¹ We use the term hidden sector to refer to all particles beyond the SM that are either stable or sufficiently long-lived to appear stable in the early Universe, as well as any other states that decay dominantly into these new particles.

² We note that for temperatures much larger than the masses of the particles in the initial state, our implicit assumption of a Maxwell-Boltzmann distribution for n^{eq} is not necessarily a good approximation. However, we do not want to limit ourselves to the case of either only fermionic or only bosonic initial states and therefore proceed with the simpler expression, noting that a more accurate treatment may change our results by up to a factor of two [3,9].

³ Hidden sector states could also be radiated from a final state, but the corresponding cross sections will be much smaller.




Fig. 1. Upper bound on the number of events at the HL-LHC for a cross section of the form $\sigma(\sqrt{s}) = \sigma_0 \sqrt{s_0} \delta(\sqrt{s} - \sqrt{s_0})$ for different initial states. For each value of $\sqrt{s_0}$ the coefficient σ_0 is fixed in such a way that the cross section saturates the non-thermalisation constraint. Shaded bands indicated pdf uncertainties at the 90% confidence level.

stringent with increasing $\sqrt{s_0}$. The reason is that in this case the non-thermalisation constraint becomes sensitive to higher energies, corresponding to earlier times and hence larger densities in the early Universe.

We can make use of the upper bound on σ_0 to substitute our ansatz into eq. (5) and calculate the maximum number of events that can be expected at the LHC for a cross section of this particular form. The result is shown in Fig. 1 for an integrated luminosity of $\mathcal{L} = 3000 \text{ fb}^{-1}$ for three different combinations of initial states: gg , $u\bar{u}$ and $d\bar{d}$. We make the following observations:

1. The predicted number of events depends sensitively on $\sqrt{s_0}$. This is a direct consequence both of the non-thermalisation constraint becoming weaker for smaller $\sqrt{s_0}$ and of the pdfs preferring smaller partonic centre-of-mass energies.
2. We find significantly larger values of N_{LHC} for the case of a gg initial state than for a $u\bar{u}$ initial state, again a direct consequence of the larger gluon pdfs at small momentum fraction.
3. In any case, the upper bounds on N_{LHC} are extremely tight. Even for $\sqrt{s_0} = 10 \text{ GeV}$ and a gg initial state, we find $N_{\text{LHC}} < 0.49$, making a discovery of this process impossible.⁴

In principle, the bound on N_{LHC} could be further relaxed by going to even smaller values of $\sqrt{s_0}$. However, at this point our description becomes increasingly questionable, because non-perturbative effects become important for the description of the initial state. Moreover, processes with such small partonic centre-of-mass energy are very challenging to observe at the LHC and would be more suitable for searches at low-energy colliders like Belle II. To avoid these complications, we will not consider values of $\sqrt{s_0}$ below 10 GeV. In fact, even for $\sqrt{s_0} = 10 \text{ GeV}$ pdf uncertainties are already substantial. We estimate these uncertainties by varying the pdfs around the central set following the procedure described in ref. [10]. The shaded bands in Fig. 1 illustrate the resulting changes in N_{LHC} at the 90% confidence level. For $\sqrt{s_0} = 10 \text{ GeV}$ the uncertainty amounts to approximately 20–30%.

We conclude that the simple ansatz chosen above leads to a very stringent bound on N_{LHC} . Nevertheless, it is far from clear that this ansatz comes close to the optimal functional form of the cross section. After all, we found that $\gamma(T) \approx 1$ only for a very small range of temperatures. We therefore expect that it should be

possible to relax the bound on N_{LHC} by considering more general forms of $\sigma(\sqrt{s})$.

2.2. General cross sections

Since it is difficult to consider completely arbitrary variations of $\sigma(\sqrt{s})$, we will instead consider a discretised version of the problem (see ref. [11] for a similar approach in the context of DM direct detection experiments). For this purpose, we define a (potentially large) number of discrete values for \sqrt{s} :

$$\sqrt{s_i} = \sqrt{s_0} \Delta^i, \quad (9)$$

where $i = 0, \dots, m$ and $\Delta > 1$ defines the step size. Since we are only interested in quantities involving the integral of $\sigma(\sqrt{s})$, we can then approximate the cross section by

$$\sigma(\sqrt{s}) \approx \sum_i \sigma_i \sqrt{s_i} (\Delta - 1) \delta(\sqrt{s} - \sqrt{s_i}) \quad (10)$$

with $\sigma_i = \sigma(\sqrt{s_i})$. This approximation can be made for any cross section that does not vary too rapidly on each interval $[\sqrt{s_i}, \sqrt{s_{i+1}}]$, and it becomes exact in the limit $\Delta \rightarrow 1$. Rather than considering cross sections of arbitrary functional form, we can therefore simply study cross sections of the form of eq. (10) with arbitrary coefficients σ_i . As we will see below, writing $\sigma(\sqrt{s})$ as a sum of δ -functions (rather than e.g. as a piecewise constant function) has the advantage that it leads to a significant simplification of the relevant equations.

Making use of the approximation introduced in eq. (10), the number of predicted events at the LHC can be written as

$$N_{\text{LHC}} = \sum_i b_i \sigma_i \quad (11)$$

with

$$b_i = (\Delta - 1) \int \frac{dx}{x} f_1(x) f_2\left(\frac{s_i}{s_{\text{tot}} x}\right) \frac{2 \mathcal{L} s_i}{s_{\text{tot}}}. \quad (12)$$

The second important simplification that we make is to assume that it is sufficient to check the non-thermalisation constraint only for discrete values of T . Specifically, we define

$$T_j = T_0 \Delta^j \quad (13)$$

for $j = 1, \dots, n$ and appropriate values $T_0 < \sqrt{s_0}$ and $n > m$, such that all relevant temperatures are covered. With this definition, we find that $\gamma_j \equiv \gamma(T_j)$ can be written as

$$\gamma_j = \sum_i a_{ji} \sigma_i \quad (14)$$

with

$$a_{ij} = (\Delta - 1) \sqrt{\frac{45}{\pi g_*}} \frac{N_c M_{\text{pl}} \sqrt{s_i}^5 K_1(\sqrt{s_i}/T_j)}{8\pi^3 T_j^4}. \quad (15)$$

We have therefore reduced the problem to a well-known problem of linear programming: the maximisation of $\mathbf{b} \cdot \mathbf{x}$ with respect to a vector $\mathbf{x} = (\sigma_i)$ with positive entries for a given vector \mathbf{b} , subject to a set of constraints given by $\mathbf{A} \cdot \mathbf{x} \leq \mathbf{1}$. This optimisation problem can be easily solved with well-known numerical methods. For sufficiently small step size Δ and sufficiently large m and n , the value of the maximum should become independent of the precise value of Δ and of the choice of T_0 . We find that this is indeed the case, and that choosing $\Delta = 1.1$ (corresponding to $m, n \sim 100$) is fully sufficient.

⁴ In principle, even a single event could allow for a discovery, if the signature is so distinctive that it cannot be missed or confused with SM backgrounds. Since we do not specify the details of the hidden sector, we cannot entirely exclude this possibility although it seems rather unlikely.

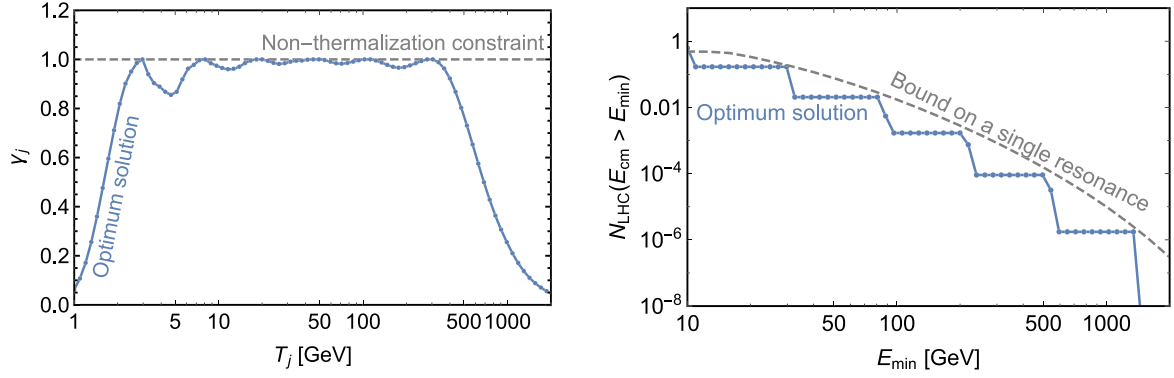


Fig. 2. Illustration of the optimum cross section, which satisfies the non-thermalisation constraint for all temperatures (visualised in the left panel) while maximizing the number of predicted events at the LHC (visualised in the right panel). This example considers a gg initial state with an integrated luminosity of $\mathcal{L} = 3000 \text{ fb}^{-1}$ at $\sqrt{s_{\text{tot}}} = 13 \text{ TeV}$. The low-energy cut-off has been set to $\sqrt{s_0} = 10 \text{ GeV}$. For easier visualisation we show continuous lines connecting the discrete values used for the optimisation. The dashed lines indicate the constraint $\gamma_i < 1$ in the left panel and the bound on N_{LHC} obtained for a single δ -function with $\sqrt{s_0} = E_{\text{min}}$ in the right panel (see Fig. 1).

2.3. Results

For concreteness, let us consider the case of a gg initial state and set $\sqrt{s_0} = 10 \text{ GeV}$ and $T_0 = 1 \text{ GeV}$. For $\sqrt{s_{\text{tot}}} = 13 \text{ TeV}$ and $\mathcal{L} = 3000 \text{ fb}^{-1}$ we find $N_{\text{LHC}} < 0.62$. This value should be compared with the value $N_{\text{LHC}} < 0.49$ that we found for the much simpler ansatz above. In fact, the similarity of the two results is no coincidence, as will become clear by inspecting the optimum solution.

The central observation is that the optimum solution has $\sigma_i = 0$ for almost all values of i . The few non-zero values of σ_i are well-separated and each (nearly) saturate the bound derived above for the case of a single δ -function. In this way, the individual contributions conspire to give $\gamma(T) \approx 1$ over a wide range of temperatures. This behaviour is illustrated in the left panel of Fig. 2, which shows that in this particular case the optimum solution consists of a sum of six separate δ -functions.

We emphasise that the number of δ -functions contributing to the optimum solution is not an artefact of the chosen step size, which is in fact much smaller than the distance between two δ -functions. Instead, it is a generic feature of the optimum solution and is robust under a change of step size. In other words, the “distance” between the individual δ -functions is optimal – introducing any additional contribution would strengthen the non-thermalisation constraint and hence lead to a stronger bound on N_{LHC} .

This observation makes it clear why the optimum solution is so close to the one obtained for a single δ -function: The dominant contribution to N_{LHC} stems from the δ -function at $\sqrt{s_0}$, while the additional δ -functions at higher energies only give sub-leading contributions. This is illustrated in the right panel of Fig. 2, which shows the predicted number of LHC events with partonic centre-of-mass energy E_{cm} above a minimum energy E_{min} as a function of that minimum energy. The dashed grey line shows for comparison the bound on N_{LHC} obtained for a single δ -function with $\sqrt{s_0} = E_{\text{min}}$ (see Fig. 1). For the specific case considered here, we find that the first three δ -functions contribute 0.45, 0.15 and 0.019 to the total of 0.62 predicted events.

To conclude this discussion, we emphasise that the numbers discussed above depend of course on the somewhat arbitrary choice of $\sqrt{s_0}$. To first approximation, reducing $\sqrt{s_0}$ by a factor r will relax the bound on N_{LHC} by the same factor. However, even if we allow $\sqrt{s_0}$ as low as 1 GeV (and consider $\sqrt{s_{\text{tot}}} = 14 \text{ TeV}$), we cannot obtain more than $\mathcal{O}(10)$ events with in full data set of the HL-LHC, almost all of which would be at extremely low energies and therefore likely unobservable.

3. Discussion

We have shown in a very general way that any process connecting the SM to a hidden sector that was not in thermal equilibrium in the early Universe is unobservable at the LHC. Turning this argument around, we conclude that any hidden sector that can be produced and observed at the LHC must have been at least partially in thermal equilibrium with the bath of SM particles in the Early Universe.

We emphasise that this does not necessarily mean that the LHC can only probe DM models in which the DM particle itself was in thermal equilibrium. For example, the DM particle may be produced from the decays of another metastable hidden sector particle, which itself was in thermal equilibrium (as in the case of gravitino DM produced in the decays of the next-to-lightest supersymmetric particle).

An immediate consequence of these observations is that it is impossible for the LHC to test most realisations of the freeze-in mechanism, in which no part of the hidden sector thermalises with the SM. More complex freeze-in models may be testable, for example if they contain a new state that couples only very weakly to the hidden sector and decays dominantly into SM particles. However, it will be very challenging for the LHC to establish a connection to the DM problem from such observations alone.

Three fundamental assumptions enter our analysis:

1. We have assumed that the early Universe reached temperatures comparable to the energies accessible for the LHC. Clearly, if the reheating temperature is significantly smaller than LHC energies, the LHC may probe processes that have never been in thermal equilibrium (see ref. [12]).
2. We have assumed that the physics relevant for the interactions of the hidden sector in the early Universe is the same as the physics relevant to the LHC. This assumption could potentially be violated by phase transitions in the early Universe or by thermal effects suppressing certain processes (see [13,14] for an example).
3. We have assumed a standard cosmological history up to energies of a few TeV. A mechanism that would significantly increase the expansion rate (e.g. by introducing a large number of additional relativistic degrees of freedom) would relax the non-thermalisation constraint and hence our bound on the number of observable LHC events [15].

While these assumptions clearly limit the generality of our results, they are essentially unavoidable when trying to connect LHC

physics to early Universe cosmology in a predictive way. At the same time they provide useful guidance for the ways in which early-Universe cosmology must be modified in order to evade our conclusions. In other words, our assumptions can be seen as a list of loopholes that may be exploited to obtain observable LHC signatures from non-thermalised hidden sectors. In any case, we emphasise that no further assumptions have been made. In particular, by considering cross sections with completely arbitrary dependence on the centre-of-mass energy, we avoid the need to specify the way in which the hidden sector interacts with the SM.

With growing interest in models of FIMP DM with non-thermal production mechanisms it becomes more and more urgent to develop new experimental strategies to probe these theories. We have shown that the LHC is not sensitive to non-thermalised hidden sectors, since the interactions are either too weak or occur dominantly at too low energies to be observable. Further work is needed to establish whether low-energy colliders or beam-dump experiments are more promising [16], or whether we need to rely on non-collider experiments, such as low-threshold direct detection experiments [17,18], to make further progress. The model-independent method proposed in this work may provide a useful tool for future studies of these questions.

Acknowledgements

FK would like to thank Stefania Gori, Andreas Goudelis, Jan Heisig, Joachim Kopp, Kai Schmidt-Hoberg, Jessie Shelton, Susanne Westhoff, Sebastian Wild and Bryan Zaldivar for valuable comments on the manuscript and all participants of the workshop “The Future of Searches for Invisible Particles” (14–15 December 2017, Aachen, Germany) for inspiring discussions. This work is supported by the DFG Emmy Noether Grant No. KA 4662/1-1.

Appendix A. Resonant thermalisation

In this appendix we take a closer look at the case where the interactions between the SM and the hidden sector are communicated by an s -channel mediator, i.e. a state that can be resonantly produced from SM initial states both at the LHC and in the early Universe. The mediator will be in equilibrium with the thermal bath if $\Gamma_R > H(T \sim M_R)$, where Γ_R and M_R denote the total width and mass of the resonance, respectively. If, however, the partial width for decays into the hidden sector is sufficiently small $\Gamma(R \rightarrow \text{hidden sector}) < H(T \sim M_R)$, decays of the mediator may populate the hidden sector without bringing it into equilibrium with the thermal bath. In this appendix we show that if $\Gamma(R \rightarrow yy) < H(T \sim M_R)$ for a given hidden sector final state yy it follows automatically that $\gamma(T \sim M_R) < 1$ for any process $xx \rightarrow R \rightarrow yy$.

Let us consider a specific SM initial state called xx . First of all, we make use of the narrow-width approximation to write

$$\sigma_{xx \rightarrow R \rightarrow yy} = \sigma_{xx \rightarrow R} \times \text{BR}(R \rightarrow yy). \quad (\text{A.1})$$

The central observation is now that the cross section $\sigma(xx \rightarrow R)$ for the production process and the partial width $\Gamma(R \rightarrow xx)$ for the

inverse process depend on the same matrix element and therefore differ only by phase space factors. This enables us to write

$$\begin{aligned} \sigma_{xx \rightarrow R \rightarrow yy} &= \pi^2 c \delta(\sqrt{s} - M_R) \frac{\Gamma(R \rightarrow xx)}{N_c M_R^2} \frac{\Gamma(R \rightarrow yy)}{\Gamma_R} \\ &= \pi^2 c \delta(\sqrt{s} - M_R) \frac{\Gamma(R \rightarrow yy)}{N_c M_R^2} \text{BR}(R \rightarrow xx) \end{aligned} \quad (\text{A.2})$$

where N_c counts the colour degrees of freedom of x and c is an order unity numerical pre-factor. For a gg initial state and a scalar resonance, one finds for example $c = 1/2$ [19].

We can now substitute eq. (A.2) into eq. (1) to calculate an upper bound on the reaction rate for the $2 \rightarrow 2$ process. For $T \sim M_R$, and making use of $\text{BR}(R \rightarrow xx) \leq 1$, we find

$$\langle \sigma_{xx \rightarrow R \rightarrow yy} v \rangle n_x^{\text{eq}} \lesssim \frac{c}{4} \Gamma(R \rightarrow yy). \quad (\text{A.3})$$

Hence, if $\Gamma(R \rightarrow yy) < H(T \sim M_R)$, it follows that the same is true for the left-hand side (since $c \sim \mathcal{O}(1)$). This is equivalent to

$$\gamma(T \sim M_R) < 1. \quad (\text{A.4})$$

In other words, for the case of an s -channel resonance the non-thermalisation constraint that we impose is necessary but not sufficient to ensure that the hidden sector does not thermalise with the SM. Imposing instead $\Gamma(R \rightarrow yy) < H(T \sim M_R)$ would lead to even stronger bounds on N_{LHC} than what is shown in Fig. 1.

References

- [1] B.W. Lee, S. Weinberg, *Phys. Rev. Lett.* **39** (1977) 165.
- [2] M. Pospelov, A. Ritz, M.B. Voloshin, *Phys. Lett. B* **662** (2008) 53, arXiv:0711.4866.
- [3] J.A. Evans, S. Gori, J. Shelton, arXiv:1712.03974, 2017.
- [4] M. Duerr, F. Kahlhoefer, K. Schmidt-Hoberg, T. Schwetz, S. Vogl, *J. High Energy Phys.* **09** (2016) 042, arXiv:1606.07609.
- [5] N. Bernal, M. Heikinheimo, T. Tenkanen, K. Tuominen, V. Vaskonen, *Int. J. Mod. Phys. A* **32** (2017) 1730023, arXiv:1706.07442.
- [6] J.L. Feng, A. Rajaraman, F. Takayama, *Phys. Rev. D* **68** (2003) 063504, arXiv:hep-ph/0306024.
- [7] L.J. Hall, K. Jedamzik, J. March-Russell, S.M. West, *J. High Energy Phys.* **03** (2010) 080, arXiv:0911.1120.
- [8] P. Gondolo, G. Gelmini, *Nucl. Phys. B* **360** (1991) 145.
- [9] G. Bélanger, F. Boudjema, A. Goudelis, A. Pukhov, B. Zaldivar, arXiv:1801.03509, 2018.
- [10] A.D. Martin, W.J. Stirling, R.S. Thorne, G. Watt, *Eur. Phys. J. C* **63** (2009) 189, arXiv:0901.0002.
- [11] B. Feldstein, F. Kahlhoefer, *J. Cosmol. Astropart. Phys.* **1408** (2014) 065, arXiv:1403.4606.
- [12] R.T. Co, F. D'Eramo, L.J. Hall, D. Pappadopulo, *J. Cosmol. Astropart. Phys.* **1512** (2015) 024, arXiv:1506.07532.
- [13] M.J. Baker, J. Kopp, *Phys. Rev. Lett.* **119** (2017) 061801, arXiv:1608.07578.
- [14] M.J. Baker, M. Breitbach, J. Kopp, L. Mittnacht, arXiv:1712.03962, 2017.
- [15] F. D'Eramo, N. Fernandez, S. Profumo, arXiv:1712.07453, 2017.
- [16] J. Alexander, et al., arXiv:1608.08632, 2016.
- [17] R. Essig, J. Mardon, T. Volansky, *Phys. Rev. D* **85** (2012) 076007, arXiv:1108.5383.
- [18] M. Battaglieri, et al., arXiv:1707.04591, 2017.
- [19] A. Djouadi, *Phys. Rep.* **457** (2008) 1, arXiv:hep-ph/0503172.