

Use of charged-track information to subtract neutral pileupMatteo Cacciari,^{1,2,3} Gavin P. Salam,^{4,*} and Gregory Soyez⁵¹*Université Paris Diderot, F-75013 Paris, France*²*Sorbonne Universités, UPMC Université Paris 06, UMR 7589, LPTHE, F-75005 Paris, France*³*CNRS, UMR 7589, LPTHE, F-75005 Paris, France*⁴*CERN, PH-TH, CH-1211 Geneva 23, Switzerland*⁵*IPhT, CEA Saclay, CNRS UMR 3681, F-91191 Gif-sur-Yvette, France*

(Received 16 March 2015; published 6 July 2015)

The use of charged pileup tracks in a jet to predict the neutral pileup component in that same jet could potentially lead to improved pileup removal techniques, provided there is a strong local correlation between charged and neutral pileup. In a Monte Carlo simulation we find that the correlation is however moderate, a feature that we attribute to characteristics of the underlying nonperturbative dynamics. Consequently, pure “neutral-proportional-to-charge” (NpC) pileup mitigation approaches do not outperform existing, area-based pileup removal methods. This finding contrasts with the arguments made in favor of a new method, “jet cleansing,” in part based on the NpC approach. We identify the critical differences between the performances of linear cleansing and trimmed NpC as being due to the former’s rejection of subjects that have no charged tracks from the leading vertex, a procedure that we name “zeroing.” Zeroing, an extreme version of the “charged-track trimming” proposed by ATLAS, can be combined with a range of pileup mitigation methods, and appears to have both benefits and drawbacks. We show how the latter can be straightforwardly alleviated. We also discuss the limited potential for improvement that can be obtained by linear combinations of the NpC and area-subtraction methods.

DOI: [10.1103/PhysRevD.92.014003](https://doi.org/10.1103/PhysRevD.92.014003)

PACS numbers: 12.38.-t, 13.87.-a, 13.87.Fh

I. INTRODUCTION

Pileup, the superposition of many soft proton-proton collisions over interesting hard-scattering events, is a significant issue at CERN’s Large Hadron Collider (LHC) and also at possible future hadron colliders. It affects many observables, including lepton and photon isolation, missing-energy determination and especially jet observables. One main technique currently in use to remove pileup from jet observables [1,2] is known as the area-median approach [3,4]. It makes an eventwise estimate of the pileup level, ρ , and then subtracts an appropriate four-momentum from each jet based on its area, i.e. its extent in rapidity and azimuth.

Detector-level information can also help mitigate the effect of pileup: for example, with methods such as particle flow [5] reconstruction, it is to some extent possible to eliminate the charged component of pileup, through the subtraction of contributions from individual charged pileup

hadrons.¹ However, even with such charged hadron subtraction (CHS), there is always a substantial remaining (largely) neutral pileup contribution, which remains to be removed. Currently, when CHS is used, area-median subtraction is then applied to remove the remaining neutral pileup.

Another approach is to use the information about charged pileup hadrons in a specific jet to estimate and subtract the remaining neutral component, without any reference to a jet area or a global event energy density ρ . Its key assumption is that the neutral energy flow is proportional to the charged energy flow and so we dub it neutral-proportional-to-charged (NpC) subtraction. An advantage that one might imagine for NpC subtraction is that, by using *local* information about the charged pileup, it might be better able to account for variations of the pileup from point to point within the event than

*On leave from CNRS, UMR 7589, LPTHE, F-75005 Paris, France.

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

¹Another experimental input that could conceivably help reject pileup in future detectors is precise timing information. One might also wonder about the potential benefit from calorimetric pointing information for photons, given that this is already being used to locate the primary vertex in Higgs decays to two photons [6]. However it seems likely that the degradation of pointing angular resolution [7] due to the lower energy of pileup photons and the higher detector occupancy would render this approach impractical. We thank Isabelle Wingerter for helpful explanations on this point.

methods that rely on eventwise pileup estimates. We understand that there has been awareness of this kind of approach in the ATLAS and CMS collaborations for some time now, and we ourselves also investigated it some years ago [8]. Our main finding was that at particle level it performed marginally worse than area subtraction combined with CHS. From discussions with colleagues in the experimental collaborations, we had the expectation that there might be further degradation at detector level. Accordingly we left our results unpublished.

Recently Ref. [9] [Krohn, Low, Schwartz and Wang (KLSW)] made a proposal for an approach to pileup removal named jet cleansing. One of the key ideas that it uses is precisely the NpC method,² applied to subjets, much in the way that area-median subtraction has in the past [12,13] been used with filtering [14] and trimming [15]. KLSW found that cleansing brought large improvements over area-median subtraction.

Given our earlier findings, KLSW’s result surprised us. The purpose of this article is therefore to revisit our study of the NpC method and also to carry out independent tests of cleansing, both to examine whether we reproduce the large improvements that they observed and to identify possible sources of differences. As part of our study, we will investigate what properties of events can provide insight into the performance of the NpC method. We will also be led to discuss the possible value of charged tracks from the leading vertex in deciding whether to keep or reject individual subjets (as in charged-track based trimming of Ref. [11]). Finally we shall also examine how one might optimally combine NpC and area-median subtraction.

II. THE NEUTRAL-PROPORTIONAL-TO-CHARGED METHOD

NpC subtraction relies on the experiments’ ability to identify whether a given charged track is from a pileup vertex, in order to measure the charged pileup entering a particular jet. To a good extent this charged component can be removed, for example as in CMS’s CHS procedure in the context of particle flow [5]. The NpC method then further estimates and subtracts the neutral pileup component by assuming it to be proportional to the charged pileup component. At least two variants can be conceived of.

If the charged pileup particles are kept as part of the jet during clustering, then the corrected jet momentum is [8]

$$p_{\mu}^{\text{jet,sub}} = p_{\mu}^{\text{jet}} - \frac{1}{\gamma_0} p_{\mu}^{\text{jet,chg-PU}}, \quad (1)$$

²They also investigated the use of a variable known as the jet vertex fraction, widely used experimentally to reject jets from a vertex other than the leading one [1,10,11].

where $p_{\mu}^{\text{jet,chg-PU}}$ is the four-momentum of the charged pileup particles in the jet and γ_0 is the average fraction of pileup transverse momentum that is carried by charged particles. Specifically, one can define

$$\gamma_0 \equiv \left\langle \frac{\sum_{i \in \text{charged particles}} P_{ti}}{\sum_{i \in \text{all particles}} P_{ti}} \right\rangle_{\text{events}}, \quad (2)$$

where the sums run over particles in a given event (possibly limited to some central region with tracking), and the average is carried out across minimum-bias events.

If the charged pileup particles are not directly included in the clustering (i.e. it is the CHS event that is provided to the clustering), then one does not have any information on which charged particles should be used to estimate the neutral pileup in a given jet. This problem can be circumvented by clustering an “emulated” CHS event, in which the charged pileup particles are kept, but with their momenta rescaled by an infinitesimal factor ϵ . In this case the correction becomes

$$p_{\mu}^{\text{jet,sub}} = p_{\mu}^{\text{jet,CHS}} - \frac{1 - \gamma_0}{\gamma_0 \epsilon} p_{\mu}^{\text{jet,rescaled-chg-PU}}, \quad (3)$$

where $p_{\mu}^{\text{jet,CHS}}$ is the momentum of the jet as obtained from the emulated CHS event, while $p_{\mu}^{\text{jet,rescaled-chg-PU}}$ is the summed momentum of the rescaled charged pileup particles that are in the jet. When carrying out NpC-style subtraction, this is our preferred approach because it eliminates any backreaction associated with the charged pileup (this is useful also for area-based subtraction), while retaining the information about charged pileup tracks.

There are multiple issues that may be of concern for the NpC method. For example, calorimeter fluctuations can limit the experiments’ ability to accurately remove the charged pileup component as measured with tracks. For out-of-time pileup, which contributes to calorimetric energy deposits, charged-track information may not be available at all. In any case, charged-track information covers only a limited range of detector pseudorapidities. Additionally there are subtleties with hadron masses: in effect, γ_0 is different for transverse components and for longitudinal components. In this work we will avoid this problem by treating all particles as massless.³ The importance of the above limitations can only be fully evaluated in an experimental context.

We will be comparing NpC to the area-median method. The latter makes a global estimate of pileup transverse-momentum flow per unit area, ρ , by dividing an event into similarly sized patches and taking the median of the

³Particle momenta are modified so as to become massless while conserving p_t , rapidity and azimuth.




FIG. 1 (color online). (Left panel) The correlation coefficient between the neutral transverse momentum in a central patch and either the charged transverse momentum in that patch (rescaling this component would not change the value of the correlation coefficient) or the prediction using the area-median method, i.e. $\rho_{\text{ntr}}A$. (Right panel) The standard deviation of the difference between neutral transverse momentum in a central patch and either the rescaled charged transverse momentum in that patch or the prediction using the area-median method, i.e. $\rho_{\text{ntr}}A$. The events are composed of superposed zero-bias collisions simulated with PYTHIA8, tune 4C, and the number of collisions per event is Poisson distributed with average μ .

transverse momentum per unit area across all the patches. It then corrects each jet using the globally estimated ρ and the individual jet's area, A_μ ,

$$p_\mu^{\text{jet,sub}} = p_\mu^{\text{jet}} - \rho A_\mu. \quad (4)$$

Like NpC, the area-median method has potential experimental limitations. They include questions of nontrivial rapidity dependence and detector nonlinearities (the latter are relevant also for NpC). These have, to a reasonable extent, been successfully overcome by the experiments [1,2]. One respect in which NpC may have advantages over the area-median method is that the latter fails to correctly account for the fact that pileup fluctuates from point to point within the event, a feature that cannot be encoded within the global pileup estimate ρ .⁵ Furthermore NpC does not need a separate estimation of the background density ρ , which can have systematics related to the event structure (e.g. $t\bar{t}$ events vs dijet events), and there is no need to include large numbers of ghosts for determining jet areas, a procedure that has a non-negligible computational cost.

Let us now proceed with an investigation of NpC's performance, focusing our attention on particle-level events for simplicity. The key question is the potential

performance gain due to NpC's use of local information. To study this quantitatively, we consider a circular patch of radius R centered at $y = \phi = 0$ and examine the correlation coefficient⁶ of the actual neutral energy flow in the patch with two estimates: (a) one based on the charged energy flow in the same patch and (b) the other based on a global energy flow determination from the neutral particles alone, ρ_{ntr} . Figure 1 (left panel) shows these two correlation coefficients, "ntr vs chg" and "ntr vs $\rho_{\text{ntr}}A$," as a function of R , for two average pileup multiplicities, $\mu = 20$ and $\mu = 100$. One sees that the local neutral-charged correlation is slightly lower, i.e. slightly worse, than the neutral- ρ_{ntr} correlation. Both correlations decrease for small patch radii, as is to be expected, and the difference between them is larger at small patch radii. The correlation is largely independent of the number of pileup events being considered, which is consistent with our expectations, since all individual terms in the determination of the correlation coefficient should have the same scaling with N_{PU} .

Quantitative interpretations of correlation coefficients can sometimes be delicate, as we discuss in Appendix C, essentially because they combine the covariance of two observables with the two observables' individual variances. We find that it can be more robust to investigate a quantity $\sigma_{\Delta p_\mu^{\text{nr}}}$, the standard deviation of

⁴With suitable adaptations, the area-median method can be applied to characteristics of jets other than the four-momentum, e.g. jet shapes [16] and moments of fragmentation functions [17].

⁵The area-median ρ determination can be adapted to use just the jet's neighborhood (e.g. as discussed in the context of heavy-ion collisions [18]), however it can never be restricted to just the jet.

⁶By correlation coefficient between two variables x and y we always intend the Pearson correlation coefficient, r_{xy} , defined as

$$r_{xy} = \frac{\langle (x - \langle x \rangle)(y - \langle y \rangle) \rangle}{\sqrt{\langle x^2 - \langle x \rangle^2 \rangle} \sqrt{\langle y^2 - \langle y \rangle^2 \rangle}},$$

where angular brackets denote averaging over the sample.

$$\Delta p_t^{\text{ntr}} = p_t^{\text{ntr}} - p_t^{\text{ntr,estimated}}, \quad (5)$$

where the estimate of neutral energy flow, $p_t^{\text{ntr,estimated}}$, may be either from the rescaled charged flow or from $\rho_{\text{ntr}}A$. The right-hand panel of Fig. 1 shows $\sigma_{\Delta p_t^{\text{ntr}}}$ for the two methods, again as a function of R , for two levels of pileup. It is normalized to $R\sqrt{\mu}$, to factor out the expected dependence on both the patch radius and the level of pileup. A lower value of $\sigma_{\Delta p_t^{\text{ntr}}}$ implies better performance, and as with the correlation we reach the conclusion that a global estimate of ρ_{ntr} appears to be slightly more effective at predicting local neutral energy flow than does the local charged energy flow. If one hoped to use NpC to improve on the performance of area-median subtraction, then Fig. 1 suggests that one will be disappointed.

In striving for an understanding of this finding, one should recall that the ratio of charged-to-neutral energy flow is almost entirely driven by nonperturbative effects. Inside an energetic jet, the nonperturbative effects are at scales $\sim \Lambda_{\text{QCD}}$ that are tiny compared to the jet transverse momentum p_t . There are fluctuations in the relative energy carried by charged and neutral particles, for example because a leading u quark might pick up a \bar{d} or a \bar{u} from the vacuum. However, because $\Lambda_{\text{QCD}} \ll p_t$, the charged and neutral energy flows mostly tend to go in the same direction.

The case that we have just seen of an energetic jet gives an intuition that fluctuations in charged and neutral energy flow are going to be locally correlated. It is this intuition that motivates the study of NpC. We should however examine if this intuition is actually valid for pileup. We will examine one step of hadronization, namely the production of short-lived hadronic resonances, for example a ρ^+ . The opening angle between the $\pi^+\pi^0$ decay products of the ρ^+ is of order $2m_\rho/p_{t,\rho}$. Given that pileup hadrons are produced mostly at low p_t , say 0.5–2 GeV, and that $m_\rho \simeq 0.77$ GeV, the angle between the charged and neutral pions ends up being of order 1 or even larger. As a result, the correlation in direction between charged and neutral energy flow is lost, at least in part. Thus, at low p_t , nonperturbative effects specifically tend to wash out the charged-neutral angular correlation.

This point is illustrated in Fig. 2. We consider zero-bias events and examine a circular patch of radius $R = 0.4$ centered at $y = \phi = 0$. The figure shows the distribution of the charged p_t fraction, r ,

$$r = \frac{p_t^{\text{chg}}}{p_t^{\text{chg+ntr}}}, \quad (6)$$

in the patch (filled histogram, broken into contributions where the patch contains one, two or more particles). The same panel also shows the distribution of the charged p_t fraction in each of the two leading anti- k_t , $R = 0.4$ jets in dijet events (dashed and dotted histograms). Whereas the charged-to-total ratio for a jet has a distribution peaked around 0.6, as one would expect, albeit with a broad distribution, the result for zero-bias events is striking: in

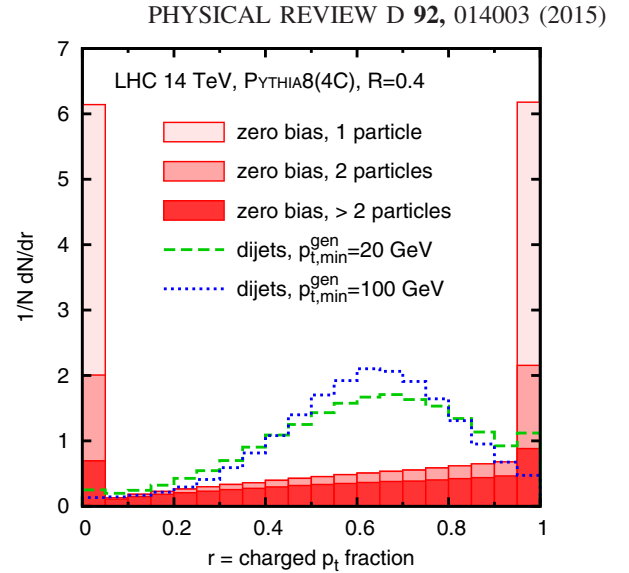


FIG. 2 (color online). The filled histogram shows the distribution, for simulated zero-bias collisions, of the fraction, r , of the transverse momentum in a central circular patch of radius $R = 0.4$ that is due to charged particles. It is separated into components according to the multiplicity of particles in the patch. The dashed and dotted histograms show the corresponding charged-fraction distributions for each of the two hardest anti- k_t , $R = 0.4$ jets in simulated dijet events, with two choices for the hard generation cut $p_{t,\text{min}}^{\text{gen}}$.

about 60% of events the patch is either just charged or just neutral, quite often consisting of just a single particle (weighting by the p_t flow in the patch, the figure goes down to 30%). This is probably part of the reason why charged information provides only limited local information about neutral energy flow in pileup events.

These considerations are confirmed by an analysis of the actual performance of NpC and area-median subtraction. We reconstruct jets using the anti- k_t algorithm [19], as implemented in FastJet⁷ [20], with a jet radius parameter of $R = 0.4$. We study dijet and pileup events generated with PYTHIA8.176 [21], in tune 4C; we assume idealized CHS, treating the charged pileup particles as ghosts. In the dijet (“hard”) event alone, i.e. without pileup, we run the jet algorithm and identify jets with absolute rapidity $|y| < 2.5$ and transverse momentum $p_t > 150$ GeV. Then in the event with superposed pileup (the “full” event) we rerun the jet algorithm and identify the jets that match those selected in the hard event⁸ and subtract them using either NpC,

⁷Results shown in this paper have been obtained in some cases with a development snapshot of version 3.1, in others with versions 3.1.0 and 3.1.1.

⁸For the matching, we introduce a quantity $p_t^{\text{shared}}(j_i^{\text{hard}}, j_j^{\text{full}})$, the scalar sum of the p_t 's of the constituents that are common to a given pair i, j of hard and full jets. For a hard jet i , the matched jet in the full event is the one that has the largest $p_t^{\text{shared}}(j_i^{\text{hard}}, j_j^{\text{full}})$. In principle, one full jet can match two hard jets, e.g. if two nearby hard jets end up merged into a single full jet due to backreaction effects. However this is exceedingly rare.

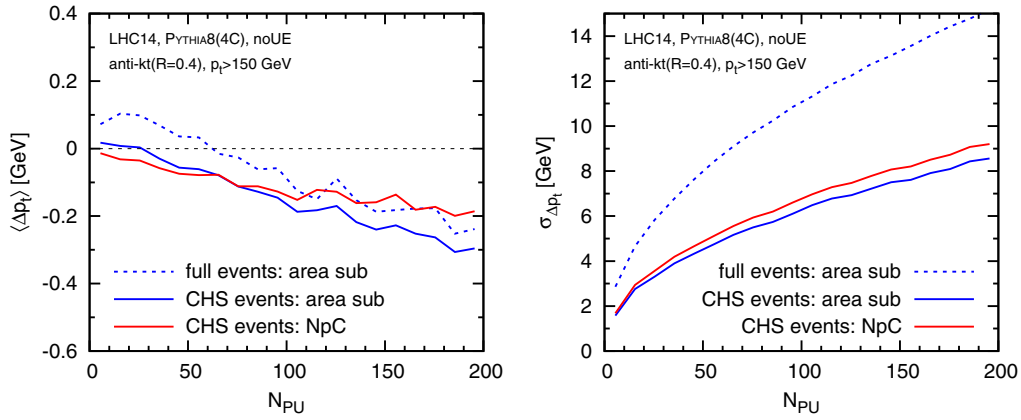


FIG. 3 (color online). A comparison of the performance of the NpC and area-median subtraction methods. The left-hand panel shows, as a function of the number of pileup vertices N_{PU} , the average difference in p_t between a jet after pileup addition and subtraction and the corresponding matched jet in the hard sample, $\Delta p_t = p_t^{\text{jet,sub}} - p_t^{\text{jet,hard}}$. The right-hand panel shows the standard deviation of Δp_t (lower values are better). NpC is shown only for CHS events, while area-median subtraction is shown both for events with CHS and for events without it (full).

Eq. (3), or the area-median method, Eq. (4), with ρ estimated from the CHS event. The hard events are generated with the underlying event turned off, which enables us to avoid subtleties related to the simultaneous subtraction of the underlying event.

Figure 3 provides the resulting comparison of the performance of the NpC and area-median subtraction methods (the latter in CHS and in full events). The left-hand panel shows the average difference between the subtracted jet p_t and the p_t of the corresponding matched hard jet, as a function of the number of pileup interactions. Both methods clearly perform well here, with the average difference systematically well below 1 GeV even for very high pileup levels. The right-hand panel shows the standard deviation of the difference between the hard and the subtracted full jet p_t . A lower value indicates better performance, and one sees that in CHS events the area-median method indeed appears to have a small, but consistent advantage over NpC. Comparing area-median subtraction in CHS and full events, one observes a significant degradation in resolution when one fails to use the available information about charged particles in correcting the charged component of pileup, as is to be expected for a particle-level study.

The conclusion of this section is that the NpC method fails to give a superior performance to the area-median method in CHS events. This is because the local correlations of neutral and charged energy flow are no greater than the correlations between local neutral energy flow and the global energy flow. We believe that part of the reason for this is that the hadronization process for low p_t particles intrinsically tends to produce hadrons separated by large angles, as illustrated concretely in the case of ρ^\pm resonance decay.

III. CLEANSING

Part of the original motivation for our work here was to cross-check a method recently introduced by KLSW and

called jet cleansing [9]. Cleansing comes in several variants. We will concentrate on linear cleansing, which was seen to perform well across a variety of observables by KLSW.⁹ It involves several elements: it breaks a jet into multiple subjets, as done for grooming methods like filtering and trimming [14,15] (cf. also the early work by Seymour [22]). In its “linear” variant, it then corrects individual subjets for pileup by a method that is essentially the same as the NpC approach described in the previous section. Cleansing may also be used in conjunction with trimming-style cuts to the subtracted subjets, specifically it can remove those whose corrected transverse momentum is less than some fraction f_{cut} of the overall jet’s transverse momentum (as evaluated before pileup removal).¹⁰

⁹Other variants of cleansing introduced by KLSW include “jet-vertex-fraction” (JVF) and “Gaussian” versions. JVF scales each subjet by $p_t^{\text{chg-LV}}/p_t^{\text{chg-total}}$, the ratio of the charged p_t from the leading vertex to the total charged p_t (including pileup) in the subjet. Gaussian is particularly interesting in that it effectively carries out a χ^2 minimization across different hypotheses for the ratio of charged to neutral energy flow, separately for the pileup and the hard event. However in KLSW’s results its performance was usually only marginally better than the much simpler linear cleansing. Accordingly we concentrate on the latter.

¹⁰In an initial preprint [23] of this article, and also in a version circulated to the authors of Ref. [9], we used $f_{\text{cut}} = 0.05$ for cleansing, reflecting our understanding of the choices made in the initial preprint version [24] of Ref. [9], which stated, “[We] supplement cleansing by applying a cut on the ratio f of the subjet p_T (after cleansing) to the total jet p_T . Subjets with $f < f_{\text{cut}}$ are discarded. [...] Where we do trim/cleanse we employ $R_{\text{sub}} = 0.3$ subjets and take $f_{\text{cut}} = 0.05$.” Subsequent to the appearance of our preprint [23], the authors of Ref. [9] clarified that the results in their Fig. 6 (Fig. 4 in the initial preprint version [24]) had used $f_{\text{cut}} = 0$. This is the choice that we adopt throughout most of this version, and it has an impact notably on the conclusions for the jet-mass performance.

The top left-hand panel of Fig. 4 shows the correlation coefficient between the dijet mass in a hard event and the dijet mass after the addition of pileup and the application of each of several pileup mitigation methods. The results are shown as a function of N_{PU} . The pileup mitigation methods include two forms of cleansing (with $f_{\text{cut}} = 0$), area-median subtraction, CHS + area subtraction and CHS + area subtraction in conjunction with trimming (also with $f_{\text{cut}} = 0$). The top right-hand panel shows the corresponding results for the jet mass. For the dijet mass we see that linear (and Gaussian) cleansing performs worse than area subtraction, while in the right-hand panel, for the jet mass, we see linear (and Gaussian) cleansing performing better than area subtraction, albeit not to the extent found in Ref. [9]. These (and, unless explicitly stated, our other Z' results) have been generated with the Z' decaying to all flavors except $t\bar{t}$, and B hadrons have been kept stable.¹¹ The lower panel shows the dijet mass for a different Z' sample, one that decays only to u , d and s quarks, but not c and b quarks. Most of the results are essentially unchanged. The exception is cleansing, which turns out to be very sensitive to the sample choice. Without stable B hadrons in the sample, its performance improves noticeably and at high pileup becomes comparable to that of area subtraction. Both of the left-hand panels in our Fig. 4 differ noticeably from Fig. 6 (top panel) of Ref. [9] and in particular they are not consistent with KLSW's observation of much improved correlation coefficients for the dijet mass with cleansing relative to area + CHS subtraction.¹²

Given our results on NpC in Sec. II, we were puzzled by the difference between the performance of area subtraction plus trimming versus that of cleansing: our expectation was that their performances should be very similar.¹³ The strong

¹¹We often find this to be useful for particle-level b -tagging studies. Experimentally, in the future, one might even imagine an “idealized” form of particle flow that attempts to reconstruct B hadrons (or at least their charged part) from displaced tracks before jet clustering.

¹²We remain puzzled also by the relative pattern of area + CHS vs area-subtracted results in Fig. 6 (top panel) of Ref. [9], since the area + CHS curve appears to tend towards the area curve at large pileup, whereas the use of CHS should significantly reduce the impact of pileup.

¹³KLSW state in [9] that fluctuations around a “best” charged fraction $\bar{\gamma}_0$ decrease with increasing N_{PU} and suggest (see also [25], pp. 16 and 17) that this will improve the determination of this fraction and therefore the effectiveness of a method like cleansing, based on a neutral-proportional-to-charged approach. However, this does not happen because, while relative fluctuations around $\bar{\gamma}_0$ do indeed decrease proportionally to $1/\sqrt{N_{\text{PU}}}$ (a result of the incoherent addition of many pileup events and of the central limit theorem), the absolute uncertainty that they induce on a pileup-subtracted quantity involves an additional factor N_{PU} . The product of the two terms is therefore proportional to $\sqrt{N_{\text{PU}}}$, i.e. the same scaling as the area-median method. This is consistent with our observations. Note that for area subtraction, the switch from full events to CHS events has the effect of reducing the coefficient in front of $\sqrt{N_{\text{PU}}}$.

sample dependence of the cleansing performance also calls for an explanation. We thus continued our study of the question.

According to the description in Ref. [9], one additional characteristic of linear cleansing relative to area subtraction is that it switches to JVF cleansing when the NpC-style rescaling would give a negative answer. In contrast, area subtraction plus trimming simply sets the (sub)jet momentum to zero. We explicitly tried turning the switch to JVF cleansing on and off and found it had a small effect and did not explain the differences.

Study of the public code for jet cleansing¹⁴ reveals an additional condition being applied to subjects: if a subject contains no charged particles from the leading vertex (LV), then its momentum is set to zero. This step appears not to have been mentioned in Ref. [9]. Since we will be discussing it extensively, we find it useful to give it a name, “zeroing.” Zeroing can be thought of as an extreme limit of the charged-track based trimming procedure introduced by ATLAS [11], whereby a JVF-style cut is applied to reject subjects whose charged-momentum fraction from the leading vertex is too low. Zeroing turns out to be crucial: if we use it in conjunction with CHS area subtraction (or with NpC subtraction) and $f_{\text{cut}} = 0$ trimming, we obtain results that are very similar to those from cleansing. Conversely, if we turn this step off in linear cleansing, its results come into accord with those from (CHS) area subtraction or NpC subtraction with $f_{\text{cut}} = 0$ trimming.

To help illustrate this, Fig. 5 shows a “fingerprint” for each of several pileup removal methods, for both the jet p_t (left panel) and mass (right panel). The fingerprint includes the average shift ($\langle\Delta p_t\rangle$ or $\langle\Delta m\rangle$) of the observable after pileup removal, shown in black. It also includes two measures of the width of the Δp_t and Δm distributions: the dispersion (i.e. standard deviation) in red and an alternative peak-width measure in blue. The latter is defined as follows: one determines the width of the smallest window that contains 90% of entries and then scales this width by a factor 0.304. For a Gaussian distribution, the rescaling ensures that the resulting peak-width measure is equal to the dispersion. For a non-Gaussian distribution the two measures usually differ and the orange shaded region quantifies the extent of this difference. The solid black, blue and red lines have been obtained from samples in which the Z' decays just to light quarks; the dotted lines are for a sample including $c\bar{c}$ and $b\bar{b}$ decays (with stable B hadrons), providing an indication of the sample dependence; in many cases they are indistinguishable from the solid lines.

Comparing $f_{\text{cut}} = 0$ grooming for NpC, area (without zeroing) and cleansing with zeroing manually disabled, all have very similar fingerprints. Turning on zeroing in the different methods leads to a significant change in the

¹⁴Version 1.0.1 from fjcontrib [26].

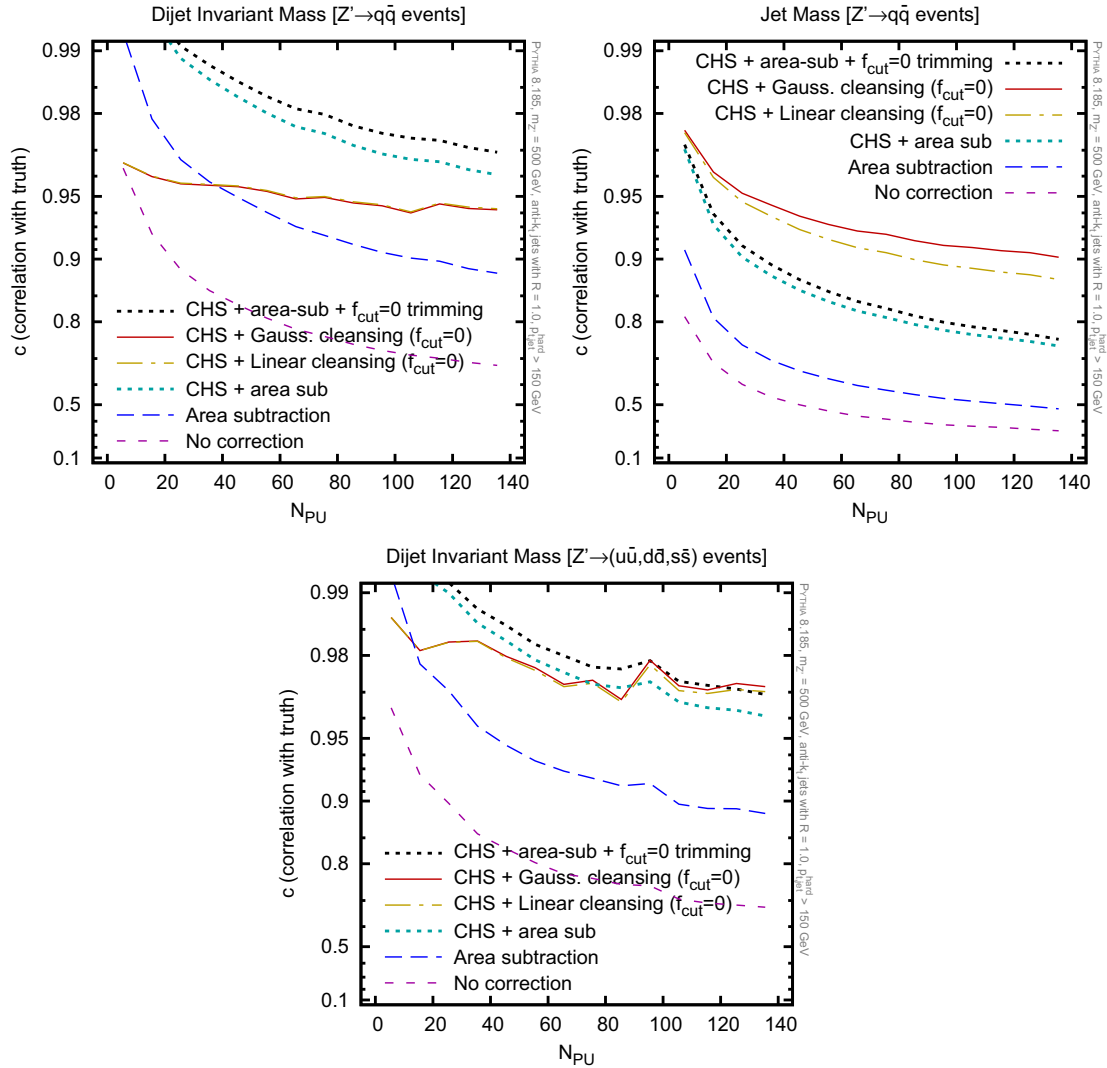


FIG. 4 (color online). (Upper-left panel) The correlation coefficient between the dijet mass in hard events and the dijet mass after adding pileup and applying various pileup-removal methods; shown as a function of the number of pileup events, N_{PU} . (Upper-right panel) Similarly for the single jet mass. Both plots are for a hadronically decaying Z' sample with $m_{Z'} = 500$ GeV. Decays are included to all flavors except $t\bar{t}$ and B hadrons are taken stable. (Lower-left panel) The correlation coefficient for the dijet mass, as in the upper-left panel, but with a sample of Z' bosons that decay only to u, d and s quarks. Jets are reconstructed, as in Ref. [9], with the anti- k_r algorithm with $R = 1$. For both trimming and cleansing, subjects are reconstructed with the k_r algorithm with $R_{sub} = 0.3$ and the f_{cut} value that is applied is $f_{cut} = 0$.

fingerprints, but again N_{pC} , area and cleansing are very similar.¹⁵

When used with $f_{cut} = 0$ trimming, and when examining quality measures such as the dispersion (in red, or the closely related correlation coefficient, cf. Appendix C),

¹⁵One exception is that for the jet mass, Gaussian cleansing does differ in the $f_{cut} = 0$ case with zeroing, and shows an advantage from its combinations of different constraints on subject momenta. Unfortunately, as we shall see later, this advantage does not seem to carry over to jet masses with $f_{cut} \neq 0$ trimming, which are phenomenologically more relevant than the full jet mass.

subject zeroing appears to be advantageous for the jet mass, but potentially problematic for the jet p_t and the dijet mass. However, the dispersion quality measure does not tell the full story regarding the impact of zeroing. Examining simultaneously the peak-width measure (in blue) makes it easier to disentangle two different effects of zeroing. On one hand we find that zeroing correctly rejects subjects that are entirely due to fluctuations of the pileup. This narrows the peak of the Δp_t or Δm distribution, substantially reducing the (blue) peak-width measures in Fig. 5. On the other hand, zeroing sometimes incorrectly rejects subjects that have no charged tracks from the LV but do have significant neutral energy flow from the LV. This can

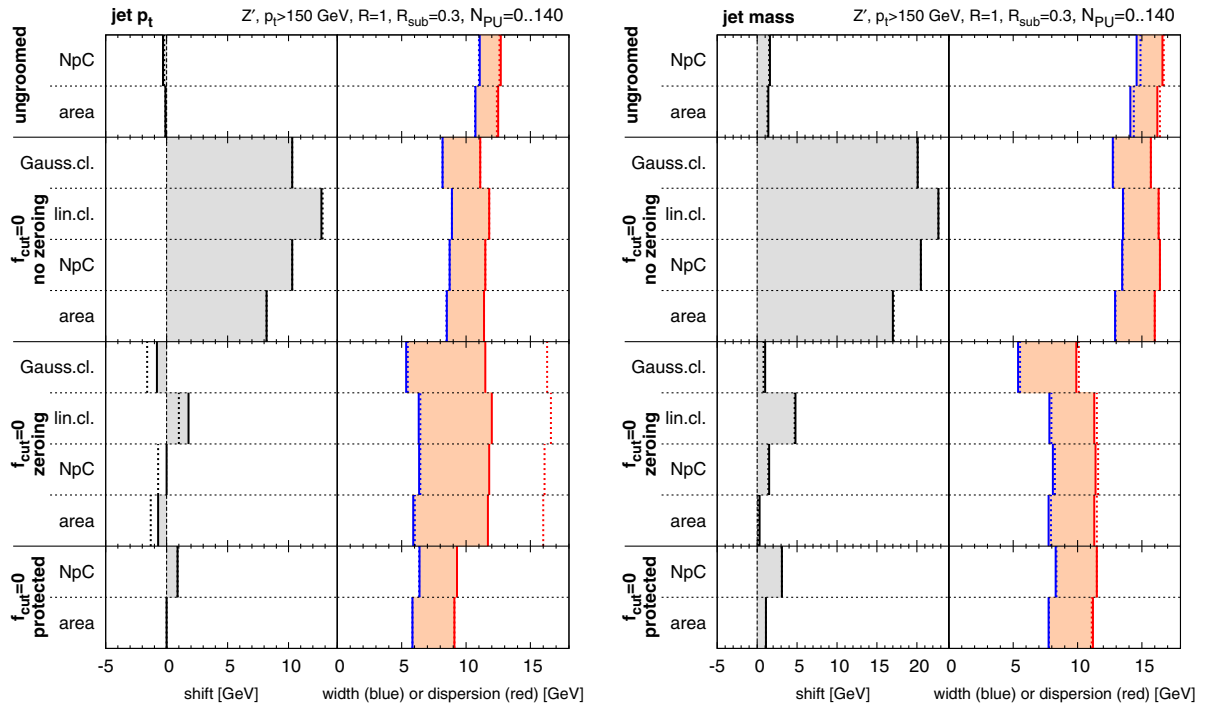


FIG. 5 (color online). The left-hand panel illustrates various characteristics of the change (Δp_t) in the jet p_t after the addition of pileup and removal by a range of methods. It shows the average shift $\langle \Delta p_t \rangle$ (in black) and the peak width (in blue) and dispersion (in red) of the Δp_t distribution. The peak width is defined as the smallest window of Δp_t that contains 90% of the Δp_t distribution, scaled by a factor ≈ 0.304 such that in the case of a Gaussian distribution the result agrees with the dispersion. The right-hand panel shows the same set of results for the jet mass. The results are obtained in a sample of events with the number of pileup vertices distributed uniformly between 0 and 140. The hard events consist of hadronic Z' decays: for the solid vertical lines the sample is $Z' \rightarrow d\bar{d}, u\bar{u}, s\bar{s}$, while for the dotted lines (sometimes not visible because they are directly over the solid lines), the sample additionally includes $Z' \rightarrow c\bar{c}, b\bar{b}$ with B hadrons kept stable. The Z' mass is $m_{Z'} = 500$ GeV and jets are reconstructed with the anti- k_r algorithm with $R = 1$. All results in this figure include charged hadron subtraction by default. The default form of cleansing, as used, e.g. in Fig. 4, is “ $f_{\text{cut}} = 0$ zeroing,”

lead to long tails for the Δp_t or Δm distributions, adversely affecting the dispersion.¹⁶ It is the interplay between the narrower peak and the longer tails that affects whether overall the dispersion goes up or down with zeroing. In particular the tails appear to matter more for the jet p_t and dijet mass than they do for the single-jet mass. Note that an accurate Monte Carlo simulation of such tails may be quite challenging: they appear to be associated with configurations where a subjet contains an unusually small number of energetic neutral particles. Such configurations are similar to those that give rise to fake isolated photons or leptons and that are widely known to be difficult to simulate correctly.

We commented earlier that the cleansing performance has a significant sample dependence. This is directly related to the zeroing: indeed Fig. 5 shows that for cleansing without zeroing, the sample dependence (dashed versus solid lines) vanishes, while it is substantial with zeroing.

¹⁶A discrepancy between dispersion and peak-width measures is to be seen in Fig. 15 of Ref. [27] for jet masses. Our fingerprint plot is in part inspired by the representation provided there, though our choice of peak-width measure differs.

Our understanding of this feature is that the lower multiplicity of jets with undecayed B hadrons (and related hard fragmentation of the B hadron) results in a higher likelihood that a subjet will contain neutral but no charged particles from the LV, thus enhancing the impact of zeroing on the tail of the Δp_t or Δm_{jj} sample.

The long tails produced by the zeroing are not necessarily unavoidable. In particular, they can correspond to the loss of subjets with tens of GeV, yet it is very unlikely that a subjet from a pileup collision will be responsible for such a large energy. Therefore we introduce a modified procedure that we call “protected zeroing”: one rejects any subjet without LV tracks *unless* its p_t after subtraction is n times larger than the largest charged p_t in the subjet from any single pileup vertex (or, more simply, just above some threshold $p_{t,\text{min}}$; however, using n times the largest charged subjet p_t could arguably be better both in cases where one explores a wide range of N_{PU} and for situations involving a hard subjet from a pileup collision). Taking $n = 10$ (or a fixed $p_{t,\text{min}} = 20$ GeV) we have found reduced tails and, consequently, noticeable improvements in the jet p_t and dijet mass dispersion (with little effect for the jet mass). This is visible for area and NpC subtraction in

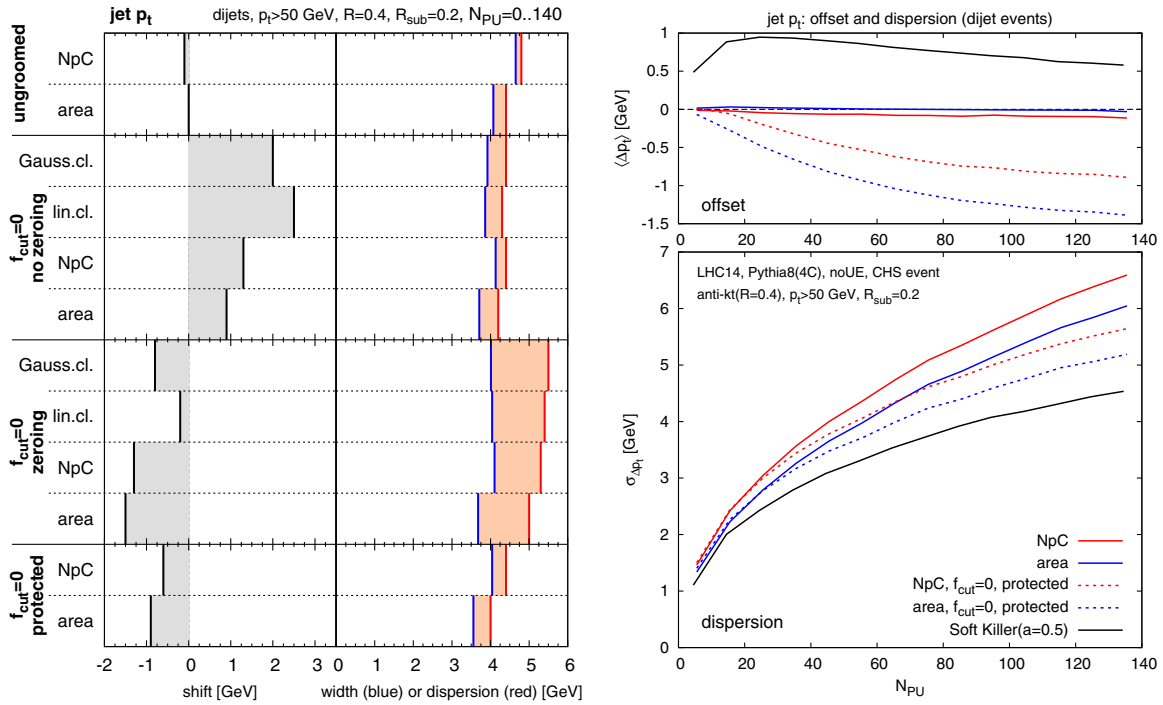


FIG. 6 (color online). (Left panel) Analogue of Fig. 5 (left panel) for a jet radius of $R = 0.4$, subjet radius (where relevant) of $R_{\text{sub}} = 0.2$ and a QCD continuum dijet sample generated with PYTHIA8. The underlying event is turned off in the sample and B hadrons decay. We consider only jets that in the hard sample have $p_t > 50$ GeV and $|y| < 2.5$. (Right panel) The dispersions for a subset of the methods, shown as a function of the number of pileup events.

Fig. 5. Protected zeroing also eliminates the sample dependence.¹⁷

Several additional comments can be made about $f_{\text{cut}} = 0$ trimming combined with zeroing. Firstly, $f_{\text{cut}} = 0$ trimming alone introduces a bias in the jet p_t , which is clearly visible in the $f_{\text{cut}} = 0$ no-zeroing shifts in Fig. 5. This is because the trimming removes negative fluctuations of the pileup, but keeps the positive fluctuations. Zeroing then counteracts that bias by removing some of the positive fluctuations, those that happened not to have any charged tracks from the LV. It also introduces further negative fluctuations for subjects that happened to have some neutral energy flow but no charged tracks. Overall, one sees that the final net bias turns out to be relatively small. This kind of cancellation between different biases is common in noise-reducing pileup-reduction approaches [28–30].

Most of the studies so far in this section have been carried out with a setup that is similar to that of Ref. [9], i.e. $R = 1$ jets in a Z' sample with $f_{\text{cut}} = 0$ trimming. This is not a common setup for most practical applications. For most uses of jets, $R = 0.4$ is a standard choice and pileup is at its most severe at low to moderate p_t . Accordingly, in

¹⁷We learned while finalizing this paper that for the identification of pileup vs nonpileup full jets in ATLAS, some form of protection is already in place, in that JVF-type conditions are not applied if $p_t > 50$ GeV. We thank David Miller for the exchanges on this point.

Fig. 6 (left panel) we show the analogue of Fig. 5's summary for the jet p_t , but now for $R = 0.4$, with $R_{\text{sub}} = 0.2$ in a QCD dijet sample, considering jets that in the hard event had $p_t > 50$ GeV. We see that qualitatively the pattern is quite similar to that in Fig. 5.¹⁸ Quantitatively, the difference between the various choices is much smaller, with about a 10% reduction in dispersion (or width) in going from ungrooved CHS area subtraction to the $f_{\text{cut}} = 0$ protected subjet-zeroing case. One should be aware that this study is only for a single p_t , across a broad range of pileup. The dispersions for a subset of the methods are shown as a function of the number of pileup vertices in the right-hand panel of Fig. 6. That panel also includes results from the SoftKiller method [29] and illustrates that the benefit from protected zeroing (comparing the solid and dashed blue curves) is about half of the benefit that is brought from SoftKiller (comparing solid blue and black curves). These panels show that protected zeroing is potentially of interest for jet p_t determinations in realistic conditions. Thus it would probably benefit from further study: one should, for example, check its behavior across a range of transverse momenta, determine optimal choices for the protection of the zeroing and investigate

¹⁸Note, however, that at even lower jet p_t 's, the difference between zeroing and protected zeroing might be expected to disappear. This is because the long negative tails are suppressed by the low jet p_t itself.

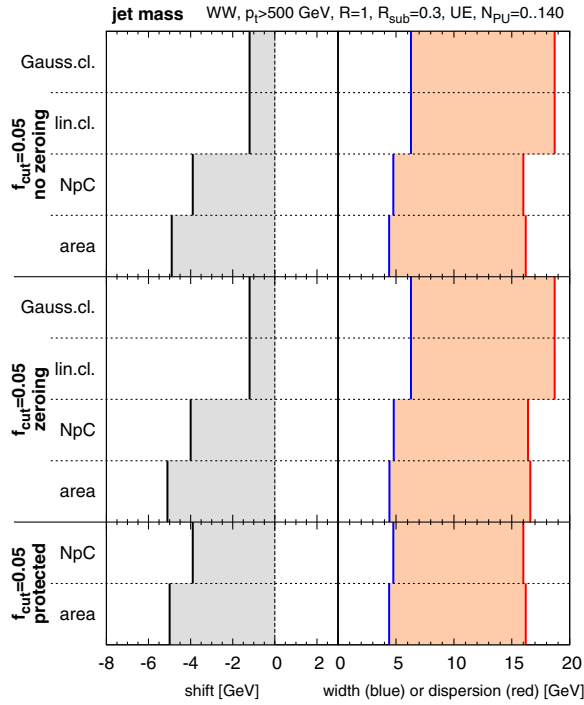


FIG. 7 (color online). Analogue of Fig. 5 (right panel), showing the performance for the jet mass, but now with $f_{\text{cut}} = 0.05$ applied to both trimming and cleansing and in a sample of hadronically decaying boosted W bosons ($pp \rightarrow W^+W^-$). The jets reconstructed after the addition and subtraction of pileup are compared to trimmed hard jets. Jets are reconstructed with a jet radius of $R = 1$ and a subjet radius of $R_{\text{sub}} = 0.3$. Only hard jets with $p_t > 500$ GeV and $|y| < 2.5$ (before trimming) are considered and we let the B hadrons decay.

also how best to combine it with particle-level subtraction methods such as SoftKiller.¹⁹

Turning now to jet masses, the use of $R = 1$ is a not uncommon choice, however most applications use a groomed jet mass with a nonzero f_{cut} (or its equivalent): this improves mass resolution in the hard event even without pileup, and it also reduces backgrounds, changing the perturbative structure of the jet [31,32] even in the absence of pileup.²⁰ Accordingly in Fig. 7 we show $f_{\text{cut}} = 0.05$ results (with shifts and widths computed relative to $f_{\text{cut}} = 0.05$ trimmed hard jets) for a hard WW sample where the hard fat jets are required to have $p_t > 500$ GeV. Zeroing, whether protected or not, appears to have little impact. One potential explanation for this fact is as follows: zeroing’s benefit comes primarily because it rejects fairly low p_t pileup subjets that happen to have no

charged particles from the leading vertex. However for a pileup subjet to pass the $f_{\text{cut}} = 0.05$ filtering criterion in our sample, it would have to have $p_t > 25$ GeV. This is quite rare. Thus filtering is already removing the pileup subjets, with little further to be gained from the charged-based zeroing. As in the plain jet-mass summary plot, protection of zeroing appears to have little impact for the trimmed jet mass.²¹ Does that mean that (protected) zeroing has no scope for improving the trimmed jet mass? The answer is, not necessarily: one could for example imagine first applying protected zeroing to subjets on some angular scale R_{zero} in order to eliminate low- p_t contamination; then reclustering the remaining constituents on a scale $R_{\text{trim}} \gtrsim R_{\text{zero}}$, subtracting according to the area or NpC methods and finally applying the trimming momentum cut (while also keeping in mind the considerations of footnote 21).

Further systematic comparisons of the different methods, with and without zeroing, are given in Appendix B.

We close this section with a summary of our findings. Based on its description in Ref. [9] and our findings about NpC vs area subtraction, cleansing with $f_{\text{cut}} = 0$ would be expected to have a performance very similar to that of CHS + area subtraction with $f_{\text{cut}} = 0$ trimming. Reference [9] however reported large improvements for the correlation coefficients of the dijet mass and the single-jet mass using $R = 1$ jets. In the case of the dijet mass we do not see these improvements, though they do appear to be there for the jet mass. The differences in behavior between (linear) cleansing and trimmed CHS + area subtraction call for an explanation, and appear to be due to a step in the cleansing code that was undocumented in Ref. [9] and that we dubbed zeroing: if a subjet contains no charged tracks from the leading vertex it is discarded. Zeroing is an extreme form of a procedure described in Ref. [11]. It can be used also with area or NpC subtraction, and we find that it brings a benefit for the peak of the Δp_t and Δm distributions, but appears to introduce long tails in Δp_t . A variant, protected zeroing, can avoid the long tails by still accepting subjets without leading-vertex tracks, if their p_t is above some threshold, which may be chosen dynamically based on the properties of the pileup. In our opinion, a phenomenologically realistic estimate of the benefit of

¹⁹An interesting feature of protected zeroing, SoftKiller and another recently introduced method, PUPPI [30], is that the residual degradation in resolution from pileup appears to scale more slowly than the $\sqrt{N_{\text{PU}}}$ pattern that is observed for area and NpC subtraction alone.

²⁰In contrast, for $f_{\text{cut}} = 0$ trimming, the jet structure is unchanged in the absence of pileup.

²¹Cleansing appears to perform slightly worse than trimming with NpC or area subtraction. One difference in behavior that might explain this is that the p_t threshold for cleansing’s trimming step is $f_{\text{cut}} p_t^{\text{full,no-CHS}}$ (even in the CHS-like `input_nc_separate` mode that we use). In contrast, for the area and NpC-based results, it is $f_{\text{cut}} p_t^{\text{full,CHS}}$. In both cases the threshold, which is applied to subtracted subjets, is increased in the presence of pileup, but this increase is more substantial in the cleansing case. This could conceivably worsen the correspondence between trimming in the hard and full samples. For the area and NpC cases, we investigated the option of using $f_{\text{cut}} p_t^{\text{area-sub,CHS}}$ or $f_{\text{cut}} p_t^{\text{NpC-sub,CHS}}$ and found that this brings a small additional benefit.

zeroing (protected or not) requires study not of $R = 1$ plain jets, but instead of $R = 0.4$ jets (for the jet p_t) or larger- R trimmed jets with a nonzero f_{cut} (for the jet mass). In a first investigation, there appear to be some phenomenological benefits from protected zeroing for the $R = 0.4$ jet p_t , whereas to obtain benefits for large- R trimmed jets would probably require further adaptation of the procedure. In any case, additional study is required for a full evaluation of protected zeroing and related procedures.

IV. COMBINING NPC AND AREA SUBTRACTION

It is interesting to further probe the relation between NpC and the area-median method, to establish whether there might be a benefit from combining them: the area-median method makes a mistake in predicting local energy flow mainly because local energy flow fluctuates from region to region. NpC makes a mistake because charged and neutral energy flow are not 100% locally correlated. The key question is whether, for a given jet, NpC and the area-median method generally make the same mistake, or if instead they are making uncorrelated mistakes. In the latter case it should be possible to combine the information from the two methods to obtain an improvement in subtraction performance.

Let p_t^{ntr} be the actual neutral pileup component flowing into a jet, while

$$\hat{p}_\mu^{\text{ntr(NpC)}} = \frac{1 - \gamma_0}{\gamma_0 \epsilon} p_\mu^{\text{jet, rescaled-chg-PU}}, \quad \hat{p}_\mu^{\text{ntr}(\rho A)} = \rho A \quad (7)$$

are, respectively, the estimates for the neutral pileup based on the local charged p_t flow and on ρA . Note that for our analysis in this section we use CHS events and in particular ρ is as determined from the CHS event, i.e. it is ρ_{CHS} . This means that it is essentially equal to ρ_{ntr} , with small

differences arising principally due to the presence of the charged component of the LV's underlying event contribution. One could of course also use the actual ρ_{ntr} instead of ρ_{CHS} to estimate $\hat{p}_\mu^{\text{ntr}(\rho A)}$, and this would in some respects be a more coherent approach, though the numerical differences will be small. For the sake of readability, in this section we simply use the symbol ρ rather than a more explicit ρ_{CHS} .

Concentrating on the transverse components, the extent to which the two estimates in Eq. (7) provide complementary information can be quantified in terms of (one minus) the correlation coefficient, r , between $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr(NpC)}}$ and $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr}(\rho A)}$. That correlation is shown as a function of R in Fig. 8 (left panel), and it is quite high, in the range 0.6 to 0.7 for commonly used R choices. It is largely independent of the number of pileup vertices.

Let us now quantify the gain to be had from a linear combination of the two prediction methods, i.e. using an estimate

$$\hat{p}_\mu^{\text{ntr}} = f \hat{p}_\mu^{\text{ntr(NpC)}} + (1 - f) \hat{p}_\mu^{\text{ntr}(\rho A)}, \quad (8)$$

where f is to be chosen so as to minimize the dispersion of $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr}}$. Given the dispersions σ_{NpC} and $\sigma_{\rho A}$, respectively, for $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr(NpC)}}$ and $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr}(\rho A)}$, the optimal f is

$$f = \frac{\sigma_{\rho A}^2 - r \sigma_{\text{NpC}} \sigma_{\rho A}}{\sigma_{\text{NpC}}^2 + \sigma_{\rho A}^2 - 2r \sigma_{\text{NpC}} \sigma_{\rho A}}, \quad (9)$$

which is plotted as a function of R in Fig. 8 (right panel), and the resulting squared dispersion for $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr}}$ is

$$\sigma^2 = \frac{(1 - r^2) \sigma_{\text{NpC}}^2 \sigma_{\rho A}^2}{\sigma_{\text{NpC}}^2 + \sigma_{\rho A}^2 - 2r \sigma_{\text{NpC}} \sigma_{\rho A}}. \quad (10)$$




FIG. 8 (color online). (Left panel) Correlation between $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr(NpC)}}$ and $p_t^{\text{ntr}} - \hat{p}_\mu^{\text{ntr}(\rho A)}$, shown as a function of R . (Right panel) Optimal weight f for combining NpC and area pileup subtraction, Eq. (9), as a function of R .




FIG. 9 (color online). Comparison of the performance of NpC, area-median and combined subtraction, as a function of the number of pileup vertices. The left-hand panel is for the jet p_t and the right-hand one for the jet mass.

Reading $r = 0.67$ from Fig. 8 (left panel) for $R = 0.4$, and $\sigma_{\text{NpC}} \approx 1.14\sigma_{\rho A}$ from Fig. 1 (right panel), one finds $\sigma \approx 0.96\sigma_{\rho A}$. Because of the substantial correlation between the two methods, one expects only a modest gain from their linear combination.

In Fig. 9 we compare the performance of pileup subtraction from the combination of the NpC and the area-median methods, using the optimal value $f = 0.31$ that can be read from Fig. 8 (right panel) for $R = 0.4$, both for the jet p_t and the jet mass. The expected small gain is indeed observed for the jet p_t , and it is slightly larger for the jet mass.²² Given the modest size of the gain, one may wonder how phenomenologically relevant it is likely to be. Nevertheless, one might still consider investigating whether the gain carries over also to a realistic experimental environment with full detector effects.

V. CONCLUSIONS

One natural approach to pileup removal is to use the charged pileup particles in a given jet to estimate the amount of neutral pileup that needs to be removed from that same jet. In this article, with the help of particle-level simulations, we have studied such a method (NpC) and found that it has a performance that is similar to, though slightly worse than, the existing, widely used area-median method. This can be related to the observation that the correlations between local charged and neutral energy flow are no larger than those between global and local energy flow. Tentatively, we believe that this is in part because the nonperturbative effects that characterize typical inelastic proton-proton collisions act to destroy local charged-neutral correlation.

²²We also examined results with other choices for f : we found that the true optimum value of f in the Monte Carlo studies is slightly different from that predicted by Eq. (9). However the dependence on f around its minimum is very weak, rendering the details of its exact choice somewhat immaterial.

The absence of benefit that we found from the NpC method led us to question the substantial performance gains quoted for the method of cleansing in Ref. [9], one of whose key differences with respect to earlier work is the replacement of the area-median method with NpC. For the dijet mass, we are unable to reproduce the large improvement observed in Ref. [9], in the correlation coefficient performance measure, for cleansing relative to area subtraction. We do however see an improvement for the jet mass. We trace a key difference in the behavior of cleansing and area subtraction (or pure NpC) to the use in the cleansing code of a step that was not documented in Ref. [9] and that discards subjets that contain no tracks from the leading vertex. This zeroing step, similar to the charged-track based trimming introduced by ATLAS [11], can indeed be of benefit. It has a drawback of introducing tails in some distributions due to subjets with a substantial neutral p_t from the leading vertex, but no charged tracks. As a result, different quality measures lead to different conclusions as to the benefits of zeroing. The tails can be alleviated by a variant of zeroing that we introduce here, protected zeroing, whereby subjets without LV charged tracks are rejected only if their p_t is below some (possibly pileup-dependent) threshold. Protected zeroing does in some cases appear to have phenomenological benefits, which are observed across all quality measures.

Given two different methods for pileup removal, NpC and area-median subtraction, it is natural to ask how independent they are and what benefit might be had by combining them. This was the question investigated in Sec. IV, where we provided a formula for an optimal linear combination of the two methods, as a function of their degree of correlation. Ultimately we found that NpC and area-median subtraction are quite highly correlated, which limits the gains from their combination to about a 5% percent reduction in dispersion. While modest, this might still be sufficient to warrant experimental investigation, as are other methods, currently being developed, that exploit constituent-level subtraction

[29,30,33,34]. A study of the integration of those methods with protected zeroing would also be of interest.

Code for our implementation of area subtraction with positive-definite mass is available as part of FastJet versions 3.1.0 and higher. Public code and samples for carrying out a subset of the comparisons with cleansing described in Sec. III, including also the NpC subtraction tools, are available from Ref. [35].

ACKNOWLEDGMENTS

Our understanding of pileup effects in the LHC experiments has benefited extensively from discussions with Peter Loch, David Miller, Filip Moortgat, Sal Rappoccio, Ariel Schwartzman and numerous others. We are grateful to David Krohn, Matthew Low, Matthew Schwartz and Liantao Wang for the exchanges about their results. This work was supported by ERC advanced grant Higgs@LHC, by the French Agence Nationale de la Recherche, under Grant No. ANR-10-CEXC-009-01, by the EU ITN grant LHCPhenoNet, Grant No. PITN-GA-2010-264564 and by the ILP LABEX (Reference No. ANR-10-LABX-63) supported by French state funds managed by the ANR within the Investissements d’Avenir program under Reference No. ANR-11-IDEX-0004-02. G. P. S. wishes to thank Princeton University for their hospitality while this work was being carried out. G. S. wishes to thank CERN for their hospitality while this work was being finalized. EU ITN grant LHCPhenoNet, Grant No. PITN-GA-2010-264564.

APPENDIX A: DETAILS OF OUR STUDY

Let us first fully specify what we have done in our study and then comment on (possible) differences relative to KLSW.

Our hard event sample consists of dijet events from pp collisions at $\sqrt{s} = 14$ TeV, simulated with PYTHIA8.176 [21], tune 4C, with a minimum p_t in the $2 \rightarrow 2$ scattering of 135 GeV and with the underlying event turned off, except for the panels presented in Figs. 4 and 5, where we use Z' events with $m_{Z'} = 500$ GeV. Jets are reconstructed with the anti- k_t algorithm [19] after making all particles massless (preserving their rapidity) and keeping only particles with $|y| < 4$. We have $R = 0.4$, except for some of the results presented in Sec. III and Appendix C, where we use $R = 1$ as in Ref. [9].

Given a hard event, we select all the jets with $p_t > 150$ GeV and absolute rapidity $|y| < 2.5$. We then add pileup and cluster the resulting full event, i.e. including both the hard event and the pileup particles, without imposing any p_t or rapidity cut on the resulting jets. For each jet selected in the hard event as described above, we find the jet in the full event that overlaps the most with it. Here, the overlap is defined as the scalar p_t sum of all the common jet constituents, as described in footnote 8. Given a pair of jets, one in the hard event and the matching one in the full event, we can apply subtraction/grooming/

cleansing to the latter and study the quality of the jet p_t or jet-mass reconstruction. For studies involving the dijet mass (cf. Fig. 4) we require that at least two jets pass the jet selection in the hard event and use those two hardest jets, and the corresponding matched ones in the full event, to reconstruct the dijet mass.²³ This approach avoids having to additionally consider the impact of pileup on the efficiency for jet selection, which cannot straightforwardly be folded into our quality measures.²⁴

Most of the studies shown in this paper use idealized particle-level CHS events. In these events, we scale all charged pileup hadrons by a factor $\epsilon = 10^{-60}$ before clustering, to ensure that they do not induce any back-reaction [4]. The jet selection and matching procedures are independent of the use of CHS or full events. When we plot results as a function of the quantity N_{PU} , this corresponds to the actual (fixed) number of zero-bias events superimposed onto the hard collision. For results shown as a function of μ , the average number of zero-bias events, the actual number of zero-bias events has a Poisson distribution. Clustering and area determination are performed with a development version of FastJet 3.1 (which for the features used here behaves identically to the 3.0.x series) and with FastJet 3.1.0 and 3.1.1 for the results in Sec. III.

Details of how the area-median subtraction is performed could conceivably matter. Jet areas are obtained using active area with explicit ghosts placed up to $|y| = 4$ and with a default ghost area of 0.01. We use FastJet’s `GridMedianBackgroundEstimator` with a grid spacing of 0.55 to estimate the event background density ρ . The ρ estimation is performed using the particles (up to $|y| = 4$) from the full or the CHS event as appropriate. When subtracting pileup from jets, we account for the

²³In the case of the Z' events used for Fig. 4, this does not exactly reflect how we would have chosen to perform a dijet (resonance) study ourselves. One crucial aspect is that searches for dijet resonances always impose a rapidity cut between the two leading jets, such as $|\Delta y| < 1.2$ [36,37]. This ensures that high dijet mass events are not dominated by low p_t forward-backward jet pairs, which are usually enhanced in QCD vs resonance production. Those forward-backward pairs can affect conclusions about pileup, because for a given dijet mass the jet p_t ’s in a forward-backward pair are lower than in a central-central pair, and so relatively more sensitive to pileup. Also the experiments do not use $R = 1$ for their dijet studies: ATLAS uses $R = 0.6$ [36], while CMS uses $R = 0.5$ with a form of radiation recovery based on the inclusion of any additional jets with $p_t > 30$ GeV and within $\Delta R = 1.1$ of either of the two leading jets (“wide jets”) [37]. This too can affect conclusions about pileup.

²⁴One alternative would have been to impose the cuts on the jets in the full event (with pileup and subtraction/grooming/cleansing) and consider as the “hard jet,” the subset of the particles in the full-event jet that come from the leading vertex (i.e. the hard event). We understand that this is close to the choice made in Ref. [9]. This can give overly optimistic results, because it neglects backreaction. However in our studies it did not appear to significantly modify the conclusions on the relative performances of different methods.

rapidity dependence of ρ , based on the rapidity dependence in a pure pileup sample (as discussed in Refs. [20,38]). We carry out 4-vector subtraction $p_\mu^{\text{sub}} = p_\mu - A_\mu \rho(y_{\text{jet}})$.

A few obviously unphysical situations need special care. For jets obtained from the full event, if $A_t \rho > p_t^{\text{jet}}$, we set p_μ^{sub} to a vector with zero transverse momentum, zero mass, and the rapidity and azimuth of the original unsubtracted jet; and if $(p^{\text{sub}})^2$ is negative, an unphysical situation since it would lead to an imaginary mass, we replace p_μ^{sub} with a vector with the same transverse components, zero mass and the rapidity of the original unsubtracted jet.²⁵ This is essentially equivalent to replacing negative squared masses with zero.

The case of CHS events is a bit more delicate. Let p_μ^{chg} denote the four-momentum of the charged component of the jet.²⁶ Then, if $p_t^{\text{sub}} < p_t^{\text{chg}}$, we set $p_\mu^{\text{sub}} = p_\mu^{\text{chg}}$, and when $(p^{\text{sub}})^2 < (p^{\text{chg}})^2$, we replace p_μ^{sub} with a vector with the same transverse components, and the mass and rapidity of p_μ^{chg} . For jets with no charged component, whenever the resulting 4-vector has an ill-defined rapidity or azimuthal angle, we use those of the original jet. Corresponding requirements that the subtracted transverse momentum and mass are non-negative are also applied in our NpC subtraction. These safety requirements have little impact on the single-jet p_t , limited impact on the dijet mass and for the single-jet mass improve the dispersion of the subtraction relative to the choice (widespread in computer codes) of taking $m \equiv -\sqrt{|m^2|}$ when $m^2 < 0$.²⁷

One difference between our study and KLSW's is that we carry out a particle-level study, whereas they project their event onto a toy detector with a 0.1×0.1 tower granularity, removing charged particles with $p_t < 0.5$ GeV and placing a 1 GeV threshold on the towers. In our original [23] studies with $f_{\text{cut}} = 0.05$ we tried including a simple detector simulation along these lines and did not find any significant modification to the pattern of our results, though CHS + area subtraction is marginally closer to the cleansing curves in this case.²⁸

²⁵In versions of FastJet prior to 3.1.0 this had to be done manually: the treatment of unphysical 4-vectors was left to the user, since the optimal treatment may depend on the context. As of version 3.1.0, FastJet provides the option of "safe" subtraction, whereby subjects with negative mass squared after subtraction are automatically assigned a zero mass (in CHS events, the safe procedure follows the description below in the text).

²⁶For our emulated CHS events, the charged contribution from pileup is negligible since it has been scaled down, and only the contribution from the hard interaction counts.

²⁷As this work was being finalized, an alternative approach to using area-median information to reconstruct the 4-vector (and shapes) of a jet was proposed in Ref. [34], based on a constituent-level assignment of the subtraction. It appears to bring further improvements in performance for observables like the jet mass.

²⁸We choose to show particle-level results here because of the difficulty of correctly simulating a full detector, especially given the nonlinearities of responses to pileup and the subtleties of particle flow reconstruction in the presence of real detector fluctuations.

Cleansing has two options. One can give it jets clustered from the full event, and then it uses an analogue of Eq. (1): this effectively subtracts the exact charged part and the NpC estimate of the neutrals. Or one can give it jets clustered from CHS events, and it then applies the analogue of Eq. (3), which assumes that there is no charged pileup left in the jet and uses just the knowledge of the actual charged pileup to estimate (and subtract) the neutral pileup. These two approaches differ by contributions related to backreaction. Our understanding is that KLSW took the former approach, while we used the latter. Specifically, our charged pileup hadrons, which are scaled down in the CHS event, are scaled back up to their original p_t before passing them to the cleansing code, in its `input_nc_separate` mode. If we use cleansing with full events, we find that its performance worsens, as is to be expected given the additional backreaction induced when clustering the full event. Were it not for backreaction, cleansing applied to full or CHS events should essentially be identical.

Regarding the NpC and cleansing parameters, our value of $\gamma_0 = 0.612$ differs slightly from that of KLSW's $\gamma_0 = 0.55$, and corresponds to the actual fraction of charged pileup in our simulated events. In our tests with a detector simulation like that of KLSW, we adjusted γ_0 to its appropriate (slightly lower) value.

Finally, for trimming we use $R_{\text{trim}} = 0.3$ and the reference p_t is taken unsubtracted, while the subjects are subtracted before performing the trimming cut, which removes subjects with p_t below a fraction f_{cut} times the reference p_t . Compared to using the subtracted p_t as the reference for trimming, this effectively places a somewhat harder cut as pileup is increased.²⁹ For comparisons with cleansing we generally use $f_{\text{cut}} = 0$ unless explicitly indicated otherwise.

APPENDIX B: FURTHER SYSTEMATIC COMPARISONS

While the fingerprint plots of Figs. 5–7 contain considerable information, they have the drawback that they do not show the dependence on the number of pileup events. To allow for further systematic comparisons, here we additionally show results for the dispersions (Fig. 10) and correlations (Fig. 11) versus the number of pileup events, N_{PU} . The main observations remain as for Figures 5–7:

- (a) Without zeroing, the linear cleansing, NpC and area-median methods show similar performance, with a slight overall preference for the area-median method.
- (b) Zeroing affects all methods in a similar way. It brings sizable improvements for the jet mass, but worsens the

²⁹And is the default behavior in FastJet if one passes an unsubtracted jet to a trimmer with subtraction, e.g. a `Filter(Rtrim, SelectorPtFractionMin(fcut), ρ)`. One may of course choose to pass a subtracted jet to the trimmer, in which case the reference p_t will be the subtracted one.

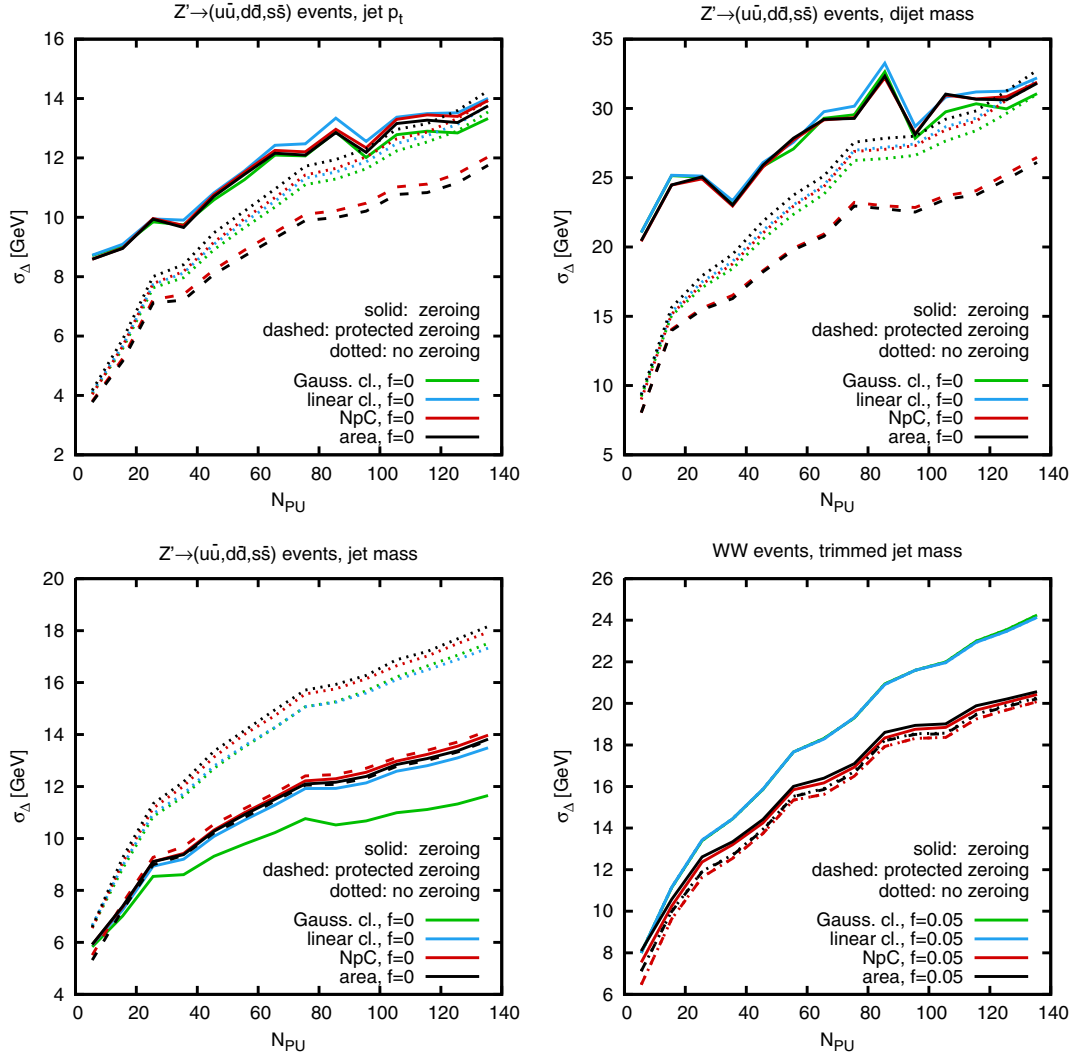


FIG. 10 (color online). Dispersions versus the number of pileup events for Gaussian and linear cleansing as well as area and NpC subtraction, showing the impact of zeroing and, where relevant, protected zeroing. All results in this figure include CHS by default. Note that in the bottom-right panel (the trimmed jet mass), all curves for linear and Gaussian cleansing, both with and without zeroing, are superposed.

reconstruction of the jet p_t and the dijet mass, notably at low and moderate values of N_{PU} .

- (c) Switching to the protected version of zeroing maintains the performance for the jet mass and also brings significant improvements for the jet p_t .
- (d) Using Gaussian cleansing instead of linear cleansing brings relevant improvements for the jet mass. We believe that this may be because Gaussian cleansing brings an advantage for very soft subjets and the jet mass has enhanced sensitivity to such subjets, notably when they are at the edge of the jet.
- (e) When used in conjunction with $f_{\text{cut}} = 0.05$ trimming (the bottom-right panels, for WW events), a standard phenomenological choice, most methods tend to perform similarly for the jet mass. The worse performance for cleansing is, we believe, partially attributable to its application of an f_{cut} condition normalized

to the full jet p_t , including neutral and charged pileup, as discussed in footnote 21.

APPENDIX C: THE CORRELATION COEFFICIENT AS A QUALITY MEASURE

In this appendix, we discuss some characteristics of correlation coefficients that affect their appropriateness as generic quality measures for pileup studies.

Suppose we have an observable v . Define

$$\Delta v = v^{\text{sub}} - v^{\text{hard}} \quad (\text{C1})$$

to be the difference, in a given event, between the pileup subtracted observable and the original hard value without pileup. Two widely used quality measures for the performance of pileup subtraction are the average offset

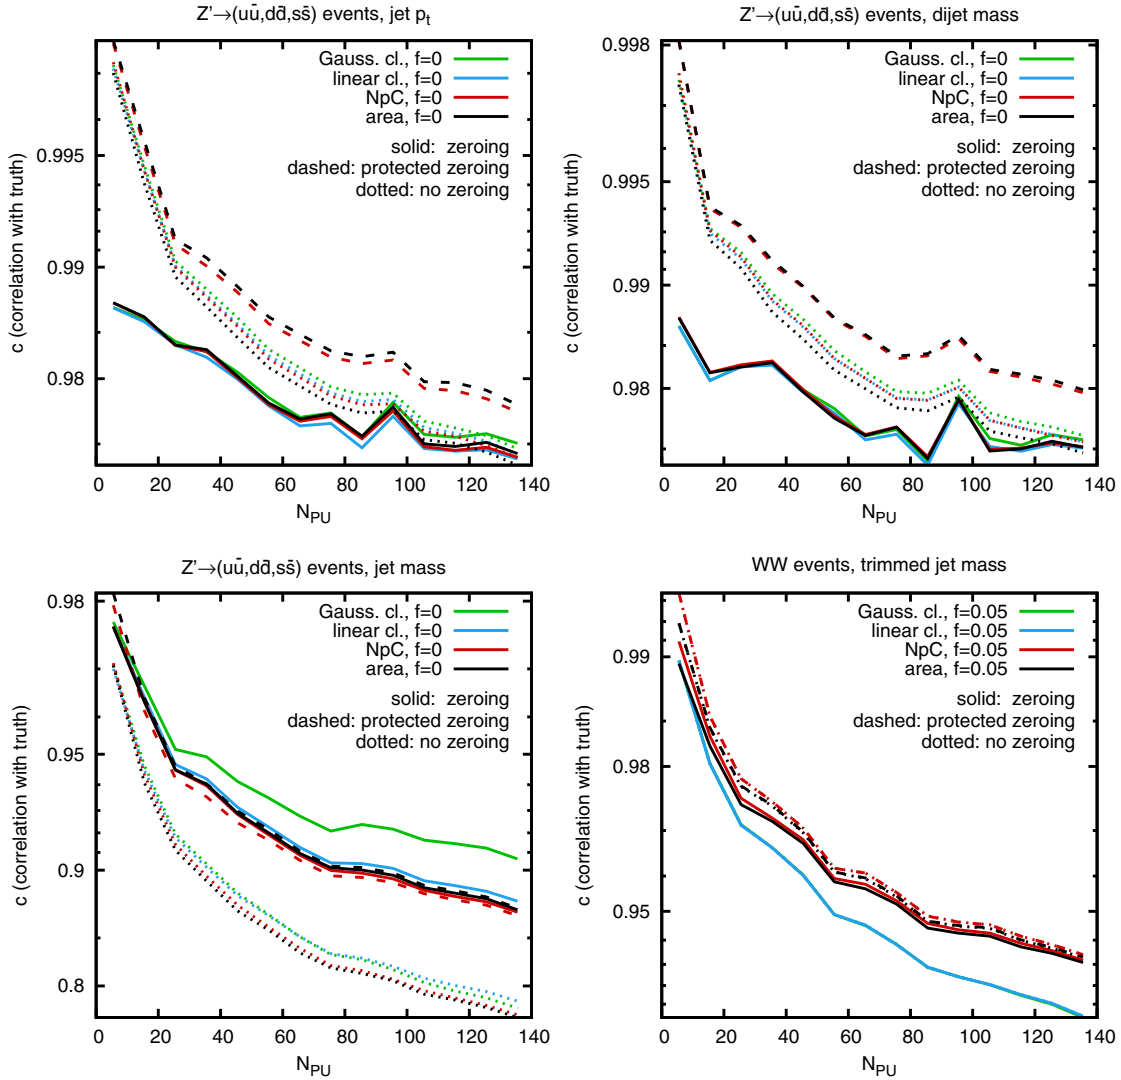


FIG. 11 (color online). As in Fig. 10, but showing correlation coefficients instead of dispersions.

of v , $\langle \Delta v \rangle$ and the standard deviation of Δv , which we write as $\sigma_{\Delta v}$. One might think there is a drawback in keeping track of two measures, in part because it is not clear which of the two is more important. It is our view that the two measures provide complementary information: if one aims to reduce systematic errors in a precision measurement then a near-zero average offset may be the most important requirement, so as not to be plagued by issues related to the systematic error in the offset. In a search for a resonance peak, then one aims for the narrowest peak, and so the smallest possible standard deviation.³⁰

The quality measure advocated in [9] is instead the correlation coefficient between v^{sub} and v^{hard} . This has the apparent simplifying advantage of providing just a

single quality measure. However, it comes at the expense of masking potentially important information: for example, a method with a large offset and one with no offset will give identical correlation coefficients, because the correlation coefficient is simply insensitive to (constant) offsets.

The correlation coefficient has a second, more fundamental flaw, as illustrated in Fig. 12. In the left-hand panel, one has a scatter plot of the dijet mass in pileup-subtracted events versus the dijet mass in the corresponding hard events, as obtained in an inclusive jet sample. There is a broad spread of dijet masses, much wider than the standard deviation of Δm_{jj} , and so the correlation coefficient comes out very high, $c = 0.988$. Now suppose we are interested in reconstructing resonances with a mass near 500 GeV, and so consider only hard events in which $450 < m_{jj} < 550$ GeV (right-hand panel). Now the correlation coefficient is 0.59, i.e. much worse. This does not reflect a much

³⁰This statement assumes the absence of tails in the Δv distribution. For some methods the long tails can affect the relevance of the standard-deviation quality measure.

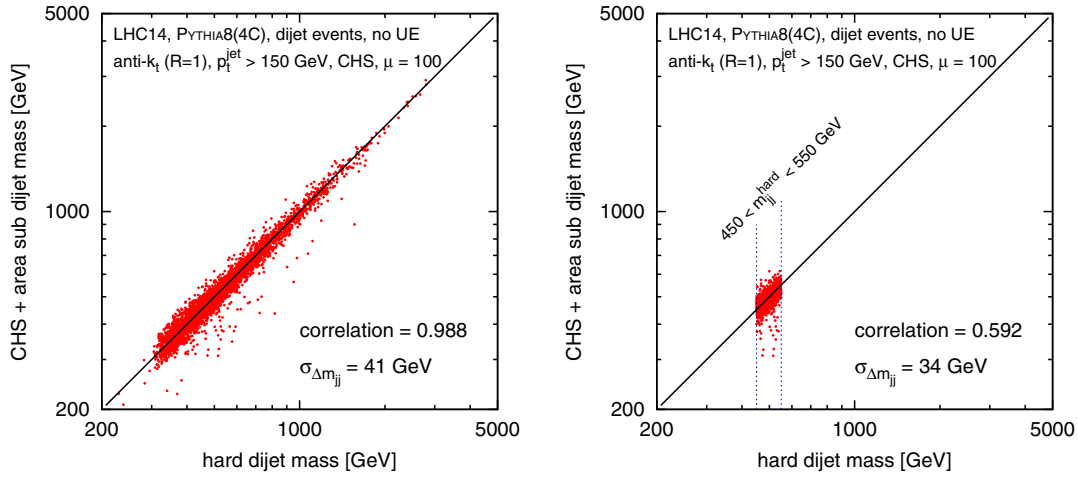


FIG. 12 (color online). (Left panel) Scatter plot of the dijet mass after the addition of an average of 100 pileup events and area-median subtraction (in CHS events) versus the dijet mass in the original hard event. The hard dijet sample and the analysis are as described in Appendix A, with a jet radius of $R = 1$. The right-hand panel is identical except for the following additional condition for the hard event: $450 < m_{jj} < 550$ GeV. Note the lower correlation coefficient, even though the lower $\sigma_{\Delta m_{jj}}$ suggests better typical subtraction in this specific mass bin.

worse subtraction: actually, $\sigma_{\Delta m_{jj}}$ is better (lower) in the sample with a limited m_{jj} window, $\sigma_{\Delta m_{jj}} = 34$ GeV, than in the full sample, $\sigma_{\Delta m_{jj}} = 41$ GeV. The reason for the puzzling decrease in the correlation coefficient is that the dispersion of m_{jj} is much smaller than before, and so the dispersion of Δm_{jj} is now comparable to that of m_{jj} : it is this, and not an actual degradation of performance, that leads to a small correlation.

This can be understood quantitatively in a simple model with two variables: let v^{hard} have a standard deviation of $\sigma_{v,\text{hard}}$, and for a given v^{hard} let v^{sub} be distributed with a mean value equal to v^{hard} (plus an optional constant offset) and a standard deviation of $\sigma_{\Delta v}$ (independent of v^{hard}). Then the correlation coefficient of v^{hard} and v^{sub} is

$$c = \frac{\sigma_{v,\text{hard}}}{\sqrt{(\sigma_{v,\text{hard}})^2 + \sigma_{\Delta v}^2}}, \quad (\text{C2})$$

i.e. it tends to zero for $\sigma_{v,\text{hard}} \ll \sigma_{\Delta v}$ and to 1 for large $\sigma_{v,\text{hard}} \gg \sigma_{\Delta v}$, in accord with the qualitative behavior seen in Fig. 12. The discussion becomes more involved if $\langle v^{\text{sub}} \rangle$ has a more complicated dependence on v^{hard} or if $\sigma_{\Delta v}$ itself depends on v^{hard} , for example as is actually the case for the dijet mass with the analysis of Appendix A.

The main conclusion from this appendix is that correlation coefficients mix together information about the quality of pileup mitigation and information about the hard event sample being studied. It is then highly nontrivial to extract just the information about the pileup subtraction. This can lead to considerable confusion, for example, when evaluating the robustness of a method against the choice of hard sample. Overall therefore, it is our recommendation that one consider direct measures of the dispersion introduced by the pileup and subtraction and not correlation coefficients. In cases with severely non-Gaussian tails in the Δv distributions it can additionally be useful to consider quality measures more directly related to the peak structure of the Δv distribution.

- [1] ATLAS Collaboration, Report No. ATLAS-CONF-2013-083, 2013.
- [2] CMS Collaboration, Report No. CMS-DP-2012-006, 2012.
- [3] M. Cacciari and G. P. Salam, Pileup subtraction using jet areas, *Phys. Lett. B* **659**, 119 (2008).
- [4] M. Cacciari, G. P. Salam, and G. Soyez, The catchment area of jets, *J. High Energy Phys.* **04** (2008) 005.

- [5] CMS Collaboration, Report No. CMS-PAS-PFT-09-001, 2009.
- [6] G. Aad *et al.* (ATLAS Collaboration), Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC, *Phys. Lett. B* **716**, 1 (2012).
- [7] J. Colas *et al.* (ATLAS Liquid Argon Calorimeter Collaboration), Position resolution and particle identification with

- the ATLAS EM calorimeter, *Nucl. Instrum. Methods Phys. Res., Sect. A* **550**, 96 (2005).
- [8] M. Cacciari, G.P. Salam, and G. Soyez, <https://cern.ch/gsalam/talks/repo/2011-CMS-week.pdf> (unpublished).
- [9] D. Krohn, M. D. Schwartz, M. Low, and L. T. Wang, Jet cleansing: Pileup removal at high luminosity, *Phys. Rev. D* **90**, 065020 (2014).
- [10] CMS Collaboration, Report No. CMS-PAS-JME-13-005, 2013.
- [11] ATLAS Collaboration, Report No. ATL-PHYS-PUB-2014-001, 2014.
- [12] M. Cacciari, J. Rojo, G.P. Salam, and G. Soyez, Quantifying the performance of jet definitions for kinematic reconstruction at the LHC, *J. High Energy Phys.* **12** (2008) 032.
- [13] A. Altheimer, A. Arce, L. Asquith, J. Backus Mayes, E. Bergeaas Kuutmann, J. Berger, D. Bjergaard, L. Bryngemark *et al.*, Boosted objects and jet substructure at the LHC, *Eur. Phys. J. C* **74**, 2792 (2014).
- [14] 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).
- [15] D. Krohn, J. Thaler, and L.-T. Wang, Jet trimming, *J. High Energy Phys.* **02** (2010) 084.
- [16] G. Soyez, G.P. Salam, J. Kim, S. Dutta, and M. Cacciari, Pileup Subtraction for Jet Shapes, *Phys. Rev. Lett.* **110**, 162001 (2013).
- [17] M. Cacciari, P. Quiroga-Arias, G.P. Salam, and G. Soyez, Jet fragmentation function moments in heavy ion collisions, *Eur. Phys. J. C* **73**, 2319 (2013).
- [18] M. Cacciari, J. Rojo, G.P. Salam, and G. Soyez, Jet reconstruction in heavy ion collisions, *Eur. Phys. J. C* **71**, 1539 (2011).
- [19] M. Cacciari, G.P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [20] M. Cacciari, G.P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [21] T. Sjostrand, S. Mrenna, and P.Z. Skands, A brief introduction to PYTHIA8.1, *Comput. Phys. Commun.* **178**, 852 (2008).
- [22] M. H. Seymour, Searches for new particles using cone and cluster jet algorithms: A comparative study, *Z. Phys. C* **62**, 127 (1994).
- [23] M. Cacciari, G.P. Salam, G. Soyez, On the use of charged-track information to subtract neutral pileup, [arXiv:1404.7353v1](https://arxiv.org/abs/1404.7353v1).
- [24] D. Krohn, M. Low, M. D. Schwartz, L.-T. Wang, Jet cleansing: pileup removal at high luminosity, [arXiv:1309.4777v1](https://arxiv.org/abs/1309.4777v1).
- [25] M. D. Schwartz, <https://indico.cern.ch/event/306155/session/2/contribution/1/material/slides/0.pdf> (unpublished).
- [26] See <http://fastjet.hepforge.org/contrib/>.
- [27] CMS Collaboration, Report No. CMS-PAS-JME-14-001, 2014.
- [28] O. L. Kodolova *et al.*, Report No. CMS-NOTE-1998-063, 1998; V. Gavrilov, A. Oulianov, O. Kodolova, and I. Vardanian, Report No. CMS-RN-2003-004, 2003; O. Kodolova, I. Vardanian, A. Nikitenko, and A. Oulianov, The performance of the jet identification and reconstruction in heavy ions collisions with CMS detector, *Eur. Phys. J. C* **50**, 117 (2007).
- [29] M. Cacciari, G.P. Salam, and G. Soyez, SoftKiller, a particle-level pileup removal method, *Eur. Phys. J. C* **75**, 59 (2015).
- [30] D. Bertolini, P. Harris, M. Low, and N. Tran, Pileup per particle identification, *J. High Energy Phys.* **10** (2014) 59.
- [31] M. Dasgupta, A. Fregoso, S. Marzani, and G.P. Salam, Towards an understanding of jet substructure, *J. High Energy Phys.* **09** (2013) 029.
- [32] M. Dasgupta, A. Fregoso, S. Marzani, and A. Powling, Jet substructure with analytical methods, *Eur. Phys. J. C* **73**, 2623 (2013).
- [33] CMS Collaboration, Underlying event subtraction for particle flow, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsDPSUESubtractionPF>.
- [34] P. Berta, M. Spousta, D. W. Miller, and R. Leitner, Particle-level pileup subtraction for jets and jet shapes, *J. High Energy Phys.* **06** (2014) 092.
- [35] M. Cacciari, G.P. Salam, and G. Soyez, public code for validation of a subset of the results in this paper, <https://github.com/npctests/1404.7353-validation>.
- [36] ATLAS Collaboration, Report No. ATLAS-CONF-2012-148, 2012.
- [37] CMS Collaboration, Report No. CMS-PAS-EXO-12-059, 2013.
- [38] J. Alcaraz Maestre *et al.* (SM and NLO Multileg and SM MC Working Groups Collaboration), The SM and NLO Multileg and SM MC Working Groups: Summary report, [arXiv:1203.6803](https://arxiv.org/abs/1203.6803).