Aspects of jets at 100 TeV

Andrew J. Larkoski and Jesse Thaler

Aspects of Jets at 100 TeV

Andrew J. Larkoski and Jesse Thaler

Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

We present three case studies at a 100 TeV proton collider for how jet analyses can be improved using new jet (sub)structure techniques. First, we use the winner-take-all recombination scheme to define a recoil-free jet axis that is robust against pileup. Second, we show that soft drop declustering is an effective jet grooming procedure that respects the approximate scale invariance of QCD. Finally, we highlight a potential standard candle for jet calibration using the soft-dropped energy loss. This latter observable is remarkably insensitive to the scale and flavor of the jet, a feature that arises because it is infrared/collinear unsafe, but Sudakov safe.

I. INTRODUCTION

The LHC has ushered in a new era of precision jet physics, with advances in both QCD calculations and jet analysis techniques. By now, jets at the LHC are robust objects, characterized not just by their overall energy and direction but also by their substructure [1–3]. On the longterm horizon is a future hadron collider at 100 TeV [4–6], which will open a new kinematic regime at high energies. Given the jet successes at the LHC, it is worth studying whether new jet techniques might improve the physics capabilities of a 100 TeV machine.

In this paper, we highlight three potentially powerful ways to define and study jets: winner-take-all (WTA) axes [7–9], soft drop declustering [10], and Sudakov safe observables [11]. Most of these methods have been introduced elsewhere, along with detailed analytical calculations. The goal here is to demonstrate the utility of these procedures in a high energy and high luminosity environment. Our focus is on jets at 100 TeV, but of course, these same studies are relevant for Run II of the LHC at 14 TeV.

Leading up to the LHC era, one of the key (if not *the* key) concept in jet physics was "infrared and collinear (IRC) safety" [12], which characterizes whether an observable can be predicted at fixed order in perturbative QCD. Looking towards a 100 TeV machine, we think that two important concepts in jet physics will be "recoil insensitivity" and "Sudakov safety".

- *Recoil insensitivity.* Jets can be affected by various kinds of contamination from uncorrelated radiation, including perturbative soft radiation and non-perturbative underlying event activity. As one goes to higher energies and luminosities, pileup (multiple proton-proton collisions per bunch crossing) becomes an increasingly problematic source of jet contamination. Just by momentum conservation, any uncorrelated radiation can significantly displace the momentum axis of a jet from the direction associated with the initiating hard parton. This effect is known as "recoil" [13–16], and can be present even for IRC safe observables. For this reason, one would like to work with "recoil-free" observables which are insensitive to this effect.

- *Sudakov safety.* IRC safety is a necessary condition for observables to be computable order-by-order in a perturbative α_s expansion. Recently it was realized that certain IRC unsafe observables can still be calculated using techniques from perturbative QCD. In particular, "Sudakov safe" observables [11] cannot be expressed as a Taylor series in α_s, but they can still be calculated perturbatively by using resummation to capture all-orders behavior in α_s. Given their different analytic structures, IRC safe and Sudakov safe observables can be sensitive to very different physics.

To highlight these concepts, we present three case studies of jets at 100 TeV. In Sec. II, we discuss the recoil-sensitivity of the standard jet axis, and show that the WTA recombination scheme [7–9] results in jets whose axis is recoil-insensitive and hence robust to high levels of pileup. In Sec. III, we show that the soft drop declustering procedure [10] is a powerful technique for pileup mitigation, and we exhibit its performance for dijet resonance reconstruction. In Sec. IV, we present a Sudakov-safe [11] and quasi-conformal observable defined via the soft drop procedure, which is only weakly-dependent on the value of α_s (and as such has weak energy scale dependence). This observable is also remarkably similar between quark and gluon jets, providing a potential "standard candle" for 100 TeV jet calibration.

The studies in this paper are based on Monte Carlo simulations of 100 TeV proton-proton collisions. All event generation, parton showering, and hadronization is done with PYTHIA 8.183 [17, 18] at Born-level only with no fixed-order corrections. Jet analyses are done with FASTJET 3.0.3 [19] at the particle-level, with no detector simulation. The algorithms and groomers used in this study are available in the Nsubjettiness and ResursiveTools FASTJET CONTRIBs [20].

II. WINNER-TAKE-ALL AT 100 TEV

High-energy jets are proxies for short-distance partons. In the standard lore, the axis of a jet corresponds roughly to the direction of the initiating parton that subsequently showered and hadronized. This is a useful picture when the jet's constituents arise primarily from final

state radiation off the initiating parton. However, at a hadron collider, there can be a significant jet contamination from (approximately) uncorrelated radiation—including underlying event, initial state radiation, and pileup—which are roughly uniformly distributed in pseudorapidity η. Because the jet axis $\hat{a}$ is typically defined by the summed three-momenta of the jet's constituents,

$$\hat{a} \propto \sum_{i \in \mathrm{jet}} \vec{p}_i, \qquad (1)$$

this uncorrelated radiation results in a significant displacement of the jet axis away from the initiating parton's direction. This effect is referred to as recoil-sensitivity [13–16], and in order to be robust against jet contamination, we would like to define and use a recoil-insensitive axis.

The first recoil-free jet axis was introduced in Refs. [8, 21, 22] where it was (eventually) called the "broadening axis". The broadening axis $\hat{b}$ of a jet is defined by minimizing the value of broadening [13, 23, 24] with respect to it:

$$\hat{b} = \underset{\hat{n}}{\mathrm{argmin}} \sum_{i \in \mathrm{jet}} p_{Ti}\, R_{i\hat{n}}, \qquad (2)$$

where the sum runs over the particles in the jet, p_{Ti} is the transverse momentum of particle i with respect to the beam direction, and $R_{i\hat{n}}$ is the angle between particle i and the axis $\hat{n}$.[1] Unlike the standard jet axis definition, the minimization procedure for the broadening axis cannot be solved exactly, and numerical procedures to estimate $\hat{b}$ suffer from significant computational costs and spurious local minima.

A more practical recoil-free jet axis was presented in Refs. [7–9], where a "winner-take-all (WTA) axis" was defined by modifying a standard jet clustering algorithm. Pairwise sequential jet algorithms are defined by two pieces: a clustering metric and a recombination scheme. The metric defines how close two particles are and whether they should be merged into a common jet. Different metrics give rise to different jet algorithms—the anti-k_T [26], Cambridge/Aachen [27–29], and k_T [30, 31] algorithms are typical examples—and the only constraint on the metric is IRC safety. The recombination scheme specifies how the momenta of two merging daughter particles should be mapped onto the momentum of the mother particle. There is considerable freedom in this mapping (up to IRC safety). The ubiquitous recombination scheme is the E-scheme [32] where the daughter four-momenta are simply summed to define the mother momentum. The E-scheme is manifestly sensitive to recoil because soft, wide-angle emissions in the jet will displace the mother from the harder of the daughter particles.

───────

[1] In this language, the standard jet axis is given approximately by minimizing thrust [25], $\hat{a} \approx \underset{\hat{n}}{\mathrm{argmin}} \sum_{i \in \mathrm{jet}} p_{Ti} R_{i\hat{n}}^2$.
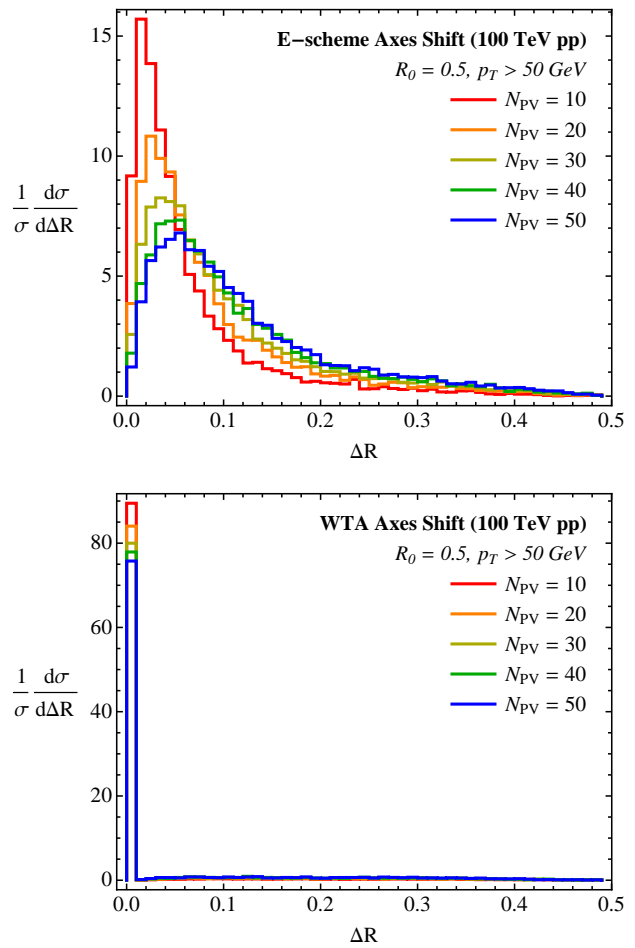

FIG. 1. The angular shift of the jet axis due to pileup, comparing the standard E-scheme jet axis (top) to the WTA jet axis (bottom), sweeping the number of pileup vertices N_{PV}.

The WTA recombination scheme removes any effect from recoil. In this scheme, the mother's p_T is given by the scalar sum of the daughters' p_T, but the mother's direction is that of the harder daughter:

$$p_{TJ} = p_{Ti} + p_{Tj},$$
$$\phi_J = \begin{cases} \phi_i, & p_{Ti} > p_{Tj}, \\ \phi_j, & p_{Tj} > p_{Ti}, \end{cases}$$
$$\eta_J = \begin{cases} \eta_i, & p_{Ti} > p_{Tj}, \\ \eta_j, & p_{Tj} > p_{Ti}, \end{cases}$$

where the daughters are i, j and the mother is J. This recombination scheme is IRC safe and is manifestly insensitive to the effects of recoil from soft, wide-angle emissions. After running a pairwise jet algorithm with the WTA scheme (and any clustering metric), the final mother's direction defines a recoil-free jet axis.

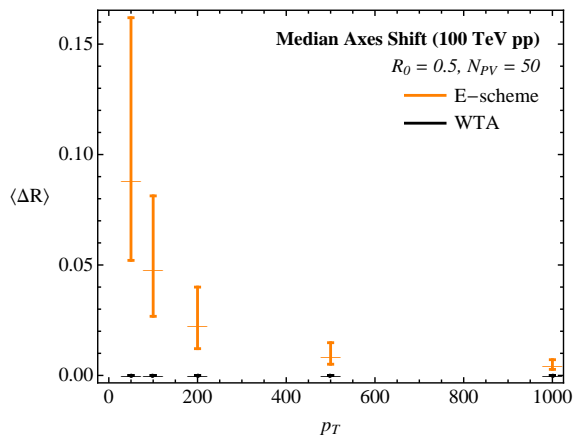To test the robustness of WTA jet axes, we generated dijet samples with $p_T > 50$ GeV at a 100 TeV proton

FIG. 2. The median angular shift due to pileup as a function of the jet p_T, comparing the E-scheme and WTA axes. We have fixed $N_{\rm PV} = 50$. The lower (upper) error bar corresponds to the first (third) quartile. For WTA, the median and the third quartile are in fact zero for the p_T range studied.

collider with various numbers of pileup vertices.[2] The direction of the jets was compared before and after the addition of pileup, and the angle between the axes was computed. Two jet algorithms were considered: anti-k_T jets with standard E-scheme recombination and anti-k_T jets with WTA recombination. The jet radius in both samples is $R_0 = 0.5$.

In Fig. 1, we plot the angle ΔR between the jet axis before and after the addition of pileup. As the number of pileup vertices $N_{\rm PV}$ increases from 10 to 50, ΔR for the E-scheme axes increases noticeably, with a long tail extending to the jet radius of 0.5. By contrast, the WTA axes are amazingly rigid, and the vast majority of the jets have an identical axis ($\Delta R = 0$) before and after the addition of pileup, even up to $N_{\rm PV} = 50$. It should be stressed that for parton-level jets with $p_T = 50$ GeV, 50 pileup vertices at a 100 TeV collider corresponds to an $\mathcal{O}(1)$ increase in the p_T of the jet, yet the WTA jet axes are still robust.

Fig. 2 illustrates the p_T dependence of the angle ΔR between the jet axis before and after the addition of pileup. Here, the number of pileup vertices is fixed at $N_{\rm PV} = 50$ and the p_T of the jets ranges from 50 to 1000 GeV. We compare the median value of ΔR for the WTA and the E-scheme axes, with the lower and upper error bars corresponding to the first and third quartile of the distribution. As the p_T of the jet increases, the median ΔR decreases for the E-scheme axes, but is still non-zero even at $p_T = 1000$ GeV. By contrast, the median of the distribution of ΔR for the WTA axes is zero over the

_____

[2] A 50 GeV jet at a 100 TeV collider—or even at Run II at the LHC—might seem a bit ridiculous, but here is used as a proof-of-concept. The WTA axes are robust to recoil effects even for jets at very low p_T.
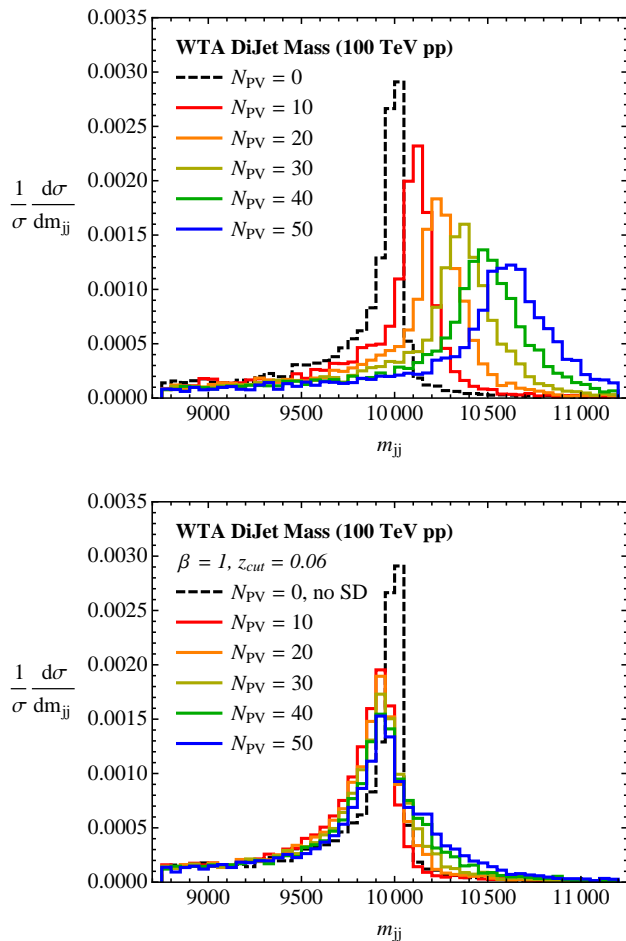
entire plotted range (as is the third quartile). In App. A, we further demonstrate the robustness of the WTA axis to the effects of non-uniform pileup.

Beyond insensitivity to pileup, it is worth mentioning that recoil-insensitive axes are powerful for other reasons as well. From a theoretical perspective, recoil-free observables are significantly easier to calculate in perturbative QCD than their recoil-sensitive counterparts [8]. Recoil-free observables also exhibit improved discrimination power between quark and gluon jets [16]. From an experimental perspective, detector noise and finite resolution can fake the addition (or subtraction) of jet radiation, and a recoil-free axis is less sensitive to such effects. Recoil-free axes can also be useful for validating pileup removal techniques, since the jet axis before and after the addition of pileup should be nearly identical.

III. SOFT DROP AT 100 TEV

While recoil-free observables offer some degree of robustness against pileup, one still needs techniques that actively identify and remove pileup contamination from a jet. Such pileup removal procedures go under the general name of "jet grooming" and are vital for, say, accurately reproducing the mass of resonances that decay to jets. Numerous grooming techniques have been developed, with filtering [33], mass drop [33], pruning [34, 35], and trimming [36] being the most widely used at the LHC. These procedures have been extensively validated on data [37, 38] and their effects on jets and jet observables are well-understood [39–41]. Here, we will review another jet grooming technique called "soft drop declustering" [10] and explain why it is well-suited for 100 TeV jets. Soft drop will also be important for the quasi-conformal observables we study in the next section.

To understand the motivation for soft drop, it is informative to review trimming [36] as an illustrative example of jet grooming. To trim a jet, one first reclusters the jet's constituents using, e.g., the k_T algorithm to form subjets of some small radius $R_{\rm sub}$ (typically about 0.3). From the set of subjets, those whose p_T fraction is less than some $z_{\rm cut}$ (usually about 0.03) are removed from the jet. This has been shown to be an effective procedure for mitigating the effects of pileup on jet observables and allows for a more accurate reconstruction of the mass of boosted resonances.

Despite its successes, one key drawback of trimming is that it does not respect the approximate scale invariance of QCD (see Ref. [40] for a discussion). As the p_T of a jet increases, jet radiation will move to smaller angular scales, such that no single $R_{\rm sub}$ will be optimal over a wide range of p_T. This is especially problematic for heavy objects like W/Z or Higgs bosons, for which, at sufficiently large boosts, all of their decay products can lie within a cone of radius $R_{\rm sub}$. Of course, realistic detectors already have finite angular granularity, so a fixed $R_{\rm sub}$ may not be too much of an issue in practice. But
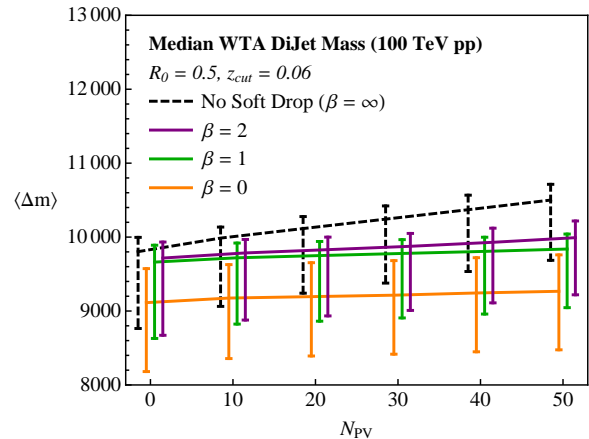
FIG. 4. Comparing the median of the dijet mass from Z' decays with different values of the soft drop grooming parameter β as a function of the number of pileup vertices N_{PV}. The lower (upper) error bar corresponds to the first (third) quartile.

FIG. 3. Distribution of the reconstructed dijet invariant mass for a 10 TeV Z' resonance, sweeping the number of pileup vertices N_{PV}. The top plot is without any pileup mitigation and the bottom plot is after the soft drop procedure.

given that jets at a 100 TeV collider can have such a large range of jet p_T values, it is worth developing grooming procedures that do not have a fundamental angular limitation.

The soft drop procedure is designed to dynamically identify the important angular range in a jet, and remove radiation outside of that range. In that sense, soft drop behaves much like the (modified) mass drop procedure [33, 40]. First, the jet is reclustered with the Cambridge/Aachen algorithm to construct an angular-ordered branching tree. Starting at the trunk of the tree, if the first branching to pseudo-jets i and j fails the soft drop criteria:

$$\frac{\min[p_{Ti}, p_{Tj}]}{p_{Ti} + p_{Tj}} > z_{\mathrm{cut}} \left(\frac{R_{ij}}{R_0} \right)^{\beta}, \qquad (3)$$

then the softer of the pseudo-jets is removed. Here, R_{ij} is the angle between the two pseudo-jets and R_0 is the jet radius. The parameter z_{cut} defines the threshold to remove soft radiation. The angular exponent β controls

how aggressive the groomer is on different angular scales, with large positive β corresponding to weak grooming and negative β to very aggressive grooming.[3] This procedure continues up the tree, following the branch with the largest p_T, until it encounters a branching that satisfies Eq. (3). At this point the recursion terminates, leaving a jet with a groomed jet radius R_g. Unlike trimming, the final angular scale R_g is determined on a per-jet basis, such that soft drop respects the approximate scale invariance of QCD. Note that the special case $\beta = 0$ corresponds to modified mass drop [33, 40].[4]

As a simple test of the efficacy of soft drop on pileup, we study the reconstruction of a massive dijet resonance. We generated events of a 10 TeV narrow Z' resonance that decays to light quarks at a 100 TeV collider with various numbers of pileup vertices. The two hardest anti-k_T jets are identified, using the WTA recombination scheme with radius $R_0 = 0.5$. The upper plot of Fig. 3 shows the dijet invariant mass as a function of number of pileup vertices. Not surprisingly, as the number of pileup vertices increases, the mass distribution both widens and drifts significantly. The lower plot of Fig. 3 shows m_{jj} after soft drop declustering with $\beta = 1$ and $z_{\mathrm{cut}} = 0.06$. As the number of pileup vertices increases, the average value of the groomed invariant mass is stable, and only

———————

[3] Strictly speaking, negative β is so aggressive that it must be used as a tagger (instead of a groomer) in order to be IRC safe. The distinction is that a groomer always return a non-zero jet (even if it only has one constituent), whereas a tagger does not return a jet if Eq. (3) fails through the whole branching tree (i.e. a jet with one remaining constituent is considered untagged).

[4] In the `ResursiveTools` FastJet contrib, the $\beta = 0$ limit of soft drop is not quite identical to modified mass drop. Soft drop defaults to grooming mode, whereas modified mass drop defaults to tagging mode.

slightly widens, nicely verifying the performance of soft drop.

The soft drop parameters used in Fig. 3 were chosen because they resulted in the best reconstruction of the resonance mass. To justify the choice of $\beta = 1$, in Fig. 4 we plot the median dijet mass as a function of the number of pileup vertices for different values of β (fixing $z_{\mathrm{cut}} = 0.06$). The upper (lower) error bars in Fig. 4 correspond to the first (third) quartiles of the mass distribution. From this figure, it is clear that $\beta = 0$ is too aggressive for this purpose, since the median value of the mass is decreased by several hundred GeV from baseline with no pileup.[5] Both $\beta = 1$ and 2 both nicely mitigate the rise in the median mass without pileup, with $\beta = 1$ performing slightly better. Also for $\beta = 1$, the locations of the first and third quartiles are relatively stable as pileup increases, comparable to the ungroomed mass distribution.

IV. STANDARD CANDLES AT 100 TEV

Beyond the utility of soft drop as a jet grooming technique, soft drop can also be used to define new classes of jet observables with unique properties. Several such observables were studied analytically in Ref. [10], including soft-dropped energy correlation functions [15, 16] and the groomed jet radius R_g itself. Here, we will study the p_T fraction of a jet removed by soft drop.

Given an original jet with transverse momentum p_{T0}, the total fractional energy loss Δ_E is[6]

$$\Delta_E = \frac{p_{T0} - p_{Tg}}{p_{T0}}, \qquad (4)$$

where p_{Tg} is the groomed jet transverse momentum. This is the definition of the groomed jet energy loss originally given in Ref. [10]. Alternatively, we can measure the maximum energy fraction of the branches removed by soft drop:

$$z_{\mathrm{max}} = \max_{\mathrm{failed\ branches}} \frac{p_{T,\mathrm{dropped}}}{p_{T,\mathrm{dropped}} + p_{T,\mathrm{kept}}}. \qquad (5)$$

The reason for considering the two different definitions of the groomed p_T loss will be explained below. We will see that both Δ_E and z_{max} are quasi-conformal observables, in the sense that their distributions are remarkably insensitive to the energy scale of the jet.

In the case of pileup mitigation, neither Δ_E nor z_{max} are particularly interesting, since they correspond to radiation that one wants to remove from a jet since it likely

comes from contamination. In the absence of pileup, though, Δ_E and z_{max} are sensitive probes of the intrinsic soft radiation captured in a jet, so one could imagine measuring them to calibrate the response of a detector to soft (perturbative) physics.

For $\beta > 0$, Δ_E and z_{max} are IRC safe, so the distributions can be computed in perturbative QCD. The calculation of the distribution of Δ_E was presented in detail in Ref. [10], including the resummation of large logarithms. To leading-logarithmic accuracy with fixed coupling α_s, the cumulative distribution was found to be

$$\Sigma(\Delta_E) = \frac{\log z_{\mathrm{cut}} - B_i}{\log \Delta_E - B_i}$$
$$+ \frac{\pi\beta}{2C_i\alpha_s} \frac{1 - e^{-2\frac{\alpha_s}{\pi}\frac{C_i}{\beta}\log\frac{z_{\mathrm{cut}}}{\Delta_E}\left(\log\frac{1}{\Delta_E} + B_i\right)}}{(\log\Delta_E - B_i)^2}, \quad (6)$$

where C_i is the color factor of the jet ($C_F = 4/3$ for quarks, $C_A = 3$ for gluons) and B_i are the subleading terms in the splitting functions ($B_q = -3/4$ for quarks, $B_g = -\frac{11}{12} + \frac{n_f}{6C_A}$ for gluons, where n_f is the number of light flavors). Of course, this distribution can be improved with running coupling effects, higher order resummation, and nonperturbative corrections. To the accuracy that we will work in this paper, the distribution for z_{max} is identical to Eq. (6), with the replacement $\Delta_E \to z_{\mathrm{max}}$.

However, there is an important distinction between Δ_E and z_{max}. Because Δ_E is determined by the sum total of all radiation that was groomed away, this observable depends on multiple emissions within the jet. The effect of multiple emissions is not included in Eq. (6), as it is formally beyond the accuracy of that expression. On the other hand, z_{max} is defined by a single groomed branch, and so the effects of multiple emissions are minimal. Therefore, we expect that Eq. (6) will give a better description of z_{max} than Δ_E.

The form of Eq. (6) is not particularly enlightening, but various expansions can be taken to illuminate its behavior. First, we can expand in α_s which results in

$$\Sigma(\Delta_E) = 1 - \frac{\alpha_s}{\pi}\frac{C_i}{\beta}\log^2\frac{z_{\mathrm{cut}}}{\Delta_E} + \mathcal{O}\left(\left(\frac{\alpha_s}{\beta}\right)^2\right), \quad (7)$$

which, for $\beta > 0$, is a Taylor series in α_s, illustrating its IRC safety. However, as $\beta \to 0$, every term in this expansion diverges, meaning that the soft-dropped energy loss is not IRC safe for $\beta = 0$. Nevertheless, from the full resummed expression we can take the $\beta \to 0$ limit first, which produces

$$\Sigma(\Delta_E)_{\beta=0} = \frac{\log z_{\mathrm{cut}} - B_i}{\log \Delta_E - B_i}, \qquad (8)$$

restricted to $\Delta_E < z_{\mathrm{cut}}$. As mentioned above, $\beta = 0$ corresponds to the (modified) mass drop procedure [33, 40]. That Δ_E for $\beta = 0$ is IRC unsafe but still calculable when all-orders effects are included means that it is a Sudakov safe observable [11].

---

[5] Of course, one can make the $\beta = 0$ groomer less aggressive by decreasing z_{cut}, but we found one had to delicately adjust z_{cut} to avoid over or under subtraction. Having $\beta > 0$ ensures that the hard core of the jet is never groomed away.

[6] For narrow jets, the distinction between energy fraction and transverse momentum fraction is negligible.
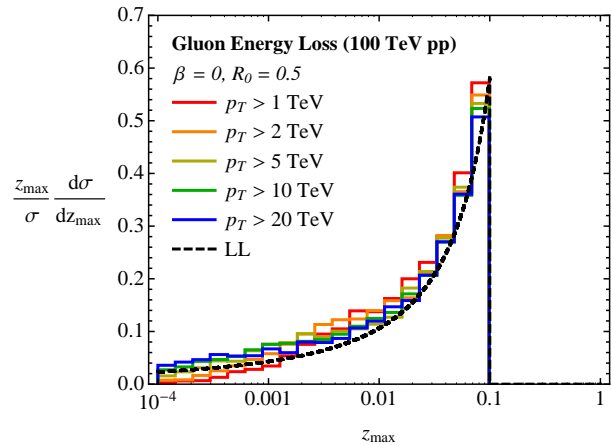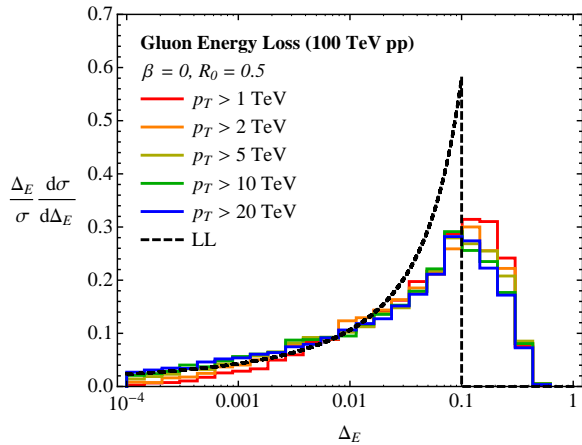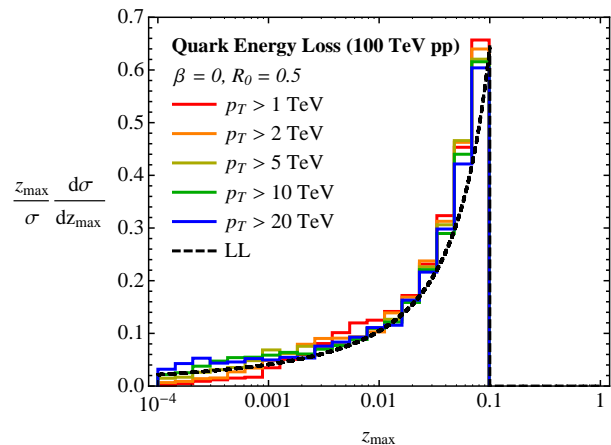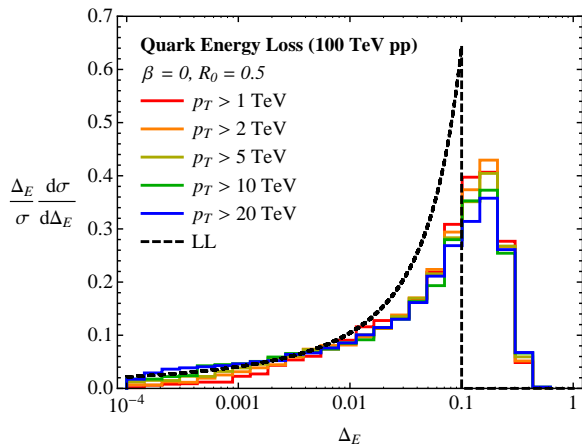
FIG. 5. Distribution of the total fractional energy loss Δ_E after soft drop for quark jet (top) and gluon jets (bottom) over a range of p_T values. "LL" is the distribution computed from Eq. (8), with the appropriate B_i factors for quark and gluon jets.

The $\beta = 0$ distribution in Eq. (8) is fascinating. In the fixed-coupling limit, it is independent of α_s. This implies that the distribution is only weakly dependent on the energy scale of the jet (i.e. it is quasi-conformal), with all dependence suppressed by the (small) β-function of QCD. It is also independent of the total color of the jet, and so the distribution should be nearly identical for quark and gluon jets, with the dependence on the flavor of the jet entering from the subleading B_i terms. This illustrates some of the surprising features of Sudakov safe observables: because their distributions are not required to be a Taylor series in α_s, they can have peculiar dependence on the coupling.

We can use parton shower simulations to test the degree to which Δ_E is independent of the jet scale and jet flavor. We generated dijet events over a range of transverse momenta at a 100 TeV collider. Unlike the previous sections, we do not include pileup in this analysis, so as to isolate the physics of the soft drop procedure on the perturbative radiation in the jet. We plot the distribu-

FIG. 6. Same as Fig. 5, but for the maximum fractional energy loss $z_{\max}$.

tion of Δ_E in Fig. 5 and $z_{\max}$ in Fig. 6 on anti-k_T jets with radius $R_0 = 0.5$ for pure quark and gluon jet samples.[7] The p_T of the jets ranges from 1 TeV to 20 TeV, and the distributions at different p_T lie on top of one another, until very small values of Δ_E or $z_{\max}$, where honest non-perturbative effects dominate. Also, the quark and gluon distributions are remarkably similar, except for large values of Δ_E where subleading perturbative effects are important. On these plots, we have also included the calculated distribution from Eq. (8), appropriate for quark or gluon jets. Especially for $z_{\max}$, the leading-logarithmic prediction nicely matches the parton shower. As expected from the discussion of multiple emissions above, Eq. (8) gives a better prediction for the distribution of $z_{\max}$ than Δ_E, especially for values near z_{cut}.

---

[7] The quark jets come from $qq \to qq$ and the gluon jets from $gg \to gg$. We ignore any subtleties regarding sample dependence or the precise definition of jet flavor.

V. CONCLUSIONS

Jets provide a unique probe into the dynamics of physics in (and beyond) the standard model. As particle physics experiments move to ever higher energies and luminosities, we have the opportunity to rethink standard approaches to jets. By incorporating new concepts like recoil insensitivity and Sudakov safety, we have the potential to increase both theoretical and experimental control over jet observables.

In this paper, we have highlighted three ways that jet analyses could be improved at a 100 TeV proton collider. The WTA recombination scheme allows the definition of a recoil-free jet axis, which improves the robustness of jet identification in the presence of pileup. The soft drop declustering procedure respects the approximate scale invariance of QCD, allowing it to groom away jet contamination over a wide dynamical range. Sudakov safe observables are not constrained to be a Taylor series in α_s, allowing for interesting probes of QCD that are not possible with standard IRC safe observables, including quasi-conformal and quasi-flavor-blind observables. While the focus of this paper was on a 100 TeV collider, these same techniques are of course relevant at the LHC, since there are similar pileup issues for 14 TeV proton collisions, and jet studies in heavy-ion collisions could benefit from increased robustness to the QCD fireball.

The most provocative proposal in this paper is using Sudakov safe observables as a standard candle for jets. Indeed, actually measuring the Δ_E or $z_{\max}$ distributions is quite challenging experimentally, since it requires a detailed understanding the energy composition of a jet. For a 10 TeV jet, measuring Δ_E down to 0.001 requires understanding 10 GeV substructure. We suspect that the scale- and flavor-independence of these observables is not unique, however, and we look forward to developing Sudakov safe observables that are more tractable experimentally. That said, this standard candle does offer an interesting and ambitious target for designing 100 TeV detectors.
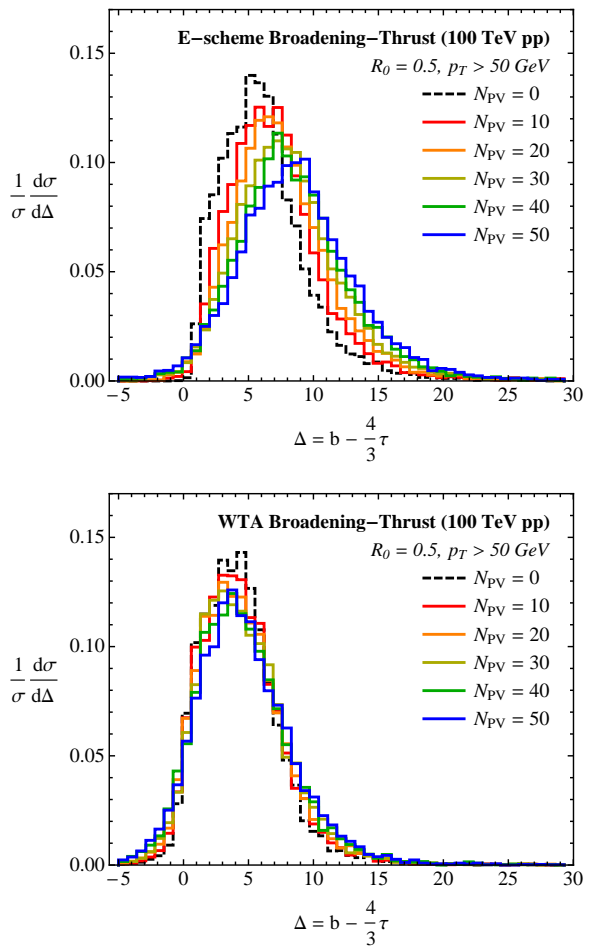
FIG. 7. The distributions for $\Delta = b - \frac{4}{3}\tau$ (i.e. broadening minus scaled thrust), comparing the standard E-scheme jet axis (top) to the WTA axis (bottom), sweeping the number of pileup vertices N_{PV}. The 4/3 scaling factor is chosen such that Δ is insensitive to uniform jet contamination, so this is a direct test of non-uniform pileup dependence.

Appendix A: Effect of Non-Uniform Pileup

We saw in Sec. II that the WTA axis was robust to pileup, but one might wonder if the same robustness could be achieved through jet area subtraction [42–44]. Crucially, the effect of recoil cannot be removed through jet area subtraction alone, since that technique assumes pileup contamination is uniformly distributed over the jet. While the jet energy can be largely corrected using a uniform subtraction, the axis shift seen in Fig. 1 really corresponds to the response of the jet axis to non-uniformities in the pileup.

To test the effect of non-uniform pileup more directly, we can study observables that, at least on average, are independent of uniform radiation in the jet. First, consider the dimensionful angularities e_α [45–47] measured

about some axis $\hat{n}$:

$$e_\alpha = \sum_{i \in \text{jet}} p_{Ti} \left(\frac{R_{i\hat{n}}}{R_0} \right)^\alpha, \qquad (A1)$$

where $R_{i\hat{n}}$ is the angle between particle i and $\hat{n}$ and the angular exponent $\alpha > 0$ for IRC safety. Different angularities have a different sensitivity to pileup because of the different angular weighting. We can exploit this fact to define an observable that is on average insensitive to uniform contamination, by taking an appropriate linear combination of two angularities e_α and e_β. For contamination with a fixed p_T, the correct linear combination is found from ($\theta \equiv R_{i\hat{n}}/R_0$)

$$\langle e_\beta - x e_\alpha \rangle \propto \int_0^1 \theta d\theta \left(\theta^\beta - x\theta^\alpha \right) = 0$$
$$\Rightarrow x = \frac{\alpha+2}{\beta+2}, \qquad (A2)$$

where $\theta d\theta$ is the angular measure for uniform radiation

in the jet. Therefore, the observable

$$\Delta = e_\beta - \frac{\alpha+2}{\beta+2} e_\alpha, \qquad (A3)$$

is insensitive to uniform contamination on average. For concreteness, we will consider the difference between jet broadening b ($\beta = 1$) and jet thrust τ ($\alpha = 2$),

$$\Delta = b - \frac{4}{3}\tau. \qquad (A4)$$

In Fig. 7, we plot Δ for the same sample of jets as in Fig. 1, comparing angularities measured with respect to the E-scheme axis and the WTA axis. For both axis choices, the distributions broaden somewhat as the number of pileup vertices increases. However, while the recoil-sensitive E-scheme axis exhibits a significant drift in the average value of Δ, the recoil-free WTA axis gives a distribution that remains centered at the same $\langle \Delta \rangle$ value. Thus, the recoil-free jet axis is more robust to non-uniform radiation in the jet than the standard recoil-sensitive jet axis.

---

[1] A. Abdesselam, E. B. Kuutmann, U. Bitenc, G. Brooijmans, J. Butterworth, *et al.*, Eur.Phys.J. **C71**, 1661 (2011), arXiv:1012.5412 [hep-ph].
[2] A. Altheimer, S. Arora, L. Asquith, G. Brooijmans, J. Butterworth, *et al.*, J.Phys. **G39**, 063001 (2012), arXiv:1201.0008 [hep-ph].
[3] A. Altheimer, A. Arce, L. Asquith, J. Backus Mayes, E. Bergeaas Kuutmann, *et al.*, (2013), arXiv:1311.2708 [hep-ex].
[4] Y. Gershtein, M. Luty, M. Narain, L. T. Wang, D. Whiteson, *et al.*, (2013), arXiv:1311.0299 [hep-ex].
[5] "Future circular collider study kick-off meeting," https://indico.cern.ch/event/282344/.
[6] "Workshop on future high energy circular colliders," http://indico.ihep.ac.cn/conferenceDisplay.py?confId=3813.
[7] D. Bertolini, T. Chan, and J. Thaler, JHEP **1404**, 013 (2014), arXiv:1310.7584 [hep-ph].
[8] A. J. Larkoski, D. Neill, and J. Thaler, JHEP **1404**, 017 (2014), arXiv:1401.2158 [hep-ph].
[9] G. Salam, Unpublished.
[10] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, JHEP **1405**, 146 (2014), arXiv:1402.2657 [hep-ph].
[11] A. J. Larkoski and J. Thaler, JHEP **1309**, 137 (2013), arXiv:1307.1699 [hep-ph].
[12] R. K. Ellis, W. J. Stirling, and B. Webber, Camb.Monogr.Part.Phys.Nucl.Phys.Cosmol. **8**, 1 (1996).
[13] S. Catani, G. Turnock, and B. Webber, Phys.Lett. **B295**, 269 (1992).
[14] Y. L. Dokshitzer, A. Lucenti, G. Marchesini, and G. Salam, JHEP **9801**, 011 (1998), arXiv:hep-ph/9801324 [hep-ph].
[15] A. Banfi, G. P. Salam, and G. Zanderighi, JHEP **0503**, 073 (2005), arXiv:hep-ph/0407286 [hep-ph].
[16] A. J. Larkoski, G. P. Salam, and J. Thaler, JHEP **1306**,

108 (2013), arXiv:1305.0007 [hep-ph].
[17] T. Sjostrand, S. Mrenna, and P. Z. Skands, JHEP **0605**, 026 (2006), arXiv:hep-ph/0603175 [hep-ph].
[18] T. Sjostrand, S. Mrenna, and P. Z. Skands, Comput.Phys.Commun. **178**, 852 (2008), arXiv:0710.3820 [hep-ph].
[19] M. Cacciari, G. P. Salam, and G. Soyez, Eur.Phys.J. **C72**, 1896 (2012), arXiv:1111.6097 [hep-ph].
[20] "Fastjet contrib," http://fastjet.hepforge.org/contrib/.
[21] H. Georgi and M. Machacek, Phys.Rev.Lett. **39**, 1237 (1977).
[22] J. Thaler and K. Van Tilburg, JHEP **1202**, 093 (2012), arXiv:1108.2701 [hep-ph].
[23] P. E. Rakow and B. Webber, Nucl.Phys. **B191**, 63 (1981).
[24] R. K. Ellis and B. Webber, Conf.Proc. **C860623**, 74 (1986).
[25] E. Farhi, Phys.Rev.Lett. **39**, 1587 (1977).
[26] M. Cacciari, G. P. Salam, and G. Soyez, JHEP **0804**, 063 (2008), arXiv:0802.1189 [hep-ph].
[27] Y. L. Dokshitzer, G. Leder, S. Moretti, and B. Webber, JHEP **9708**, 001 (1997), arXiv:hep-ph/9707323 [hep-ph].
[28] M. Wobisch and T. Wengler, (1998), arXiv:hep-ph/9907280 [hep-ph].
[29] M. Wobisch, DESY-THESIS-2000-049 (2000).
[30] S. Catani, Y. L. Dokshitzer, M. Seymour, and B. Webber, Nucl.Phys. **B406**, 187 (1993).
[31] S. D. Ellis and D. E. Soper, Phys.Rev. **D48**, 3160 (1993), arXiv:hep-ph/9305266 [hep-ph].
[32] G. C. Blazey, J. R. Dittmann, S. D. Ellis, V. D. Elvira, K. Frame, *et al.*, , 47 (2000), arXiv:hep-ex/0005012 [hep-ex].
[33] J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, Phys.Rev.Lett. **100**, 242001 (2008), arXiv:0802.2470 [hep-ph].

[34] S. D. Ellis, C. K. Vermilion, and J. R. Walsh, Phys.Rev. **D80**, 051501 (2009), arXiv:0903.5081 [hep-ph].

[35] S. D. Ellis, C. K. Vermilion, and J. R. Walsh, Phys.Rev. **D81**, 094023 (2010), arXiv:0912.0033 [hep-ph].

[36] D. Krohn, J. Thaler, and L.-T. Wang, JHEP **1002**, 084 (2010), arXiv:0912.1342 [hep-ph].

[37] G. Aad *et al.* (ATLAS Collaboration), JHEP **1309**, 076 (2013), arXiv:1306.4945 [hep-ex].

[38] S. Chatrchyan *et al.* (CMS Collaboration), JHEP **1305**, 090 (2013), arXiv:1303.4811 [hep-ex].

[39] J. R. Walsh and S. Zuberi, (2011), arXiv:1110.5333 [hep-ph].

[40] M. Dasgupta, A. Fregoso, S. Marzani, and G. P. Salam, JHEP **1309**, 029 (2013), arXiv:1307.0007 [hep-ph].

[41] M. Dasgupta, A. Fregoso, S. Marzani, and A. Powling, Eur.Phys.J. **C73**, 2623 (2013), arXiv:1307.0013 [hep-ph].

[42] M. Cacciari and G. P. Salam, Phys.Lett. **B659**, 119 (2008), arXiv:0707.1378 [hep-ph].

[43] M. Cacciari, G. P. Salam, and G. Soyez, JHEP **0804**, 005 (2008), arXiv:0802.1188 [hep-ph].

[44] G. Soyez, G. P. Salam, J. Kim, S. Dutta, and M. Cacciari, Phys.Rev.Lett. **110**, 162001 (2013), arXiv:1211.2811 [hep-ph].

[45] C. F. Berger, T. Kucs, and G. F. Sterman, Phys.Rev. **D68**, 014012 (2003), arXiv:hep-ph/0303051 [hep-ph].

[46] L. G. Almeida, S. J. Lee, G. Perez, G. F. Sterman, I. Sung, *et al.*, Phys.Rev. **D79**, 074017 (2009), arXiv:0807.0234 [hep-ph].

[47] S. D. Ellis, C. K. Vermilion, J. R. Walsh, A. Hornig, and C. Lee, JHEP **1011**, 101 (2010), arXiv:1001.0014 [hep-ph].