Editors' Suggestion

Flavored jets with exact anti-k_t kinematics and tests of infrared and collinear safety

Fabrizio Caola,[1] Radosław Grabarczyk,[1] Maxwell L. Hutt,[1,2] Gavin P. Salam,[1,3] Ludovic Scyboz,[1] and Jesse Thaler[4]

[1]*Rudolf Peierls Centre for Theoretical Physics, Parks Road, Oxford OX1 3PU, United Kingdom*
[2]*The Blackett Laboratory, Imperial College London, Prince Consort Road, London, SW7 2AZ, United Kingdom*
[3]*All Souls College, Oxford OX1 4AL, United Kingdom*
[4]*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

(Received 11 July 2023; accepted 6 October 2023; published 6 November 2023)

We propose extensions of the anti-k_t and Cambridge/Aachen hierarchical jet clustering algorithms that are designed to retain the exact jet kinematics of these algorithms, while providing an infrared-and-collinear-safe definition of jet flavor at any fixed order in perturbation theory. Central to our approach is a new technique called interleaved flavor neutralization (IFN), whereby the treatment of flavor is integrated with, but distinct from, the kinematic clustering. IFN allows flavor information to be meaningfully accessed at each stage of the clustering sequence, which enables a consistent assignment of flavor both to individual jets and to their substructure. We validate the IFN approach using a dedicated framework for fixed-order tests of infrared and collinear safety, which also reveals unanticipated issues in earlier approaches to flavored jet clustering. We briefly explore the phenomenological impact of IFN with anti-k_t jets for benchmark tasks at the Large Hadron Collider.

DOI: 10.1103/PhysRevD.108.094010

I. INTRODUCTION

The use of jet clustering algorithms is essential and ubiquitous at colliders. Jet algorithms relate collimated sprays of energetic hadrons to the underlying concept of hard, perturbative quarks and gluons (or, more generally, partons). In the vast majority of cases, only the kinematics of the resulting jets are used for analysis. Insofar as jets are meant to represent the underlying partonic structure of an event, though, it is natural to ask whether jets can also reflect the flavor of the underlying partons, for example, their quark or gluon nature. The question of how to formulate a jet algorithm where the flavors assigned to jets are infrared and collinear (IRC) safe was first posed in 2006 [1,2]. The algorithm developed there, flavor k_t, based on a modification of the k_t algorithm [3–5], appeared to be successful in this task. However, one of the characteristics of flavor k_t was that the kinematics of the resulting jets depended on the flavor of the underlying constituents being clustered.

In modern jet usage, where the subsequently developed anti-k_t algorithm [6] has found widespread applications, a flavor-induced modification of the jets' kinematics is undesirable. Notably, it has been found to complicate unfolding corrections [7]. Nevertheless, there are situations where IRC-safe flavored jet algorithms would be highly beneficial. For example, the question of IRC-safe jet flavor has recently come to the fore in the context of heavy-flavor jets [8–11]. IRC safety in this instance ensures that flavored jet cross sections do not contain any logarithms of the ratio of the jet transverse momentum p_t to the quark mass m_q. It also makes it possible to use an $m_q = 0$ approximation in fixed-order perturbative calculations [7,12–21], with an expectation that any missing contributions are suppressed by powers of m_q/p_t.

In this article, we present a new strategy for flavored jet finding called interleaved flavor neutralization (IFN), which is designed to combine an IRC-safe definition of jet flavor with the IRC-safe kinematics of sequential clustering. We will study IFN with two generalized-k_t-style jet algorithms: the anti-k_t algorithm, used extensively at the LHC, and the Cambridge/Aachen (C/A) algorithm [22,23], widely favored for jet substructure studies. In the case of the anti-k_t algorithm, our objectives are similar to those of the recent "flavored anti-k_t" [10] and "flavor dressing" [11] algorithms, which respectively achieve approximate and

exact anti-k_t kinematics. Like flavor dressing, IFN yields exact anti-k_t (or C/A) kinematics, but because it integrates flavor information at each stage of the clustering sequence, it is a viable candidate for jet substructure studies. We also carry out a much more extensive set of IRC safety tests than in any prior work, which support the conclusion that IFN is IRC safe, at least through order α_s^6. These tests also reveal unexpected and subtle issues in the default formulations of all prior flavored jet algorithms.

We focus on the theoretical definition of jet flavor, leaving a study of experimental issues to future work. The extent to which any IRC-safe flavor algorithm can be adopted experimentally is an open question. Even when identifying heavy-flavor jets, where collinear singularities are regulated by a nonzero m_q, such algorithms would typically require the identification of all heavy-flavored hadrons in an event. That is challenging when there are multiple heavy-flavored hadrons in a single jet or when some of the heavy-flavored hadrons have low momenta.[1] Despite these experimental subtleties, the underlying question of IRC-safe flavor identification remains conceptually important. Jet flavor can provide a valuable tool in a range of theoretical work, for example, in matching parton showers and fixed-order calculations [24]. One can further anticipate that it will be useful in testing logarithmic accuracy for flavor-related aspects of parton showers [25,26].

The remainder of this article is organized as follows. In Sec. II, we review the key features of widely adopted jet algorithms and some of the issues that arise when flavor tagging is sought. We also briefly describe existing proposals for those flavored jet algorithms that aim to achieve all-order IRC safety [1,2,10,11]. In Sec. III, we outline our general design aims for a modern flavored jet algorithm and present a concrete realization via IFN. That section also includes a discussion of some of the subtle considerations brought about by IRC-safety requirements. In Sec. IV, we present the framework that we developed to explore IRC-safety issues in some depth (a substantial extension of the approach developed some time ago for testing the SISCone jet algorithm [27]), which we apply both to our IFN proposal and to earlier flavored jet algorithms. These tests expose unanticipated issues in earlier proposals, many of them connected with the treatment of initial-state radiation in a hadron collider context. In some cases, we identify simple adaptations of the original algorithms that should make them IRC safe. In Sec. V, we perform three benchmark phenomenological studies to illustrate the behavior of various flavored jet algorithms, restricting ourselves to the ones that pass our IRC-safety tests. In Sec. VI, we briefly present the adaptation of our approach to e^+e^- colliders. We conclude in Sec. VII.

[1]Though, as we will see in Sec. V B for the $t\bar{t}$ process, this may be less of an issue than one might fear.

Additional material is presented in the Appendixes. In Appendix A, we review some features of the double-soft quark emission current that we used for our analyses. In Appendix B, we perform numerical tests to justify some of the design decisions we made for the IFN algorithm. In Appendix C, we provide detailed analyses of the main IRC-safety issues that we encountered in this work. In Appendix D, we present summary plots of IRC-safety tests for those algorithms that we expect to be IRC safe.

II. REMINDERS ABOUT EXISTING JET ALGORITHMS

In this section, we briefly review standard jet algorithms and their interplay with jet flavor, including the original flavor-k_t approach [1]. To avoid confusion, we refer to the flavored anti-k_t algorithm of Ref. [10] as "CMP" and the flavor dressing algorithm of Ref. [11] as "GHS." Throughout this section and most of this article, we concentrate on longitudinally invariant hadron-collider algorithms, with a brief mention of an e^+e^- adaptation in Sec. VI.

A. Flavorless kinematic clustering

Let us start with a reminder of how the generalized-k_t algorithm works. It employs distances d_{ij} between each pair of pseudojets[2] i and j and d_{iB} between each pseudojet i and the beam:

$$d_{ij} = \min(p_{ti}^{2p}, p_{tj}^{2p}) \frac{\Delta R_{ij}^2}{R^2}, \tag{1a}$$

$$\Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2, \tag{1b}$$

$$d_{iB} = p_{ti}^{2p}, \tag{1c}$$

where p_{ti}, y_i, and ϕ_i are, respectively, the transverse momentum, rapidity, and azimuth of i ($y_i = \frac{1}{2}\ln\frac{E_i+p_{zi}}{E_i-p_{zi}}$). The algorithm has two parameters, the jet radius R, which sets the angular reach of the jets, and the power p, which sets the nature of the algorithm: -1, 0, 1, respectively, for the anti-k_t [6], Cambridge/Aachen [22,23], and k_t algorithms [4,5]. The algorithm starts with all event particles and proceeds as follows:

(1) Identify the smallest of the d_{ij} and d_{iB} among all i and j at this stage of the clustering.
(2) If it is a d_{ij}, recombine i and j into a single new pseudojet, and return to step 1.
(3) If it is a d_{iB}, declare i to be a jet, and remove it from the list of pseudojets to be considered at subsequent clustering steps; return to step 1.

[2]Recall that a pseudojet may be either a single particle or the combination of more than one particle arising from an earlier stage of the clustering.

The clustering stops once no pseudojets are left to be clustered. Given the resulting jets, it is common to consider only the subset that pass minimum p_t (and maximum rapidity or pseudorapidity[3]) constraints.

B. Flavor via recombination scheme

A crucial element of the jet definition is the choice of recombination scheme. The most common is the (somewhat inappropriately named) E scheme, in which 4-momenta are simply added. Flavor is usually not considered within standard jet algorithms, but it is useful to introduce three potential flavor recombination schemes:

(i) *Any-flavor scheme*. This scheme is relatively close to typical experimental practice for b and c tagging. Here, any recombination that involves nonzero flavor, e.g., $q+g$, $\bar{q}+g$, or $q+\bar{q}$, yields a flavored result. From a theoretical point of view, this scheme is collinear unsafe for massless quarks due to the collinear divergence of $g \to q\bar{q}$ splitting. For massive quarks, as in the case of b and c production, this scheme is logarithmically sensitive to the quark mass. We will further consider this "any-flavor" scheme only in a phenomenological context in Sec. V B.

(ii) *Net-flavor scheme*. This is a theoretically better-motivated scheme that considers the net flavor in the recombination. In this scheme, a q carries flavor, a $\bar{q}$ carries antiflavor, and a $q\bar{q}$ carries no flavor. This "net-flavor" scheme resolves the collinear unsafety for $g \to q\bar{q}$ splitting.

(iii) *Flavor modulo-2 scheme*. Typically for heavy flavor at hadron level, it is not conceptually possible to distinguish flavor from antiflavor, e.g., because of $B_0 - \bar{B}_0$ oscillations. In such a situation, one may consider a "flavor modulo-2" scheme (see, e.g., Ref. [2]). Specifically, b and $\bar{b}$ are treated as equivalent, while $b\bar{b}$, bb, and $\bar{b}\bar{b}$ are all considered to be flavorless. This scheme also resolves the issue of collinear unsafety for $g \to q\bar{q}$ splitting.

While the net flavor and modulo-2 options ensure that the jet flavor is unaffected by collinear divergences for $g \to q\bar{q}$ splittings, they still exhibit IRC safety issues for jet flavor at higher orders, at least when used with standard jet algorithms. This occurs at next-to-next-to-leading order (NNLO), as discussed in Ref. [1] and illustrated in Fig. 1 (see Appendix A for further discussion about the matrix element for this process). Specifically, when a soft gluon splits to a large-angle $q\bar{q}$ pair, one or another of the resulting soft quarks can be clustered with a hard jet, and the net-flavor and modulo-2 recombination schemes result in an IRC-unsafe flavor for hard jets, with the divergence appearing as $\alpha_s^2 \ln p_{t,\mathrm{jet}}/m_q$ for a finite quark mass m_q. This is the classic problem when attempting to obtain IRC-safe jet flavor.

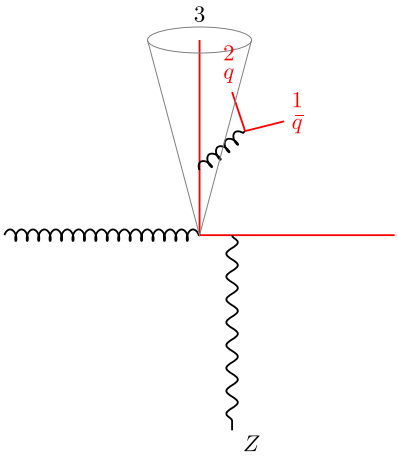

FIG. 1. Classic problematic flavor configuration at NNLO. A soft gluon at large angle splits to a $\bar{q}q$ pair (labeled 1 and 2), and the flavor of the hard jet (numbered 3) is polluted by the flavor of 2, while 1 ends up outside the jet.

When considering more than one flavor (e.g., all of $udscb$), flavor recombination is typically applied separately for each flavor. This may be done either within a single run of the algorithm or (for algorithms where the flavor does not affect the jet kinematics) applying the flavor part of the jet clustering in one separate run of the algorithm for each kind of flavor.

C. Existing flavored jet algorithms

We now review three jet-flavor definitions that aim to achieve all-order IRC safety (see Refs. [8,9,29] for alternative definitions of jet flavor).

1. Flavor k_t

The flavor-k_t [1] algorithm took the approach of using a net or modulo-2 flavor scheme, while modifying the clustering distances relative to Eq. (1). Specifically, it modifies the standard k_t ($p=1$) distance when the softer of i and j is flavored

$$d_{ij}^{\text{flav-}k_t} = [\max(p_{ti}, p_{tj})]^{\alpha}[\min(p_{ti}, p_{tj})]^{2-\alpha}\frac{\Delta R_{ij}^2}{R^2}, \quad \text{if the softer of } i \text{ and } j \text{ is flavored}, \tag{2}$$

with the parameter α usually taken to be 1 or 2.[4] This has the consequence that the d_{ij} for the clustering of a soft

[3]The jets may be massive, and as a result, pseudorapidity is not advised [28].

[4]The $\alpha = 1$ variant evokes a longitudinally invariant extension of the classic JADE (squared invariant mass) clustering distance [30,31]. The well-known drawback of the JADE distance, namely, that early in the sequence it can cluster soft pairs going in opposite directions, is precisely the behavior needed to resolve the classic jet-flavor IRC safety issue of Fig. 1.

flavored particle with a significantly harder particle is much larger than the d_{ij} for two similarly soft particles. As a result, the soft particles cluster first, resolving the original IRC safety issue of Fig. 1. Note that flavor k_t also uses a modified d_{iB} distance for flavored particles. The details are best obtained from the original article [1], however, the essence of the modified beam distance is that one uses the same kind of construction as in Eq. (2),

$$d_{iB}^{\text{flav-}k_t} = [\max(p_{ti}, p_{tB}(y_i))]^{\alpha}[\min(p_{ti}, p_{tB}(y_i))]^{2-\alpha}, \quad (3)$$

with $p_{tB}(y)$ a rapidity-dependent hardness scale. In the central region, $p_{tB}(y)$ is of the same order as the overall event hardness.

Relative to the standard k_t algorithm, the flavor-k_t algorithm can significantly alter the kinematics of the clustering of hard flavored jets. For example, in the presence of a hard $b\bar{b}$ pair, the flavor-k_t algorithm can cluster them even when $\Delta R_{b\bar{b}} > R$, as observed, e.g., in Ref. [32] (see also the discussion in Sec. V A).

2. Flavor anti-k_t (CMP)

The algorithm of Ref. [10], there called "flavor anti-k_t," will be referred to here as CMP, to avoid ambiguity with other flavor anti-k_t algorithms. As in the flavor-k_t algorithm, it is to be used with net-flavor or modulo-2 flavor recombination. It modifies the anti-k_t ($p = -1$) d_{ij} distance when i and j are oppositely flavored,

$$d_{ij}^{\text{flav-anti-}k_t} = d_{ij}^{\text{anti-}k_t} \times \mathcal{S}_{ij}, \quad \text{if } i \text{ and } j \text{ are oppositely flavored}, \quad (4)$$

where

$$\mathcal{S}_{ij} = 1 - \Theta(1-\kappa)\cos\left(\frac{\pi}{2}\kappa\right), \quad \kappa \equiv \frac{1}{a}\frac{p_{ti}^2 + p_{tj}^2}{2p_{t,\max}^2}, \quad (5)$$

and $p_{t,\max}$ would typically be a hard scale (see Ref. [10] for further details). Throughout this paper, we use $p_{t,\max} \equiv p_{t,\text{global-max}}$, where $p_{t,\text{global-max}}$ is the transverse momentum of the hardest pseudojet across the event at the given stage of the clustering.[5] In addition to the jet radius, the algorithm has one parameter, a, taken in the range 0.01–0.5 in the original publication [10]. Unlike the flavor-k_t algorithm, the CMP algorithm uses a beam distance that is identical to that of the plain anti-k_t algorithm.

The CMP algorithm resolves the problem in Fig. 1 because when particles 1 and 2 are both soft, κ is very small. Specifically, taking dimensions such that $p_{t,\text{global-max}} = 1$, a soft ij quark pair has $\mathcal{S}_{ij} \sim \kappa^2 \sim \max(p_{ti}^4, p_{tj}^4)$, leading to an overall $d_{ij} \sim \max(p_{ti}^2, p_{tj}^2)\Delta R_{ij}^2$. This is much smaller than the anti-k_t clustering distance of a soft quark with a hard parton, which is of order ΔR_{ij}^2. As a result, the soft $q\bar{q}$ pair clusters first, and there is no IRC-safety issue in Fig. 1. Note that when one or other of i and j is hard the use of a small value for the parameter a results in κ being large and thus $\mathcal{S}_{ij} = 1$. As a result, the CMP algorithm behaves like the anti-k_t algorithm for hard particles. For $a \to 0$, the algorithm reduces to anti-k_t. However, for finite a, the algorithm does sometimes yield jets whose kinematics differ from those of the anti-k_t algorithm.

3. Flavor dressing (GHS)

The algorithm of Ref. [11], there called "flavor dressing," will be referred to here as GHS. This algorithm involves three stages: a standard clustering stage in which flavor is not considered; an "accumulation" stage in which flavored particles accumulate momentum from nonflavored ones; and a "dressing" stage, which assigns the flavor to the original anti-k_t jets. Here, we limit ourselves to sketching the main features of each of the steps and refer the reader to the original reference for the full details.

In the first step, the event is clustered with the standard anti-k_t algorithm. In this step, one also applies standard jet cuts, e.g., on transverse momentum and rapidity, to the resulting jets.

In the second step, the algorithm runs an accumulation stage, which follows a version of C/A clustering [i.e., $p = 0$ in Eq. (1a)] with a radius of R_{cut}, with two modifications: (i) clustering of flavored objects with nonflavored ones discards the nonflavored one if the clustering fails to pass a SoftDrop kinematic cut [33],

$$\frac{\min(p_{ti}, p_{tj})}{(p_{ti} + p_{tj})} > z_{\text{cut}}\left(\frac{\Delta R_{ij}}{R_{\text{cut}}}\right)^{\beta}, \quad (6)$$

where z_{cut} and β are the usual SoftDrop parameters; and (ii) When two flavored objects would normally cluster, they are instead both removed from the accumulation clustering process, and each is treated as a "flavor cluster," to be used as an input to the third step of the algorithm. Any flavored clusters that remain at the end of the modified C/A clustering also serve as inputs to the third step.

The third step is the flavor dressing itself. It evaluates flavor-k_t distances (a) between pairs of flavor clusters ($d_{\hat{f}_i\hat{f}_j}$); (b) between each flavor cluster and the anti-k_t jet, j_k, to which the flavored particle in the cluster belonged ($d_{\hat{f}_i j_k}$); and (c) with the beam ($d_{\hat{f}_i B_\pm}$). When the smallest distance is a $d_{\hat{f}_i\hat{f}_j}$, the flavors annihilate, and $\hat{f}_i$ and $\hat{f}_j$ are removed from further consideration; when it is a $d_{\hat{f}_i j_k}$, the flavor of i is assigned to jet j_k, and $\hat{f}_i$ is removed from further consideration; and when it is a $d_{\hat{f}_i B_\pm}$, $\hat{f}_i$ is simply discarded. Distance measures involving any flavor clusters

[5] We are grateful to the authors of Ref. [10] for discussions on this point.

$\hat{f}_i$ or $\hat{f}_j$ that were annihilated, assigned, or discarded are then removed from the list, and the procedure repeats until no flavored clusters remain. Besides the standard jet radius, the algorithm has four parameters: R_{cut}, associated with the C/A clustering; β and z_{cut} for the SoftDrop condition; and the α of the flavor-k_t distances.

In the configuration of Fig. 1, we would have three flavor clusters (1, 2, 3), with 2 and 3 associated with a hard jet. The third step of the algorithm would annihilate the $\bar{q}$ and q flavors of 1 and 2, because they have the smallest flavor-k_t distance, and attribute the flavor of 3 to the hard jet. The flavor dressing algorithm never modifies the kinematics of the original anti-k_t jets, only their flavor. Note that for events where every anti-k_t jet consists of a single particle, i.e., events where there has been no kinematic recombination, the flavor of each jet is the same as for the anti-k_t algorithm. This is a property that we will seek also in our IFN algorithm.

4. Multiflavored events

A final comment concerns clustering of events with more than one flavor (e.g., tracking both b and c flavor). The flavor-k_t algorithm is to be run with all flavors for which one wants information in the final jets. The CMP and GHS algorithms are designed for a single flavor at a time (e.g., just the b and $\bar{b}$ flavor in the event). However, we note that for the CMP algorithm, it is straightforward to identify potential ways of extending it, for example, by using the distance in Eq. (4) whenever a pair has the potential for at least some cancellation of flavor. As concerns GHS, since it does not modify the anti-k_t jets' kinematics, one can simply rerun it again for a second flavor, and so forth.

III. ANTI-k_t AND C/A JETS WITH INTERLEAVED FLAVOR NEUTRALIZATION

In this section, we present the motivation for, and description of, our new flavor neutralization algorithm.

A. Design aims and core concept

If we consider what is needed for broad usage of a jet flavor algorithm, we can identify at least four criteria that are necessary, or at least highly desirable:

(i) *IRC safety.* Both the kinematics and the flavors of any hard jets should be IRC safe.

(ii) *Preserved kinematics*. For a given member of the generalized-k_t algorithm family, the flavor algorithm should not modify the jets' kinematics.

(iii) *Multiscale flavor resolution*. The flavors of the pseudojets should be well defined at any step of the clustering, so as to leave open the possibility of using flavor information with the full cluster sequence, e.g., for jet substructure studies.

Additionally, as mentioned at the end of Sec. II C 3, it can be beneficial to have the following property:

(iv) *Single parton consistency.* For events in which each jet contains exactly one parton, the flavor algorithm should not modify the jets' flavors relative to the simple generalized-k_t algorithm. This ensures that typical leading-order calculations will give the same results for the generalized-k_t algorithm and its flavored extension. It notably means that one cannot *a priori* decide to neglect some subset of flavor in an event without specifying the jet kinematics.

To achieve these aims, the core novel idea that we introduce here is that of maintaining the standard clustering procedure but modifying the flavor-related aspects of the recombination scheme at each step of kinematic recombination. In particular, our approach uses a global, event-wide treatment of flavor at each pairwise clustering step. By construction, the resulting jets will have identical kinematics as compared to the original jet algorithm, and we aim to arrange for the flavor labels associated with the jets to be IRC safe at any stage in the clustering sequence.

There is quite some freedom in such an approach. The fundamental principle of flavor neutralization, which we believe can be applied in a variety of ways, is illustrated in Fig. 2. When a pseudojet with nonzero flavor is about to undergo a kinematic clustering [soft q (2), clustering with hard $\bar{q}$ (3) in Fig. 2(a)], the algorithm needs to establish whether the flavors of 2 and 3 should be combined as per usual net-flavor summation or instead whether the flavor of either of the particles should be "neutralized" by some other particle(s) in the event before allowing the kinematic $(2+3)$ clustering to proceed. For example, in Fig. 2(b), with a soft $\bar{q}$ (particle 1) in the vicinity of the soft q (2), the algorithm may decide to first neutralize the flavors of particles 1 and 2, before moving ahead with the $2+3$ clustering. If that neutralization happens, then particles 1 and 2 become flavorless, as illustrated by the black dashed lines in Fig. 2(c). This is then followed by the kinematic clustering in Fig. 2(d), resulting in a $2+3$ jet that retains the $\bar{q}$ flavor of hard particle 3, as needed for IRC safety.

In general, the IRC safety (or otherwise) of the algorithm resides in the criteria used to decide whether to neutralize a given pseudojet's flavor and, if so, then with which other pseudojet(s). As with earlier flavored clustering algorithms, such a procedure will need to rely on some measure of the likelihood that a given flavored pair came from an effective parent gluon's splitting, versus the flavor originating from a genuine hard parton.

B. Introducing the IFN algorithm

We now construct a concrete algorithm based on Fig. 2 that integrates jet clustering with flavor neutralization: interleaved flavor neutralization. The core of our algorithm is the search for neutralization candidates at any given stage of the clustering. Among the ingredients of that search is a measure of flavor neutralization distance u_{ij} between any pair of particles i and j, the softer of which will always be

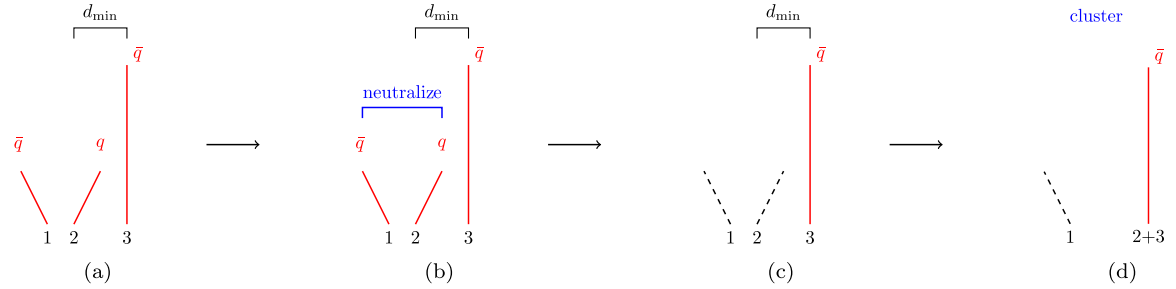

FIG. 2. Illustration of the flavor-neutralization approach. The event displayed here (a) has the property that there is a soft $\bar{q}q$ pair (particles 1 and 2) and a hard $\bar{q}$ (particle 3) with $p_{t1} \sim p_{t2} \ll p_{t3}$. Additionally, we have all ΔR distances of order one, but with the constraint that $\Delta R_{23} < R$, while $\Delta R_{12} > R$, so that within the anti-k_t algorithm 2 and 3 cluster into one jet, while 1 would form a separate soft jet. In (b), just before the $2+3$ clustering, the flavor of 1 is used to neutralize the flavor of 2, which results in the intermediate stage shown in (c), where particles 1 and 2 have lost their flavor (as represented by the black dashed lines). Finally, in (d), the (now) flavorless pseudojet 2 is clustered with 3 into a pseudojet $2+3$ with the $\bar{q}$ flavor of just particle 3.

flavored. For now, the reader may wish to think of u_{ij} as being a flavor-k_t type distance, cf. Eq. (2), though there are important further subtleties, discussed below in Sec. III C.

In defining the algorithm in the next few paragraphs, we shall frequently make reference to Fig. 2 to illustrate the function of the different steps, keeping in mind that the flavor of the final hard jet (made of particles 2 and 3) should ultimately just be that of the hard particle $3_{\bar{q}}$ without contamination from the flavors of the soft $1_{\bar{q}}2_{q}$ pair.

We write the core neutralization search part of the algorithm in the style of a computer subroutine $N(i, u_{\max}, C, E)$, taking a number of arguments as inputs, specifically:

(i) the index i of the pseudojet for which to identify potential neutralization partner(s) [e.g., $i = 2$ in Fig. 2(a)],
(ii) a threshold $u_{\max}$ above which to ignore neutralization candidates [e.g., in the context of the $2+3$ kinematic clustering in Fig. 2(a), this would be $u_{\max} = u_{23}$],
(iii) a list C of all potential neutralization candidates, i.e., all currently flavored pseudojets in the event [$C = \{1, 2, 3\}$ in Fig. 2(a)], and
(iv) a subset E among those flavored pseudojets to be excluded in the neutralization search because they have already been considered in some prior step of the algorithm [$E = \{2, 3\}$ in Fig. 2(a), because particles 2 and 3 have already been considered in that they set $u_{\max} = u_{23}$].

The $N(i, u_{\max}, C, E)$ algorithm is formulated as follows:

N1. Create a list L of u_{ik} distances for all k among the candidates C that satisfy $u_{ik} < u_{\max}$, excluding those in the exclusion set E.
N2. Identify the k that corresponds to the smallest u_{ik} in the list.
N3. If k contains no flavor that can neutralize flavor in i (e.g., k is a b-quark and i is a c-quark), remove the corresponding u_{ik} from list L, and loop back to step N2.
N4. Before using k to neutralize flavor in i, check to see whether there are other pseudojets that could more naturally be paired with k in order to neutralize k's flavor. Do so through a recursive use of flavor neutralization, searching for neutralization partners of k by running $N(k, u_{ik}, C, E \cup \{k\})$. Section III D explains the importance of recursion for IRC safety.
N5. For each flavor currently in i, neutralize as much of that flavor as one can with any flavor that is still present in k.[6] For example, if i has flavor $c\bar{b}$ and k has flavor bb, use k to cancel the $\bar{b}$ flavor, so that the updated i has flavor c and the updated k has flavor b.
N6. If i is now flavorless, exit.
N7. Otherwise, remove the current u_{ik} from list L. If any entries are still left in list L, loop back to step N2. Otherwise, exit.

In our IFN formulation, the flavor neutralization search is triggered whenever a clustering is about to occur for which the softer pseudojet is flavored, specifically:

I1. When pseudojets i and j recombine in the standard kinematic clustering sequence, let i be the pseudojet with lower p_t. If i is flavorless, then, $i+j$ simply takes the flavor of j, and one moves on to the next kinematic jet clustering step.
I2. Otherwise, identify all pseudojets that currently carry flavor, including any flavored jets declared earlier according to a d_{iB} step, and put them into a list C of potential neutralization candidates. Initialize the set $E = \{i, j\}$ of particles to be excluded from the search for neutralization candidates.
I3. Call the flavor-neutralization search, $N(i, u_{ij}, C, E)$, which may use one or more flavored particles in set C to neutralize some or all of the flavor contained in i.

[6]If working with flavor modulo-2, then initial flavors are always to be understood as being modulo-2, and each comparison and/or combination is also to be performed in a modulo-2 sense.

I4. For any remaining flavor in i, apply the standard net-flavor (or flavor modulo-2) summation of i with j and move on to the next kinematic jet clustering step.

Interleaving flavor neutralization at each step of the clustering is important from the point of view of collinear safety. To illustrate this, it is helpful to suppose that particles i, j, and k all have comparable transverse momenta and interparticle distances $\Delta R \sim R$. In this situation, $u_{ij} \sim u_{ik}$. Consider the case where j undergoes a collinear splitting, $j \to j_a, j_b$ with $\Delta R_{j_a,j_b} \ll R$. If one ran flavor neutralization without clustering, one could find oneself in a situation where $u_{ik} < u_{ij}$, but $u_{ik} > u_{ij_a}$, thus changing the neutralization sequence.

Now, let us examine how this changes if neutralization is interleaved with clustering. The clustering algorithms that we consider are the anti-k_t and C/A algorithms. They both have the property that when all particles have similar transverse momenta clustering of the collinear j_a, j_b pair will precede the ij clustering step. At the j_a, j_b clustering, if the neutralization search gets triggered, then j_a and j_b will cluster with normal net-flavor recombination, since $u_{j_a j_b}$ is much smaller than all other u's. When the clustering reaches the ij step, all distances will see the kinematics of j, rather than that of the underlying j_a and j_b, thus ensuring that the algorithm is collinear safe.[7]

C. Choice of neutralization distance

Let us now turn to the u_{ik} flavor neutralization distance between a pair of particles i and k. Recall that the softer of the two will always be flavored, while the harder one may or may not be.

We write the u_{ik} distance generically with two parameters, α and ω,

$$u_{ik} \equiv [\max(p_{ti}, p_{tk})]^{\alpha} [\min(p_{ti}, p_{tk})]^{2-\alpha} \times \Omega_{ik}^2, \tag{7a}$$

$$\Omega_{ik}^2 \equiv 2\left[\frac{1}{\omega^2}(\cosh(\omega \Delta y_{ik}) - 1) - (\cos \Delta\phi_{ik} - 1)\right], \tag{7b}$$

where $\Delta y_{ik} = y_i - y_k$ and analogously for $\Delta\phi_{ik}$. Let us start with the part related to the transverse momenta. This is identical to that used in the flavor-k_t algorithm, cf. Eq. (2), with the same parameter α. As in typical flavor-k_t studies, we assume $0 < \alpha \leq 2$ and in particular concentrate on $\alpha = 1$ and $\alpha = 2$.

[7]When considering collinear splitting in events with a hierarchy of energies, the different members of the generalized-k_t family may perform the soft and the collinear clusterings in different orders. However, when the neutralization search is, say, comparing neutralization distances involving two soft particles i and k and a hard particle j ($u_{ik} \ll u_{ij}, u_{kj}$), a collinear splitting of any of the soft or hard particles will only modify the u's by a factor of order 1, and it will leave the hierarchies untouched, and correspondingly also the resulting neutralization pattern.

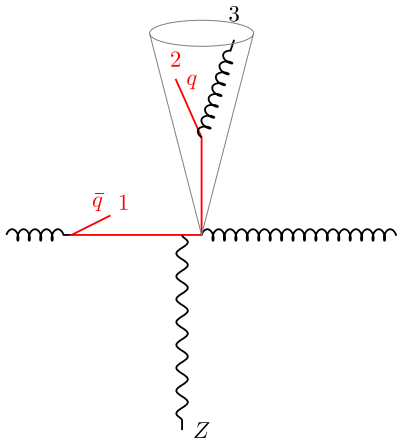

FIG. 3. NNLO contribution to the $pp \to Z +$ jet process, that helps illustrate the origin of the condition, Eq. (9), on the ω parameter in the angular part of the u_{ij} distance, Eq. (7a). It involves a hard jet with a final-state splitting (where the jet constituents, q and a gluon, are labeled 2 and 3, respectively), as well as an initial-state collinear splitting ($g \to q\bar{q}$, with the $\bar{q}$ labeled 1). When $\alpha + \omega < 2$, the initial-state collinear $\bar{q}$ (1) neutralizes the flavor of the q (2).

Next, we examine the angular part of the distance, Ω_{ik}^2, which involves a parameter ω. For any ω of order 1, in the limit of small Δy_{ik} and small $\Delta\phi_{ik}$, Ω_{ik}^2 reduces to the standard $\Delta R_{ik}^2 = \Delta y_{ik}^2 + \Delta\phi_{ik}^2$. The reason for using Ω_{ik}^2 rather than the standard ΔR^2 is to ensure IRC safety as concerns the interplay between collinear initial-state splittings and splittings elsewhere in the event. This is best explained with the help of Fig. 3. In the anti-k_t and C/A algorithms, particles 2 and 3 will cluster first.[8] When $p_{t2} < p_{t3}$, the $2+3$ clustering triggers a flavor neutralization search. The only candidate for flavor neutralization is particle 1, and one should compare the u_{12} and u_{23} distances. We will suppose that particles 2 and 3 have similar p_t's and are at central rapidity. The initial-state collinear splitting that creates particle 1 typically results in $y_1 = \ln p_{t3}/p_{t1} + \mathcal{O}(1)$. Neglecting $\mathcal{O}(1)$ factors, we then have

$$u_{12} \sim p_{t2}^{\alpha} p_{t1}^{2-\alpha} \left(\frac{p_{t3}}{p_{t1}}\right)^{\omega} \sim p_{t3}^{(\alpha+\omega)} p_{t1}^{(2-\alpha-\omega)}, \tag{8a}$$

$$u_{23} \sim p_{t3}^{\alpha} p_{t2}^{2-\alpha} \Delta R_{23}^2 \sim p_{t3}^2 \Delta R_{23}^2, \tag{8b}$$

where in the rightmost part of each equation we have exploited $p_{t2} \sim p_{t3}$. One immediately observes that if $\alpha + \omega < 2$ then in the initial-state collinear limit, where $p_{t1} \ll p_{t3}$, one has $u_{12} \ll u_{23}$. This causes particle 1 to

[8]This would not be the case for the k_t algorithm, and an investigation of the interplay of k_t clustering with IFN is left to future work.

neutralize the flavor of particle 2, even when 1 is arbitrarily collinear, resulting in a flavorless hard jet. In contrast, when the initial-state splitting is absent, the hard jet will be flavored. Thus, the algorithm would be unsafe with respect to initial-state collinear splittings. On the other hand, if we take

$$\alpha + \omega > 2, \tag{9}$$

then u_{12} will always be parametrically larger than u_{23} in the limit $p_{t1} \to 0$, thus effectively forbidding neutralization of 1 and 2; see Appendix B 1 for further discussion.[9] In practice, we will nearly always take

$$\text{default: } \omega = 3 - \alpha, \tag{10}$$

and where not explicitly stated in plots, this will be the choice that we adopt.

IRC-safety subtleties connected with the large Δy_{ij} behavior of normal ΔR_{ij}^2 distances are relevant for all flavor algorithms, though sometimes the issues appear only at orders beyond α_s^2. Further discussion of this point is provided in Appendixes B 1, C 1, and C 3. Note also that the original formulation of the k_t algorithm for hadron colliders [4] foresaw the possibility of an angular distance Ω_{ik}^2 with $\omega = 1$, though this does not have IRC safety implications for the kinematic aspects of normal jet clustering.

D. Need for recursion

A key element of IFN is the recursion in step N4 above. The need for recursion can be illustrated with the help of Fig. 4. Again considering the anti-k_t or the C/A clustering algorithms, the first clustering step is that of particles 3 and 4. If $p_{t3} < p_{t4}$, their clustering triggers a flavor neutralization search. That search will identify particle 2 (from a large-angle soft pair) as a potential neutralization candidate. With $\alpha = 1$, we will have $u_{23} \ll u_{34}$, while for $\alpha = 2$, $u_{23} \sim u_{34}$. Either way, without recursion, it would be possible for 2 to neutralize the flavor of 3, which would ultimately result in the hard jet being flavorless. In the absence of the (1, 2) pair, the hard jet would be flavored. This would induce an infrared divergence.

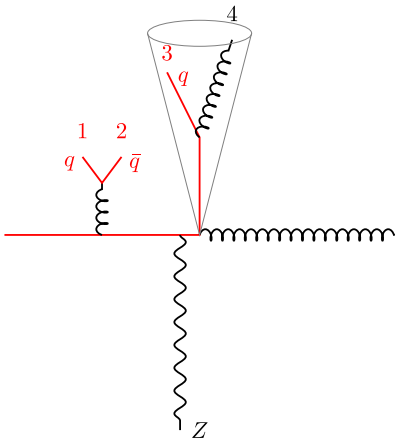

FIG. 4. N^{3}LO contribution to the Z + jet process that helps illustrate the need for recursion in step N4 of the flavor neutralization search. It involves a hard jet with a noncollinear splitting (flavored 3 and flavorless 4) and a flavored initial-state double-soft pair (labeled 1 and 2). Without recursion, particle 2 can end up neutralizing the flavor of 3.

The recursive aspect of the algorithm resolves this problem as follows: when 2 is identified as a neutralization candidate for 3, the recursive search that is triggered in step N4 identifies particle 1 as a neutralization candidate for 2. For both $\alpha = 1$ and $\alpha = 2$, we have $u_{12} \ll u_{23}$, and so particles 1 and 2 will neutralize. When the algorithm exits the recursion step, there are no longer any remaining flavored particles to neutralize the flavor of particle 3. Thus, the hard jet will retain its net flavor, resolving this IRC safety issue (see, e.g., Appendix B 2).

E. Further comments

We conclude this section with a few general comments about the IFN algorithm.

A first comment concerns single parton consistency, as discussed in Sec. III. A potentially useful characteristic of IFN, shared with GHS, is that for configurations where each jet contains no more than one particle the flavors of those jets are identical to those in standard anti-k_t. This is trivial because for such configurations there is never a situation where two particles would cluster together and so the flavor neutralization part of the algorithm is never triggered. Thus, any leading-order jet calculation, for an arbitrary number of final-state jets, will give identical jets and flavors for those jets in the anti-k_t and its IFN extension.

A second comment concerns the fact that, unlike the flavor-k_t algorithm, the flavor-related part of our IFN algorithm has no specific treatment of beam distances for flavored particles (the CMP algorithm has similarities in that it leaves the anti-k_t beam distance untouched for flavored particles). This means that particular care is needed around the potential for long-distance clusterings,

[9]We have also explored the border case of $\alpha + \omega = 2$ and find that it diverges. This is relevant in particular to the case of $\alpha = 1$ and $\omega = 1$, for which u_{ik} coincides with the ik squared invariant mass when i and k are massless, i.e., a JADE-like distance [30,31]. An issue to be aware of with an invariant-mass distance in a hadron collider context is that the invariant mass between an energetic initial-state collinear emission and a hard final-state particle is commensurate with that between two well-separated hard final-state particles. Furthermore, a potential solution to this issue, i.e., clustering initial-state collinear emissions early, via their small invariant mass with the beam, involves ambiguities in the identification of the beam energy.

as discussed in Sec. III C. Nevertheless, even algorithms with beam distances can suffer from long distance clustering when using standard ΔR_{ij}^2-type angular measures, as discussed in Appendix C 1.

A third comment concerns events with more than one flavor, e.g., both c and b flavor. One possibility is to consider all flavors within a single IFN run. Suppose i has flavor b and is about to cluster with j. This triggers a search for candidates to neutralize i's flavor. The search may find a particle k with flavor $c\bar{b}$ (which could have arisen, for example, through earlier clusterings). The recursion of the IFN algorithm may then identify some other particle with flavor $\bar{c}$, which neutralizes the c component of k's flavor. Thus, c flavor elsewhere in the event is affected by the b flavor in the $i+j$ clustering. Alternatively, one could choose to run the IFN algorithm first for the b flavor, then for the c flavor. In that case, the flavor neutralization steps for b flavor have no side effects on those for c flavor. Consequently, the output of the algorithm can be different according to whether one runs it for all flavors at once or separately a flavor at a time. In those of our studies below that include multiple flavors (the IRC safety tests of Sec. IV for the IFN algorithms and the phenomenological study of Sec. V C), we treat all flavors at once.

In the discussion so far, we have always described the IFN algorithm as happening at the same time as the kinematic clustering. However, because IFN preserves the kinematic clustering sequence, the neutralization steps can also be run as an add-on. Here, one loops again through each step of the kinematic clustering and updates the flavor information. This may be more convenient in cases where one already has a jet collection (and associated clustering sequence) defined.

A final comment concerns the "bland" option of flavor k_t [1], which sets to infinity any clustering distances that would lead to flavors that are inconsistent with a single partonic flavor (e.g., bb or $\bar{c}b$).[10] One could imagine a similar bland extension for our flavor neutralization distances, but we leave the study of this question to future work.

IV. IRC SAFETY: DISCUSSION AND TESTS

Given the considerable subtlety of IRC safety for jet flavor, it is important to design tests to help build confidence in the IRC safety of any new algorithm. Subtle IRC-safety problems have arisen in the past in the context of cone-type jet algorithms, which ultimately led to the construction of an automated testing framework, used to verify the IRC safety of the SISCone algorithm [27]. Here, we adapt and substantially extend that framework. The framework is available on request from the authors.

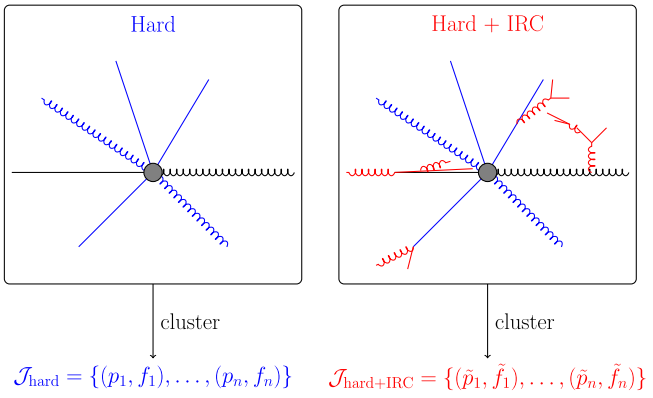

FIG. 5. On the left, hard particles (in blue) are generated, some with flavor, at central rapidities. The event is clustered with a given jet algorithm, resulting in a set of "hard" jets $\mathcal{J}_{\text{hard}}$, with kinematics $\{p_i\}$ and associated flavors $\{f_i\}$. On the right, additional IRC radiation is added to the event as explained in the main text. This modified event is then clustered with the same jet algorithm, and the resulting set of "hard + IRC" jets $\mathcal{J}_{\text{hard+IRC}}$ is compared against the original set of hard jets (and similarly for each hard step of the underlying clustering sequence). The sets agree if both the kinematics and the flavors of the various jets (and hard clustering steps) are identical. In the limit where the extra radiation becomes soft and collinear, the rate of failed events (where $\mathcal{J}_{\text{hard}} \neq \mathcal{J}_{\text{hard+IRC}}$) should go to zero for an algorithm that is IRC safe. The right-hand figure also serves to illustrate some of the classes of IRC additions that we make, though in practice, we do not go beyond sixth order in α_s; i.e., we do not simultaneously add as many emissions as are shown.

A. Methodology

Our approach is illustrated in Fig. 5, which goes beyond the tests performed in the more recent literature. We begin by generating a random hard event, with some number of particles (flavored or not), and run the clustering with the jet definition that we wish to test. This results in a set of hard jets, $\mathcal{J}_{\text{hard}} = \{(p_1, f_1), \ldots, (p_n, f_n)\}$ with kinematics $\{p_1, \ldots, p_n\}$ and associated flavors $\{f_1, \ldots, f_n\}$. Note that here we do not force the total 4-momentum (or even transverse momentum) of the hard event to be balanced; i.e., it is as if the events have neutrinos, leptons, or isolated photons that would balance the momentum but do not take part in the clustering. We then construct a modified hard + IRC event, where we add soft emissions and collinear splittings up to some given order in α_s. We cluster that modified event and verify whether the hard jets in the modified event, $\mathcal{J}_{\text{hard+IRC}}$ coincide with the hard jets in the original event, both in terms of kinematics and flavor.[11]

[10] This approach was adopted also in the IRC-unsafe "QCD-aware" clustering algorithm [29], without any clustering distance modification.

[11] In the modified event, we also identify each step in the clustering sequence that involves clustering of two hard particles, and compare its kinematics and flavor to that of the corresponding step for the unmodified event.

We then examine the rate of failure as a function of the logarithmic momentum range (L) of IRC additions. For an IRC-safe algorithm, we expect that failure rate to vanish as a (possibly fractional) power of the momentum scale of the IRC additions.

Ideally, we would consider all possible IRC insertions. There are two logarithms per order in α_s, and we have found that it is important to explore configurations at least up to α_s^4. The smallest non-IRC-safe contribution would be a term independent of L, and at α_s^4, that would imply identifying one event in L^8 that fails. We will return to the question of the meaning and range of L below, but for now, let us consider $L = 30$. That would imply identifying failures at the level of one event in $30^8 \simeq 6.6 \times 10^{11}$, which is prohibitive. Note, however, that the only contributions that give the maximum number of logarithms are those that exclusively involve the emission of simultaneously soft and collinear gluons, which are not the most likely configuration for triggering flavor-related IRC safety issues.

Consequently, we take a more targeted approach, in which we allow up to one logarithm per order in α_s, prioritizing configurations that are potentially nontrivial from the point of view of flavor. We do so by omitting single soft-gluon divergences unless they involve a subsequent splitting to a pair of commensurate-angle partons.

B. Classes of IRC emissions

The specific IRC emissions included in our testing framework are shown in Fig. 6 and described below:

(i) *Final-state hard-collinear (FHC) emission*. We perform a (hard) collinear splitting of a randomly chosen final-state particle. We uniformly sample the logarithm of the transverse momentum of the splitting. We also uniformly sample the longitudinal momentum fraction of the splitting. This is consistent with our choice not to include the soft gluon emission divergence as part of the FHC class. For all flavor combinations ($q \to qg$, $g \to q\bar{q}$, $g \to gg$), an FHC branching is associated with one power of α_s and one logarithm of the IRC scale. For readers in the habit of using a Lund-diagram [34] representation of soft-collinear phase space, this corresponds to a strip close to the hard-collinear boundary in the Lund leaf of the emitter [Fig. 6(a); that figure shows a shaded logarithmic transverse momentum range, which we further discuss below]. Note that sensitivity to soft gluon emission will still be present in the analysis but will be obtained through the double-soft (FDS/IDS) contributions below.

(ii) *Initial-state hard-collinear (IHC) emission*. We perform a hard-collinear splitting of the beam [Fig. 6(b)]. Again, we sample the longitudinal momentum fraction of the splitting uniformly, bringing one order of α_s and one logarithm.

(iii) *Final-state double-soft (FDS) pair, i.e., the addition of a $g \to gg$ or $g \to q\bar{q}$ pair.* We choose one emitter randomly among the final-state hard particles and uniformly sample the logarithm of the transverse momentum $\ln p_{t,g}$ and the rapidity y_g of the intermediate gluon with respect to the emitter. This corresponds to uniform sampling of the bulk of the Lund leaf for that emitter [Fig. 6(c)] and brings one power of the coupling and two powers of the logarithm. We allow the intermediate gluon to split and distribute the kinematics of the resulting pair in such a way as to correctly reproduce the asymptotic behavior of the $q\bar{q}$ double-soft matrix element in those kinematic regions where the splitting is asymmetric (either substantially different momentum fractions or rapidities, as elaborated upon in Appendix A). Note that, even for $g \to gg$ splittings, we use the asymptotic matrix element for $g \to q\bar{q}$.[12] With the branching to a pair, we gain an extra power of the coupling and no logarithms, giving in total two powers of the coupling and two powers of the logarithm. We do not include the collinear divergence when $\theta_{q\bar{q}} \ll \theta_{hq} \sim \theta_{h\bar{q}}$, where h is the hard particle. That would bring three powers of logarithm for two powers of the coupling. That said, we still have configurations with $\theta_{q\bar{q}} \ll \theta_{hq}$, but those are generated by a different sequence, namely, FDS production of a pair of soft gluons, followed by nested FHC branching of one of those gluons to a $q\bar{q}$ pair.

(iv) *Initial-state double-soft (IDS) pair.* These are generated similarly to FDS, but with respect to the forward or backward beam [Fig. 6(d)]. Note that both the IDS and FDS mechanisms include a subset of phase space where the double-soft pair is not collinear but instead at large angles, i.e., the configuration of Fig. 1. The matrix element for large-angle double-soft production is the same as collinear double-soft production (up to complications of color factors), a consequence of longitudinal boost invariance of soft production in the (leading-N_C) color dipole rest frame. As a result, the IDS component in particular is guaranteed to fully cover the soft large-angle double-soft phase space.

[12]This might seem surprising at first sight, since the production of a soft gg pair has a qualitatively different structure from that of a soft $q\bar{q}$ pair when the pair is well separated. For example, for emission of a double soft gluon pair from a quark line, the component with the C_F^2 color factor corresponds to independent emission, with logarithmic divergences both in the ratio of the gluon transverse momenta and the rapidity separation. However, these contributions would be associated with the double (soft-collinear) logarithms that we are deliberately leaving out. Similarly, the soft singularity in the $C_F C_A$ term would also bring an extra soft-gluon logarithm that is beyond what we aim to sample.

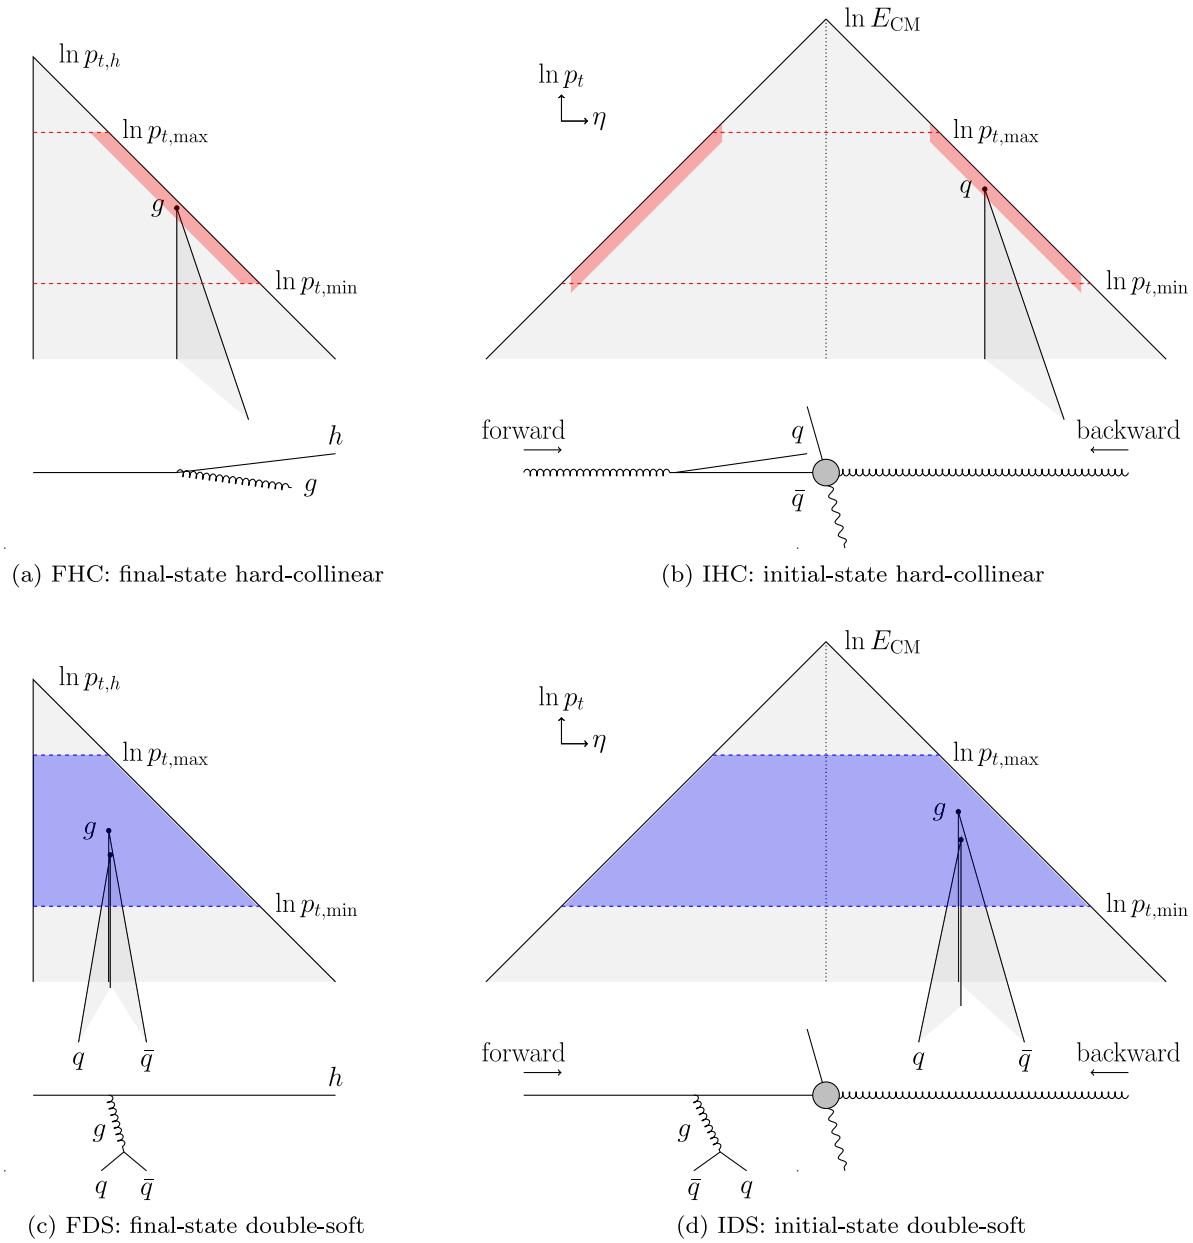

(a) FHC: final-state hard-collinear

(b) IHC: initial-state hard-collinear

(c) FDS: final-state double-soft

(d) IDS: initial-state double-soft

FIG. 6. Illustration of the emissions sampled in the hard + IRC event on a Lund diagram [34]. The left column shows emissions from a final-state jet. The right column shows initial-state radiation from the beams, where an emission collinear to the forward beam (coming from the left on the diagram) ends up at positive rapidities (right-hand half of the Lund plane) and vice-versa for an emission collinear to the backward beam. The top row shows hard-collinear splittings from [(a) FHC] hard final-state particles or [(b) IHC] the beams. The bottom row shows double-soft pairs, flavored or not, being emitted from [(c) FDS] hard final-state particles or [(d) IDS] the beams. In a bin defined by $\ln p_{t,\min} < \ln p_t < \ln p_{t,\max}$, we sample any additional radiation in slices in the Lund planes (both for initial- and final-state radiation). We typically choose $\ln p_{t,\min} = 3 \ln p_{t,\max}$, with $\ln p_{t,\max} < 0$.

For each emission, we need to choose the range in $\ln p_t$, which we define as $\ln p_{t,\min} < \ln p_t < \ln p_{t,\max}$ (we choose our dimensions such that $\ln p_{t,\max}$ and $\ln p_{t,\min}$ are always negative). One potential difficulty is that of proximity (or overlap) between the momentum scale of the IRC additions and the momentum scales of the hard event because when there is proximity or overlap the IRC additions can legitimately modify the hard event. If one keeps the range of hard scales fixed and takes IRC scales to zero, one expects the rate of such legitimate modifications to vanish. There should also be a suitably generous factor between the upper ($\ln p_{t,\max}$) and lower ($\ln p_{t,\min}$) edges of the IRC additions' $\ln p_t$ range, so as to be sensitive to IRC-unsafety mechanisms that work across a hierarchy of scales.

TABLE I. Summary of the IRC safety test results. Red crosses (×) indicate a clear failure of IRC safety. Checkmarks (✓) signify that the algorithm passes numerical tests at that order or for that configuration. The tilde (∼) for flavor k_t (and by extension GHS, which uses flavor-k_t distances) indicates marginal convergence, though one expects divergent behavior at higher orders. For algorithms that fail or are marginal at a given order, we display long dashes at higher orders, since those higher orders are also bound to fail. In a few cases, we have identified a new class of problem that only arises at higher order, and we explicitly mark these with a red cross. The GHS parameters here are set to $\alpha = 2$, $\beta = 2$. The IFN procedure is tested both for the anti-k_t and C/A algorithms, and the IFN parameters are chosen as $\alpha \in \{1, 2\}$ with $\omega = 3 - \alpha$ (tests are successful for both sets of parameters). Detailed discussions of the issues identified are linked to from the relevant table cells. Plots in support of the IRC safety conclusion for the IFN combinations are to be found in Appendix D, specifically Figs. 24 and 25, as are plots (Figs. 26 and 27) supporting the IRC safety of our modified versions of the flavor-k_t and CMP algorithms, respectively, flavor-$k_{t,\Omega}$ and CMP_Ω, which are discussed in the text. (They are not shown in the table because we have run them with lower statistics.).

Order relative to born		Anti-k_t	Flavor k_t ($\alpha = 2$)	CMP	$\text{GHS}_{\alpha,\beta}$ (2, 2)	Anti-k_t + IFN_α	C/A + IFN_α
α_s	FHC	✓	✓	✓	✓	✓	✓
	IHC	✓	✓	✓	✓	✓	✓
α_s^2	FDS	$\times_{\text{II B}}$	✓	✓	✓	✓	✓
	IDS	$\times_{\text{II B}}$	✓	✓	✓	✓	✓
	FHC × IHC	✓	✓	✓	✓	✓	✓
	IHC^2	✓	✓	$\times_{\text{C 2}}$	✓	✓	✓
	FHC^2	✓	✓	✓	$\times_{\text{C 4}}$	✓	✓
α_s^3	IHC × IDS	—	$\sim_{\text{C 1}}$	$\times_{\text{C 3}}$	$\sim_{\text{C 1}}$	✓	✓
	Rest		—	—	—	✓	✓
α_s^4	IDS × FDS	—	—	—	$\times_{\text{C 5}}$	✓	✓
	Rest				—	✓	✓
α_s^5		—	—	—	—	✓	✓
α_s^6		—	—	—	—	✓	✓

C. Implementation details

In our testing framework, the FHC, IHC, FDS, and IDS emissions are applied in a nested way (with some caveats; see below) so as to generate configurations up to order α_s^6. We could have imposed angular ordering between nested emissions but choose not to, so as to ensure that we retain sensitivity to cases where strict angular ordering might miss a potentially divergent part of phase space. We also allow for FHC emissions from the individual q, $\bar{q}$, or g of the FDS and IDS pairs (again with no angular ordering condition).

Nesting is essential in order to test collinear splitting chains. It also partially alleviates worries of missing relevant IRC-unsafe configurations as a result of only considering configurations with one logarithm per power of the coupling. For example, as indicated above, $g \to q\bar{q}$ collinear splitting from a soft-collinear gluon ($\alpha_s^2 L^3$) is not generated directly, but the underlying L^3 divergence does appear in our framework, in the context of FDS (double gluon) emission followed by the nested collinear splitting of one of the soft gluons to $q\bar{q}$. This is not exactly the same configuration as the $\alpha_s^2 L^3$ divergence, since it involves an extra soft gluon. In effect, if there is a problem in an $\alpha_s^2 L^3$ $h \to hg \to hq\bar{q}$ configuration, we are making the assumption that we will still detect it in at least some of the $h \to hgg \to hgq\bar{q}$ configurations. For technical reasons, there are some nestings that we are missing: (a) nesting of one double-soft emission from a prior double-soft emission and (b) insertion of double-soft emissions on more than one of the descendants of a collinear splitting. These limitations may be addressed in future work.

One final potential concern is that by allowing strongly angular-disordered configurations we might mistakenly declare an algorithm to be IRC unsafe. To guard against that risk, when we identify an IRC-unsafe configuration, we further investigate it to establish whether it is genuine. All IRC failure classes that we identified were genuine, as documented in Appendix C.

We close this description with some final technical details:

(i) We generate a random number of hard particles and randomly sample their flavors. The maximum number of hard particles we consider is 8. In the figures targeting specific IRC-unsafe configurations, the hard particles' momenta are chosen randomly uniformly between 100 GeV and 1 TeV, and their rapidities uniformly in the range $|y| < 1.5$. In our final IRC safety tests for the IFN algorithms, we use a broader range, with the hard particles' momenta chosen randomly uniformly between 1 GeV and 1 TeV and their rapidities taken uniformly within $|y| < 2.5$. Plots with results will always indicate the ranges used.

(ii) The $\ln p_{t,\text{max}}$ scale is typically scanned in the range -3 down to -42.5 (with p_t expressed in units of GeV). When we show a failure rate, it will always include the $\ln^n p_{t,\text{max}}/p_{t,\text{min}}$ measure that arises from

the n logarithmic integrations at a given order n, but it will not include overall constants such as color factors. We will refer to this as a phase-space weighted failure rate.

(iii) We generally use $\ln p_{t,\min} = 3 \ln p_{t,\max}$. However, in the tests of our IFN algorithms, we have also carried out a subset of tests with a larger ratio $\ln p_{t,\min} = 5 \ln p_{t,\max}$. Results are consistent between the two sets of tests.

(iv) The jet radius is sampled randomly between 0.3 and 1.57.

(v) The jet algorithms are coded as plugins to FastJet [35] version 3.4.1 using techniques that allow for more accurate handling of large rapidities and very small rapidities. Ultimately, however, we found that in order to fully explore the phase space it was also necessary to use higher-precision numerical types from the QD package [36], up to four times normal double precision. This was achieved with the help of a suitably converted version of the FJCORE form of the FastJet package. Many of the techniques that we explored were inspired by, adapted from, and sometimes fed back to the PanScales project [26,37,38].

(vi) In practice, the framework can operate in two modes: it can sample randomly across the available configurations at any given order, which is useful to systematically check whether there are any IRC-unsafe configurations for a given algorithm, or alternatively, it can focus on a specific class of configuration, which is useful when trying to understand the detailed nature of any IRC safety issues that have appeared.

(vii) For flavor jet algorithms where the flavor of the cluster sequence is meaningful (i.e., all algorithms except GHS), we test not just the flavor and kinematics of the final jets but additionally those of all steps in the hard + IRC clustering sequence that correspond to steps in the hard clustering sequence.

(viii) Some algorithms (such as flavor-k_t and IFN) can naturally handle multiple flavors at a time, while others (such as CMP and GHS) are designed around a single flavor at a time. Most of our tests will be carried out with one flavor. For the higher-order IFN tests that go into our summary, Table I, we use six light flavors, so as to ensure that we do not accidentally miss IRC issues that would arise only for multiflavored configurations. Plots will always be labeled with the number of light flavors used.

D. Results

Sample results from our numerical IRC-safety tests are illustrated in Fig. 7. The left-hand plot shows the α_s^2 FDS contribution to the phase-space weighted failure rate for the plain anti-k_t algorithm. The failure rate grows linearly with $\ln p_{t,\max}$, consistent with the expectation of an $\alpha_s^2 L$ divergence.

The middle plot of Fig. 7 shows the phase-space weighted failure rate for the anti-k_t + IFN algorithm (both with $\alpha = 1$ and $\alpha = 2$) at order α_s^2. These results are summed over all sampled configurations (IDS, FDS, IHC^2, FHC^2, and IHC × FHC). One sees that the failure rate vanishes as $p_{t,\max} \to 0$, consistently with a power law, implying no IRC divergence. The right-hand plot shows the results for anti-k_t + IFN at order α_s^4. Again, the plot indicates that the phase-space weighted failure rate vanishes as $p_{t,\max}$ is reduced, consistently with a power law.

For more negative values of $\ln p_{t,\max}$, no points are shown simply because we observed no failures with anti-k_t + IFN. The gray band at the bottom of the plots shows how the test is broken up into different regions with the number of events used for each region (5×10^9 in the lowest region). The regions each involve a different underlying numerical precision type in the code, and one of the limiting factors in our tests is the speed of the code in the lowest region where we are using four times the precision of a standard IEEE double type.[13] Overall, the results provide a strong indication of IRC safety for the anti-k_t + IFN algorithm developed in this paper.

Table I summarizes the results of our testing framework applied to a range of jet algorithms. At lowest order, we organize the results according to the class of divergence being probed, as indicated in the second column of the table, while at higher orders, we limit the breakdown to configurations that have turned out to be of specific interest. The corresponding failure rate plots for the IFN algorithms (with the anti-k_t and C/A algorithms) are given in Appendix D.

Algorithms whose failure rate goes down as the extra radiation becomes softer/more collinear are indicated by a checkmark (✓). Algorithms that develop a divergence (a nonvanishing integrated failure rate as $\ln p_{t,\max} \to -\infty$)

[13] In interpreting these results, one should keep in mind that the integration volume at order α_s^n is effectively $(\ln p_{t,\max} - \ln p_{t,\min})^n$. Defining $\ln p_{t,\max} \equiv L$ and $\ln p_{t,\min} = (1 + c)L$ ($c = 2$ in Fig. 7), that corresponds to $(cL)^n$. Assuming that the failure rate for a L^p divergence goes as $(cL)^p$, then with N events in a given bin, we expect to observe a nonzero failure rate down to $L \sim -N^{1/(n-p)}/c$. Figure 7 was generated with $N \sim 10^9$ events per point and $c = 2$. We therefore expect that up to $n = 4$ we should observe all main classes of failure (i.e., any case with $p \geq 0$) for $L \gtrsim -90$, i.e., over the full range of $L = \ln p_{t,\max}$ in Fig. 7 (keeping in mind that order 1 factors that we have neglected can have a significant impact on the range). For $n = 5$, failures with a $p = 0$ structure would be observed only for $L \gtrsim -30$, while for $n = 6$, this would reduce to $L \gtrsim -15$. Note that it is still useful to explore the full range of L, even at high orders n, because there can be cases where an IRC divergence appears for the first time already with some number of logarithmic enhancements (e.g., the L^3 divergence that appears at α_s^4 in the GHS algorithm, as discussed in Appendix C 5). A final comment concerns the overall normalization of the result of the phase-space weighted failure rate. This can appear large at moderate values of L where there are still failures because it includes the $(cL)^n$ integration volume.

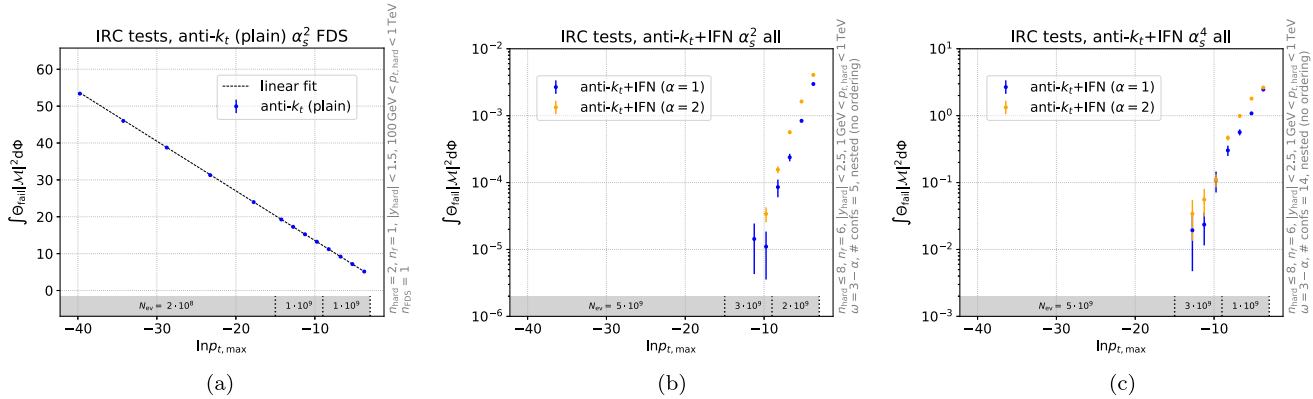

(a) (b) (c)

FIG. 7. Example results from our numerical IRC safety tests showing (a) the IRC unsafety of standard anti-k_t at order α_s^2, and (b, c) the IRC safety of anti-k_t with Interleaved Flavor Neutralization (IFN) at orders α_s^2 and α_s^4, respectively. In all plots, the phase-space weighted failure rate is shown as a function of the maximum hardness of the extra radiation $\ln p_{t,\max}$, where we sample values $3 \ln p_{t,\max} < \ln p_t < \ln p_{t,\max}$ for any bin of $\ln p_{t,\max}$. In (a), the integrated failure rate is plotted on a linear scale for standard anti-k_t in the FDS configuration from Fig. 1 with one double-soft pair and 2 hard particles in the event. The classic anti-k_t IRC safety issue is confirmed numerically by the linear divergence $\alpha_s^2 \ln p_{t,\max}$, from one soft gluon splitting to a flavored pair at large angle. In (b) and (c), the integrated failure rate is plotted on a logarithmic scale for anti-k_t with IFN for parameters $(\alpha = 1, \omega = 2)$ and $(\alpha = 2, \omega = 1)$, for all configurations that contribute at (b) order α_s^2 and (c) order α_s^4, with up to 8 hard particles in the event and up to 6 flavors. The total failure rate goes to zero as $\ln p_{t,\max} \to -\infty$, implying the IRC safety of anti-k_t with IFN to the tested order and accuracy. The absence of points below $\ln p_{t,\max} \simeq -15$ signals no IRC safety failures out of the 5×10^9 events studied for lower $\ln p_{t,\max}$ values (the number of events in each of three regions of $\ln p_{t,\max}$ is indicated in the shaded bands at the bottom of the plot). The plots also include moderately small values of $\ln p_{t,\max}$ so as to better show the overall scaling behavior.

for a given target configuration, are marked by a red cross (×). For each case that has shown a divergence, we have examined a few events where there is a clear failure and developed an analytic understanding of the nature of the problem. We will briefly discuss each issue here, while the table also links to the relevant part of Appendix C with further analytic and numerical studies.

The first two rows of Table I emphasize that at order α_s with just one emission (FHC or IHC) there are no divergences for any jet algorithm with IRC safe kinematics, even without a special treatment of flavor. The classic IRC safety problem of standard anti-k_t only shows up at order α_s^2, as highlighted in the next two rows of Table I for a configuration with one double-soft pair (see Fig. 1). That problem arises in both the FDS and the IDS channels, and in each case, it appears for the subset of events where the FDS or IDS pair is at large angles. From the table, it is clear that all flavor jet algorithms solve that original IRC safety issue.

However, the tests reveal new issues for all algorithms other than our IFN-based procedure. In two cases, flavor-$k_{t,\Omega}$ and CMP_Ω, we will propose modifications that seem to resolve the problem(s). For the interested reader, the summary of the issues is as follows:

(i) *Initial-state* (IHC × IDS) *subtlety at* α_s^3 *for flavor-*k_t *and GHS.* The "∼" for $\alpha = 2$ flavor-k_t and GHS at α_s^3 (IHC × IDS) indicates a borderline case. It arises, for example, for a hard event consisting of a single energetic parton (and resulting hard jet), supplemented with a hard-collinear initial-state splitting and a large-angle double-soft pair, which may be IDS or FDS (see Fig. 15, together with the complete details in Appendix C 1). When one (anti) quark from the double-soft pair is somewhat softer than the other, its d_{ij} distance with the hard-collinear particle can be smaller than that with the other (anti) quark from the soft pair, essentially because the ΔR_{ij}^2 distance goes as Δy_{ij}^2, which is only logarithmically large. The large-angle soft (anti)quark and the initial-state collinear quark cluster, leaving a lone large-angle soft quark, which can contaminate the flavor of the hard jet. At $\mathcal{O}(\alpha_s^3)$, one ends up with an integral that goes as $\int d \ln p_t / (\ln p_t)^2$. This integral converges for $p_t \to 0$; however, the way in which the integrand (multiplying $d \ln p_t$) vanishes as $p_t \to 0$ is not a power-law in p_t. One may thus consider the algorithm to be marginally safe at this order; however, at the next order, one would expect to see additional logarithmic enhancements. These might arise, e.g., from the running of the QCD coupling or evolution of the parton distribution functions (PDF), and would ultimately cause the integral to diverge. Indeed, our study identified a problem in the $\text{IHC}^2 \times$ IDS channel at order α_s^4. However, a conclusive understanding of this configuration requires inclusion also of the virtual and PDF-counterterm contributions, which is beyond the scope of this study. A similar problem arises with $\alpha = 1$, but with extra

logarithms in the denominator of the corresponding integral. This generic class of problem can be solved by replacing $\Delta R_{ij}^2 \to \Omega_{ij}^2$, and, as before, we will use Eq. (10) as our default choice for its ω parameter. We refer to the modified algorithm as flavor-$k_{t,\Omega}$. This simple adaptation is possible because the issue is not with the original underlying strategy but rather with the subtleties that arise in distance measures with QCD initial-state radiation (the same comment holds for related issues in other algorithms). As a consequence, we do not expect to have to make any modifications to the e^+e^- version of the flavor-k_t algorithm.

(ii) *Initial-state* (IHC^2) *issue at* α_s^2 *for CMP.* This issue arises, for example, for a hard (Born) event consisting of a single hard parton, supplemented with two collinear initial-state quark and antiquark emissions, one on each beam (see Fig. 17 and Appendix C 2). Those initial-state emissions cluster in the first step of the algorithm, producing a large-mass, low-p_t flavorless pseudojet at central rapidities, which can then cluster with the hard parton, modifying its kinematics. The problem arises because in the CMP distance Eq. (4), the small factor from the transverse-momenta dominates over the (only logarithmically large) factor from the rapidity separation between the pair. Ultimately, this leads to an $\alpha_s^2 L^2$ divergence. It can be resolved by replacing

$$\mathcal{S}_{ij} \to \bar{\mathcal{S}}_{ij} = \mathcal{S}_{ij} \frac{\Omega_{ij}^2}{\Delta R_{ij}^2} \tag{11}$$

for oppositely flavored pairs and requiring the parameter $\omega > 1$ in the Ω_{ij} distance. In practice, we find that this modification has almost no impact on the phenomenological behavior of the algorithm (e.g., $\lesssim 1\%$ in Fig. 9). We refer to this modified version of the CMP algorithm as CMP$_\Omega$, and unless otherwise specified, we use $\omega = 2$.

(iii) *Initial-state* (IHC × IDS) *issue at* α_s^3 *for CMP.* This issue involves the same configuration and sequence that led to a marginal issue for flavor-k_t (Fig. 15), but here it brings an $\alpha_s^3 L$ divergence (Appendix C 3), rather than α_s^3/L. Recall that the problem here is that the quark from the large-angle double-soft pair can cluster with the IHC ($\bar{q}$) particle, leaving the large-angle soft $\bar{q}$ to contaminate the hard jet flavor. The different IRC behavior of CMP versus flavor-k_t can in part be attributed to the fact that the CMP algorithm retains the standard anti-k_t form of the beam distance, causing the beam clustering of the IHC particle to come at the end of the algorithm. In contrast, in the flavor-k_t algorithm, the IHC beam clustering nearly always comes before the IHC particle can cluster with the soft large-angle quark, reducing the phase-space associated with IRC problems for this configuration. The fix of Eq. (11), i.e., CMP$_\Omega$, also solves this problem.

(iv) *Final-state* (FHC^2) *issue at* α_s^2 *for GHS.* This problem appears with four hard particles such that each of two hard jets contains one flavored (q or $\bar{q}$) and one nonflavored particle, as could arise in semileptonic $t\bar{t}$ decays when considering only b's to be flavored. The α_s^2 modification of the event involves the hard-collinear splitting of one of the energetic quarks, followed by the hard-collinear splitting of that gluon into a $q'\bar{q}'$ pair (see Fig. 20 and Appendix C 4). The accumulation stage leaves the collinear $q'\bar{q}'$ as separate flavor clusters, and relative to the original Born event, the energy of the hard q's cluster is modified. During the dressing stage, the collinear $q'\bar{q}'$ annihilate immediately, but the modification of the energy of the hard-q cluster means that that cluster can behave differently during the dressing stage. We envisage that this problem could be solved by accounting for energies within each jet during the dressing stage but have yet to formulate a concrete modification of the algorithm. This issue is present independently of the parameters of the algorithm and leads to an $\alpha_s^2 L^2$ divergence.

(v) *Mixed initial/final-state* (IDS × FDS) *issue at* α_s^4 *for GHS.* This problem involves a hard event with one Born gluon (h) leading to a single hard jet.[14] The issue arises with a final-state soft-collinear $q_1\bar{q}_2$ pair emitted inside the jet and a large-angle $q_3\bar{q}_4$ pair (see Fig. 22 and Appendix C 5). If $\theta_{h1} < \theta_{12}, \theta_{h2}$, then the $h1$ clustering will be the first step of the accumulation stage and may pass the SoftDrop condition, resulting in a pseudojet with the energy of h but the flavor of 1, which goes on to form a flavor cluster separate from that of q_2. The two flavor clusters in the jet now have a large hierarchy of energy, and the softer one (q_2) may ultimately annihilate with a large-angle soft quark (3 or 4) if the latter has a similar (or larger) p_t with respect to the beam, resulting in the flavor of the hard jet being set by the $h1$ cluster. This issue gives an $\alpha_s^4 L^3$ divergence for $\alpha\beta > 2$, and numerical results are consistent with an $\alpha_s^4 L$ divergence for $\alpha\beta = 2$. The analytical study of Appendix C 5 indicates that the problem should be resolved when one takes $\alpha\beta < 2$ (if one additionally replaces $\Delta R_{ij}^2 \to \Omega_{ij}^2$), though the above FHC^2 issue remains.

The discovery of the above issues highlights the importance of having a systematic framework for testing IRC safety. Indeed, some of these issues were first discovered with the testing framework, as were the identification of

[14] We are grateful to Simone Marzani for discussions that first led us to investigate this configuration analytically.

$\alpha\beta < 2$ as a potential solution for the GHS α_s^4 (IDS × FDS) issue, and the requirement Eq. (9) for the IFN algorithm. Such tests also led us to suspend study of more general u_{ij} distances involving $u_{ij} = [\max(p_{ti}, p_{tj})]^{2p} \times [\min(p_{ti}, p_{tj})]^{2q} \Omega_{ij}^2$, specifically a dimensionless form with $p = -q = 1$.

A final comment is that it is important to remember that the IRC tests cover many cases but are not totally exhaustive. Specifically, as discussed in Sec. IV A, we have at most one power of L per power of α_s, and only up to α_s^6, the events that we have generated have a band gap,[15] and a couple of nestings are still missing, notably as regards double-soft emissions. Thus, our tests should not be considered an ultimate proof of IRC safety but merely a strong indication.

V. PHENOMENOLOGICAL ILLUSTRATIONS

In this section, we present three phenomenological test cases, intended to convey some of the main features of our IFN algorithms. We include comparisons to standard anti-k_t clustering and also to those prior flavor algorithms for which we have been able to identify an IRC-safe adaptation, namely, flavor-$k_{t,\Omega}$ and CMP_Ω.

The first two tests will be specific to heavy flavor, which is the main experimental application of flavored jet algorithms. The third test will be for generic flavor and can be seen as a stress test of the algorithm's practical performance with light flavor at parton level.

A. Heavy flavor in $pp \to WH(\to \mu\nu b\bar{b})$

We begin with the case of Higgs production in association with a W boson at hadron colliders, $pp \to WH$, where the Higgs boson decays to a pair of b-quarks and the W decays leptonically. This process is of interest for obvious phenomenological reasons, e.g., because of the sensitivity to the HWW and $Hb\bar{b}$ couplings, and it has been measured by both ATLAS and CMS [39,40]. Additionally, it is one of the processes in which one can probe high-p_t Higgs production [41,42], especially in conjunction with jet substructure tools [43,44], bringing particular sensitivity to new physics. For a long time, calculations at NNLO QCD were performed with massless b quarks, which prohibited the use of the standard anti-k_t algorithm to cluster the final state. Only recently [32] was the calculation performed with massive b-quarks.

Here, we examine a classic resolved-jet analysis of this process, similar to that of Ref. [32]. We use PYTHIA8.306 [45,46] with the 4C tune [47] to generate $pp \to W(\to \mu\nu_\mu)H(\to b\bar{b})$. Following Ref. [32], we require the presence of a muon satisfying

$$|\eta_\mu| < 2.5, \qquad p_{t\mu} > 15 \text{ GeV}. \tag{12a}$$

We cluster the event with a given jet algorithm, using a jet radius of $R = 0.4$, and identify b-flavored jets that satisfy

$$|y_{j_b}| < 2.5, \qquad p_{tj_b} > 25 \text{ GeV}. \tag{12b}$$

We require the event to have at least two such jets. Finally, the reconstructed Higgs boson is defined as the 4-momentum sum of the two b-jets whose invariant mass is closest to the Higgs mass.

The distribution of the transverse momentum of the reconstructed Higgs boson is presented in Fig. 8 at hadron level (with multiparton interactions turned on), for four algorithms:

(i) standard anti-k_t with net flavor summation (red),
(ii) anti-k_t with our IFN algorithm ($\alpha = 2$, in green),
(iii) the CMP_Ω algorithm [$a = 0.1$, where the angular part of the distance measure is corrected as in Eq. (11), in black], and
(iv) the flavor-$k_{t,\Omega}$ algorithm ($\alpha = 2$, in gold).

The flavor-$k_{t,\Omega}$ algorithm leads to a reconstructed Higgs spectrum that is markedly different from that of the anti-k_t algorithm. In particular, for $p_{tH} \gtrsim 300$ GeV, the distribution starts to drop relative to that with anti-k_t, reaching about 60% of the latter's value at $p_{tH} \sim 600$ GeV. As noted in Ref. [32], this occurs because the flavor-k_t algorithm starts clustering the b and $\bar{b}$ together at lower values of p_{tH} than for the anti-k_t algorithm. When the b and $\bar{b}$ end up in a single jet, the event fails the selection requirement of having at least two b-jets. Specifically for the decay of a scalar particle with invariant mass m and transverse momentum p_t, for small R and in the limit of $p_t R \ll m$, the efficiency for having two separate jets (without any p_t or rapidity cut on the jets) is 1 at low p_t. Above some threshold in $x = p_t R/m > x_{\min}$, it becomes

$$\text{gen-}k_t\text{:}\ 1 - \frac{\sqrt{x^2-4}}{x}, \qquad x_{\min} = 2, \tag{13a}$$

$$\alpha = 1\ \text{flav-}k_t\text{:}\ \frac{2}{x^2}, \qquad x_{\min} = \sqrt{2}, \tag{13b}$$

$$\alpha = 2\ \text{flav-}k_t\text{:}\ \frac{2}{1+x^2}, \qquad x_{\min} = 1. \tag{13c}$$

[15] By band gap, we mean the white region between the upper hard tip of the Lund diagrams in Fig. 6 and the upper edge of the shaded region. The concern that one might have is that of an IRC unsafety mechanism whereby an emission at some momentum scale ϵ clusters with an emission at a still soft, but much larger scale $\epsilon^{1/p}$ (for some $p > 1$) in the white band-gap region, and only after that clustering can it cause the IRC unsafety. Our test procedure is only sensitive to values of $p < \ln p_{t,\min}/\ln p_{t,\max}$. One would therefore like to take as large a value of that ratio as possible, keeping in mind, however, that larger values of $\ln p_{t,\min}/\ln p_{t,\max}$ are numerically more challenging, both because of the higher precision and the need for higher statistics; cf. footnote 13.

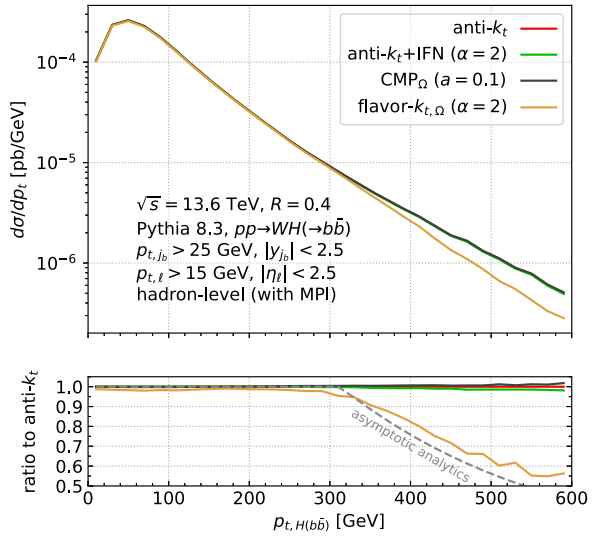

FIG. 8. The transverse momentum spectrum of the reconstructed Higgs boson in $WH(\to \mu\nu b\bar{b})$ at center-of-mass energy $\sqrt{s} = 13.6$ TeV, at hadron level (with stable B hadrons). The upper panel shows the spectrum for four jet algorithms: anti-k_t with net flavor of the jet constituents (red), our IFN version of anti-k_t (with $\alpha = 2$, green), the CMP$_\Omega$ algorithm [as adapted from Ref. [10] with a fix of the angular measure; see Eq. (11), black], and the flavor-$k_{t,\Omega}$ algorithm (with $\alpha = 2$, gold). The lower panel shows the ratio to standard anti-k_t. CMP$_\Omega$, and our IFN algorithms all give results very similar to those from the plain anti-k_t algorithm. In contrast, as already pointed out in Ref. [32], flavor-$k_{t,\Omega}$ jets can differ significantly from anti-k_t kinematics at large transverse momentum because they start clustering the b and $\bar{b}$ together into a single jet well before the scale of $p_t \simeq 2m_H/R = 625$ GeV where this occurs with the normal anti-k_t algorithm. This is reflected in Eq. (13c), which is used to generate the "asymptotic analytics" curve in the lower panel.

The last of these, in particular, explains the qualitative behavior seen in Fig. 8, cf. the "asymptotic analytics" dashed line in the lower panel. Note that Eqs. (13b) and (13c) are independent of the $\Delta R_{ij}^2 \to \Omega_{ij}^2$ change because they are evaluated in a small-angle limit, where the two distance measures are identical.

In Fig. 8, the CMP$_\Omega$ algorithm and anti-k_t + IFN give results that are very similar to those of plain anti-k_t jets, to within about a percent. This result is not entirely trivial; while it was expected that the new generation of flavor algorithms should be kinematically more similar or identical to anti-k_t, there was still a possibility that flavor assignments could modify cross sections, e.g., because the original anti-k_t jets' flavors would have been subject to modifications from soft $b\bar{b}$ pairs, while any such effect should be substantially reduced for the new algorithms. The absence of a numerically significant difference between jet algorithms other than flavor-$k_{t,\Omega}$ suggests that for signal processes like that shown here, with $R = 0.4$ jets, the contribution of soft $g \to b\bar{b}$ contamination is relatively small. On one hand, this means that certain experimental signal measurements with standard anti-k_t jets may not require much unfolding in order to be compared to the flavored jet definitions. On the other hand, if one is to perform a higher-order calculation with the approximation of massless b quarks, a flavored jet algorithm will still be required in order to obtain a finite result.

B. Heavy flavor in $pp \to t\bar{t} \to \ell\nu$ + jets

For our second test case with heavy flavor, we consider top-quark pair production in hadron collisions with semileptonic top decays $pp \to t\bar{t} \to \ell$ + jets. We select events by requiring at least one muon with

$$p_{t\mu} > 30 \text{ GeV}, \qquad |\eta_\mu| < 2.4, \tag{14}$$

and additionally $p_{t,\text{miss}} > 30$ GeV. We then run a jet algorithm and examine the p_t distribution of jets that are considered b flavored according to a given jet algorithm, with a requirement of $p_t > 20$ GeV applied to the jets. We again use PYTHIA8.306, but now with the Monash13 tune [48]. It will be instructive to examine results at both tree level (where we use b quarks as the flavored inputs to the algorithms) and hadron level including parton showering (where we use stable B hadrons as the flavored inputs).

The results are shown in Fig. 9. The inclusive b-jet p_t spectrum is shown in the upper panels, on the left at partonic tree level, i.e., without shower or hadronization, and on the right at hadron level including showering. We examine the same algorithms as in Fig. 8, and since we are again interested in the similarity of the distributions to that of standard anti-k_t with net flavor summation (in red), the lower panels show the ratio to that result. Additionally, we include a line for standard anti-k_t clustering with "any-flavor" tagging, i.e., counting a $b\bar{b}$-jet as a b-tagged jet, which is more in line with experimental procedures than net flavor recombination.

Let us start by examining the tree-level results in Fig. 9(a). When each jet contains at most one parton, the IFN algorithm is, by design, intended to give identical results to the plain net-flavor anti-k_t algorithm. Note that the tree-level PYTHIA sample does not guarantee this property, since our analysis does not require each of the four tree-level partons to be in four separate jets and, sometimes, two tree-level partons may cluster together. Nevertheless, we see that the IFN algorithm (green) gives results that are essentially indistinguishable from the plain net-flavor anti-k_t net-flavor (red) results. The flavor-$k_{t,\Omega}$ algorithm (gold) is expected to show differences, but these are relatively modest, typically a few percent. Finally, the CMP$_\Omega$ results [with $a = 0.1$, including the IRC fix as in Eq. (11), in black] show a few-percent depletion of low-p_t b-jets. We believe that this is a consequence of the small clustering distance for pairs of two low-p_t b-flavored

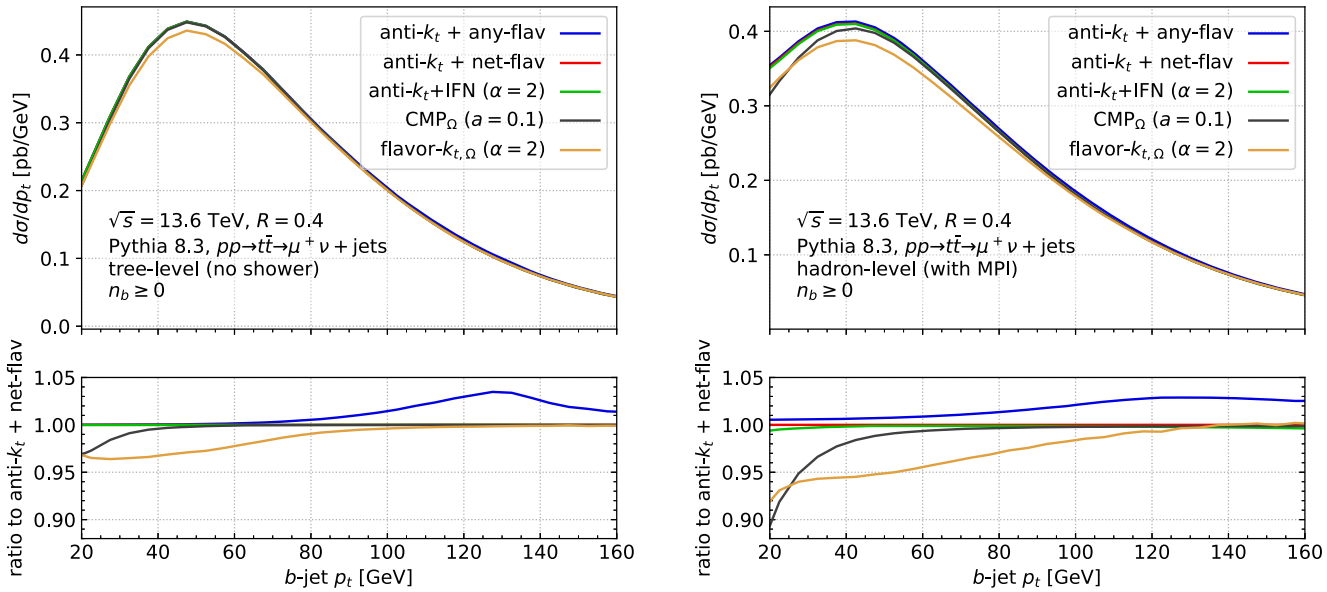

(a) (b)

FIG. 9. Inclusive b-jet spectrum from PYTHIA8.3 in $pp \to t\bar{t} + X \to b\mu^+\nu\bar{b}q\bar{q}' + X$ events at $\sqrt{s} = 13.6$ TeV, at (a) partonic tree level (i.e., no showering or hadronization) and (b) hadron level (with stable B hadrons). The distribution is shown in the upper panels, for four jet algorithms (as in Fig. 8) and additionally for anti-k_t with any-flavor recombination (i.e., a $b\bar{b}$ jet counts as b-tagged). The lower panels show the ratio to anti-k_t jets with net flavor summation. The anti-k_t + IFN algorithm yields a b-jet spectrum that is almost identical to that from the net-flavor anti-k_t algorithm, across the whole p_t range. The closeness to anti-k_t holds both at tree level and after showering and hadronization (with the spectrum differing maximally by less than a percent at $p_t = 20$ GeV, at hadron level). See the text for further details.

particles, which enhances the likelihood that such pairs will cluster, even when well separated.

At hadron-level in Fig. 9(b), including parton showering and multiparton interactions, the qualitative pattern is broadly similar, but with some effects enhanced relative to what is seen at tree level. Now, there are very small differences between anti-k_t and IFN, below a percent—this once again suggests that soft $g \to b\bar{b}$ induced contamination is a small effect, as we noted in the WH case of Sec. V A. In contrast, the relative differences of flavor-$k_{t,\Omega}$ and CMP$_\Omega$ as compared to anti-k_t now reach 8%–10%. Further examining the results, we have identified two effects that contribute to this: 1) a small reduction of the ratios due the shower and hadronization, in events where there are two B-hadrons, perhaps because fragmentation of the b quarks enhances the impact of modified clustering distances for flavored particles, and 2) substantially smaller ratios, especially at low p_t, in the ∼8% of events with an additional $b\bar{b}$ pair from the showering.

Finally, regarding anti-k_t with any-flavor recombination (in blue), we see that it differs only by a few percent from the net-flavor tagging and the IFN algorithms. The difference appears mostly to be associated with events where the b and $\bar{b}$ from the t and $\bar{t}$ decays end up in a single jet. Insofar as any-flavor recombination is a good stand-in for standard experimental tagging, the similarity of net-flavor and any-flavor recombination indicates only a limited need for unfolding corrections in order for experimental $t\bar{t}$ results to be presented unfolded to an IFN-style flavor truth level. Note, nevertheless, that there are other processes for which this would not be true, e.g., inclusive b-jet production [2], and a case-by-case study is needed to establish whether any-flavor and net-flavor recombinations are similar for a given process.

C. Full flavor at parton level in $pp \to Z + j$

Our final hadron-collider test is carried out at parton level (after showering) and applies jet flavor algorithms to all flavors of partons in the context of events with a hard jet recoiling against a high-p_t Z boson. This study is not intended to be of direct experimental relevance but rather to test the flavor algorithm's performance and limitations for addressing more theoretical questions such as the fractions of quark vs gluon jets. In particular, knowledge of the quark vs gluon fractions in a given sample is important when assessing the performance of approaches that attempt to distinguish quark vs gluon-induced jets from jet substructure and energy flow observables [49]. To do so, we study $pp \to Z + j$ events. We focus here on the $Z(\to \mu^+\mu^-) + q$ final state, where we require exactly two muons to reconstruct a high-p_t Z candidate:

$$|\eta_\mu| < 2.4, \qquad p_{t\mu} > 20 \text{ GeV}, \tag{15a}$$

$$p_{t,\mu^+\mu^-} > 1 \text{ TeV}, \qquad m_{\mu^+\mu^-} \in [80, 102] \text{ GeV}. \tag{15b}$$

We find qualitatively consistent results for the $Z + g$ case.

We use PYTHIA8.306 with the Monash13 tune to generate the events and specifically consider its $pp \to Z + q$ process. We cluster the events with a given jet algorithm and examine the flavor of the leading-p_t jet. At leading order, we expect the hard recoiling jet to always carry the flavor of the underlying quark or antiquark, and the question that we examine is that of how often the leading jet in the full showered sample has a flavor other than that of a single quark or antiquark.

Schematically, it is useful to think of two mechanisms that can cause the flavor to differ. One is that the quark can split to $q + g$ with a separation $\Delta R_{qg} > R$. If the gluon carries more energy than the quark, then the leading jet will actually be a gluon jet. The rate for this to happen is logarithmically enhanced in the small-R limit [50]. The second mechanism to keep in mind is the contamination of the flavor of a hard quark jet from a soft $g \to q\bar{q}$ splitting (i.e., the issue of Fig. 1, which flavored jet algorithms are supposed to mitigate against). This can have two effects: if the soft $q\bar{q}$ pair's flavor coincides with that of the jet, then it can cancel the jet's flavor; much more often, a fraction $\sim 1 - 1/(2n_f)$ of the time, it will lead to a multiflavor jet. To a first approximation, this effect is expected to grow with increasing jet radius. We show results both with and without multiparton interactions (MPIs), and we expect the flavor contamination to be worsened by MPI, insofar as it adds significant numbers of additional low-p_t $q\bar{q}$ pairs.

In Fig. 10, we show the fraction of leading-p_t jets that are flavorless (green), singly flavored (quark or antiquark, blue), or multiflavored (neither flavorless or singly flavored, red), as a function of the jet radius parameter R used in the clustering. We perform this comparison with PYTHIA at parton level, where the underlying event is turned off (upper row), and with MPI turned on (lower row). From left to right, the columns show results with the standard anti-k_t algorithm, flavor-$k_{t,\Omega}$ ($\alpha = 2$), and anti-k_t with our IFN algorithm for two values of $\alpha = \{1, 2\}$ (and $\omega = 3 - \alpha$). A first point to observe is the large multiflavored contribution for the plain anti-k_t algorithm, about 14% at $R = 0.4$ without MPI, increasing to 19% with MPI. Increasing R substantially worsens the situation with over 40% multiflavored jets for $R = 1$ when MPI is on.

Flavor-$k_{t,\Omega}$ improves the situation somewhat, giving a multiflavored contribution of 5% (10%) with MPI off (on) at $R = 0.4$. The anti-k_t algorithm with IFN brings a more

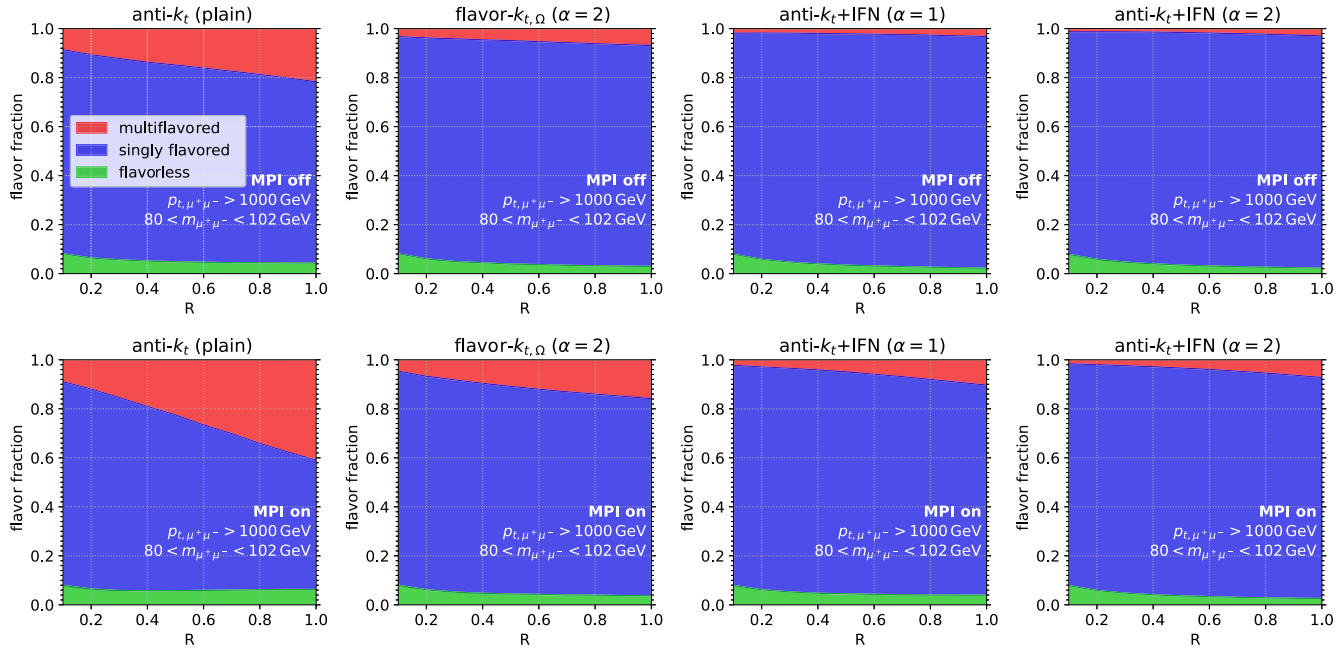

FIG. 10. Stress tests of the performance of the plain anti-k_t algorithm (with net flavor summation, left column), the flavor-$k_{t,\Omega}$ algorithm (middle left column), and the anti-k_t algorithm with flavor neutralization (with $\alpha = 1$, middle right column, and $\alpha = 2$, right column). The stress tests are performed in $pp \to Z + q$ collisions with $p_{tZ} > 1$ TeV, as simulated with PYTHIA8.3 at parton level with multiparton interactions disabled (enabled) on the upper row (lower row). As a function of the jet radius parameter R, the plots show the fraction of leading jets that are multiflavored, i.e., whose flavor is neither that of a gluon nor a single quark or antiquark (red band), singly flavored (blue band), and flavorless (green band). The key observation is the large fraction of multiflavored jets with the standard anti-k_t algorithm, which occur due to contamination of the hard jet flavor from low-momentum particles. With the flavor-$k_{t,\Omega}$ algorithm, we see some reduction, while anti-k_t with IFN shows a further reduced rate, especially for $\alpha = 2$.

substantial improvement, yielding 2% (4%) for $\alpha = 1$ and 1.5% (3%) for $\alpha = 2$.[16]

Examining instead the unflavored ("gluon") jet fractions, we find that all flavor algorithms give a ∼4% gluon-jet fraction at $R = 0.4$, relatively unaffected by the presence of MPI. This figure is important to keep in mind for quark/gluon discrimination studies [49]; the fact that a jet was initiated by a quark in PYTHIA does not mean that the corresponding jet observed after showering is always a quark jet. In particular, Fig. 10 implies that if one is attempting to tag gluon jets and reject quark jets, and one is using PYTHIA's $Z + q$ and $Z + g$ samples as the sources of quark and gluon jets, then even a perfect gluon tagger will still show an acceptance of about 4% on the $Z + q$ sample.

Ultimately, we would argue that the truth flavor labels should be derived not from the generation process but by running a jet flavor algorithm such as anti-k_t + IFN. Nevertheless, the anti-k_t + IFN labeling remains subject to some ambiguities, and the multiflavored jet fraction discussed above is probably a good measure of those ambiguities. As a future direction, one might wish to investigate whether one can develop jet flavor algorithms that further reduce the multiflavored jet fraction, while maintaining other good properties.

VI. EXPLORATION OF IFN ALGORITHM FOR e^+e^- COLLISIONS

The IFN algorithm for e^+e^- collisions follows the same set of rules as the pp version in Sec. III, adapting the u_{ij} distances so that they coincide with the e^+e^- flavor-k_t distances, specifically

$$u_{ik} \equiv [\max(E_i, E_k)]^{\alpha}[\min(E_i, E_k)]^{2-\alpha} \times 2(1 - \cos\theta_{ik}). \tag{16}$$

In contrast with the hadron-collider case, there is no need to use a modified angular distance in the e^+e^- form of the u_{ij}. Additionally in step 1 of the part of the algorithm in Sec. III B, we define i to be the particle with lower energy.

We have not explicitly performed the same full set of IRC safety tests on the e^+e^- algorithm that we carried out in the pp case. The issues identified across various algorithms in the pp case fall into two classes: those involving initial-state hard-collinear radiation, which is irrelevant in the e^+e^- case, and those involving just the interplay between final-state soft large-angle and hard-collinear branchings. The analyses of these latter issues in Appendix C are expected to be insensitive to the differences between Eqs. (7) and (16). For this reason, we do not anticipate IRC safety issues in the e^+e^-, though a detailed study would ultimately be desirable.

The relatively clean environment of e^+e^- collisions allows for a further exploration of the performance of IFN-style algorithms. Specifically, we take a parton shower simulation of $e^+e^- \to q\bar{q}$ events, cluster each event into two jets, and examine the flavors of those two jets. Typically, one expects each jet to have the flavor of the parton in the Born event. In analogy with our pp study of Fig. 10, we examine the fraction of events where this does not happen, breaking it into two components, one where each jet is flavorless ("gg") and the other where the two jets have neither the original $q\bar{q}$ flavor nor gg flavor ("other"). We plot these fractions as a function of $\sqrt{s}$. For a well-behaved flavor algorithm, we expect the rate of gg configurations to decrease with increasing $\sqrt{s}$, as $\alpha_s(\sqrt{s})$, associated with the probability of producing a $q\bar{q}g$ configuration where the $q\bar{q}$ pair ends up back to back with respect to the gluon. Similarly, the rate of other configurations should decrease as $\alpha_s^2(\sqrt{s})$, since one expects to have to generate a hard $q\bar{q}q'\bar{q}'$ configuration to obtain an other flavor.

We carry out the parton shower simulation with PYTHIA8.306 (4C tune) at parton level. The jet clustering is performed as follows: for the k_t and C/A algorithms, we use a large radius $R = 2\pi$ and then decluster the event back to two jets by undoing the last stage of the clustering in order to obtain the two hard jets. This is equivalent to asking for two hard jets in the normal exclusive k_t or Cambridge algorithms. For the anti-k_t algorithm, we use a jet radius of $R = 3\pi/4$ and take the two highest energy jets. The specific jet radius choice is designed to be large enough that events with multiple well-separated hard partons still only give two hard jets but small enough that opposite hemispheres of the event do not cluster together.

Figure 11 shows the gg (left column) and other (right column) rates as a function of $\sqrt{s}$ in the range 100 to 10^5 GeV. This broad (and today unrealistic) energy range is intended to help visualize the scaling behavior of the rates. Each row corresponds to one underlying jet algorithm, with different curves showing results for different flavor approaches. Let us start by examining the gg rate for the anti-k_t algorithm and its IFN variants (top-left panel). Over the energy range being considered, α_s decreases by almost a factor of 2. The anti-k_t + IFN algorithms show a gg rate that is more or less consistent with $\alpha_s(\sqrt{s})$ scaling. In contrast, the plain anti-k_t algorithm (with net flavor summation) features a rate that slowly increases. It is natural to ascribe the growth to the IRC unsafety of the algorithm; however, the differences between the safe (IFN) and unsafe (plain anti-k_t) variants remain relatively modest.

The situation becomes clearer when looking at the other flavor combinations (top-right panel). Here, the plain anti-k_t algorithm gives a rate that increases from about 10% at $\sqrt{s} = 100$ GeV to almost 30% at $\sqrt{s} = 10^5$ GeV.

[16]For the CMP_Ω algorithm, there is freedom in how one extends it to multiflavored events, and accordingly, we defer study of multiflavored events with that algorithm to future work.

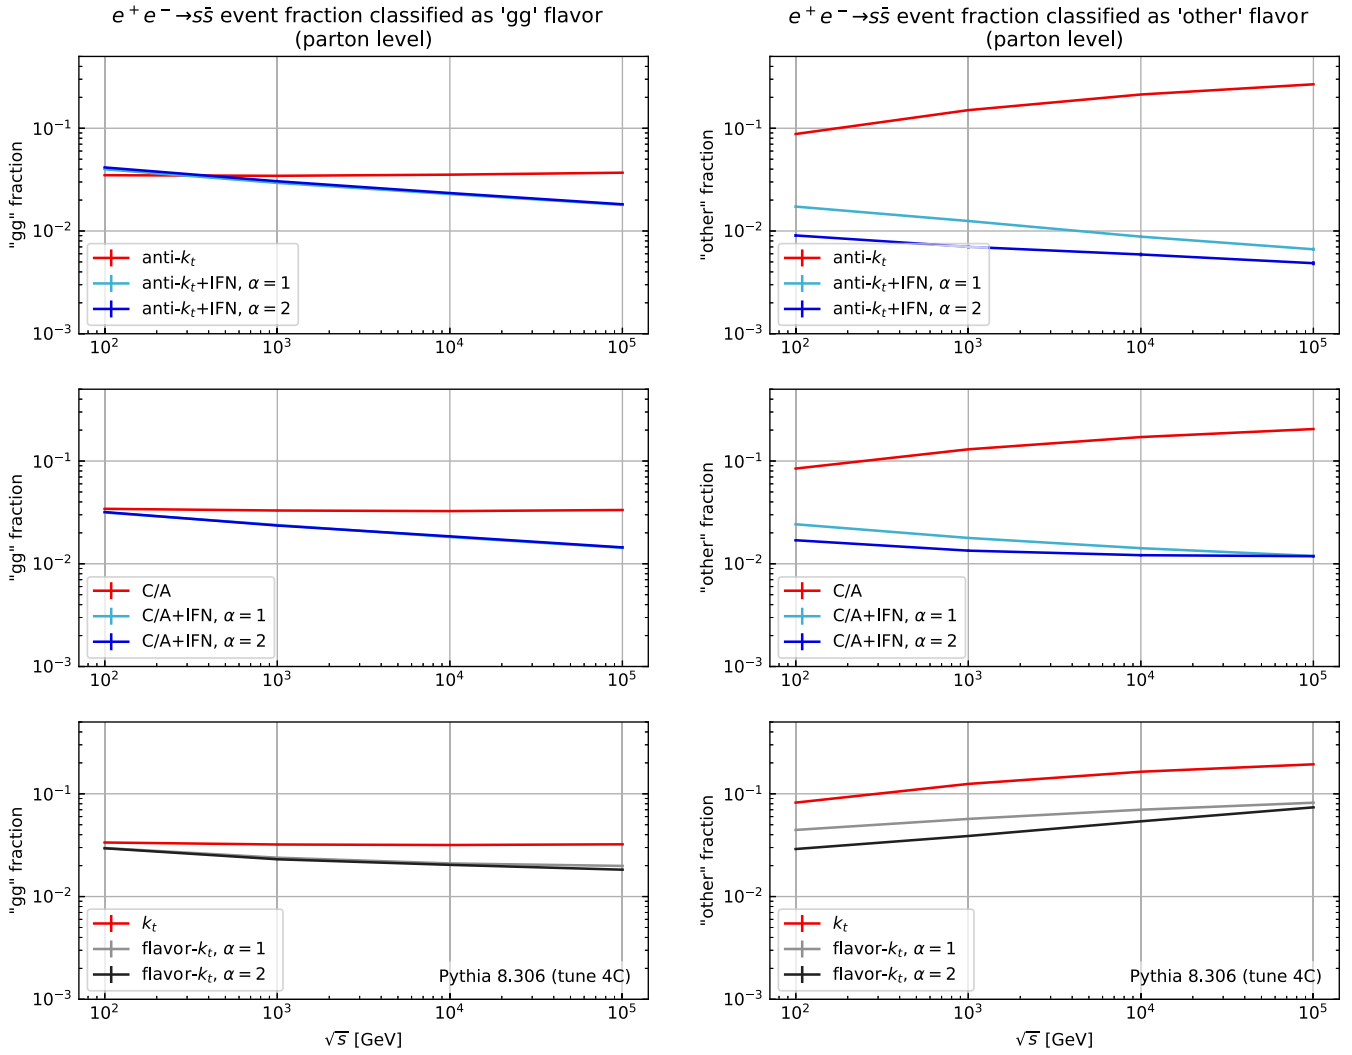

FIG. 11. The fraction of $e^+e^- \to q\bar{q}$ events (after parton showering and clustering to two jets), in which the flavor of the two jets is classified as being gg (left column) or any other combination that is not the original $q\bar{q}$ (right column). The results are shown as a function of $\sqrt{s}$ for algorithms in the anti-k_t family (top row), the C/A family (middle row), and k_t family (lower row). The results have been obtained using PYTHIA8.306 at parton level, with tune 4C.

In contrast, the IRC-safe IFN variants give much smaller rates, well below 2% across the whole energy range, i.e., a huge improvement on the plain anti-k_t algorithm. With the $\alpha = 1$ IFN choice, the rate decreases more or less consistently with the expectation of α_s^2 scaling. The $\alpha = 2$ IFN algorithm shows a lower rate, but also a slower scaling. The situation is broadly similar for the C/A algorithm, with the IFN rate a little higher. For the k_t family, we show only the plain k_t algorithm and the flavor-k_t algorithm, since we have not conclusively validated the IRC safety of the k_t + IFN combination. Interestingly, the flavor-k_t algorithm shows only modest improvement in the other rate relative to the plain k_t algorithm and a scaling that is no better.

It is intriguing that different IRC safe algorithms lead to other rates that have varying degrees of consistency with the expected α_s^2 scaling. While we do not yet have a complete understanding of this phenomenon, detailed investigations into the events have revealed all-order mechanisms that operate across multiple scales and that, in some situations, cause soft flavor to be successively associated with harder and harder momenta, ultimately transferring soft flavor to the hard jets. Further study would require detailed analysis of the interplay between the main jet algorithm's clustering sequence and the IFN flavor neutralization scales. Still, despite these observations, the IFN algorithms clearly perform much better than plain flavor unsafe ones, indicating the substantial benefits for detailed flavor studies in using a suitably chosen flavor-safe algorithm.

VII. CONCLUSIONS AND OUTLOOK

In this article, we introduced an approach to jet clustering that maintains the kinematics of the original anti-k_t and

C/A algorithms, while also providing IRC-safe jet-flavor identification. Our IFN algorithm has passed a battery of fixed-order IRC safety tests, which revealed a number of unexpected and subtle issues in prior jet-flavor proposals. While not an absolute guarantee, these tests do provide a reasonable degree of confidence in the IRC safety of our approach. On three benchmark jet flavor tasks, IFN exhibits the desired phenomenological behavior. These studies suggest that IFN can yield a theoretically sound meaning to the concept of a flavored jet in the majority of heavy-flavor related applications that can be envisaged at the LHC.

There are various experimental considerations that should be noted before deploying IFN in a full analysis. Our algorithm, like all other attempts at IRC-safe flavor jet algorithms, requires the complete flavor information in the event for those flavors under consideration, e.g., b flavor. This is highly challenging in an experimental environment because of the difficulties of tagging low-momentum single B hadrons, as well as quasicollinear pairs of B hadrons. The question remains, however, whether recent advances in machine learning can help reveal the information that is needed and, more generally, whether experimentally one can unfold detector-level results to particle-level jet definitions such as ours. Furthermore, for certain signal processes, the practical impact of this issue may only be moderate, cf. the $\lesssim 3\%$ difference between the any-flavor and net-flavor anti-k_t results for $t\bar{t}$ (Fig. 9 right).

Theoretically, we stress that the concept of jet flavor remains subtle also beyond the scope of the discussion in this article. We focused on the fixed-order behavior, but there can be nontrivial interplay with the still perturbative but complex structure induced by all-order showered events. Beyond a perturbative analysis, there are even more difficult issues of jet flavor in the presence of the high densities of flavored particles that result from hadronization. These questions warrant more investigation. Nevertheless, the IFN algorithm developed here already shows clear and substantial benefits both with respect to standard unflavored algorithms and to prior incarnations of flavored algorithms.

Code implementing the IFN algorithm is available from https://github.com/jetflav/IFNPlugin, in the form of a FastJet Plugin.

ACKNOWLEDGMENTS

We are grateful to Andrea Banfi, Simone Marzani, and Giulia Zanderighi for helpful discussions and comments on the manuscript. We are also grateful to the authors of Refs. [10,11] for discussions about their respective algorithms and more generally on these topics. G. P. S. and L. S. would also like to thank their PanScales collaborators for joint work on the high-precision adaptations of the FJCORE code used in the IRC safety tests. G. P. S. would also like to thank Matteo Cacciari and Gregory Soyez for joint work on updates to FastJet to facilitate the inclusion of flavor in plugins. This work has been funded by a Royal Society Research Professorship (RP\R1\180112) (G. P. S. and L. S.), by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme [Grant No. 788223, PanScales (G. P. S., L. S., and M. H.) and Grant No. 804394, HipQCD (F. C.), and by the Science and Technology Facilities Council (STFC) under Grant No. ST/T000864/1 and by Somerville College (L. S.). J. T. is supported by the U.S. DOE Office of High Energy Physics under Grant Contract No. DE-SC0012567. R. G. is supported by the STFC, a Wolfson Harrison UK Research Council Physics Scholarship, and a Clarendon Scholarship and in the early stages of this work benefited from support from Merton College. Part of this work benefited from the support and hospitality of the Munich Institute for Astro-, Particle and BioPhysics (MIAPbP), which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC-2094–390783311.

APPENDIX A: ASYMMETRIC DOUBLE-SOFT BRANCHING

As discussed in Fig. 1, the classic IRC safety issue involves configurations with a double-soft $g \to q\bar{q}$ pair. One surprising feature that we will encounter in Appendixes B 1, C 1, and C 3 is that the structure of higher-order IRC divergences in some algorithms is sensitive to configurations where one of the quarks is significantly softer than the other. Consequently, it is important to understand the asymptotic behavior of double soft $q\bar{q}$ production in such limits.

We know that collinear $g \to q\bar{q}$ splitting comes with a $P_{qg}(z)dz = T_R(z^2 + (1-z)^2)dz$ structure, which is finite and nonzero when $z \to 0$. One question one might ask, though, is whether the splitting probability remains finite and nonzero for $z \to 0$ when the pair is not collinear but instead separated by an angle that is commensurate with the emission angle of the parent gluon. In this appendix, we find that the splitting probability is indeed finite in that limit, and we derive a simple approximate expression for its behavior.

Consider a process with n hard massless QCD partons with momenta $\{p_i\}$, and study the emission of two additional soft quarks with momenta $q_{1,2}$. The double-soft emission probability can be written as a sum over dipole contributions as

$$d\mathcal{P}_{\text{d.s.}} = -\sum_{i\neq j=1}^{n} \mathbf{T}_i \cdot \mathbf{T}_j d\mathcal{P}^{(i,j)}_{\text{d.s.}}, \tag{A1}$$

where $\mathbf{T}_i$ are the standard color operators (see, e.g., Ref. [51] for details) and $d\mathcal{P}^{(i,j)}$ only depends on the

momenta of the soft quarks and of the hard partons i and j. To write an explicit representation for $d\mathcal{P}^{(i,j)}$, it is convenient to work in the dipole center-of-mass frame. Specifically, we write

$$\begin{aligned} p_i &= E(1;0,0,1),\\ p_j &= E(1;0,0,-1),\\ q_1 &= p_{t,1}(\cosh y_1;1,0,\sinh y_1),\\ q_2 &= p_{t,2}(\cosh y_2;\cos\Delta\phi,\sin\Delta\phi,\sinh y_2). \end{aligned} \tag{A2}$$

We stress that $p_{t,1,2}$, $y_{1,2}$, and $\Delta\phi$ are dipole specific and not global variables. In terms of these variables, the double-soft emission probability reads

$$d\mathcal{P}^{(i,j)}_{\text{d.s.}} = \left(\frac{\alpha_s}{2\pi}\right)^2 4T_R dp_{t,1}dp_{t,2}dy_1 d\Delta y \frac{d\Delta\phi}{2\pi}\frac{2p_{t,1}p_{t,2}-(p_{t,1}^2+p_{t,2}^2)\cos\Delta\phi+|\vec{p}_{t,1}-\vec{p}_{t,2}|^2\cosh\Delta y}{(p_{t,1}^2+p_{t,2}^2+2p_{t,1}p_{t,2}\cosh\Delta y)^2(\cosh\Delta y-\cos\Delta\phi)^2}, \tag{A3}$$

where $\Delta y = y_2 - y_1$ and $T_R = 1/2$. In the limit when the soft quark pair is also collinear to parton i, Eq. (A1) simplifies to

$$d\mathcal{P}_{\text{d.s.}} = -\sum_{i\neq j=1}^{n} \mathbf{T}_i\cdot\mathbf{T}_j d\mathcal{P}^{(i,j)}_{\text{d.s.}} \to C_i d\mathcal{P}^{(i,j)}_{\text{d.s.}}, \tag{A4}$$

with $d\mathcal{P}^{(i,j)}_{\text{d.s.}}$ still given by Eq. (A3) and where $C_i = C_A$ if parton i is a gluon and $C_i = C_F$ if it is a quark.

We can now study the asymmetric $p_{t,2} \ll p_{t,1}$ configuration. We write $p_{t,2} = zp_{t,1}$. In the small-z region, Eq. (A3) becomes

$$d\mathcal{P}^{(i,j)}_{\text{d.s.}} \sim \left(\frac{\alpha_s}{2\pi}\right)^2 \frac{4T_R dp_{t,1}dzdy_1 d\Delta y d\Delta\phi}{2\pi p_{t,1}(\cosh\Delta y-\cos\Delta\phi)}. \tag{A5}$$

We see that, as for the P_{qg} splitting function, this probability is finite and nonzero for $z \to 0$. For our analysis, we also find it useful to consider Eq. (A3) in the limit of large rapidity separation between the two soft quarks, $\Delta y \gg 1$. We obtain

$$\begin{aligned} d\mathcal{P}^{(i,j)}_{\text{d.s.}} \sim \left(\frac{\alpha_s}{2\pi}\right)^2 8T_R \frac{dp_{t,1}}{p_{t,1}}\frac{dp_{t,2}}{p_{t,2}}dy_1 d\Delta y\frac{d\Delta\phi}{2\pi}\\ \times \frac{p_{t,1}p_{t,2}e^{-\Delta y}}{(p_{t,1}+e^{\Delta y}p_{t,2})^2}. \end{aligned} \tag{A6}$$

In the asymptotic regime, the second line of Eq. (A6) is well approximated by the expression

$$\frac{p_{t,1}p_{t,2}e^{-\Delta y}}{(p_{t,1}+e^{\Delta y}p_{t,2})^2} \approx \min\left[\frac{p_{t,1}}{p_{t,2}}e^{-3\Delta y}, \frac{p_{t,2}}{p_{t,1}}e^{-\Delta y}\right], \tag{A7}$$

which interpolates between the $1 \ll \ln(p_{t,1}/p_{t,2}) \ll \Delta y$ and $1 \ll \Delta y \ll \ln(p_{t,1}/p_{t,2})$ limits.

In practice, we find that this interpolation works well across the whole phase space. As shown in Fig. 12, we find that

$$\begin{aligned} d\mathcal{P}^{(i,j)}_{\text{d.s.,approx}} \equiv \left(\frac{\alpha_s}{2\pi}\right)^2 8T_R \frac{dp_{t,1}}{p_{t,1}}\frac{dp_{t,2}}{p_{t,2}}dy_1 d\Delta y\frac{d\Delta\phi}{2\pi}\\ \times \min\left[\frac{p_{t,1}}{p_{t,2}}e^{-3\Delta y}, \frac{p_{t,2}}{p_{t,1}}e^{-\Delta y}\right] \end{aligned} \tag{A8}$$

provides a good approximation of the exact $d\mathcal{P}^{(i,j)}_{\text{d.s.}}$ result down to values of $\ln(p_{t,1}/p_{t,2})$ and Δy of order 1. Furthermore, it is free of the collinear divergence when the q and $\bar{q}$ go close in angle, a collinear divergence

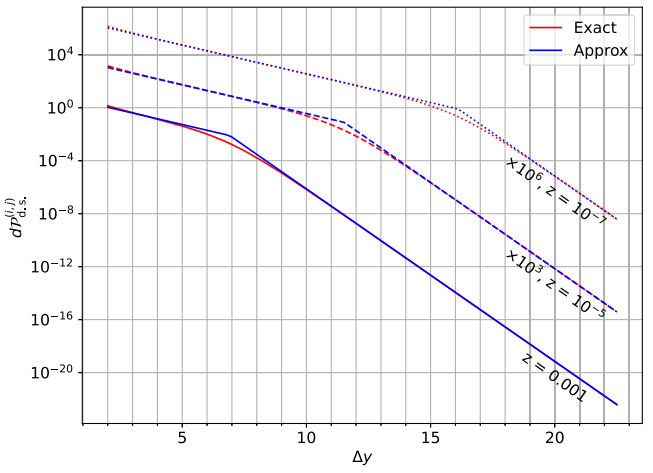

FIG. 12. The double-soft matrix element $d\mathcal{P}^{(i,j)}_{\text{d.s.}}$ in (red) its exact form from Eq. (A3) and in (blue) the approximate form from Eq. (A8) that we use in our numerical IRC-safety tests, as a function of the rapidity separation Δy of the two soft partons in the dipole center-of-mass frame. The matrix element is plotted for $p_{t,1} = 1$ GeV, and $\Delta\phi = 0$, for several values of z (shown as full, dashed, and dotted lines). The two regimes with $\sim e^{-\Delta y}$ and $\sim e^{-3\Delta y}$ scaling are clearly visible.

that we deliberately wish to leave out, because of our approach of allowing at most one divergence (or logarithm) per power of α_s. In our studies, we therefore use the convenient interpolation Eq. (A8) rather than the exact result Eq. (A3).

APPENDIX B: NUMERICAL TESTS OF IFN

In the process of developing the IFN algorithm, we tested a number of possible variants. In this appendix, we provide numerical support for the analytic arguments made in Sec. III to justify our design choices.

1. Relation between α and ω

According to the arguments in Sec. III C, we need to take particular care in choosing the values of the parameters α and ω in the IFN algorithms. A potentially dangerous configuration is that presented in Fig. 3. In that diagram, two partons (one flavored and one flavorless) at central rapidity are clustered together by the anti-k_t algorithm. In our IFN algorithms, the IRC safety issue arises from an initial-state hard-collinear splitting, which can act as a possible neutralization partner for the flavored hard particle. As argued in Eq. (9), the condition $\alpha + \omega > 2$ ensures that such a neutralization does not happen.

To test this argument numerically, we integrate uniformly over the momentum of each of a central hard quark and hard gluon (each in the range 1 GeV to 1 TeV) and sample an IHC emission as described in Sec. IV A. The results are presented in Fig. 13, for various values of the parameters α and ω. As expected, in cases where $\alpha + \omega < 2$, as well as for IFN variants that use a ΔR_{ij}^2 type angular distance instead of our Ω_{ij}^2, the failure rate typically diverges for $p_{t,\max} \to 0$ and conversely falls off as a power law when $\alpha + \omega > 2$ (green and blue curves). We observe numerically that the border cases, $\alpha + \omega = 2$, are all unsafe.

Let us see analytically why $\alpha + \omega = 2$ is problematic for the specific case of $\alpha = 2$ (and $\omega = 0$). We note that in the limit where $\omega \to 0$ the angular factor Ω_{ik}^2 in Eq. (7b) differs from ΔR_{ik}^2 at most by a factor of $\mathcal{O}(1)$, which we can typically neglect in the discussions below. We take the configuration shown in Fig. 3 with $p_{t2} = z_2 p_{t3}$ with $z_2 \ll 1$, $p_{t3} = 1$. There are two competing distances in the neutralization step,

$$u_{12} = z_2^2 \Delta y_{12}^2 \simeq z_2^2 \ln^2 1/p_{t1}, \tag{B1a}$$

$$u_{23} = \Delta R_{23}^2 \sim 1. \tag{B1b}$$

The IFN algorithm will neutralize the flavors of 1 and 2 when $z_2 \ln 1/p_{t1} < \Delta R_{23}$. If we integrate over the momentum of 2 and assume a dz_2 distribution (see, e.g., Ref. [52]) for finite ΔR_{23} and take $z_2 \to 0$, then the resulting integral is given by $\int d\ln p_{t1} \int_0^{1/\ln p_{t1}} dz_2$, which diverges.

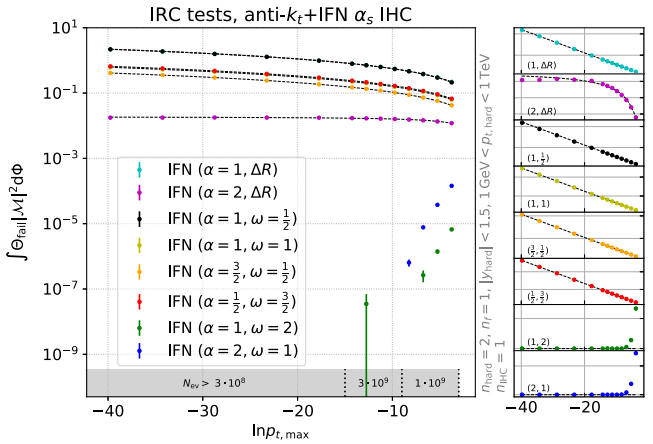

FIG. 13. IRC safety tests of anti-k_t + IFN for variants with different angular scaling factors. The tested configuration from Fig. 3 features two hard partons clustered together by the anti-k_t algorithm and one initial-state hard-collinear splitting. With a ΔR_{ij}^2 angular factor, the IFN algorithms diverge for all choices of α. Switching to the Ω_{ij}^2 angular distance, the cases where $\alpha + \omega \leq 2$ also diverge, whereas for $\alpha + \omega > 2$, they converge to zero as a power law, as expected from Eq. (9). The right-hand-side panels show the results on a linear scale, to help visualize the scaling for the IRC-unsafe variants. Note that here and in some of the following figures certain curves differ in the number of generated events. Because the multiple curves feature different scaling behaviors, they require a varying number of events to make a conclusive statement about the form of the divergence (or the absence thereof). The smallest number of events generated among all runs performed for a given figure is shown in the three shaded regions at the bottom of the plot.

The analytic argument shown here does not apply to generic values of α and ω, but as mentioned above, we find numerically that all cases that we have tested with $\alpha + \omega = 2$ diverge.

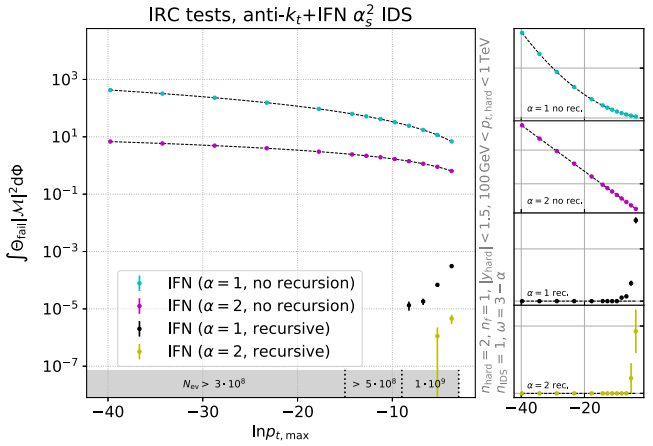

FIG. 14. IRC safety test of anti-k_t + IFN for variants with and without the recursion step. The tested events consist of two hard partons supplemented with one initial-state double soft pair, as in Fig. 4.

2. Recursive vs nonrecursive

In Sec. III D, we presented an analytic argument to explain why the IFN algorithms need a recursion step. Figure 14 shows the failure rate events with two hard partons and one IDS pair, which includes configurations such as that of Fig. 4. It clearly shows that without recursion the algorithm shows a growing failure rate for $p_{t,\max} \to 0$, while the failure vanishes for $p_{t,\max} \to 0$ with the recursive step turned on. The side figures help illustrate that the failure rate goes as $\ln^2 p_{t,\max}$ for $\alpha = 1$ and as $\ln p_{t,\max}$ for $\alpha = 2$. The stronger power for $\alpha = 1$ arises because failures can happen even when the IDS pair is collinear to the beams.

APPENDIX C: IRC-UNSAFE CONFIGURATIONS

In this appendix, we analyze the specific IRC-unsafe configurations identified in Sec. IV D for the flavor-k_t, CMP and GHS algorithms. For each of the configurations that we have identified, we present both analytic and numerical results to demonstrate why they are problematic. Throughout this section, we define $p_{ti} \lesssim p_{tj}$ to mean that $p_{ti} < p_{tj}$ but that they are of similar orders of magnitude.

1. IHC × IDS subtlety at α_s^3 for flavor-k_t

The flavor-k_t (and GHS) algorithms encounter a problematic configuration at order α_s^3, shown in Fig. 15, associated with the choice of angular measure. There is a hard parton (1, with flavor q) that produces the only hard jet in the event, together with a soft gluon g that splits to a soft large-angle $q\bar{q}$ pair (2 and 3), and additionally an initial-state collinear gluon splitting that produces an energetic small-angle quark of flavor $\bar{q}$ (4). For the sake of the discussion, we assume that the transverse momentum of 2 is smaller than that of 3, by a factor z_{23},

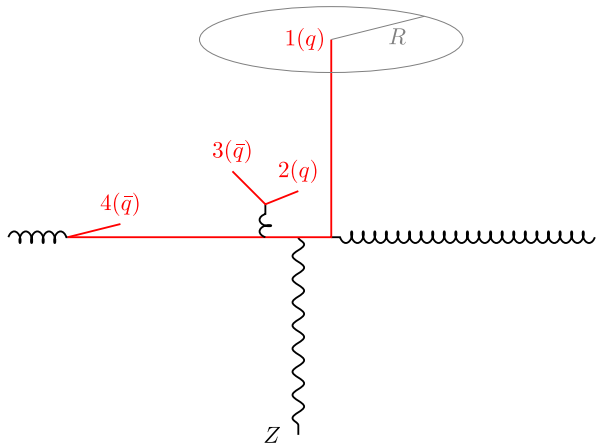

FIG. 15. Example configuration to illustrate issues that arise across multiple algorithms when using a standard ΔR-type angular measure in interparticle distances.

$$p_{t2} = z_{23} p_{t3}. \tag{C1}$$

We take the rapidity and azimuth of 1 to be zero, $y_1 = \phi_1 = 0$. Additionally, we take $E_4 = z_{41} E_1$, which implies

$$y_4 = \ln \frac{2 z_{41} p_{t1}}{p_{t4}}. \tag{C2}$$

For concreteness, we work with the $\alpha = 2$, $R = 1$ variant of the flavor-k_t algorithm. The interparticle distances for the flavor-k_t algorithm were given in Eq. (2), and additionally, we will need to take into account a beam distance with the right-moving beam. When i is flavored (as it will always be in the example here),

$$d_{iB}^{\text{flav-}k_t} = [\max(p_{ti}, p_{tB}(y_i))]^{\alpha} [\min(p_{ti}, p_{tB}(y_i))]^{2-\alpha}, \tag{C3}$$

with

$$p_{tB}(y) = \sum_i p_{ti} [\Theta(y - y_i) e^{y_i - y} + \Theta(y_i - y)], \tag{C4}$$

so that

$$p_{tB}(0) \simeq p_{t1}, \tag{C5a}$$

$$p_{tB}(y_4) \simeq p_{t4} + p_{t1} e^{-y_4} \simeq p_{t4} \left(1 + \frac{1}{2 z_{41}}\right), \tag{C5b}$$

where we have used the fact that $p_{t4} \ll p_{t1}$. In the absence of the soft quark pair, there are three distances:

$$d_{14} = p_{t1}^2 \Delta R_{14}^2 \simeq p_{t1}^2 y_4^2, \tag{C6a}$$

$$d_{1B} \simeq p_{t1}^2, \tag{C6b}$$

$$d_{4B} \simeq p_{t4}^2 \left(1 + \frac{1}{2 z_{41}}\right)^2. \tag{C6c}$$

The smallest is d_{4B}, since $p_{t4}/z_{41} \ll p_{t1}$ and so initial-state collinear particle 4 clusters first, leaving a flavored jet consisting of particle 1.

Now, we examine the additional distances that arise when the soft $q\bar{q}$ (23) pair is present. The beam distances d_{2B} and d_{3B} are both similar to $d_{1B} = p_{t1}^2$, since they are at central rapidities where $p_{tB} \sim p_{t1}$. The distances that will matter for the clustering are

$$d_{23} = p_{t3}^2 \Delta R_{23}^2 \sim \frac{p_{t2}^2}{z_{23}^2}, \tag{C7a}$$

$$d_{24} = \max(p_{t4}^2, p_{t2}^2) \Delta R_{24}^2 \sim \max(p_{t4}^2, p_{t2}^2) \ln^2 \frac{z_{41}^2 p_{t1}^2}{p_{t4}^2}, \tag{C7b}$$

where $\sim$ implies that we leave out factors of $\mathcal{O}(1)$, e.g., from $\Delta R_{23}^2 \sim 1$. We neglect d_{34}, since in the moderately small z_{23} limit where we will be working [cf. Eq. (C1)] $d_{34} > d_{24}$.

If d_{23} is the smallest of the distances across Eqs. (C6) and (C7), particles 2 and 3 annihilate, then 4 clusters with the beam, and the hard jet has flavor q. If d_{4B} is the smallest, 4 clusters with the beam, then 2 and 3 annihilate, and the hard jet has flavor q. The problematic situation is when d_{24} is the smallest of the distances, causing 2 and 4 to annihilate. This leaves 3, which can cluster with 1, resulting in a flavorless hard jet.

To understand the likelihood of this occurring, we first introduce the shorthand

$$\ell_{ij} = \ln \frac{p_{ti}}{p_{tj}}. \tag{C8}$$

Let us first consider $p_{t4} < p_{t2}$ ($\ell_{24} > 0$). The d_{24} will be the smallest one when

$$d_{24} < d_{23} \to \ell_{14} + \mathcal{O}(\ln z_{41}) < \frac{1}{z_{23}}, \tag{C9a}$$

$$d_{24} < d_{4B} \to \ell_{14} + \mathcal{O}(\ln z_{41}) < \frac{e^{-\ell_{24}}}{z_{41}}. \tag{C9b}$$

Ignoring azimuthal integrals (and the rapidity of 2 and 3), we now have to integrate over four phase-space variables, which we take to be ℓ_{14}, ℓ_{24}, z_{23}, and z_{41}. We have found that we can ignore the $\mathcal{O}(\ln z_{41})$ terms in Eq. (C9) and rewrite the limits as

$$z_{23} < \frac{1}{\ell_{14}}, \qquad z_{41} < \frac{e^{-\ell_{24}}}{\ell_{14}}, \qquad [\ell_{24} > 0]. \tag{C10}$$

Since both z fractions will be small, we will perform the integrations over z_{23} and z_{41} using the constant small-z limit of $P_{g\to q\bar{q}}(z)dz$. The overall rate of 24 clustering (with $l_{24} > 0$) is then given by

$$\begin{aligned} N_{24,\text{flav-}k_t}^{(\ell_{24}>0)} &\sim \alpha_s^3 \int_0^\infty d\ell_{14} \int_0^{\ell_{14}} d\ell_{24} \int_0^{\frac{1}{\ell_{14}}} dz_{23} \int_0^{\frac{e^{-\ell_{24}}}{\ell_{14}}} dz_{41} \\ &\sim \alpha_s^3 \int^\infty \frac{d\ell_{14}}{\ell_{14}^2}, \end{aligned} \tag{C11}$$

where in setting the lower limits of the ℓ_{14} and ℓ_{24} integrals to zero, we are ignoring any constraints from the interplay, e.g., with the z_{23} integral. A similar analysis can be carried out for $\ell_{24} < 0$ (or equivalently $\ell_{42} > 0$), giving

$$z_{41} < \frac{1}{\ell_{14}}, \qquad z_{23} < \frac{e^{-\ell_{42}}}{\ell_{14}}, \qquad [\ell_{42} > 0], \tag{C12}$$

and yielding

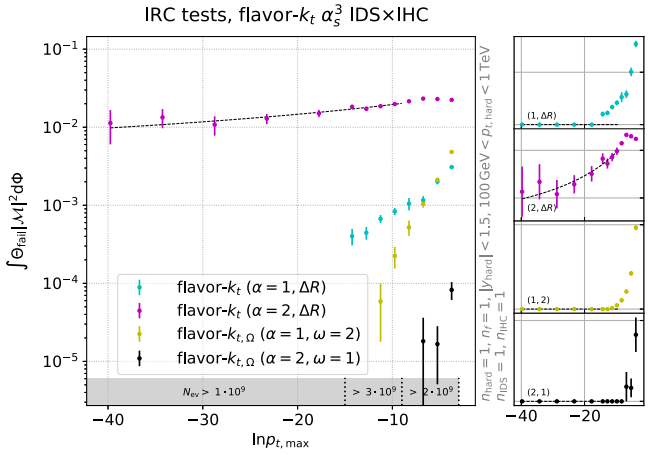

FIG. 16. Failure rate of the flavor-k_t algorithm for the configuration of Fig. 15, in particular illustrating results for $\alpha = 2$ (magenta points) that are qualitatively consistent with the expected $1/\ln p_{t,\max}$ behavior (dashed line). Also shown are results for $\alpha = 1$, as well as the results for flavor-$k_{t,\Omega}$, i.e., the adaptation with an Ω_{ij} angular distance, illustrating the much faster drop of the failure rate.

$$N_{24,\text{flav-}k_t}^{(\ell_{42}>0)} \sim \alpha_s^3 \int^\infty \frac{d\ell_{14}}{\ell_{14}^2}. \tag{C13}$$

This and Eq. (C11) both converge in the infrared, i.e., for $\ell_{14} \to \infty$; however, this convergence is extremely slow. In particular, if one places an upper limit $p_{t4} > \epsilon$, the result converges as $(\ln 1/\epsilon)^{-1}$, which is consistent also with what we find in our numerical tests; cf. the magenta ($\alpha = 2$) points in Fig. 16. While this is strictly IRC safe at this order, one should worry that at the next order there may be logarithmic enhancements proportional to ℓ_{14} (for example from running-coupling effects), which would be sufficient to make the integral diverge. Accordingly, it would seem wise for future uses of the flavor-k_t algorithm to adopt the same kind of $\Delta R_{ij}^2 \to \Omega_{ij}^2$ replacement as used in our IFN algorithm, and similarly for any other algorithms that make use of similarly defined distances, e.g., the GHS flavor-dressing algorithm.

A final comment concerns the $\alpha = 1$ case. The analysis is somewhat more involved than for $\alpha = 2$, and it is also clear from Fig. 16 that the issue is reduced with $\alpha = 1$. Our investigations are consistent with a $1/\ln^p \epsilon$ scaling, with a larger value of p than for the $\alpha = 2$ case. One might wish to investigate this point further; however, it would anyway seem wise to use the $\Delta R_{ij}^2 \to \Omega_{ij}^2$ replacement also for $\alpha = 1$.

2. IHC^2 issue at α_s^2 for CMP

An issue arises in the CMP algorithm at order α_s^2 for a configuration like the one shown in Fig. 17. We consider two initial-state hard-collinear emissions, a q and a $\bar{q}$

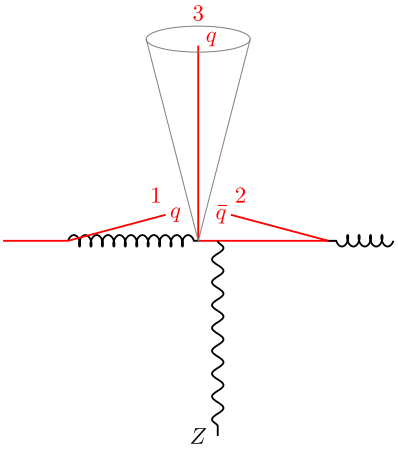

FIG. 17. Example $\mathcal{O}(\alpha_s^2)$ configuration that yields an issue for the CMP algorithm. There are two oppositely flavored initial-state hard-collinear splittings (b and $\bar{b}$, labeled 1 and 2) and a hard particle 3 at central rapidity.

(labeled 1 and 2), from the forward and backward beams, respectively, and additionally one hard large-angle particle (numbered 3 in the figure). The initial-state hard-collinear emissions have a very small transverse momentum ($p_{t1}, p_{t2} \ll p_{t3}$) but large energies ($E_1, E_2 \sim E_3$). Let us assume $p_{t3} = E_3 = 1$, and particle 3 is simply aligned along the x axis. Then, we have that $y_1 \sim -\ln p_{t1}$ and $y_2 \sim \ln p_{t2}$ in the hard-collinear limit. For simplicity, we will work with $R = 1$.

The CMP algorithm will strongly favor clustering 1 and 2 together first. The global scale [$p_{t,\text{global-max}}$ in Eq. (5)] is set by the p_t of the hardest pseudojet currently available, p_{t3}, so the value of κ_{12} is small,

$$\kappa_{12} = \frac{1}{2a}\frac{p_{t1}^2 + p_{t2}^2}{p_{t3}^2} \ll 1, \tag{C14}$$

and the distance between the oppositely flavored particles 1 and 2 is thus given by

$$d_{12} = \frac{1}{\max(p_{t1}^2, p_{t2}^2)}\Delta R_{12}^2\left(1 - \cos\left(\frac{\pi}{2}\kappa_{12}\right)\right) \tag{C15a}$$

$$\simeq \frac{1}{\max(p_{t1}^2, p_{t2}^2)}\Delta R_{12}^2\frac{1}{2}\left(\frac{\pi}{2}\kappa_{12}\right)^2 \tag{C15b}$$

$$\simeq \frac{\pi^2}{32a^2}\Delta R_{12}^2\frac{\max(p_{t1}^2, p_{t2}^2)}{p_{t3}^4}, \tag{C15c}$$

$$\simeq \frac{\pi^2}{32a^2}\left(\ln\frac{p_{t3}}{p_{t1}} + \ln\frac{p_{t3}}{p_{t2}}\right)^2\frac{\max(p_{t1}^2, p_{t2}^2)}{p_{t3}^4}, \tag{C15d}$$

where ΔR_{12} is dominated by the large rapidity difference. The other distances are

$$d_{iB} = \frac{1}{p_{ti}^2}, \qquad i = \{1, 2, 3\}, \tag{C16a}$$

$$d_{i3} \sim \frac{y_i^2}{p_{t3}^2}, \qquad i = \{1, 2\}. \tag{C16b}$$

When $p_{t1}, p_{t2} \ll p_{t3}$, it is straightforward to see that $d_{12} < d_{1B}, d_{2B}$ and $d_{12} < d_{13}, d_{23}$ [the logarithms in Eq. (C15d) have no impact on this]. Therefore, the first step of the algorithm will be to cluster particles 1 and 2, giving a flavorless pseudojet with transverse momentum, rapidity, and squared invariant mass of

$$p_{t,(1+2)} \sim \max(p_{t1}, p_{t2}), \tag{C17a}$$

$$y_{(1+2)} \simeq \frac{1}{2}\ln\frac{E_1}{E_2}, \tag{C17b}$$

$$m_{(1+2)}^2 \simeq 4E_1E_2. \tag{C17c}$$

From the point of view of standard jet clustering, the $(1 + 2)$ pseudojet is unusual because its transverse momentum is much smaller than its invariant mass.

The $(1 + 2)$ pseudojet will cluster with particle 3 if it is within a distance $\Delta R_{(1+2),3} < R = 1$. For any $p_{t1}, p_{t2} \ll p_{t3}$, there is always a finite azimuthal and E_1, E_2 phase space region such that that condition is satisfied, and the resulting $1 + 2 + 3$ cluster will have significantly different kinematics than particle 3, because of the extra energy brought by particles 1 and 2. Thus, the jets can differ between the 1,2,3 event and the event with just particle 3. The rate for this to happen is given by constants from the azimuthal and energy integrations multiplying divergent integrals over p_{t1} and p_{t2},

$$N \sim \alpha_s^2\int_\epsilon^{p_{t3}}\frac{dp_{t1}}{p_{t1}}\int_\epsilon^{p_{t3}}\frac{dp_{t2}}{p_{t2}} = \alpha_s^2\ln^2\epsilon, \tag{C18}$$

where we have explicitly included a cutoff scale ϵ in order to make the nature of the divergence manifest.

In Sec. IV D, we proposed a modification of the CMP algorithm, Eq. (11). For a generic ω, Eq. (C15d) in particular is replaced by

$$d_{12}^{(\Omega)} \sim \left(\frac{p_{t3}^2}{p_{t1}p_{t2}}\right)^\omega\frac{\max(p_{t1}^2, p_{t2}^2)}{p_{t3}^4}, \tag{C19a}$$

$$\sim \frac{\max(p_{t1}^2, p_{t2}^2)}{p_{t1}p_{t2}p_{t3}^2}\cdot\left(\frac{p_{t3}^2}{p_{t1}p_{t2}}\right)^{\omega-1}. \tag{C19b}$$

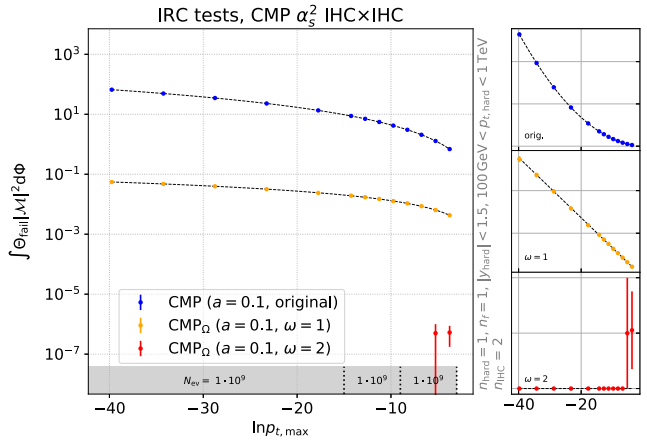

FIG. 18. Results of the numerical tests for CMP for a configuration as in Fig. 17, with the ΔR^2 angular factor in the distance measure, as in the original algorithm (in blue), and with a corrected angular factor as in Eq. (11) where $\omega = 1$ (in orange), and $\omega = 2$ (in red). Miniature plots on the right depict the integrated failure rate on a linear scale to help read the functional form of the divergence. The original algorithm suffers from a quadratic divergence, the corrected measure with $\omega = 1$ from a linear divergence, and the fix with $\omega = 2$ is IRC safe.

When $\omega = 1$, $d_{12}^{(\Omega)}$ will be of same order as d_{3B} if $p_{t1} \sim p_{t2}$. Schematically, that suggests that there can still be a divergence of the form

$$N \sim \alpha_s^2 \int_\epsilon^{p_{t3}} \frac{dp_{t1}}{p_{t1}} \int_\epsilon^{p_{t3}} \frac{dp_{t2}}{p_{t2}} \delta(\ln p_{t2} - \ln p_{t1}) = \alpha_s^2 \ln \epsilon, \tag{C20}$$

i.e., with one power of the logarithm. With a more complete calculation, one can verify that that divergence is indeed present for $\omega = 1$. For $\omega > 1$, the second factor in Eq. (C19b) instead ensures that $d_{12}^{(\Omega)} \gg d_{3B}$, thus ensuring particle 3 becomes a jet before any $1+2$ clustering and resolving the IRC safety issue with the kinematics of jet 3.[17]

Results of the numerical IRC safety tests are presented in Fig. 18 for the original algorithm (with an angular factor ΔR^2 in the distance measure, in blue) and for the algorithm with an IRC-safe angular distance (for values of the parameter $\omega = 1$, in orange, and $\omega = 2$, in red). The results confirm the above analyses, showing a quadratic divergence in the original algorithm, while the divergence is linear for $\omega = 1$ and fully lifted for $\omega = 2$.[18]

Finally, note that the flavor-k_t algorithm does not suffer from the issue presented here thanks to the form of its beam distance: the initial-state emissions 1 and 2 would be declared as beam jets and so be removed from further consideration early in the clustering sequence.

3. IHC × IDS issue at α_s^3 for CMP

The subtlety from Appendix C 1 for the flavor-k_t and GHS algorithms has an interesting manifestation in the CMP algorithm. We consider again the scenario of Fig. 15, with the same set of variables and in a configuration where $p_{t4} \ll p_{t2}$. We will concentrate on the two distances that are smallest, which, neglecting $\mathcal{O}(1)$ factors, read

$$d_{23} \sim \frac{p_{t2}^2}{z_{23}^2 p_{t1}^4}, \qquad d_{24} \sim \frac{p_{t2}^2}{p_{t1}^4} y_4^2. \tag{C21}$$

As before, the probability for a $2+4$ recombination must be finite for IRC safety. The $2+4$ recombination will occur if

$$z_{23} \lesssim \frac{1}{y_4} = \frac{1}{\ell_{14} + \ln 2 z_{41}}. \tag{C22}$$

We neglect the $\ln 2 z_{41}$ term in the denominator (and take the z_{41} integral to give a constant), integrate over all p_{t4} values, over $p_{t2} > p_{t4}$ and over the allowed z_{23} range (with the same constant splitting function approximation as in Appendix C 1). We then obtain the probability for a $2+4$ clustering,

$$N_{24} \sim \alpha_s^3 \int_0^\infty d\ell_{14} \int_0^{\ell_{14}} d\ell_{24} \int_0^{1/\ell_{14}} dz_{23}, \tag{C23}$$

which is divergent. If we regulate the upper integration region of the ℓ_{14} integral with $\infty \to \ln \frac{1}{\epsilon}$, the probability scales as

$$N_{24} \sim \alpha_s^3 \ln \frac{1}{\epsilon}. \tag{C24}$$

[17]Once particle 3 has been declared a jet and removed from the clustering, the $1+2$ cluster, if formed, would become a jet in its own right. Standard jet analyses place a cut on the jet p_t, which ensures that a residual lone massive, but low-p_t $1+2$, cluster does not count as a hard jet. However, one could imagine a scenario where one cuts not on p_t but on $p_t^2 + m^2$, and in this case, the $1+2$ cluster would count as an additional hard jet. In such a case, one could envisage that this would cause IRC unsafety even for $\omega > 1$. We have not explored this question further, insofar as analyses do not normally cut on $p_t^2 + m^2$.

[18]Note that Fig. 1 (left) of Ref. [10] has studied $pp \to \ell^+\ell^- b\bar{b}\bar{b}$ (with massless b's), which should include the configuration of our Fig. 17. That figure does not appear to show a divergence as the technical cutoff is reduced. At first sight, that may seem surprising; however, two considerations should be kept in mind. First, the configuration of Fig. 17 requires an initial-state gluon to originate from a $b \to gb$ splitting, which is responsible for only a small fraction of incoming gluons. Second, the impact of the IRC unsafety is to smear the rapidity distribution of the b jet, and a smearing of broad distribution tends to have a limited impact on the integral of the broad distribution within some window (Ref. [10] used $|y_b| < 2.4$). As a result, it is conceivable that the expected squared logarithmic divergence in Fig. 1 (left) of Ref. [10] might be too small to clearly see in that figure.

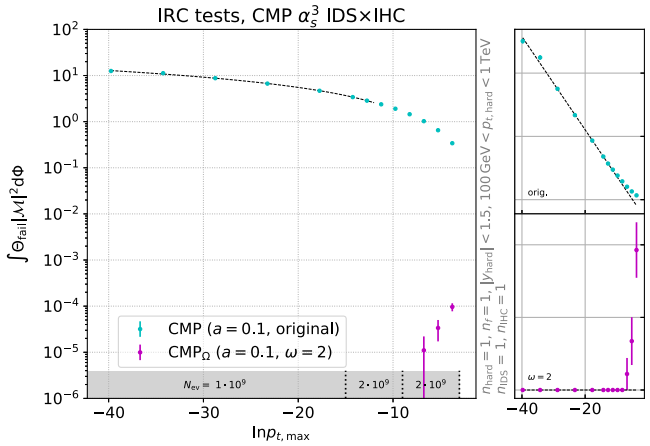

FIG. 19. Failure rate of the CMP algorithm for the configuration of Fig. 15, both for the original formulation and our modification, Eq. (11), showing a divergence for the former and none for the latter.

As with the other occurrences of this kind of issue, the replacement Eq. (11) solves the problem, as can be seen from the numerical IRC safety tests for this configuration shown in Fig. 19.

4. FHC^2 issue at α_s^2 for GHS

For the discussion here and in Appendix C 5, we assume a version of the GHS algorithm with an Ω_{ij}^2 style angular distance in the dressing (flavor-k_t-like) phase, since we know from Appendix C 1 that this is required for the flavor-k_t algorithm, which is the basis of the algorithm's dressing step. To make this clear, in plots where this modification is used, we refer to the algorithm as GHS$_\Omega$.

Let us consider a hard event as that in Fig. 20. The event has four particles, which we will call g_1, q, $\bar{q}$, and g_2, starting from the left. We will work through the algorithm to see what happens if a hard, but collinear gluon g, is emitted from the q and splits collinearly to $q'\bar{q}'$ (with $\Delta R_{q'\bar{q}'} \ll 1$ much smaller than any other scale in the problem).

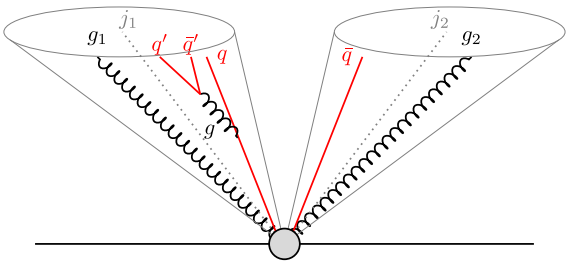

FIG. 20. Example $\mathcal{O}(\alpha_s^2)$ configuration that yields an issue for the GHS algorithm. There are four hard particles (that one can imagine recoiling against a hard gluon or electroweak system on the other side of the event), a collinear emission of a hard gluon g from one of the flavored particles (the q), which then splits collinearly to a flavored pair $q'\bar{q}'$.

Focusing on the hard event first (i.e., without the emission of the hard collinear gluon g), we assume that the anti-k_t algorithm clusters the four particles into two jets (j_1 and j_2), as indicated in the figure. We can further assume that $\Delta R_{g_1 q}, \Delta R_{g_2 \bar{q}} > R_{\text{cut}} \sim 0.1$, so that the hard gluons g_1 and g_2 are not accumulated into q and $\bar{q}$ in that phase of the algorithm.

First, we will consider the case $\alpha < 2$. For any angular structure of the event satisfying the above limits, we take the momenta of g_1, g_2, q, and $\bar{q}$ such that the event without the $g \to q'\bar{q}'$ emission has the following properties:

$$d_{qj_1} > d_{q\bar{q}} > d_{\bar{q}j_2}, \tag{C25a}$$

$$p_{tq} < p_{t\bar{q}}. \tag{C25b}$$

As a result, the first dressing step is for the $\bar{q}$ flavor to be assigned to jet j_2, followed by the q flavor being assigned to jet j_1. Thus, both j_1 and j_2 are flavored. Note that for a full analysis one should also take into account d_{iB} beam distances for all flavored particles i. To help understand why we can ignore it, suppose that all the hard particles have rapidities close to zero, which results in $p_{tB}(0)$ in Eq. (C4) being approximately the scalar sum of all the particles' transverse momenta. That scale will tend to be a few times larger than the transverse momenta of any of the individual particles, which ensures that the distance of any cluster to its jet will be smaller than the d_{iB}, as will the $d_{q\bar{q}}$ if the two jets are not too far away in angle.

Next, we consider the impact of the emission of the collinear hard gluon from q with $p_{tg} = z p_{tq}$, followed by its splitting into a collinear $q'\bar{q}'$ pair. Recall that we work with $\Delta R_{q'\bar{q}'} \ll R_{\text{cut}}$ so that it is the smallest angular distance in the event. The algorithm goes through the accumulation step and will identify four flavor clusters: $\hat{q}'$, $\hat{\bar{q}}'$, and the original $\hat{q}$ and $\hat{\bar{q}}$. The angular structure is otherwise unchanged, so we get no further flavor accumulation. To lighten the notation, below we will leave out the explicit "hats" for the flavor clusters, especially as the flavor clusters coincide with the original particles.

The final step is the flavor dressing: the $q', \bar{q}'$ pair will annihilate first, as it should because the pair came from a common parent gluon. These flavor clusters (including their kinematics) are discarded from further consideration, and any distance involving them is removed from the list. The remaining distances (d') for the event with the $g \to q'\bar{q}'$ splitting are then given in terms of the hard event's distances (d) as

$$d'_{qj_1} = (1-z)^{2-\alpha} d_{qj_1}, \tag{C26a}$$

$$d'_{q\bar{q}} = (1-z)^{2-\alpha} d_{q\bar{q}}, \tag{C26b}$$

$$d'_{\bar{q}j_2} = d_{\bar{q}j_2}, \tag{C26c}$$

where the $(1-z)^{2-\alpha}$ factor arises because of the reduction in transverse momentum of the q after emission of the $g \to q'\bar{q}'$ (which carries a fraction z of its original q momentum). The potentially dangerous scenario is that where the ordering of distances, Eq. (C25a), is modified,

$$d'_{q\bar{q}} < \min\left(d'_{qj_1}, d'_{\bar{q}j_2}\right), \tag{C27}$$

because then q and $\bar{q}$'s flavors will annihilate, leaving flavorless hard jets, associated with a squared logarithmic divergence from the two nested hard collinear divergences. There is a finite range of z in which this occurs,

$$1 - z < \left(\frac{d_{\bar{q}j_2}}{d_{q\bar{q}}}\right)^{\frac{1}{2-\alpha}}, \tag{C28}$$

thus confirming the presence of IRC unsafety from the configuration of Fig. 20 for $\alpha < 2$.

When $\alpha = 2$, we instead consider a hard event satisfying $p_{tq} > p_{t\bar{q}}$ rather than the inequality in Eq. (C25b), in which case we have

$$d'_{qj_1} = d_{qj_1}, \tag{C29a}$$

$$d'_{q\bar{q}} = \max\left((1-z)^2, \frac{p^2_{t\bar{q}}}{p^2_{tq}}\right) d_{q\bar{q}}, \tag{C29b}$$

$$d'_{\bar{q}j_2} = d_{\bar{q}j_2}. \tag{C29c}$$

Again, there is the possibility of $d'_{q\bar{q}}$ becoming the smallest of the three distances, with the outcome that the q and $\bar{q}$ flavors would annihilate, leaving flavorless hard jets, with a squared logarithmic divergence associated with the collinear splittings.

The set of distances in the argument above is perhaps somewhat complicated, with angular factors to consider, the beam distances and the extra subtleties of the $\alpha = 2$ case. Therefore, in Fig. 21, we show the outcome of our IRC safety tests, illustrating that the divergence is indeed present for the two combinations $\alpha = 1, \beta = 1$ and $\alpha = 2, \beta = 2$. We leave to future work the possibility of identifying a concrete modification of the algorithm that solves this problem; nevertheless, we anticipate that one line of investigation could be to allow accumulation of kinematics within a jet during the dressing stage.

A final comment is that this configuration can appear at NNLO for a process such as fully hadronic $t\bar{t}$ production, however only if one asks for two massive b-tagged jets. It also appears at N^{4}LO for a process such as $Zb\bar{b}$ production.

5. IDS × FDS issue at α_s^4 for GHS

The GHS algorithm exhibits an interesting interplay between initial-state and final-state double-soft emissions at order α_s^4 if $\alpha\beta \geq 2$. The configuration that we consider here is that represented in Fig. 22, involving a hard Born event with one or more unflavored jets. That event is then supplemented with a double-soft pair (1, 2) that is collinear to an (originally) unflavored jet and an additional large-angle double-soft pair (3, 4) outside the jet.

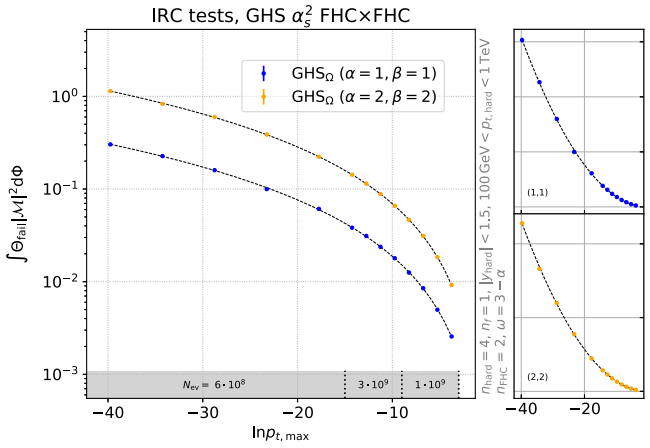

FIG. 21. Failure rate of the GHS algorithm for the FHC × FHC configuration of Fig. 20, illustrating the quadratic divergence, specifically for $\alpha = 1, \beta = 1$ and $\alpha = 2, \beta = 2$. Other parameters are $R_{\text{cut}} = 0.1$, $z_{\text{cut}} = 0.1$, and $p_{t,\text{cut}} = 100$ GeV. The jet radius R has been sampled in the range 0.3–1.57. The version of the GHS algorithm used is one where ΔR_{ij}^2 in the dressing stage has been replaced with Ω_{ij}^2 using $\omega = 3 - \alpha$ (the original ΔR_{ij}^2 similarly gives a squared logarithmic divergence).

We are specifically interested in the situation where

$$\theta_{1g} \lesssim \{\theta_{2g}, \theta_{12}\} \ll 1, \tag{C30a}$$

$$p_{t1} \sim p_{t2} \ll p_{tg}, \tag{C30b}$$

$$p_{t3} \sim p_{t4} \ll p_{tg}, \tag{C30c}$$

$$\theta_{23} \sim \theta_{34} \sim 1. \tag{C30d}$$

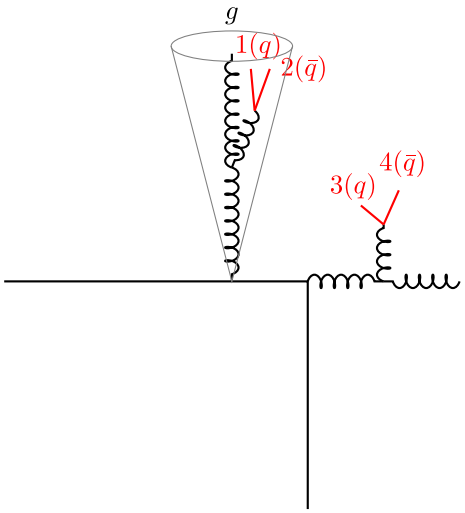

FIG. 22. An FDS × IDS kinematic configuration that causes problems for GHS algorithms for $\alpha\beta \geq 2$.

During the accumulation step, there is a possibility that 1 clusters with g, giving a hard $\widehat{g1}$ flavor cluster, leaving an unclustered, much softer $\hat{2}$ flavor cluster. If during the subsequent dressing phase, $\hat{2}$ goes on to annihilate with $\hat{3}$ rather than with $\widehat{g1}$, then the resulting hard jet will be flavored.

The SoftDrop condition for 1 to cluster with g is given by

$$\frac{p_{t1}}{p_{tg}} > \theta_{1g}^{\beta}, \tag{C31}$$

where throughout our discussion here we neglect factors of order 1 (e.g., R_{cut} and z_{cut}). There is no further accumulation since all particles are now flavored and the flavored clusters will be $\widehat{g1}$, $\hat{2}$, $\hat{3}$, and $\hat{4}$. The angle between the $\widehat{g1}$ and the jet direction will be given by

$$\theta_{j\widehat{g1}} \sim \frac{p_{t2}}{p_{tg}}\theta_{2g}. \tag{C32}$$

Without loss of generality, we can consider the case where $d_{\hat{2}\hat{3}} < d_{\hat{2}\hat{4}}$, and then the distances to take into account during the dressing phase are

$$d_{j\widehat{g1}} \sim p_{t2}^2\theta_{2g}^2, \tag{C33a}$$

$$d_{g\widehat{1}\hat{2}} \sim d_{j\hat{2}} \sim p_{tg}^2\left(\frac{p_{t2}}{p_{tg}}\right)^{2-\alpha}\theta_{2g}^2, \tag{C33b}$$

$$d_{\hat{2}\hat{3}} \sim \max(p_{t2}, p_{t3})^{\alpha}\min(p_{t2}, p_{t3})^{2-\alpha}, \tag{C33c}$$

$$d_{\hat{3}\hat{4}} \sim p_{t3}^2, \tag{C33d}$$

where, again we have ignored factors of order 1, e.g., from angular distances. The hard jet will acquire a flavor if the SoftDrop condition of Eq. (C31) is satisfied and if additionally $\hat{2}$ fails to annihilate the flavor of the $\widehat{g1}$ cluster. This will occur if $d_{23} < d_{34}$ and $d_{23} < d_{j\hat{2}} \simeq d_{g\widehat{12}}$.[19] In determining whether these conditions are satisfied, it is helpful to introduce shorthands

$$\ell_i = \ln\frac{p_{tg}}{p_{ti}}, \tag{C34a}$$

$$\ell_\theta = \ln\frac{1}{\theta_{2g}} \tag{C34b}$$

and to observe that in the SoftDrop condition (C31), we can replace $1 \to 2$, since this only affects $\mathcal{O}(1)$ terms.

[19]Note that the $d_{23} < d_{34}$ condition implies that p_{t2} cannot be substantially larger than p_{t3}, which leads to $d_{j\widehat{g1}}$ being the smallest of all the distances. Consequently, the first step of the dressing is that the flavor of $\widehat{g1}$ is assigned to the jet and the $\widehat{g1}$ cluster is removed from consideration.

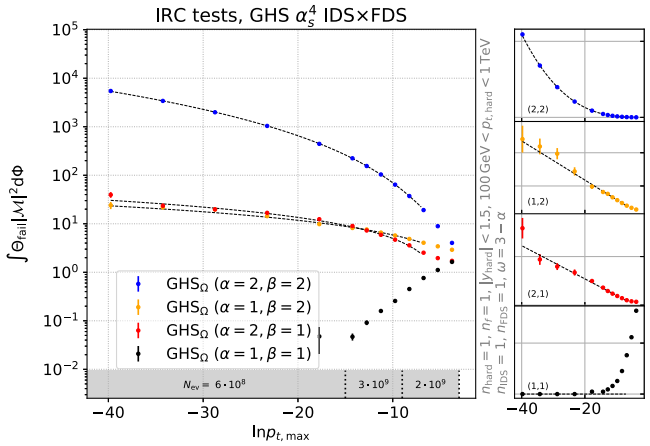

FIG. 23. Failure rate of the GHS algorithm for the α_s^4 FDS × IDS configuration of Fig. 22, illustrating the cubic divergence for parameter choices involving $\alpha\beta > 2$, the linear divergence for $\alpha\beta = 2$, and convergence for $\alpha\beta < 2$. In its dressing stage, the GHS implementation for these runs uses an Ω_{ij}^2 angular distance instead of ΔR_{ij}^2, where $\omega = 3 - \alpha$. Other parameters are $R_{\text{cut}} = 0.1$, $z_{\text{cut}} = 0.1$, and $p_{t,\text{cut}} = 100$ GeV.

Our conditions then become [to within $\mathcal{O}(1)$ offsets]

$$\text{SD:}\ \ell_2 < \beta\ell_\theta, \tag{C35a}$$

$$d_{34} > d_{23}:\ \ell_2 > \ell_3, \tag{C35b}$$

$$d_{g2} > d_{23}:\ (2-\alpha)\ell_2 + 2\ell_\theta < \alpha\ell_3 + (2-\alpha)\ell_2, \tag{C35c}$$

where the last line already underwent some simplification (using the second line) and can then be further simplified to read

$$2\ell_\theta < \alpha\ell_3. \tag{C36}$$

Assembling all inequalities, we obtain

$$2\ell_\theta < \alpha\ell_3 < \alpha\ell_2 < \alpha\beta\ell_\theta. \tag{C37}$$

We immediately see that if $\alpha\beta < 2$ there is no available logarithmic integration region, and so no IRC divergence from this configuration. Conversely, if $\alpha\beta > 2$, we expect to see a cubic logarithmic divergence from integrals over ℓ_θ, ℓ_3, and ℓ_2. For $\alpha\beta = 2$, the $\mathcal{O}(1)$ factors become critical, and it is easiest to carry out a numerical study, but it is reasonable to expect a divergence with a single logarithm.

The results of the numerical study are shown in Fig. 23 for four combinations of α and β. They confirm our expectations and suggest that if one wishes to employ a GHS-style algorithm one should use it with $\alpha\beta < 2$. Nevertheless, one would still need to find a solution to the separate issue identified in Appendix C 4 (which cannot be resolved just through parameter choices) and then verify that the resulting algorithm passes a full set of IRC safety tests.

APPENDIX D: SUMMARY PLOTS FOR IRC-SAFE ALGORITHMS

In this appendix, we present summary plots from our IRC safety tests for the three approaches that have passed all those tests: IFN, CMP_Ω, and flavor-$k_{t,\Omega}$. In Figs. 24 and 25, we show the results from the IRC-safety tests for the anti-k_t and C/A algorithms with IFN, i.e., each of the algorithms labeled as safe in Table I. Figures 26 and 27 show corresponding results for our adaptations of the flavor-k_t and CMP algorithms, supporting the conclusion that they, too, are IRC safe.

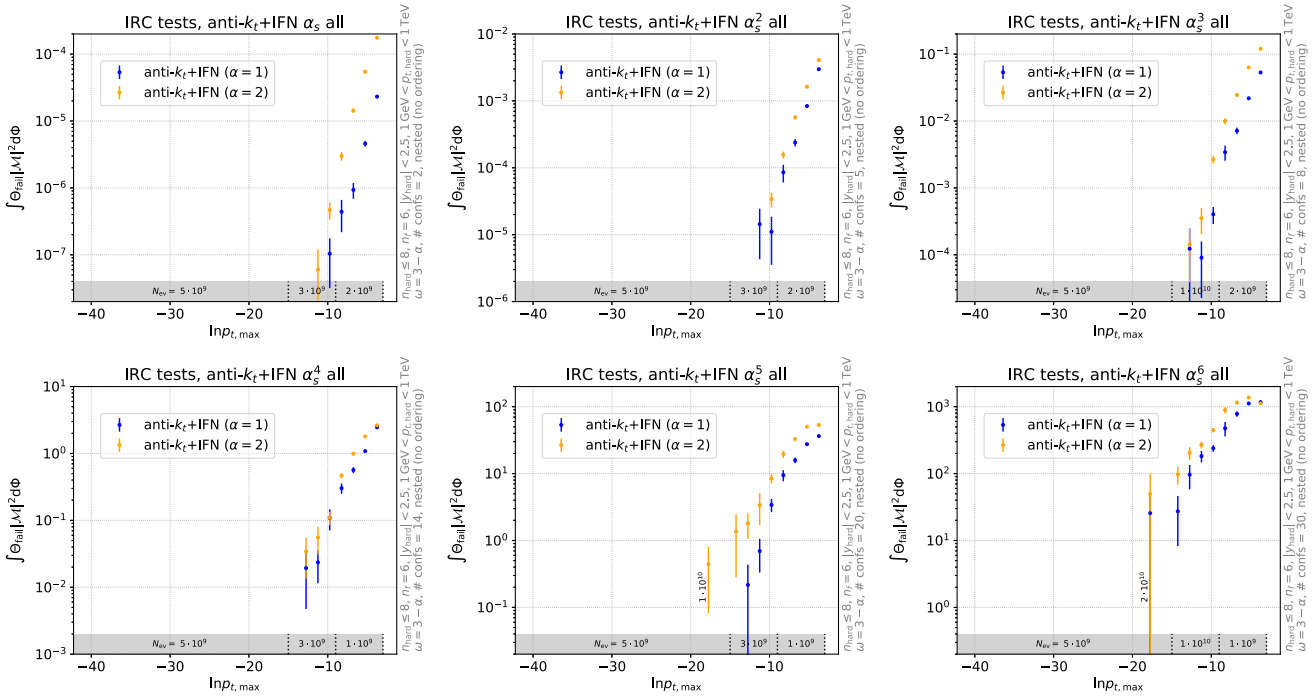

FIG. 24. Summary of IRC safety test results at orders α_s to α_s^6 for the anti-k_t algorithm with IFN.

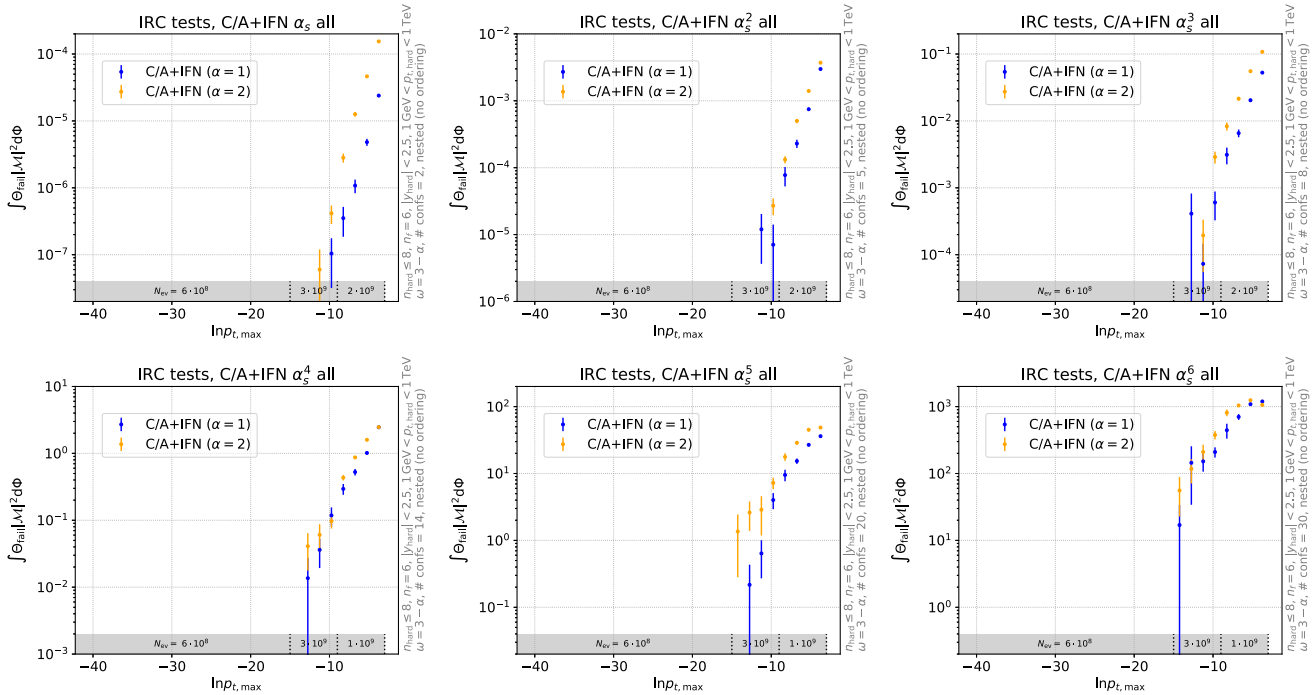

FIG. 25. Same as Fig. 24, for the C/A algorithm with IFN.

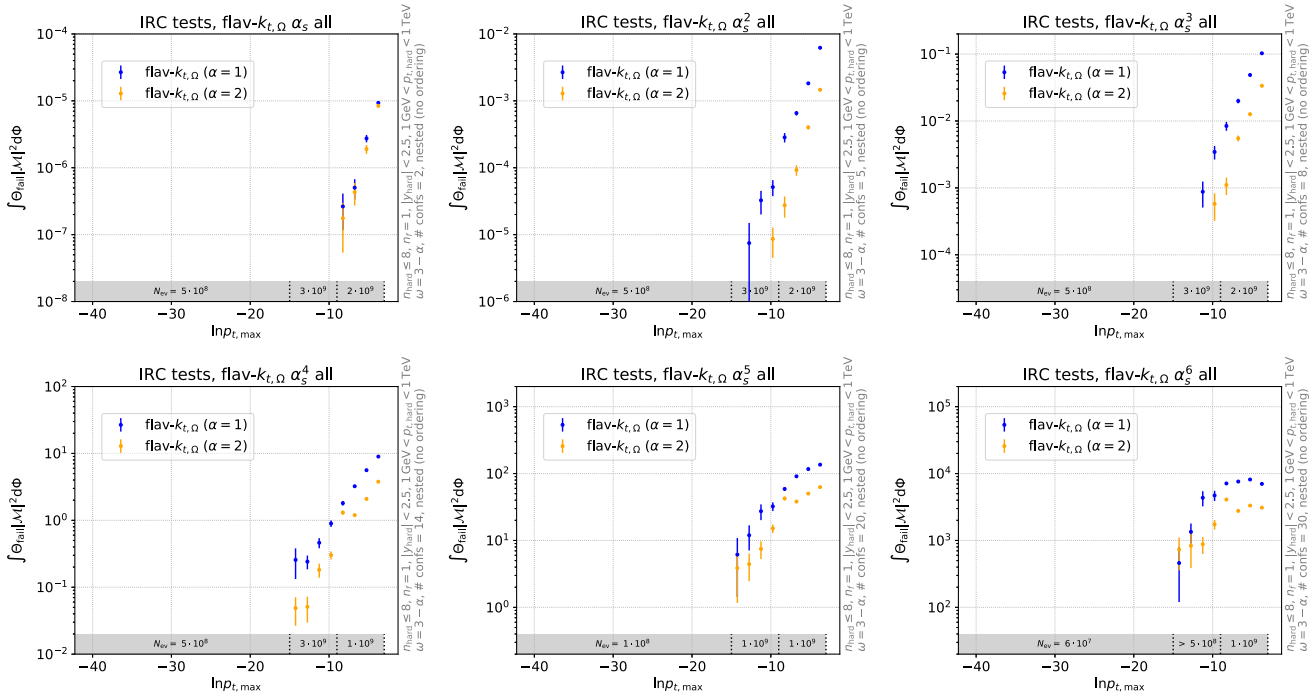

FIG. 26. Same as Fig. 24, for the flavor-$k_{t,\Omega}$ algorithm.

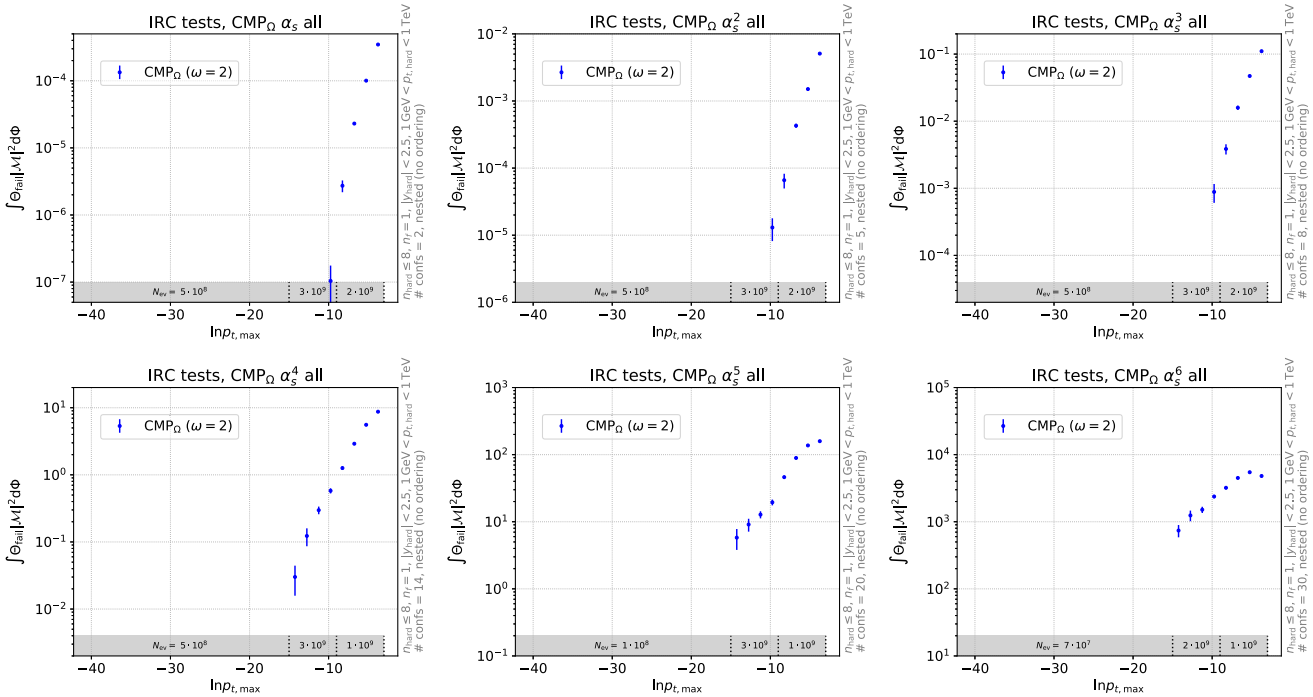

FIG. 27. Same as Fig. 24, for the CMP$_\Omega$ algorithm.

[1] A. Banfi, G. P. Salam, and G. Zanderighi, Infrared safe definition of jet flavor, Eur. Phys. J. C **47**, 113 (2006).

[2] A. Banfi, G. P. Salam, and G. Zanderighi, Accurate QCD predictions for heavy-quark jets at the Tevatron and LHC, J. High Energy Phys. 07 (2007) 026.

[3] S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock, and B. R. Webber, New clustering algorithm for multi-jet cross-sections in e^+e^- annihilation, Phys. Lett. B **269**, 432 (1991).

[4] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, Longitudinally invariant K_t clustering algorithms for hadron hadron collisions, Nucl. Phys. **B406**, 187 (1993).

[5] S. D. Ellis and D. E. Soper, Successive combination jet algorithm for hadron collisions, Phys. Rev. D **48**, 3160 (1993).

[6] M. Cacciari, G. P. Salam, and G. Soyez, The anti-k_t jet clustering algorithm, J. High Energy Phys. 04 (2008) 063.

[7] R. Gauld, A. Gehrmann-De Ridder, E. W. N. Glover, A. Huss, and I. Majer, Predictions for Z-boson production in association with a b-jet at $\mathcal{O}(\alpha_s^3)$, Phys. Rev. Lett. **125**, 222002 (2020).

[8] S. Caletti, A. J. Larkoski, S. Marzani, and D. Reichelt, Practical jet flavour through NNLO, Eur. Phys. J. C **82**, 632 (2022).

[9] S. Caletti, A. J. Larkoski, S. Marzani, and D. Reichelt, A fragmentation approach to jet flavor, J. High Energy Phys. 10 (2022) 158.

[10] M. Czakon, A. Mitov, and R. Poncelet, Infrared-safe flavoured anti-k_T jets, J. High Energy Phys. 04 (2023) 138.

[11] R. Gauld, A. Huss, and G. Stagnitto, Flavor identification of reconstructed hadronic jets, Phys. Rev. Lett. **130**, 161901 (2023).

[12] S. Weinzierl, The forward-backward asymmetry at NNLO revisited, Phys. Lett. B **644**, 331 (2007).

[13] Z. Trócsányi, G. Somogyi, and F. Tramontano, Fully differential decay rate of a standard model Higgs boson into two b-jets at NNLO, Acta Phys. Pol. B **46**, 2097 (2015).

[14] G. Ferrera, G. Somogyi, and F. Tramontano, Associated production of a Higgs boson decaying into bottom quarks at the LHC in full NNLO QCD, Phys. Lett. B **780**, 346 (2018).

[15] F. Caola, G. Luisoni, K. Melnikov, and R. Röntsch, NNLO QCD corrections to associated WH production and $H \to b\bar{b}$ decay, Phys. Rev. D **97**, 074022 (2018).

[16] R. Gauld, A. Gehrmann-De Ridder, E. W. N. Glover, A. Huss, and I. Majer, Associated production of a Higgs boson decaying into bottom quarks and a weak vector boson decaying leptonically at NNLO in QCD, J. High Energy Phys. 10 (2019) 002.

[17] M. Czakon, A. Mitov, M. Pellen, and R. Poncelet, NNLO QCD predictions for W+c-jet production at the LHC, J. High Energy Phys. 06 (2021) 100.

[18] H. B. Hartanto, R. Poncelet, A. Popescu, and S. Zoia, Next-to-next-to-leading order QCD corrections to $Wb\bar{b}$ production at the LHC, Phys. Rev. D **106**, 074016 (2022).

[19] H. B. Hartanto, R. Poncelet, A. Popescu, and S. Zoia, Flavour anti-k_T algorithm applied to $Wb\bar{b}$ production at the LHC, arXiv:2209.03280.

[20] M. Czakon, A. Mitov, M. Pellen, and R. Poncelet, A detailed investigation of $W + c$-jet at the LHC, J. High Energy Phys. 02 (2023) 241.

[21] R. Gauld, A. Gehrmann-De Ridder, E. W. N. Glover, A. Huss, A. R. Garcia, and G. Stagnitto, NNLO QCD predictions for Z-boson production in association with a charm jet within the LHCb fiducial region, Eur. Phys. J. C **83**, 336 (2023).

[22] Y. L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber, Better jet clustering algorithms, J. High Energy Phys. 08 (1997) 001.

[23] M. Wobisch and T. Wengler, Hadronization corrections to jet cross-sections in deep inelastic scattering, in *Proceedings of the Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)* (DESY, Hamburg, Germany, 1998), pp. 270–279, arXiv:hep-ph/9907280.

[24] S. Hoeche, F. Krauss, S. Schumann, and F. Siegert, QCD matrix elements and truncated showers, J. High Energy Phys. 05 (2009) 053.

[25] A. Karlberg, G. P. Salam, L. Scyboz, and R. Verheyen, Spin correlations in final-state parton showers and jet observables, Eur. Phys. J. C **81**, 681 (2021).

[26] M. van Beekveld, S. Ferrario Ravasio, K. Hamilton, G. P. Salam, A. Soto-Ontoso, G. Soyez, and R. Verheyen, PanScales showers for hadron collisions: All-order validation, J. High Energy Phys. 11 (2022) 020.

[27] G. P. Salam and G. Soyez, A practical seedless infrared-safe cone jet algorithm, J. High Energy Phys. 05 (2007) 086.

[28] J. Gallicchio and Y.-T. Chien, Quit using pseudorapidity, transverse energy, and massless constituents, arXiv:1802.05356.

[29] A. Buckley and C. Pollard, QCD-aware partonic jet clustering for truth-jet flavour labelling, Eur. Phys. J. C **76**, 71 (2016).

[30] W. Bartel *et al.* (JADE Collaboration), Experimental studies on multi-jet production in e^+e^- annihilation at PETRA energies, Z. Phys. C **33**, 23 (1986).

[31] S. Bethke *et al.* (JADE Collaboration), Experimental investigation of the energy dependence of the strong coupling strength, Phys. Lett. B **213**, 235 (1988).

[32] A. Behring, W. Bizoń, F. Caola, K. Melnikov, and R. Röntsch, Bottom quark mass effects in associated WH production with the $H \to b\bar{b}$ decay through NNLO QCD, Phys. Rev. D **101**, 114012 (2020).

[33] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, Soft Drop, J. High Energy Phys. 05 (2014) 146.

[34] B. Andersson, G. Gustafson, L. Lonnblad, and U. Pettersson, Coherence effects in deep inelastic scattering, Z. Phys. C **43**, 625 (1989).

[35] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, Eur. Phys. J. C **72**, 1896 (2012).

[36] Y. Hida, X. S. Li, and D. H. Bailey, Quad-double arithmetic: Algorithms, implementation, and application, in *Proceedings of the 15th IEEE Symposium on Computer Arithmetic* (IEEE Computer Society, Washington DC, USA, 2000), pp. 155–162.

[37] M. Dasgupta, F. A. Dreyer, K. Hamilton, P. F. Monni, G. P. Salam, and G. Soyez, Parton showers beyond leading logarithmic accuracy, Phys. Rev. Lett. **125**, 052002 (2020).

[38] K. Hamilton, R. Medves, G. P. Salam, L. Scyboz, and G. Soyez, Colour and logarithmic accuracy in final-state parton showers, J. High Energy Phys. 03 (2021) 041.

[39] A. M. Sirunyan *et al.* (CMS Collaboration), Evidence for the Higgs boson decay to a bottom quark–antiquark pair, Phys. Lett. B **780**, 501 (2018).

[40] G. Aad *et al.* (ATLAS Collaboration), Measurements of WH and ZH production in the $H \to b\bar{b}$ decay channel in pp collisions at 13 TeV with the ATLAS detector, Eur. Phys. J. C **81**, 178 (2021).

[41] G. Aad *et al.* (ATLAS Collaboration), Measurement of the associated production of a Higgs boson decaying into b-quarks with a vector boson at high transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Lett. B **816**, 136204 (2021).

[42] A. M. Sirunyan *et al.* (CMS Collaboration), Inclusive search for highly boosted Higgs bosons decaying to bottom quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV, J. High Energy Phys. 12 (2020) 085.

[43] J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, Jet substructure as a new Higgs search channel at the LHC, Phys. Rev. Lett. **100**, 242001 (2008).

[44] S. Marzani, G. Soyez, and M. Spannowsky, *Looking Inside Jets: An Introduction to Jet Substructure and Boosted-Object Phenomenology* (Springer, New York, 2019), Vol. 958.

[45] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, Comput. Phys. Commun. **191**, 159 (2015).

[46] C. Bierlich *et al.*, A comprehensive guide to the physics and usage of PYTHIA 8.3, arXiv:2203.11601.

[47] R. Corke and T. Sjostrand, Interleaved parton showers and tuning prospects, J. High Energy Phys. 03 (2011) 032.

[48] P. Skands, S. Carrazza, and J. Rojo, Tuning PYTHIA 8.1: The Monash 2013 tune, Eur. Phys. J. C **74**, 3024 (2014).

[49] P. Gras, S. Höche, D. Kar, A. Larkoski, L. Lönnblad, S. Plätzer, A. Siódmok, P. Skands, G. Soyez, and J. Thaler, Systematics of quark/gluon tagging, J. High Energy Phys. 07 (2017) 091.

[50] M. Dasgupta, F. Dreyer, G. P. Salam, and G. Soyez, Small-radius jets to all orders in QCD, J. High Energy Phys. 04 (2015) 039.

[51] S. Catani and M. H. Seymour, A general algorithm for calculating jet cross-sections in NLO QCD, Nucl. Phys. **B485**, 291 (1997); **B510**, 503(E) (1998).

[52] M. van Beekveld, W. Beenakker, E. Laenen, and C. D. White, Next-to-leading power threshold effects for inclusive and exclusive processes with final state jets, J. High Energy Phys. 03 (2020) 106.