

Simulating glueball production in $N_f = 0$ QCD

David Curtin^{*} and Caleb Gemmell[†]*Department of Physics, University of Toronto, Toronto M5S 1A7, Canada*Christopher B. Verhaaren[‡]*Department of Physics and Astronomy, Brigham Young University, Provo, Utah 84602, USA*

(Received 25 March 2022; accepted 20 September 2022; published 13 October 2022)

In an $SU(N_c)$ gauge theory with zero light quark flavors $N_f = 0$, the only hadronic states that form below the confinement scale are composite gluon states called glueballs. These minimal confining sectors arise in many hidden valley extensions of the Standard Model, including scenarios that could hold the solution to the dark matter question and the hierarchy problem. Quantitative study of dark glueball phenomenology requires an understanding of pure glue hadronization, which to date is severely lacking. In this work we show that significant progress can be made by combining a perturbative pure glue parton shower with a self-consistent and physically motivated parametrization of the unknown nonperturbative physics, thanks to the modest hierarchy between the glueball mass and the confinement scale. We make our simulation code available as the public GlueShower package, the first glueball generator for hidden valley theories, and perform preliminary studies of several glueball production observables, with theoretical uncertainties that take the full range of possible hadronization scenarios into account. We hope this will enable new studies of dark sector phenomenology that were previously inaccessible.

DOI: [10.1103/PhysRevD.106.075015](https://doi.org/10.1103/PhysRevD.106.075015)

I. INTRODUCTION

The ongoing mysteries of the nature of dark matter (DM) and the electroweak hierarchy problem have long been driving forces for extending the Standard Model (SM). The hierarchy problem motivates searches for new particles below the TeV scale but current experiments have not yet found evidence for their existence. To solve these problems, increasingly interesting and complex dark sectors are being considered. A popular framework that aims to address these issues are hidden valley (HV) models [1]. These often include SM-singlet particles charged under a confining $SU(N_c)$ group; see e.g., [2–5]. Couplings to the SM are possible via portal interactions [6–8] which tend to be very weak, allowing the possibility of GeV-scale states that can evade current experimental bounds.

Some realizations of the HV framework, generally referred to as neutral naturalness models, solve the little hierarchy problem by canceling quadratic SM contributions

to the Higgs mass with particles uncharged under SM color. Important examples include mirror twin Higgs [9], fraternal twin Higgs [10], folded supersymmetry [11], and many more [12–17]. Additionally, HV models can produce rich and diverse phenomenology, distinct from any SM processes, such as long-lived particles (LLPs) [18–20], soft unclustered energy patterns (SUEPs) [21–23], and “dark showers” [24], leading to semivisible or emerging jet signatures [25–31].

An important special case of hidden valleys is the pure Yang-Mills QCD case with $N_f = 0$. The only hadronic states in the dark sector are a spectrum of dark glueballs [32–37], which can decay to SM states via dimension 6 or 8 operators [38,39] and have potentially long lifetimes on collider or even cosmological scales. $N_f = 0$ QCD-like sectors appear commonly in neutral naturalness models; for example, in the fraternal twin Higgs [10] only the third generation of SM fermions is mirrored in the dark sector, leaving no strongly interacting states below the confinement scale. Dark glueballs can then be the lightest hadronic states in the twin spectrum. Dark glueballs have also been considered as potential DM candidates [40–50], with their relatively strong self-interaction giving rise to interesting astrophysical signatures [51,52].

Clearly, studying this scenario in detail is highly motivated. However, to date there is no reliably way of simulating dark glueball production in high-energy

*dcurtin@physics.utoronto.ca

†caleb.gemmell@mail.utoronto.ca

‡verhaaren@physics.byu.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

processes. Previous studies have resorted to making very simplistic conservative assumptions, like assuming exotic Higgs decays to just two mirror glueballs in studies of LLP signals in neutral naturalness [53,54]. Another approach is the use of analytical approximations for the final glueball distributions after dark hadronization [55], but this involves some *ad hoc* parameter choices as well as being inconsistent for all but very high initial energies, due to the relatively high mass of glueballs compared to the confinement scale. Clearly, the absence of a reliable event generator for $N_f = 0$ hidden valleys severely hampers their phenomenological and experimental study [24]. Our work addresses this shortcoming, opening the door to a large variety of new and detailed investigations.

The difficulty in simulating $N_f = 0$ QCD arises from the unknown nature of hadronization without light quarks. For $N_f > 0$ with some dark quark masses below the dark confinement scale, the Lund string model [56] can in principle be used to describe hadronization in the dark sector, with the existence of light-colored states allowing tubes of color flux to break via light quark pair production.¹ This is implemented in the hidden valley module [60,61] of the PYTHIA 8 [62] generator. In pure $SU(N_c)$ Yang-Mills theory, on the other hand, no existing hadronization model has been implemented so far.

In this work we present a simulation strategy for obtaining dark glueball final states from pair-produced dark gluons with some initial center-of-mass energy M , implemented as the public PYTHON package GlueShower [63] for $N_f = 0$ $SU(N_c)$ QCD with $N_c = \{2, \dots, 12\}$. Because dark gluon production and decay are highly dependent on the specific beyond-SM (BSM) model, we do not specify those aspects of the dark shower process, instead focusing on the perturbative dark gluon shower and hadronization into dark glueballs. This can be combined with other event generators for production and decay to give a complete signal simulation for a given dark sector.²

Apart from the practical usefulness of assembling a usable event generator for $N_f = 0$ QCD, the main novelty of GlueShower is our parametrization of different possible hadronization mechanisms, given that the underlying non-perturbative dynamics are even less well understood from first principles than hadronization with light quarks. Our hadronization model is simple and physically motivated, with enough built-in variation to span the space of reasonably possible “jetlike” and “plasmalike” final outcomes. Despite faithfully incorporating our large theoretical

¹Alternative hadronization schemes such as the cluster model [57] implemented by HERWIG++ [58] also only apply for the case with light quarks, as does preconfinement [59].

²A follow-up paper applying this work to study the possible indirect detection signals from dark matter annihilating to dark glueballs in our Galaxy is currently in progress [64].

ignorance of pure Yang-Mills hadronization, the resulting predictions are of sufficient precision to make them highly useful for dark sector searches and constraints.

In Sec. II, we briefly review the known properties of glueballs as obtained from lattice QCD. (For simplicity, when referring to gluon and glueballs in this paper we refer to the pure Yang-Mills case, explicitly specifying when we instead refer to SM gluons or states.) The perturbative aspects of QCD relevant to our Monte Carlo generator are reviewed very briefly for completeness in Sec. III. Section IV covers the hadronization process of GlueShower. In Sec. V we simulate glueball production for a variety of hadronization assumptions, define a set of eight benchmark hadronization parameters to span the range of physically reasonable possible outcomes, and make new predictions for observables of glueball production with theoretical uncertainties derived by the variation across these hadronization benchmarks. We conclude in Sec. VI.

II. DARK GLUEBALLS

The properties of $SU(N_c)$ glueballs have been studied on the lattice for decades [32–37], establishing a spectrum of 12 stable states in the absence of external couplings, as shown in Fig. 1. These states are distinguished by their J^{PC} quantum numbers, and their masses can be parametrized entirely in terms of the confinement scale Λ .

In this work we use the lattice values calculated in [36]. Across the range of $N_c = \{2, \dots, 12\}$ values we consider, the lightest glueball mass m_0 is approximately 6Λ . This is much heavier than e.g., SM hadrons and important for our characterization of glueball hadronization. Additionally they also provide the glueball masses for the next nine states and relate the string tension $\sqrt{\sigma}$ to the three-loop confinement scale in the \overline{MS} scheme. Since $\Lambda_{\overline{MS}}$ is quoted to three-loop accuracy we also use the three-loop running coupling as given in [65], which has been used with tree-level splitting functions in other Monte Carlo generators




FIG. 1. Glueball mass m_G spectrum for pure $SU(3)$ Yang-Mills theory [32] in terms of the lightest glueball mass m_0 , plot taken from [38].

[58].³ These values are provided for N_c between 2 and 12, and because the perturbative QCD calculations are easily adjusted for a general N_c , GlueShower generates glueball showers for this same range of $SU(N_c)$ confining sectors. This opens the door for simulating the phenomenology of a wide range of interesting exotic dark sectors and hopefully encourages study of confining dark sectors beyond the SM-like $SU(3)$ case [41,44,66,67].

Glueball wave functions have also been studied on the lattice and by other methods. These have determined the average size of glueballs, with the r_{rms} value for the 0^{++} glueball typically found to be of order $\sim\Lambda^{-1}$ [68–70]. The next heaviest state, 2^{++} , is approximately twice the size of the 0^{++} state [70]. There is still uncertainty in these measurements and questions regarding whether quenched QCD lattice studies will agree with the pure Yang-Mills results. Still, these small radii support focusing on local physics rather than considering glueballs as truly extended objects, which guides our discussion of hadronization below.

III. PERTURBATIVE SHOWER REVIEW

Perturbative QCD is an extensively studied and established field within quantum field theory. Following [71,72] we briefly review the salient details of perturbative parton showers in this section for completeness and to establish notation. More complete documentation of how this is implemented in GlueShower is found in the Appendix.

For $N_f = 0$ HV models, the only parton is the gluon. Our code works in the center-of-mass frame of a two-gluon initial state with invariant mass M .⁴ This amounts to simulating gluon production from the decay of a massive scalar particle of mass M but can be easily generalized to, for example, direct digluon production via effective operators or intermediate states. These initial gluons are produced with large virtualities (effective mass-squared) t and as they split into more gluons produce an increasing ensemble of lower-virtuality gluons. The energies and virtualities of these evolving gluons are described by perturbative QCD and can be simulated as a parton shower.

The probability that a gluon splits into two gluons with energies z and $(1 - z)$, where z is the energy fraction of the mother dark gluon, is determined only by the gluon-to-gluon splitting function:

$$P_{gg}(z) = 2C_A \left[\frac{z}{1-z} + \frac{1-z}{z} + z(1-z) \right], \quad (3.1)$$

where $C_A = N_c$. This splitting function is also used to define the Sudakov form factor, which gives the probability that the gluon evolves from an initial virtuality t_0 to a lower virtuality t without splitting:

$$\Delta(t) = \exp \left[- \int_{t_0}^t \frac{dt'}{t'} \int dz \frac{\alpha_s}{2\pi} P_{gg}(z) \right]. \quad (3.2)$$

Note that this Sudakov form factor only accounts for the leading-log collinear gluon enhancements, which is unable to reproduce the correct scaling of parton multiplicity with center-of-mass energy. However, this is easily rectified in standard implementations of parton showers by imposing angular ordering on subsequent splittings, which accounts for soft gluon interference effects at leading order in $1/N_c^2$ [73]; see the Appendix for further details.

Monte Carlo parton shower evolution can be framed in the following way: Given a gluon with virtuality and energy, (t_1, z_1) , after some step in the evolution, what is its new virtuality and energy, (t_2, z_2) ? Note that a gluon can only decrease its virtuality by splitting. Thus, finding a value for t_2 implies the gluon split in the intermediate step. To calculate t_2 a random number $\mathcal{R} \in [0, 1]$ is generated and t_2 is found by solving

$$\Delta(t_2) = \frac{\Delta(t_1)}{\mathcal{R}}. \quad (3.3)$$

If there is no solution, then the gluon does not split. In this case the gluon can only hadronize, and its shower terminates. If a t_2 is found, z_2 is determined by generating another random number $\mathcal{R}' \in [0, 1]$ and solving

$$\int_{z_{\min}}^z dz' \frac{\alpha_s}{2\pi} P_{gg}(z') = \mathcal{R}' \int_{z_{\min}}^{1-z_{\min}} dz' \frac{\alpha_s}{2\pi} P_{gg}(z'), \quad (3.4)$$

where z_{\min} is set by kinematic thresholds of the possible gluon splittings. This Monte Carlo generation is implemented in GlueShower as detailed in the Appendix. We also note that since $\alpha_s \propto N_c^{-1}$ at one-loop order, and the perturbative shower is only dependent on $\alpha_s P_{gg}(z)$, changes to N_c only affect the shower due to the slight dependence of glueball masses on N_c . In our numerical studies below we therefore only show the $N_c = 3$ case, with other numbers of colors giving similar behavior.

Lastly, as is standard in parton shower MC algorithms, we work in the leading color limit, which is equivalent to taking $N_c \rightarrow \infty$ with $\alpha_s N_c$ kept constant [74]. This amounts to ignoring higher-order color interference effects in $1/N_c^2$ expansions. Additionally in this limit we can use the simple representation of gluons being the direct product of a fundamental and antifundamental, since the weight of

³Using the three-loop running coupling is simply to be consistent with the lattice results and not indicative of a higher level of accuracy, as any improvement compared to the one-loop result is subdominant to the uncertainties of our hadronization model.

⁴This is the only initial state we consider, as it is of the most use to BSM physics studies. We leave the study of other interesting cases, such as a gas of gluons that is cooled until it undergoes confinement, for future investigations.

the singlet in $N_c \otimes \bar{N}_c = (N_c^2 - 1) \oplus 1$ vanishes. It is in this limit that the t' Hooft double line notation can be used to trace color flow, as shown in Fig. 2. From SM QCD we know this approximation works well for $N_c = 3$, and thus also for $N_c > 3$. It is possible that this approximation introduces larger errors for $N_c = 2$, but we leave the problem of including the subleading color corrections for future investigations.

IV. HADRORIZATION

While the perturbative aspects of the parton shower have been extensively studied, the nonperturbative aspects of QCD have not been computed from first principles. We follow the usual approach, outlined in the previous section, whereby we iteratively evolve gluons from their initial virtuality and undergo splittings above some hadronization scale Λ_{had} , which is typically taken to be of the same order as the confinement scale Λ (though it turns out this assumption must be modified for our case). Some hadronization model capturing the nonperturbative dynamics of glueball formation is then needed to turn the gluon final states of the shower into physical final states. In SM QCD, phenomenological hadronization models [56,57] can be tuned to reproduce the observed data, but these rely on the existence of light quarks, which either allow a flux-tube string to fragment or allow color-singlet clusters to form during preconfinement [59]. For pure Yang-Mills, the corresponding hadronization models have not been formulated and at any rate cannot be tuned to data prior to the discovery of a corresponding dark sector.⁵ We therefore provide a novel parametrization of the various possibilities for $N_f = 0$ QCD hadronization, relying on simple physically motivated arguments to capture the range of both jetlike or plasmalike possibilities for how pure glue might hadronize, producing a consistent spectrum of simulated glueball final states obeying full energy-momentum conservation. This accurately captures present theoretical uncertainties while allowing for the quantitative study of pure glue hidden valleys with sufficient precision for searches and constraints.

A. An intuitive cartoon

We begin with a representative cartoon to guide our intuition for pure glue hadronization. In this cartoon we assume there is only one glueball species and consider the simple example of producing two gluons that eventually hadronize into two glueballs. Because the glueballs are color-singlet final states, and we begin with two back-to-back color-octet gluons, long-range color exchange must occur at some point along the shower or hadronization *if the*

⁵A highly exciting possibility is understanding pure glue hadronization from first principles on the lattice. This is very challenging, but there has been significant recent progress for $N_f > 0$; see e.g., [75].




FIG. 2. Comparison between Feynman-diagram-like (left) and flux-string-like (right) cartoons of the evolution of color-singlet separation in $N_f = 0$ QCD.

*center-of-mass energy is high enough.*⁶ Thus, before the glueballs can form, at minimum one of the original gluons must split into three gluons, with one new gluon joining the other branch and allowing color-singlet states to form, referred to as color-octet neutralization [76].⁷ This illustrative cartoon is depicted in a Feynman-diagram-like way on the left side of Fig. 2, showing gluon splittings and flow of color charge.

Alternatively, one can consider how tubes of color flux in the fundamental representation evolve throughout the shower. This equivalent representation is depicted on the right side of Fig. 2. The initial state is an overall color-singlet loop, with the gluons behaving as localized energy or “kinks” in the loop. Forming two final-state glueballs can only be accomplished by the initial loop fragmenting in two. As the loop evolves, gluon splitting introduces new kinks and new colors for the flux-tube segments in the loop. Toward the end of the shower, step 3, enough gluons have formed to allow two flux-tube segments in the loop to have the same color, making fragmentation into two color-singlet glueballs possible by crossing the same-color flux-tube segments. *Thus, we see that glueball hadronization can be pictured as crossing color-fundamental flux tubes “pinching off” one flux-tube loop into two separate loops, which becomes possible even after a small number of gluon splittings.*

The tube crossing picture is supported by the fact that loop fragmentation is hugely energetically preferred once it is possible, since flux-tube crossing has little energy cost compared to the energy stored in flux tubes of length $\gg \Lambda^{-1}$. Making the crossing possible only requires the exchange of arbitrarily soft IR gluons carrying color information, which carries no particularly relevant

⁶The minimum initial energy required for separation into multiple color singlets will be made more precise in the next subsection.

⁷Other detailed splitting histories are of course possible but do not change our qualitative argument.

suppression and is in fact enhanced by α_s running to large values at low energies. This is also consistent with glueball radii being of order the confinement scale, suggesting their formation is dominated by local physics.

We can compare the glueball flux tubes with the Lund string model used for QCD. For the simple case of quark pair production, the shower and subsequent hadronization process essentially follows the evolution of an open flux tube with the quarks at either end. These flux tubes are easily broken through light quark pair production, which increases the open flux-tube multiplicity. The final-state flux-tube states with quarks and antiquarks on either end are then associated with a variety of mesons. $N_f > 0$ QCD is predominantly determined by the fragmentation of open strings, while the final state for $N_f = 0$ QCD HV theories is determined by closed loops pinching and fragmenting.

In summary, formation of individual glueballs requires crossing flux tubes, which seems to be entirely unsuppressed once a small number of (possibly very soft) gluons have been exchanged between different branches of the shower. This required gluon exchange in turn becomes unsuppressed once the virtualities in the shower approach the confinement scale. *This motivates the intuition that glueball formation should occur around the confinement scale, qualitatively similar to SM jetlike behavior.*

B. GlueShower hadronization implementation

We now describe our pure glue hadronization algorithm in detail, first for the default jetlike assumption as motivated above and then for more exotic plasmalike behavior, to cover all physically reasonable possibilities.

1. Jetlike glueball hadronization

The assumption that glueball production proceeds in a qualitatively similarly jetlike fashion as for SM QCD is well motivated by physical arguments, as explained in the previous section. However, the detailed question still remains how to map some arrangement of gluons at the end of the perturbative shower to a set of final-state glueballs. One could implement a full Lund-like string model of closed flux tubes and their crossing and vary its parameters to obtain a range of possibilities for glueball hadronization. This would be quite involved, and we leave this for future investigations. However, thanks to the significant separation between the confinement scale Λ and the lightest glueball mass m_0 , we can already make great progress with much simpler physical arguments.

We begin by assuming that there is just a single glueball species of mass m_0 , and consider a gluon in the perturbative shower that has virtuality $\sqrt{\hat{t}} > 2m_0$ and does not undergo further splittings above scale $2m_0$. If the gluon were allowed to evolve further down in virtuality, any subsequent splittings cannot result in more than one on-shell glueball unless there is significant momentum exchange with other branches of the shower. Instead, subsequent splittings result

in lower-energy gluons physically clustered around the original gluon 4-momentum in a decay cone that is at most as wide as one arising from the decay of a particle with mass below $2m_0$. As $\sqrt{\hat{t}} \rightarrow \Lambda$, hadronization must combine these soft gluons (plus some soft IR gluons to exchange color information with other branches of the shower) into an on-shell, color-singlet glueball with mass m_0 .

Therefore, we argue that for a given shower history, an *upper bound* on the number of glueball states produced can be obtained by setting the hadronization scale to $\Lambda_{\text{had}} = 2m_0$ and simply turning the gluons at the end of the shower into on-shell glueballs of mass m_0 . Note that $2m_0 \sim 12\Lambda$ is well within the perturbative regime, so the simulated momenta of gluons at the end of the shower are highly reliable. However, in converting these gluons to glueballs some soft gluons must be exchanged with other branches of the shower to form color singlets. This suggests momentum transfers of order $\Lambda/2m_0$, but the relatively high mass of glueballs makes this correction factor smaller than $\approx 10\%$. Therefore, naively turning gluons with $\sqrt{\hat{t}} = \Lambda_{\text{had}}$ into on-shell glueballs is likely to be a good approximation.⁸

How could this upper bound on glueball multiplicity be violated? First, it may be possible for two neighboring branches of the shower to exchange gluons with momenta between $2m_0$ and Λ , allowing two gluons with virtuality $2m_0$ each to result in three instead of two final-state glueballs. This can still be regarded as a (marginally) perturbative process in the shower's regime of validity. On the other hand, this assumption can be violated entirely in the nonperturbative regime if the flux strings between two branches contain enough energy to produce a third glueball.

We first discuss the perturbative possibility, where the size of the momentum exchange gives us some hope of using the perturbative shower and simple phase space arguments to estimate this rate $P_{2 \rightarrow 3}$ of turning two gluons at the end of the shower into three glueballs. We overestimate $P_{2 \rightarrow 3}$ to show that it is small enough to ignore at our current level of precision.

Consider two gluons 1 and 2 at the end of the shower, both with virtuality $2m_0$, energies $E_{1,2}$ and angle θ between their momenta. To overestimate $P_{2 \rightarrow 3} = P_{2 \rightarrow 3}(E_1, E_2, \theta)$, we assume that a third glueball is formed if the gluons split $1 \rightarrow 1'3$ and $2 \rightarrow 2'4$ such that $m_{1',2'} > m_0$, the daughters have sufficient invariant mass to form the glueball $m_{34} \geq m_0$ and the two daughter momenta are close in phase space compared to the confinement scale, i.e., $\Delta p_{34} \equiv |\vec{p}_3 - \vec{p}_4| < a\Lambda$ for $a \sim \mathcal{O}(1)$, to enable the merger. We choose

⁸This method brushes over the specifics of how color information is exchanged to create color singlets, reminiscent of early independent fragmentation models [77]. However, these early models, proved to be historically useful in SM QCD, are even more useful for pioneering BSM studies in our case due the larger mass of the glueball hadrons in $N_f = 0$ compared to the typical momentum transfer involved in color exchange.

$a = 2$ but the precise value does not significantly affect our result. We can therefore estimate

$$P_{2 \rightarrow 3}(E_1, E_2, \theta) = \int dt_1 dt_2 dz_1 dz_2 \times P_{1 \rightarrow 1'3, 2 \rightarrow 2'4, m_{1', 2'} > m_0}(E_1, E_2, \theta) \times P_{m_{34} > m_0, \Delta p_{34} < a\Lambda}(|p_3|, |p_4|, \theta). \quad (4.1)$$

This integrates over all possible splittings (t_1, z_1) , (t_2, z_2) of the two parent gluons that produce daughters 3 and 4. The first term (splitting probability) just evaluates the Sudakov to give the splitting probability of both parent glueballs producing two daughters with momentum p_3 , p_4 such that the parents still have sufficient virtuality to form their own glueballs. The second (merger probability) term is the probability, given a random emission angle for each daughter in the transverse plane of the parent momentum, that the two daughters could in principle combine to form a third glueball according to our above criteria.

To simplify evaluation of this integral, we overestimate both terms separately. The splitting probability (first term) is overestimated by letting both gluons 1 and 2 run from $\sqrt{t} = 2m_0$ down to $\sqrt{t} = 2\Lambda$, allowing for emission of daughters with virtualities as low as Λ . This is a huge overestimate, since we are not enforcing $m_{1', 2'} > m_0$, and turns $P_{1 \rightarrow 1'3, 2 \rightarrow 2'4, m_{1', 2'} > m_0}(E_1, E_2, \theta; p_3, p_4)$ into a constant $P_{\text{split}}^{\text{max}}$ that sits outside the integral. We then set the virtualities of the daughters 3,4 to be the largest possible masses allowed by a given splitting to define their 4-momenta p_3 , p_4 , which in turn maximizes the merger probability (second term). This allows us to define

$$P_{2 \rightarrow 3}^{\text{max}}(E_1, E_2, \theta) \equiv P_{\text{split}}^{\text{max}}(E_1, E_2) \int dt_1 dt_2 dz_1 dz_2 P_{\text{merge}}(|p_3|, |p_4|, \theta), \quad (4.2)$$

satisfying $P_{2 \rightarrow 3}^{\text{max}} > P_{2 \rightarrow 3}$. Having thus obtained a function $P_{2 \rightarrow 3}^{\text{max}}(E_1, E_2, \theta)$, it is then straightforward to generate events using the perturbative pure-gluon shower for a variety of initial center-of-mass energies M , terminate at hadronization scale $\Lambda_{\text{had}} = 2m_0$, and evaluate the chance of obtaining an extra glueball compared to our naive expectation of turning gluons with virtuality $2m_0$ into glueballs:

$$P_{\text{extra glueball}} = 1 - \prod_{i,j} (1 - P_{2 \rightarrow 3}^{\text{max}}(E_i, E_j, \theta_{ij})), \quad (4.3)$$

where the product is over all gluon pairs (i, j) in the event. Averaged over all events, this probability is shown in Fig. 3 for $N_c = 3$, but the result is nearly identical for other numbers of colors. As expected, the chance of producing additional glueballs beyond our upper bound increases with center-of-mass energy since this gives more gluons and more chances for the required splittings and mergers.




FIG. 3. Upper bound on the probability of producing more glueballs of mass m_0 in gluon pair production events than the naive upper multiplicity bound obtained by terminating the pure glue shower at $\Lambda_{\text{had}} = 2m_0$, as a function of the center-of-mass energy M .

However, for $M \lesssim 100m_0$, the error introduced by ignoring $2 \rightarrow 3$ production of glueballs is smaller than 5%, and we are justified in ignoring it for our current implementation.

We now consider nonperturbative effects that could produce additional glueballs, orienting ourselves in the SM analogues of Schwinger pair production in QED [78], or pion production in the “snapping” of color-fundamental strings [79]. The latter case is of most physical interest. The energy density of the color string is given by the string tension σ , and the production rate of hadrons in string fragmentation scales as

$$P_{\text{hadron}} \propto \exp\left[-\frac{\pi m_{\text{hadron}}^2}{\sigma}\right], \quad (4.4)$$

since the constituent masses of the partons have to be provided by the color background field. Ignoring any additional suppressions that may arise from the required geometrical arrangement of two color flux tubes to allow for the creation of a glueball, the fact that the glueball mass is much larger than the mass of SM pions compared to the string tension should result in a large suppression on the number of glueballs produced from the dynamics of the nonperturbative color strings. For $N_f = 0$, $N_c = 3$, the ratio of the exponential factors is

$$\frac{P_{\text{GB}}}{P_{\pi}} \sim 10^{-16}, \quad (4.5)$$

where the numerator was evaluated using $\Lambda_{\overline{MS}}/\sqrt{\sigma} = 0.5424$ and the lightest glueball mass, $m_{\text{GB}} = 6.28\Lambda_{\overline{MS}}$ [36]. Note that for different $SU(N)$ groups, the constants change slightly but the extreme suppression persists. We therefore conclude that both perturbative and nonperturbative effects should not invalidate our jetlike approach.

If setting $\Lambda_{\text{had}} = 2m_0$ and simply turning final-state gluons into on-shell glueballs gives the largest possible glueball multiplicity per event, how can we take into account the possibility that the real number of produced glueballs could be lower? Physically, this would correspond to color singlets forming via gluon exchange at a higher virtuality scale than naively expected, which is certainly a possibility given the unknown details of non-perturbative $N_f = 0$ QCD. We can obtain consistent events representing this scenario by simply terminating the shower at a higher scale before turning final-state gluons into on-shell glueballs. In other words, we set

$$\Lambda_{\text{had}} = c \times 2m_0, \quad (4.6)$$

where $c > 1$ is a dimensionless parameter encoding our assumption of the higher scale Λ_{had} where the shower fragments into disconnected color singlets that each yield one glueball. Varying $c \sim \mathcal{O}(1)$ gives us a controlled way to parametrize different possible assumptions on jetlike glueball hadronization and explore the effect of this uncertainty on observables while still producing fully consistent shower histories that conserve energy and momentum for arbitrary choices of $c \geq 1$ and $M/m_0 > 2$.

So far, we have only considered a single glueball species, but in reality there is a spectrum of roughly a dozen different glueball states with different J^{PC} quantum numbers. While our arguments support the notion that the *inclusive* glueball multiplicity and *momentum distributions* are dominated by local jetlike physics, the same is not true for *exclusive* distributions for each glueball species. Soft gluon exchange at the scale $\sim \Lambda$ (also the scale of mass differences between glueball species) can easily exchange angular momentum and other quantum numbers to turn a given candidate gluon into a variety of different glueball species, and rigorously analyzing these nonlocal effects is far beyond our scope. We therefore adopt a much simpler approach of assuming that the relative multiplicities of different glueballs follows a thermal distribution [8] in the absence of other threshold effects, with the probability for producing glueball state J given by

$$P_J \propto (2J + 1) \left(\frac{m_J}{m_0} \right)^{3/2} e^{-(m_J - m_0)/T_{\text{had}}}, \quad (4.7)$$

where the glueball masses for different N_c are known from the lattice, and we define a hadronization temperature

$$T_{\text{had}} = dT_c, \quad (4.8)$$

which is related by a dimensionless coefficient $d \sim \mathcal{O}(1)$ to the critical temperature of the $N_f = 0$ QCD phase transition T_c . This is justified by investigations of final-state distribution from closed string emission [80] that support a thermal model, where T_c is taken to be the Hagedorn

temperature. Consequently, we assume the Hagedorn temperature is the critical temperature of deconfinement [81,82]. In SM QCD, the critical temperature is 150 MeV [83], smaller than the confinement scale. In $N_f = 0$ Yang-Mills theories the relation between the critical temperature and string tension has been studied on the lattice [84–87], with [87] finding the relation

$$\frac{T_c}{\sqrt{\sigma}} = 0.5949 + \frac{0.458}{N_c^2}, \quad (4.9)$$

giving a critical temperature that is slightly larger than the confinement scale. By combining this result with the three-loop relation between the confinement scale Λ and the string tension, the relative glueball multiplicities are entirely determined by the number of colors N_c and the nuisance parameter d .

We incorporate these probabilities into our hadronization routine in the following manner. Once it is determined that a gluon can no longer split during the perturbative shower, it remains with some virtuality $\sqrt{t} \geq \Lambda_{\text{had}} = 2cm_0$. This gluon then selects a random glueball final state with on-shell mass below its current virtuality, weighted by the probabilities in Eq. (4.7). This treats the thermal probabilities P_J as fundamental and introduces some additional threshold effects that favor light glueball production, since a high-virtuality gluon that does not split before termination of the shower can have more glueball final states kinematically accessible to it than a gluon that was produced during the shower with virtuality close to $2cm_0$. The relative multiplicity of glueball species produced by this hadronization routine will therefore skew toward lighter flavors than the thermal $\{P_J\}$ alone, but we believe this is a physically reasonable prediction of the perturbative shower based on little more than phase space arguments.

Additionally, color rope hadronization [88] in SM QCD can affect the final-state relative multiplicities by enhancing strangeness production [89] via an increased string tension. This could be relevant for the $N_f = 0$ QCD case since the two fundamental flux strings between gluons would behave collectively when this effect is included, changing the relative multiplicity of glueballs. Including these dynamics is beyond our scope, but we account for their possible effect by not treating T_{had} as firmly determined, and vary d to parametrize the theory uncertainty in relative glueball species multiplicity. Higher hadronization temperatures favor production of heavier glueball states. Depending on the full details of the theory and the operator by which glueballs decay to SM states [38,39], this can significantly affect the visible phenomenology, and it is important to treat $d \sim \mathcal{O}(1)$ as a nuisance parameter in quantitative analyses. We demonstrate this by studying some benchmarks below.

While we expect our approach of randomly assigning glueball species identity based on thermal probabilities to be reasonable on average over many simulated events, this

simple approximation is unlikely to give accurate intraevent correlations between separations in momentum space and relative species assignments of different glueballs, for example. Even so, this simple parametrization should be sufficient for many first phenomenological investigations.

2. Plasmalike glueball hadronization

While the jetlike showering and hadronization behavior is highly physically motivated, we want `GlueShower` to cover the largest range of physically possible showering behavior. We therefore also consider a much more exotic plasmalike or SUEP-like regime.

Within the jetlike assumption, the fragmentation of closed flux tubes is viewed as an IR process that results in the immediate formation of on-shell glueballs. However, if fragmentation into color singlets occurs at a significantly higher scale, then a population of high-mass closed flux tubes could be produced.⁹ These states could be treated as a collection very excited glueballs or a hot ball of gluon plasma. Much like a quark-gluon plasma that evaporates via quasi-isotropic emission of pions, this pure glue plasma would evaporate by emitting glueballs approximately isotropically with thermal momenta in its rest frame.

To implement this possibility within `GlueShower`, we introduce boolean parameter `plasma_mode` which is false by default (jet behavior) but can be set to true to enable plasma behavior. In plasma mode, the shower still terminates at a scale determined by setting $c > 1$ just as in jet mode:

$$\Lambda_{\text{had}} = c \times 2m_0, \quad (4.10)$$

but this scale is now interpreted as the scale below which the shower separates into singlet plasma balls, and each final-state gluon is turned into a singlet of mass $m_{\text{plasma}} = cm_0$ instead of a glueball J of mass m_J .

The evaporation of these plasma balls into glueballs is treated analogously to dark hadron production in SUEP scenarios [21,22,90], assuming isotropic thermal glueball emission. We assume that the thermal glueball energy distribution and their relative species probabilities are dictated by the same temperature $T_{\text{had}} = d \cdot T_c$; see Eq. (4.7). We borrow the SUEP-simulation methods used in [22,90] to generate isotropic glueball momenta in the rest frame of each plasma. In short, isotropic glueball momenta are successively generated with a thermal energy distribution until additional gluon emission would cause the total invariant mass of all gluon momenta to exceed the plasma mass m_{plasma} . For each glueball, its species is randomly picked weighted by the thermal probabilities P_J in Eq. (4.7). Once glueball emission is completed for a given plasma ball, the entire system of daughter glueballs is slightly

boosted and its kinetic energies rescaled to exactly equalize their rest frame and invariant mass with the original plasma ball.

In plasma mode, glueballs are therefore produced in a manner that lies between the jetlike behavior of QCD and the purely isotropic behavior of pseudoconformal theories that realize the SUEP scenario: The shower generates branches of total invariant mass $m_{\text{plasma}} = \Lambda_{\text{had}}/2$, which become plasma balls that evaporate via isotropic thermal glueball emission in their respective rest frames.

We emphasize that this possibility is highly exotic: We regard jetlike hadronization as far more physically motivated, since realizing plasmalike behavior requires very unusual long-distance nonperturbative effects that arrest further fragmentation well within the naive regime of reliability for the perturbative shower. However, we include it in our code to make sure that even highly exotic hypotheses for the nonperturbative dynamics of $N_f = 0$ QCD can be qualitatively accommodated.

3. Input parameters for `GlueShower`

In summary, `GlueShower` has two purely physical parameters: $N_c \in (2, 3, \dots, 12)$ to specify the dark gauge group and m_0 to specify the mass of the lightest glueball.¹⁰ This determines the confinement scale Λ , critical temperature T_c , and the three-loop running coupling $\alpha_S(\mu)$. For a given run, one also specifies the initial center-of-mass energy M for the digluon initial state (corresponding to digluon pair production in the decay of a scalar particle of mass M) as well as the number of shower histories to simulate. Theoretical uncertainties of dark hadronization are captured in two-and-a-half nuisance parameters: the boolean parameter `plasma_mode`, which specifies whether gluons hadronize in the physically motivated jetlike or the more exotic plasmalike fashion, and multiplicative parameters $c = \Lambda_{\text{had}}/(2m_0)$ and $d = T_{\text{had}}/T_c$, which set the hadronization or singlet formation scale and the hadronization (and plasma ball, if in plasma mode) temperature, respectively.

V. SIMULATION OF DARK GLUEBALL FINAL STATES

In this section we study events generated by `GlueShower`. We first explore the basic jetlike shower case, using parameters `plasmamode = false`, $\Lambda_{\text{had}}/(2m_0) = c = 1$, and $T_{\text{had}}/T_c = d = 1$ to demonstrate the code and its output, as well as comparing the multiplicity scaling and shape of fragmentation functions to analytical approximations in their expected regime of validity. We then comment on the qualitative differences in the plasmalike case, using parameters `plasmamode = true`, $\Lambda_{\text{had}}/(2m_0) = 4$,

⁹We thank Matthew Strassler for bringing this possibility to our attention.

¹⁰The large N_c regime can be well covered by simulating $N_c = 12$.




FIG. 4. Average glueball multiplicity $\langle N \rangle$ for $N_c = 3$ as a function of center-of-mass energy M relative to the lightest glueball mass m_0 . The jetlike case with the highest possible glueball multiplicity is shown in blue, a representative plasmalike case in orange (see legend). The bands show the 16%–84% percentile range of multiplicities in the event distributions for each M/m_0 (jagged since multiplicities are small integers). The analytical expectation of Eq. (5.1), normalized to the $M/m_0 = 100$ jetlike simulation, is shown as the dashed red line.

and $T_{\text{had}}/T_c = 1$. Finally, we argue that a set of eight benchmark values for the nuisance parameters covers the physically motivated range of hadronization possibilities for glueball production (plasmamode = false{true}, $c = 1, 2 \{4, 6\}$, and $d = 1, 2$) and study the resulting range of physical predictions for some observables.

A. Jetlike hadronization

We first demonstrate how the inclusive glueball multiplicity scales with initial center-of-mass energy M in Fig. 4 (blue). This case of $c = \Lambda_{\text{had}}/(2m_0) = 1$ represents the largest multiplicity possible in the jetlike case, but, even for large $M \sim 100m_0$, the sizable hierarchy between m_0 and Λ_{had} results in only a handful of produced glueballs per event. This differs greatly from the high-multiplicity production of pions in high-energy QCD jets. The blue band shows the range of multiplicities produced in simulated events, which is much more sharply peaked than a Poisson distribution. We show the $N_c = 3$ case but the results are similar for other numbers of colors.

As a consistency check, we compare the average multiplicity predicted by GlueShower to the analytical expectation for average hadron multiplicity scaling in perturbative QCD. The standard result, for example found in [71] for $N_f = 0$, is

$$\langle N(E_{CM}^2) \rangle \propto \exp \left[\frac{12\pi}{11C_A} \sqrt{\frac{2C_A}{\pi\alpha(E_{CM}^2)}} + \frac{1}{4} \ln(\alpha(E_{CM}^2)) \right]. \quad (5.1)$$

This is normalized to the absolute multiplicity of our simulation for $M = 100m_0$ and shown as the red dashed line in Fig. 4. Note the good agreement for $M \gg m_0$,¹¹ but as expected, the scaling relation breaks down for M closer to the glueball mass, demonstrating that finite-mass effects make analytical approximations of glueball distributions unreliable even for modest initial energies.

We now compare the simulated glueball energy spectra to analytical expectations. In the SM, fragmentation functions are *a priori* unknown; thus, a physically motivated functional form is used as an ansatz and fit to data. We compare our jetlike output to one of the commonly used Colangelo and Nason function [91]

$$D_g^{GB}(x) \propto x^\alpha (1-x)^\beta, \quad (5.2)$$

where $x = 2E_{GB}/M$ and E_{GB} is the glueball energy; thus, $x \in [0, 1]$. This function has been used for heavy quark fragmentation [92] and is equivalent to the gluon-to-kaon fragmentation function parametrization in [93], which has also been studied using lattice QCD [94]. Using a heavy quark fragmentation function for a pure glue shower might seem surprising, but this simply encodes that the energies of heavy hadronic final states are dominated by the heavy quark energies. This process is determined by the perturbative shower since it occurs significantly above the confinement scale, much like our treatment of jetlike glueball hadronization, making such a functional form a reasonable ansatz. A similar approach was also used to analytically estimate dark glueball momenta in [55].

We compare our simulated events to an analytical fragmentation function in the high-energy regime $M/m_0 = 100$ and for $0.1 < x < 1$ [93], where finite-mass effects are less significant. As shown in Fig. 5 (left), we find very good agreement in this regime of applicability, provided we can find the required values of α and β from data (i.e., simulation). In the same figure we show fragmentation functions for lower M , which demonstrates that the produced glueball spectrum becomes dominated by finite-mass effects as $M \rightarrow 10m_0$, since small x values become kinematically forbidden and a significant fraction of events only produce two glueballs with $x = 1$. This again demonstrates the limitations of using analytical approximations to estimate glueball distributions and underlines the importance of using a self-consistent Monte Carlo simulation like GlueShower.

We also show the exclusive fragmentation functions of the lightest ten glueball states for $M/m_0 = 100$ in Fig. 5 (right). For $T_{\text{had}} = T_c$, the final states predominantly consist of 0^{++} and 2^{++} glueballs, followed by the next three heaviest states in roughly equal proportion. Heavier state production is suppressed by an order of magnitude.

¹¹Note that some deviation is expected since our shower uses three-loop running of the coupling.

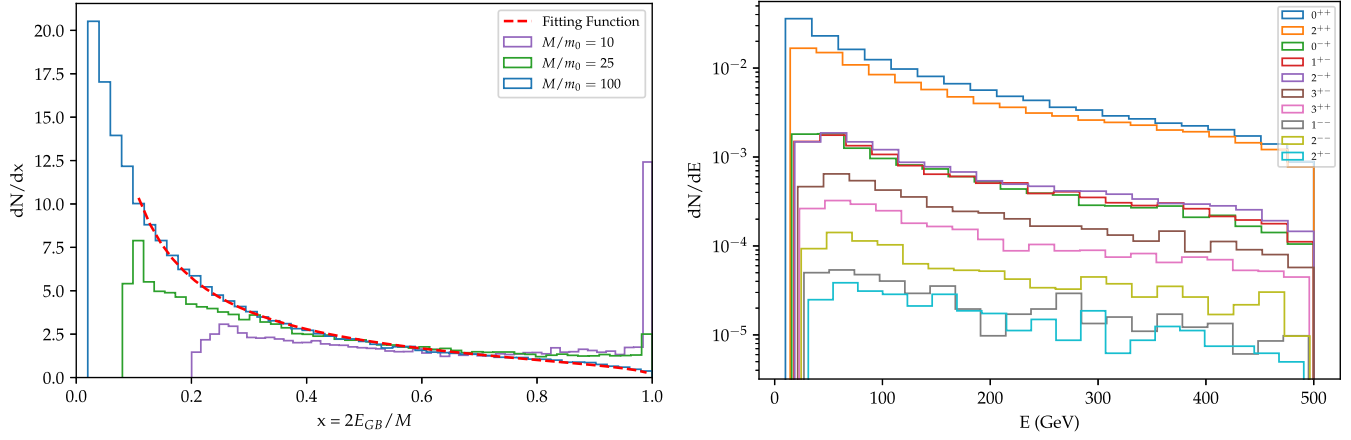


FIG. 5. Glueball energy spectra for an example jetlike shower, with plasmamode = false, $\Lambda_{\text{had}}/(2m_0) = 1$, and $T_{\text{had}}/T_c = 1$. Left: representative 0^{++} spectra for a range of M values. The $M/m_0 = 100$ spectrum is fit to Eq. (5.2) (red line) on the range $0.1 < x < 1$. Right: spectra for the lightest ten glueball states for initial center-of-mass energy $M/m_0 = 100$.

B. Plasmalike hadronization

We now discuss some of the same GlueShower outputs for a basic plasmalike shower case. The orange curve in Fig. 4 shows how multiplicity scales with initial center-of-mass energy for $\Lambda_{\text{had}}/(2m_0) = c = 4$. Overall we find that the plasmalike case leads to higher multiplicity events compared to the jetlike case but also that the inclusive multiplicity scales differently in various energy regimes.

- (i) In the large energy limit $M \gg \Lambda_{\text{had}}$, we recover the same multiplicity scaling as the jetlike case, but larger by a constant since each produced plasma ball evaporates into at least two glueballs. This does not mean, however, that the plasmalike case converges to the jetlike case in the high-energy limit: While the overall multiplicity scales similarly with energy, the absolute multiplicity is higher in the plasmalike case, resulting in lower characteristic energies for the final-state glueballs.¹²
- (ii) For $M \rightarrow \Lambda_{\text{had}}$ the multiplicity asymptotes to a value larger than 2, in this case roughly 5. This is the regime in which the initial state only splits into two singlet plasma balls, which then each evaporate to at least two glueballs each.
- (iii) At $M = \Lambda_{\text{had}}$ there is a discontinuity as the shower enters a new regime in which there is insufficient center-of-mass energy to form two separate plasma balls. For $M \leq \Lambda_{\text{had}}$, the initial digluon production

¹²It is helpful to consider a SM QCD analogy: One could imagine a hypothetical hadronization model that dominantly produces heavy B mesons (analogous to the plasma balls), which then decay to lighter hadrons. Certainly, various multiplicity and momentum scalings of this B -meson shower would be similar to the scalings of pion multiplicity and momenta in realistic SM hadronization, but the different hadronization assumption would also introduce fundamental differences in the final states no matter what the initial energy.

event is therefore taken to form just a single color-singlet plasma ball of mass M , which evaporates by glueball emission. Since this leads to a larger m_{plasma} just below Λ_{had} than just above, there is a small spike in produced glueball multiplicity, but this is a reasonable threshold effect. In this regime, glueball production is entirely SUEP-like, leading to a very different multiplicity scaling until the absolute minimum of $N = 2$ is reached for $M < 3m_0$.

Inclusive and exclusive glueball fragmentation functions in the plasmalike case are shown in Fig. 6. As demonstrated in the left plot, the analytical functional form of the inclusive fragmentation function is still a fair approximation in the high-energy large- x regime, though significantly worse than for the jetlike case, and for modest or low energies, the differences are even more pronounced. On the right we show exclusive energy spectra for the lightest ten glueball species, which are produced with very similar relative multiplicities as in the jetlike case. This is what we would expect, as relative glueball multiplicities are determined by the hadronization temperature which is kept constant between the cases considered in both Figs. 5 and 6. Additionally this shows that threshold effects in the jetlike case that favor the production of light glueballs have minimal effect.

C. Defining hadronization benchmarks

Having discussed how some simple observables behave at different energies in both the jetlike and plasmalike cases, we now systematically examine their dependence on two nuisance parameters of our simulation: the hadronization scale $\Lambda_{\text{had}} = c \cdot 2m_0$ and the hadronization temperature $T_{\text{had}} = d \cdot T_c$. This allows us to argue for a small set of benchmark hadronization parameters that should span the range of physically reasonable possible outcomes for glueball production and which, hence, define theory uncertainties for physical predictions.

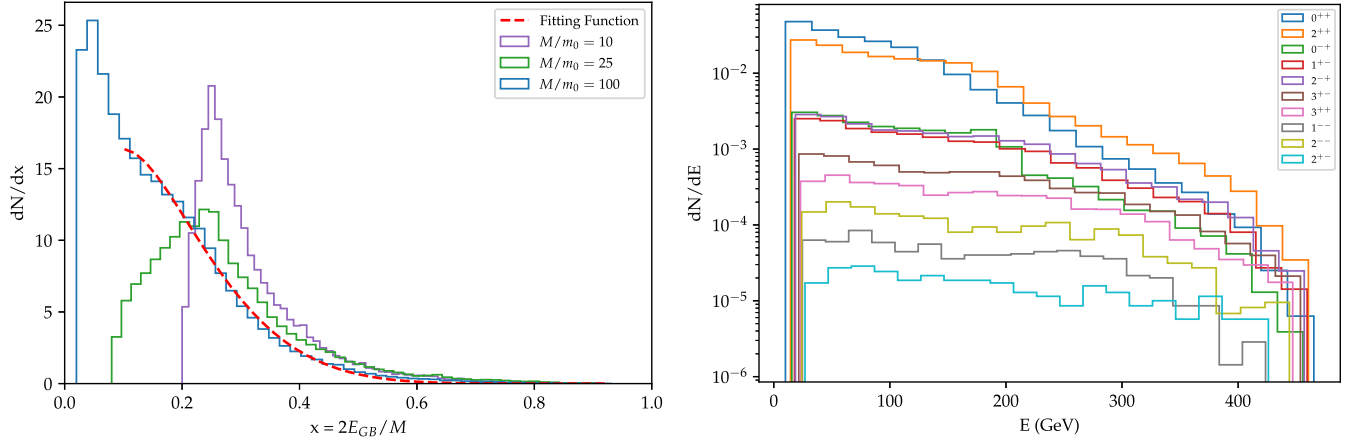


FIG. 6. Glueball energy spectra for an example plasmalike shower, with $\text{plasmamode} = \text{true}$, $\Lambda_{\text{had}}/(2m_0) = 4$, and $T_{\text{had}}/T_c = 1$. Left: representative 0^{++} spectra for a range of M values. The $M/m_0 = 100$ spectrum is fit to Eq. (5.2) (red line) on the range $0.1 < x < 1$. Right: spectra for the lightest ten glueball states for initial center-of-mass energy $M/m_0 = 100$.

Figure 7 shows how the inclusive glueball multiplicity in both the jet- and plasmalike cases depends on Λ_{had} for $T_{\text{had}} = T_c$. In the jetlike case, shown in blue, we see that as Λ_{had} increases the multiplicity decreases and asymptotes to 2. The plasmalike case (shown in orange) requires a minimum value of $c = 2$ to allow each plasma ball to emit at least two glueballs. At this limit, only the lightest glueball can be produced. Naively we would expect multiplicity to increase with c , because even as increasing c logarithmically suppresses the number of splittings in the shower, it linearly increases the number of glueballs emitted by the larger plasma states. Somewhat surprisingly, from $c = 2$ to ~ 3 the multiplicity decreases. This is due to the fact the plasma still only decays to two glueballs most of the time across this range, but now with greater access to the higher-mass glueball states. Not until $c \sim 4$ is the entire glueball spectrum sampled and produced in plasma ball decays, and further increasing c now does lead to an increase in glueballs produced by the plasma. Finally, as discussed in the previous subsection, when $\Lambda_{\text{had}} \rightarrow M$ only a single plasma ball is produced in the event, resulting in fully SUEP-like behavior.

While Fig. 7 depicts both hadronization interpretations across the full range of c values, each case is physically sensible in separate regimes. The jetlike interpretation assumes the initial flux-tube fragments directly into final-state glueballs, which is only sensible for values of c close to 1. Larger values force gluons with virtualities $\sqrt{t} \gg m_0 \sim 6\Lambda$ into a single low-mass glueball, which seems very implausible given the *a priori* reliable perturbative prediction for a much higher jet mass of that branch of the shower. We therefore adopt $c = 1, 2$ as benchmark values for the hadronization scale in the jetlike regime.

In contrast, the plasmalike case assumes fragmentation into high-mass plasma states, and small values $c \lesssim 4$ seem to imply that these plasma balls have a large bias of evaporating

only into the light glueball states. Sampling the full range of plasmalike behavior therefore motivates picking slightly larger values of c . On the other hand, the plasmalike scenario already defies the most reasonable physical expectation based on our understanding of perturbative QCD and flux-tube dynamics, and a very large value of the plasma ball mass pushes this scenario into even more implausible regimes. For example, $c = 6$ corresponds to ending the perturbative shower at $\sim 70\Lambda$, where $SU(N_c)$ should be entirely perturbative. It should therefore serve as a suitable ceiling for the possible range of plasma masses produced under this exotic assumption for the nonperturbative




FIG. 7. Effect of $c = \Lambda_{\text{had}}/(2m_0)$ on average glueball multiplicity for exotic Higgs decays ($M = 125$ GeV) into glueballs with $m_0 = 10$ GeV. In the jetlike case, Λ_{had} is the hadronization scale at which the shower is terminated and gluons turned into glueballs. In the plasmalike case, it is the twice the mass of color-singlet plasma balls produced in the shower, which then evaporate into glueballs. In both cases we take $T_{\text{had}}/T_c = 1$. The red dashed line indicates $\Lambda_{\text{had}} = M$, resulting in the production of just two glueballs in the jetlike case and a single plasma ball of mass M in the plasmalike case.

behavior of $N_f = 0$ QCD. We therefore adopt $c = 4, 6$ as benchmark values for the hadronization scale in the plasmalike regime.

The hadronization temperature $T_{\text{had}} = d \cdot T_c$ is less constrained by physical arguments and relatively unimportant for inclusive observables across glueball species. However, T_{had} dominantly determines the exclusive predictions for each type of glueball. Given the exponential dependence of both glueball momentum in the plasmalike case and relative multiplicities in both jet- and plasmalike cases on T_{had} , sampling $d = 1, 2$ should span a wide range of physically plausible predictions.

In summary, careful phenomenological studies involving dark glueball production should compute physical predictions for a variety of different values of the nuisance parameters in GlueShower. For the jetlike case, $c = 1, 2$ and $d = 1, 2$ should be simulated. For the plasmalike case, $c = 4, 6$ and $d = 1, 2$. To be conservative, all eight benchmark points should be used to define the systematic error bar on predictions.

D. New predictions for glueball production

We now have in hand a physically motivated simulation of the glueball production process, as well as eight hadronization benchmarks which span the range of physically reasonable possible outcomes. This allows us to make fully self-consistent predictions for glueball production with accurate theoretical uncertainties included for the first time.

Naively, we would expect collider signals of dark glueball production to be most sensitive to the multiplicity and decay mode of the shortest-lived or most visibly decaying glueball state, for example the 0^{++} if decay proceeds via the Higgs portal. Conversely, indirect detection of dark matter annihilating into dark glueballs [64] will be affected by the relative distributions and decays of all unstable glueball species. To give a feeling for how each of these two types of studies might be affected by theoretical uncertainties, we show in Fig. 8 how average multiplicity and energy predictions for all glueballs inclusively and for the 0^{++} exclusively change across our range of possible

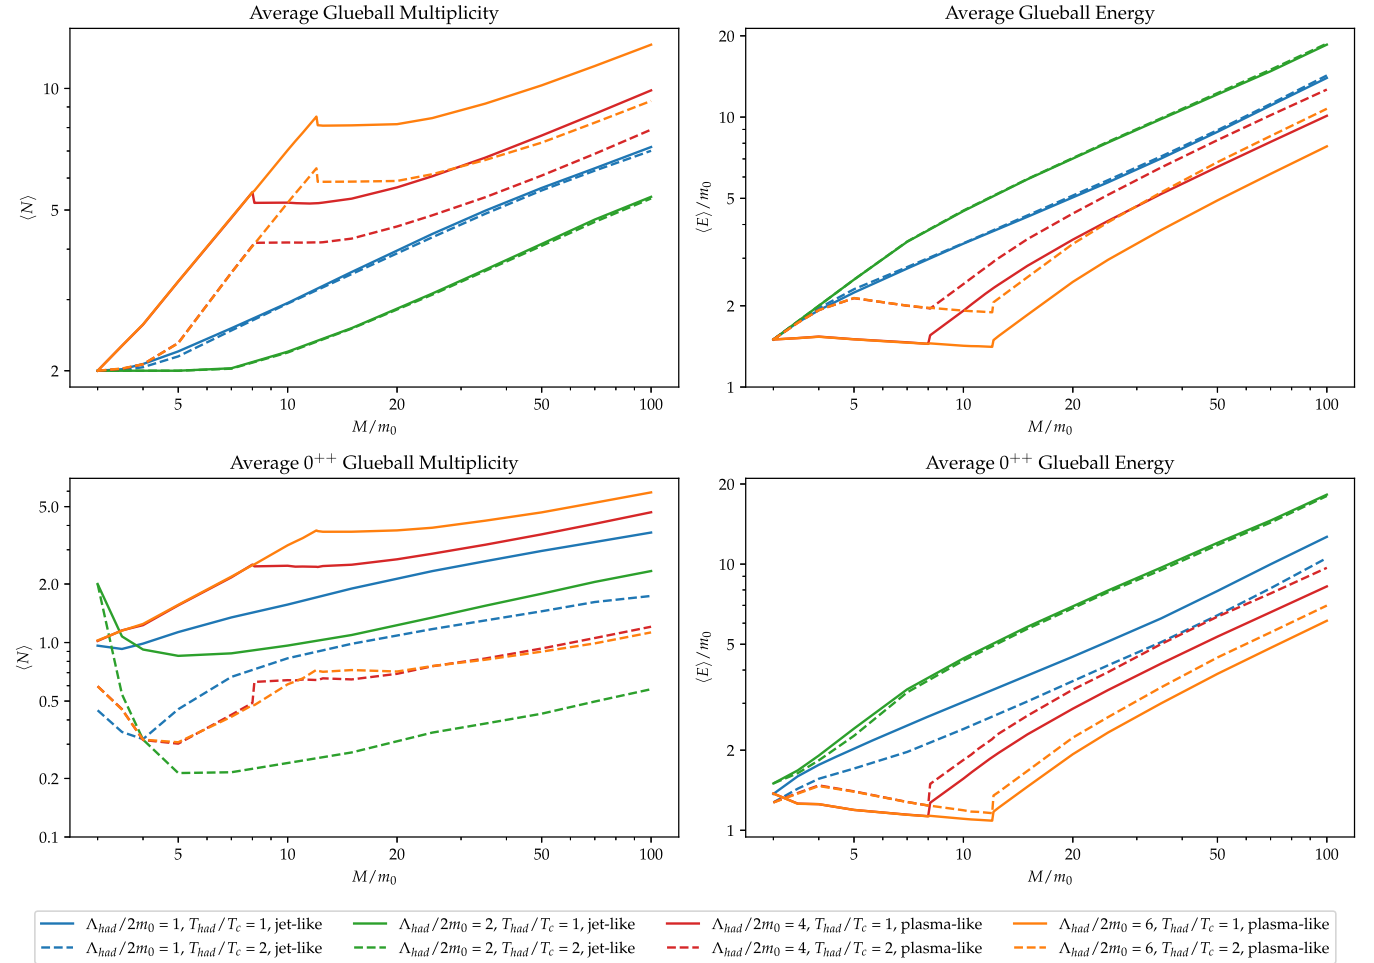


FIG. 8. Dependence of average glueball multiplicity and energy on M/m_0 for the eight hadronization benchmarks with $N_c = 3$. The range of predictions can be interpreted as our current theoretical uncertainty on glueball production: a factor of a few on average glueball energy and inclusive multiplicity and a factor of 10 on exclusive 0^{++} multiplicity.




FIG. 9. 0^{++} glueball energy spectra for $M/m_0 = 10$ (left) and 100 (right) and $N_c = 3$, for our eight hadronization benchmarks.

hadronization benchmarks. We take the variation across all benchmarks to indicate the theoretical uncertainty for each observable.

We find that average glueball multiplicity, average glueball energy and average 0^{++} energy have an uncertainty of about a factor of 3 across the range of considered center-of-mass energies M/m_0 . On the other hand, the exclusive multiplicity of the 0^{++} state has a much larger spread of possible predictions, roughly a factor of 10 across the hadronization benchmarks. Slightly more can be said if one is willing to ascribe different priors to the default jetlike versus the more exotic plasmlike hadronization hypotheses. Jetlike showers produce lower multiplicity jets of higher-energy glueballs, while plasmlike showers produce fatter jets with higher multiplicities of softer glueballs.

Figure 9 compares the 0^{++} energy spectra for different hadronization benchmarks. At low center-of-mass energies, we find very significant differences between the jetlike and plasmlike showers. While the plasmlike case favors low-energy glueballs, the jetlike case is roughly flat, with dominant contributions by simple two-body glueball production. At high center-of-mass energies, the spectral shapes become much more similar, with low-energy final states favored, to varying degrees, across the benchmarks.

VI. CONCLUSION

Hidden valleys are an extremely popular hypothesis for BSM physics. They may address fundamental mysteries like the hierarchy problem or the nature of dark matter, and their signatures are targeted by many new LHC searches (e.g., [18–20,30,31]) and proposed detectors [95–98]. However, the seemingly simple and minimal case of $N_f = 0$ hidden

QCD has undergone very little quantitative study, despite its high theoretical motivation within frameworks like neutral naturalness [9–17]. This can be traced back to our lack of understanding of pure glue hadronization.

In this work, we show that significant progress can be made by combining a perturbative pure glue parton shower with a self-consistent and physically motivated parametrization of the unknown nonperturbative physics. This is in large part enabled by the modest hierarchy between the glueball mass and the confinement scale $m_0/\Lambda \sim 6$ in $SU(N_c)$ theories. We make our simulation code available as the public PYTHON code `GlueShower`, the first glueball generator for hidden valley theories [63]. This can be combined with other Monte Carlo generators for dark gluon production and dark glueball decay [38,39] for a complete simulation of glueball production within a given dark sector scenario.

Our approach incorporates what we consider to be an almost maximally wide range of possible hadronization scenarios, to make sure no physically reasonable possibility escapes our study. We argue that glueball production proceeds largely similar to jetlike hadron production in SM QCD, but we also consider the possibility that highly exotic nonperturbative physics of pure Yang-Mills theories somehow favors the production of high-mass color-singlet gluon-plasma states, which evaporate via isotropic emission of thermal glueballs in their rest frame. In both jetlike and plasmlike modes, a range of color-singlet formation scales and hadronization temperatures can be selected to generate a range of possible hadronization behaviors. We intend `GlueShower` to be a starting point for $N_f = 0$ dark sector studies and point out future improvements that could be implemented. A careful treatment of color confinement

or a study of gluons in the ropelike picture [88] could lead to a more sophisticated hadronization model, but the current version represents a physically reasonable starting point for phenomenological studies.

We define a set of eight hadronization benchmark scenarios, four in each of the above two modes, which we argue span the range of physically remotely reasonable hadronization outcomes for glueball production. We then use `GlueShower` to make some preliminary studies of glueball production observables, with theoretical uncertainties obtained from the variation across the different hadronization benchmarks. Uncertainties range from modest $\mathcal{O}(1)$ factors to orders of magnitude, depending on the observable.

Accurately taking these uncertainties into account will be important for designing and interpreting future searches and constraints. In particular, our work could highlight which observables are more robust with respect to unknown details of glueball hadronization, encouraging a variety of collider and astrophysical probes to assemble a complete picture of the dark sector fundamental parameters and nonperturbative dynamics. It is our hope that this work enables many new studies and searches of dark sectors that were previously intractable.

ACKNOWLEDGMENTS

We especially thank Matt Strassler for many insightful conversations and discussions. We also thank Jared Barron and Benjamin Fuks for helpful conversations. The research of D. C. and C. G. was supported in part by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada, the Canada Research Chair program, the Alfred P. Sloan Foundation, and the Ontario Early Researcher Award. The work of C. G. was also supported by the University of Toronto Connaught International Scholarship.

APPENDIX: PERTURBATIVE QCD SHOWER

We provide a step-by-step description of how `GlueShower` generates perturbative showers of gluons. This is a completely standard lowest-order perturbative shower with angular ordering, implemented following the `PYTHIA` manual [72], but we include this summary here for completeness and pedagogical purposes.

The two initial gluons are treated separately to ensure energy-momentum conservation between their initial virtualities and energies. The shower is then evolved iteratively for all the subsequent daughter dark gluons.

1. Initial gluon setup

- (1) Find virtuality of gluon 1 (t_1) assuming gluon 2 is on shell. This is done by starting at the maximum allowed value $(M - m_{\min})^2$ and evolving down using the Monte Carlo method described in Sec. III.

There is some ambiguity regarding what on shell means when the final-state glueball species is still unspecified. For consistency across the shower, we define the minimum mass as half the hadronization scale, $c \cdot m_0$.

- (2) Once t_1 is found, if this allows the gluon to split, z_1 is determined using the value of t_1 .
- (3) Steps 1 and 2 are repeated for gluon 2, to find t_2 , and possibly z_2 if gluon 2 also splits.
- (4) Using these values, a range of vetoes are checked before continuing with the shower.
 - (a) First check that $t_1 + t_2 < M$; if false, then evolve the larger of the two virtualities to a smaller value.
 - (b) $z_{1,2}$ are found assuming that the other gluon was on shell, but now they have virtualities of $t_{1,2}$; check that $z_{1,2}$ remain within the updated allowed range. If a gluon fails this check, its virtuality is evolved to smaller values.
 - (c) Above vetoes are repeated until passed.
- (5) We now have values for the virtuality and splitting energy of each initial gluon, which are then used to evolve the daughters. Note that the 4-vectors of the initial gluons are now entirely determined.

2. Evolving the daughters

We now outline the steps applied for general gluon evolution at any point along the shower, excluding the initial gluons. We label the produced daughter gluons in this step 3 and 4.

- (1) First, the coordinate system of this new splitting is established:
 - (a) The z axis is aligned with the direction of the parent gluon's momentum.
 - (b) A random angle is chosen in the $x - y$ plane for the perpendicular momentum of the daughters.
- (2) Find initial t and z guess for the daughters.
 - (a) Initial energies are given by $E_3 = z_1 * E_1$ and $E_4 = (1 - z_1) * E_1$.
 - (b) The starting virtuality $t_{\text{start},i}$ for each daughter is given by $\min(E_i^2, m_1^2)$.
 - (c) Evolve the daughters down to find first guesses for t_3 and t_4 . If splitting is allowed, use these values to find z_3 and z_4 .
- (3) Apply unconstrained or constrained evolution for each daughter. In unconstrained evolution, the maximum possible kinematic z range is given when assuming the daughters have virtuality m_{\min}^2 . Thus, when the daughters are evolved and their actual virtualities are found, $t_{3,4}$, the originally assumed z value may lie outside the newly determined range. The solution is to redefine the z value. This is done by taking the original four momenta of the daughters and boosting them to their center-of-mass frame, rescaling the vectors to reflect their determined evolved virtuality, and then boosting them back to

the lab frame. This leads to a new z value, matching the new energies of the rescaled dark gluons, that lies within the required range and is equivalent to using Eq. (10.11) given in [72].

For constrained evolution, the final masses of the daughter gluons are constrained by the requirement that z remain in the originally determined range. The default option enabled in the `GlueShower` code is unconstrained evolution, as used in `PYTHIA`, and leads to a higher rate of splittings in the shower.

- (4) Impose angular ordering. This is the result of soft gluon coherence effects which cause the opening angle between the daughter gluons to be smaller than the opening angle of the parent gluon. This can

be described as the suppression of wide-angle gluon emission. Practically it is enforced by determining the absolute maximum opening angle of the daughters, by assuming their daughters have mass m_{\min} , and if this angle is bigger than the parent's opening angle, the daughter is evolved to lower virtuality. Note that if a gluon is evolved to satisfy angular ordering, step 4 is repeated to ensure that the unconstrained evolution condition still holds.

- (5) With the t , z values of the daughter gluons determined, the steps of this section are repeated iteratively until all dark gluons in the shower have reached the hadronization scale.

-
- [1] M. J. Strassler and K. M. Zurek, *Phys. Lett. B* **651**, 374 (2007).
- [2] J. Kang and M. A. Luty, *J. High Energy Phys.* **11** (2009) 065.
- [3] Y. Bai and P. Schwaller, *Phys. Rev. D* **89**, 063522 (2014).
- [4] S. Renner and P. Schwaller, *J. High Energy Phys.* **08** (2018) 052.
- [5] H. Mies, C. Scherb, and P. Schwaller, *J. High Energy Phys.* **04** (2021) 049.
- [6] B. Holdom, *Phys. Lett.* **166B**, 196 (1986).
- [7] B. Patt and F. Wilczek, [arXiv:hep-ph/0605188](https://arxiv.org/abs/hep-ph/0605188).
- [8] A. Falkowski, J. Juknevich, and J. Shelton, [arXiv:0908.1790](https://arxiv.org/abs/0908.1790).
- [9] Z. Chacko, H.-S. Goh, and R. Harnik, *Phys. Rev. Lett.* **96**, 231802 (2006).
- [10] N. Craig, A. Katz, M. Strassler, and R. Sundrum, *J. High Energy Phys.* **07** (2015) 105.
- [11] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, *J. High Energy Phys.* **02** (2007) 009.
- [12] R. Barbieri, T. Gregoire, and L. J. Hall, [arXiv:hep-ph/0509242](https://arxiv.org/abs/hep-ph/0509242).
- [13] Z. Chacko, Y. Nomura, M. Papucci, and G. Perez, *J. High Energy Phys.* **01** (2006) 126.
- [14] H. Cai, H.-C. Cheng, and J. Terning, *J. High Energy Phys.* **05** (2009) 045.
- [15] D. Poland and J. Thaler, *J. High Energy Phys.* **11** (2008) 083.
- [16] T. Cohen, N. Craig, G. F. Giudice, and M. McCullough, *J. High Energy Phys.* **05** (2018) 091.
- [17] H.-C. Cheng, L. Li, E. Salvioni, and C. B. Verhaaren, *J. High Energy Phys.* **05** (2018) 057.
- [18] J. Alimena *et al.*, *J. Phys. G* **47**, 090501 (2020).
- [19] G. Aad *et al.* (ATLAS Collaboration), *J. Instrum.* **8**, P07015 (2013).
- [20] G. Aad *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **80**, 450 (2020).
- [21] M. J. Strassler, [arXiv:0801.0629](https://arxiv.org/abs/0801.0629).
- [22] S. Knapen, S. Pagan Griso, M. Papucci, and D. J. Robinson, *J. High Energy Phys.* **08** (2017) 076.
- [23] J. Barron, D. Curtin, G. Kasieczka, T. Plehn, and A. Spurdalakis, *J. High Energy Phys.* **12** (2021) 129.
- [24] S. Knapen, J. Shelton, and D. Xu, *Phys. Rev. D* **103**, 115013 (2021).
- [25] T. Cohen, M. Lisanti, and H. K. Lou, *Phys. Rev. Lett.* **115**, 171804 (2015).
- [26] T. Cohen, M. Lisanti, H. K. Lou, and S. Mishra-Sharma, *J. High Energy Phys.* **11** (2017) 196.
- [27] T. Cohen, J. Doss, and M. Freytsis, *J. High Energy Phys.* **09** (2020) 118.
- [28] P. Schwaller, D. Stolarski, and A. Weiler, *J. High Energy Phys.* **05** (2015) 059.
- [29] D. Linthorne and D. Stolarski, *Phys. Rev. D* **104**, 035019 (2021).
- [30] A. M. Sirunyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **02** (2019) 179.
- [31] A. Tumasyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **06** (2022) 156.
- [32] C. J. Morningstar and M. J. Peardon, *Phys. Rev. D* **60**, 034509 (1999).
- [33] B. Lucini and G. Moraitis, *Phys. Lett. B* **668**, 226 (2008).
- [34] M. J. Teper, [arXiv:hep-th/9812187](https://arxiv.org/abs/hep-th/9812187).
- [35] B. Lucini, A. Rago, and E. Rinaldi, *J. High Energy Phys.* **08** (2010) 119.
- [36] A. Athenodorou and M. Teper, *J. High Energy Phys.* **12** (2021) 082.
- [37] N. Yamanaka, A. Nakamura, and M. Wakayama, *Proc. Sci., LATTICE2021* (**2022**) 447 [[arXiv:2110.04521](https://arxiv.org/abs/2110.04521)].
- [38] J. E. Juknevich, *J. High Energy Phys.* **08** (2010) 121.
- [39] J. E. Juknevich, D. Melnikov, and M. J. Strassler, *J. High Energy Phys.* **07** (2009) 055.
- [40] A. E. Faraggi and M. Pospelov, *Astropart. Phys.* **16**, 451 (2002).
- [41] K. K. Boddy, J. L. Feng, M. Kaplinghat, and T. M. P. Tait, *Phys. Rev. D* **89**, 115017 (2014).

- [42] K. K. Boddy, J. L. Feng, M. Kaplinghat, Y. Shadmi, and T. M. P. Tait, *Phys. Rev. D* **90**, 095016 (2014).
- [43] I. Garcia Garcia, R. Lasenby, and J. March-Russell, *Phys. Rev. D* **92**, 055034 (2015).
- [44] A. Soni and Y. Zhang, *Phys. Rev. D* **93**, 115025 (2016).
- [45] A. Soni, H. Xiao, and Y. Zhang, *Phys. Rev. D* **96**, 083514 (2017).
- [46] L. Forestell, D. E. Morrissey, and K. Sigurdson, *Phys. Rev. D* **95**, 015032 (2017).
- [47] L. Forestell, D. E. Morrissey, and K. Sigurdson, *Phys. Rev. D* **97**, 075029 (2018).
- [48] N. Yamanaka, H. Iida, A. Nakamura, and M. Wakayama, *Phys. Lett. B* **813**, 136056 (2021).
- [49] N. Yamanaka, H. Iida, A. Nakamura, and M. Wakayama, *Phys. Rev. D* **102**, 054507 (2020).
- [50] B. Jo, H. Kim, H. D. Kim, and C. S. Shin, *Phys. Rev. D* **103**, 083528 (2021).
- [51] D. N. Spergel and P. J. Steinhardt, *Phys. Rev. Lett.* **84**, 3760 (2000).
- [52] D. H. Weinberg, J. S. Bullock, F. Governato, R. Kuzio de Naray, and A. H. G. Peter, *Proc. Natl. Acad. Sci. U.S.A.* **112**, 12249 (2015).
- [53] D. Curtin and C. B. Verhaaren, *J. High Energy Phys.* **12** (2015) 072.
- [54] Z. Chacko, D. Curtin, and C. B. Verhaaren, *Phys. Rev. D* **94**, 011504 (2016).
- [55] G. Burdman and G. Lichtenstein, *J. High Energy Phys.* **08** (2018) 146.
- [56] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Phys. Rep.* **97**, 31 (1983).
- [57] B. Webber, *Nucl. Phys.* **B238**, 492 (1984).
- [58] M. Bahr *et al.*, *Eur. Phys. J. C* **58**, 639 (2008).
- [59] D. Amati and G. Veneziano, *Phys. Lett.* **83B**, 87 (1979).
- [60] L. Carloni and T. Sjostrand, *J. High Energy Phys.* **09** (2010) 105.
- [61] L. Carloni, J. Rathsmann, and T. Sjostrand, *J. High Energy Phys.* **04** (2011) 091.
- [62] C. Bierlich *et al.*, [arXiv:2203.11601](https://arxiv.org/abs/2203.11601).
- [63] <https://github.com/davidrcurtin/GlueShower>.
- [64] D. Curtin and C. Gemmell (to be published).
- [65] G. M. Prospero, M. Raciti, and C. Simolo, *Prog. Part. Nucl. Phys.* **58**, 387 (2007).
- [66] B. Batell, W. Hu, and C. B. Verhaaren, *J. High Energy Phys.* **08** (2020) 009.
- [67] C. Kilic, C. B. Verhaaren, and T. Youn, *Phys. Rev. D* **104**, 116018 (2021).
- [68] W.-S. Hou, C.-S. Luo, and G.-G. Wong, *Phys. Rev. D* **64**, 014028 (2001).
- [69] N. Ishii, H. Suganuma, and H. Matsufuru, in *Workshop on Lepton Scattering, Hadrons and QCD* (2001), pp. 252–259, [arXiv:hep-lat/0106004](https://arxiv.org/abs/hep-lat/0106004).
- [70] M. Loan and Y. Ying, *Prog. Theor. Phys.* **116**, 169 (2006).
- [71] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and Collider Physics* (Cambridge University Press, Cambridge, England, 2011), Vol. 8.
- [72] T. Sjostrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [73] G. Marchesini and B. R. Webber, *Nucl. Phys.* **B238**, 1 (1984).
- [74] G. 't Hooft, *Nucl. Phys.* **B72**, 461 (1974).
- [75] J. Bulava, B. Hörz, F. Knechtli, V. Koch, G. Moir, C. Morningstar, and M. Peardon, *Phys. Lett. B* **793**, 493 (2019).
- [76] P. Minkowski and W. Ochs, *Phys. Lett. B* **485**, 139 (2000).
- [77] R. Field and R. Feynman, *Nucl. Phys.* **B136**, 1 (1978).
- [78] J. Schwinger, *Phys. Rev.* **82**, 664 (1951).
- [79] A. Casher, H. Neuberger, and S. Nussinov, *Phys. Rev. D* **20**, 179 (1979).
- [80] J. L. Manes, *Nucl. Phys.* **B621**, 37 (2002).
- [81] P. Blanchard, S. Fortunato, and H. Satz, *Eur. Phys. J. C* **34**, 361 (2004).
- [82] J. Noronha-Hostler, J. Noronha, and C. Greiner, *J. Phys. G* **37**, 094062 (2010).
- [83] P. Petreczky, *J. Phys. G* **39**, 093002 (2012).
- [84] B. Lucini, M. Teper, and U. Wenger, *J. High Energy Phys.* **01** (2004) 061.
- [85] G. Boyd, J. Engels, F. Karsch, E. Laermann, C. Legeland, M. Lutgemeier, and B. Petersson, *Nucl. Phys.* **B469**, 419 (1996).
- [86] B. Lucini, M. Teper, and U. Wenger, *J. High Energy Phys.* **02** (2005) 033.
- [87] B. Lucini, A. Rago, and E. Rinaldi, *Phys. Lett. B* **712**, 279 (2012).
- [88] T. S. Biro, H. B. Nielsen, and J. Knoll, *Nucl. Phys.* **B245**, 449 (1984).
- [89] C. Bierlich, *EPJ Web Conf.* **171**, 14003 (2018).
- [90] M. Swisdak, *Phys. Plasmas* **20**, 062110 (2013).
- [91] G. Colangelo and P. Nason, *Phys. Lett. B* **285**, 167 (1992).
- [92] P. A. Zyla *et al.* (Particle Data Group), *Prog. Theor. Exp. Phys.* **2020**, 083C01 (2020).
- [93] D. de Florian, M. Epele, R. J. Hernandez-Pinto, R. Sassot, and M. Stratmann, *Phys. Rev. D* **95**, 094019 (2017).
- [94] A. Salas-Chavira, Z. Fan, and H.-W. Lin, [arXiv:2112.03124](https://arxiv.org/abs/2112.03124).
- [95] D. Curtin *et al.*, *Rep. Prog. Phys.* **82**, 116201 (2019).
- [96] C. Alpigiani *et al.* (MATHUSLA Collaboration), [arXiv:2009.01693](https://arxiv.org/abs/2009.01693).
- [97] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **97**, 035001 (2018).
- [98] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, *Phys. Rev. D* **97**, 015023 (2018).