Recent progress in defining jets

Grégory Soyez[a*]

[a]Brookhaven National Laboratory, Building 510, Upton NY 11973, USA

From dedicated QCD studies to new physics background estimation, jets will be everywhere at the LHC. In these proceedings, we discuss two important recent series of improvements. In the first one, we introduce new algorithms and new implementations of previously existing algorithms, in order to cure limitations of their predecessors and to satisfy fundamental requirements. In the second part, we show that it is of prime importance to carefully choose the jet definition — algorithm and parameters — to optimise kinematic reconstructions at the LHC. Noticeably, we show that while at scales around 100 GeV, $R \simeq 0.5$ is an appropriate choice, clustering at the TeV scale requires $R \simeq 1$ for optimal efficiency. We finally show that our results are valid in the presence of pileup, provided that a subtraction procedure is applied.

1. Introduction

In collider physics, as soon as the final state involves hadronic particles, jets become fundamental objects present in many studies. Even if the idea of a jet as a bunch of collimated partons or hadrons is what one always keeps in mind, the concept of a parton itself is ambiguous and as a consequence different jet definitions exist, possibly leading to different sets of jets for the same hadronic event.

In practice, many jet definitions, which we shall review in Section 2, have been used for recent analysis. It turns out however that, among them, some fail to satisfy the fundamental requirements agreed upon in 1990 (Section 2.2).

The first series of results presented in these proceedings addresses those failures and give solutions that have been proposed recently. This includes the FastJet implementation of the k_t algorithm, Section 2.3, as well as the the SISCone and anti-k_t algorithms, Sections 2.4 and 2.5. Finally, in Section 2.6, we will briefly discuss a filtering technique using jet substructure that has been recently proposed to reduce sensitivity to the underlying event.

Since one has the choice between different jet definitions to perform jet analysis, a legitimate question is which of them is best suited for a given analysis one wants to perform. In Section 3 of these proceedings, we answer that question in the case of simple kinematic reconstructions at the LHC. To quantify the efficiency of a jet definition, we introduce a figure of merit that is directly related to an effective luminosity ratio. We will see that not being careful enough in the choice of the jet algorithms and its parameters (basically the radius R) can lead to important consequences for potential discoveries at the LHC.

2. Meeting fundamental requirements

2.1. Defining jets in the 20[th] century

To begin with, let us briefly review the algorithms that have been widely used over the past two decades for jet reconstruction in pp collisions. Generally speaking, they fall in two categories that we discuss hereafter.

Successive recombinations. The first family of jet clustering algorithms works by defining a distance between any pair of objects and a beam distance for every object. One identifies the smallest distance; if it is a beam distance, the object is called a jet and removed from the event, otherwise, the two objects are recombined into a single one. The procedure is repeated until no objects are left in the event. Two well-known examples of recombination algorithms are the k_t

[1, 2] and Cambridge/Aachen (C/A) [3] algorithms, using the distance

$$d_{ij} = \min(k_{t,i}^{2p}, k_{t,j}^{2p}) \left(\Delta y_{ij}^2 + \Delta \phi_{ij}^2 \right),$$
$$d_{iB} = R^2 \, k_{t,i}^{2p},$$

where $p = 1$ ($p = 0$) corresponds to the k_t (C/A) case.

Cone. The cone algorithms aim at defining jets as dominant directions of energy flow. To achieve that goal, one defines the concept of *stable cone* as a circle of fixed radius in the (y, ϕ) plane such that the sum of the 4-momenta of the particles inside it points in the direction of its centre. Most of the cone algorithms used so far are *seeded* in the sense that the search for stable cones starts from a given set of seeds from which one iterates the cone contents until it is stable. Since stable cones might overlap, one cannot define them as jets directly. Two different techniques are in use to overcome this overlapping problem.

The first option, referred to as *cone algorithms with split–merge*, first identifies the set of all stable cones, then performs a split–merge procedure on them. The latter repeatedly identifies the two hardest (originally in E_t) overlapping cones; if their overlap passes a threshold fraction (the overlap parameter f) of the softer of the selected cones, they are merged. Otherwise, they are split by associating every particle to the cone to whose centre it is closer.

For the seeded versions of the cone algorithms with split–merge, one usually starts with all the particles in the event (usually with a p_t threshold) as seeds, as is the case for the CDF JetClu [4] and ATLAS Cone algorithms. A step forward is to add as new seeds the midpoints between all pairs of stable cones found after this first pass. This is the case for the majority of the recently-used cone algorithms [5], noticeably the CDF MidPoint, D0 run II Cone and PxCone algorithms.

The second solution to the problem of overlapping stable cones, referred to as *iterative cone algorithms with progressive removal (IC-PR)*, starts by iterating the stable cone search from the hardest seed in the event. The resulting stable cone is called a jet and its contents are removed from the event. One then proceeds by iterating from the hardest remaining seed, until all the particles are clustered. The characteristic feature of this type of cone algorithm is that it produces hard jets that are circular and soft-resilient. In other words, the addition of soft particles does not modify the shape of the hard jets, which is sometimes seen as an advantage for calibrating the jets. The CMS Iterative Cone algorithm [6] falls in this sub-category.

2.2. Fundamental requirements

Back in 1990, a list of fundamental requirements that every jet definition has to fulfil was agreed upon [7]. It is known as the *SNOWMASS accord* and consists of the 5 following criteria:

1. simple to implement in an experimental analysis;
2. simple to implement in the theoretical calculations;
3. defined at any order of perturbation theory;
4. yields finite cross section at any order of perturbation theory;
5. yields a cross section that is relatively insensitive to hadronisation.

As consequences, because of constraints 1 and 5, we want the implementation of the algorithm to be fast enough and as insensitive as possible to the underlying event (UE) so that it can be used in experimental analysis. Also, the requirement that the cross section remains finite at any order of perturbation theory implies that the algorithm has to be infrared and collinear (IRC) safe. Indeed, if it was not the case, cancellation between real emissions and virtual corrections would not happen properly, leading to divergences in perturbative cross-sections.

2.3. Speeding up the k_t algorithm

Beside its rather large sensitivity to the UE, one argument that was sometimes used against the k_t algorithm was its relatively long computing time. The *KtJet* [8] and *KtClus* [1] implementations, of complexity $\mathcal{O}\left(N^3\right)$ where N is the number of particles in the event, have a clustering time around 1 second for $N = 1000$, compared to about 0.2 seconds for the MidPoint cone with a 1 GeV seed threshold.
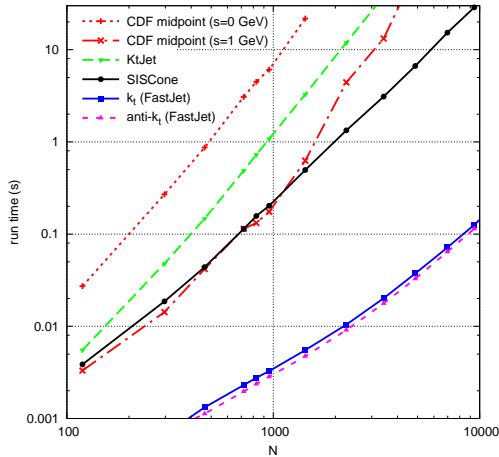
Remarkably, the complexity can be reduced

Figure 1. Clustering time for various algorithms.

Figure 2. Stable cones found by the MidPoint algorithm for (a) a 3-particle event, (b) the same event with an additional infinitely soft gluon.

to $\mathcal{O}\left(N\log(N)\right)$ using computational-geometry techniques[2] [9]. As seen on figure 1 (k_t (FastJet) curve), this leads to a considerable improvement.

2.4. SISCone and the IR unsafety of the MidPoint algorithm

It had already been noticed that the JetClu algorithm had some IR unsafety problems, *e.g.* when two hard particles were distant by more that R and less that $2R$, an additional infinitely soft gluon added between the particles could change the clustering from 2 jets to 1 jet, yielding unreliable cross-sections at NLO in the inclusive jet cross-section. The MidPoint algorithm was then introduced to cure that problem.

Unfortunately, this is not the end of the story. Even if situations with 2 hard particles in a common vicinity (plus another one, hadronic or electroweak, to balance p_t) are IR safe provided one uses the MidPoint algorithm instead of JetClu, the problem has just been shifted to situations with 3 hard particles in a common vicinity. This is illustrated in figure 2, where we have clustered a 3-hard-particle event twice — with, fig. 2(b), and without, fig. 2(a), an additional infinitely soft gluon — and find different stable cones in
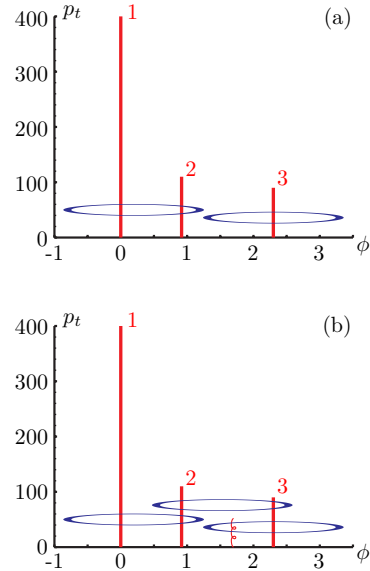
the two cases. This means that the MidPoint algorithm is also IR unsafe, though one order further in the perturbative expansion in the strong coupling than the JetClu algorithm.

To solve that problem to all orders in the perturbative expansion, we first notice that the IR unsafety comes from the fact that in the event without the soft gluon, the stable cone enclosing particles 2 and 3 has been missed. Since the mathematically well-defined set of stable cones is IRC safe, *i.e.* it changes neither when splitting a particle collinearly nor when adding infinitely soft particles (up to harmless stable cones made only of those soft particles), finding a method that provably identifies all stable cones, guarantees IRC safety[3].

We then observe that every circular enclosure

---

[2]For $N \lesssim 5000$, there exists an $\mathcal{O}\left(N^2\right)$ implementation that turns out to be faster than the $N\log(N)$ one [9].

[3]We also need to be careful about the IRC safety of the split–merge procedure. Actually, a couple of technical issues, such as the choice of the variable used to order the overlapping cones, have to be dealt with. We will not address them here (see [10] for details).

of given radius R in the (y, ϕ) plane can be translated in any direction until it touches one point, then rotated around that point until it touches a second one, without changing its contents. Therefore, enumerating all pairs of points, and for every pair of points considering the two circles of radius R they define and the four possible inclusion/exclusion states of the edge particles, we enumerate all possible enclosures. For each of them we can then test if it is stable or not, which solves our problem. The complexity of this algorithm is $\mathcal{O}\left(N^3\right)$ (a factor N^2 coming from the enumeration of the pairs of parent points and an additional factor of N to test the stability of every enclosure). It is actually possible, using extra geometric observations, to improve that complexity to $\mathcal{O}\left(N^2 \log(N)\right)$. This has been implemented [10] in a new algorithm named *SISCone* (Seedless Infrared Safe Cone). Fig. 1 shows that SISCone runs faster than the IR-unsafe MidPoint algorithm (of order N^3), even when a seed threshold of 1 GeV is applied.

2.5. Anti-k_t and the collinear unsafety of the IC-PR algorithm

In a similar way as for the case of the MidPoint algorithm, we can show that the iterative cone algorithm with progressive removal (IC-PR) also suffers from divergences in the perturbative series, at the same order as MidPoint, this time due to collinear unsafety. In order to see that, first consider the event of fig. 3(a). Iterating from the hardest seed gives one jet containing all particles. If one splits the hardest of these particles into two collinear ones, fig. 3(b), iteration starts with the leftmost particle and 2 jets are found. This means that the IC-PR is collinear unsafe at the level of 3 particles (+1 to balance p_t).

To address this issue, we will go back to the recombination-type algorithms. We have already mentioned that setting $p = 0$ or 1 in eq. (1) reproduces the C/A and k_t algorithms respectively. We now introduce the *anti-k_t algorithm*, corresponding to $p = -1$ in eq. (1) [11].

At first sight, it is not obvious what this new, IRC-safe, algorithm has to do with the IC-PR. However, since hard particles will be associated a small anti-k_t distance, they will grow in circles,
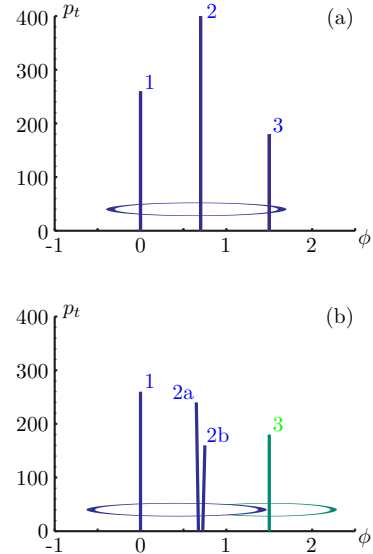

Figure 3. Jets found by the iterative cone for (a) a 3-particle event, (b) the same event with a collinear splitting.

clustering softer particles in their vicinity up to a distance R. This thus leads to soft-resilient hard jets, *i.e.* the boundary of the hard jets is not affected by soft radiation, the precise characteristic feature of the IC-PR.

Finally, the anti-k_t algorithm is also amenable to a fast implementation (see fig. 1) using the same techniques as for the k_t algorithm (see Section 2.3).

2.6. Filtering

One of the promising potential improvement of jet clustering is to make use of the jet substructure. For example, we can use the following filtering technique to reduce the contamination due to the underlying event:

1. Cluster the event using, *e.g.*, the C/A algorithm with a radius R,
2. For each jet, recluster it using a smaller radius R_{sub} and keep only the n_{sub} hardest jets as part of the initial jet, throwing the other subjets.

The aim of filtering is to remove contamination due to soft background like the underlying event while keeping as much as possible of the perturbative radiation. This has already proven [12] to be efficient in Higgs searches in the $b\bar{b}$ channel.

In what follows we shall use $R_{\mathrm{sub}} = R/2$ and $n_{\mathrm{sub}} = 2$, though a more extensive study of the effects of these parameters would be interesting.

3. Quantifying kinematic reconstruction efficiency

Now that we dispose of 5 IRC-safe jet algorithms — the k_t, C/A, anti-k_t, SISCone and C/A+filtering algorithms, all available from Fast-Jet [13] — we may ask, given an analysis involving jets we want to perform, which jet definition, *i.e.* the jet algorithm and its parameters, is best suited. In this Section we address that question for the case of kinematic reconstructions at the LHC.

We first introduce a series of benchmark processes we will investigate, then a figure of merit that allows one to quantify the performance of a jet definition and finally present our results, both with and without including pileup. For a more extensive discussion, see [14, 15].

3.1. Benchmark processes

We study the following 3 processes:

- $Z' \rightarrow q\bar{q}$ as a source of quark jets. We reconstruct the Z' from the 2 hardest jets in the event (imposing a maximal rapidity difference $|\Delta y| \leq 2$ between them).
- $H \rightarrow gg$ as a source of gluon jets. We reconstruct the H from the 2 hardest jets in the event (imposing again $|\Delta y| \leq 2$).
- a $t\bar{t}$ pair with fully hadronic decay into 6 jets (4 jets from the 2 W bosons and 2 jets from the b and $\bar{b}$ quarks). The two W's and then the 2 tops are reconstructed from the 6 hardest jets in the event[4].

---

[4]Practically, b-jets are tagged assuming the b mesons are stable; then the two W are reconstructed by pairing the 4 remaining jets so as to minimise $(M_{i_1 i_2} - M_W)^2 + (M_{i_3 i_4} - M_W)^2$; finally, the two top jets are reconstructed by matching the b-jets with the W so as to minimise the mass difference between the 2 top candidates.

For the case of quark and gluon jets, we can study the scale dependence by varying the mass of the Z'/H boson (in practice, between 100 GeV and 4 TeV). The last case, fully hadronic $t\bar{t}$ decay, allows us to study the relevance of our results in more complex environments, where one risks tensions between resolving the jets and capturing the perturbative radiation. In each of these situations we have generated a sample of events using Pythia 6.4 tune DWT.

3.2. Figure of merit

In order to quantify which jet definition is performing better than another, we need a figure of merit. Since jets are expected to represent an original parton, one might be tempted to quantify the efficiency by comparing the jets to the initial partons. But because partons are an ill-defined concept (especially at NLO), this is not robust enough.

Another option would be to fit a given distribution, *e.g.* a Gaussian, to the peak in the mass spectrum. But since the shape of the peak can be asymmetric, this would not produce a reliable result either.

For these reasons, we will quantify the efficiency of a jet definition based on the observation that, if two peaks have similar number of events then the narrower is the better one. We therefore introduce the quality measure $Q^w_{f=z}$ as the *width of the smallest possible window that contains a fraction $f = z$ of the events*. The smallest values of $Q^w_{f=z}$ should then correspond to the most efficient jet definitions.

This intuitively does what we want in the sense that a "better" jet definition should contain a given fraction of the events in a smaller window and therefore have a smaller $Q^w_{f=z}$.

This quality measure can be related to a variation of luminosity needed to maintain a constant significance for a signal relative to background. For a jet definition JD, the latter is defined as $\Sigma(\mathrm{JD}) \equiv N^{\mathrm{JD}}_{\mathrm{signal}}/\sqrt{N^{\mathrm{JD}}_{\mathrm{bkgd}}}$. Assuming a flat background, we thus have

$$\frac{\Sigma(\mathrm{JD}_1)}{\Sigma(\mathrm{JD}_2)} = \left[\frac{N^{\mathrm{JD}_2}_{\mathrm{bkgd}}}{N^{\mathrm{JD}_1}_{\mathrm{bkgd}}}\right]^{1/2} = \left[\frac{Q^w_{f=z}(\mathrm{JD}_2)}{Q^w_{f=z}(\mathrm{JD}_1)}\right]^{1/2}$$
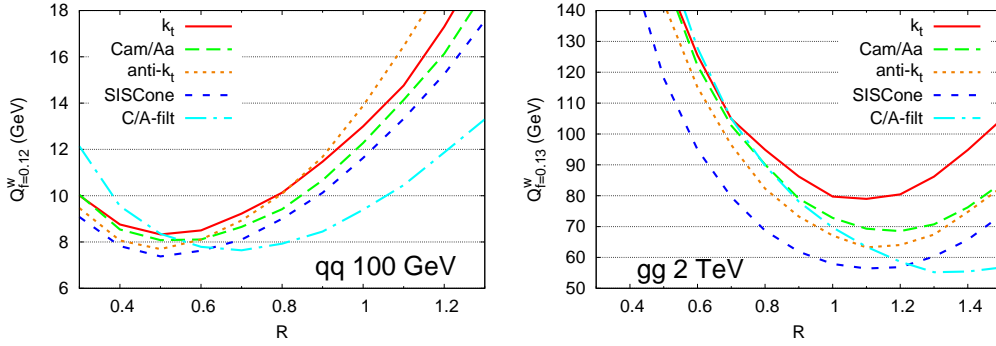
Figure 4. Result for the quality measure for quark jets at 100 GeV (left) and gluon jets at 2 TeV (right). The results are displayed as a function of R and each curve corresponds to a given algorithm.

for two jet definitions JD_1 and JD_2. This relation comes from the fact that the number of signal events is fixed by the fraction $f = z$ in the quality measure, while the background is directly proportional to the width of the window *i.e.* to $Q^w_{f=z}$. A "better" definition, *i.e.* a smaller $Q^w_{f=z}$, thus corresponds to a larger discriminating power.

We can then define an effective luminosity ratio

$$
\begin{aligned}
\rho_{\mathcal{L}}(JD_2/JD_1) &\equiv \frac{\mathcal{L}(\text{needed with } JD_2)}{\mathcal{L}(\text{needed with } JD_1)} \quad (1) \\
&= \left[\frac{\Sigma(JD_1)}{\Sigma(JD_2)}\right]^2 = \frac{Q^w_{f=z}(JD_2)}{Q^w_{f=z}(JD_1)}.
\end{aligned}
$$

This means that a jet definition JD_1 with a quality measure twice as large as JD_2 will need twice the integrated luminosity in order to achieve the same discriminating power as JD_2.

3.3. Results without pileup

For the processes presented in Section 3.1, we have clustered our event samples using the 5 IRC safe jet algorithms and varying the parameter[5] R between 0.1 and 1.5. We can then compute the quality measure in each of those cases (see figure 4), which allows us to (i) compare different algorithms, and (ii), for a given algorithm, find the optimal value R_{best} of the parameter R.

Figure 5 summarises our results in a more compact and physical way: we have plotted the effective luminosity ratios corresponding to a series of selected processes: quark and gluon jets at 100 GeV and 2 TeV, and top reconstruction in the $t\bar{t}$ sample. For every process, $\rho_{\mathcal{L}}$ has been normalised to the best jet definition for that process.

We observe a few important features: first of all, in general, SISCone and C/A with filtering tend to perform slightly better than the k_t, C/A and anti-k_t algorithms, especially at higher scales (with the exception of top reconstruction, where all algorithms have similar performances). For example, for the 2 TeV gluon case, choosing the k_t algorithm instead of SISCone or C/A with filtering which are the preferred choices, translates into a cost of nearly 50% in the effective luminosity ratio.

Furthermore, the efficiency of a jet definition strongly depends on R and not choosing the preferred value R_{best} can be very costly. On top of that, R_{best} varies significantly from one process to another. It increases from $R_{\text{best}} \simeq 0.5$ at small scales up to $R_{\text{best}} \simeq 1$ at TeV scales. R_{best} is also bigger for gluon jets than for quark jets. Practically, using SISCone with $R = 0.5$, the preferred choice for quark jets at 100 GeV, to cluster gluon jets at 2 TeV costs a factor of 2 in $\rho_{\mathcal{L}}$ compared to the best definition for that process. Conversely, using the preferred value of R for gluon jets at
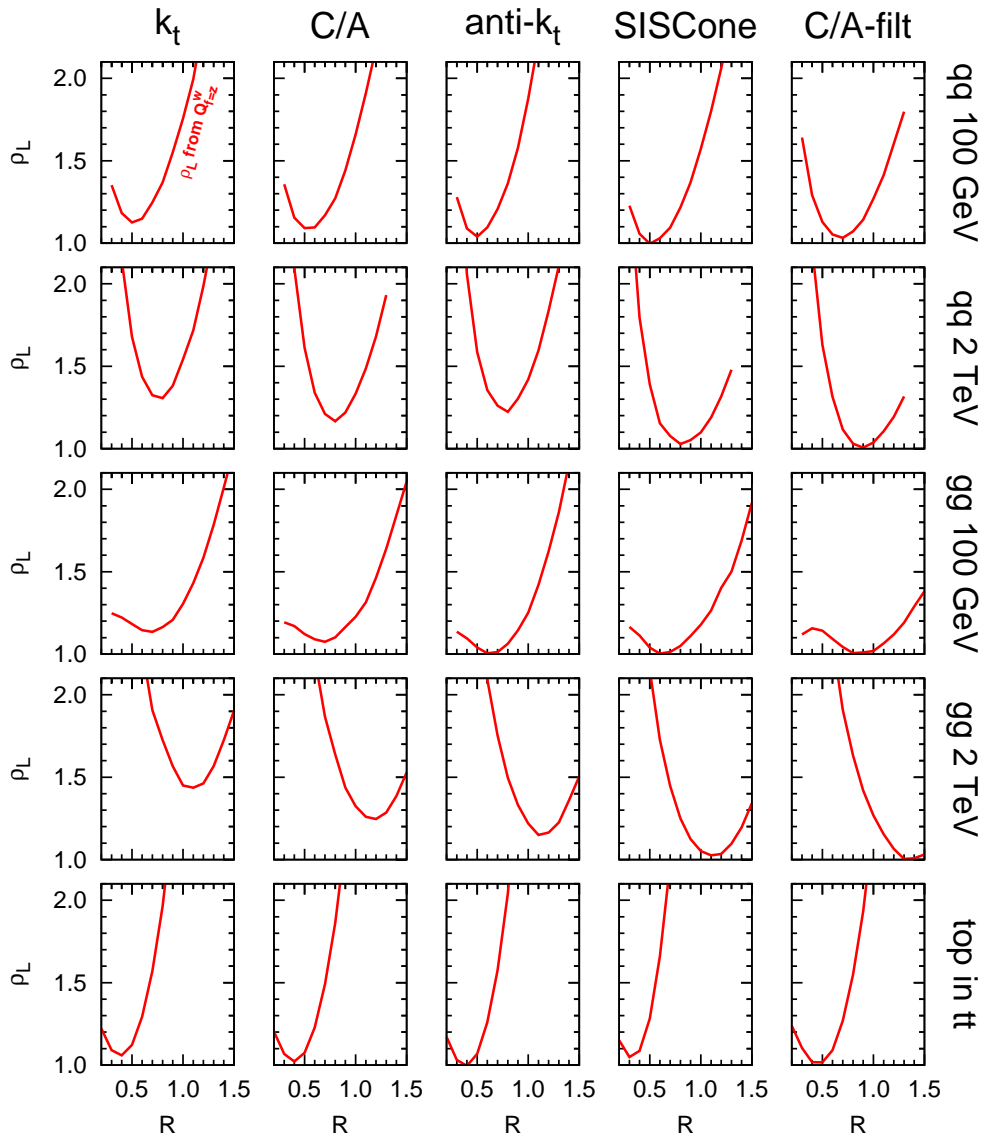
---

[5]In the case of SISCone, the overlap threshold f has been fixed to 0.75, the preferred choice at the time being.

Figure 5. Effective luminosity ratios as a function of the parameter R. Each line corresponds to a different process and each column to a different algorithm. For each process, $\rho_{\mathcal{L}}$ is normalised to the best possible definition for that process.

2 TeV, *i.e.* SISCone with $R = 1.1$ or C/A with filtering and $R = 1.3$, to cluster quark jets at 100 GeV also leads to an increase in $\rho_\mathcal{L}$ by at least 50%.

3.4. Results with pileup

Pileup corresponds to the fact that multiple pp interactions can happen at the same time. For the case of the LHC at the design luminosity, one typically has an average of about 25 collisions per bunch crossing. This produces a large number of additional soft particles that form a reasonably uniform background over the detector.

For our concerns, pileup has two consequences. Because it adds background to every jet, the p_t of the jets will be overestimated, and the position of a mass peak in our kinematic reconstruction will be shifted towards larger masses. Also, because the amount of pileup varies from one event to another, the reconstructed peak will be smeared, an effect that directly affects our quality measure.

We therefore want a (simple and generic enough) method to subtract the contamination from the pileup background on an event-by-event basis. In practice, we will follow the subtraction method suggested in [16]: for each jet j in an event, we first compute its (4-vector) area [17] A_j^μ. The 4-momentum of the subtracted jet is then obtained using

$$p_{j,\mathrm{sub}}^\mu = p_j^\mu - \rho\, A_j^\mu, \qquad (2)$$

where ρ is the average density of pileup per unit area. Since the density of background is reasonably uniform, we estimate ρ for each event using [16] $\rho = \mathrm{median}\{p_{t,j}/A_{t,j}\}$, where all the jets (up to a maximal rapidity) are included in the computation of the median.

In practice, we have considered the same benchmark processes as in the study without pileup. Pileup is added to each event by choosing a random (Poisonian) number of minimum bias events, which are also generated with Pythia 6.4 tune DWT. We apply the same reconstruction procedure as above and compute the quality measure $Q_{f=z}^w$ with pileup effects subtracted or not.

We illustrate our results by showing on figure 6 the effective luminosity ratio $\rho_\mathcal{L}$ as obtained for a few representative processes. We compare the results with pileup (subtracted or not) with the corresponding quality measures obtained without pileup (solid curves).

The first observation is that if pileup is not subtracted (dashed curves), it causes a large degradation of the quality measure. The preferred R is also shifted to smaller values, which one can explain by the fact that the contamination due to pileup background is smaller at small R. The second element of information comes from the case where we apply our subtraction method (dotted curves). Obviously, even if the subtraction is not perfect — the quality measure is still larger than before pileup addition — one sees a significant improvement compared to the situation without subtraction. This means that the subtraction method gives narrower peaks, *i.e.* the smearing due to pileup fluctuations between events is strongly reduced by the subtraction. On top of that, after the subtraction has been performed, the preferred value for R obtained in the situation without pileup is no longer strongly disfavoured to the profit of a smaller R. This is important, since it implies that our conclusions from Section 3.3 are still valid in the presence of pileup provided one uses subtraction.

4. Conclusions

The first point we have addressed concerns the jet algorithms themselves. As summarised in table 1, some of the commonly used algorithms fail to satisfy the fundamental requirements of infrared and collinear safety. We have introduced the SISCone and the anti-k_t algorithm to cure infrared unsafeties of the MidPoint and collinear unsafeties of the IC-PR algorithm. We have also mentioned a new implementation of the k_t algorithm leading to a sizeable improvement in speed.

We have then addressed the problem of quantifying the performances of the jet definitions in the case of invariant mass reconstruction at the LHC. From that study one learns important messages showing that optimisation of the jet definition can significantly improve the potential for an early discovery at the LHC. We have shown that the preferred jet definition varies significantly from
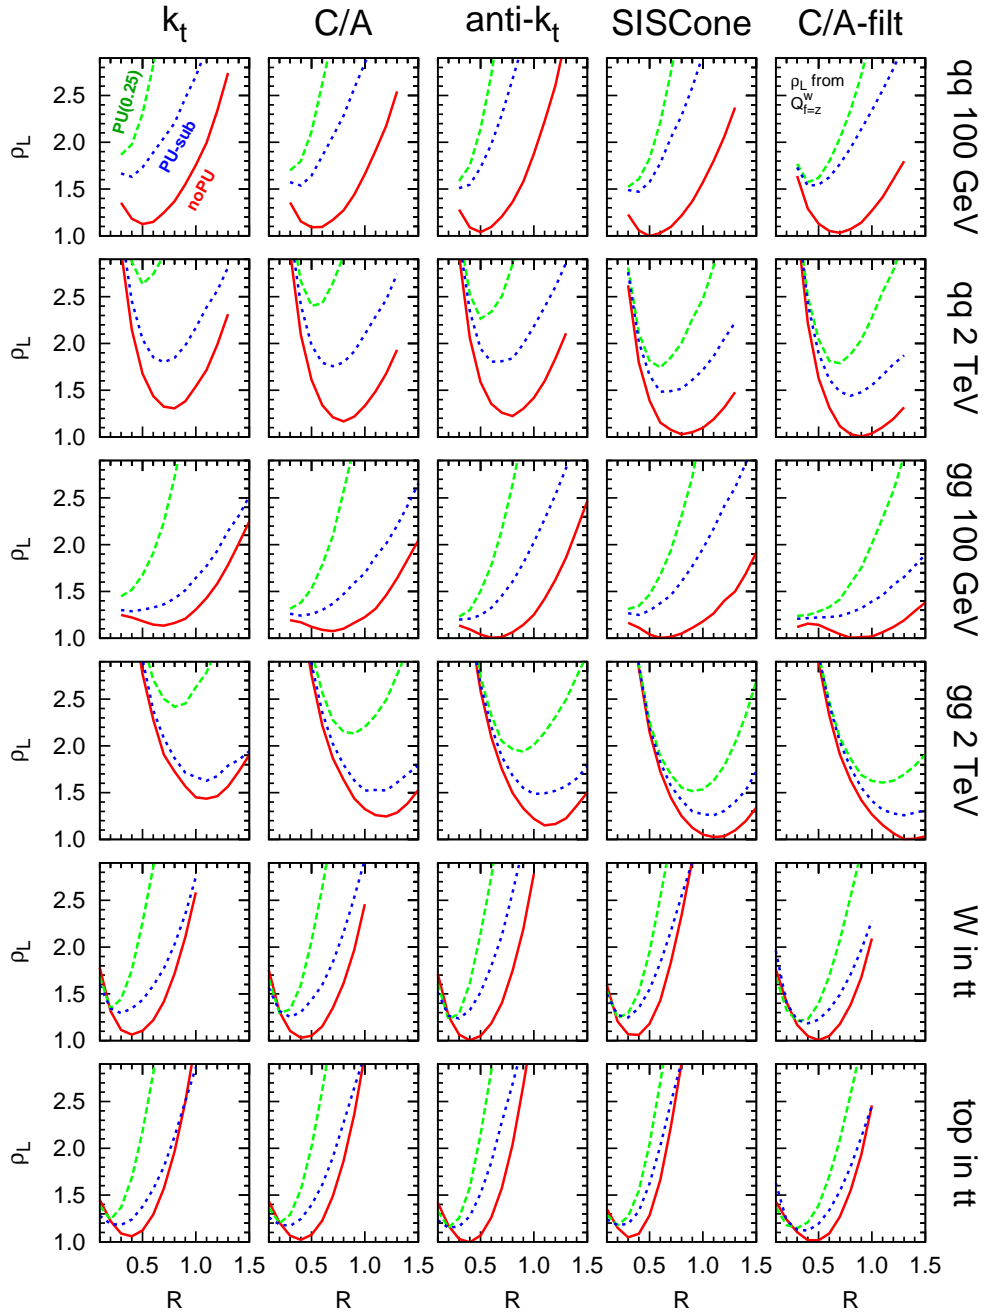
Figure 6. Same as for figure 5, now including the effects of pileup. The solid (red) curve corresponds to the situation without pileup, the dashed (green) one to the case of unsubtracted high-luminosity pileup and the dotted (blue) curve to the situation where pileup is subtracted.

Grégory Soyez

Algorithm	Type	IRC
inclusive k_t [1, 2]	$SR_{p=1}$	OK
Cambridge/Aachen [3]	$SR_{p=0}$	OK
anti-k_t [11]	$SR_{p=-1}$	OK
SISCone [10]	SC-SM	OK
CDF JetClu [4]	IC_r-SM	IR_{2+1}
CDF MidPoint [5]	IC_{mp}-SM	IR_{3+1}
D0 Run II cone [5]	IC_{mp}-SM	IR_{3+1}
ATLAS Cone	IC-SM	IR_{2+1}
CMS Iterative Cone [6]	IC-PR	$Coll_{3+1}$

Table 1
Overview of some jet algorithms used recently in experimental or theoretical work: $SR_{p=x}$ = sequential recombination (with $x = 0, \pm 1$, see (1)); SC = seedless cone (finds all cones); IC = iterative cone (with midpoints mp, ratcheting r), using either split–merge (SM) or progressive removal (PR) in order to deal with overlapping stable cones. Regarding IRC status, given n hard particles in a common neighbourhood, IR_{n+1} indicates that the addition of 1 extra soft particle can modify the number of final hard jets, while $Coll_{n+1}$ means that the collinear splitting of one of the particles can modify the number of final hard jets.

one process to another. SISCone and C/A with filtering tend to perform slightly better than the other algorithms, and the preferred value for R goes from $R \simeq 0.5$ for 100 GeV jets to $R \simeq 1$ for TeV jets with a slightly larger value for gluon jets compared to quark jets. This indicates that a single choice for R is not sufficient to cover the whole kinematic range at the LHC.

Finally, if one includes the effects of pileup, we have seen that provided we use an adequate subtraction procedure, the conclusions obtained in the analysis without pileup hold once pileup is added, noticeably, there is no need to take a smaller value of R.

In the more complex case of top reconstruction, where jet clustering involves tension between catching enough radiation to reconstruct the jet energy and resolving the 3 jets coming from the top decay products, techniques like the subjet analyses might be of crucial importance and definitely deserve further studies.

REFERENCES

1. S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, Nucl. Phys. B **406** (1993) 187;
2. S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock, and B. R. Webber, Phys. Lett. B **269** (1991) 432; S. D. Ellis and D. E. Soper, Phys. Rev. D **48** (1993) 3160 [`hep-ph/9305266`].
3. Y. L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber, JHEP **08** (1997) 001 [`hep-ph/9707323`]; M. Wobisch and T. Wengler, [`hep-ph/9907280`].
4. F. Abe *et al.* [CDF Collaboration], Phys. Rev. D **45** (1992) 1448.
5. G. C. Blazey *et al.*, [`hep-ex/0005012`].
6. G. L. Bayatian *et al.* [CMS Collaboration], CERN-LHCC-2006-001; A. Bhatti *et al.* [CMS Collaboration], CMS PAS JME-07-003 in preparation (2008).
7. J. E. Huth *et al.*, FERMILAB-CONF-90-249-E (1990).
8. J. M. Butterworth, J. P. Couchman, B. E. Cox and B. M. Waugh, Comput. Phys. Commun. **153** (2003) 85 [hep-ph/0210022].
9. M. Cacciari and G. P. Salam, Phys. Lett. B **641** (2006) 57 [`hep-ph/0512210`].
10. G. P. Salam and G. Soyez, JHEP **05** (2007) 086 [`arXiv:0704.0292`]; see also `http://projects.hepforge.org/siscone`.
11. M. Cacciari, G. P. Salam and G. Soyez, JHEP **0804** (2008) 063 [`arXiv:0802.1189`].
12. J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, Phys. Rev. Lett. **100** (2008) 242001 [`arXiv:0802.2470`].
13. M. Cacciari, G. P. Salam, and G. Soyez, "FastJet." `http://www.fastjet.fr`.
14. C. Buttar *et al.*, [arXiv:0803.0678].
15. M. Cacciari, J. Rojo, G. P. Salam and G. Soyez, JHEP **0812** (2008) 032 [arXiv:0810.1304], see also `http://quality.fastjet.fr`.
16. M. Cacciari and G. P. Salam, Phys. Lett. B **659** (2008) 119 [arXiv:0707.1378 [hep-ph]].
17. M. Cacciari, G. P. Salam and G. Soyez, JHEP **0804** (2008) 005 [arXiv:0802.1188].