

Thèse d'Habilitation à Diriger des Recherches

*Préparée à l'Institut de Physique Théorique du CEA Saclay
et au Centre National de la Recherche Scientifique (CNRS)
Soutenue à l'Université Pierre et Marie Curie*

Pileup mitigation at the LHC

A theorist's view

Grégory Soyez

Institut de Physique Théorique
CEA Saclay
CNRS UMR 3681

Jury composé de:

Fawzi Boudjema (rapporteur)

Jon Butterworth (rapporteur)

Bryan Webber (rapporteur)

Abdelhak Djouadi

Marumi Kado

Samuel Wallon



Abstract

Given the high luminosities at which CERN's Large Hadron Collider (LHC) operates, multiple proton-proton collisions occur during the same bunch crossing. Hard collisions are therefore contaminated by several soft, zero-bias, ones. This effect, known as *pileup* pollutes the reconstruction of the final state of the collision.

This document is meant to provide a theoretical review of the methods employed during Run I of the LHC and at the beginning of Run II to mitigate pileup effects. I will start with a in-depth presentation of the *area–median* used for the vast majority of applications, including several refinements of the original idea, their practical (numerical) implementation and an assessment of their efficiency and robustness. I will then focus on several theoretical calculations that can provide both quantitative and qualitative information on the area–median approach.

In the specific case of boosted jets, a field that has seen a wide interest recently, a set of methods, known as *grooming* techniques has also been used. I will describe these techniques, address their performance and briefly show that they are amenable to a theoretical, analytic, understanding.

The last part of this review will focus on ideas oriented towards future pileup mitigation techniques. This includes a couple of new methods that have recently been proposed as well as a large series of alternative methods. The latter are yet unpublished and have not received the same amount of investigation than the former but they have the potential to bring new developments and further improvement over existing techniques in a future where pileup mitigation will be crucial.

Acknowledgements

The body of work presented here is the result of many scientific collaborations and human interactions. Elements of this document date back from the early days of 2007 when I started to be involved in jet physics. Trying to look back, it seems to me like the path between then and now has been paved with a pile of positive encounters which have left a trace on who I am today.

First and above all, a very special thanks goes to Matteo Cacciari and Gavin Salam for a decade worth of fruitful team work. I hardly see how any of this could have existed without them. It has been, and still is, a great adventure during which I have learned a lot about physics and human aspects of a longstanding collaboration. I can only look forwards to more of this.

There are many other scientists who contributed to various degrees in some of the results presented here. This includes direct collaborators like Jihun Kim, Paloma Quiroga-Arias, Juan Rojo, Jesse Thaler, Simone Marzani, Andrew Larkoski or Souvik Dutta. They each carry a part of my memories associated with the work detailed below.

This list can be extended to a long series of people I have interacted with and talked to over the same period. This obviously includes colleagues at the IPhT, LPTHE, BNL and CERN but also countless people met in conferences, seminars, visits and schools. This includes in particular Nestor Armesto, Mrinal Dasgupta, David Krohn, Matthew Low, Michelangelo Mangano, Guilherme Milhano, Carlos Salgado, Sebastian Sapeta, Matthew Schwartz, Liantao Wang and Urs Wiedemann.

Next, even though the title of this document might suggest otherwise, I have greatly benefited from many interactions with the experimental colleagues: Helen Caines, David d'Enterria, Guenther Dissertori, Phil Harris, Peter Jacobs, David Miller, Filip Moortgat, Matt Nguyen, Matheus Ploskon, Sal Rappoccio, Guenther Roland, Christophe Royon, Sevil Salur, Ariel Schwartzman, and, last but not least, Peter Loch from whom I have learned a great deal of experimental details in lively discussions during the Les-Houches workshop series.

There are 6 other scientists who deserve a special thanks, namely Fawzi Boudjema, Jon Butterworth, Abdelhak Djouadi, Marumi Kado, Samuel Wallon and Bryan Webber, who have accepted to be members of the jury. I know it is a heavy task and I hope they nevertheless managed to get something out of this lengthy document.

As I said above this is both a scientific and a human story. All this work has been built on a firm ground of family and friends. Their continuous presence is the crucial part which has allowed me to enjoy life at least as much as I have enjoyed the science. This includes a special thought for a few close family member who have sadly passed away during the past ten (or fifteen) years. On a much brighter side, I am infinitely grateful (and many of you should probably be as well) to Emilie who has lived with me through all of this journey, for example supporting my work at odd times and in odd places. My final thanks go to Thiago and Neil, together with apologies for having more than once traded playing games with them for research. I have no doubt that the years to come will bring many moments of great fun with both of them.

Contents

1	Introduction	7
I	The area–median approach	11
2	Description of the area–median method	13
2.1	Basic characterisation of the pileup effects	13
2.2	Effects of average shifts and resolution degradation on physical observables	16
2.3	The jet catchment areas	17
2.3.1	Passive Area	18
2.3.2	From passive to Voronoi area	19
2.3.3	Active Area, catchment area for pileup subtraction	20
2.4	Area–median pileup subtraction	22
2.4.1	Basic recipe for pileup subtraction	22
2.4.2	Positional dependence	24
2.4.3	Particle masses	25
2.4.4	Safe subtraction and positivity constraints	26
2.4.5	Extension to Charged-Hadron Subtracted (CHS) events	27
2.5	Jet shapes	27
2.6	Jet fragmentation function	29
2.7	Numerical Implementation	33
2.7.1	Clustering with jet areas	33
2.7.2	Background estimation and subtraction	36
2.7.3	Subtraction for jet shapes	39
2.7.4	Subtraction for fragmentation function moments	40
2.8	Recommendations for practical usage	41
3	Monte-Carlo validation of approach	43
3.1	Properties of the pileup density	44
3.2	Widest application: the jet transverse momentum	47
3.2.1	Subtraction techniques	47
3.2.2	Testing framework	48
3.2.3	Generic performance and rapidity dependence	50
3.2.4	Digression: back-reaction effects	53
3.2.5	Robustness and Monte-Carlo dependence	55
3.2.6	PU v. UE subtraction: an analysis on Z +jet events	57
3.2.7	Summary and discussion	58

3.3	Jet mass and jet shapes	60
4	Applications to heavy-ion collisions	65
4.1	Challenges of jet reconstruction in heavy-ion collisions	65
4.2	Details of the study	66
4.2.1	Simulation and analysis framework	66
4.2.2	Jet reconstruction: definition and subtraction	68
4.3	Results	69
4.3.1	Matching efficiency	70
4.3.2	Choice of background-estimation range	71
4.3.3	Choice of algorithm	72
4.3.4	Centrality dependence	76
4.3.5	Quenching effects	78
4.4	Further details and analytic estimates	79
4.4.1	Back reaction contribution to the dispersion	79
4.4.2	Quality of area determination	80
4.4.3	Relative importance of average shift and dispersion	80
4.4.4	Fluctuations in extracted ρ	82
4.4.5	Hard-jet bias in extracted ρ	83
4.4.6	Subtraction bias due to filtering	83
4.5	The issue of fakes	84
4.5.1	Inclusive analyses	84
4.5.2	Exclusive analyses	86
4.6	Summary and discussions	87
4.7	Jet fragmentation function	89
5	Analytic insight	93
5.1	Properties of jet areas	93
5.1.1	Passive areas	94
5.1.2	Active areas	100
5.1.3	Back reaction	111
5.1.4	Comparison to Monte-Carlo simulations	118
5.1.5	Summary	122
5.1.6	Future investigations	123
5.2	Estimation of the pileup properties	125
5.2.1	A toy model for pileup	125
5.2.2	Pileup distribution impact on ρ estimation	126
5.2.3	Perturbative impact of the hard event	127
5.2.4	Perturbative radiation and soft pileup	128
5.2.5	Fluctuations in the estimated ρ	129
5.2.6	Comparison to Monte-Carlo studies	131
II	Grooming techniques for fat jets	133
6	Main methods	135
6.1	Grooming and tagging fat jets	135

6.2	Mass-drop, filtering, trimming, pruning, soft drop	137
7	Monte-Carlo studies	143
7.1	Jet grooming configurations	144
7.2	Jet substructure performance	144
7.3	Summary	148
8	Analytical insight	151
8.1	Foreword	151
8.2	Groomed Jet Radius	152
8.2.1	Analytic calculation	152
8.2.2	Comparison to Monte Carlo	154
8.3	Non-Perturbative Contributions	156
III	Towards more efficient techniques	157
9	Challenges beyond the area–median subtraction	159
9.1	Motivation	159
9.2	Challenges of “noise-reduction” techniques	160
10	Grooming as a generic tool for pileup subtraction	163
10.1	Foreword and motivation	163
10.2	Grooming configurations and workflow	164
10.3	Results	164
10.4	Concluding remarks and future prospects	166
11	Using information from charged tracks	169
11.1	Pileup subtraction methods for CHS events	169
11.2	Monte-Carlo studies and physical discussions	173
11.2.1	Performance of NpC v. area–median	173
11.2.2	Performance of cleansing v. area–median	176
11.3	Digression: combining different methods	183
11.4	Concluding remarks and discussions	185
12	Noise reduction with the SoftKiller	187
12.1	The SoftKiller method	187
12.2	SoftKiller performance	189
12.3	Adaptation to CHS events and calorimetric events	197
12.4	Computing time	199
12.5	Discussion and conclusions	201
13	A (rich) forest of ideas to explore	203
13.1	Alternative noise-reduction methods based on particle areas	203
13.1.1	List of methods	204
13.1.2	Monte-Carlo studies	205
13.2	SoftKiller improvements	212
13.2.1	Improved R dependence	212

13.2.2 Combination with (protected) zeroing	215
14 Conclusion and perspectives	219
A Jet algorithms and definitions	227
A.1 Recombination algorithms	227
A.2 The SISCone algorithm	227
B Representations of the jet fragmentation	229
C The correlation coefficient as a quality measure	231
D More analytics results for jet areas	233
D.1 Transition from one-particle jet to soft jet.	233
D.2 Fluctuations of the active area	233
E Details of our NpC and cleansing study	237
F SoftKiller collinear safety issued	239
G Details of the summary simulations	241

Chapter 1

Introduction

Collider physics has played a central role in fundamental research over the past decades. Today, most earth-based experiments exploring particle physics at high energies are beam-beam collider experiments with the Large Hadron Collider (LHC) occupying the front row. We will therefore focus our attention on experiment colliding two beams of hadrons, primarily protons but also heavy nuclei, and studying the products of the collisions to extract information on fundamental interactions. With the level of our current knowledge, most of the analyses focus on high-energy scales, typically at the TeV scale or above, and on rare phenomena, typically below a picobarn i.e. more than 11 orders of magnitude smaller than the total proton-proton cross-section.

To have access to rare events, lots of statistics is crucial. This can only be achieved with high collision rates i.e. high luminosities. The LHC is designed to reach a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and hence gather several inverse femtobarns worth of data every year. This requires storing many proton bunches in each beam, storing as many protons as possible in each bunch and optimising the beam optics. A consequence of the last two requirements is that when two proton bunches cross at an interaction point, several proton-proton collisions will occur on the same time. For a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and a bunch spacing of 25 ns, one would have an average of 25 coincidental proton-proton collisions at each bunch crossing. This effect is known as **pileup** and is what we shall focus our interest on in this document.

Pileup is a major concern in collider experiments: whenever a hard proton-proton interaction occurs it will be seen, in the detectors, as the superposition of its final state and the products of the additional pileup, zero-bias, interactions that occurred during the same bunch crossing. This complicates significantly the observed final-state of the collision and, therefore, the reconstruction of the objects in that final-state. Since the zero-bias interactions making the pileup are almost entirely hadronic, pileup will compromise the reconstruction of all the final-state objects involving hadrons. This obviously impacts jet reconstruction but also many other quantities like isolation criteria for leptons and photons, which place cuts on the hadronic energy deposits in their vicinity, or missing E_T reconstruction. Pileup thus affects the majority of the objects reconstructed in the final state and compromises the precision at which one can measure the kinematic properties of these objects. In other words, pileup is one of the dominant source of systematic uncertainties in many analyses at the LHC. It is therefore of a generic importance to develop methods to correct for pileup effects as efficiently as possible. This is mandatory to make precision measurements at the LHC.

This is not really a new issue. Pileup was already present at the Tevatron. However, in that case, the average pileup multiplicity — i.e. the average number of pileup interactions per bunch crossing — was much lower, around 5, and the CDF and D0 collaborations had developed techniques, mostly based on active cells in the calorimeter, to correct for its effects. With the LHC, pileup multiplicity

has significantly increased. It has been around 20 during Run I [1], it will soon reach multiplicities around 30 for Run II and is expected to go up to 60 in the coming years. For the High-Luminosity LHC, this number might reach 140 or even 200 [2, 3]. Finding efficient methods to handle pileup is becoming increasingly important and the failure to do so might compromise key measurements like precision Higgs coupling measurements or double Higgs production.

Over the past decade, several pileup mitigation techniques have been proposed. One of the first proposed methods for used at the LHC relies on the ability to reconstruct quite efficiently pileup vertices using charged tracks and applies a correction directly proportional to the number of reconstructed pileup vertices. Ultimately, Run I of the LHC has adopted the area–median technique, initially proposed [4] by Matteo Cacciari and Gavin Salam, and significantly extended by Matteo Cacciari, Gavin Salam and myself [5, 6, 7, 8, 9]. As we will show at length in this document, this method provides an efficient and robust way to correct for pileup. However, with the perspective for yet higher pileup multiplicities in the future, several groups have started to investigate yet more powerful techniques. This has led to several new methods like cleansing [10], the SoftKiller [11] or PUPPI [12].

This document serves several purposes. The main one is to provide a review of the area–median method that is used at the LHC. We will provide an extensive discussion of the method itself, and its application to several observables, including the ubiquitous jet transverse momentum, the jet mass and jet shapes, and the jet fragmentation function. We will also perform an in-depth Monte-Carlo validation of the approach for applications at the LHC and highlight our recommendation for practical use of the area–median pileup subtraction technique. We will also show that many of the ingredients of the area–median method are amenable to an analytic understanding. This is especially true for the concept of jet areas [13] that extends beyond its usage in pileup mitigation techniques. Since pileup subtraction is most often done numerically, we will provide a description of the area–median interface in `FastJet` [14, 15].

Besides a thorough description of the area–median pileup subtraction method, we also wish to present a series of other results and techniques. These will cover two main directions. The first concerns the use of *grooming techniques* which are methods relying on jet substructure to limit the impact of pileup and the Underlying Event on fat jets, i.e. on large-radius jets often used for the reconstruction of boosted objects. This is a field that has been extremely active in the past years, with recent results showing that grooming and tagging techniques are amenable to some degree of analytic understanding. We will therefore dedicate the second part of this document to these techniques. Then, in the third and last part of this document, we will discuss methods, like the SoftKiller, that have been proposed more recently and which aim at reducing the sensitivity to pileup fluctuations and their negative impact on the jet energy resolution. These methods are mostly meant for use during Run II of the LHC and beyond, where the pileup conditions will significantly worsen. Amongst these new techniques we shall discuss the applicability of the area–median subtraction method to charged-hadron-subtracted (CHS) type of events where one can associate charged tracked to a given interaction vertex and hence remove tracks not associated directly to the primary hard interaction (or leading vertex).

Our last goal is to document a series of yet unpublished results, all discussed in the third part of this work. These mostly cover ideas and candidates for new pileup mitigation methods. the first series of studies concerns the possibility to use grooming techniques as a generic tool for pileup subtraction (not confined to their applications to fat jets). We will then provide a brief study of a series event-wide pileup mitigation methods that we have come across when developing the SoftKiller. In all these cases, the idea is to explore different methods that can lead to a reduction of the pileup fluctuation effects on jet energy resolution. Only a few methods have been proposed so far and our hope here is

that the series of preliminary results presented below can help towards building more powerful tools in the future.

It is worth mentioning already now that, throughout this document, we shall consider a simplified “theoretical” framework. Our main goal, as indicated in the title of this report, is to concentrate on theoretical aspects of pileup mitigation. This means that we shall not at all discuss additional complications related to detector effects¹ or other effects like out-of-time pileup which comes from the fact that the detectors readout often spans over the time of several bunch crossings. To be fully addressed, these effects need a dedicated experimental analysis, with details specific to each experiments, which goes beyond the scope of this document. It is however worth mentioning that, as far as the area–median pileup subtraction technique is concerned, we are confident that the picture described here is not going to be strongly affected by detector effects. This is mainly due to that fact that the area–median method is extremely robust against the details of both the hard and the pileup events.² For the same reason, we shall concentrate our discussion on the reconstruction of the properties of the jets. Other quantities like lepton/photon isolation criteria or missing- E_T reconstruction are much more dependent on the details of the detectors and will not be investigated here. Readers who are interested in experimental aspects of pileup subtraction can for example consult Ref. [16] for the ATLAS experiment or Ref. [17] for CMS.

Reading guide and organisation of this document. Even without entering into the details of detector effects, this document remains rather long. We therefore close this introduction with a reader’s guide through the following pages.

We first note that this document only discusses pileup mitigation. A brief description of the commonly-used jet algorithms is given in Appendix A and we refer to Ref. [18] for a comprehensive introduction to jet physics in general. That said, we have separated the text in three main parts: first the extensive description of the area–median method, followed by a discussion of grooming techniques in the second part and, finally, opening towards more recent techniques in the third part of this document. I strongly suggest to all readers to go through the first two Sections of Chapter 2 which provide a generic description of pileup effects, generic considerations about pileup mitigation and the net effects on physical observables. These two Sections also introduce basic pileup characteristics that will be used throughout the whole document e.g. to provide physical interpretations of our results.

The discussion of the area–median method has been organised in three different steps: we describe the method and its several extensions in Chapter 2, we provide a validation of the method based on Monte Carlo simulations in Chapters 3 and 4 and introduce a series of analytic results in Chapter 5. Again, the reader interested in the area–median method is invited to go first through Chapter 2 and in particular through Section 2.4, which gives the basic description of the method, and Section 2.8, which lists in details our practical recommendations for use of the method at the LHC. Some of the more technical aspects like the definition of jet areas (Section 2.3), extensions to jet shapes (Section 2.5) and jet fragmentation function moments (Section 2.6), or the practical details of the `FastJet` interface (Section 2.7) can be skipped if one is just interested in the main lines. Chapter 3 then demonstrates the performance of the method, justifying the recommendations made in Section 2.8, with Chapter 4 departing from pure pileup mitigation to concentrate on the application of the area–median method to the subtraction of the large Underlying Event in heavy-ion collisions. Our overview of the area–

¹Except for our description of the `SoftKiller` method in Chapter 12 where we shall briefly investigate the robustness of the method in an experimental setup.

²One specific example of a detector effect that we shall not discuss is the, usually small, non linearities that will arise from calorimeter thresholds.

median method ends with Chapter 5 where we discuss at length analytic properties of jet areas, and provide a brief translation of the analytic results of [19], initially discussing the analytic estimation of the Underlying Event activity, to the estimation of the pileup properties.

The second and third parts of this document cover more specialised topics. For the discussion of grooming techniques applied to boosted jets (the second part of this document), we have kept the same structure as for the first part: we first provide a detailed description of a series of grooming methods (Chapter 6), then move on to a series of Monte Carlo studies (Chapter 7) and finish with a description of a few analytic results for groomers (Chapter 8). The reader who is not familiar with grooming is obviously invited to read first the descriptions in Chapter 6. The Monte Carlo studies (Chapter 7) mostly revisit a study done in the context of the Boost 2012 workshop [20], studying how groomers mitigate pileup for boosted jets. However, the study presented here extends the one originally done for the Boost proceedings to include a more extensive and modern set of grooming techniques. At last, Chapter 8 will provide a calculation of the grooming radius, directly related to the sensitivity to pileup, for the SoftDrop [21] grooming procedure. This is merely an example illustrating how jet substructure methods are amenable to an analytic treatment, a field which has seen a lot of progress very recently. This Chapter is more appropriate for the reader who is curious about recent developments aiming towards an analytic understanding of jet substructure.

The third part of this document will discuss several methods going beyond the area–median approach. We encourage all readers to explore Chapter 9 which provides an introductory discussion of the motivations to seek for more powerful techniques as well as the associated fundamental challenges. For readers interested in continuing their journey in this direction, we have then split the results in four separate chapters that can be read more-or-less independently. We start in Chapter 10 with a preliminary study of the usage of grooming techniques as a generic tool for pileup mitigation, beyond the scope of boosted jets. then, in Chapter 11, we discuss the use of CHS events and charged tracks information to help mitigating pileup contamination for neutral particles. This includes a discussion of the area–median method and the “Neutral-proportional-to-Charged” method [9] as well as an in-depth comparison to the recently proposed cleansing technique [10]. In Chapter 12 we introduce the SoftKiller method that shows significant improvements over the area–median method in terms of energy resolution, while keeping limited average biases. We believe that it can be a powerful method for use in high-pileup environment, especially in view of its remarkable speed properties. Finally, in Chapter 13 we list a series of preliminary studies including several event-wide subtraction methods that we have only briefly investigated when we were developing the SoftKiller method, as well as possible extensions of the SoftKiller method itself. These are still preliminary results but I believe that they can prove helpful in the development of future pileup mitigation methods.

Finally, as a conclusion, we will provide a summary of the key performances of several methods, published and preliminary, using the common framework that had been introduced in the context of the PileUp WorkShop held at CERN in May 2014. This part is also appropriate for a generic reader who can then go back and dig more into the additional details he or she would be interested in.

Part I

The area–median approach

Chapter 2

Description of the area–median method

2.1 Basic characterisation of the pileup effects

Before discussing how to remove the pileup contamination from a jet it is helpful to characterise that contamination in terms of a few simple physical effects. This simple characterisation has its limitations but provides a simple physical understanding of many situations we shall discuss throughout this document.

The most fundamental case we will focus on is the reconstruction of the jet transverse momentum, p_t . This is by far the most used jet characteristics. Applications to more generic jet properties will be discussed in Sections 2.5 and 2.6.

The main characteristic of pileup that the area–median approach makes use of is the fact that, in each event, the energy (or, rather, the transverse momentum) deposit coming from pileup is to a large extent uniform in rapidity and azimuth. In that approximation the pileup activity is characterised by a single number: **the pileup transverse momentum density**, ρ . In this framework, ρ is defined per unit of area in the rapidity–azimuth plane. It is crucial to point out that, here, ρ is defined *event by event*. In practice, ρ has a additional relatively small rapidity dependence, smooth on average. For the sake of this discussion, let us temporarily neglect that dependence. We will come back to it in Section 2.4.2.

Given the pileup transverse momentum density ρ , the pileup effect on a jet will be to shift its p_t by an amount proportional to ρ . Since ρ is defined per unit area in the rapidity–azimuth plane, the proportionality constant will be the “area” of the jet, A_{jet} . If one thinks of a calorimeter with towers of size $a \times a$, this is equivalent to saying that the p_t of a jet will be shifted by the average pileup p_t in each tower, ρ_{tower} , times the number of towers in the jet, N_{towers} .¹ In our language, this corresponds to $\rho_{\text{tower}} = \rho a^2$ with a jet areas $A_{\text{jet}} = N_{\text{towers}} a^2$.

We will discuss at length the generic definition of the area of a jet in the next section. In a nutshell, the idea is to add to a given event a dense and uniform coverage of infinitely soft particles, *ghosts*, each carrying a quantum of area a_{ghost} . We then define the **active area** of a jet as the number of ghosts in the jet, multiplied by a_{ghost} . Physically, this active area mimics the effect of pileup in the sense that pileup particles are uniformly distributed, in the limit where their p_t becomes infinitely small. This last point ensures that ghosts do not affect the clustering provided one uses an infrared-safe jet algorithm.

¹One needs to be careful how to define the number of active towers in a jet. Even if we count the active towers that include pileup, there is always a probability that a tower inside a jet remains empty. We will come back to this at the beginning of Section 2.3, where we introduce a proper definition for jet areas.

For a given pileup density ρ and jet area A_{jet} , the average effect due to the clustering of pileup particles in a jet is to shift the jet p_t by ρA_{jet} .

A series of additional effects come on top of the simple transverse momentum shift discussed above. The fact that the pileup transverse momentum deposit is uniform in rapidity and azimuth is just an approximation only true on average. If one takes an event with a given ρ one can imagine looking at the energy deposit in a square patch of size 1×1 in the rapidity-azimuth plane. As this patch moves in the event the energy deposit will vary. This results in a full distribution which is interesting in itself² but, for most of the simple discussions in this document, it is sufficient to see it as a Gaussian of average ρ and standard deviation σ . This new pileup property, σ , **characterises the pileup fluctuations within an event** and will appear each time we discuss the uncertainty of pileup subtraction techniques.

The effect of the pileup fluctuations on the jet is to smear its transverse momentum. Indeed, if the jet sits on an upwards (resp. downwards) pileup fluctuation, the contamination due to the capture of pileup particle in the jet will be more (resp. less) than the average ρA_{jet} . These fluctuations around ρA_{jet} are given by $\sigma \sqrt{A_{\text{jet}}}$. The proportionality to σ is expected but the appearance of $\sqrt{A_{\text{jet}}}$ may be a bit more of a surprise. This is actually the fact that fluctuations, i.e. dispersion of statistical distributions, add in quadrature, which means that it is σ^2 which scales like the area of the jet. Another, perhaps more concrete, way to see this is to come back to the case of a calorimeter. If the pileup p_t in towers fluctuates by an amount σ_{tower} , the p_t fluctuations in a jet of N_{towers} towers will be $\sigma_{\text{tower}} \sqrt{N_{\text{towers}}}$. If we take a patch of area 1 in which there is $1/a^2$ towers, this gives $\sigma_{\text{tower}} = \sigma a$ and, remembering $A_{\text{jet}} = N_{\text{towers}} a^2$, we get fluctuations in our jet given by $\sigma \sqrt{A_{\text{jet}}}$.

Before turning to the discussion of how pileup affects the jet p_t , there is a third effect that one has to introduce: **back-reaction** (BR). This relates to the fact that pileup does not just add extra soft particles, from pileup interactions, to jets: since pileup particles are clustered at the same time as particles from the hard interaction, **the particles from the hard interactions clustered in a given jet can be affected by the presence of the pileup particles**. Back-reaction is the difference between the clustering done with and without pileup particles, computed only using the particles from the hard interactions within a given jet. Again, the effect of back-reaction on the jet transverse momentum is both an average shift and a dispersion. We will provide detailed calculations and simulations of back-reaction effects in Section 5.1.3, but all that we need to know at this stage is that if one uses the anti- k_t algorithm [22] to cluster the jets, the average p_t shift due to back-reaction is negligible and its smearing effects are proportional to $\sqrt{\rho p_t}$ with a relatively small coefficient which make them relevant only at large p_t .

The discussion so far can be summarised under the form of a pocket formula which will be extremely useful for many discussions throughout this report: the effect of pileup on the jet transverse momentum is given by

$$\boxed{p_{t,\text{full}} = p_{t,\text{hard}} + \rho A_{\text{jet}} \pm \sigma \sqrt{A_{\text{jet}}} + \text{BR},} \quad (2.1)$$

where $p_{t,\text{full}}$ is the transverse momentum of the jet when clustering the event with particles from both the hard interaction and pileup, while $p_{t,\text{hard}}$ is the jet p_t with only the hard particles, i.e. the quantity we really want to reconstruct. In the rhs of this equation, the second term (the ρ term) is the average shift, the third term (the σ term) is the pileup fluctuations, responsible for degrading the jet energy resolution, and the fourth is the back-reaction effects, all introduced above. We will see these effects directly in Monte-Carlo simulations in Chapter 3. Similar considerations also apply to

²We shall come back to that later in this document, e.g. to the fact that this distribution has longer tails towards large p_t than a Gaussian distribution.

other jet properties, with ρ shifting their average value and σ smearing their value (see e.g. Section 2.5 for additional details).

For the sake of future discussions, it is helpful to introduce a few additional quantities. When averaging over an ensemble of events, we will denote the average of ρ by $\langle \rho \rangle_{\text{ev}}$ and its standard deviation by σ_ρ . It is also customary to see the pileup as a superposition of N_{PU} minimum bias events. Each minimum bias event has itself a corresponding $\rho \equiv \rho_{\text{mb}}$ and $\sigma = \sigma_{\text{mb}}$. We shall denote its average by $\langle \rho_{\text{mb}} \rangle_{\text{ev}}$ and its standard deviation by $\sigma_{\rho, \text{mb}}$. For a class of events with the same N_{PU} , we would have, neglecting the fluctuations across the event sample of all the various σ ,

$$\langle \rho \rangle_{\text{ev, fixed } N_{\text{PU}}} = N_{\text{PU}} \langle \rho_{\text{mb}} \rangle_{\text{ev}}, \quad \sigma_{\text{ev, fixed } N_{\text{PU}}} = \sqrt{N_{\text{PU}}} \sigma_{\text{mb}} \quad \text{and} \quad \sigma_{\rho, \text{ fixed } N_{\text{PU}}} = \sqrt{N_{\text{PU}}} \sigma_{\rho, \text{mb}}. \quad (2.2)$$

Of course, if one averages over all events, one also needs to take into account the variations of N_{PU} across events. The dispersion of ρ will then pick an extra contribution coming from the variation of N_{PU} . This gives

$$\langle \rho \rangle_{\text{ev}} = \mu \langle \rho_{\text{mb}} \rangle_{\text{ev}}, \quad \sigma = \sqrt{\mu} \sigma_{\text{mb}} \quad \text{and} \quad \sigma_\rho^2 \simeq \mu \sigma_{\rho, \text{mb}}^2 + \sigma_{N_{\text{PU}}}^2 \langle \rho_{\text{mb}} \rangle_{\text{ev}}^2 \quad (2.3)$$

where, we have defined μ as the average number of pileup interactions and $\sigma_{N_{\text{PU}}}$ its standard deviation. For a Poisson-distributed pileup, we have $\sigma_{N_{\text{PU}}}^2 = \mu$.

From these considerations, we can reformulate Eq. (2.1) in several ways:³

$$p_{t, \text{full}} = p_{t, \text{hard}} + \langle \rho_{\text{mb}} \rangle_{\text{ev}} N_{\text{PU}} A_{\text{jet}} \pm \sigma_{\rho, \text{mb}} \sqrt{N_{\text{PU}}} A_{\text{jet}} \pm \sigma \sqrt{A_{\text{jet}}} + \text{BR} \quad (2.4)$$

$$= p_{t, \text{hard}} + \langle \rho_{\text{mb}} \rangle_{\text{ev}} \mu A_{\text{jet}} \pm \langle \rho_{\text{mb}} \rangle_{\text{ev}} \sqrt{\mu} A_{\text{jet}} \pm \sigma_{\rho, \text{mb}} \sqrt{\mu} A_{\text{jet}} \pm \sigma \sqrt{A_{\text{jet}}} + \text{BR} \quad (2.5)$$

These expressions differ in the level at which we compute the average shift and fluctuations:

- in (2.1), the average shift is characterised by ρ , defined event by event, leaving the intra-event fluctuations σ ;
- in (2.4) the average shift is proportional to the number of pileup vertices and we therefore get an extra source of fluctuations, $\sigma_{\rho, \text{mb}} \sqrt{N_{\text{PU}}} A_{\text{jet}}$, coming from the fact that the fluctuations in the overall pileup deposit of each minbias event will translate in fluctuations in ρ ;
- in (2.5), the average shift is written as a constant $p_t = \langle \rho_{\text{mb}} \rangle_{\text{ev}} \mu A_{\text{jet}}$ for each jet of area A_{jet} and we pick yet an extra source of smearing coming from the fluctuations of N_{PU} , $\langle \rho_{\text{mb}} \rangle_{\text{ev}} \sqrt{\mu} A_{\text{jet}}$.

As we shall see in this Chapter and the following ones, various pileup mitigation techniques subtract from the full jet transverse momentum and estimate of the average shift. One would then be left with residual fluctuations depending on the level at which this is done. These various expressions will prove useful when we discuss the net bias and residual resolution degradation after we apply a pileup subtraction technique.

These basic considerations are almost enough to introduce the area–median subtraction method but a few subtleties arise from our definition of the jet area. We shall therefore introduce that concept once and for all in Section 2.3 and return to the area–median subtraction method in Section 2.4

Note finally that for all the applications that we will consider in this document, we can assume some form of ordering between the various quantities introduced above. Typically, the scale of the hard jet will be larger than the background contamination, itself larger than the pileup fluctuations, namely, $p_{t, \text{hard}}/A_{\text{jet}} > \rho > \sigma$. This approximation is valid for the case of pileup at the LHC.

³We can also write a series of similar expressions where the averages are done at the level of the jets rather than at the level of transverse momentum densities. In that case, we would pick an extra contribution to the dispersion coming from the fluctuations of the area of the jet.

2.2 Effects of average shifts and resolution degradation on physical observables

We have seen that the main effects of pileup are to shift and smear the transverse momentum of the jets, as well as their other kinematic properties. The goal of pileup mitigation is, to a large extent, to correct for the average shift and to reduce as much as possible the smearing effects. At many stages in this document, we will characterise pileup subtraction techniques using these properties as quality measures.

In order to put things in perspective and to facilitate the understanding of several arguments used when we will discuss our findings, it is helpful to provide a simple description of the effects that an average shift and a dispersion would have on physical observables.

For the sake of the argument, let us consider the differential jet cross-section as a function of the jet transverse momentum, $d\sigma/dp_t$. Imagine this has a truth distribution $d\sigma_{\text{truth}}/dp_t$ corresponding to a case without pileup, usually denoted by $d\sigma_{\text{hard}}/dp_t$ throughout this document. Let us now add pileup to that initial distribution and, optionally, subtract it, to obtain a “reconstructed” jet cross-section $d\sigma_{\text{reco}}/dp_t$. This is the result of the convolution of the initial, truth, spectrum with the spectrum characterising the residual pileup effects. We want to understand how this will differ from the expected result $d\sigma_{\text{truth}}/dp_t$.

For the vast majority of the results discussed here, several simplifications can be made. First, the effect of pileup and the optional additional subtraction can, to a large extent, be approximated by a Gaussian of average Δ and of dispersion σ . The former corresponds to the average shift, or bias, or offset, while the latter corresponds to the residual smearing, or resolution degradation, or fluctuations. As long as the residual effects of pileup do not show long tails⁴ the Gaussian approximation is appropriate to capture the physics.

Then, let us approximate the initial distribution by a decreasing exponential:

$$\frac{d\sigma_{\text{truth}}}{dp_t} = \frac{\sigma_0}{\mu} e^{-p_t/\mu}. \quad (2.6)$$

This is representative of a steeply-falling distribution and, as long as Δ and σ are not too large, it is always a good approximation, at least locally.

If we convolute that distribution with a Gaussian, we find

$$\frac{d\sigma_{\text{reco}}}{dp_t} \approx \int dk_t \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(p_t - k_t - \Delta)^2}{2\sigma^2}} \frac{d\sigma_{\text{truth}}}{dk_t} \quad (2.7)$$

$$\approx \frac{d\sigma_{\text{truth}}}{dp_t} \exp\left(\frac{\Delta}{\mu} + \frac{\sigma^2}{2\mu^2}\right). \quad (2.8)$$

The above result shows that both the average shift and the dispersion have an impact on the final reconstructed distribution. The effect of the shift is rather trivial: an offset Δ on p_t translates immediately in a factor $\exp(\Delta/\mu)$. The effect of the dispersion is a bit more subtle. It is related to the fact that, on a steeply-falling distribution, positive fluctuations will be given a larger weight than negative fluctuations.

An alternative way to see this is to look, for a given reconstructed $p_{t,\text{reco}}$, at the distribution of the

⁴In particular, long positive tails.

initial, truth, $p_{t,\text{truth}}$. This is given by

$$P(p_{t,\text{truth}}|p_{t,\text{reco}}) = \left[\frac{d\sigma_{\text{reco}}}{dp_{t,\text{reco}}} \right]^{-1} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(p_{t,\text{reco}}-p_{t,\text{truth}}-\Delta)^2}{2\sigma^2}} \frac{d\sigma_{\text{truth}}}{dp_{t,\text{truth}}} \quad (2.9)$$

$$= \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2\sigma^2} \left(p_{t,\text{truth}} - p_{t,\text{reco}} + \Delta + \frac{\sigma^2}{\mu} \right)^2 \right]. \quad (2.10)$$

This shows that the most probable truth p_t is given by

$$p_{t,\text{truth}}^{(\text{most likely})} = p_{t,\text{reco}} - \Delta - \frac{\sigma^2}{\mu}. \quad (2.11)$$

The presence of the Δ shift here is again expected and we see that we get an extra contribution coming from the dispersion. This extra contribution increases if the initial distribution becomes steeper (i.e. when μ decreases) and disappears for a flat distribution ($\mu \rightarrow \infty$).

From the above perspective, one sees that if we want to mitigate pileup, the ideal situation would be to have both Δ and σ as small as possible. From Eq. (2.8), we see that having $\Delta + \sigma^2/(2\mu) = 0$ would also translate into an unbiased reconstructed distribution but this requires some level of fine-tuning and would be process-dependent. In practice, we typically target pileup mitigation techniques that result in $\Delta \approx 0$, and try to reduce σ as much as possible. In an experimental context, the effect of σ would ultimately have to be unfolded together with the other detector effects.

The discussion in this Section has only taken the jet p_t and the jet cross-section as an example. The very same arguments can be made for other processes and for other jet properties. As long as the residual effect of pileup are small enough, they can reasonably be approximated by a Gaussian distribution and the physical observable can locally be approximated by a decreasing exponential. In these conditions, the results (2.8) and (2.10) would therefore apply and capture most of the relevant effects to estimate pileup effects and to assess the quality of pileup mitigation techniques.

Finally, this discussion also justifies that the average shift and dispersion are the two basic quality measures to study and, up to very few exceptions, we will do so in all our studies in this document. The robustness of a method can be assessed by checking that the average shift Δ remains close to 0 regardless of the physical process under consideration, the scale of the hard process, the pileup conditions and other jet properties like its rapidity. Once the robustness of a method is established, we can discuss its efficiency in terms of how small the residual smearing σ is.

2.3 The jet catchment areas

The first ingredient we will need for the area–median pileup subtraction method is the definition of the concept of *jet areas*. This concept can be seen as an additional property of each jet, like their transverse momentum or mass. Note that, for hadronic collisions, we work with the transverse momentum (p_t), rapidity (y), azimuthal angle (ϕ) coordinate system⁵ and by “jet area” we really mean the area in the $y - \phi$ plane.

In the context of a (simple) calorimeter, it is natural to think of the area of a jet as the sum of the areas of each of the cells included in the jet. In a sense, our goal in this Section is to provide a

⁵One could also use the transverse energy and pseudo-rapidity. It is however more natural to keep the longitudinal-boost-invariant p_t and y since the clustering itself uses that set of coordinates. Also, rapidity differences are invariant along longitudinal boosts while pseudo-rapidity differences are not, except for massless particles.

well-defined concept of jet areas which would also work at particle-level and remain well-defined in the case of a calorimeter with empty towers in a jet. Naively, one could just say that the area of a jet of radius R should be πR^2 which would be biased because not all jet algorithms produce circular jets and the boundary of a jet can be influenced by the surrounding jets. Another tempting definition would be to take the convex hull of all the particles constituting the jet. Again, this is not an ideal choice since convex hulls could overlap and would not pave the rapidity-azimuth plane. In a nutshell, our solution is to define the area of a jet as a *catchment area*, i.e. the region of the rapidity-azimuth plane where the jet would catch infinitely soft particles. There are several ways to implement this in practice, hence we will introduce three definitions: the *passive area*, the *active area* and the *Voronoi area*.

For reasons that will soon become clear, **the active area is the most appropriate for pileup subtraction** but we will start our discussion with the passive area since it is the most simple one. Finally, jet areas are amenable to analytic calculations and we shall discuss this in great details in Section 5.1.

2.3.1 Passive Area

Suppose we have an event composed of a set of particles $\{p_i\}$ which are clustered into a set of jets $\{J_i\}$ by some infrared-safe jet algorithm.

Imagine then adding to the $\{p_i\}$ a single *infinitely soft* ghost particle g at rapidity y and azimuth ϕ , and repeating the jet-finding. As long as the jet algorithm is infrared safe, the set of jets $\{J_i\}$ is not changed by this operation: their kinematics and hard particle content will remain the same, the only possible differences being either the presence of g in one of the jets or the appearance of a new jet containing only g .

The passive area of the jet J can then either be defined as a scalar

$$a(J) \equiv \int dy d\phi f(g(y, \phi), J) \quad f(g, J) = \begin{cases} 1 & g \in J \\ 0 & g \notin J \end{cases}, \quad (2.12)$$

which corresponds to the area of the region where g is clustered with J , or as a 4-vector,

$$a_\mu(J) \equiv \int dy d\phi f_\mu(g(y, \phi), J) \quad f_\mu(g, J) = \begin{cases} g_\mu/g_t & g \in J \\ 0 & g \notin J \end{cases}, \quad (2.13)$$

where g_t is the ghost transverse momentum. For a jet with a small area $a(J) \ll 1$, the 4-vector area has the properties that its transverse component satisfies $a_t(J) = a(J)$, it is roughly massless and it points in the direction of J . For larger jets, $a(J) \sim 1$, the 4-vector area acquires a mass and may not point in the same direction as J . We shall restrict our attention here to scalar areas because of their greater simplicity. Nearly all results for scalar areas carry over to 4-vector areas, modulo corrections suppressed by powers of the jet radius (usually accompanied by a small coefficient).⁶

⁶ The above definitions apply to jet algorithms in which each gluon is assigned at most to one jet. For a more general jet algorithm (such as the ‘‘Optimal’’ jet finder of [23] or those which perform $3 \rightarrow 2$ recombination like ARCLUS [24]), then one may define the 4-vector area as

$$a_\mu(J) = \lim_{g_t \rightarrow 0} \frac{1}{g_t} \int dy d\phi (J_\mu(\{p_i\}, g(y, \phi)) - J_\mu(\{p_i\})), \quad (2.14)$$

where $J_\mu(\{p_i\}, g(y, \phi))$ is the 4-momentum of the jet as found on the full set of event particles $\{p_i\}$ plus the ghost, while $J_\mu(\{p_i\})$ is the jet-momentum as found just on the event particles.

2.3.2 From passive to Voronoi area

The only algorithm for which one can make any statement about the passive area for a general n -particle configuration is the k_t algorithm [25].

Because of the k_t distance measure, the single ghost will cluster with one of the event particles before any other clustering takes place. One can determine the region in which the ghost will cluster with a given particle, and this is a definition of the area $a_{k_t,R}(p_i)$ of a particle p_i . Since the ghost-particle clustering will occur before any particle-particle clustering, the jet area will be the sum of the areas of all its constituent particles:

$$a_{k_t,R}(J) = \sum_{p_i \in J} a_{k_t,R}(p_i). \quad (2.15)$$

Can anything be said about the area of a particle? The ghost will cluster with the event particle to which it is closest, as long it is within a distance R . This corresponds to a geometrical construction known as the Voronoi diagram, which subdivides the plane with a set of vertices into cells around each vertices.⁷ Each cell has the property that all points in the cell have as their closest vertex the cell's vertex. Thus the Voronoi cell is equivalent similar to the region in which a ghost will cluster with a particle up to the limitation that the ghost should be within a distance R of the particle. This leads us to define the Voronoi area of particle i , $a_R^\mathcal{V}(p_i)$, to be the area of its Voronoi cell \mathcal{V}_i intersected with a circle of radius R , $\mathcal{C}_{i,R}$, centred on the particle:

$$a_R^\mathcal{V}(p_i) \equiv \text{area}(\mathcal{V}_i \cap \mathcal{C}_{i,R}). \quad (2.16)$$

Thus given a set of momenta, the passive area of a k_t jet can be directly determined from the Voronoi diagram of the event,⁸ using Eq. (2.15) and the relation

$$a_{k_t,R}(p_i) = a_R^\mathcal{V}(p_i). \quad (2.17)$$

The Voronoi construction of the k_t -algorithm passive area is illustrated in fig. 2.1. One sees both the Voronoi cells and how their intersection with circles of radius $R = 1$ gives the area of the particles making up those jets.

Note that it is not possible to write passive areas for jet algorithms other than k_t in the form Eq. (2.15). One can however introduce a new type of area for a generic algorithm, a *Voronoi area*, in the form

$$a_{\text{JA},R}^\mathcal{V}(J) = \sum_{p_i \in J} a_R^\mathcal{V}(p_i). \quad (2.18)$$

While for algorithms other than k_t (for which, as we have seen, $a_{k_t,R}(J) = a_{k_t,R}^\mathcal{V}(J)$), this area is not in general related to the clustering of any specific kind of background radiation, it can nevertheless be a useful tool, because its numerical evaluation is efficient [26, 14] and as we shall discuss later when we discuss analytic properties of jet areas (section 5.1.2.4), for dense events its value coincides with both passive and active area definitions.

⁷It is this same geometrical construction that was used to obtain a nearest neighbour graph that allowed k_t jet clustering to be carried out in $N \ln N$ time [14].

⁸Strictly speaking it should be the Voronoi diagram on a $y - \phi$ cylinder, however this is just a technical detail.

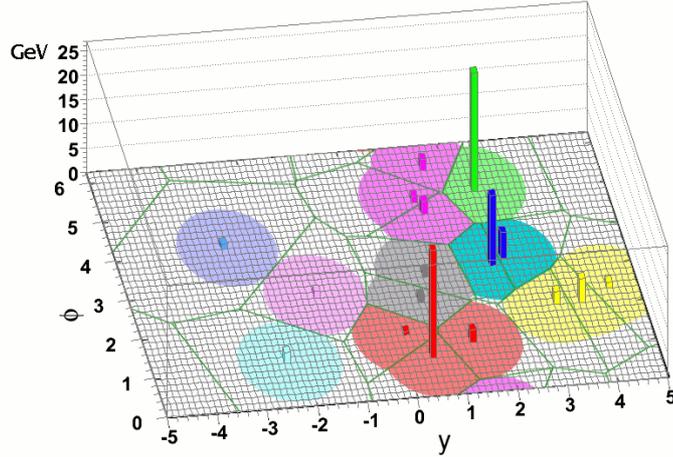


Figure 2.1: The passive area of jets in a parton-level event generated by Herwig and clustered with the k_t algorithm with $R = 1$. The towers represent calorimeter cells containing the particles, the straight (green) lines are the edges of the Voronoi cells and the shaded regions are the areas of the jets.

2.3.3 Active Area, catchment area for pileup subtraction

To define an active area, as for the passive area, we start with an event composed of a set of particles $\{p_i\}$ which are clustered into a set of jets $\{J_i\}$ by some infrared-safe jet algorithm. However, instead of adding a single soft ghost particle, we now add a dense coverage of ghost particles, $\{g_i\}$, randomly distributed in rapidity and azimuth, and each with an infinitesimal transverse momentum.⁹ The clustering is then repeated including the set of particles plus ghosts.

During the clustering the ghosts may cluster both with each other and with the hard particles. This more ‘active’ participation in the clustering is the origin of the name that we give to the area defined below. It contrasts with the passive area definition (section 2.3.1) in which the single ghost acted more as a passive spectator.

Because of the infrared safety of any proper jet algorithm, even the addition of many ghosts does not change the momenta of the final jets $\{J_i\}$. However these jets do contain extra particles, ghosts, and we use the number of ghosts in a jet as a measure of its area. Specifically, if the number of ghosts per unit area (on the rapidity-azimuth cylinder) is ν_g and $\mathcal{N}_g(J)$ is the number of ghosts contained in jet J , then the (scalar) active area of a jet, given the specific set of ghosts $\{g_i\}$ is

$$\boxed{A(J | \{g_i\}) = \frac{\mathcal{N}_g(J)}{\nu_g}}. \quad (2.19)$$

An example of jet areas obtained in this way is shown in figure 2.2. One notes that the boundaries of the jets obtained with the k_t algorithm (the left plot) are rather ragged. Clustering with a different set of ghosts would lead to different boundaries. This is because the ghosts can cluster among themselves to form macroscopic subjects, whose outlines inevitably depend on the specific set of initial ghosts, and these then subsequently cluster with true event particles. This can happen for any density of ghosts, and thus the jet boundaries tend to be sensitive to the randomness of the initial sets of ghosts. This is true for most jet algorithms. One noticeable exception is the case of the anti- k_t algorithm (right plot in Fig. 2.2). In that case, ghosts will cluster with the particles in the event before they cluster

⁹In most cases the distribution of those transverse momenta will be irrelevant, at least in the limit in which the density of ghosts is sufficiently high.

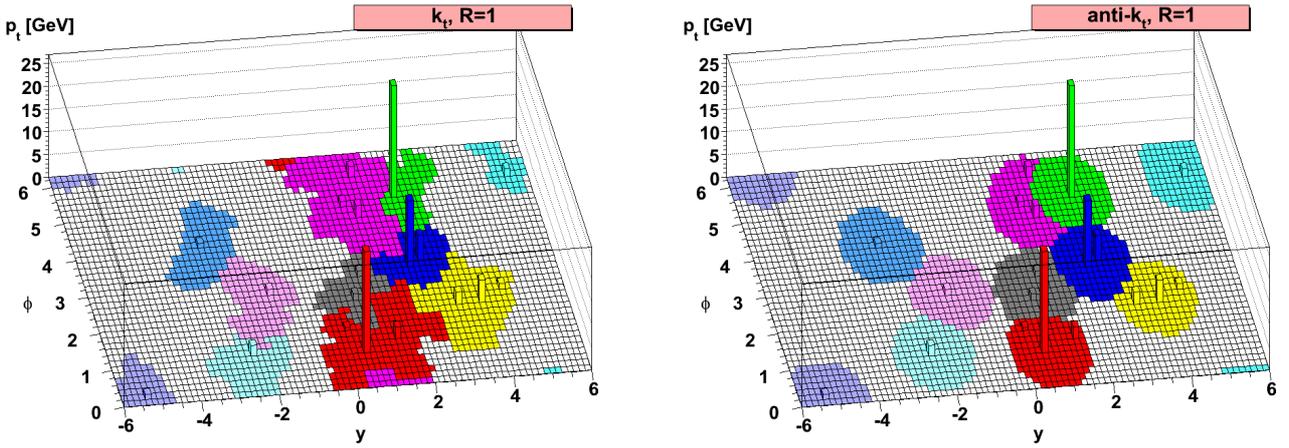


Figure 2.2: Active area for the same event as in figure 2.1, once again clustered with the k_t algorithm (left) and the $\text{anti-}k_t$ algorithm (right) both using $R = 1$. Only the areas of the hard jets have been shaded — the pure ‘ghost’ jets are not shown.

among themselves. The boundaries of the jets are thus defined by the structure of the event, (almost) independently of the initial set of ghosts.¹⁰ In particular, the hard jets will be circular, as seen in Fig. 2.2.

This randomness propagates through to the number of ghosts clustered within a given jet, even in the limit $\nu_g \rightarrow \infty$, resulting in a different area each time. To obtain a unique answer for active area of a given jet one must therefore average over many sets of ghosts, in addition to taking the limit of infinite ghost density,¹¹

$$A(J) = \lim_{\nu_g \rightarrow \infty} \langle A(J | \{g_i\}) \rangle_g. \quad (2.20)$$

Note that as one takes $\nu_g \rightarrow \infty$, the ghost transverse momentum density, $\nu_g \langle g_t \rangle$, is to be kept infinitesimal.

The active area bears a close resemblance to the average susceptibility of the jet to a high density of soft radiation like pileup, since the many soft particles will cluster between each other and into jets much in the same way as will the ghosts. In that sense, the ‘‘area of the jet’’, A_{jet} , referred to in our introductory description of pileup effects (see Section 2.1), has to be seen as an active area.

A feature that arises when adding many ghosts to an event is that some of the final jets contain nothing but ghost particles. They did not appear in the original list of $\{J_i\}$ and we refer to them as *pure ghost jets*. These pure ghost jets (not shown in fig. 2.2), fill all of the ‘empty’ area, at least in jet algorithms for which all particles are clustered into jets. They will be similar to the jets formed from purely soft radiation in events with minimum-bias pileup, and so are interesting to study in their own right and will play a crucial role in the area–median subtraction method.

In analogy with the 4-vector passive area, we can define a 4-vector version of the active area. It is

¹⁰Because of the finite ghost density and the randomness of the ghosts, the boundaries will slightly fluctuate between different sets of ghosts. In a limit of an infinite ghost density, the boundary would become independent of the set of ghost.

¹¹One may wonder if the averaged area (and its dispersion) depends on the specific nature of the fluctuations in ghost positions and momenta across ensembles of ghosts — for a range of choices of these fluctuations, no significant difference has been observed (except in the case of pure ghost jets with SISCone, whose split–merge step introduces a strong dependence on the microscopic event structure).

given by

$$A_\mu(J|\{g_i\}) = \frac{1}{\nu_g \langle g_t \rangle} \sum_{g_i \in J} g_{\mu i}, \quad A_\mu(J) = \lim_{\nu_g \rightarrow \infty} \langle A_\mu(J|\{g_i\}) \rangle_g. \quad (2.21)$$

The sum of the $g_{\mu i}$ is to be understood as carried out in the same recombination scheme as used in the jet clustering.¹²

One may also define the standard deviation $\Sigma(J)$ of the distribution of a jet’s active area across many ghost ensembles,

$$\Sigma^2(J) = \lim_{\nu_g \rightarrow \infty} \langle A^2(J|\{g_i\}) \rangle_g - A^2(J). \quad (2.22)$$

This provides a measure of the variability of a given jet’s contamination from (say) pileup and is closely connected with the momentum resolution that can be obtained with a given jet algorithm.

2.4 Area–median pileup subtraction

2.4.1 Basic recipe for pileup subtraction

Now that we have a well-defined concept of the area of a jet in our toolkit, we can finally introduce the area–median approach.

Coming back to our original discussion of the pileup effects, we have noticed that pileup introduces an average bias appearing as a shift of the jet transverse momentum, as well as a smearing of its resolution. We aim at a method that would remove the p_t shift on average, i.e. be unbiased on average, and limit the residual resolution degradation effects to a minimum. Looking at equations (2.1), (2.4) and (2.5), it is preferable to apply a correction that avoids making too many averages which would each time result in additional contributions to the resolution degradation.

The main idea of the area–median pileup subtraction technique is to subtract from the jet transverse momentum $p_{t,\text{full}}$ its average shift ρA_{jet} in (2.1). This requires two ingredients: the calculation of the jet area, which we already have under the form of the *active area* introduced in the previous Section, and a method to obtain an estimate of ρ , ρ_{est} .

The estimation of the average pileup activity in an event is based on the observation that, if we break an event into patches, most of the patches will not contain particles from the hard jets and for those patches $p_{t,\text{patch}}/A_{\text{patch}}$ is, on average, a good estimate of ρ . Averaging the $p_{t,\text{patch}}/A_{\text{patch}}$ over all the patches would however not give a good estimate of ρ since it will be strongly biased by the few patches containing particles from the hard jets. In practice, it was found [4] that the median of $p_{t,\text{patch}}/A_{\text{patch}}$ provides a good estimate of ρ , not significantly affected by the “outlying” patches containing particles from the hard jets.¹³

In the end, the area–median subtraction method works as follows:

¹²Though we do not give the details it is simple to extend the 4-vector active area definition to hold also for a general IR safe jet algorithm, in analogy with the extension of the passive area definition in Eq. (2.14).

¹³Several alternative approaches can be envisaged, like excluding (recursively or not) patches more than $n\sigma$ away from the global average, similarly to what ATLAS has used for background subtraction in heavy-ion collisions [27]. We found that the use of a median was equally efficient and simpler than alternative methods which require additional parameter(s).

- Find an estimate of ρ :

1. Break the event into patches of similar area. Typically, one can break the event into grid cells in the $y-\phi$ plane (“*grid–median*” estimator). Alternatively, one can use k_t [25] or Cambridge/Aachen [28] jets¹⁴ obtained with a jet radius potentially different from the one used to cluster the physical jets (“*jet–median*” estimator).
2. Compute the momentum and area of each patch. When patches are grid cells, one could use the scalar p_t and scalar, geometric, area, while for jet patches, one would rather use their 4-vector p_t together with the transverse component of the 4-vector active area.
3. The estimate of ρ is given by

$$\rho_{\text{est}} = \text{median}_{i \in \text{patches}} \left\{ \frac{p_{t,i}}{A_i} \right\} \quad (2.23)$$

where the empty cells or the pure-ghost jets are included in the calculation.

- Subtract pileup from physical jets:

1. cluster the jets with your favourite jet definition, computing their active area
2. correct their 4-momentum using

$$p_{\mu}^{(\text{sub})} = p_{\mu} - \rho_{\text{est}} A_{\text{jet},\mu}. \quad (2.24)$$

Modulo possible offsets in the estimation of ρ (see Section 5.2 for details), the area–median method would subtract from each jet its ρA_{jet} contamination, leaving a jet with, on average, an unbiased transverse momentum¹⁵, with a p_t resolution degradation given by $\sigma \sqrt{A_{\text{jet}}}$. There could also be an additional bias and resolution degradation coming from back-reaction although, apart from the resolution degradation at large p_t , this can safely be ignored when subtracting pileup from anti- k_t jets which covers the vast majority of the applications.¹⁶ Overall, we therefore expect the area–median approach to provide jets with better resolution than alternative unbiased approaches like subtracting an average pileup contamination per pileup vertex or an overall average pileup contamination, since these would result in additional sources of resolution degradation. The key element behind this is the fact that the pileup activity, ρ , is estimated fully dynamically on an event-by-event basis with the jet area computed individually for each jet. The left-over resolution degradation due to pileup fluctuations are of the order of the much-smaller σ .

When implementing this procedure in a practical context, several technical details have to be fixed like the rapidity coverage for particles and ghosts or the size of the patches. We will come back to these in Section 2.8. Here we rather discuss the physics behind the area–median approach and carry

¹⁴These algorithms have the advantage that the resulting jets tend to have reasonably uniform areas. Conversely, anti- k_t and SISCone suffer from jets with near zero areas or, for SISCone, sometimes huge, “monster” jets, biasing the ρ determination. They are therefore not recommended.

¹⁵To make things clear, by “on average” we mean that if we take a jet with a given p_t (or a family of jets with the same p_t), and repeatedly add pileup to it (or them), the average transverse momentum of the jets obtained after applying the area–median subtraction procedure would be the original p_t .

¹⁶At large p_t the resolution degradation due to back-reaction would grow like $\sqrt{\alpha_s \rho p_t A_{\text{jet}}}$ (see Section 5.1.3) mostly coming from a long negative tail. This would likely have very little effect in practical applications, especially when detector effects are also taken into account.

on with extensions in the next Sections.

The area–median method would work basically with any jet algorithm giving well-defined active jet areas. Apart from some fancy grooming techniques, this mostly means that it works with any infrared-safe jet algorithm.

A potentially non-trivial point is why empty patches (or pure-ghost jets) have to be included in the calculation of the median. In an event with a large pileup activity, where the whole event is densely populated, this would make no difference. However, when the event becomes sparse, e.g. when applying our procedure to parton-level Monte-Carlo simulations or with events with no or little pileup activity, most of the non-empty patches would be made of hard jets. Including empty patches would then avoid having ρ_{est} set to the scale of the hard jets.¹⁷

Note that, conceptually, one can view the area–median pileup subtraction method as part of the jet definition. In that viewpoint, jets are defined by applying the anti- k_t algorithm with a given radius (or any other jet definition one wishes to use) and subtracting the result using the area–median method. It is the combination of these two steps which can be seen as a jet definition.

Furthermore, the Underlying Event produced in proton-proton collisions also corresponds to a largely uniform energy deposit. It will therefore also be subtracted by the area–median method. We will discuss this in more details in Section 3.2.6.

2.4.2 Positional dependence

So far, we have assumed that, apart from small up and down fluctuations of order σ , the pileup energy deposit is uniform in rapidity and azimuthal angle. We know however that this is not true in practice. The most striking example is that the energy deposit in a minimum bias collision is, on average, larger in central rapidities than in the forward regions, translating into an average pileup transverse momentum ρ larger in the central region than at forward rapidities. Another example is the case of heavy-ion collisions where flow [30, 31] would also generate a non-trivial ϕ dependence, mostly dominated by elliptic flow and its $\cos(2\phi)$ dependence. In both cases, the background density ρ will vary smoothly across an event.

From the perspective of pileup subtraction, this means that we have to deal with a situation where ρ depends on the position of the jet j that one wants to subtract. The area–median approach proposes two methods to handle the positional dependence of ρ , both based on adaptations of Eq. (2.23).

Local range. The first method is to estimate ρ only from the patches which are in the local vicinity of the jet one wishes to subtract:

$$\rho_{\text{est},\mathcal{R}(j)} = \text{median}_{i \in \text{patches} \in \mathcal{R}(j)} \left\{ \frac{p_{t,i}}{A_i} \right\}, \quad (2.25)$$

where $\mathcal{R}(j)$ denotes a *local range* around jet j . Typical choices for $\mathcal{R}(j)$ are the following:

- the *Strip* range, $\mathcal{S}_\Delta(j)$, includes the jets j' satisfying $|y_{j'} - y_j| < \Delta$,
- the *Circular* range, $\mathcal{C}_\Delta(j)$, includes the jets j' satisfying $\sqrt{(y_{j'} - y_j)^2 + (\phi_{j'} - \phi_j)^2} < \Delta$,

¹⁷Experimentally, this can come with additional complications. For example, since the pileup energy deposit is smaller in forward regions, it can be cut by the calorimeter noise threshold. This can result in ρ being underestimated or even estimated as 0 more often than it should [29, 16]. One way around that, used by ATLAS [16], is to estimate the pileup activity in the central region and to extrapolate it to forward regions using a profile depending on N_{PU} and μ (the two are needed because of out-of-time pileup).

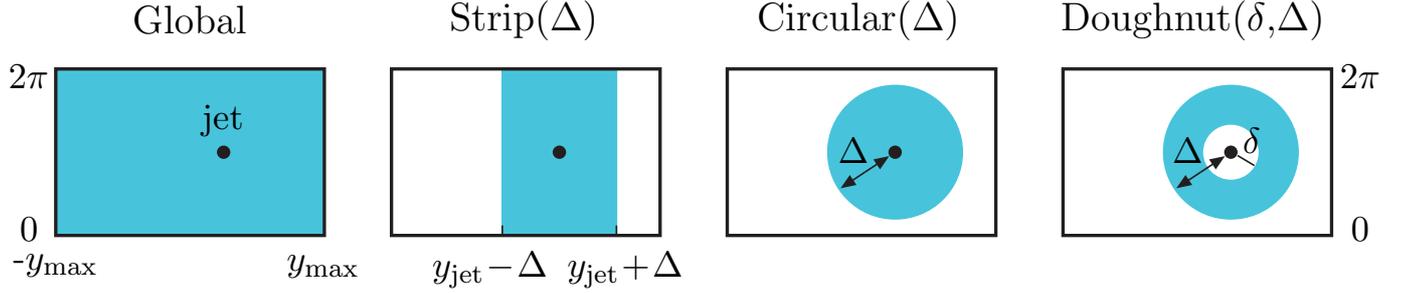


Figure 2.3: Typical local ranges that can be used for the estimation of the positional-dependent pileup density $\rho(j)$, compared to a global range (left).

- the *Doughnut* range, $\mathcal{D}_{\delta,\Delta}(j)$, includes the jets j' satisfying $\delta < \sqrt{(y_{j'} - y_j)^2 + (\phi_{j'} - \phi_j)^2} < \Delta$.

These three options are depicted in Figure 2.3.

Note that, when using a local range, a compromise has to be found between taking it large enough so that it contains enough patches for the estimation of the median to be reliable, and not too large to capture the positional dependence of ρ . Based on analytic estimates discussed in Section 5.2, we can show that for jets of radius R , if we want the uncertainty on the determination of ρ to contribute at most for a fraction ϵ of the natural pileup fluctuations $\sigma\sqrt{A_{\text{jet}}}$, the range needs to cover at least an area $A_{\mathcal{R}}$ such that

$$A_{\mathcal{R}} \gtrsim \frac{\pi^2 R^2}{4\epsilon}. \quad (2.26)$$

For $\epsilon = 0.1$ and $R = 0.4$, this gives $A_{\mathcal{R}} \gtrsim 4$. We will discuss this in more details in our Monte Carlo studies of heavy-ion collisions in Chapter 4.

Rescaling. The alternative approach is to correct for the positional dependence before taking the median. In practice, if $f(j)$ captures the positional dependence of the pileup energy deposit as a function of the position j (arbitrarily normalised), we can use this as a rescaling function and compute the background density at the position j using

$$\rho_{\text{est},f}(j) = f(j) \operatorname{median}_{i \in \text{patches}} \left\{ \frac{p_{t,i}}{A_i f(i)} \right\}, \quad (2.27)$$

where $f(i)$ is the positional dependence computed at the position of “patch i ”.

Compared to the local-range approach, this requires the rescaling function f to be computed before using Eq. (2.27). In the situation where this is used to correct for the smooth rapidity dependence of pileup, this function can be taken from an average over minimum bias events. For elliptic flow in heavy-ion collisions, $f(j)$ can be parametrised as a function of the event plane and v_2 determined event-by-event.

We shall discuss in more details these local estimates later, when we validate the area-median subtraction using Monte-Carlo simulations (Section 3.2). We shall also discuss the local range approach when we discuss heavy-ion collisions (Chapter 4).

2.4.3 Particle masses

In order to obtain a full pileup subtraction of the jet 4-momentum, there is a last ingredient that needs to be discussed: the effect of particle masses.

To see this, let us go back to the basic area–median method described in Section 2.4.1 and to equation (2.24) in particular. The 4-vector jet area is nothing but a 4-vector sum over the ghosts inside the jets (see e.g. Eq. (2.21)). The main formula (2.24) can then be seen as subtracting from the jet 4-vector, j^μ , an amount $a\rho g^\mu$ for each ghost in the jet. In this expression, $a = 1/\nu$ is the fundamental area carried by each ghost and ρ is the appropriate conversion factor to correct for the jet transverse momentum. In other words, (2.24) makes a parallel between the pileup particles contaminating the jet and the ghosts used to calculate the jet area, using the latter to subtract the former through an estimate of the proportionality constant ρ .

This analogy between the pileup particles and the ghosts is however no longer exact when pileup particles become massive since the ghosts are massless. On top of its rapidity and azimuth, a generic 4-vector has two additional degrees of freedom: its transverse momentum p_t and its *transverse mass* $m_t = \sqrt{p_t^2 + m^2}$, with m its mass. Since ρ is determined from the transverse-momentum component of the patches, subtracted a 4-vector proportional to a jet area made of massless ghosts will result in a small mismatch for the components of the jet sensitive to the mass of the particles.

To put this discussion on a more formal ground, let us take an infinitesimal patch of the event of size $\delta y \times \delta\phi$. The 4-vector corresponding to a uniform pileup deposit in that patch will take the generic form

$$(p_x, p_y, p_z, E) \equiv (\rho \cos(\phi), \rho \sin(\phi), (\rho + \rho_m) \sinh(y), (\rho + \rho_m) \cosh(y)) \delta y \delta\phi \quad (2.28)$$

where we have used the fact that ρ is defined as the pileup transverse momentum per unit area — estimated using Eq. (2.23) — and the longitudinal components, p_z and E , can receive an extra contribution coming from the mass of the pileup particle. The new term ρ_m is an extra pileup contamination that comes from the fact that the transverse mass is larger than the transverse momentum for massive particles. A uniform pileup deposit is therefore characterised by 2 numbers: ρ and ρ_m .

The correction for the ρ contamination can be carried over using the method already presented in Section 2.4.1 and correcting from the new ρ_m term can be straightforwardly done following the same logic. We first estimate ρ_m using a median approach:

$$\rho_{m,\text{est}} = \text{median}_{i \in \text{patches}} \left\{ \frac{m_{\delta,i}}{A_i} \right\}, \quad (2.29)$$

with $m_{\delta,\text{patch}} = \sum_{i \in \text{patch}} \sqrt{p_{t,i}^2 + m_i^2} - p_{t,i}$, where the sum runs over all the particles in the patch. We then correct the jet 4-vector using

$$p_\mu^{(\text{sub})} = p_\mu - (\rho_{\text{est}} A_{\text{jet},x}, \rho_{\text{est}} A_{\text{jet},y}, (\rho_{\text{est}} + \rho_{m,\text{est}}) A_{\text{jet},z}, (\rho_{\text{est}} + \rho_{m,\text{est}}) A_{\text{jet},E}). \quad (2.30)$$

These expressions can easily be extended to include positional dependence either using local ranges or using a rescaling function. In the second case, note that for all the applications we have seen, a common rescaling function for ρ and ρ_m is good enough (see also Section 3.1).

2.4.4 Safe subtraction and positivity constraints

There are a few cases where the area–median subtraction formula could lead to unphysical jets.

The first one is when ρA is larger than the (unsubtracted) p_t of the jet. This could typically happen for low- p_t jets sitting on a downwards pileup fluctuation and would lead to jets with “negative” p_t

after subtraction.¹⁸ In this case, the most reasonable prescription is usually to set the subtracted jet to a vanishing 4-momentum.

The second case would be a jet with a positive transverse momentum but a negative squared mass. In this case, a good prescription is to keep the transverse momentum and azimuth of the jet obtained after subtraction of the ρ contamination, take the rapidity of the original jet and set its mass to 0.

We will explicitly refer to these prescriptions as “safe” subtraction and discuss them later in our Monte Carlo tests (see Chapter 11).

2.4.5 Extension to Charged-Hadron Subtracted (CHS) events

In several experimental contexts, it is possible to associate charged tracks either to the leading vertex or to a pileup interaction. In that context, one builds a *Charged-Hadron-Subtracted (CHS) event* where one subtracts the tracks associated to pileup interactions keeping only the tracks associated with the leading vertex and the neutral particles (or energy deposits) coming either from the leading vertex or from pileup. Optionally, the charged tracks associated with a pileup vertex can be retained as ghosts, i.e. with their momentum scaled down by a large factor, in order to retain information about their presence in the original event. Modulo experimental complications¹⁹, this considerably reduces the pileup contamination (both the average shift and the dispersion), since charged tracks account for about a fraction $f \approx 0.62$ of the pileup energy deposit.

The area–median method can be straightforwardly applied to CHS events, with the pileup properties estimated either on the full CHS event or only on their neutral component.²⁰ One should expect a reduction of the average pileup contamination by a factor $1/(1-f) \approx 2.6$ and of the resolution smearing by a factor $1/\sqrt{1-f} \approx 1.6$. Although results from Chapter 3 directly extend to CHS events, we will discuss this at length in Chapter 11 where we compare the area–median approach to alternative methods that make use of charged track information.

2.5 Jet shapes

So far, we have introduced the area–median method for subtracting pileup contamination from the jet 4-momentum. It is now natural to ask if the method can be extended to more general properties of the jets. In this Section, we show that it is possible to obtain a generic area–median subtraction method for jet shapes, i.e. infrared-and-collinear safe functions of the jet constituents.

Let us therefore consider a generic infrared-and-collinear-safe function $v_{\text{jet}}(\{p_i\})$ of the jet constituents $\{p_i\}$, including a uniform pileup contamination characterised by the scales ρ and ρ_m . We also include in these constituents a set of ghosts each carrying an area $A_g = 1/\nu$ in the $y-\phi$ plane and to which we associate an infinitesimal transverse momentum $p_{t,g}$ and a (new) infinitesimal transverse mass $m_{t,g} = p_{t,g} + m_{\delta,g}$ with $m_{\delta,g} = \sqrt{p_{t,g}^2 + m_g^2} - p_{t,g}$. As discussed in Section 2.4.3, these ghosts mimic the effect of a uniform pileup deposit and the ghost scales $p_{t,g}$ and $m_{\delta,g}$ correspond to the pileup densities ρ and ρ_m .

¹⁸Strictly speaking, a 4-vector subtraction would result in a jet pointing in the opposite direction in ϕ which is not any better than a negative p_t .

¹⁹In an experimental context, this comes with the difficulty of having to subtract the energy deposit left by charged tracks in the calorimeter. Additional complications also come from the smaller rapidity coverage of the tracker.

²⁰The only small difference being that for “safe” subtraction, one would use the rapidity and mass of the charged tracks associated to the leading vertex in the cases where the initial area–median subtraction leads to squared mass smaller than the squared mass of the charged tracks coming from the leading vertex.

For any given jet we can then consider the derivatives of the jet shape w.r.t. to the ghost transverse scales $p_{t,g}$ and $m_{\delta,g}$,

$$v_{\text{jet}}^{(m,n)} = A_g^{m+n} \partial_{p_{t,g}}^m \partial_{m_{\delta,g}}^n v_{\text{jet}}(\{p_i\}). \quad (2.31)$$

In a sense, these derivatives probe the susceptibilities of the shape to the ρ and ρ_m components of pileup and can be used to correct the shape v by extrapolating its value down to zero pileup:

$$v^{(\text{sub})} = \sum_{m=0}^{\infty} \frac{1}{m!} \sum_{n=0}^{\infty} \frac{1}{n!} (-)^{m+n} \rho^m \rho_m^n v_{\text{jet}}^{(m,n)} \quad (2.32)$$

$$\approx v_{\text{jet}} - \rho v_{\text{jet}}^{(1,0)} - \rho_m v_{\text{jet}}^{(0,1)} + \frac{1}{2} \rho^2 v_{\text{jet}}^{(2,0)} + \rho \rho_m v_{\text{jet}}^{(1,1)} + \frac{1}{2} \rho_m^2 v_{\text{jet}}^{(0,2)} + \dots, \quad (2.33)$$

where the second line is expanded to second order in the pileup properties. This second-order subtraction is the main formula we shall use to subtract jet shapes.²¹

It is illustrative to consider two specific examples of this formula: the jet transverse momentum and the jet mass squared. For the jet transverse momentum, the uniform pileup distribution will shift the original hard transverse momentum $p_{t,\text{hard}}$ (without pileup), by an amount $A_{\perp} \rho$, with A_{\perp} the transverse momentum component of the jet area 4-vector. On top of that, the ghosts will add a contribution $p_{t,g} A_{\perp} / A_g$ to the total p_t , giving successively

$$p_t = p_{t,\text{hard}} + (\rho + p_{t,g}/A_g) A_{\perp}, \quad (2.34)$$

$$p_t^{(1,0)} = A_{\perp}, \quad (2.35)$$

with all other derivatives vanishing, and finally,

$$p_t^{(\text{sub})} = p_t - \rho p_t^{(1,0)} = p_t - \rho A_{\perp} = p_{t,\text{hard}} + p_{t,g} A_{\perp} / A_g = p_{t,\text{hard}} \quad (2.36)$$

where, for the last equality, we have dropped the negligible contribution coming from the infinitesimal p_t of the ghosts. This simple exercise shows that (2.33) reduces to the expression $p_t^{(\text{sub})} = p_t - \rho A_{\perp}$ at the core of the area–median subtraction. It also provides an interesting interpretation of the “jet area” as the susceptibility of the jet transverse momentum to variations of the background density ρ .

Albeit a bit more involved technically, the case of the squared jet mass flows similarly. Let us consider a hard jet of transverse momentum $p_{t,\text{hard}}$ and mass m_{hard} with additional uniform pileup contamination and ghosts. We already know that the hard jet transverse momentum is shifted by $(\rho + p_{t,g}/A_g) A_{\perp}$ and, in a similar way, the jet transverse mass will be shifted by $(\rho + p_{t,g}/A_g) A_{\perp} + (\rho_m + m_{\delta,g}/A_g) A_{\text{jet}}$, where we have separated the contribution from the transverse momentum and the one from the particle masses. This gives

$$\begin{aligned} m^2 &= \left[\sqrt{p_{t,\text{hard}}^2 + m_{\text{hard}}^2} + (\rho + p_{t,g}/A_g) A_{\perp} + (\rho_m + m_{\delta,g}/A_g) A_{\text{jet}} \right]^2 - [p_{t,\text{hard}} + (\rho + p_{t,g}/A_g) A_{\perp}]^2, \\ m^{2(1,0)} &= 2A_{\perp} \left[\sqrt{p_{t,\text{hard}}^2 + m_{\text{hard}}^2} - p_{t,\text{hard}} + (\rho_m + m_{\delta,g}/A_g) A_{\text{jet}} \right], \\ m^{2(0,1)} &= 2A_{\text{jet}} \left[\sqrt{p_{t,\text{hard}}^2 + m_{\text{hard}}^2} + (\rho + p_{t,g}/A_g) A_{\perp} + (\rho_m + m_{\delta,g}/A_g) A_{\text{jet}} \right], \\ m^{2(2,0)} &= 0, \\ m^{2(1,1)} &= 2A_{\perp} A_{\text{jet}}, \\ m^{2(0,2)} &= 2A_{\text{jet}}^2, \end{aligned} \quad (2.37)$$

²¹Theoretically, there is no limitation to go to higher orders but, in practice, the derivatives will be estimated numerically and third-order derivatives become delicate to compute for common applications.

with higher derivatives vanishing. Substituting this in (2.33), and neglecting the infinitesimal ghosts, we find after easy manipulations

$$m^{2(\text{sub})} = m_{\text{hard}}^2. \quad (2.38)$$

Note that, in order to recover exactly the hard jet without pileup contamination, it is important to apply the subtraction procedure to the jet mass squared. This is mostly because, as can be seen from (2.37), the squared mass m^2 is explicitly quadratic in $p_{t,g}$ and $m_{\delta,g}$. If we were to apply the subtraction formula to the jet mass, the extra square root would translate in non-vanishing higher derivatives and subtracting pileup contamination up to second-order would leave some (parametrically small) residual corrections (proportional to ρ^3/m^2 at large p_t).

These examples show not too surprisingly that our second-order expansion would work better if the quadratic approximation works fine. As we shall see in practice in Monte Carlo simulations, there are some notorious cases where this approximation tends to break. One typical example is when a shape has an endpoint and the effect of pileup is to bring the value of the shape close to that endpoint.

In some situations, it is however possible to work around these limitations by splitting the shape into components and subtracting these components individually before reconstructing the final subtracted shape from the subtracted components. The case of the N -subjettiness ratio $\tau_{21} = \tau_2/\tau_1$ illustrates this very well: while τ_{21} itself has a complicated dependence on the pileup density, (unnormalised) τ_1 and τ_2 are linear in ρ . We should therefore apply the subtraction procedure to τ_1 and τ_2 and reconstruct the subtracted value of τ_{21} from those. The only delicate point here is to properly handle the situation where the subtracted value of τ_1 becomes 0 or negative, e.g. by setting the subtracted value of τ_{21} to 0.

Finally, let us note that it is possible to rewrite (2.33) in terms of a single variable. We can for example introduce a variable x_g such that $p_{t,g} = x_g$ and $m_{\delta,g} = \rho_m/\rho x_g$ and define

$$v_{\text{jet}}^{[k]} = A_g^k \partial_{x_g}^k v_{\text{jet}}(\{p_i\}). \quad (2.39)$$

The subtraction formula then takes the simpler form

$$v^{(\text{sub})} = \sum_{k=0}^{\infty} \frac{1}{k!} (-)^k \rho^k v_{\text{jet}}^{[k]} \approx v_{\text{jet}} - \rho v_{\text{jet}}^{[1]} + \frac{1}{2} \rho^2 v_{\text{jet}}^{[2]} + \dots \quad (2.40)$$

2.6 Jet fragmentation function

At this stage, we have seen that the area–median pileup subtraction method can be applied to all infrared and collinear safe jet observables. In this section, we shall add to that list yet another observable, the jet fragmentation function, which is collinear unsafe.

Part of the idea behind the area–median subtraction of the jet fragmentation function is to subtract its moments rather than the direct distribution. We provide a discussion of the different representations of a jet fragmentation function in Appendix B, including some additional physics motivations to consider the moments of the jet fragmentation function. Here we shall just recall the basic definitions.

Definition. The jet fragmentation function (FF) is defined as the distribution dN_h/dz of the momentum fraction

$$z = \frac{p_{t,h}}{\tilde{p}_t} \quad (2.41)$$

of all the hadrons h in a jet. In the definition of z , $p_{t,h}$ is the transverse momentum of the hadron and \tilde{p}_t is the scalar p_t of the jet, defined as the scalar sum of the transverse momenta of the constituents of the jet, $\tilde{p}_t = \sum_{h \in \text{jet}} p_{t,h}$. In future discussions it will also be interesting to define $\xi = \log(1/z)$. In practice, the distribution dN_h/dz is usually binned in z or ξ and normalised to the number N_{jet} of jets in the sample.

Instead of applying the pileup subtraction directly to dN_h/dz , we will rather do it to its moments defined as

$$M_N = \frac{1}{N_{\text{jet}}} \int_0^1 dz z^N \frac{dN_h}{dz} = \frac{1}{N_{\text{jet}}} \int_0^\infty d\xi e^{-N\xi} \frac{dN_h}{d\xi}. \quad (2.42)$$

In practice, one usually computes the N th moment for a single jet:

$$M_N^{\text{jet}} = \frac{\sum_{h \in \text{jet}} p_{t,h}^N}{\tilde{p}_t^N}, \quad (2.43)$$

and average over the jet sample, $M_N = \langle M_N^{\text{jet}} \rangle$.

Area–median subtraction. We can use an adaptation of the area–median approach to directly subtract the moments M_N of the fragmentation function.

The procedure to subtract the FF moments is quite similar to the main area–median approach. One first determines the expected background contribution per unit area to a given moment (or rather to its numerator in Eq. (2.43)):

$$\rho_N = \text{median}_{\text{patches}} \left\{ \frac{\sum_{i \in \text{patch}} p_{t,i}^N}{A_{\text{patch}}} \right\}. \quad (2.44)$$

For example, for $N = 0$ this will be the median particle multiplicity per unit area. The subtracted FF moment is then obtained by separately taking the numerator and denominator of Eq. (2.43), as measured in the full event, and respectively subtracting $\rho_N A$ and ρA :

$$M_N^{\text{sub}} = \frac{\sum_{i \in \text{jet}} p_{t,i}^N - \rho_N A}{(p_{t,\text{full}} - \rho A)^N}. \quad (2.45)$$

Improved background-subtracted fragmentation function. In Section 4.7, we will apply this subtraction method to FF measurements in heavy-ion collisions and observe that it is insufficient to fully correct the jet FF moments. We believe this can largely be attributed to background fluctuations and we describe below an improved subtraction method that corrects for the dominant fluctuation effects.

Background fluctuations mean that neither the jet's p_t nor the numerator of the FF moment are perfectly reconstructed for any given jet. When selecting jets above some p_t threshold for a process with a steeply falling jet spectrum, it is favourable to select jets slightly below the p_t threshold, but which have an upwards background fluctuation (cf. the discussion of Section 2.2). One consequence of this is that the denominator in Eq. (2.45) tends to be larger than the actual jet p_t . The larger the value of N , the greater the impact of this effect. One can also understand it as resulting in an underestimate of the z fraction, leading to the FF being spuriously shifted to lower z . Such an effect is already well known: CMS [32], for example, when comparing with pp FFs, explicitly applies a correction to the pp FFs to account for the smearing of the denominator of the z variable that is

expected when the jet’s momentum is reconstructed in a HI environment. ATLAS [33] carries out an unfolding to account for this.

A second consequence of fluctuations affects mostly the low- z , or low- N region of the FF: the upwards fluctuations of the background p_t that cause a jet to pass the p_t cut even when it is below threshold also tend to be associated with upwards fluctuations of the multiplicity of soft background particles. It is this effect that causes the FF to be overestimated for low values of N .²²

One way of verifying the above interpretation is to consider γ +jet events, selecting events based on the photon p_t and normalising the fragmentation function z also to the photon p_t . In this case, the background-induced fluctuations of the jet’s reconstructed p_t are of no relevance, and in explicit simulations, we have found that plain subtraction is already quite effective.

For the dijet case, in the limit where the background fluctuations are reasonably small compared to the transverse momentum of the jet, it is possible to devise a simple correction in moment space for both kinds of fluctuation-induced bias. In order to do so, we start by rewriting the subtracted moments of the jet FF in Eq. (2.45) as

$$M_N^{\text{sub}} = \frac{\sum_i p_{t,i}^N - \rho_N A}{(p_{t,\text{full}} - \rho A)^N} \equiv \frac{S_N}{S_1^N} \quad (2.46)$$

where we have introduced the shorthands S_N for the subtracted numerator and S_1 for $p_{t,\text{full}}^{\text{sub}}$, as defined by the area–median subtraction, Eq. (2.24). Given our assumption that background fluctuations are moderate, we can locally approximate the hard jet cross-section by an exponential (as already done in Section 2.2)

$$H(p_t) \equiv \frac{d\sigma}{dp_t} = \frac{\sigma_0}{\mu} \exp(-p_t/\mu). \quad (2.47)$$

We then take a Gaussian approximation for the spectrum of background transverse-momentum fluctuations from one jet to the next. Denoting the fluctuation by q_t , the distribution of q_t is then

$$B(q_t) \equiv \frac{dP}{dq_t} = \frac{1}{\sqrt{2\pi A}\sigma} \exp\left(-\frac{q_t^2}{2\sigma^2 A}\right), \quad (2.48)$$

where σ is a parameter that describes the size of fluctuations from one patch of area 1 to another. For patches of area A , the variance of the q_t across the patches is $\sigma^2 A$.

We further introduce the variable Q_N to denote the difference between the actual background contribution to S_N in a specific jet and the expected contribution, $\rho_N A$, i.e.

$$Q_N = \left(\sum_{i \in \text{jet (bkgd)}} k_{t,i}^N \right) - \rho_N A, \quad (2.49)$$

where the sum runs just over the background constituents $k_{t,i}$ of the jet. By construction, $Q_1 = q_t$. In practice Q_N cannot be determined for a single jet, since we don’t know which particles are the background ones, but its statistical properties can be determined by looking at many jets, and we denote the standard deviation of its fluctuations from one patch of area 1 to another by σ_N , in direct analogy to σ . The fluctuations of Q_N are not independent of the fluctuations q_t of the background’s transverse momentum: there is a correlation coefficient r_N between them, defined as

$$r_N = \frac{\text{Cov}(q_t, Q_N)}{\sqrt{\text{Var}(q_t)\text{Var}(Q_N)}}, \quad (2.50)$$

²²ATLAS correct for this [33] by multiplying the expected background contribution by a z -independent factor that is a function of the jet p_t .

where $\text{Var}(X)$ is the variance of the variable X and $\text{Cov}(X, Y)$ the covariance of X and Y . Using the fact that $\text{Var}(q_t) = \sigma^2 A$, $\text{Var}(Q_N) = \sigma_N^2 A$, we have that the average value for Q_N as a function of q_t is

$$\langle Q_N \rangle(q_t) = \frac{\text{Cov}(q_t, Q_N)}{\text{Var}(q_t)} q_t = r_N \frac{\sigma_N}{\sigma} q_t. \quad (2.51)$$

With these ingredients we can now correct for the fluctuation effects as follows. If we measure a certain value S_N in a jet, then as a function of the q_t fluctuation, the expected true hard contribution to it is

$$S_N^{\text{hard}} = S_N - \langle Q_N \rangle(q_t) = S_N - r_N \frac{\sigma_N}{\sigma} q_t, \quad (2.52)$$

where we have averaged over possible Q_N values, given the q_t fluctuation. To obtain our estimate for M_N^{hard} , this should be normalised by the N^{th} power of the true hard p_t of the jet, $(S_1 - q_t)^N$:

$$M_N^{\text{hard}} = \frac{S_N^{\text{hard}}}{(S_1 - q_t)^N} = \frac{S_N}{S_1^N} + N \frac{S_N q_t}{S_1^{N+1}} - r_N \frac{\sigma_N q_t}{\sigma S_1^N} + \mathcal{O}(q_t^2). \quad (2.53)$$

One subtlety here is that this is an estimate for M_N^{hard} in hard jets with true $p_t = S_1 - q_t$. However because M_N is a slowly varying function of p_t , by taking the result as contributing to $M_N^{\text{hard}}(S_1)$ rather than $M_N^{\text{hard}}(S_1 - q_t)$ we make only a small mistake, of the same order as other terms that we shall neglect.²³

Following Eq. (2.10), the average fluctuation q_t for a given reconstructed S_1 is

$$\langle q_t \rangle = \frac{\sigma^2 A}{\mu}, \quad (2.54)$$

therefore, retaining the terms linear in q_t in Eq. (2.53) and averaging now over possible q_t values, leads to the following prescription for an ‘‘improved’’ subtracted $M_N(S_1)$, corrected for fluctuations effects up to first order in q_t/S_1 :

$$\boxed{M_N^{\text{sub,imp}} = M_N^{\text{sub}} \times \left(1 + N \frac{\sigma^2 A}{S_1 \mu} \right) - r_N \frac{\sigma \sigma_N A}{\mu S_1^N}}. \quad (2.55)$$

This is simpler than the corresponding correction would be directly in z space, in particular because in z space the correction to one bin of the fragmentation function depends in a non-trivial way on the contents of nearby bins. One can think of the advantages of moment space as being that the correction to a given N value does not depend on M_N at all other values of N , and that it is straightforward to account for correlations between fluctuations in the jet- p_t and in the moments.

Note that in a real experimental context, calorimeter fluctuations of the reconstructed jet and background p_t 's would have an effect akin to increasing σ and decreasing the correlation coefficient r_N . The in-situ methods that we use for the determination of σ , σ_N and r_N would automatically take this into account. Noise-reduction methods in the reconstruction of the jet p_t , as used by CMS [34],

²³This can be seen by observing that M_N satisfies a DGLAP-style equation for its evolution $dM_N/d \ln p_t \sim \alpha_s M_N$. Given that q_t is itself small compared to p_t , M_N for jets with $p_t = S_1 - q_t$ differs from that for jets with $p_t = S_1$ by a relative amount $\sim \alpha_s q_t/S_1$, which we can neglect in the same way that we neglect $\mathcal{O}(q_t^2)$ terms. If we wanted to improve on this approximation, then one approach might be, for each event, to use the middle expression in Eq. (2.53) and assign it to M_N at $p_t = S_1 - \langle q_t \rangle$. We have not, however, investigated this option in detail and other improvements would probably also be necessary at a similar accuracy, e.g. taking into account deviations from the simple exponential and Gaussian approximations that we have used for $H(p_t)$ and $B(q_t)$.

would have the effect of reducing σ (and probably also r_N). However noise reduction is likely to complicate the meaningful determination of r_N , since it acts differently on pure background jets as compared to jets with a hard fragmenting component.

The correction in Eq. (2.55) can be applied jet-by-jet to correct for the fluctuation effects. We shall discuss the practical application of these techniques in the case of heavy-ion collisions in Section 4.7.

2.7 Numerical Implementation

All of the methods described above are made numerically available either directly in `FastJet` (for the standard application to p_t and 4-vector subtraction) or in various contributions to `fjcontrib` (for the more advanced applications). We describe in this Section some details of their implementation and leave to the next Section the recommendations on how to use the area–median pileup subtraction technique in practical applications.

Following the same logic as above, the first ingredient that is needed is a computation of the jet areas and the second ingredient is an implementation of the estimation of the background density ρ .

In what follows, we assume that the reader is familiar with the basic usage of `FastJet` for jet clustering, i.e. is familiar with the usage of the three basic `FastJet` classes: `PseudoJet`, `JetDefinition`, `ClusterSequence`. We refer to the `FastJet` manual [15] for more details.

A complete description of the `FastJet` classes related to jet areas and background subtraction can be found in the `FastJet` manual. Here we just give a quick overview of the relevant classes and their typical usage.

2.7.1 Clustering with jet areas

All three main types of jet areas (active, passive and Voronoi) are directly available from `FastJet`. Compared to standard clustering, the user is exposed to two new classes:

```
class fastjet::AreaDefinition;
class fastjet::ClusterSequenceArea;
```

with input particles, a jet definition and an area definition being supplied to a `ClusterSequenceArea` in order to obtain jets with area information. Typical usage would be as follows:

```
// include the appropriate header
#include "fastjet/ClusterSequenceArea.hh"

// define an area with ghosts up to the particle acceptance
double ghost_maxrap = 5.0; // e.g. if particles go up to y=5
AreaDefinition area_def(active_area_explicit_ghosts, GhostedAreaSpec(ghost_maxrap));

// cluster using that definition of jet areas
ClusterSequenceArea clust_seq(input_particles, jet_def, area_def);

// retrieve the jets and show how to get their area
vector<PseudoJet> jets = sorted_by_pt(clust_seq.inclusive_jets());
double area_hardest_jet = jets[0].area();
```

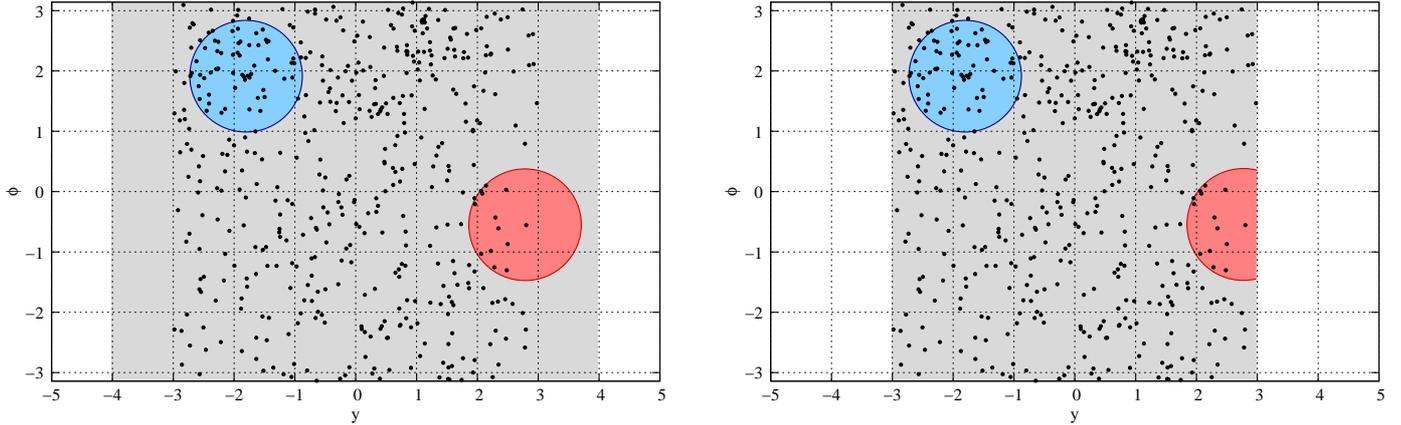


Figure 2.4: Two choices for ghost placement. The grey area in each plot indicates the region containing ghosts, while the dots are particles, here accepted up to $|y| < 3$. The circular regions indicate the areas that will be found for two particular jets. In the left-hand case, with ghosts that extend well beyond the acceptance for particles, jet areas are unaffected by the particle acceptance; in the right-hand case, with ghosts placed only up to the acceptance limit, jet areas are clipped at the edge of the acceptance.

Defining areas. For active areas, pertinent for pileup subtraction, one would use

```
AreaDefinition(fastjet::AreaType area_type, fastjet::GhosedAreaSpec ghost_spec);
```

with an `area_type` being one of the two active area choices:

```
enum AreaType{ [...], active_area, active_area_explicit_ghosts, [...]};
```

The two variants give identical results for the areas of hard jets. The second one explicitly includes the ghosts when the user requests the constituents of a jet and also leads to the presence of “pure ghost” jets.

The second argument used to define an `AreaDefinition` specifies how to distribute the ghosts. This is done via the class `GhosedAreaSpec` whose constructor is

```
GhosedAreaSpec(double ghost_maxrap,
               int   repeat          = 1,   double ghost_area      = 0.01,
               double grid_scatter   = 1.0, double pt_scatter     = 0.1,
               double mean_ghost_pt  = 1e-100);
```

The ghosts are distributed on a uniform grid in y and ϕ , with small random fluctuations to avoid clustering degeneracies.

The `ghost_maxrap` variable defines the maximum rapidity up to which ghosts are generated. If one places ghosts well beyond the particle acceptance (at least R beyond it), then jet areas also stretch beyond the acceptance, giving a measure of the jet’s full extent in rapidity and azimuth. If ghosts are placed only up to the particle acceptance, then the jet areas are clipped at that acceptance and correspond more closely to a measure of the jet’s susceptibility to contamination from accepted soft particles. The two choices are illustrated in fig. 2.4. To define more complicated ghost acceptances it is possible to replace `ghost_maxrap` with a `Selector`, which must be purely geometrical and have finite rapidity extent.

The two other relevant arguments to construct a `GhosedAreaSpec` object are `ghost_area` and `n_repeat`. The `ghost_area` parameter is the area associated with a single ghost. Smaller values will mean more accurate determinations of the jet area at the cost of a longer CPU-time for clustering.

We have found the default ghost area given above to be adequate for most applications, with smaller ghost areas beneficial mainly for high-precision determinations of areas of jets with small R .

Since the ghost positions and p_t 's include small random fluctuations, clustering with different sets of ghosts will give different jet areas. The `n_repeat` parameter controls how many sets of ghosts will be used. With `repeat > 1`, the final area is averaged over the different sets.²⁴ For most applications, including those related to pileup mitigation, keeping the default value `n_repeat = 1` is a good default. Finally, the other parameters of `GhostedAreaSpec` control the size of the position and p_t fluctuations and can safely be kept to their default values.

Note that, for the sake of reproducibility, facilities are also available to set and retrieve the seeds for the random-number generator. This is done through the `set_random_status(...)` and `get_random_status(...)` members of `GhostedAreaSpec`.

Clustering with areas. As mentioned above, once an `AreaDefinition` has been constructed, the clustering with jet-area calculation is obtained using the `ClusterSequenceArea` class which is constructed using

```
template<class L> ClusterSequenceArea(const std::vector<L> & input_particles,
                                     const JetDefinition & jet_def,
                                     const AreaDefinition & area_def);
```

As long as an instance of this class is in scope, a user can access information about the area of its jets using the following methods of `PseudoJet`:

```
/// Returns the scalar area associated with the given jet
double area = jet.area();

/// Returns the error (uncertainty) associated with the determination of the
/// scalar area of the jet; gives zero when the repeat=1 and for passive/Voronoi areas
double area_error = jet.area_error();

/// Returns the 4-vector area associated with the given jet
PseudoJet area_4vector = jet.area_4vector();

/// When using active_area_explicit_ghosts, returns true for jets made
/// exclusively of ghosts and for ghost constituents.
bool is_pure_ghost = jet.is_pure_ghost();
```

It is also possible to specify explicitly the set of ghosts to be used provided one directly uses the more basic `ClusterSequenceActiveAreaExplicitGhosts` class:

```
template<class L> ClusterSequenceActiveAreaExplicitGhosts
  (const std::vector<L> & pseudojets,
   const JetDefinition & jet_def_in,
   const std::vector<L> & ghosts,
   double ghost_area,
   const bool & writeout_combinations = false);
```

This is useful e.g. to keep area information in tools which require a re-clustering.

²⁴And a pseudo-Monte Carlo estimate of the error is available through the `jet.area_error()` call.

2.7.2 Background estimation and subtraction

FastJet natively provide implementations of the two main methods to estimate the background density ρ using equation (2.23).

2.7.2.1 Background estimation.

The simplest workflow for background estimation is first, typically outside the event loop, to create a background estimator. For the grid based method one creates a `fastjet::GridMedianBackgroundEstimator`,

```
#include "fastjet/tools/GridMedianBackgroundEstimator.hh"
// ...
GridMedianBackgroundEstimator bge(double max_rapidity,
                                   double requested_grid_spacing);
```

where `max_rapidity` is the maximal $|y|$ covered by the grid and `requested_grid_spacing` the (requested) grid-cell size.

For the jet-based method, one creates a `JetMedianBackgroundEstimator`,

```
#include "fastjet/tools/JetMedianBackgroundEstimator.hh"
// ...
JetMedianBackgroundEstimator bge(const Selector & selector,
                                   const JetDefinition & jet_def,
                                   const AreaDefinition & area_def);
```

where the selector is used to indicate which jets are used for background estimation (for most use cases, one just specifies a rapidity range, e.g. `SelectorAbsRapMax(4.5)` to use all jets with $|y| < 4.5$), together with a jet definition and an area definition.²⁵ An active area with explicit ghosts is recommended.²⁶

Both of the above background estimators derive from a `fastjet::BackgroundEstimatorBase` class and the remaining functionality is common to both. In particular, for each event, one tells the background estimator about the event particles,

```
bge.set_particles(event_particles);
```

where `event_particles` is a vector of `PseudoJet`, and then extracts the background density and a measure of its fluctuations with the two following calls

```
// the median of (pt/area) for grid cells, or for jets that pass the selection cut,
// making use also of information on empty area in the event (in the jets case)
rho = bge.rho();

// an estimate of the fluctuations in the pt density per unit  $\sqrt{A}$ ,
sigma = bge.sigma();
```

²⁵It is also possible to construct a `JetMedianBackgroundEstimator` from a `ClusterSequenceAreaBase` object and a `Selector`. The latter is useful e.g. when one uses several estimators sharing the same clustering to define their patches.

²⁶With the k_t algorithm one may also use a Voronoi area (`effective_Rfact = 0.9` is recommended), which has the advantage of being deterministic and faster than ghosted areas. In this case however one must use a selector that is geometrical and selects only jets well within the range of event particles. E.g. if particles are present up to $|y| = y_{\max}$ one should only use jets with $|y| \lesssim y_{\max} - R$. When using ghosts, the selector can instead go right up to the edge of the acceptance, if the ghosts also only go right up to the edge, as in the right-hand plot of fig. 2.4.

2.7.2.2 Background subtraction.

The `Subtractor` class, defined in `include/tools/Subtractor.hh`, directly implements the subtraction method (2.24). Its constructor takes a pointer to a background estimator:²⁷

```
GridMedianBackgroundEstimator bge(...); // or a jet-based estimator
Subtractor subtractor(&bge);
```

The subtractor can then be used as follows:

```
PseudoJet jet;
PseudoJet subtracted_jet = subtractor(jet);
vector<PseudoJet> jets;
vector<PseudoJet> subtracted_jets = subtractor(jets);
```

The subtractor normally returns `jet - bge.rho(jet)*jet.area_4vector()`. If `jet.pt()` is smaller than `bge.rho(jet)*jet.area_4vector().pt()`, then the subtractor instead returns a jet with zero 4-momentum (so that `(subtracted_jet==0)` returns `true`). In both cases, the returned jet retains the user and structural information of the original jet.

Note that the `Subtractor` has an option

```
void set_known_selectors(const Selector & sel_known_vertex,
                        const Selector & sel_leading_vertex);
```

The idea here is that there are contexts where it is possible, for some of a jet’s constituents, to identify which vertex they come from. In that case it is possible to provide a user-defined a `Selector` that indicates whether a particle comes from an identified vertex or not and a second user-defined `Selector` that indicates whether that vertex was the leading vertex. The 4-momentum from the non-leading vertex is then discarded, that from the leading vertex is kept, and subtraction is applied to component that is not from identified vertices. It follows that ρ must correspond only to the momentum flow from non-identified vertices. With this setup, the jet p_t is bounded to be at least equal to that from the leading vertex, as is the mass if the “safe” subtraction option is enabled (see below). The typical application would be CHS events where charged particles would have an identified vertex (see Chapter 11 for more details).

2.7.2.3 Positional dependence of background

The positional dependence can be treated using one of the two techniques introduced in Section 2.4.2 where, now, the properties ρ and σ depend on the position of the jet.

In `FastJet`, the properties of the background are to be obtained through the methods (available for both `JetMedianBackgroundEstimator` and `GridMedianBackgroundEstimator`)

```
double rho (const PseudoJet & jet); // pt density per unit area A near jet
double sigma(const PseudoJet & jet); // fluctuations in the pt density near jet
```

Rescaling method. To account for positional dependence of the background using the “rescaling” technique, one first need to parametrise the average shape of the rapidity dependence from some number of pileup events. Then for subsequent event-by-event background determinations, one carries out a global ρ determination and then applies the previously determined average rescaling function to that global determination to obtain an estimate for ρ in the neighbourhood of a specific jet.

²⁷Alternatively, it is also possible to construct the `Subtractor` with a fixed value for ρ .

The rescaling approach is available for both grid and jet-based methods. To encode the background shape, the user defines a rescaling function derived from `FunctionOfPseudoJet<double>` (see the `FastJet` manual for details). `FastJet` provides a 4th order polynomial in y as a simple illustrative example:

```
// gives rescaling(y) = 1.16 + 0 · y - 0.023 · y2 + 0 · y3 + 0.000041 · y4
fastjet::BackgroundRescalingYPolynomial rescaling(1.16, 0, -0.023, 0, 0.000041);
```

Then one tells the background estimator (whether jet or grid based) about the rescaling with the call

```
bge.set_rescaling_class(&rescaling);
```

Subsequent calls to `rho()` will return the median of the distribution $p_t/A/\text{rescaling}(y)$ (rather than p_t/A). Any calls to `rho(jet)` and `sigma(jet)` will include an additional factor of `rescaling(yjet)`. Note that any overall factor in the rescaling function cancels out for `rho(jet)` and `sigma(jet)`, but not for calls to `rho()` and `sigma()` (which are in any case less meaningful when a rapidity dependence is being assumed for the background).

Local estimation (jet based). The “local estimation” technique is available for now only in the case of the jet-based estimator, involves constructing a `JetMedianBackgroundEstimator` with a `Selector` retaining only jets in a local range. Such a `Selector` should be able to take a reference, like `SelectorStrip(Δy)` which keeps jets in a strip of half-width Δy centred on the reference jet. When the user calls either `rho(jet)` or `sigma(jet)`, `jet` will be used as the reference for the `Selector` and only the jets in the event that pass the cut specified by this newly positioned `selector` are used to estimate ρ or σ . This method is adequate if the number of jets that pass the selector is much larger than the number of hard jets in the range (otherwise the median p_t/A will be noticeably biased by the hard jets). One can attempt to remove some given number of hard jets before carrying out the median estimation, e.g. with a `selector` such as

```
selector = SelectorStrip( $\Delta y$ ) * (!SelectorNHardest(2))
```

which removes the 2 hardest jets globally and then, of the remainder, takes the ones within the strip.²⁸ This is however not always very effective, because one may not know how many hard jets to remove.

Estimation in regions (grid based). The grid-based estimator does not currently provide for local estimates, in the sense that the set of tiles used for a given instance of the grid-based estimator is always fixed. However, as of `FastJet` 3.1, it is possible to obtain relatively fine control over which fixed set of tiles a grid-based estimator uses. This is done with the help of the `RectangularGrid` class

```
RectangularGrid grid(rap_min, rap_max, rap_spacing, phi_spacing, selector);
GridMedianBackgroundEstimator bge(grid);
```

A given grid tile will be used only if a massless `PseudoJet` placed at the centre of the tile passes the `selector`. So, for example, to obtain an estimate for ρ based on the activity in the two forward regions of a detector, one might set `rap_min` and `rap_max` to cover the whole detector regions and then supply a `SelectorAbsRapRange(rap_central_max, rap_max)` to select just the two forward regions.

²⁸If you use non-geometric selectors such as this in determining ρ , the area must have explicit ghosts in order to simplify the determination of the empty area. If it does not, an error will be thrown.

2.7.2.4 Handling masses and safe subtraction

There are several subtleties in handling masses in subtraction. The first is related to the hadrons masses which, as we have seen in Section 2.4.3, requires the introduction of the ρ_m property of the background. Since `FastJet` 3.1, the grid and jet-median background estimators automatically determine ρ_m by default. It is returned from a call to `rho_m()` or `rho_m(jet)`, with fluctuations accessible through `sigma_m()` or `sigma_m(jet)`. The determination of ρ_m involves a small speed penalty and can be disabled with a call to `set_compute_rho_m(false)` for both background estimators.

To avoid changes in results relative to version 3.0, by default `FastJet` 3.1 and 3.2 do not use ρ_m in the `Subtractor`, i.e. it uses Eq. (2.24) instead of the more complete Eq. (2.30). However, for a given subtractor, a call to `set_use_rho_m(true)`, will cause it to instead use Eq. (2.30) for all subsequent subtractions. We *strongly recommend* switching this on if your input particles have masses, and in future versions of `FastJet` we may change the default so that it is automatically switched on.²⁹ An alternative is to make the input particles massless.

A second issue, relevant both for Eqs. (2.24) and (2.30), is that sometimes the resulting squared jet mass is negative.³⁰ This is obviously unphysical. By default the 4-vector returned by the subtractor is left in that unphysical state, so that the user can decide what to do with it. For most applications a sensible behaviour is to adjust the 4-vector so as to maintain its p_t and azimuth, while setting the mass to zero. This behaviour can be enabled for a given subtractor by a call to its `set_safe_mass(true)` function (available since v.3.1). In this case the rapidity is then taken to be that of the original unsubtracted jet.

For situations where “known selectors” have been specified, `set_safe_mass(true)` will impose stronger constraints: if the transverse momentum of the jet after subtraction is smaller than the transverse momentum from all the known particles from the leading vertex, we set the subtracted 4-momentum to the sum of the known particles coming from the leading vertex. If the subtracted jet has larger transverse momentum but smaller mass, we set the transverse momentum and azimuthal angle of the subtracted jet to the result of the pileup subtraction and set the rapidity and mass from the sum of the known particles coming from the leading vertex. For jets with no known particles coming from the leading vertex, whenever the subtracted 4-vector has an ill-defined rapidity or azimuthal angle, we use those of the original, unsubtracted, jet.

In future versions of `FastJet`, the default behaviour may be changed so that “safe” subtraction is automatically enabled.

2.7.3 Subtraction for jet shapes

The area–median subtraction procedure for jet shapes introduced in Section 2.5 is not directly included in `FastJet` but is made available as `GenericSubtractor` in `fjcontrib`.

This introduces a class named `GenericSubtractor`. In practice, the code varies the ghost transverse momentum in order to estimate the derivatives (2.31). Note that to do that, it brings the ghost p_t scale from 10^{-100} to a value where its effect on the shape is tractable within machine precision.

The basic usage of this class is as follows:

```
GenericSubtractor gensub(&some_background_estimator);
FunctionOfPseudoJet<double> shape; // whatever shape you are using
double subtracted_shape_value = gensub(shape, jet);
```

²⁹It is also possible to construct a `Subtractor subtractor(ρ , ρ_m)` with explicit ρ and ρ_m values; if this is done, then ρ_m use *is enabled* by default.

³⁰If $m^2 < 0$, `m()` returns $-\sqrt{-m^2}$, to avoid having to introduce complex numbers just for this special case.

By default, this only subtracts the transverse momentum density ρ (obtained from the background estimator provided to the class). Two options allow also for the inclusion of the ρ_m term: one can either instruct the `GenericSubtractor` object to compute ρ_m from the same background estimator using³¹

```
gensub.use_common_bge_for_rho_and_rhom();
```

or explicitly construct `gensub` with two background estimators³²

```
GenericSubtractor gensub(& bge_for_rho, & bge_for_rhom);
```

In all cases, extra information (first order subtraction, derivatives, ...) can be retrieved using

```
GenericSubtractorInfo info;
double subtracted_shape_value = gensub(shape, jet, info);
```

(see directly the class definition in `GenericSubtractor.hh` for details)

Finally, there are some cases where it is possible to use clever tricks to improve the subtraction or its efficiency. The first case is the one when a shape is made of several components and one wishes to subtract each of these components independently rather than subtracting the shape as a whole. (A typical case would be ratios of shapes.) To do that, one derives the shape from `ShapeWithComponents`. The second case is the one where the calculation of a shape first partitions the jet into subjets before proceeding to an actual calculation of the shape. To avoid recomputing that partition every time `GenericSubtraction` varies the ghost scale, one can derive such a shape from the `ShapeWithPartition`. We refer to the complete documentation of the `GenericSubtractor` contrib for more details.

2.7.4 Subtraction for fragmentation function moments

The technique to subtract jet fragmentation function moments introduced in Section 2.6 has been implemented in the `JetFFMoments` contribution to `fjcontrib`.

It introduces a class, `JetFFMoments`, which computes the subtracted FF moments for given values of N . It is constructed in one of the two following ways:

```
/// constructor from a vector of N values, with an optional background estimator
JetFFMoments(const std::vector<double> & ns, JetMedianBackgroundEstimator *bge=0);

/// constructor using nn regularly-spaced values of N between nmin and nmax.
JetFFMoments(double nmin, double nmax, unsigned int nn,
              JetMedianBackgroundEstimator *bge=0);
```

specifying the values of N either explicitly or as a number of values between a minimum and a maximum value. The class can be used to compute just the FF moments, without any form of background subtraction, by not providing a `JetMedianBackgroundEstimator` to the constructor.

To obtain the FF moments corresponding to a given jet, one uses

```
virtual vector<double> operator()(const PseudoJet &jet) const;
virtual vector<double> operator()(const PseudoJet &jet, Info & info) const;
virtual vector<vector<double> > operator()(const vector<PseudoJet> &jets) const;
```

³¹Since `FastJet` 3.1, this option would work with any background estimator that internally computes ρ_m (and has that computation enabled). For `FastJet` 3.0, it is only available for `JetMedianBackgroundEstimator`.

³²In that case, ρ_m will be estimated using the `rho()` method of `_bge_for_rhom` — and not `rho_m()` — unless one explicitly calls `use_rhom_from_bge_rhom(bool value=true);`. Also, if ρ and, optionally, ρ_m are known independently, the subtractor can be constructed directly as `GenericSubtractor gensub(ρ, ρ_m)`.

The first version just returns a vector with the FF moments computed at the requested values of N . The third version does the same for a vector of jets. The second version requires an extra argument of the type `JetFFMoments::Info` which stores potentially useful intermediate information used during the calculation (see the internal contrib’s documentation for details).

We can control how FF moments are normalised using

```
void use_scalar_sum(bool use_scalar_sum = true);
void set_return_numerator(bool return_numerator);
void set_denominator(double norm);
```

The first decides whether the normalisation uses the scalar p_t of the jet, \tilde{p}_t of the jet p_t itself, with the former being the default; the second only computes the numerator (i.e. sets the denominator to 1); and the third sets the denominator to a fixed value, as one would do in, say, γ +jet events when one wants to use the p_t of the photon to normalise the jet FF moments.

Finally, improved subtraction is turned on by calling³³

```
/// mu is the slope  $\mu$  of  $d\sigma/dp_t$  parametrised as  $\sigma_0 \exp(-p_t/\mu)$ 
/// in the vicinity of the jet
void set_improved_subtraction(double mu);
```

2.8 Recommendations for practical usage

We have introduced all the concepts and ingredients needed to apply the area–median pileup subtraction method. However, at the end of the day, several details have to be specified for a practical use. We give below our recommendations for applications at the LHC and we will provide extensive studies and tests of these settings in the next chapters. For simplicity, we assume that we have events where particles are kept up to a rapidity $|y| = y_{\max}$.

Jet area type. To define the area of the jets, an active area has to be used. We strongly advise the use of explicit ghosts (i.e. an area of the `active_area_explicit_ghosts` type) which allows for an explicit computation of the contribution from the “empty patches” to (2.23). Furthermore, having ghosts explicitly kept makes possible the calculation of jet areas in more elaborate jet processing like substructure tools. Note that `FastJet` provides a `PseudoJet::is_pure_ghost()` method which tells whether a jet is only made of ghosts or not, and a `SelectorIsPureGhosts` selector which can be used e.g. to eliminate the ghosts when asking for the constituents of a jet.

Ghosts used for jet areas. Using active areas means that one has to specify the placement of the ghosts. With particles up to $|y| = y_{\max}$, we advise to place the ghosts up to a rapidity y_{\max} too (see the right of Figure 2.4). We have found that in practice, using a unique set of ghosts, i.e. setting `n_repeat=1`, is sufficient. It also allows for the set of explicit ghosts to be manipulated in jet tools without ambiguity.³⁴ For many applications, the default fundamental ghost area, `ghost_area`, of 0.01 is sufficient. For the production of the final results, it might nevertheless be useful to set this to a smaller value. Although this has a potentially large impact on timings, it also tends to

³³This requires at least version 3.1 of `FastJet`. For the `FastJet` 3.0 series, additional arguments have to be provided. With improved subtraction switched on, the computation of FF moments automatically includes the correction (2.55).

³⁴When using explicit ghosts with `n_repeat` larger than 1, the last set of ghosts is retained. For full usage with `n_repeat` larger than 1 one should retain explicitly all the sets of ghosts, a feature which is not implemented in `FastJet`.

improve the resolution we obtain at the end of the subtraction procedure (see e.g. Fig.4.12). Finally, if one wants to ensure reproducibility, one can always retrieve and set the random seed using `GhostedAreaSpec::get_random_status(...)` and `GhostedAreaSpec::set_random_status(...)`.

ρ determination. We recommend the use of the `GridMedianBackgroundEstimator`, i.e. the estimation of the background properties based on a grid-based computation of the patches, mainly for speed reasons.³⁵ We have found that, in practice, a grid size in the range 0.5-0.7 leads to good estimates of ρ , with 0.55 (our recommended default) providing a good default across a wide range of processes and pileup conditions.

Positional dependence. To handle the non-negligible rapidity dependence of pileup (and in the case of heavy-ion collisions, also the azimuthal dependence of the underlying event), we tend to prefer the rescaling method, with local estimation also a useful option e.g. in the case of heavy-ion collisions. Note that, for the sake of convenience, `FastJet` provides a parametrisation of a 4th order polynomial which has been sufficient for our own studies. This specific form is not meant to be a recommendation and user-defined rescaling functions should be used if the simple function in `FastJet` fails to describe the positional dependence correctly.

Particle masses. If you are interested in jet masses and your input particles have non-zero masses, make sure you use the ρ_m component in Eq. (2.30) by calling your subtractor's `set_use_rho_m()` method.. You should also pay attention to what happens with negative squared masses, and consider calling the subtractor's `set_safe_mass()` option. For backwards compatibility reasons, both of these options are disabled by default (this may change in versions of `FastJet` larger than 3.2).

CHS events. The area–median pileup subtraction technique can also be used in conjunction with Charged-Hadron-Subtracted (CHS) types of events. In that case, all the pileup properties have to be estimated from the neutral particles or from the CHS events directly. Note that, optionally, charged tracks from pileup vertices can be kept as ghosts, i.e. with their momentum scaled down by a large factor. Alternatively, one can cluster the event including the (unmodified) charged tracks from pileup interactions and use something like `subtractor.set_known_selectors(SelectorIsCharged(), SelectorIsLeadingVertex())` so that the subtractor itself discards the charged tracks from pileup vertices.³⁶ CHS events will be discussed at length in Chapter 11.

Subtraction for groomed jets. For pileup mitigation in jet substructure studies, one usually works with fat jets — i.e. jets of large radius — for which one applies a grooming procedure to clean the soft contamination to the jets. Whenever possible, we strongly recommend to use also the area–median pileup subtractions technique. The main physical reason behind this is that grooming mostly works by reducing the jet area, replacing, say the original A_{jet} into a groomed area A_{groomed} . The contamination (2.1) still applies with $A = A_{\text{groomed}}$ and the area–median pileup subtraction would correct for it. Many of the common grooming tools have an implementation that supports attaching a `Subtractor` to them and provide automatic internal pileup subtraction. We will come back to this in our studies in Chapter 7.

³⁵The `JetMedianBackgroundEstimator`, with jets of radius in the 0.4-0.6 range, performs equally well but tends to be slower due to the extra clustering it involves.

³⁶This is not our recommendations because keeping the charged tracks from pileup interactions in the initial clustering would slightly increase the back-reaction effects.

Chapter 3

Monte-Carlo validation of approach

In this Chapter, we provide an extensive validation of the area–median pileup subtraction technique based on Monte-Carlo studies. Our goal is to validate two key aspects of the area–median method: (i) that it indeed corrects for pileup contamination and does it for a wide range of jet observables, (ii) that it is a robust method, fairly insensitive to variations of the pileup conditions and details of the hard event over wide kinematic ranges. We also want to justify our recommendations outlined in Section 2.8.

Before entering into this vast territory, we shall first briefly discuss the basic pileup properties ρ and σ . This is meant to build a cleaner quantitative picture of pileup effects and is done in Section 3.1. We shall then perform Monte Carlo studies of the area–median performance for several observables and several environments.

- *The jet p_t* (Section 3.2). This is by far the most common application, relevant for the vast majority of LHC analyses.
- *Jet shapes* (Section 3.3). This includes the jet mass as well as a series of other (infrared-and-collinear safe) jet observables, used in more specific LHC analyses.
- *Heavy-ion Underlying Events* (Chapter 4). We will show that the area–median method is also suited to subtract the overwhelming Underlying Event created in heavy-ion collisions at RHIC and at the LHC. This also includes a series of interesting analytic estimates in Section 4.4 which could also be relevant for applications to pileup subtraction in proton-proton collisions.
- *Jet fragmentation function* (Section 4.7). Still using heavy-ion collisions, we validate our subtraction method for fragmentation function moments.

For all these studies, we will explore various kinematic dependencies so as to assess the robustness of the area–median subtraction method. This includes varying the pileup multiplicity, the jet transverse momentum and rapidity, but also checking that the jet multiplicity in the hard event does not introduce large biases. On top of that, we also want to make sure that our findings do not depend on the specific details of Monte-Carlo generators. In the end, this translates into rather CPU-expensive simulations. Even though some of the simulations performed in the context of the original publications are slightly out of date (e.g. use old versions of Monte Carlo generators and tunes, or use different collision energies), we have decided to keep the original results and not update them to a more modern setup *except* for our study of the jet transverse momentum. There are several reasons justifying this approach. First, as already mentioned, the jet transverse momentum is by far the most important use-case and so we consider that it is important to revisit the original study done in [6] with a more

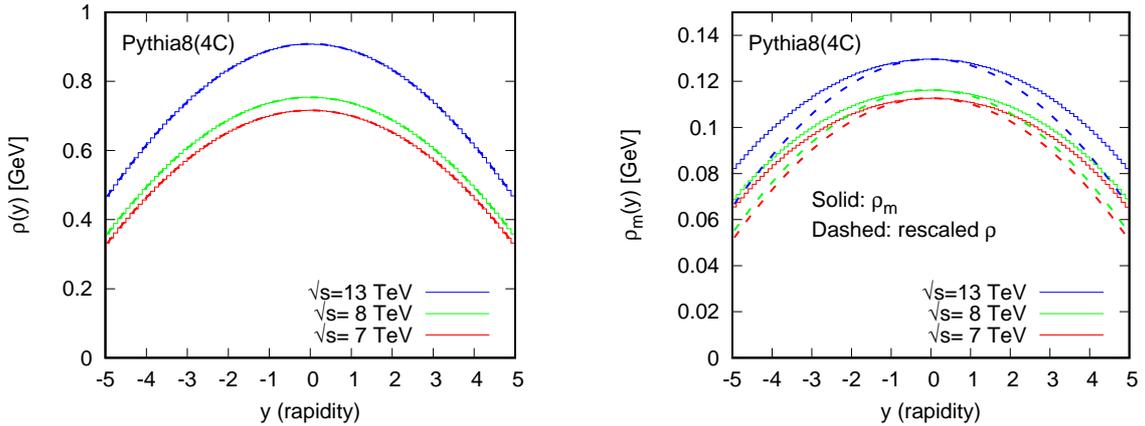


Figure 3.1: Rapidity dependence of the background density ρ calculated on minimum bias events as obtained by Pythia 8 with the tune 4C. Left: ρ profile; right: ρ_m profile. For all plots, we show the results for the 3 LHC energies, $\sqrt{s} = 7, 8$ and 13 TeV. The dashed lines on the rightmost plots represent the result of fitting ρ with a fourth order polynomial as in Eq. (3.1) and rescaling it so that it matches at $y = 0$.

modern setup.¹ This is also the study that provides most of the bases for our recommendations in Section 2.8, again justifying an up-to-date series of Monte-Carlo studies. Once this important baseline has been established, we will have a clear indication that the area–median subtraction procedure is robust and insensitive to the details of the simulation. We have therefore estimated that it was not critical to update the other Monte Carlo studies. We will discuss more at the beginning of each section the differences between the original simulations, the current (experimental) situation and the results presented in this document.

3.1 Properties of the pileup density

Before digging into the systematic validation of the area–median pileup subtraction technique, it is useful to have a quantitative idea of how large we should expect the pileup contamination to be.

In the case of pileup in pp collisions, we can estimate directly the pileup contamination from minimum bias collisions.

The first result of such a study is presented in Fig. 3.1. We have used Pythia 8.186 [35, 36], with tune 4C [37], to generate minimum bias events for the three main LHC centre-of-mass energies: 7, 8, and 13 TeV.² Let us first look at the leftmost plot, which shows the average transverse momentum density ρ in the event as a function of rapidity. In this case, it has been computed exactly, in each rapidity bin of width 0.1, from the minimum bias events. We see that, in central rapidity, each pileup event would, on average bring around 700 MeV at 7-8 TeV and around 900 MeV at 13 TeV. This drops down by about a factor of 2 if we go to forward rapidities, $|y| = 5$. The next average property of pileup that we have discussed is ρ_m , the contribution from particle masses, or, more precisely, from

¹The interested reader can also consult the original write-up with our initial findings and check that our conclusions are unaffected.

²The 4C tune gives reasonable agreement with a wide range of minimum-bias data, as can be seen by consulting MCPlots [38].

\sqrt{s}	a	b	c
7 TeV	0.7161921	-0.0159483	1.01133×10^{-5}
8 TeV	0.7538506	-0.0164948	1.44286×10^{-5}
13 TeV	0.9084412	-0.0189777	4.05981×10^{-5}

Table 3.1: Coefficients of the fit of the rapidity profile $\rho_{\text{seen}}(y)$, Eq. (3.1) for the three centre-of-mass energies under consideration.

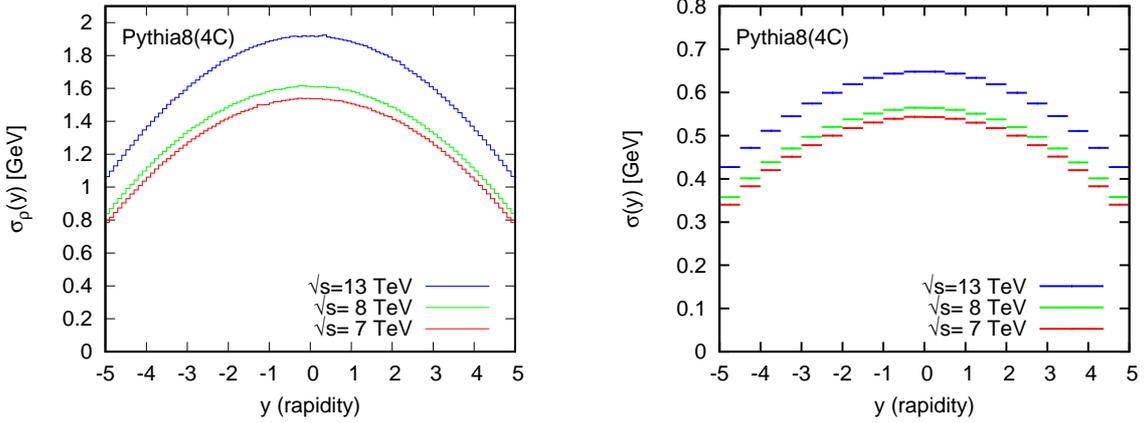


Figure 3.2: Rapidity dependence of the background fluctuations between different events (σ_ρ , left plot) and inside an event (σ , right plot).

$m_t - p_t$. This is shown on the right plot of Fig. 3.1 and it indicates that it is only about 15% of the transverse component ρ .

To handle the rapidity dependence of pileup when using the area–median subtraction method, we have seen that one approach is to use a rescaling function. The rescaling function itself can be readily extracted from minimum bias data. The profiles we obtain on Fig. 3.1 can very well be approximated by a fourth degree polynomial of the form

$$\rho(y) = (a + by^2 + cy^4) \text{ GeV}. \quad (3.1)$$

The fitted parameters for the different centre-of-mass energies are given in Table 3.1.

On the rightmost plot of Fig. 3.1, we have shown, using dashed lines, the result of rescaling the ρ rapidity profile so that it matches ρ_m at $y = 0$. We see that it reproduces the ρ_m rapidity dependence within about 15%.

Let us now move to the other fundamental properties of pileup, namely measures of its fluctuations: σ_ρ (between events) and σ (inside an event). The former is the simplest to measure: we just look at the dispersion of $\rho(y)$ across our event sample. The result is plotted on the left panel of Fig. 3.2 and it shows large fluctuations, reaching about 1.5 GeV for Run I energies and nearly 2 GeV at 13 TeV. Next, to estimate σ in a given rapidity range $y_1 < y < y_2$, we split the azimuthal angle range in several (25) cells and compute the dispersion of the cell scalar p_t divided by the square root of the cell area is an estimate of σ . The resulting rapidity distribution for σ is plotted on the right panel of Fig. 3.2. It shows a similar shape than ρ , and reaches ~ 500 (resp. ~ 650) MeV at central rapidities for 7 TeV (resp. 13 TeV) collisions.

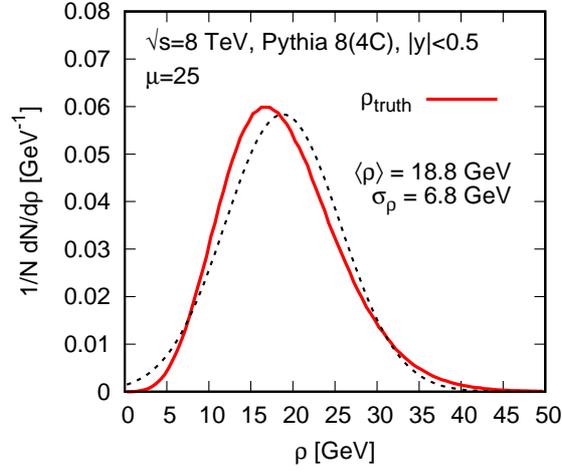


Figure 3.3: Distribution of the true pileup density ρ computed for $|y| < 0.5$ with Poisson-distributed events with $\mu = \langle N_{\text{PU}} \rangle = 25$. The dotted line represents a Gaussian fit with the average and dispersion indicated on the plot.

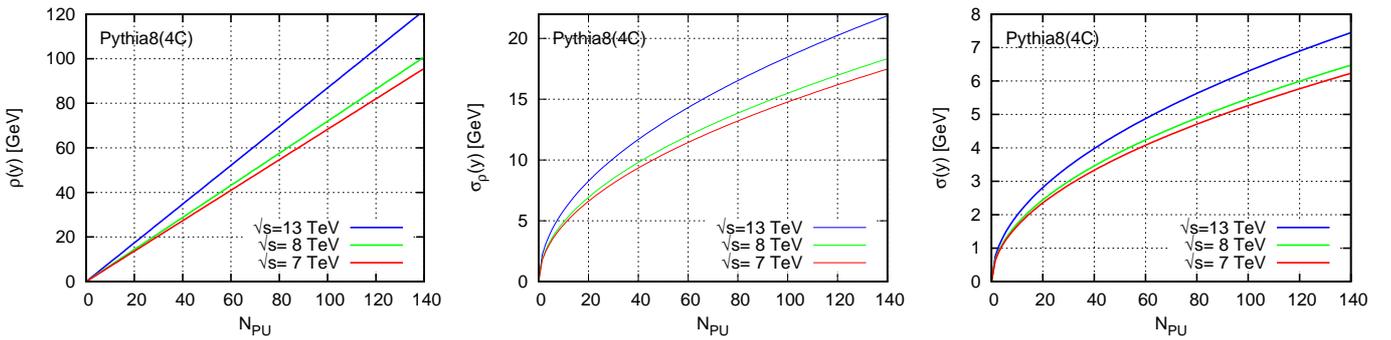


Figure 3.4: Evolution of the pileup fundamental properties as a function of N_{PU} . The properties are averaged for $|y| < 2.5$.

One might also wonder how precisely the distribution of the pileup density ρ in a given event is characterised by its average density $\langle\rho\rangle$ and its fluctuations σ_ρ . This is illustrated in Fig. 3.3 where we looked at the ρ distribution for a sample of Poisson-distributed pileup events. The dotted line, corresponding to a Gaussian fit, shows a reasonable description of the Monte-Carlo results. The Gaussian approximation appears therefore good enough for the simple analytic estimates we do in this document.

It is a trivial exercise to infer from the above numbers the pileup dependence of ρ , σ_ρ and σ . Based on our introductory discussion, we know that ρ will scale like N_{PU} while σ_ρ and σ scale like $\sqrt{N_{\text{PU}}}$. To illustrate this, we have averaged the minimum bias properties in the central region, between $|y| = 0$ and $|y| = 2.5$.³ The expected N_{PU} scaling is plotted on Fig. 3.4 for ρ , σ_ρ and σ respectively. If, for the sake of the argument, we assume that a jet of radius R has an area πR^2 , which is a good approximation for the standard anti- k_t jets, the pileup contamination into a jet can be inferred from these curves by multiplying ρ and σ_ρ by $A = \pi R^2$ and σ by $\sqrt{A} = \sqrt{\pi}R$. For a typical jet obtained with $R = 0.4$, this gives $A \approx 0.5$ and $\sqrt{A} \approx 0.7$.

3.2 Widest application: the jet transverse momentum

By far the most widely used property of a jet for which pileup subtraction is mandatory is its transverse momentum. It is used in a majority of analyses at the LHC. We shall therefore begin our systematic study of the area–median method by a study of its performance for the jet p_t .

This study serves different purposes. First, it is meant to be a validation of the area–median approach. To that aim, we shall also consider an alternative approach, based on the number of observed pileup vertices. Next, we want to study the various approaches presented in the previous Chapter to handle the positional dependence of ρ . Finally, we want to assess the robustness of the method so we shall use a wide range of hard processes — this is crucial since it is supposed to be used across the whole range of analyses at the LHC — and test different Monte-Carlo generators in order to estimate the uncertainty on our conclusions.

Most of the study presented below comes from a contribution to the proceedings of the 2011 Physics at TeV colliders workshop in Les Houches. However, since the reconstruction of the jet transverse momentum is used in such a wide variety of LHC analysis, we have renewed our studies with more realistic conditions for the run I of the LHC. Compared to the original study, we have moved from 7 to 8 TeV for the centre-of-mass collision energy, extended our range of N_{PU} to 50 (instead of 30) and used $\mu = 25$ as an average pileup multiplicity instead of 10. Note that, while conducted at 8 TeV, the features observed below are not expected to be significantly altered at 13 TeV. We will see some example studies carried at 13 TeV in the last part of this document where we shall discuss more elaborate techniques for Run II and beyond.

3.2.1 Subtraction techniques

Since we want to validate the various aspects of the area–median method, we shall test both the jet-based and the grid-based approaches, as well as several approaches to handling the positional dependence. In order to provide a baseline for comparison and discussion, we shall also test a simpler alternative based on the number of (seen) pileup vertices.

³For the fluctuations, we have actually averaged the square.

In all cases, the subtraction of the jet transverse momentum will be done through the central formula, Eq. (2.24), and define the *subtracted jet* transverse momentum using

$$p_{t,\text{sub}} = p_t - \rho_{\text{est}} A \quad (3.2)$$

where ρ_{est} is the estimated value for the background density per unit area. The various methods we put under scrutiny here only depend in how they estimate ρ_{est} .

Using *seen vertices*. Since experimentally it might be possible — within some level of accuracy that goes beyond the scope of this discussion — to count the number of pileup vertices using charged tracks, one appealing way to estimate the background density in a given event would be to count these vertices and subtract a pre-determined number for each of them:

$$\rho_{\text{est}}^{(N_{\text{PU}})}(y) = f(y) N_{\text{PU,seen}}, \quad (3.3)$$

where we have made explicit the fact that the proportionality constant $f(y)$ can carry a rapidity dependence. $f(y)$ can be studied from minimum bias collisions (see the previous Section) and can take into account the fact that not all the PU vertices will be reconstructed.

Area–median subtraction. This estimates the background using the central formula (2.23):

$$\rho_{\text{est}}^{(\text{global})} = \text{median}_{i \in \text{patches}} \left\{ \frac{p_{t,i}}{A_i} \right\} \quad (3.4)$$

We shall test both jets and grid cells as patches.

Area–median with a local range. This uses a local estimate of the area–median background-density as obtained from Eq. (2.25). A typical example, that we shall study later on, is the case of a *strip range* where only the jets with $|y - y_j| < \Delta$ are included.

Area–median with rescaling. This is the other option to correct for the positional dependence. It uses Eq. (2.27). The rescaling $f(y)$ can be taken directly from Section 3.1.

3.2.2 Testing framework

Embedding. The remainder of this Section is devoted to an in-depth comparison of the subtraction methods proposed above. Our testing framework goes as follows. First, we embed a hard event into a pileup background, either with a fixed number N_{PU} of minimum-bias events or with a Poisson-distributed number of pileup interactions, with μ pileup events on average. Then, we reconstruct and subtract the jets in both the hard and full events.⁴ Jets in the hard (resp. full) event will be referred to as hard jets (resp. full jets) and the kinematic cuts are applied on the hard jets.

⁴One may argue whether or not one should subtract the jets in the hard event. We decided to do so to cover the case where the hard event contains Underlying Event which, as a relatively uniform background, will also be subtracted together with the pileup.

Matching. Testing Eq. (3.5) requires that we match the jets in the full event (optionally subtracted) with the jets in the initial hard event. For each of the full jets, we do that by finding a hard jet such that their common constituents contribute for at least 50% of the transverse momentum of the hard jet. Compared to a matching where we use distances between the jet axes, this has the advantage that a given hard jet is matched to at most one jet in the full event, and it avoids the arbitrariness of choosing a threshold for the matching distance. Also, this matching is independent of the subtraction procedure.

Quality measures. For a matched pair of jets, we shall use the notation p_t^{hard} and p_t^{full} for the transverse momentum of the jet in the hard and full event respectively, and $p_t^{\text{hard,sub}}$ and $p_t^{\text{full,sub}}$ for their subtracted equivalents. We can then compute the shift

$$\Delta p_t = p_t^{\text{full,sub}} - p_t^{\text{hard,sub}}, \quad (3.5)$$

i.e. the difference between the reconstructed-and-subtracted jet with and without pileup. A positive (resp. negative) Δp_t would mean that the PU contamination has been underestimated (resp. overestimated). Though in principle there is some genuine information in the complete Δp_t distribution — e.g. it could be useful to deconvolute the extra smearing brought by the pileup, see e.g. [5] and [39] and the discussion in Chapter 4 below — we shall mainly focus on two simpler quantities: the average shift $\langle \Delta p_t \rangle$ and the dispersion $\sigma_{\Delta p_t}$. While the first one is a direct measure of how well one succeeds at subtracting the pileup contamination on average, the second quantifies the remaining effects on the resolution. One thus wishes to have $\langle \Delta p_t \rangle$ close to 0 and $\sigma_{\Delta p_t}$ as small as possible. Note that these two quantities can be studied as a function of variables like the rapidity and transverse momentum of the jets or the number of pileup interactions. In all cases, a flat behaviour would indicate a robust subtraction method.

Instead of these quality measures, some studies in the literature (see e.g. [10]), prefer to use the correlation coefficient between $p_t^{\text{full,sub}}$ and $p_t^{\text{hard,sub}}$.⁵ We find that quantitative interpretations of correlation coefficients can sometimes be delicate, essentially because they combine the covariance of two observables with the two observables' individual variances. This is further discussed in Appendix C). To the exception of Chapter 11, we will thus use $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ as indicators of the performance of pileup mitigation techniques.

Robustness. The robustness of our conclusions can be checked by varying many parameters:

- one can study various hard processes to check that the PU subtraction is not biased by the hard event. In what follows we shall study dijets with p_t ranging from 20 GeV to 1 TeV, as well as fully hadronic $t\bar{t}$ events as a representative of busier final states.
- The Monte-Carlo used to generate the hard event and PU can be varied. For the hard event, we have used Pythia 6.4.28 [35] with the Perugia 2011 tune [40], Pythia 8.186⁶ [36] with tune 4C [37] and Herwig 6.5.10 [41] with the ATLAS tune and we have switched multiple interactions on (our default) or off. For the minimum bias sample used to generate PU, we have used Pythia 8, tune 4C.⁷

⁵Or between $\mathcal{O}^{\text{full,sub}}$ and $\mathcal{O}^{\text{hard,sub}}$ for any observable \mathcal{O} .

⁶The Z +jet studies in Section 3.2.6 have used version 8.150.

⁷In the context of the original study, carried at 7 TeV in the context of the Les-Houches Physics at TeV Colliders workshop, we had also checked that our conclusions remain unchanged when using Herwig++ [42] (tune LHC-UE7-2). Since we have no good reason to believe that it would be any different in this slightly different case, we have not reiterated that part of the study.

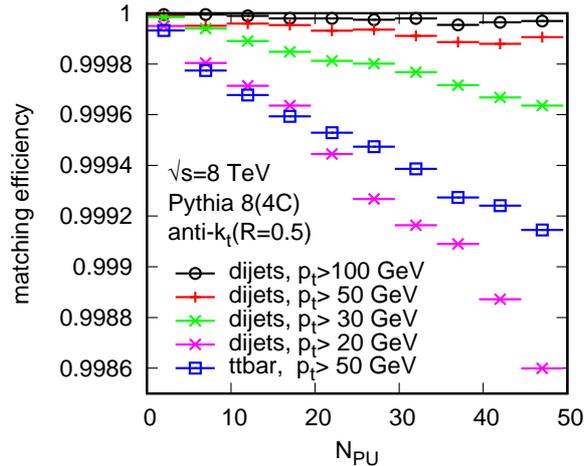


Figure 3.5: fraction of the jets in the hard event, passing the p_t cut, which have been matched to a jet in the full event. Different curves correspond to different processes. Efficiencies are shown as a function of the number of pileup interactions.

Additional details of the analysis. For the sake of completeness, we list here the many other details of how the Δp_t analysis has been conducted. We have considered particles with $|y| \leq 5$ with no p_t cut or detector effect; jets have been reconstructed with the anti- k_t algorithm with $R = 0.5$ keeping jets with $|y| \leq 4$; for area computations, we have used active areas with explicit ghosts with ghosts placed⁸ up to $|y| = 5$. For jet-based background estimations, we have used the k_t algorithm with $R = 0.4$ though other options will be discussed (and the 2 hardest jets in the set have been excluded from the median computation to reduce the bias from the hard event); for grid-based estimations the grid extends up to $|y| = 5$ with cells of edge-size 0.55 (other sizes will be investigated). For estimations using a local range, a strip range of half-width 1.5 has been used (again excluding the two hardest jets in the strip⁹) and for estimations using rapidity rescaling, we have used the results obtained in the previous Section (see e.g. Table. 3.1). Jet reconstruction, area computation and background estimation have all been carried out using FastJet (v3.1) [14, 15]. Whenever we quote an average pileup multiplicity, pileup is generated as a superposition of a Poisson-distributed number of minimum bias events and we will vary the average number of pileup interactions. We shall always assume pp collisions with $\sqrt{s} = 8$ TeV.

For the “seen vertices” pileup subtraction, we need to provide a proper definition of what we call “seen vertices”. In what follows, we shall define that as a minimum bias interaction that has at least 2 charged tracks with $|y| \leq 2.5$ and $p_t \geq 100$ MeV, which corresponds to 69.7% of the events.¹⁰

3.2.3 Generic performance and rapidity dependence

The first thing we want to look at is the matching efficiency between the hard and full jets. This is shown in Fig. 3.5 where we plot, as a function of N_{PU} and for different processes, the fraction of hard jets which have been matched to a jet in the full event. We see that even with busy events and high pileup multiplicities, the matching efficiencies remain very high, always above 99.85% and often much

⁸Note that we have used the ghost placement of FastJet 3 which differs slightly from the one in v2.4.

⁹It would actually be more appropriate to exclude instead the 2 hardest jets in the event.

¹⁰This is a bit optimistic but does not affect in any way our discussion.

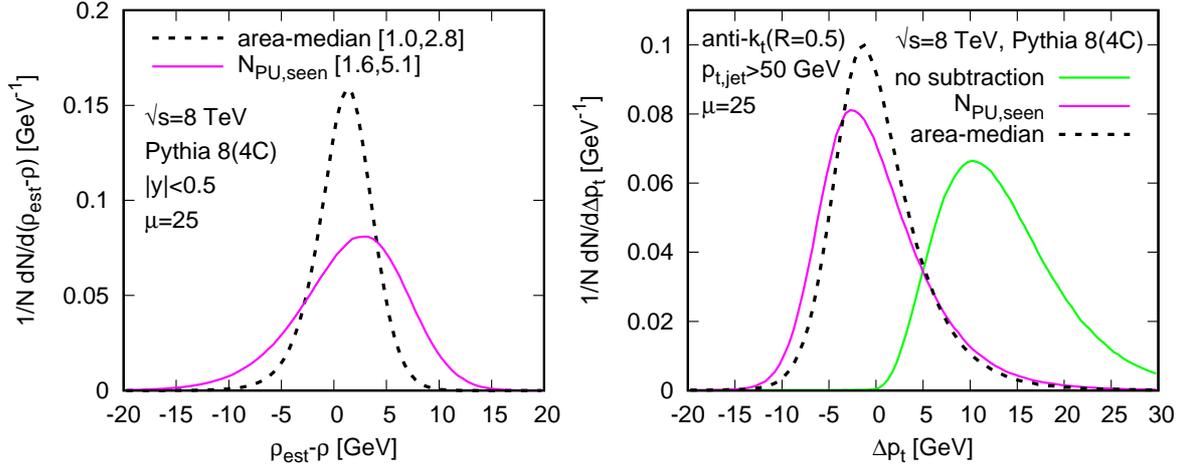


Figure 3.6: Left: distribution of difference between the estimated ρ , ρ_{est} , and its true value in the $|y| < 0.5$ rapidity interval. The calculation is done using an average of 25 Poisson-distributed pileup events for a dijet sample generated with a generator p_t cut of 40 GeV. The dashed (black) line corresponds to the area–median estimation (using a rapidity-rescaled grid estimator) while the solid (magenta) line corresponds to the “seen vertices” estimate. The numbers indicated inside brackets in the legend correspond to the average and dispersion of the corresponding distributions. Right: the corresponding Δp_t distribution for anti- $k_t(R = 0.5)$ jets above 50 GeV. We have also added the distribution obtained when no subtraction are performed (solid, green, line).

higher. As for the p_t dependence, the matching efficiency quickly goes close to 1 when p_t increase, with again large matching efficiencies all the way down to jets of 20 GeV.

Before we focus on the quality measured, $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$, averaged over event samples, let us first have a brief look at global distributions. A few illustrative examples are presented in Fig. 3.6. On the left plot, we directly plot the quality of the background estimation at central rapidity. We see that both the area–median estimate and the “seen vertices” estimate are close to the actual, true, ρ . The width of these distributions show that the area–median tends to give a better estimate than the “seen vertices” approach. On the right plot, we show how this translates in a more practical situation where we subtract pileup contamination from a sample of jets above 50 GeV. In this case, we have also included the unsubtracted distribution which is clearly biased and wider than both subtracted ones. Overall, all these distributions can be well-approximated by Gaussian distributions and we can safely use the averaged quantities $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ to assess the performance and robustness of the various subtractions methods.

Let us therefore carry on with more detailed performance benchmarks, starting with the study of the rapidity dependence of PU subtraction. Fig. 3.7 shows the residual average shift ($\langle \Delta p_t \rangle$) as a function of the rapidity of the hard jet. These results are presented for different hard processes, generated with Pythia 8 (v8.186) and using Poisson-distributed pileup with $\mu = 25$. Robustness w.r.t. that choice will be discussed in the next Section but does not play any significant role for the moment.

The first observation is that the subtraction based on the number of seen PU vertices does a very good job in all 3 cases. Then, global area–median (using jets or grid cells) estimations of ρ , i.e. the (red) square symbols, do a fair job on average but, as expected, fail to correct for the rapidity dependence of the PU contamination. If one now restricts the median to a rapidity strip around the jet, the (blue) triangles, or if one uses rapidity rescaling, the (black) circles, the residual shift is very close to 0, typically a few hundreds of MeV, and flat in rapidity.

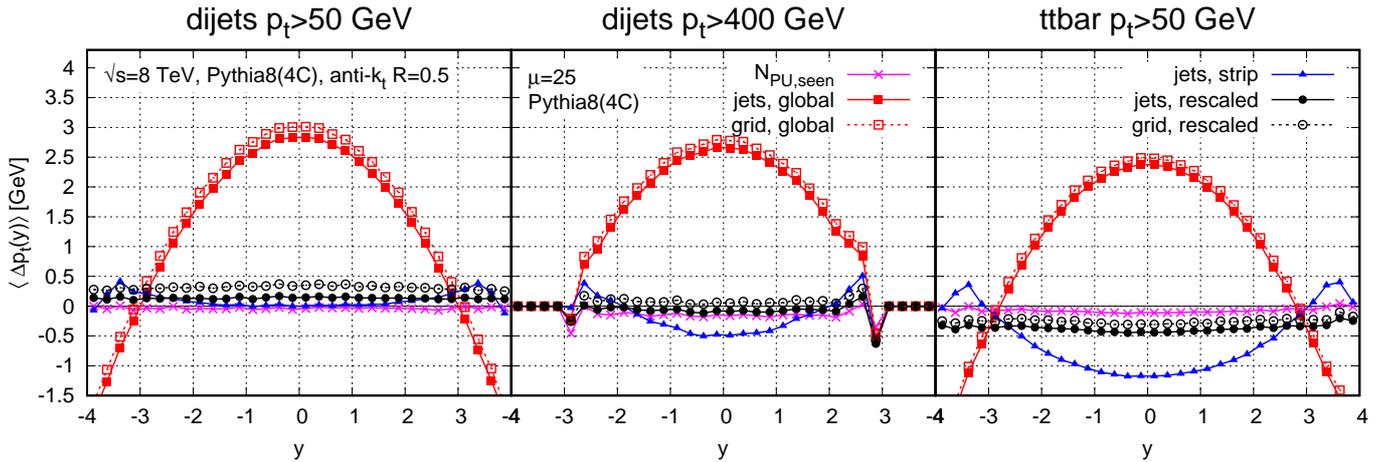


Figure 3.7: Residual average shift as a function of the jet rapidity for all the considered subtraction methods. For the left (resp. centre, right) plot, the hard event sample consists of dijets with $p_t \geq 50$ GeV (resp. dijets with $p_t \geq 400$ GeV, and jets above $p_t \geq 50$ GeV in $t\bar{t}$ events), generated with Pythia 8 (tune 4C) in all cases. The typical PU contamination (for unsubtracted jets) is around 14 GeV at central rapidities.

Note that the strip-range approach seems to have a small residual rapidity dependence and overall offset for high- p_t processes or multi-jet situations. That last point may be due to the fact that smaller ranges tend to be more affected by the presence of the hard jets (see also Section 5.2), an effect which is reinforced for multi-jet events. The fact that the residual shift seems a bit smaller for grid-based estimates will be discussed more extensively in the next Section.

Next, we turn to the dispersion of Δp_t , a direct measure of the impact of PU fluctuations on the p_t resolution of the jets. Our results are plotted in Fig. 3.8 as a function of the rapidity of the hard jet (left panel), the number of PU vertices (central panel) and the transverse momentum of the hard jet (right panel). All subtraction methods have been included as well as the dispersion one would observe if no subtraction were performed.

The results show a clear trend: first, a subtraction based on the number of seen PU vertices brings an improvement compared to not doing any subtraction; second, area–median estimations of ρ give a more significant improvement; and third, all area–median approaches perform similarly well (with a little penalty for global approaches).

The reason why the area–median estimations of ρ outperform the estimation based on the number of seen PU vertices is simply because minimum bias events do not all yield the same energy deposit and this leads to an additional source of fluctuations in the “seen vertices” estimation compared to all area–median ones. This is the main motivation for using an event-by-event determination of ρ based on the energy deposited in the event. This motivation is further strengthened by the fact that additional issues like vertex resolution or out-of-time PU would affect both $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ if estimated simply from the number of seen vertices while area–median approaches are more robust.

Note that even though local ranges and rapidity rescaling do correct for the rapidity dependence of the PU on average, the dispersion still depends on rapidity. The increase with the number of PU vertices is in agreement with the expected $\sqrt{N_{PU}}$ behaviour and the increase with the p_t of the hard process can be associated with *back-reaction*, see the next Section as well as Section 4.4.1. These number can also be compared to the typical detector resolutions which would be ~ 10 GeV for 100 GeV jets and ~ 20 GeV at $p_t = 400$ GeV [43, 44].

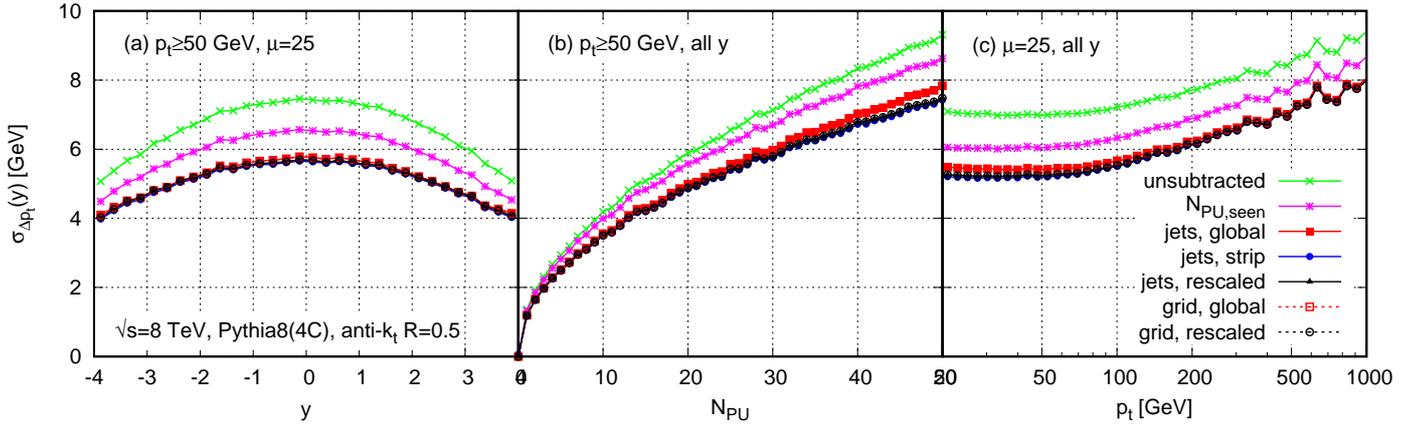


Figure 3.8: Dispersion $\sigma_{\Delta p_t}$. Each curve corresponds to a different subtraction method and the results are presented as a function of different kinematic variables: left, as a function of the rapidity of the hard jet for a sample of jets with $p_t \geq 100$ GeV and assuming an average of 25 PU (Poisson-distributed) events; centre: as a function of the number of PU events for a sample of jets with $p_t \geq 100$; right: as a function of the p_t of the hard jet, assuming again an average of 25 PU events

3.2.4 Digression: back-reaction effects

Remember that back-reaction is the effect of pileup on the clustering of the particles from the hard interaction themselves. This effect is independent of the subtraction method. In this short Section, we discuss a few of the quantitative effects of back-reaction on the jet transverse momentum and compare them to the average shift and dispersion quoted in the previous Section.

The determination of the back reaction works as follows: for each pair of matching hard and full jets, we look at the particles from the full jet that come from the hard interaction, say $p_t^{\text{hard-in-full}}$.¹¹ The Back-reaction is then computed for each jet as

$$\Delta p_t^{\text{back-reaction}} = p_t^{\text{hard-in-full}} - p_t^{\text{hard}}. \quad (3.6)$$

We can then, as before, look at the average shift and dispersion due to back-reaction. These are shown in Figure 3.9. We see that, as expected, the effect of back-reaction on the average shift is very small (less than 200 MeV). We also see that it comes with a non-negligible dispersion effect, in particular for busy events ($t\bar{t}$) or high- p_t jets.

To investigate this a little further, it is interesting to look at the contribution of back-reaction to the overall dispersion observed in Fig. 3.8. This is shown in Fig. 3.10 for dijet events as a function of the jet p_t and for $t\bar{t}$ events as a function of the number of pileup interactions. In both cases, the total dispersion (the black curve) is as obtained with the rescaled grid subtraction in Fig. 3.8. The back reaction is as shown in Fig. 3.9. The difference, computed as $[\sigma_{\Delta p_t}^2 - (\sigma_{\Delta p_t}^{\text{BR}})^2]^{\frac{1}{2}}$, can be directly associated with background fluctuations and misestimation of ρ . One sees that the background-fluctuation component dominates over the back-reaction component although the latter is certainly non-negligible. In particular, it is the back-reaction component which is mainly responsible for the rise of the dispersion with the jet p_t , with a mostly- p_t -independent background-fluctuation contribution.

This is the case because the back-reaction dispersion is dominated by rare events in which two similarly hard subjets are separated by a distance close to R (specifically by $R \pm \epsilon$ with $\epsilon \ll 1$).

¹¹This is doable in our Monte Carlo setup where we have access to the true origin of each particle in the full event.

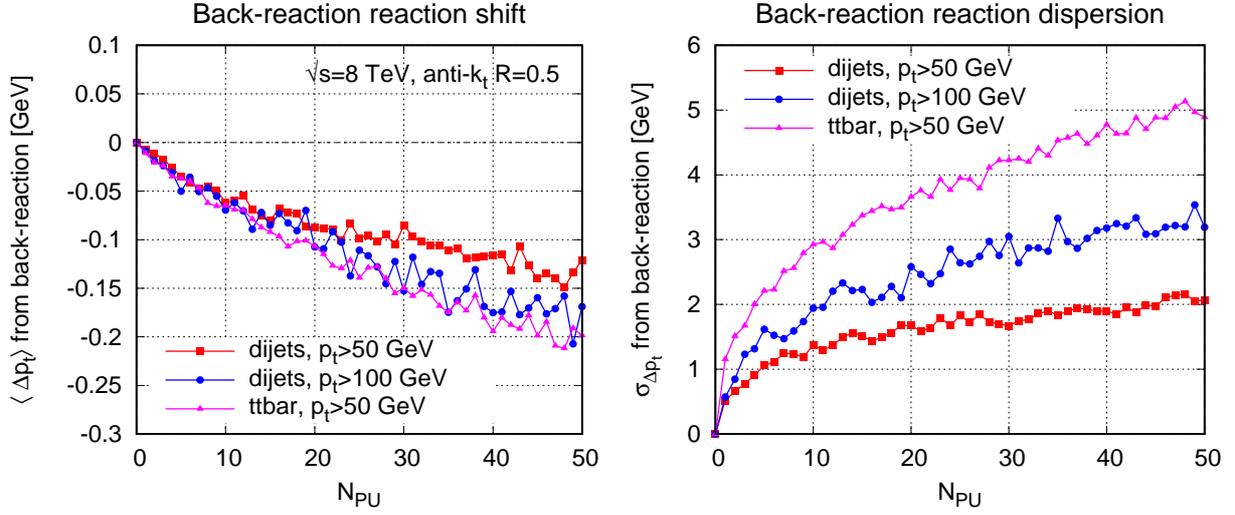


Figure 3.9: Contribution of the back-reaction to the pileup average shift and dispersion.

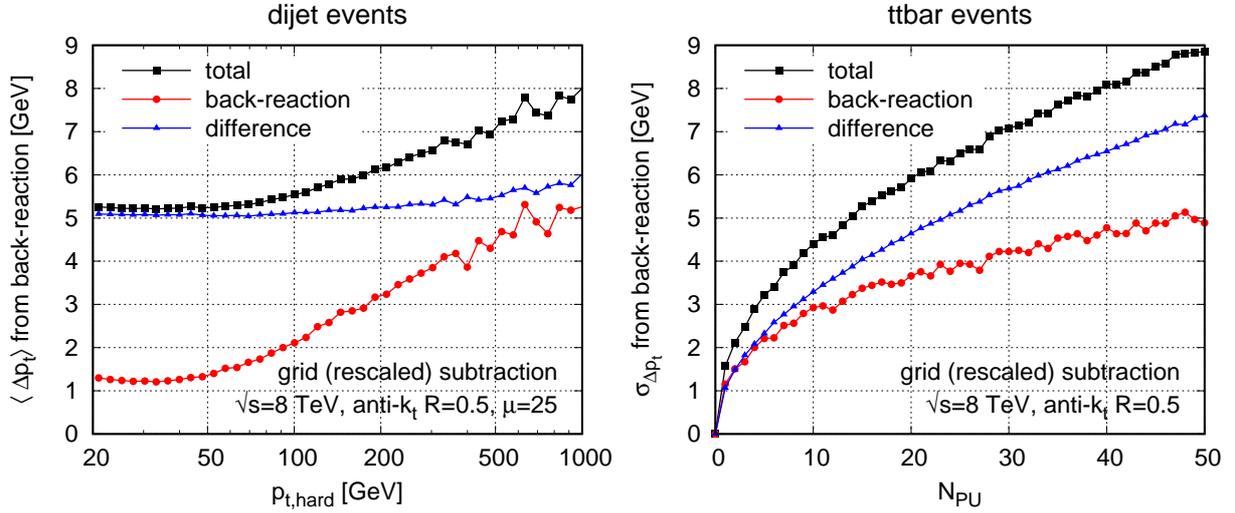


Figure 3.10: Dispersion $\sigma_{\Delta p_t}$ broken into two contributions: the back-reaction contribution and the remaining contribution due to residual pileup fluctuations. The left plot shows dijet events as a function of the jet p_t , the right plot focuses on $t\bar{t}$ events as a function of the number of pileup interactions.

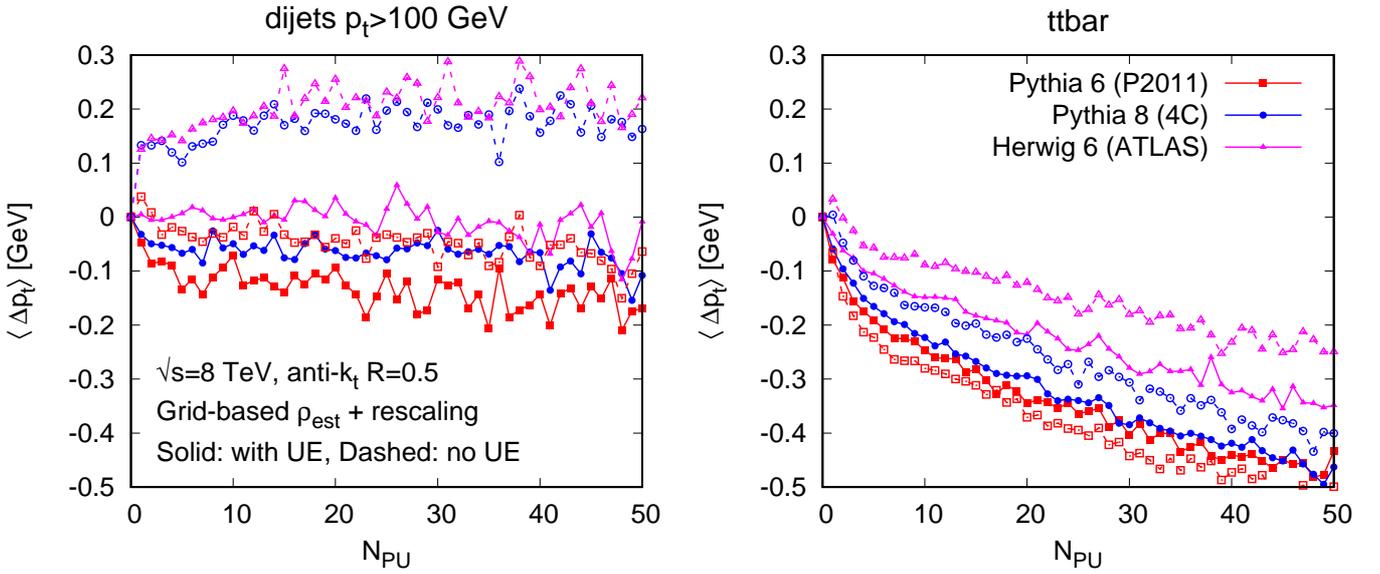


Figure 3.11: Dependence of the average p_t shift as a function of the number of PU vertices for various Monte-Carlo generators. For the left plot, the hard sample is made of dijets with $p_t \geq 100$ GeV while for the right plot, we have used a hadronic $t\bar{t}$ sample. For each generator, we have considered both the case with the Underlying Event switched on (filled symbols) and off (open symbols). All results have been obtained using a grid-based median estimation of ρ using rapidity rescaling.

In such a configuration, the background's contribution to the two subjects can affect whether they recombine and so lead to a large, $\mathcal{O}(p_t)$, change to the jet's momentum. In the limit of a uniform background, $\sigma/\rho \ll 1$, this can be shown (see Section 5.1.3 for details) to occur with a probability of order $\alpha_s \rho R^2 / p_t$. Thus the contribution to the average shift $\langle \Delta p_t \rangle$ is proportional to $\alpha_s \rho R^2$, while the contribution to $\langle \Delta p_t^2 \rangle$ goes as $\alpha_s \rho R^2 p_t$, and so leads to a dispersion that should grow asymptotically as $\sqrt{p_t}$ (fig. 3.10 is, however, probably not yet in the asymptotic regime).

It is worth keeping in mind that even though rare but large back-reaction dominates the overall dispersion, it will probably not be the main contributor in distorting the reconstructed jet spectrum. Such distortions come from upwards Δp_t fluctuations, whereas large back-reaction tends to be dominated by downwards fluctuations. The reason is simple: in order to have an upwards fluctuation from back-reaction, there must be extra p_t near the jet in the original pp event. This implies the presence of a harder underlying $2 \rightarrow 2$ scattering than would be deduced from the jet p_t , with a corresponding significant price to pay in terms of more suppressed matrix elements and PDFs. We shall show a similar behaviour for heavy-ion collisions and discuss analytic results in more details in Section 5.1.3.

3.2.5 Robustness and Monte-Carlo dependence

The last series of results we want to present addresses the stability and robustness of the area–median estimation of the PU density per unit area.

To do that, the first thing we shall discuss is the Monte-Carlo dependence of our results. In Fig. 3.11 we compare the different Monte-Carlo predictions for the $\langle \Delta p_t \rangle$ dependence on the number of PU vertices in the case of a grid-based median estimate of ρ with rapidity rescaling. For each of the

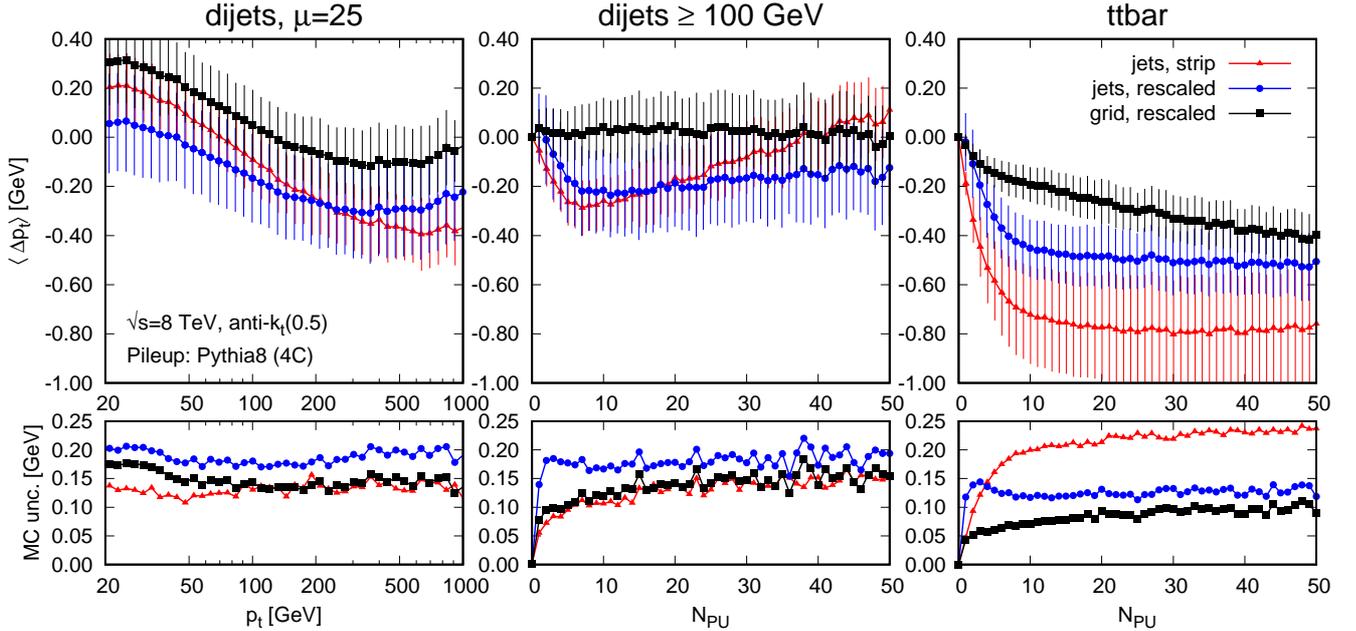


Figure 3.12: Average shift, $\langle \Delta p_t \rangle$, plotted as a function of the jet p_t for an average of 25 PU events (left), or as a function of the number of PU vertices for dijets with $p_t \geq 100$ GeV (centre) and for $t\bar{t}$ events (right). In all cases, we compare 3 methods: the rapidity-strip range, (red) triangles, the jet-based approach with y rescaling, (blue) circles, and the grid-based approach with y rescaling, (black) squares. Each curve is the result of averaging over the various Monte-Carlo generator options. Their dispersion is represented both as error bars on the top row and directly on the bottom row.

three considered Monte-Carlos, we have repeated the analysis with and without Underlying Event (UE) in the hard event. The first observation is that all the results span a range of 300-400 MeV in Δp_t and have a similar dependence on the number of PU vertices. The dependence on N_{PU} is flat for dijet events but shows a small decrease for busier events. The 300-400 MeV shift splits into a 100-200 MeV effect when changing the generator, which is likely due to the small but non-zero effect of the hard event on the median computation, and a 100-200 MeV effect coming from the switching on/off of the UE.

This question of subtracting the UE deserves a discussion: since the UE is also a soft background which is relatively uniform, it contributes to the median estimate and, therefore, one expects the UE, or at least a part of it, to be subtracted together with the PU. Precisely for that reason, when we compute Δp_t , our subtraction procedure is not applied only on the “full jet” (hard jet+PU) but also on the hard jet, see Eq.(3.5). The 100-200 MeV negative shift observed in Fig. 3.11 thus means that, when switching on the UE, one subtracts a bit more of the UE in the full event (with PU) than in the hard event alone (without PU). This could be due to the fact (see Section 5.2 for details) that for sparse events, as is typically the case with UE but no PU, the median tends to slightly underestimate the “real” ρ , e.g. if half of the event is empty, the median estimate would be 0. This is in agreement with the fact that for $t\bar{t}$ events, where the hard event is busier, switching on the UE tends to have a smaller effect. Note finally that as far as the size of the effect is concerned, this 100-200 MeV shift has to be compared with the ~ 1 GeV contamination of the UE in the hard jets.

Finally, we wish to compare the robustness of our various subtraction methods for various processes i.e. hard events and PU conditions. In order to avoid multiplying the number of plots, we shall treat the Monte-Carlo (including the switching on/off of the UE) as an error estimate. That is, an average

measure and an uncertainty will be extracted by taking the average and dispersion of the 6 Monte-Carlo setups. The results of this combination are presented on Fig. 3.12 for various situations and subtraction methods. For example, the 6 curves from the left plot of Fig. 3.11 have been combined into the (black) squares of the central panel in Fig. 3.12.

Two pieces of information can be extracted from these results. First of all, for dijets, the average bias of PU subtraction remains very small and, to a large extent, flat as a function of the p_t of the jets and the number of PU vertices. When moving to multi-jet situations, we observe an additional residual shift in the 100-300 MeV range, extending to ~ 500 MeV for the rapidity-strip-range method. This slightly increased sensitivity of the rapidity-strip-range method also depends on the Monte-Carlo. While in all other cases, our estimates vary by ~ 100 MeV when changing the details of the generator, for multi-jet events and the rapidity-strip-range approach this is increased to ~ 200 MeV.

Overall, the quality of the subtraction is globally very good. Methods involving rapidity rescaling tends to perform a bit better than the estimate using a rapidity strip range, mainly a consequence of the latter's greater sensitivity to multi-jet events. In comparing grid-based to jet-based estimations of ρ , one sees that the former gives slightly better results, though the differences remain small.

Since the grid-based approach is considerably faster than the jet-based one, as it does not require an additional clustering of the event¹², the estimation of ρ using a grid-based median with rapidity rescaling comes out as a very good default for PU subtraction. One should however keep in mind local-range approaches for the case where the rapidity rescaling function cannot easily be obtained.

3.2.6 PU v. UE subtraction: an analysis on Z +jet events

To give further insight on the question of what fraction of the Underlying Event gets subtracted together with the pileup, we have performed an additional study of Z +jet events. We look at events where the Z boson decays into a pair of muons. We have considered 5 different situations: events without PU or UE, events with UE but no PU with and without subtraction, and events with both UE and PU again with and without subtraction. Except for the study of events without UE, this analysis could also be carried out directly on data.

Practically, we impose that both muons have a transverse momentum of at least 20 GeV and have $|y| \leq 2.5$, and we require that their reconstructed invariant mass is within 10 GeV of the nominal Z mass. As previously, jets are reconstructed using the anti- k_t jet algorithm with $R = 0.5$ and the pileup subtraction is performed using the grid-based-median approach with rapidity rescaling and a grid size of 0.55. All events have been generated at $\sqrt{s} = 7$ TeV with Pythia 8 (tune 4C) and we have assumed an average PU multiplicity of 20 events.¹³ Note that for this study, no matching between the hard and full events are required. In practice, we only focus on the hardest jet in the event, with the cut applied after the area-median subtraction for the subtracted cases.

In Fig. 3.13, we have plotted the ratio $p_{t,\text{jet}}/p_{t,Z} - 1$, with $p_{t,\text{jet}}$ the transverse momentum of the leading jet, for the various situations under considerations. Compared to the ideal situation with no PU and no UE, the (black) triangles, one clearly sees the expected effect of switching on the UE, the empty (green) circles, or adding PU, the empty (red) squares: the UE and PU add to the jet ~ 1.2 and 13 GeV respectively.

¹²Note that the clustering of the main event still needs to include the computation of jet areas since they are needed in Eq. (3.2).

¹³Compared to the in-depth validation of the subtraction quality for the jet transverse momentum presented earlier in this Section, the simulations in this Section have not been updated to 8 TeV and correspond to the original results from the Les Houches studies.

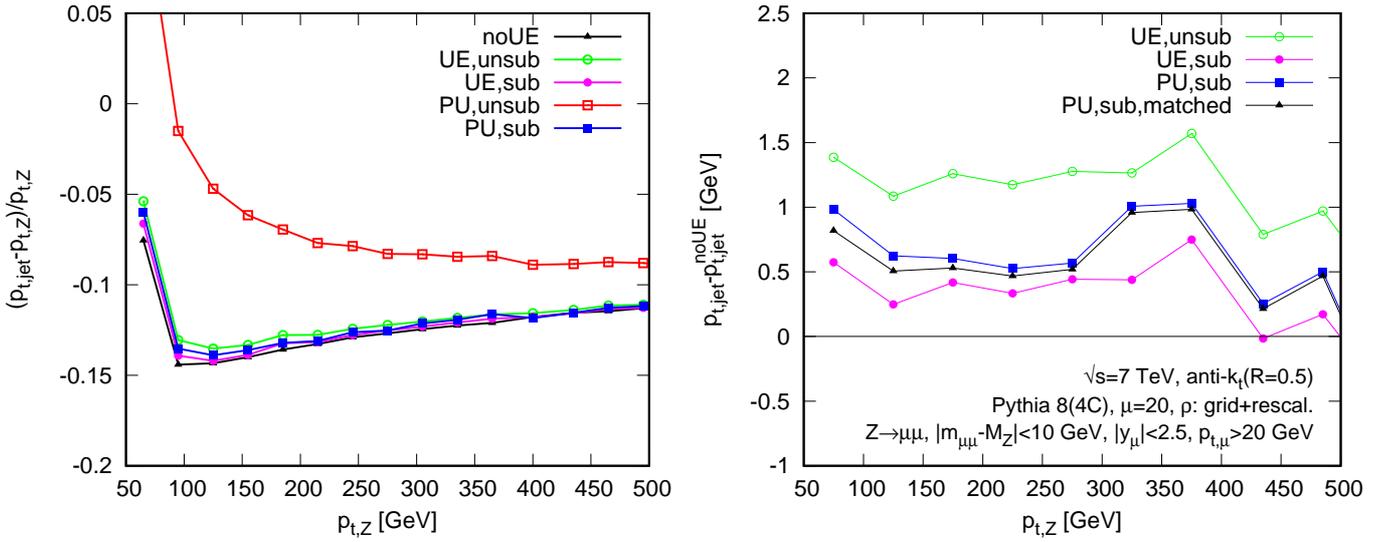


Figure 3.13: Left: relative difference between the reconstructed jet and the reconstructed Z boson transverse momenta. Right: at a given p_t of the reconstructed Z boson, difference between the reconstructed p_t of the jet and the ideal p_t with no UE or PU, i.e. p_t shift w.r.t. the “noUE” curve, the (black) triangles, on the left panel. See the text for the details of the analysis.

We now turn to the cases where the soft background is subtracted, i.e. the filled (blue) squares and (magenta) circles, for the cases with and without PU respectively. There are two main observations:

- with or without PU, the UE is never fully subtracted: from the original 1-1.5 GeV shift, we do subtract about 800 MeV to be left with a ~ 500 MeV effect from the UE. That effect decreases when going to large p_t .
- in the presence of PU, the subtraction produces results very close to the corresponding results without PU and where only the UE is subtracted. This nearly perfect agreement at large $p_{t,\text{jet}}$ slightly degrades into an additional offset of a few hundreds of MeV when going to smaller scales. This comes about for the following reason: the non-zero p_t resolution induced by pileup (even after subtraction) means that in events in which the two hardest jets have similar p_t , the one that is hardest in the event with pileup may not correspond to the one that is hardest in the event without pileup. This introduces a positive bias on the hardest jet p_t (a similar bias would be present in real data even without pileup, simply due to detector resolution). The “matched” curve in Fig. 3.13 (right) shows that if, in a given hard event supplemented with pileup, we explicitly use the jet that is closest to the hardest jet in that same event without pileup, then the offset disappears, confirming its origin as due to resolution-related jet mismatching.

3.2.7 Summary and discussion

In this Section, we have validated the area–median approach as a powerful and robust method for pileup mitigation at the LHC. We have compared it to an alternative approach which applies a correction based on the observed number of pileup vertices. We have also studied different methods to correct for the positional dependence of the pileup energy deposit.

The subtraction efficiency has been studied by embedding hard events into PU backgrounds and investigate how jet reconstruction was affected by measuring the remaining p_t shift after subtrac-

tion ($\langle\Delta p_t\rangle$) as well as the impact on resolution ($\sigma_{\Delta p_t}$). We have investigated the robustness of the area–median approach by varying several kinematic parameters and studying different Monte-Carlo generators and setups.

The first important message is that, though all methods give a very good overall subtraction ($\langle\Delta p_t\rangle \approx 0$), event-by-event methods should be preferred because their smaller PU impact on the p_t resolution (see Fig. 3.8). This is mostly because the “seen vertices” method has an additional smearing coming from the fluctuations between different minimum bias collisions as expected from our discussion in Section 2.1. This does not happen in event-by-event methods that are only affected by point-to-point fluctuations in an event. Note also that event-by-event methods are very likely more robust than methods based on identifying secondary vertices when effects like vertex identification and out-of-time PU are taken into account.

The next observation is that the two techniques to correct for the positional dependence of the pileup energy deposit work (see Fig. 3.7). The median approach using a local range (with jets as patches) or rapidity rescaling (using jets or grid cells as patches) all give an average offset below 500 MeV, independently of the rapidity of the jet, its p_t or the number of PU vertices (see Fig. 3.12) and are thus very suitable methods for PU subtraction at the LHC. Pushing the analysis a bit further one may argue that the local-range method has a slightly larger offset when applied to situations with large jet multiplicity like $t\bar{t}$ events (the right panel of Fig. 3.12) though this argument seems to depend on the Monte-Carlo used to generate the hard-event sample. Also, since it avoids clustering the event a second time, the grid-based method has the advantage of being faster than the jet-based approach.

Our observations are to a large extent independent of the details of the simulations, with the average bias remaining small when we vary the kinematic properties of the jet (its transverse momentum and rapidity), the pileup activity N_{PU} , or the Monte-Carlo generator used for the simulation.

At the end of the day, this study is the base for our recommendations in Section 2.8: *we can recommend the area–median subtraction method with rapidity rescaling and using grid cells as patches* as a powerful default PU subtraction method at the LHC. One should nonetheless keep in mind that the use of jets instead of grid cells also does a very good job and that local-ranges can be a good alternative to rapidity rescaling if the rescaling function cannot be computed. Also, though we have not discussed that into details, a grid cell size of 0.55 is a good default as is the use of k_t jets with $R = 0.4$.

To conclude, let us make a few general remarks. First, our suggested method involves relatively few assumptions, which helps ensure its robustness. Effects like in-time v. out-of-time PU or detector response should not have a big impact. Many of the studies performed here can be repeated with “real data” rather than Monte-Carlo simulations. The best example is certainly the Z +jet study of Section 3.2.6 which could be done using LHC data samples with different PU activity. Also, the rapidity rescaling function can likely be obtained from minimum bias collision data and the embedding of a hard event into pure PU events could help quantifying the remaining $\mathcal{O}(100 \text{ MeV})$ bias.

Finally, we could also investigate hybrid techniques where one would discard the charged tracks that do not point to the primary vertex (CHS events) and apply the subtraction technique described here to the rest of the event. This would have the advantage to further reduce fluctuation effects (roughly by a factor $\sim \sqrt{1/(1-f_{\text{chg}})} \approx 1.6$, where $f_{\text{chg}} \approx 0.62$ is the fraction of charged particles in an event). We will discuss this option at length in Chapter 11, together with other possible approaches using information from charged tracks.

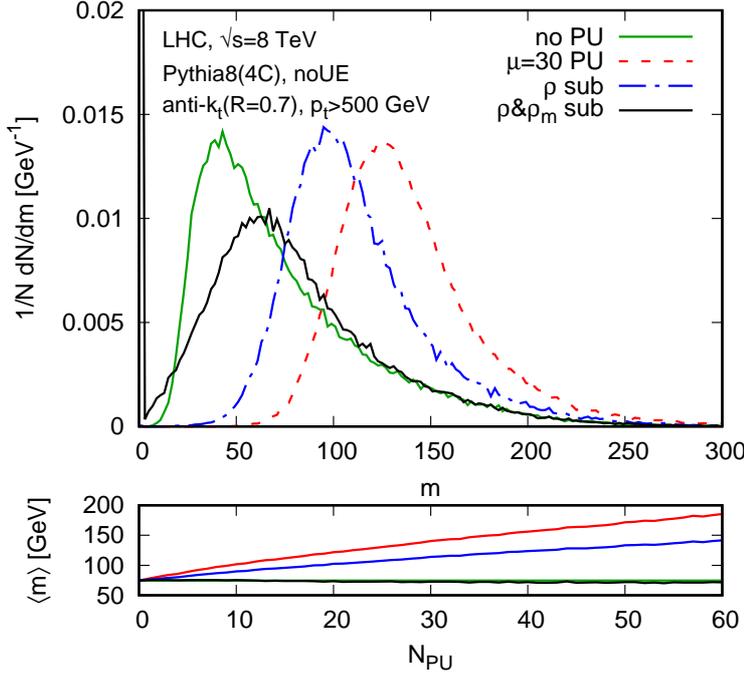


Figure 3.14: Area–median subtraction for the jet mass. Top panel: mass distribution plotted for the hard events with no pileup (solid green line), the full events including (unsubtracted) pileup with $\mu = 30$ (dashed red line), the subtracted full events including only the ρ term in the correction (dash-dotted blue line) and the subtracted events including both the ρ and ρ_m terms (solid black line). Bottom panel: the corresponding average mass shift as a function of the number of pileup vertices.

3.3 Jet mass and jet shapes

Let us first briefly discuss the subtraction of the jet mass and the necessity to include the ρ_m term in (2.30) when particles are massive. This is illustrated on Fig. 3.14 where we have measured the jet-mass spectrum for anti- $k_t(R = 0.7)$ jets with $p_t > 500$ GeV, together with the shift of the average mass as a function of N_{PU} on the lower panel. Compared to the distribution observed without pileup (the solid green line), we see that adding pileup (dashed red line) significantly shifts the distribution to larger masses. If we apply the area–median subtraction Eq. (2.24), the dash-dotted line in Fig. 3.14, we still observe a large mass shift. Only once we also include the ρ_m correction do we recover the right average mass, explicitly proving the importance of that correction with massive particles.

It is also interesting to notice that the mass peak is not quite well reproduced at small jet masses. In that region, the jet mass is rather sensitive to pileup fluctuations resulting in a smeared peak.

We now turn to the validation of the pileup subtraction method for jet shapes introduced in Section 2.5. To investigate the performance of our correction procedure, we consider a number of jet shapes:

- Angularities [45, 46], adapted to hadron-collider jets as $\theta^{(\beta)} = \sum_i p_{ti} \Delta R_{i,\text{jet}}^\beta / \sum_i p_{ti}$, for $\beta = 0.5, 1, 2, 3$; $\theta^{(1)}$, the “girth”, “width” or “broadening” of the jet, has been found to be particularly useful for quark/gluon discrimination [47, 48].
- Energy-energy-correlation (EEC) moments, advocated for their resummation simplicity in [49], $E^{(\beta)} = \sum_{i,j} p_{ti} p_{tj} \Delta R_{i,j}^\beta / (\sum_i p_{ti})^2$, using the same set of β values. EEC-related variables have been studied recently also in Refs. [50, 51, 52].
- “Subjettiness” ratios, designed for characterising multi-pronged jets [53, 54, 55]: one defines the subjettiness $\tau_N^{(\text{axes},\beta)} = \sum_i p_{ti} \min(\Delta R_{i1}, \dots, \Delta R_{iN})^\beta / \sum_i p_{ti}$, where ΔR_{ia} is the distance between particle i and axis a , where a runs from 1 to N . One typically considers ratios such as $\tau_{21} \equiv \tau_2/\tau_1$ and $\tau_{32} \equiv \tau_3/\tau_2$ (the latter used e.g. in a recent search for R-parity violating gluino decays [56]);

we consider $\beta = 1$ and $\beta = 2$, as well as two choices for determining the axes: “kt”, which exploits the k_t algorithm to decluster the jet to N subjets and then uses their axes; and “1kt”, which adjusts the “kt” axes so as to obtain a single-pass approximate minimisation of τ_N [55].

- A longitudinally invariant version of the planar flow [57, 46], involving a 2×2 matrix $M_{\alpha\beta} = \sum_i p_{ti}(\alpha_i - \alpha_{\text{jet}})(\beta_i - \beta_{\text{jet}})$, where α and β correspond either to the rapidity y or azimuth ϕ ; the planar flow is then given by $\text{Pf} = 4\lambda_1\lambda_2/(\lambda_1 + \lambda_2)^2$, where $\lambda_{1,2}$ are the two eigenvalues of the matrix.

One should be aware that observables constructed from ratios of shapes, such as $\tau_{n,n-1}$ and planar flow, are not infrared and collinear (IRC) safe for generic jets. In particular Pf and τ_{21} are IRC safe only when applied to jets with a structure of at least two hard prongs, usually guaranteed by requiring the jets to have significant mass; τ_{32} requires a hard three-pronged structure,¹⁴ a condition not imposed in previous work, and that we will apply here through a cut on τ_{21} .

For the angularities and EEC moments we have verified that the first two numerically-obtained derivatives agree with analytical calculations in the case of a jet consisting of a single hard particle. For variables like τ_N that involve a partition of a jet, one subtlety is that the partitioning can change as the ghost momenta are varied to evaluate the numerical derivative. The resulting discontinuities (or non-smoothness) in the observable’s value would then result in nonsensical estimates of the derivatives. We find no such issue in our numerical method to evaluate the derivatives, but were it to arise, one could choose to force a fixed partitioning.

To test the method in simulated events with pileup, we use Pythia 8.165, tune 4C. We consider 3 hard event samples: dijet, WW and $t\bar{t}$ production, with hadronic W decays, all with underlying event (UE) turned off (were it turned on, the subtraction procedure would remove it too). We use anti- k_t jets with $R = 0.7$, taking only those with $p_t > 500$ GeV (before addition of pileup). All jet-finding is performed with FastJet 3.0. The determination of ρ and ρ_m for each event follows the area–median approach using Eqs. (2.23) and (2.29) to estimate the background properties and (2.30) to correct the jet 4-momentum. Here, the estimation of ρ and ρ_m is done using jets¹⁵ obtained with the k_t algorithm with $R = 0.4$ and includes rapidity rescaling.

We have 17 observables and 3 event samples. Fig. 3.15 gives a representative subset of the resulting 51 distributions, showing in each case both the distribution (top panel) and average shift (bottom panel, plotted as function of N_{PU}) for the shape without pileup (solid green line), the result with pileup (dashed line) and the impact of subtracting first and second derivatives (dotted and solid black lines respectively). The plots for the distributions have been generated using a Poisson distribution of pileup events with an average of $\mu = 30$ events.

For nearly all the jet shapes, the pileup has a substantial impact, shifting the average values by up to 50 – 100% (as compared to a 5 – 10% effect on the jet p_t). The subtraction performs adequately: the averaged subtracted results for the shapes usually return very close to their original values, with the second derivative playing a small but sometimes relevant role. For the distributions, tails of the distributions are generally well recovered; however intrajet pileup fluctuations cause sharp peaks to be somewhat broadened. These cannot be corrected for without applying some form of noise reduction,

¹⁴Consider a jet consisting instead of just two hard particles with $p_t = 1000$ GeV, with $\phi = 0, 0.5$ and two further soft particles with $p_t = \epsilon$, at $\phi = 0.05, 0.1$, all particles having $y = 0$. It is straightforward to see that τ_{32} is finite and independent of ϵ for $\epsilon \rightarrow 0$, which results in an infinite leading-order perturbative distribution for τ_{32} . Note however that some level of calculability can be reached for these, so-called Sudakov-safe, observables [58] (see also [52]).

¹⁵These results were obtained with version 3.0 of FastJet for which ρ_m estimation was only available for jet-based methods. Since FastJet v3.1, it is now also possible to use grid-based methods for ρ_m .

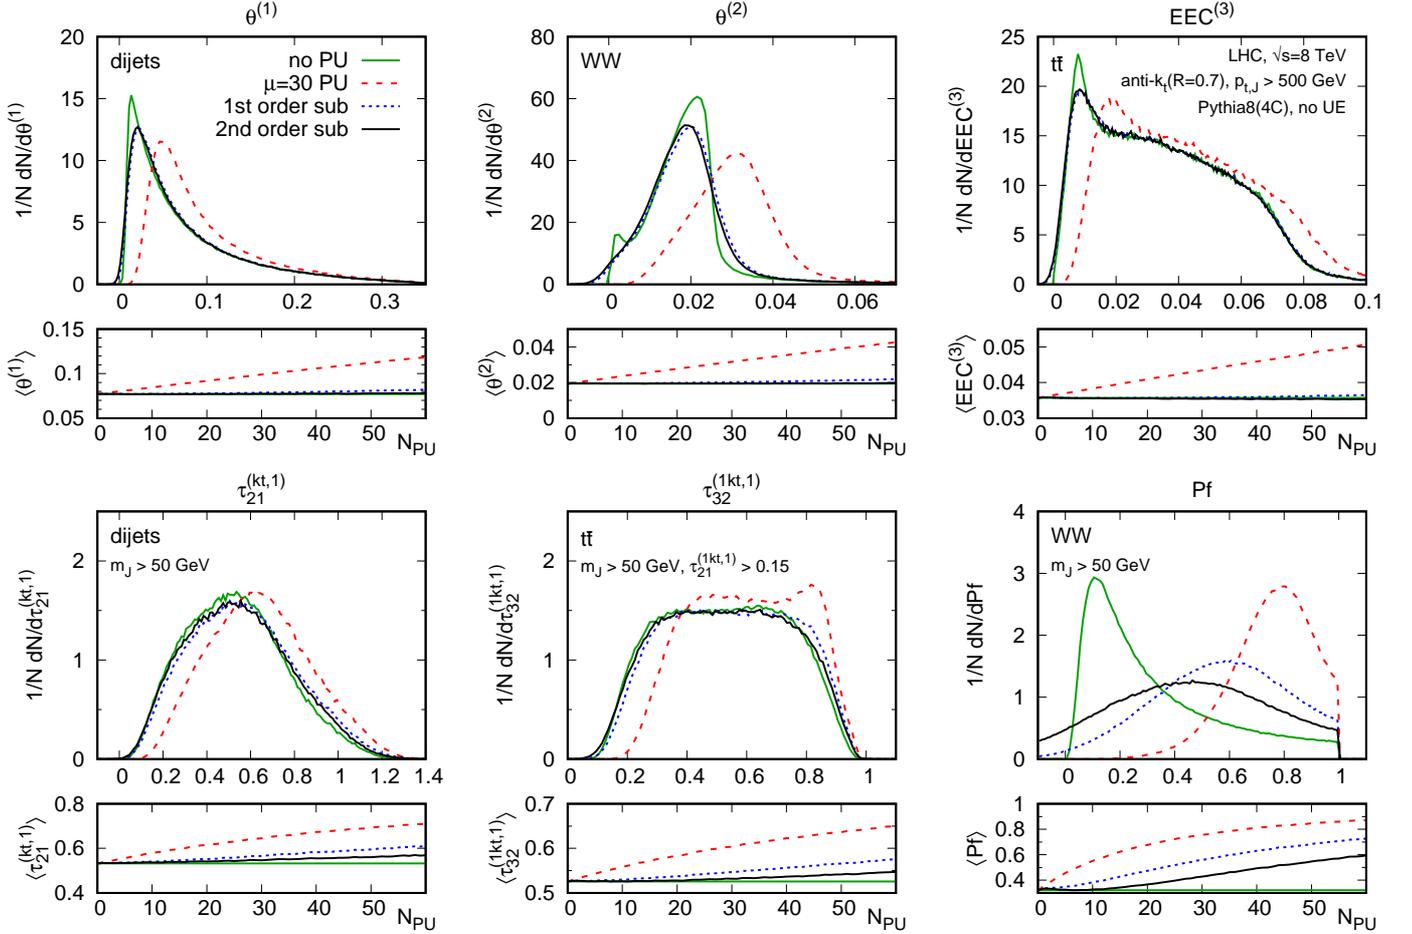


Figure 3.15: Impact of pileup and subtraction on various jet-shape distributions and their averages, in dijet, WW and $t\bar{t}$ production processes. The distributions are shown for Poisson distributed pileup (with an average of 30 pileup events) and the averages are shown as a function of the number of pileup events, N_{PU} . The shapes are calculated for jets with $p_t > 500$ GeV (the cut is applied before adding pileup, as are the cuts on the jet mass m_J and subjettness ratio τ_{21} where relevant).

which would however also tend to introduce a bias.¹⁶ Of the 51 combinations of observable and process that we examined, most were of similar quality to those illustrated in Fig. 3.15, with the broadening of narrow peaks found to be more extreme for larger β values. The one case where the subtraction procedure failed was the planar flow for (hadronic) WW events: here the impact of pileup is dramatic, transforming a peak near the lower boundary of the shape’s range, $\text{Pf} = 0$, into a peak near its upper boundary, $\text{Pf} = 1$ (bottom-right plot of Fig. 3.15). This is an example where one cannot view the pileup as simply “perturbing” the jet shape, in part because of intrinsic large non-linearities in the shape’s behaviour; with our particular set of p_t cuts and jet definition, the use of the small- ρ expansion of Eq. (2.33) fails to adequately correct the planar flow for more than about 15 pileup events.

Next, we want to show the behaviour of the area–median subtraction for jet shapes in more practical applications.

We first consider the use of the subtraction approach in the context of quark/gluon discrimination. In a study of a large number of shapes, Ref. [47] found the jet girth or broadening, $\theta^{(1)}$, to be the

¹⁶See e.g. Chapter 12.

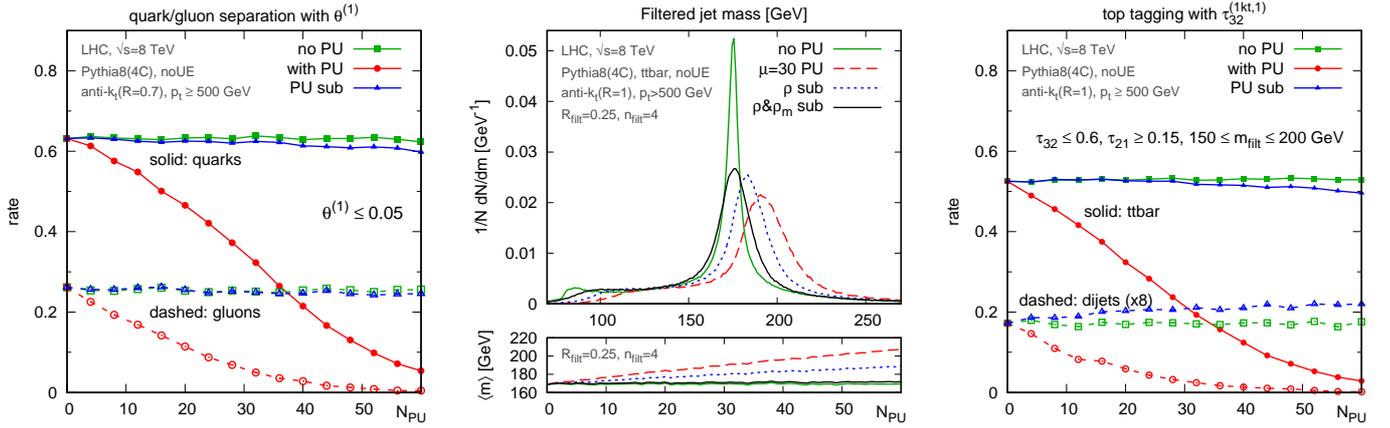


Figure 3.16: Left: rate for tagging quark and gluon jets using a fixed cut on the jet width, shown as a function of the number of pileup vertices. Middle: filtered jet-mass distribution for fat jets in $t\bar{t}$ events, showing the impact of the ρ and ρ_m components of the subtraction. Right: tagging rate of an N -subjettiness top tagger for $t\bar{t}$ signal and dijet background as a function of the number of pileup vertices. All cuts are applied after addition (and possible subtraction) of pileup. Subtraction acts on τ_1, τ_2 and τ_3 individually. See text for further details.

most effective single infrared and collinear safe quark/gluon discriminator. Fig. 3.16 (left) shows the fraction of quark and gluon-induced jets that pass a fixed cut on $\theta^{(1)} \leq 0.05$ as a function of the level of pileup — pileup radically changes the impact of the cut, while after subtraction the q/g discrimination returns to its original behaviour.

Our last test involves top tagging, which we illustrate on $R = 1$, anti- k_t jets using cuts on the “filtered” jet mass and on the τ_{32} subjettiness ratio. The filtering selects the 4 hardest $R_{filt} = 0.25$, Cambridge/Aachen subjects after pileup subtraction. The distribution of filtered jet mass is shown in Fig. 3.16 (middle), illustrating that the subtraction mostly recovers the original distribution and that ρ_m is as important as ρ (specific treatments of hadron masses, e.g. setting them to zero, may limit the impact of ρ_m in an experimental context). The tagger itself consists of cuts on $\tau_{32} < 0.6, \tau_{21} \geq 0.15$ and a requirement that the filtered [59] jet mass be between 150 and 200 GeV. The rightmost plot of Fig. 3.16 shows the final tagging efficiencies for hadronic top quarks and for generic dijets as a function of the number of pileup events. Pileup has a huge impact on the tagging, but most of the original performance is restored after subtraction.

Chapter 4

Applications to heavy-ion collisions

In this Chapter, we shall review the performance of the area–median subtraction method applied to heavy-ion collisions.

Prior to any physics discussion, we should mention that the Monte Carlo studies presented in this Chapter were made before any heavy-ion collision actually happened at the LHC, at a time where jet reconstruction in such a busy environment was still in its early days.

Our studies mostly correspond to $\sqrt{s} = 5.5$ TeV, except for the jet fragmentation studies presented in Section 4.7 where we have used $\sqrt{s} = 2.76$ TeV. Since no LHC data were available at the time of our study, the Monte Carlo generators we used to simulate $PbPb$ collision had to be extrapolated from RHIC energies. In practice, it seems that the tunes we have adopted are doing a decent job at capturing what is now known from LHC data. This is at the very least good enough for the purpose of the indicative studies in this Chapter and we have not deemed necessary to update our results to the lower centre-of-mass energy or to more modern generators and tunes. A discussion comparing the tune we used for $\sqrt{s} = 2.76$ TeV collisions and first LHC measurements can be found in Ref. [60].

4.1 Challenges of jet reconstruction in heavy-ion collisions

Jets in HI collisions are produced in an environment that is far from conducive to their detection and accurate measurement. Monte Carlo simulations (and real RHIC data) for gold-gold collisions at $\sqrt{s_{NN}} = 200$ GeV (per nucleon-nucleon collision) show that the transverse momentum density ρ of final-state particles is about 100 GeV per unit area (in the rapidity-azimuth plane). For lead-lead collisions at $\sqrt{s_{NN}} = 2.36$ TeV or 5.5 TeV at the LHC this figure was expected to increase by some factor $\sim 2-3$. This means that jets returned by jet definitions with a radius parameter of, e.g. , $R = 0.4$, will contain background contamination of the order of $\pi R^2 \rho \simeq 50$ and $100 - 150$ GeV respectively.

A related, and perhaps more challenging obstacle to accurate jet reconstruction¹ is due to the fluctuations both of the background level (from event to event, but also from point to point in a single event) and of the jet area itself: knowing the contamination level ‘on average’ is not sufficient to accurately reconstruct each individual jet.

Our framework has the following characteristics:

¹Note that, as for the pp studies carried out earlier in this chapter, we are considering the reconstruction of the jets exclusively at the *particle level*. Detector effects can of course be relevant, and need to be considered in detail, but are beyond the scope of our analysis. Note also that we shall be using the terms ‘(background-)subtracted momentum’ and ‘reconstructed momentum’ equivalently.

- We shall consider a wider range of jet algorithms, still restricting ourselves to IRC safe algorithms. In particular, the jet algorithms we use will not be limited to those yielding jets of regular shape. This point probably has to be put in the historical context: at the time we conducted this study, the LHC had not taken any $PbPb$ data and jet reconstruction in heavy-ion collisions was still in its infancy. Opening our study to a variety of jet algorithms is meant to provide a broader perspective on the topic. We shall see that it indeed comes with a few interesting observations.
- Our analysis will avoid excluding small transverse momentum particles from the clustering as some original approaches were leaning towards. Doing so is collinear-unsafe and inevitably biases the reconstructed jet momenta, which must then be corrected using Monte Carlo simulations, which can have substantial uncertainties in their modelling of jet quenching and energy loss.² Instead, we shall try to achieve a bias-free reconstructed jet, working with all the particles in the event.³
- The analysis presented here was achieved before the grid-based approach to estimating ρ was introduced and before rescaling was suggested to correct for positional dependence of ρ . The study presented below will therefore be limited to jet-based estimates of ρ , with local ranges to control the positional dependence. This is however perfectly sufficient to illustrate our the points we want to make. Unpublished preliminary studies show that similar performance can be reached with rapidity rescaling. In particular a rescaling in ϕ can be used to handle the elliptic flow. We will not include these results here.

Some of the features observed here are also closely connected to some analytic considerations discussed in the next Chapter of this document.

4.2 Details of the study

4.2.1 Simulation and analysis framework

Embedding. As for the pp study from Section 3.2 above, we shall proceed by embedding a “truth” hard jet in a heavy-ion background and study how precisely area–median subtraction corrects the jets back to their truth.

This approach is however less motivated in the case of heavy-ion collisions where the physical separation between the jets and the heavy-ion Underlying Event (HI-UE) is not well-defined as it was the case for pileup in pp collisions. This is due to the intrinsic interaction between the products of the “hard” collision and the dense heavy-ion medium it travels through. In today’s Monte-Carlo event generators for heavy-ion collisions, such a separation is however possible and we shall use that to gain some knowledge of how efficient the area–median subtraction subtracts that unwanted contribution. In that context, it is important to check that our conclusions are not affected neither by the details of the HI-UE, nor by the details of the embedded hard jets, including the fact that they are quenched

²It would still be interesting to test the SoftKiller method (see Chapter 12) in the context of heavy-ion collisions.

³Our framework can, of course, also accommodate the elimination of particles with low transverse momenta, and this might help reduce the dispersion of the reconstructed momentum, albeit at the expense of biasing it. Whether one prefers to reduce the dispersion or instead the bias depends on the specific physics analysis that one is undertaking. Note also that detectors may effectively introduce low- p_t cutoffs of their own. These detector artefacts should not have too large an impact on collinear safety as long as they appear at momenta of the order of the hadronisation scale of QCD, i.e. a few hundred MeV.

of not. In that context, the subtraction procedure can be seen as directly belonging to our definition of jets.

In practice, if the Monte Carlo program used to simulate heavy-ion events explicitly provides a separation between hard events and a soft part — e.g. as does Hydjet [61, 63] — one can extract one of the former as the single hard event and take the complete event as the full one. Alternatively, one can generate a hard event independently and embed it in a heavy-ion event (obtained from a Monte Carlo or from real collisions) to obtain the full event. This second approach, which we have adopted for the bulk of results presented here because of its greater computational efficiency, is sensible as long as the embedded hard event is much harder than any of the semi-hard events that tend to be present in the background. For the transverse momenta that we consider, this condition is typically fulfilled. Note however that for studies in which the presence of a hard collision is not guaranteed, e.g. the evaluation of fake-jet rates discussed below, one should use the first approach.

Matching and quality measures. We shall use the same procedure as for our earlier study of the jet p_t in pp collisions, namely, we will reconstruct the jets in both the hard and the full events, — yielding the “hard” and “full” jets — apply the kinematic cuts on the hard jets to guarantee a clean sample, and consider a full jet as matched to a hard jet if their shared constituents make up at least 50% of the transverse momentum of the constituents of the hard jet.⁴ Events without any matched jets are simply discarded and contribute to the evaluation of matching inefficiencies (see section 4.3.1). For what follows, we shall keep only the pairs of jets which have been matched to one of the two hardest (subtracted) jets in the hard event.

For a matched pair of jets we can compute the corresponding shift Δp_t as introduced in Eq. (3.5). We shall mostly focus our study on same two quality measures used in Section. 3.2: the average p_t shift, $\langle \Delta p_t \rangle$, and its dispersion, $\sigma_{\Delta p_t} \equiv \sqrt{\langle \Delta p_t^2 \rangle - \langle \Delta p_t \rangle^2}$, where the average is made over the events. Small (absolute) values of both $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ will be the sign of a good subtraction.

Monte Carlo simulations. Quantifying the quality of background subtraction using Monte Carlo simulations has several advantages. Besides providing a practical way of generating the hard “signal” separately from the soft background, one can easily check the robustness of one’s conclusions by changing the hard jet or the background sample.

One difficulty that arises in gauging the quality of jet reconstruction in heavy-ion collisions comes from the expectation that parton fragmentation in a hot medium will differ from that in a vacuum. This difference is often referred to as *jet quenching* [64] (for reviews, see [65]). The details of jet quenching are far less well established than those of vacuum fragmentation and can have an effect on the quality of jet reconstruction. Here we shall examine the reconstruction of both unquenched and quenched jets. For the latter it will be particularly important to be able to test more than one quenching model, in order to help build confidence in our conclusions about any reconstruction bias that may additionally exist in the presence of quenching.

In practice, for this study we have used both the Fortran (v1.6) [61] and the C++ (v2.1) [63] versions of Hydjet to generate the background. Hard jets have been generated with Pythia 6.4 [35], either running it standalone, or using the version embedded in Hydjet v1.6.⁵ The quenching effects

⁴Actually, in our implementation, the condition was that the common part should be greater than 50% of the p_t of the hard jet after UE subtraction as in section 4.2.2. In practice, this detail has negligible impact on matching efficiencies and other results.

⁵The events have been generated using the following Hydjet v1.6 program parameters: for RHIC, $\text{nh} = 9000$, $\text{y1f1} = 3.5$, $\text{ytf1} = 1.3$ and $\text{ptmin} = 2.6$ GeV; for LHC, $\text{nh} = 30000$, $\text{y1f1} = 4$, $\text{ytf1} = 1.5$ and $\text{ptmin} = 10$ GeV. In

have been studied using both QPythia [66] and Pyquen [61, 62].

4.2.2 Jet reconstruction: definition and subtraction

Jet definitions. Since jet reconstruction in heavy-ion collisions is far more complex than in the pp case, we shall widen a bit our perspective and consider a series of jet definitions. This will also allow us to highlight some of the key differences between the algorithms, in particular in their response to soft backgrounds.

We shall therefore use the k_t , Cambridge/Aachen (C/A), and anti- k_t algorithms. In an attempt to further reduce the sensitivity of the jets to soft backgrounds, we have also investigated grooming techniques (see the next part of this document for more on this topic). In practice, we have considered “filtering” [59] which works by reclustering each jet with a radius R_{filt} smaller than the original radius R and keeping only the n_{filt} hardest subjects (see also Chapter 6).⁶ Background subtraction is applied to each of the subjects before deciding which ones to keep. In this study, we have used filtering applied to Cambridge/Aachen (C/A(filt)) with $R_{\text{filt}} = R/2$ and $n_{\text{filt}} = 2$. We have not used SIS Cone, as its relatively slower speed compared to the sequential recombination algorithms makes it less suitable for a HI environment.

In all cases, we have taken the radius parameter $R = 0.4$. We have adopted this value as it is the largest used at RHIC and a reasonable choice for the LHC.⁷ Note that the effect of the background fluctuations of the jet energy resolution increases linearly with R , disfavoring significantly larger choices. On the other hand, too small a choice of R may lead to excessive sensitivity to the details of parton fragmentation, hadronisation and detector granularity.

Background determination and subtraction. This will be carried using the now familiar area–median approach. The background density ρ will be obtained from a jet-based estimation and the choice of an algorithm and a jet radius R_ρ will be discussed below.

To handle the positional dependence, we shall use the local range technique. Besides the *global* estimate of ρ , independent of the jet position, we shall consider three options for a local range introduced in Section 2.4.2: the strip range $\mathcal{S}_\Delta(j)$, the circular range $\mathcal{C}_\Delta(j)$, and the doughnut range $\mathcal{D}_{\delta,\Delta}(j)$. Whenever this is necessary, we shall denote by $\mathcal{R}(j)$ a range around jet j .

When using a local range, a compromise needs to be found between choosing a range small enough to get a valid local estimation, but also large enough to contain a sufficiently large number of background (soft) jets for the estimation of the median to be reliable. Two effects need to be considered: (1) statistical fluctuations in the estimation of the background and (2) biases due to the presence of a hard jet in the region used to estimate the background.

1. If we require that the dispersion in the reconstructed jet p_t coming from the statistical fluctuations in the estimation of the background (Eq. (5.103) multiplied by the area of the jet) does not amount to more than a fraction ϵ of the overall $\sigma_{\Delta p_t} \approx \sigma \sqrt{A_{\text{jet}}}$, then — see Section 4.4.4 for details — the range needs to cover an area $A_{\mathcal{R}}$ such that

$$A_{\mathcal{R}} \gtrsim \frac{\pi^2 R^2}{4\epsilon}. \quad (4.1)$$

both cases quenching effects are turned on in Hydjet, `nhse1 = 2`, even when they are not included for the embedded pp event. The corresponding Pyquen parameters we have used are `iengl = 0`, `ianglu = 0`, `T0 = 1.0 GeV`, `tau0 = 0.1 fm` and `nf = 0`.

⁶Note that the filtering that we use here is unrelated to the Gaussian filter approach of [67] (not used in this paper as no public code is currently available).

⁷The LHC experiments tend however to adopt smaller values in the 0.2-0.3 range.

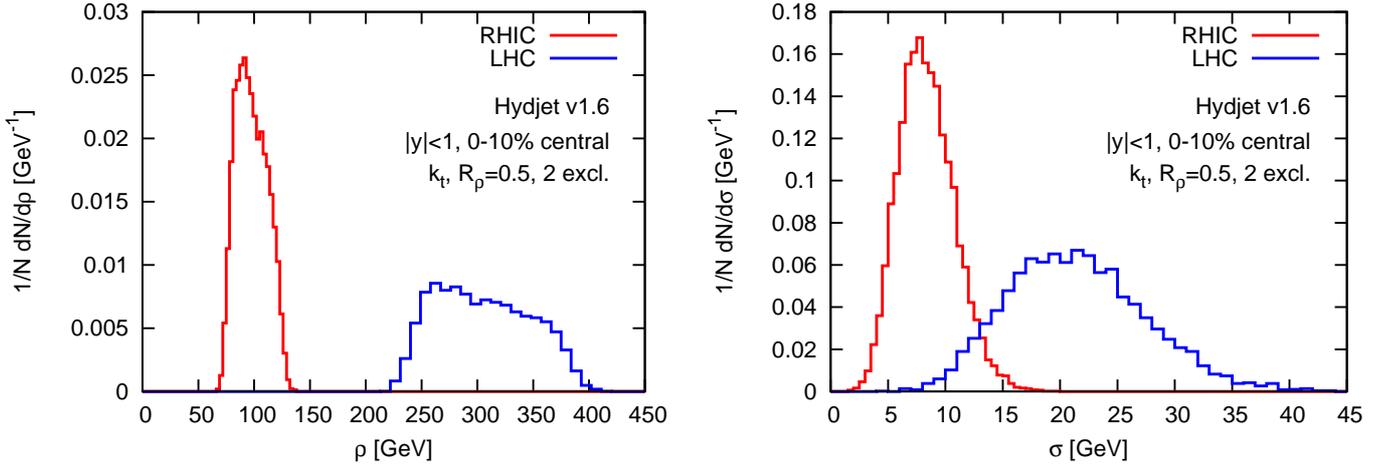


Figure 4.1: Distribution of the background density ρ per unit area (left) and its intra-event fluctuations σ (right). It has been obtained from 5000 Hydjet events with RHIC (AuAu, $\sqrt{s_{NN}} = 200$ GeV) and LHC (PbPb, $\sqrt{s_{NN}} = 5.5$ TeV) kinematics. The background properties have been estimated using the techniques presented in section 4.2.2, using the k_t algorithm with $R_\rho = 0.5$, and keeping only the jets with $|y| < 1$ (excluding the two hardest).

Taking $\epsilon = 0.1$ and $R = 0.4$ this corresponds to an area $A_{\mathcal{R}} \gtrsim 25R^2 \simeq 4$. For the applications below, when clustering with radius R , we have used the rapidity-strip ranges \mathcal{S}_{2R} and \mathcal{S}_{3R} , the circular range \mathcal{C}_{3R} and the doughnut range $\mathcal{D}_{R,3R}$. One can check that their areas are compatible with the $25R^2$ lower-limit estimated above.

2. The bias in the estimate of ρ due to the presence of n_b hard jets in the range is given roughly by

$$\langle \Delta\rho \rangle \simeq 1.8\sigma R_\rho \frac{n_b}{A_{\mathcal{R}}}, \quad (4.2)$$

as discussed in Section 4.4.5. The ensuing bias on the p_t can be estimated as $\langle \Delta\rho \rangle \pi R^2$ (for anti- k_t jets). For $R = 0.4$, $R_\rho = 0.5$, $A_{\mathcal{R}} \simeq 4$ and $n_b = 1$, the bias in the reconstructed jet p_t is $\simeq 0.1\sigma$. Given σ in the 10 – 20 GeV range (as we will find in section 4.3) this corresponds to a 1 – 2 GeV bias. In order to eliminate this small bias, we will often choose to exclude the two hardest jets in each event when determining of ρ .⁸

A third potential bias discussed in ref. [19] is that of underestimating the background when using too small a value for R_ρ . Given the high density of particles in HI collisions, this will generally not be an issue as long as $R_\rho \sim 0.5$.

4.3 Results

Let us first discuss the main characteristics of the background properties. Our setup for the background, using Hydjet v1.6 with 0-10% centrality leads, for RHIC (AuAu, $\sqrt{s_{NN}} = 200$ GeV), to an

⁸We deliberately choose to exclude the two hardest jets in the event, not simply the two hardest in the range. Note, however, that for realistic situations with limited acceptance, only one jet may be within the acceptance, in which case the exclusion of a single jet might be more appropriate. Excluding a second one does not affect the result significantly, so it is perhaps a good idea to use the same procedure regardless of any acceptance-related considerations.

average background density per unit area at central rapidity of $\langle\rho\rangle \simeq 99$ GeV with average fluctuations in a single event of $\langle\sigma\rangle \simeq 8$ GeV and event-to-event fluctuations $\sigma_\rho \equiv \sqrt{\langle\rho^2\rangle - \langle\rho\rangle^2} \simeq 14$ GeV. For the LHC (PbPb, $\sqrt{s_{NN}} = 5.5$ TeV) the corresponding values are $\langle\rho\rangle \simeq 310$ GeV, $\langle\sigma\rangle \simeq 20$ GeV and event-to-event fluctuations $\sigma_\rho \simeq 45$ GeV. Fig. 4.1 shows the distributions obtained from the simulations for ρ and σ .

In the case of the RHIC simulation, the result for $\langle\rho\rangle$ is somewhat higher than the experimental value of 75 GeV quoted by the STAR collaboration [68], however this is probably in part due to limited tracking efficiencies at low p_t at STAR,⁹ and explicit STAR [69] and PHENIX [70] results for $dE_t/d\eta$ correspond to somewhat higher ρ values, about 90 GeV. The multiplicity of charged particles ($dN_{ch}/d\eta \simeq 660$ for $\eta = 0$ and 0–6% centrality) and the pion p_t spectrum in our simulation are sensible compared to experimental measurements at RHIC [71, 72, 73]. For the LHC our charged particle multiplicity is $dN_{ch}/d\eta \simeq 1600$ for $\eta = 0$ and 0–10% centrality, which is comparable to many of the predictions reviewed in fig. 7 of [74].¹⁰

An independent control analysis has also been performed with Hydjet++ 2.1 [63] (with default parameters) for the background. The results at RHIC are similar, while for LHC the comparison is difficult because the default tune of Hydjet++ 2.1 predicts a much higher multiplicity, $dN_{ch}/d\eta \simeq 2800$ for $\eta = 0$ and 0–10% centrality.

Most of the results of this section will be obtained without quenching, though in section 4.3.5 we will also consider the impact on our conclusions of the Pyquen and QPythia simulations of quenching effects.

For the results presented below, we have employed a selection cut of $|y| \leq y_{\max}$ on the jets with $y_{\max} = 1$ for RHIC and $y_{\max} = 2.4$ for the LHC.¹¹ We only consider full jets that are matched to one of the two hardest jets in the hard event. The computation of the jet areas in **FastJet**, needed both for the subtraction and the background estimation, has been performed using active areas (including explicit ghosts), with ghosts up to $y_{\max} + 1.8$, a single repetition and a ghost area of 0.01. The determination of the background density ρ has been performed using the k_t algorithm with $R_\rho = 0.5$. Though the estimate of the background depends on R_ρ [19] (see also Section 5.2), we have observed that choices between 0.3 and 0.5 lead to very similar results (e.g. differing by at most a few hundred MeV at RHIC).

4.3.1 Matching efficiency

Let us start the presentation of our results with a brief discussion of the efficiency of reconstructing jets in the medium. As explained in Section 4.2.1, the jets in the medium are matched to a “bare” hard jet when their common particle content accounts for at least 50% of the latter’s transverse momentum.

The matching efficiencies we observe depend to some extent on the details of the Monte Carlo

⁹We thank Helen Caines for discussions on this point.

¹⁰For reference, at 2.76 TeV with 0 – 5% centrality, our Hydjet simulation gives $dN_{ch}/d\eta \simeq 1520$ for $|\eta| < 0.5$, which can be compared to the recent ALICE result of $dN_{ch}/d\eta \simeq 1584 \pm 4 \pm 76$ [75], which appeared subsequent to version 1 of our article. The two numbers are in agreement within systematic errors, and we note that our conclusions are in any case largely independent of fine details of the background.

¹¹This corresponds roughly to the central region for the ATLAS and CMS detectors. Note that we also kept particles beyond y_{\max} in the jet clustering and apply the acceptance cut only to the resulting jets.

For ALICE, the acceptance is more limited [76, 77]. Some adaptation of our method will be needed for estimating ρ in that case, in order for information to be derived from jets near the edge of the acceptance and thus bring the available area close to the ideal requirements set out in Section 4.4.4. One option that we have recommended based on preliminary quick tests is to limit the ghosts to the detector acceptance and to include all the resulting patches in the determination of ρ .

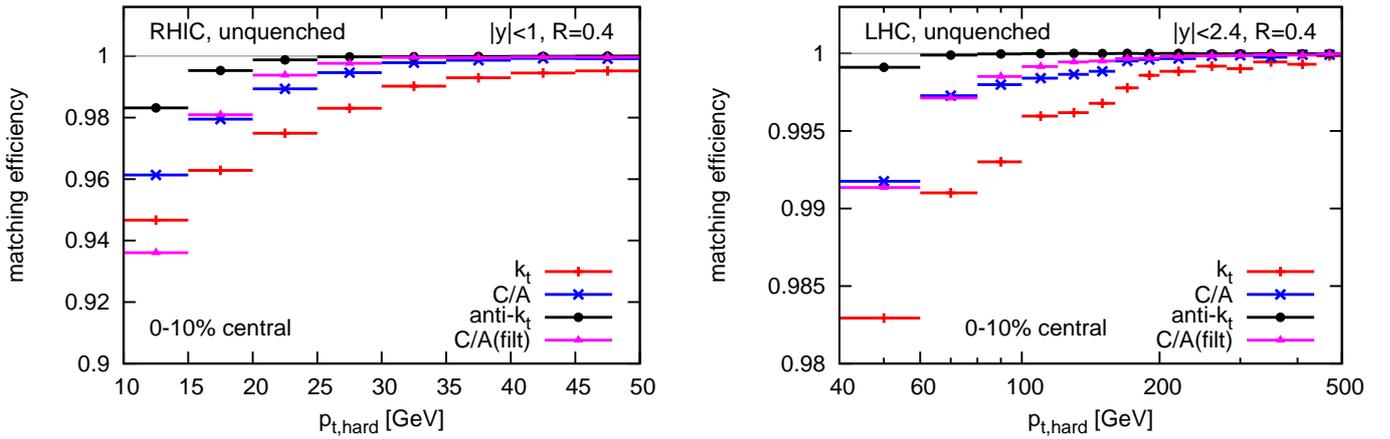


Figure 4.2: Matching efficiency for reconstructed jets as a function of the jet p_t . Left: RHIC, right: LHC. These results are independent of the choice of background-subtraction range in the heavy-ion events, since background subtraction does not enter into the matching criterion. Here and in later figures, the label “unquenched” refers to the embedded pp event; the background is always simulated including quenching.

used for the background so our intention is just to illustrate the typical behaviour we observe and highlight that these efficiencies tend to be large. We observe from fig. 4.2 that we successfully match at least 95% of the jets above $p_t \simeq 15$ GeV at RHIC, and at least 99% of the jets above $p_t \simeq 60$ GeV at the LHC. It is also interesting to notice that the anti- k_t algorithm performs best, likely as a consequence of its ‘rigidity’, namely the fact that anti- k_t jets tend to have the same (circular) shape, independently of the soft-particles that are present.

4.3.2 Choice of background-estimation range

We now turn to the results concerning the measurement of the background density and the reconstruction of the jet transverse momentum. We first concentrate on the impact of the choice of a local range and/or of the exclusion of the two hardest jets when determining ρ .¹²

In fig. 4.3 we show the average shift $\langle \Delta p_t \rangle$ for the list of ranges mentioned in section 4.2.2. The results presented here have been obtained with the anti- k_t algorithm with $R = 0.4$, but the differences among the various range choices have been seen to be similar with other jet definitions. The label “2 excl” means that the two hardest jets in the event have been excluded from the estimation of the background. We have found that this improves the precision of the subtractions whenever expected, i.e. for all choices of range except the doughnut range, where its central hole already acts similarly to the exclusion of the hardest jets. To keep the figure reasonably readable, we have only explicitly shown the effect of removing the two hardest jets for the global range. The change of 0.4–0.6 GeV (both for RHIC and the LHC) is in reasonable agreement with the analytic estimate of about 0.6 GeV for RHIC and the LHC obtained from $\langle \Delta p_t \rangle = \pi R^2 \langle \Delta \rho \rangle$ with $\langle \Delta \rho \rangle$ calculated using Eq. (4.2). Note that at LHC the exclusion of the two hardest jets for the global range appears to worsen the subtraction,

¹²Independently of the choice made for the full event, we always use a global range up to $|y| = y_{\max}$ for the determination of ρ in the hard event, without exclusion of any jets. This ensures that the reference jet p_t is always kept the same. The impact of subtraction in the hard event is in any case small, so the particular choice of range is not critical.

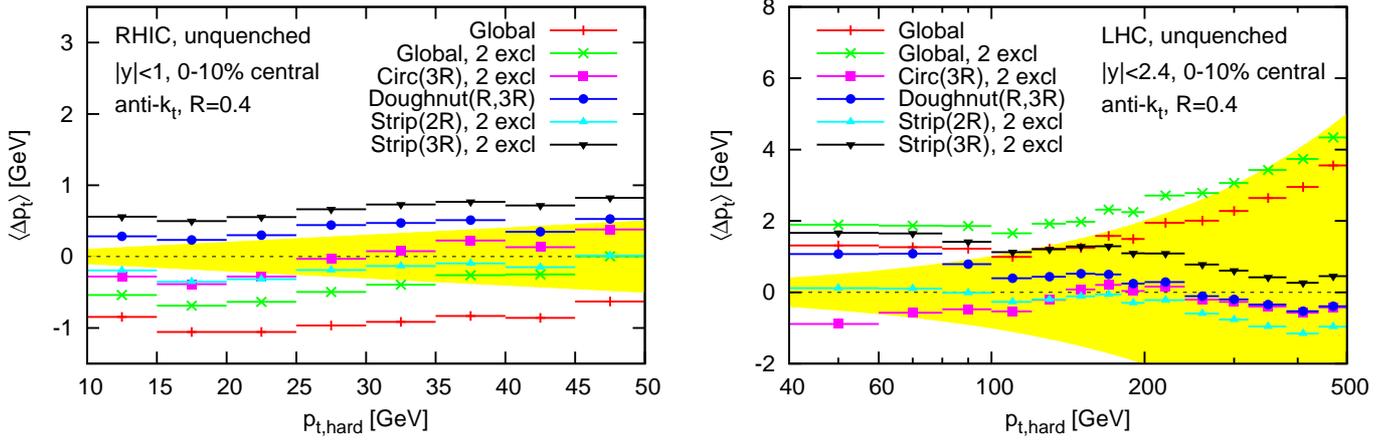


Figure 4.3: Effect of the choice of range on the average p_t shift, Δp_t , as defined in Eq. (3.5). Left: RHIC, right: LHC. In this figure and those that follow, the yellow band corresponds to 1% of the p_t of the hard jet.

however what is really happening is that the removal of the two hardest jets exacerbates a deficiency of the global range, namely the fact that its broad rapidity coverage causes it to underestimate ρ , leading to a positive net $\langle \Delta \rho \rangle$.

Other features that can be understood qualitatively include for example the differences between the two strip and the global (2 excl) range for RHIC: while the rapidity width of the global range lies in between that of the two strip ranges, the global range gives a lower $\langle \Delta p_t \rangle$ than both, corresponding to a larger ρ estimate, which is reasonable because the global range is centred on $y = 0$, whereas the strip ranges are mostly centred at larger rapidities where the background is lower.

The main result of the analysis of fig. 4.3 is the observation that all choices of a local range lead to a small residual Δp_t offset: the background subtraction typically leaves a $|\langle \Delta p_t \rangle| \lesssim 1$ GeV at both RHIC and LHC, i.e. better than 1-2% accuracy over much of the p_t range of interest. It is not clear, within this level of accuracy, if one range is to be preferred to another, nor is it always easy to identify the precise origins of the observed differences between various ranges.¹³ Another way of viewing this is that the observed differences between the various choices give an estimate of the residual subtraction error due to possible misestimation of ρ . For our particular analysis, at RHIC this comment also applies to the choice of the global range (with the exclusion of the two hardest jets in the event). This is a consequence of the limited rapidity acceptance, which effectively turns the global range into a local one, a situation that does not hold for larger rapidity acceptances, as we have seen for the LHC results. In what follows we will use the Doughnut($R, 3R$) choice, since it provides a good compromise between simplicity and effectiveness.

4.3.3 Choice of algorithm

The next potential systematic effect that we consider is the choice of the jet algorithm used for the clustering.¹⁴ Fig. 4.4 shows the distribution of Δp_t for each of our four choices of jet algorithm, k_t , C/A, anti- k_t and C/A(filt), given for RHIC collisions and a specific bin of the hard jets' transverse momenta, $30 < p_{t,\text{hard}} < 35$ GeV. One sees significant differences between the different algorithms.

¹³Furthermore, the differences may also be modified by jet-medium interactions.

¹⁴Recall that in all cases, the k_t algorithm with $R_\rho = 0.5$ is used for the estimation of the background.

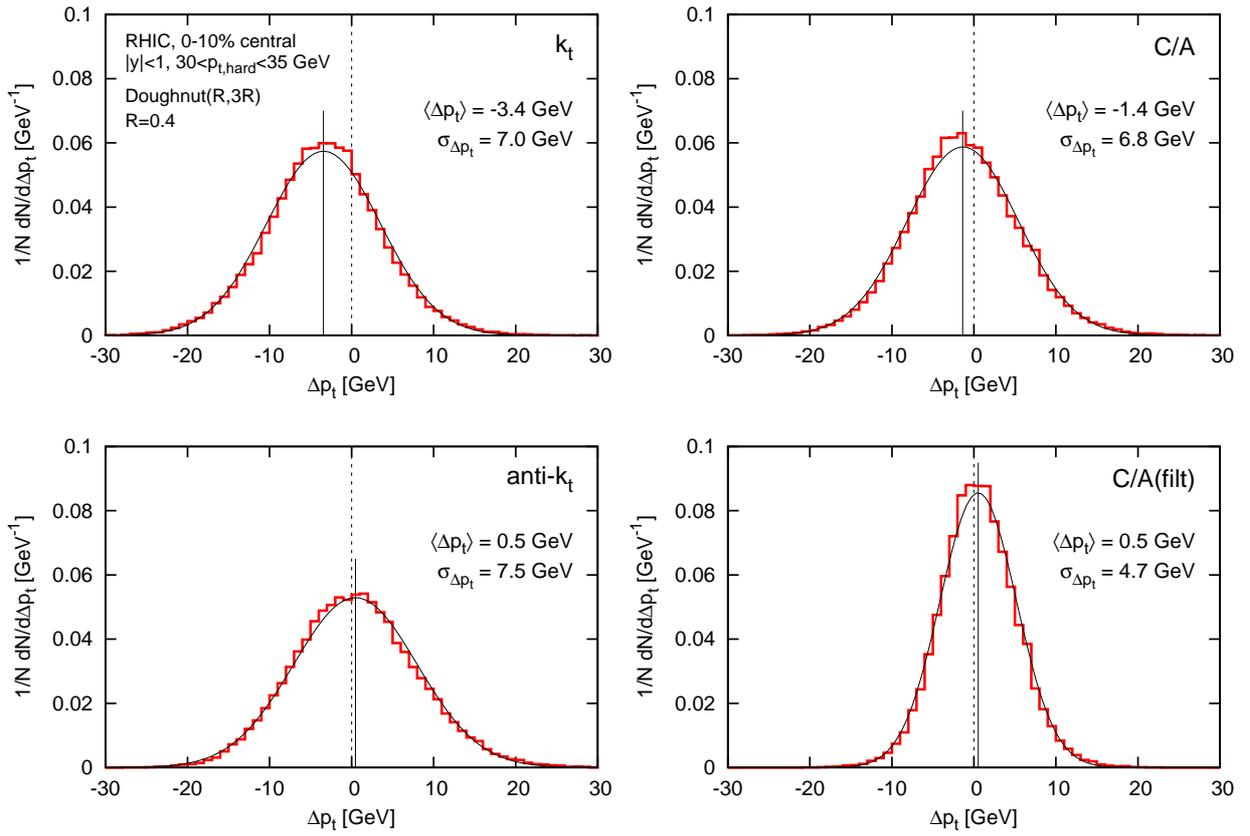


Figure 4.4: Distribution of Δp_t (red histograms) for each of our 4 jet algorithms, together with a Gaussian (black curve) whose mean (solid vertical line) and dispersion are equal to $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ respectively.

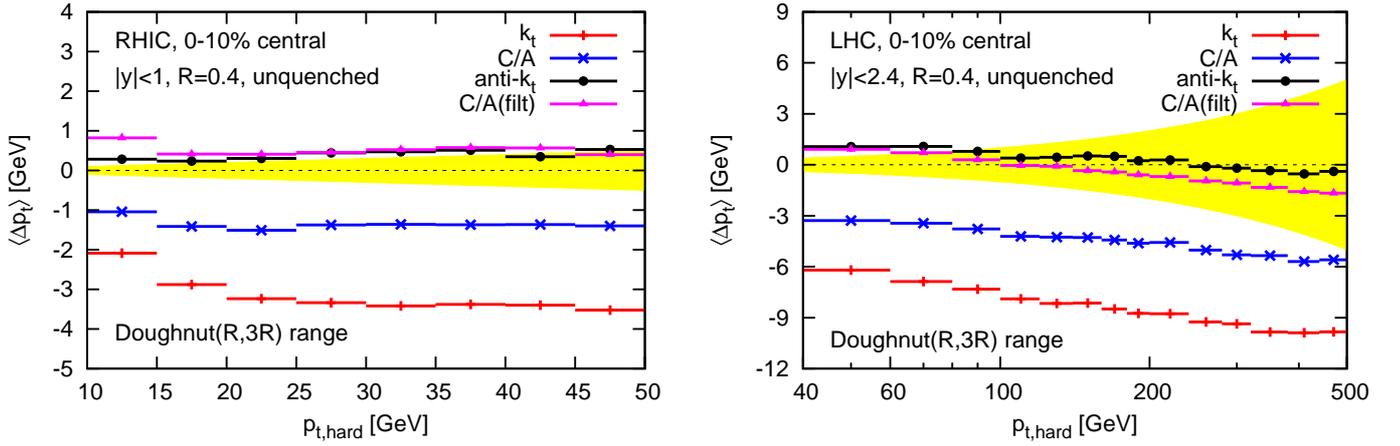


Figure 4.5: Average shift $\langle \Delta p_t \rangle$, as a function of $p_{t,\text{hard}}$, shown for RHIC (left) and the LHC (right).

One also observes that Gaussians with mean and dispersion set equal to $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ provide a fair description of the full histograms. This validates our decision to concentrate on $\langle \Delta p_t \rangle$ and $\sigma_{\Delta p_t}$ as quality measures. One should nevertheless be aware that in the region of high $|\Delta p_t|$ there are deviations from perfect Gaussianity, which are more visible if one replicates fig. 4.4 with a logarithmic vertical scale (not shown, for brevity).

Average Δp_t . The first observable we analyse is the average p_t shift. We show in fig. 4.5 the $\langle \Delta p_t \rangle$ results for the four algorithms listed in Section 4.2.2, as a function of $p_{t,\text{hard}}$. We use the doughnut range to estimate the background. The first observation is that, while the anti- k_t and C/A(filt) algorithms have a small residual $\langle \Delta p_t \rangle$, the C/A and k_t algorithms display significant offsets. The reason for the large offsets of k_t and C/A is well understood, related to the effect of *back-reaction*. This is the fact that the addition of a soft background can alter the clustering of the particles of the hard event: some of the constituents of a jet in the hard event can be gained by or lost from the jet when clustering the event with the additional background of the full event. This happens, of course, on top of the simple background contamination that adds background particles to the hard jet. Even if this latter contamination is subtracted exactly, the reconstructed p_t will still differ from that of the original hard jet as a consequence of the back-reaction.

The effect of the back-reaction can be studied in detail, since in Monte Carlo simulations it is possible to identify which hard-event constituents are present in a given jet before and after inclusion of the background particles in the clustering. The average p_t shift due to back-reaction can be seen in fig. 4.6 for the different jet algorithms. As expected (see Section 5.1.3 or Refs. [22, 13]), it is largest for k_t , and smallest (almost zero, in fact) for anti- k_t . By comparing fig. 4.6 and fig. 4.5 one can readily explain the difference between the $\langle \Delta p_t \rangle$ offsets of the various algorithms in terms of their back-reaction. The rigidity (and hence small back-reaction) of the anti- k_t jets manifestly gives almost bias-free reconstructed jets, while the large back-reaction effects of the k_t algorithm and, to a smaller extent, of the C/A algorithm translates into a worse performance in terms of average shift. The p_t dependence of the back-reaction is weak. This is expected based on the interplay between the $\ln \ln p_t$ dependence found in Section 5.1.3 [13] and the evolution with p_t of the relative fractions of quark and gluon jets.

The case of the C/A(filt) algorithm is more complex: its small net offset, comparable to that of the anti- k_t algorithm, appears to be due to a fortuitous compensation between an under-subtraction

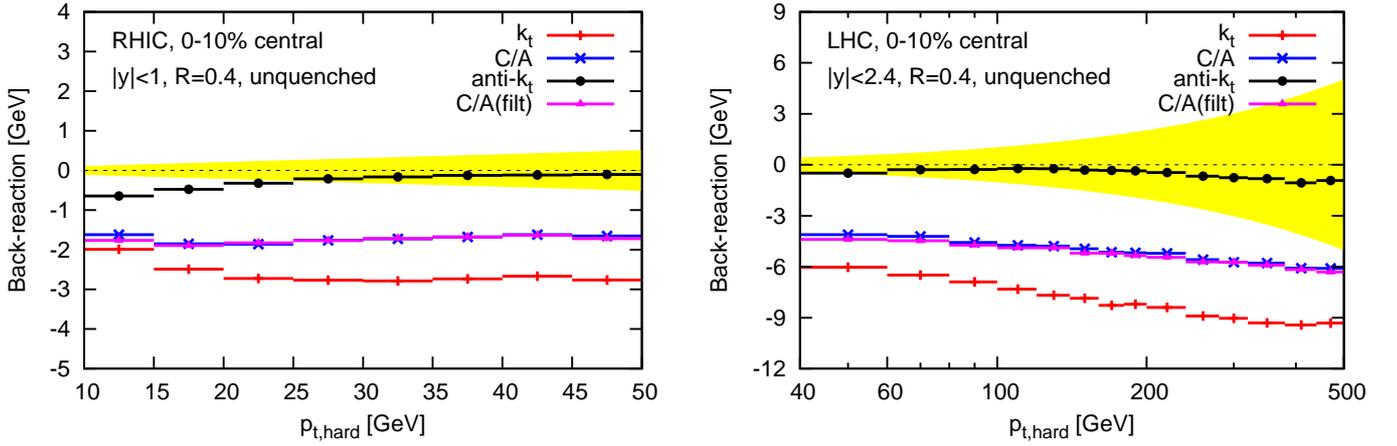


Figure 4.6: Contribution to $\langle \Delta p_t \rangle$ due to back-reaction. Note that these results are independent of the range used for estimating ρ in the heavy-ion event. Left: RHIC, right: LHC.

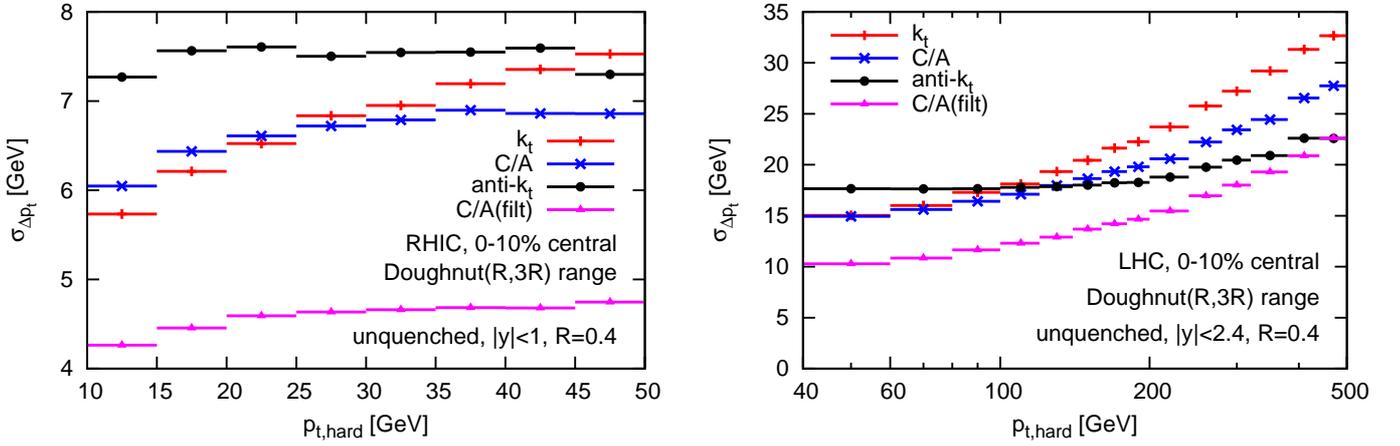


Figure 4.7: Dispersion $\sigma_{\Delta p_t}$. Left: RHIC, right: LHC.

of the background and a negative back-reaction. The negative back-reaction is very similar to that of C/A without filtering, while the under-subtraction is related to the fact that the selection of the hardest subjects introduces a bias towards positive fluctuations of the background. This effect is discussed in Section 4.4.6, where we obtain the following estimate for the average p_t shift (specifically for $R_{\text{filt}} = R/2$):

$$\langle (\Delta p_t)_{\text{filt}} \rangle \simeq 0.56 R \sigma, \quad (4.3)$$

yielding an average bias of 2 GeV for RHIC and 4.5 GeV at the LHC, which are both in good agreement with the differences observed between C/A with and without filtering in fig. 4.5. Note that while the bias in Eq. (4.3) is proportional to σ , the back-reaction bias is instead mainly proportional to ρ . It is because of these different proportionalities that the cancellation between the two effects should be considered as fortuitous. Since it also depends on the substructure of the jet, one may also expect that the cancellation that we see here could break down in the presence of quenching.

Dispersion of Δp_t . Our results for the Δp_t dispersion, $\sigma_{\Delta p_t}$, are shown in fig. 4.7, again using the doughnut range. (Our conclusions are essentially independent of the particular choice of range.)

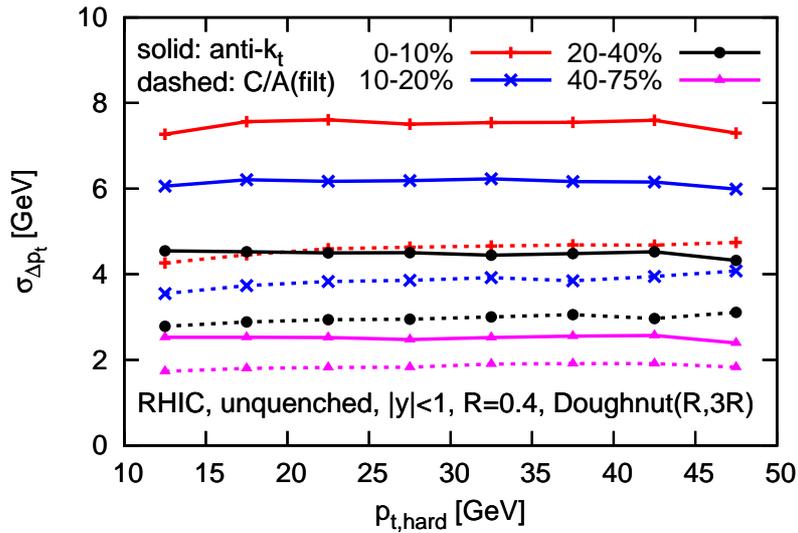


Figure 4.8: p_t dependence dependence of the Δp_t dispersion at RHIC, for different centrality classes.

We first discuss the case of RHIC kinematics. For k_t and anti- k_t , the observed dispersions are similar to the result of 6.8 GeV quoted by STAR [68] (though the number from STAR includes detector resolution effects, so that the true physical $\sigma_{\Delta p_t}$ may actually be somewhat lower). Of note, the advantage enjoyed by anti- k_t in terms of smallest $\langle \Delta p_t \rangle$ does not hold at the level of the dispersion: C/A and k_t tend to behave slightly better at small transverse momentum. The algorithm which performs best in terms of dispersion over all the p_t range is now C/A with filtering, for which the result is smaller than that of the other algorithms by a factor of about $1/\sqrt{2}$. This reduction factor can be explained because the dispersion $\sigma_{\Delta p_t}$ is expected to be proportional to the square-root of the jet area: the C/A(filt) algorithm with $R_{\text{filt}} = R/2$ and $n_{\text{filt}} = 2$ produces jets with an area of, roughly, half that obtained with C/A; hence the observed reduction of $\sigma_{\Delta p_t}$.

In the LHC setup, the conclusions are quite similar at the lowest transverse momenta shown. As p_t increases, the dispersion of the anti- k_t algorithm grows slowly, while that of the others grows more rapidly, so that at the highest p_t 's shown, the k_t and C/A algorithms have noticeably larger dispersions than anti- k_t , and C/A(filt) becomes similar to anti- k_t . The growth of the dispersions can be attributed to an increase of the back-reaction dispersion. The latter is dominated by rare occurrences, where a large fraction of the jet's p_t is gained or lost to back-reaction, hence the noticeable p_t dependence (*cf.* Section 4.4.1). An additional effect, especially for the k_t algorithm, might come from the anomalous dimension of the jet areas, i.e. the growth with p_t of the average jet area.

Note that the dispersion has some limited dependence on the choice of ghost area — for example, reducing it from 0.01 to 0.0025 lowers the dispersions by about 0.2–0.4 GeV at RHIC. This is discussed further in Section 4.4.2.

4.3.4 Centrality dependence

So far, we have only considered central collisions. Since it is known that non-central collisions give rise to elliptic flow [30, 31], one might worry that this leads to an extra source of background fluctuations and/or non-uniformities, potentially spoiling the subtraction picture discussed so far. One can study this on azimuthally averaged jet samples (as we have been doing so far) or as a function of the

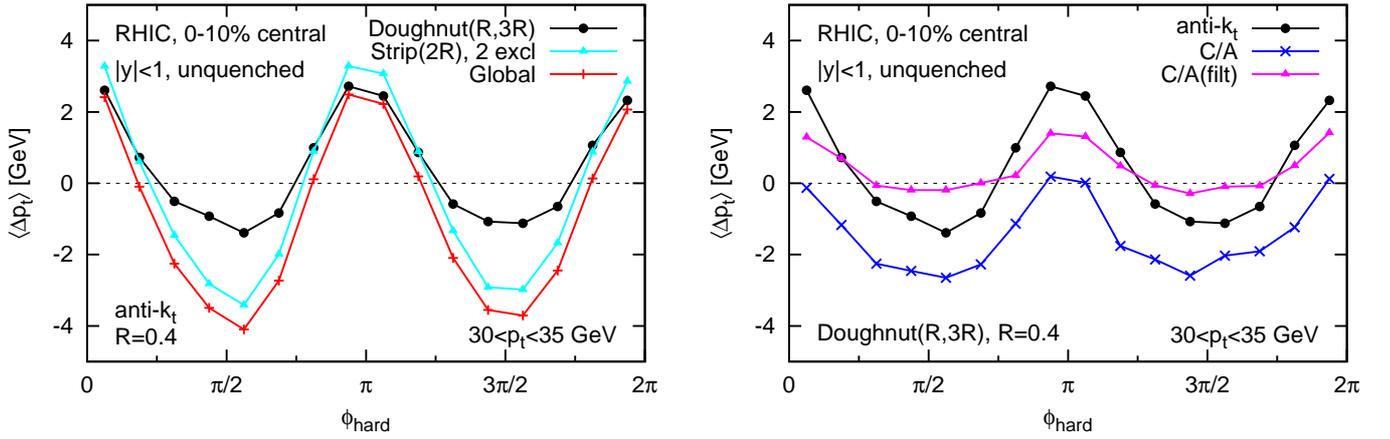


Figure 4.9: ϕ dependence of the Δp_t shift at RHIC, for the 0 – 10% centrality bin. Left: for three different ranges for the anti- k_t algorithm; right: for three different jet algorithms for the doughnut($R, 3R$) range. The absolute size of the ϕ dependence is similar for centralities up to 40% and then decreases beyond.

azimuthal angle, $\Delta\phi$, between the jet and the reaction plane. As above, we use Hydjet v1.6, whose underlying HYDRO component includes a simulation of elliptic flow [78].

We have generated heavy-ion background events for RHIC in four different centrality bins: 0-10% (as above), 10-20%, 20-40% and 40-75%, with v_2 values respectively of 1.7%, 3.3%, 5.0% and 5.3%.¹⁵

We first examine azimuthally averaged results, repeating the studies of the previous sections for each of the centrality bins. We find that the results for the average shift, $\langle \Delta p_t \rangle$, are largely independent of centrality, as expected if the elliptic flow effects disappear when averaged over ϕ . The results for the dispersion are shown in fig. 4.8. We observe that the dispersion decreases with increasing non-centrality. Even though one might expect adverse effects from elliptic flow, the heavy-ion background decreases rapidly when one moves from central to peripheral collisions, and this directly translates into a decrease of $\sigma_{\Delta p_t}$.

The first conclusion from this centrality-dependence study is therefore that the subtraction methods presented in this paper appear to be applicable also for azimuthally averaged observables in non-central collisions.

We next consider results as a function of $\Delta\phi$, which is relevant if one wishes to examine the correlation between jet quenching and the reaction plane. An issue in real experimental studies is the determination of the reaction plane, and the extent to which it is affected by the presence of hard jets. In the Hydjet simulations, this problem does not arise because the reaction plane always corresponds to $\phi = 0$. Figure 4.9 (left) shows the average Δp_t as a function of $\Delta\phi$ for the anti- k_t algorithm and several different background-estimation ranges, for the 0 – 10% centrality bin for RHIC. The strip range shows significant $\Delta\phi$ dependence, which is because a determination of ρ averaged over all ϕ cannot possibly account for the local ϕ -dependence induced by the elliptic flow. Other ranges, such as the doughnut range, instead cover a more limited region in ϕ . They should therefore be able to provide information on the ϕ -dependence of the background.¹⁶ However, since their extent in ϕ tends

¹⁵ v_2 was determined as the average of $\cos 2\phi$ over all particles with $|\eta| < 1$ (excluding the additional hard pp event).

¹⁶At the expense of being more strongly affected by jet-medium interactions that could manifest themselves as broad enhancement of the energy flow in the vicinity of the jet.

to be significantly larger than that of the jet, and ρ varies relevantly over that extent, some residual ϕ dependence remains in $\langle \Delta p_t \rangle$ after subtraction. The right-hand plot of figure 4.9 shows that the effect is reduced with filtering, as is to be expected since its initial background contamination is smaller. The conclusion from this part of the study is that residual ϕ -dependent offsets may need to be corrected for explicitly in any studies of jets and their correlations with the reaction plane. The investigation of extensions to our background subtraction procedure to address this issue will be the subject of future work.

4.3.5 Quenching effects

An important issue we wish to investigate is how the phenomenon of jet quenching (i.e. medium effects on parton fragmentation) may affect the picture developed so far. The precise nature of jet quenching beyond its basic analytic properties (see e.g. [64]) is certainly hard to estimate in detail, especially at the LHC, where experimental data from, say, flow or particle spectra measurements are not yet available for the tuning of the Monte Carlo simulations. Additionally, the implementation of Monte-Carlo generators that incorporate the analytic features of jet quenching models is currently a very active field [66, 79, 80]. We may therefore expect a more robust and complete picture of jet quenching in the near future, together with the awaited first PbPb collisions at the LHC.

In this section, we examine the robustness of our HI background subtraction in the presence of (simulated) jet quenching.¹⁷ For this purpose we have used two available models which allow one to simulate quenched hard jets, Pyquen [61], which is used by Hydjet v1.6, and QPythia [66]. Pyquen has been run with the parameters listed in footnote 5 for the LHC, and with $T_0 = 0.5$ GeV, $\tau_{a0} = 0.4$ fm and $n_f = 2$ for RHIC¹⁸. For QPythia we have tested two options for the values of the transport coefficient and the medium length ($\hat{q} = 3$ GeV²/fm, $L = 5$ fm and $\hat{q} = 1$ GeV²/fm, $L = 6$ fm), with similar results. No serious attempt has been made to tune the two codes with each other or with the experimental data, beyond what is already suggested by the code defaults: in the absence of strong experimental constraints on the details of the quenching effects, this allows us to verify the robustness of our results for a range of conditions.

As in sections 4.3.2 and 4.3.3, we have embedded the hard Pyquen or QPythia events in a Hydjet v1.6 background and tested the effectiveness of the background subtraction for different choices of algorithm.¹⁹ We shall restrict our attention to the anti- k_t and C/A(filt) algorithms, as they appear to be the optimal choices from our analysis so far.

We have found that the jet-matching efficiencies are still high, with essentially no changes at RHIC, and at LHC a doubling of the (small) inefficiencies that we saw in fig. 4.2(right). The dispersions $\sigma_{\Delta p_t}$ are also not significantly affected within our sample of jet-quenching simulations. We therefore concentrate on the $\langle \Delta p_t \rangle$ offset, which is plotted in fig. 4.10 for Pyquen. The results are shown for both RHIC and the LHC. In the case of RHIC, and for the whole p_t range up to about 50 GeV, quenching can be seen not to significantly affect the subtraction offset $\langle \Delta p_t \rangle$ (within the usual uncertainty related to the choice of range, which was shown in fig. 4.3). In the LHC case, instead, while the shift obtained using the anti- k_t algorithm is largely similar to the unquenched case, the C/A(filt) algorithm performance can be seen to deteriorate slightly when quenching is turned on, all the more so at very

¹⁷ Our focus here is therefore not the study of quenching itself, but merely how it may affect our subtraction procedure.

¹⁸The parameters for RHIC are taken from [61]. The difference with the default parameters does not appear to be large for the purpose of our investigations and, in any case, a systematic study of quenching effects is not among the goals of this paper.

¹⁹The effect of the choice of range remains as in section 4.3.2 for the unquenched case. We will therefore keep employing the doughnut range.

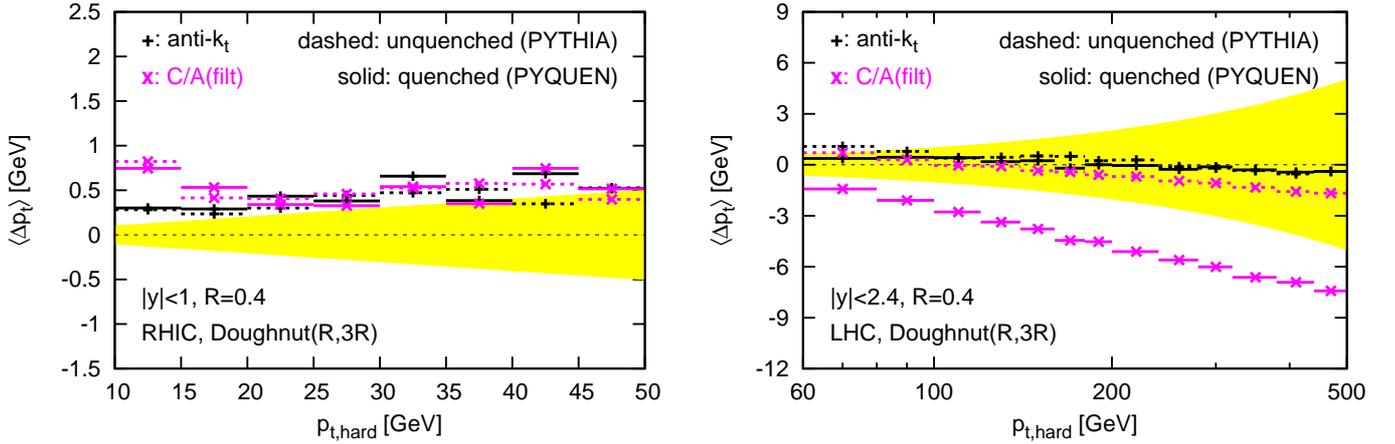


Figure 4.10: Average p_t shift for background-subtracted jets with the anti- k_t and C/A(filt) jet algorithms. The dashed lines correspond to unquenched hard jets (Pythia) and the solid ones to quenched hard jets (Pyquen). Results are shown for RHIC kinematics on the left plot and for the LHC on the right one.

large transverse momentum. The k_t and the C/A algorithms are not shown for clarity, but they share the behaviour of C/A(filt). This deterioration of the quality of the subtraction can be traced back to an increased back-reaction compared to the unquenched jets. Anti- k_t jets do not suffer from this effect as a consequence of the usual rigidity of this algorithm. In the case of C/A(filt), one should nevertheless emphasise that an error of (at most) 10 GeV on the reconstruction of a 500 GeV jet is still only a 2% effect. This is modest, both relative to the likely experimental precision and to the expected effect of quenching on the overall jet p_t , predicted by Pyquen to be at the level of 10% at this p_t .

Though for brevity we have not explicitly shown them, the results with QPythia are very similar.

Before closing this section, we reiterate that we have only investigated simple models for quenching and that our results are meant just to give a first estimate of the effects that one might have to deal with in the case of quenched jets. The expected future availability of new “quenched” Monte Carlo programs, together with specific measurements in the early days of heavy-ion collisions at the LHC, will certainly allow one to address this question more extensively.

4.4 Further details and analytic estimates

4.4.1 Back reaction contribution to the dispersion

In the context of heavy-ion collisions, where the residual resolution-degradation effects on the jets can be relatively large, it is interesting to look a bit deeper in the various sources that contribute to the dispersion. The natural candidates are the UE fluctuations themselves and residual uncertainties on the estimation of ρ (which have been made small by choosing large-enough a range). Here we discuss two additional sources: the dispersion in the back-reaction which starts kicking in at large p_t and the uncertainty in our calculation of the jet area.

We stated, in section 4.3.3, that the increase of the dispersion at high p_t seen in fig. 4.7 was mainly due to back-reaction. This is made explicit in fig. 4.11, which decomposes the dispersion into its two components: that associated with the back-reaction, $\sigma_{\Delta p_t}^{\text{BR}}$ and that associated with background

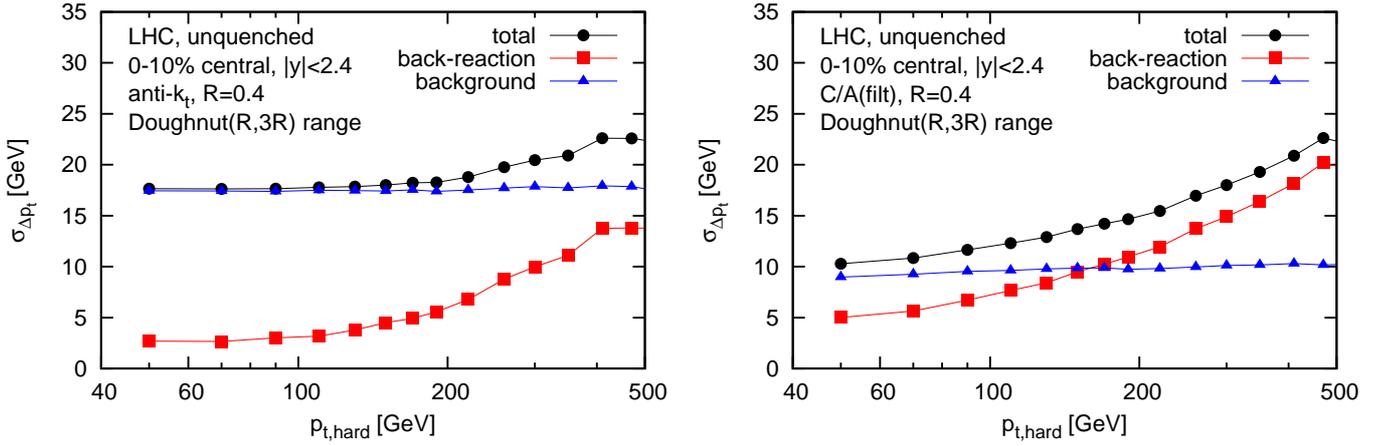


Figure 4.11: The decomposition of the dispersion into back-reaction and “background” components (including misestimation of ρ). The left-hand plot is for the anti- k_t algorithm and the right-hand one for C/A(filt). Both correspond to LHC collisions at $\sqrt{s_{NN}} = 5.5$ TeV.

fluctuations and misestimation of ρ (defined as $[\sigma_{\Delta p_t}^2 - (\sigma_{\Delta p_t}^{\text{BR}})^2]^{\frac{1}{2}}$). As already observed in pp collisions (see Section 3.2.4), one sees that the background-fluctuation component is essentially independent of p_t , while the back-reaction dispersion has a noticeable p_t dependence. The flatness of the residual, non-back-reaction, component is again observed and is expected since the anomalous dimension of the jet area is zero for anti- k_t and small for C/A (with or without filtering), and in any case leads to a weak scaling with p_t , as $\ln \ln p_t$. Furthermore, there is roughly a factor of $1/\sqrt{2}$ between anti- k_t and C/A(filt), as expected based on a proportionality of the dispersion to the square-root of the jet area.

4.4.2 Quality of area determination

One further source of Δp_t fluctuations can come from imperfect estimation of the area of the jets. We recall that throughout this article we have used soft ghosts, each with area of 0.01, in order to establish the jet area. That implies a corresponding finite resolution on the jet area and related poor estimation of the exact edges of the jets, which can have an impact on the amount of background that one subtracts from each jet, and, consequently, on the final dispersion. It is therefore interesting to see, figure 4.12, that the dispersion $\sigma_{\Delta p_t}$ is reduced by about 0.2 – 0.4 GeV if one lowers the ghost area to 0.0025.

While this doesn’t affect any of the conclusions of our paper, it does suggest that for a full experimental analysis there are benefits to be had from using a ghost area that is smaller than the default FastJet setting of 0.01. These observations are behind one the points made in our final recommendations, Section 2.8.

4.4.3 Relative importance of average shift and dispersion

To close this section, we examine the relative importance of the average shift and its dispersion, taking the illustrative example of their impact on the inclusive jet cross-section as a function of p_t . This follows our simple description introduced in Section 2.2 where we have seen that the true p_t

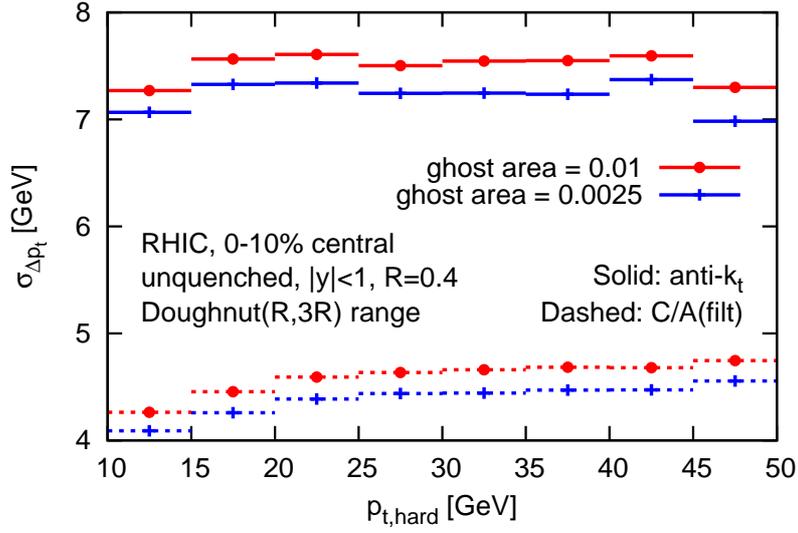


Figure 4.12: Dispersion, $\sigma_{\Delta p_t}$, for RHIC, as a function of the jet p_t , with two different choices for the ghost area, 0.01 and 0.0025.

spectrum decays exponentially i.e.

$$\frac{d\sigma^{pp}}{dp_t} = \sigma_0 \lambda e^{-\lambda p_t}. \quad (4.4)$$

The reconstructed spectrum, after embedding in a heavy-ion background and applying subtraction will be multiplied by a factor

$$\exp\left(\lambda \langle \Delta p_t \rangle + \frac{\lambda^2 \sigma_{\Delta p_t}^2}{2}\right), \quad (4.5)$$

as obtained in Eq. (2.8). For a given reconstructed p_t , the most likely original true transverse momentum is given by Eq. (2.10) and reads

$$\text{most likely } p_t^{pp} \simeq p_t^{AA, \text{sub}} - \langle \Delta p_t \rangle - \lambda \sigma_{\Delta p_t}^2, \quad (4.6)$$

where we have neglected the small impact of subtraction on the pp jets.

To illustrate these effects quantitatively, let us first take the example of RHIC, where between 10 and 60 GeV, the cross-section is well approximated by Eq. (4.4) with $\lambda = 0.3 \text{ GeV}^{-1}$. Both the anti- k_t and C/A(filt) have $\langle \Delta p_t \rangle \simeq 0$, leaving only the dispersion effect. In the case of the anti- k_t (respectively C/A(filt)) algorithm, we see from fig. 4.7 that $\sigma_{\Delta p_t} \simeq 7.5 \text{ GeV}$ (4.8 GeV), which gives a multiplicative factor of about 12 (3). For a given reconstructed p_t , the most likely true p_t is about 17 GeV (7 GeV) smaller. In comparison, for C/A (k_t), with $\langle \Delta p_t \rangle \simeq -1.5 \text{ GeV}$ (-3.5 GeV) (fig. 4.5) and $\sigma_{\Delta p_t} \simeq 6.5 \text{ GeV}$ (similar for k_t) there is a partial compensation between factors of 0.64 (0.35) and 6.7 coming respectively from the shift and dispersion, yielding an overall factor of about 4 (2.3), while the most likely true p_t is about 14 GeV (12 GeV) smaller than the reconstructed p_t .

At the LHC ($\sqrt{s_{NN}} = 5.5 \text{ TeV}$), Eq. (4.4) is a less accurate approximation. Nevertheless, for $p_t \sim 100 - 150 \text{ GeV}$, it is not too unreasonable to take $\lambda = 0.05 \text{ GeV}^{-1}$ and examine the consequences. For anti- k_t (respectively C/A(filt)), we have $\langle \Delta p_t \rangle \simeq 0$ (also for C/A(filt)) and $\sigma_{\Delta p_t} \simeq 18 \text{ GeV}$ (13 GeV), giving a multiplicative factor of 1.5 (1.2), i.e. far smaller corrections than at RHIC. For a given reconstructed p_t , the most likely true p_t is about 16 GeV (8 GeV) smaller, rather similar to

the values we found at RHIC (though smaller in relative terms, since the p_t 's are higher), with the increase in σ being compensated by the decrease in λ .

In the LHC case, it is also worth commenting on the results for the k_t algorithm, since this is what was used in the original work from Ref. [81]: we have $\langle \Delta p_t \rangle \simeq -8 \text{ GeV}$ and $\sigma_{\Delta p_t} \simeq 18 \text{ GeV}$, giving a multiplicative factor of 1.05, which is consistent with the near perfect agreement that was seen there between the pp and subtracted AA spectra. That agreement does not however imply perfect reconstruction, since the most likely p_t^{pp} is about 7 GeV lower than $p_t^{AA, \text{sub}}$.

Though the above numbers give an idea of the relative difficulties of using different algorithms at RHIC and LHC, experimentally what matters most will be the systematic errors on the correction factors (for example due to poorly understood non-Gaussian tails of the Δp_t distribution). Note also that a compensation between shift and dispersion factors, as happens for example with the C/A algorithm, is unlikely to reduce the overall systematic errors.

4.4.4 Fluctuations in extracted ρ

In section 4.2.2, we gave an estimate of the minimum size of a range one should require for determining ρ , given a requirement that fluctuations in the determination of ρ should be moderate. We give the details of the computation in this Section.

We start from the fact that the error made on the estimation of the background density ρ will translate into an increase of the dispersion $\sigma_{\Delta p_t}$: on one hand, the dominant contribution to $\sigma_{\Delta p_t}$ comes from the intra-event fluctuations of the background i.e. is of order $\sigma \sqrt{A_{\text{jet}}}$ with A_{jet} the jet area; on the other hand, the dispersion $S_{\Delta \rho}$ of the misestimation of ρ leads to an additional dispersion on the reconstructed jet p_t of $S_{\Delta \rho} A_{\text{jet}}$. As we will show in the next Chapter, $S_{\Delta \rho}$ can be estimated analytically, see Eq. (5.103). Adding these two sources of dispersion in quadrature and using (5.103), with $A_{\text{tot}} = A_{\mathcal{R}}$ the area of the range under consideration, we get

$$\sigma_{\Delta p_t} \simeq \sigma \left(A_{\text{jet}} + \frac{\pi A_{\text{jet}}^2}{2A_{\mathcal{R}}} \right)^{\frac{1}{2}}. \quad (4.7)$$

If we ask e.g. that the contribution to the total dispersion coming from the misestimation of the background be no more than a fraction ϵ of the total $\sigma_{\Delta p_t}$, then we obtain the requirement

$$A_{\mathcal{R}} \gtrsim A_{\text{min}} \simeq \frac{\pi A_{\text{jet}}}{4 \epsilon}. \quad (4.8)$$

For anti- k_t jets of radius R , with $A_{\text{jet}} \simeq \pi R^2$, this translates to

$$A_{\mathcal{R}} \gtrsim \frac{\pi^2 R^2}{4 \epsilon} \simeq 25 R^2, \quad (4.9)$$

where the numerical result has been given for $\epsilon = 0.1$. For $R = 0.4$, it becomes $A_{\mathcal{R}} \gtrsim 4$.

We can also cast this result in terms of the number of jets that must be present in \mathcal{R} . Assuming that the jets used to estimate ρ have a mean area of $0.55\pi R_\rho^2$,²⁰ we find a minimal number of jets,

$$n_{\text{min}} \simeq \frac{A_{\text{min}}}{0.55\pi R_\rho^2} \simeq \frac{1.4 R^2}{\epsilon R_\rho^2}, \quad (4.10)$$

where, as before, we have taken $A_{\text{jet}} \simeq \pi R^2$. Taking the numbers quoted above and $R_\rho = 0.5$, as used in the main body of the article, this gives $n_{\text{min}} \simeq 9$.

²⁰This is the typical area one would obtain using a (strongly recommended) jet definition like the k_t or C/A algorithms for the background estimation, see Section 5.1.2.

4.4.5 Hard-jet bias in extracted ρ

As we will show in our detailed analytic studies, see Section 5.2.3, the presence of hard jets and initial-state radiation leads to a bias in the extraction of ρ of

$$\langle \Delta\rho \rangle \simeq \sigma R_\rho \sqrt{\frac{\pi c_J}{2}} \frac{\langle n_h \rangle}{A_{\mathcal{R}}} \quad (4.11)$$

where $c_J \simeq 2$ is a numerical constant and $\langle n_h \rangle$ is the average number of “hard” jets (those above the scale of the background fluctuations, including initial-state radiation); $\langle n_h \rangle$ is given by

$$\frac{\langle n_h \rangle}{A_{\mathcal{R}}} \simeq \frac{n_b}{A_{\mathcal{R}}} + \frac{C_i L}{\pi^2 b_0}, \quad L = \ln \frac{\alpha_s(\sqrt{c_J} \sigma R_\rho)}{\alpha_s(p_t)}, \quad (4.12)$$

where n_b is the number of “Born” partons from the underlying $2 \rightarrow 2$ scattering that enter the region \mathcal{R} , while $b_0 = (11C_A - 2n_f)/(12\pi)$ is the first coefficient of the QCD β -function. Contrarily to the original case of UE studies in pp collisions [19], in the case of a large background like the UE in AA collisions, one can safely ignore the impact of the second term in (4.12). We thus arrive at the result

$$\langle \Delta\rho \rangle \simeq \sigma R_\rho \sqrt{\frac{\pi c_J}{2}} \frac{n_b}{A_{\mathcal{R}}} \simeq 1.8 \sigma R_\rho \frac{n_b}{A_{\mathcal{R}}}. \quad (4.13)$$

Note that the presence of hard jets and initial-state radiation also affects the fluctuations in the misestimation of ρ and this should in principle have been included in the estimates of Section 4.4.4. However, while the effect is not completely negligible, to within the accuracy that is relevant for us (a few tens of percent in the estimation of a minimal $A_{\mathcal{R}}$) it does not significantly alter the picture outlined there.

4.4.6 Subtraction bias due to filtering

We have seen from fig. 4.5 in section 4.3.3 that the subtraction differs when we use the C/A algorithm with and without filtering. Since this difference is not due to back-reaction (see fig. 4.6), it has to be due to the subtraction itself.

The difference comes from a bias introduced by the selection of the two hardest subjects during filtering. The dominant contribution comes when only one subject, that we shall assume harder than all the others, contains the hard radiation, all the other subjects being pure background. In that case, the selection of the hardest of these pure-background subjects as the second subject to be kept tends to pick positive fluctuations of the background. This in turn results in a positive offset compared to pure C/A clustering, as observed in section 4.3.3.

To compute the effect analytically, let us thus assume that we have one hard subject and n_{bkg} pure-background subjects of area $\mathcal{A}_g = 0.55\pi R_{\text{filt}}^2$ (see Section. 5.1.2). After subtraction, the momentum of each of the pure-background subject can be approximated as having a Gaussian distribution of average zero and dispersion $\sigma\sqrt{\mathcal{A}_g}$. Assuming that the “hard” subject’s transverse momentum remains larger than that of all the background jets, the 2 subjects that will be kept by the filter are the hard subject (subtracted) and the hardest of all the subtracted background jets. The momentum distribution of the latter is given by the maximum of the n_{bkg} Gaussian distributions.²¹ We are only interested here

²¹Within `FastJet`’s filtering tools, when the subtracted transverse momentum of a subject is negative, the subject is assumed to be pure noise and so discarded. This means that the momentum distribution of the hardest subtracted background jet is really given by the distribution of the maximum of the n Gaussian-distributed random numbers, but with the result replaced by zero if all of them are negative. In the calculations here we ignore this subtlety, since we will have $n = 3$ and only $1/8^{\text{th}}$ of the time are three Gaussian-distributed random numbers all negative.

in computing the average bias introduced by the filtering procedure, which is then given by

$$\langle(\Delta p_t)_{\text{filt}}\rangle \simeq \int \prod_{k=1}^{n_{\text{bkg}}} \left(dp_{t,k} \frac{1}{\sqrt{2\pi\mathcal{A}_g\sigma}} e^{-\frac{p_{t,k}^2}{2\mathcal{A}_g\sigma^2}} \right) \max(p_{t,1}, \dots, p_{t,n_{\text{bkg}}})$$

For the typical case $R_{\text{filt}} = R/2$ and $n_{\text{bkg}} = 3$, one finds

$$\langle(\Delta p_t)_{\text{filt}}\rangle \simeq \frac{3\sqrt{\mathcal{A}_g}\sigma}{2\sqrt{\pi}} \simeq 0.56 R\sigma. \quad (4.14)$$

If we insert in that expression the typical values for the fluctuations quoted in section 4.3 and $R = 0.4$, we find average biases of 2 GeV for RHIC and 4.5 GeV at the LHC, which are in good agreement with the differences observed between C/A with and without filtering in fig. 4.5.

4.5 The issue of fakes

While the goal of this paper is not to discuss the issue of “fake-jets” in detail, it is a question that has been the subject of substantial debate recently (see for example [67, 82]). Here, therefore, we wish to devote a few words to it and discuss how it relates to our background-subtraction results so far.

In a picture in which the soft background and the hard jets are independent of each other, one way of thinking about a fake jet is that it is a reconstructed jet (with significant p_t) that is due not to the presence of an actual hard jet, but rather due to an upwards fluctuation of the soft background. The difficulty with this definition is that there is no uniquely-defined separation between “hard” jets and soft background. This can be illustrated with the example of how Hydjet simulates RHIC collisions: one event typically consists of a soft HYDRO background supplemented with ~ 60 pp collisions, each simulated with a minimum p_t cut of 2.6 GeV on the $2 \rightarrow 2$ scattering. To some approximation, the properties of the full heavy-ion events remain relatively unchanged if one modifies the number of pp collisions and corresponding p_t cut and also retunes the soft background. The fact that this changes the number of hard jets provides one illustration of the issue that the soft/hard separation is ill-defined. Additionally, while there are ~ 60 semi-hard pp collisions (~ 120 mostly central semi-hard jets²²) in an event, there is only space within (say) the acceptance of RHIC for $\mathcal{O}(40)$ jets. Thus there is essentially no region in an event which does not have a semi-hard jet. From this point of view, every reconstructed jet corresponds to a (semi-)hard pp jet and there are no fake jets at all.

4.5.1 Inclusive analyses

For inclusive analyses, such as a measurement of the inclusive jet spectrum, this last point is particularly relevant, because every jet in the event contributes to the measurement. Then, the issue of fakes can be viewed as one of unfolding. In that respect it becomes instructive, for a given bin of the reconstructed heavy-ion p_t , to ask what the corresponding matched pp jet transverse momenta were. Specifically, we define a quantity $O(p_t^{AA,\text{sub}}, p_t^{pp})$, the distribution of the pp “origin”, p_t^{pp} , of a heavy-ion jet with subtracted transverse momentum $p_t^{AA,\text{sub}}$. If the origin $O(p_t^{AA,\text{sub}}, p_t^{pp})$ is dominated by a region of p_t^{pp} of the same order as $p_t^{AA,\text{sub}}$, then that tells us that the jets being reconstructed are truly hard. If, on the other hand, it is dominated by p_t^{pp} near zero, then that is a sign that apparently hard heavy-ion jets are mostly due to upwards fluctuations of the background superimposed on low- p_t pp jets, making the unfolding more delicate.

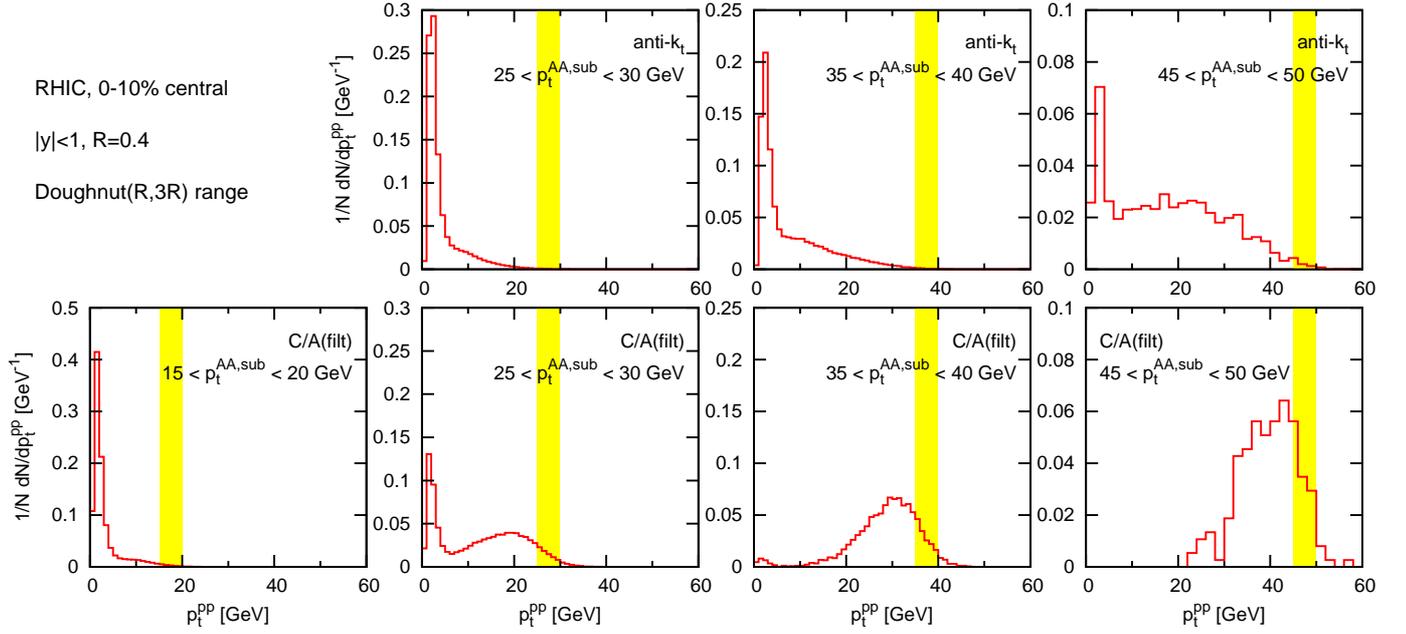


Figure 4.13: The p_t distribution of the pp jet corresponding to a given bin of reconstructed heavy-ion jet $p_t^{AA,sub}$ at RHIC, i.e. $O(p_t^{AA,sub}, p_t^{pp})$ as a function of p_t^{pp} for a given bin of $p_t^{AA,sub}$. The upper row is for the anti- k_t algorithm, while the lower row is for C/A(filt). Each column corresponds to a different $p_t^{AA,sub}$ bin, as indicated by the vertical band in each plot. Cases in which the histogram is broad or peaked near 0 are indicative of the need for special care in the unfolding procedure. These plots were generated using approximately 90 million events. Each plot has been normalised to the number of events in the corresponding $p_t^{AA,sub}$ bin.

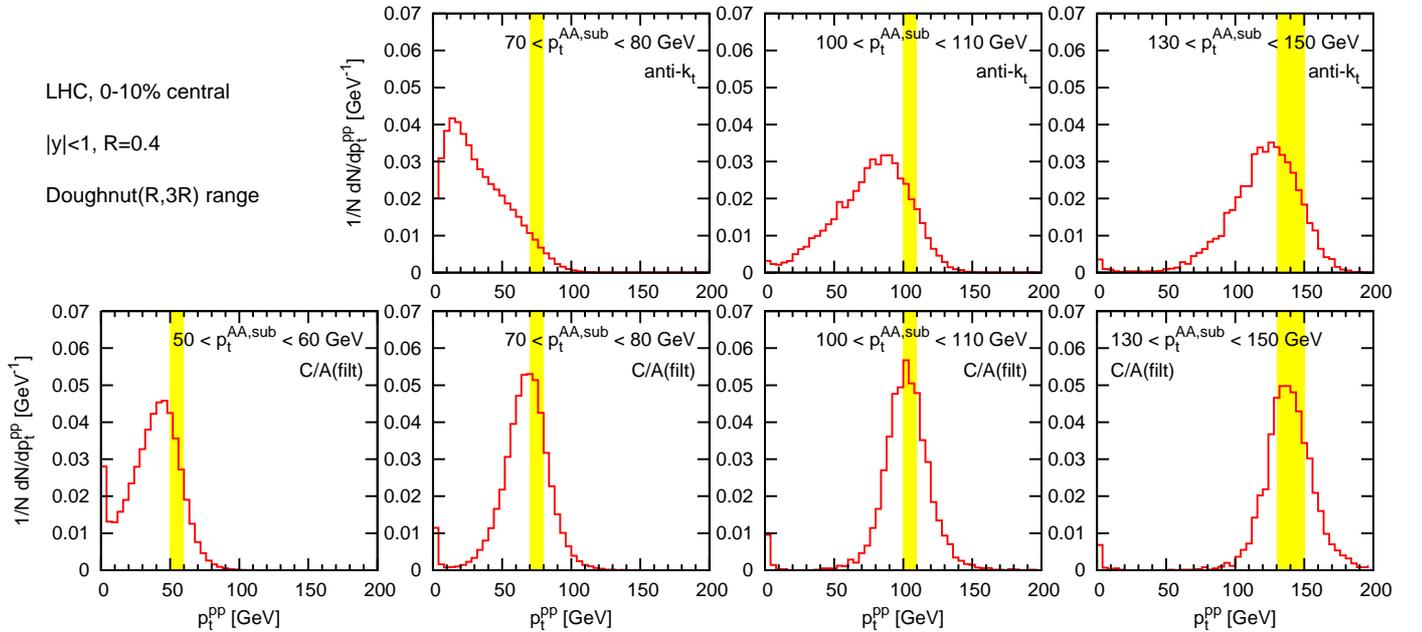


Figure 4.14: Same as fig. 4.13 for LHC kinematics (PbPb, $\sqrt{s} = 5.5$ TeV), generated with approximately 16 million events. Note the use of a smaller rapidity range here, $|y| < 1$, compared to the earlier LHC plots.

In figure 4.13 we show the origins of heavy-ion jets as determined in our Hydjet simulations.²³ The upper row provides the origin plots for anti- k_t jets at RHIC. Each plot corresponds to one bin of $p_t^{AA,sub}$, and shows $O(p_t^{AA,sub}, p_t^{pp})$ as a function of p_t^{pp} . At moderate $p_t^{AA,sub}$, the 25 – 30 GeV bin, the origin is dominated by low p_t^{pp} . This is perhaps not surprising, given the result in subsection 4.4.3 that the p_t^{pp} origin is expected to be ~ 17 GeV lower than $p_t^{AA,sub}$ for anti- k_t jets — additionally, that result assumed an exponential spectrum for the inclusive jet distribution, whereas the distribution rises substantially faster towards low p_t^{pp} . As $p_t^{AA,sub}$ increases one sees that the contribution of high p_t^{pp} jets increases, in a manner not too inconsistent with the expected ~ 17 GeV shift, though the p_t^{pp} distribution remains rather broad and a peak persists at small p_t^{pp} . These plots suggest that an inclusive jet distribution measurement with the anti- k_t algorithm at RHIC is not completely trivial since, up to rather large $p_t^{AA,sub}$, one is still sensitive to the jet distribution at small values of p_t^{pp} where the separation between “hard” jets and the soft medium is less clear. Nevertheless, two points should be kept in mind: firstly, the upper row of fig. 4.13 shows that different $p_t^{AA,sub}$ have complementary sensitivities to different parts of the p_t^{pp} spectrum. Thus it should still be possible to “unfold” the $p_t^{AA,sub}$ distribution to obtain information about the p_t^{pp} , unfolding being in any case a standard part of the experimental correction procedure. Secondly STAR quotes [68] a 10% smaller value for $\sigma_{\Delta p_t}$ than the 7.5 GeV that we find in Hydjet. Such a reduction can make it noticeably easier to perform the unfolding.

The impact of a reduction in $\sigma_{\Delta p_t}$ is illustrated in the lower row of fig. 4.13, which shows the result for C/A(filt). Here, even the 25 – 30 GeV bin for $p_t^{AA,sub}$ shows a moderate- p_t peak in the distribution of p_t^{pp} , and in the 35 – 40 GeV bin the low- p_t “fake” peak has disappeared almost entirely. Furthermore, the p_t^{pp} peak is centred about 7 GeV lower than the centre of the $p_t^{AA,sub}$ bin, remarkably consistent with the calculations of section 4.4.3. Overall, therefore, unfolding with C/A(filt) will be easier than with anti- k_t .

Corresponding plots for the LHC are shown in fig. 4.14. While C/A(filt)’s lower dispersion still gives it an advantage over anti- k_t , for $p_t^{AA,sub} \gtrsim 80$ GeV, anti- k_t does now reach a domain where the original pp jets are themselves always hard.

Procedures to reject fake jets have been proposed, in [83, 84]. They are based on a cut on (collinear unsafe) jet shape properties and it is thus unclear how they will be affected by quenching and in particular whether the expected benefit of cutting the low- p_t^{pp} peak in Figs. 4.13 and 4.14 outweighs the disadvantage of potentially introducing extra sources of systematic uncertainty at moderate p_t .

One final comment is that experimental unfolding should provide enough information to produce origin plots like those shown here. As part of the broader discussion about fakes it would probably be instructive for such plots to be shown together with the inclusive-jet results.

4.5.2 Exclusive analyses

An example of an exclusive analysis might be a dijet study, in which one selects the two hardest jets in the event, with transverse momenta p_{t1} and p_{t2} , and plots the distribution of $\frac{1}{2}H_{T,2} \equiv \frac{1}{2}(p_{t1} + p_{t2})$. Here one can define “fakes” as corresponding to cases where one or other of the jets fails to match to one of the two hardest among all the jets from the individual pp events. This definition is insensitive

²²Specifically, keeping in mind the Hydjet simulation, one can cluster each pp event separately to obtain a long list of pp jets from all the separate hard events.

²³A point to be aware of is that multiple pp jets can match a single heavy-ion jet, i.e. have at least half their p_t contained in the heavy-ion jet. In evaluating $O(p_t^{AA,sub}, p_t^{pp})$ we take only the highest- p_t matched jet. If there is no matched jet (this occurs only rarely) then we fill the bin at $p_t^{pp} = 0$. Note also that since we are not explicitly embedding hard jets, all pp jets in the events have undergone Hydjet’s quenching.

to the soft/hard boundary in a simulation such as Hydjet, because it naturally picks out hard pp jets that are far above that boundary. And, by concentrating on just two jets, it also evades the problem of high occupancy from the large multiplicity of semi-hard pp collisions. This simplification of the definition of fakes is common to many exclusive analyses, because they tend to share the feature of identifying just one or two hard reference jets.

The specific case of the exclusive dijet analysis has the added advantage that it is amenable to a data-driven estimation of fakes. One divides the events into two groups, those for which the two hardest jets are on the same side (in azimuth) of the event and those in which they are on opposite sides (a related analysis was presented by STAR in ref. [85]). For events in which one of the two jets is “fake,” the two jets are just as likely to be on the same side as on the opposite side. This is not the case for non-fake jets, given that the two hardest “true” jets nearly always come from the same pp event and so have to be on opposite sides.²⁴ Thus by counting the number of same-side versus opposite-side dijets in a given $\frac{1}{2}H_{T,2}$ bin, one immediately has an estimate of the fake rate.²⁵

This is illustrated in fig. 4.15, which shows the distribution at RHIC of the full, subtracted $\frac{1}{2}H_{T,2}$ result, together with its separation into opposite-side and same-side components. One sees that in the peak region, the opposite and same-side distributions are very similar, indicating a predominantly “fake” origin for at least one of two hardest jets (they are not quite identical, because there is less phase-space on the same side for a second jet than there is on the away side). However above a certain full, reconstructed $\frac{1}{2}H_{T,2}$ value, about 30 GeV for anti- k_t and 20 GeV for C/A(filt) the same-side distribution starts to fall far more rapidly than the opposite-side one, indicating that the measurement is now dominated by “true” pairs of jets.

The LHC results, fig. 4.16, are qualitatively similar, with the same-side spectrum starting to fall off more steeply than the opposite-side one around 70 GeV for the anti- k_t algorithm and 50 GeV for C/A(filt).

One can also examine origin plots for $H_{T,2}$, in analogy with the Monte Carlo analysis of section 4.5.1. For brevity, we refrain from showing them here, and restrict ourselves to the comment that in the region of $H_{T,2}$ where the result is dominated by opposite-side pairs, the origin plots are consistent with a purely hard origin for the dijets.

4.6 Summary and discussions

We have shown that the area–median subtraction method gives good performance to subtract the large heavy-ion UE. We have tested different choices of range to estimate the rapidity dependence and observed little difference between various ranges, as long as they are chosen to be localised in the vicinity of the jet of interest and of sufficient size (at least 4 units of area for jets with $R = 0.4$).

As for the quality measures, the offset can be brought close to zero by using the anti- k_t algorithm, while the k_t algorithm has the largest offset; the Cambridge/Aachen (C/A) algorithm with filtering also gives a small offset, however this seems to have been due to a fortuitous cancellation between two only partially related effects. The dispersion is comparable for anti- k_t , C/A and k_t , but significantly smaller for C/A(filt) (except at high transverse momenta for LHC), as a consequence of its smaller jet

²⁴At RHIC energies, above $p_t \sim 10 - 15$ GeV, it is nearly always the case that the two hardest jets come from the same pp event. At the LHC, this happens above 20 – 30 GeV.

²⁵Note that for plain pp events, if one has only limited rapidity acceptance then the same-side/opposite-side separation is not infrared safe, because of events in which only one hard jet is within the acceptance and the other “jet” is given by a soft gluon emission. Thus to examine the same-side/opposite-side separation in plain pp events with limited acceptance, one would need to impose a p_t cut on the second jet, say $p_{t,2} > \frac{1}{2}p_{t,1}$.

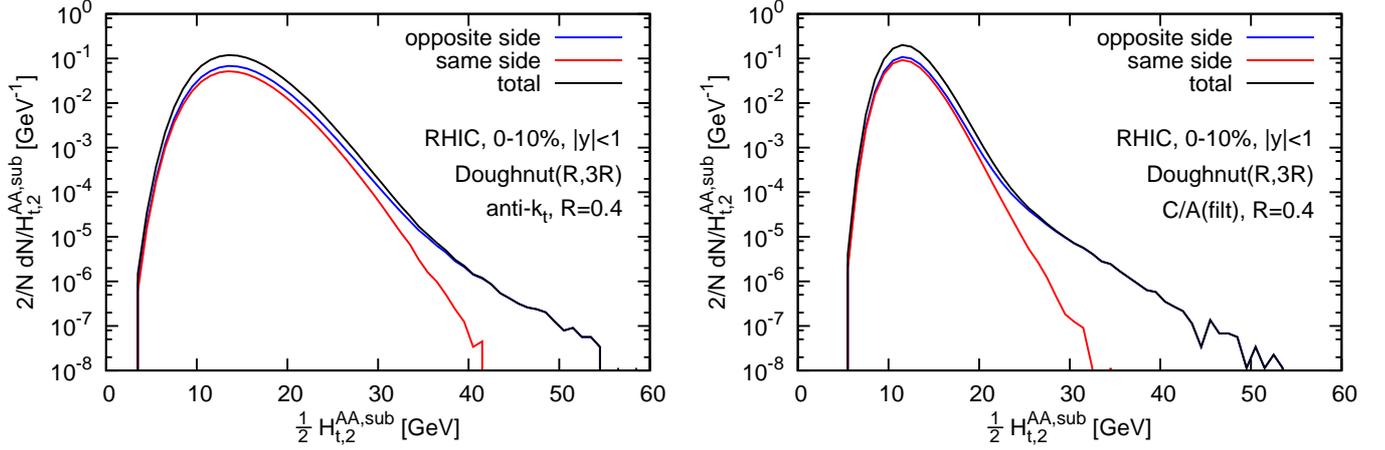


Figure 4.15: The distribution of $\frac{1}{2}H_{T,2}$ obtained from the two hardest full, subtracted heavy-ion jets in each event at RHIC, as obtained from simulations with Hydjet 1.6. The left-hand plot is for anti- k_t and the right-hand plot for C/A(filt). The same-side curves give an approximate measure of (half of) the contribution of “fake” jets to the dijet spectrum.

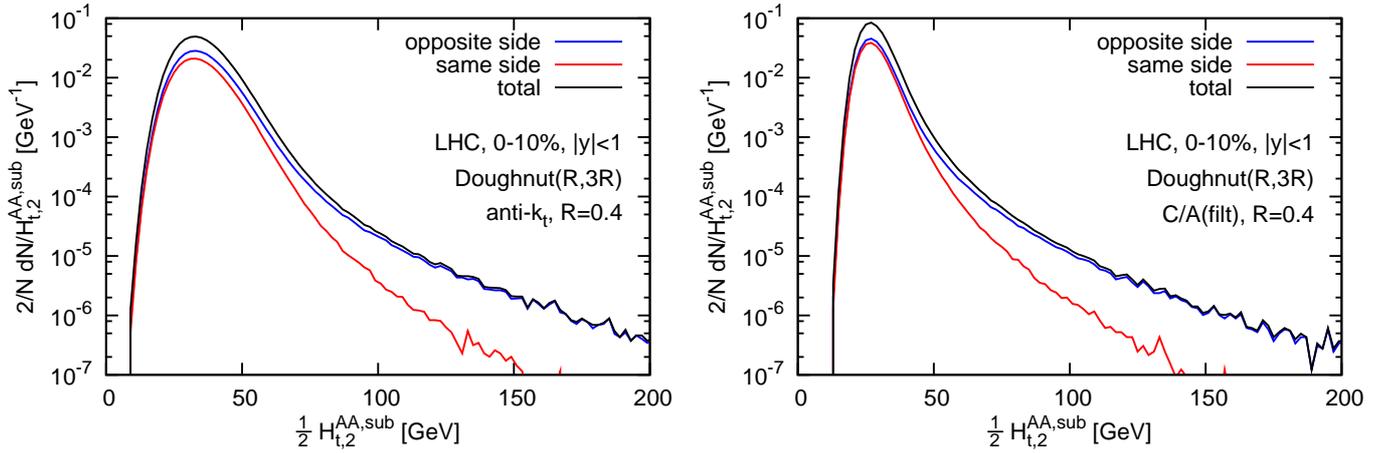


Figure 4.16: Same as fig. 4.15 for LHC kinematics.

area. Among the different algorithms, anti- k_t is the most robust with respect to quenching effects, and C/A(filt) seems reasonably robust at RHIC, though a little less so at the LHC. The precise numerical results for offset and dispersion can depend a little on the details of the simulation and the analysis, however the general pattern remains.

Overall our results indicate that the area-based subtraction method seems well suited for jet reconstruction in heavy-ion collisions. Two jet-algorithm choices were found to perform particularly well: anti- k_t , which has small offsets but larger fluctuations, and C/A with filtering, for which the offsets may be harder to control, but for which the fluctuations are significantly reduced, with consequent advantages for the unfolding of experimentally measured jet spectra. Ultimately, we suspect that carrying out parallel analyses with these two choices may help maximise the reliability of jet results in HI collisions.

4.7 Jet fragmentation function

The last extension of the area–median approach we want to validate is its application to the subtraction of pileup contamination to the fragmentation function (moments), as described in Section 2.7.4. We shall apply it here to heavy-ion collisions, where it is considered as one of the first observables that one would measure to study jet-quenching effects [86, 87, 88, 89, 90, 91, 92].²⁶

Details of the simulation. Before proceeding with the results, it is useful to detail the terminology and simulation tools that we have used.

What we call ‘hard’ jets in QCD are simulated in proton-proton (pp) collisions using Pythia 6.425 [35] in dijet mode with the DW tune [93].²⁷ We have not included any pp underlying event (UE), its effect being minimal in this context anyway. Jets including quenching effects have been generated using Pyquen, v1.5.1, as included in Hydjet [78, 63]. We consider all hadrons, not just charged tracks, and take π_0 ’s to be stable, so that we are still considering genuine hadron distributions.

The heavy-ion background (also called underlying event) is simulated using Hydjet v1.6 for 0-10% centrality and, where needed, it is superimposed to the hard event generated by Pythia. This has been found [60] to be in good agreement with the experimental measurements of background fluctuations by ALICE [94]. The jets observed in this combined event will be denoted as ‘full’ jets.

All events are generated for either pp or lead-lead (PbPb) collisions at the LHC, with a centre-of-mass energy of 2.76 TeV per nucleon-nucleon collision. Jets are reconstructed using the anti- k_t algorithm with $R = 0.4$, as implemented in FastJet. In estimating the jet’s p_t , the HI background is subtracted using the area–median method using k_t jets with $R = 0.4$ to estimate ρ . At most the two hardest jets passing a hard cut on the subtracted transverse momentum (see below) are subsequently used for the fragmentation function analysis. The cuts that we shall use are 100 GeV and 200 GeV.

Effect of the background. The addition of the heavy-ion background has the potential to modify a measured fragmentation function in two ways. Firstly, the jet’s p_t is modified, affecting the normalisation of z in Eq. (2.41). Secondly, the heavy-ion background adds many extra particles to the jet, predominantly at low momenta.

To appreciate the impact of the extra particles in the jet from the heavy-ion background, it is instructive to first examine the ‘‘Pythia+Hydjet’’ dashed red curves of Fig. 4.17. These show the FF extracted in heavy-ion collisions, without any subtraction of the background contribution to the FF, but always using a z value defined such that the jet’s p_t has been corrected for the expected HI background contamination (as is standard in the experimental measurements). At this stage we will not perform any unfolding to account for fluctuations in the HI background.

One sees how the FF acquires a ‘bump’ in the soft region, which lies at larger ξ (smaller z) than the maximum in the original pp result (blue solid line) and is up to two orders of magnitude higher. For the moments, this presence of the background is seen as a steep increase in the small- N region, taking the curve far off the scale.

At large z (small ξ), the impact of the addition of the background is only barely visible, as might be expected given that it is dominated by soft particles. However, this visually small effect is partially an artefact of the logarithmic scale used to show the FF versus z . Considering instead the moments,

²⁶The generic underlying idea is that jet quenching generates additional medium-induced radiation [64]. One therefore expects a reduction of the fragmentation function at large z and an increase at small z . In terms of the moments of the fragmentation, this means an increase for $N < 1$, including e.g. for multiplicity ($N = 0$), and a decrease for $N > 1$.

²⁷This tune is not the most up to date; however it is not unrealistic for the LHC and has the characteristic that it is based on Pythia’s virtuality ordered shower, which is a prerequisite for use with Pyquen.

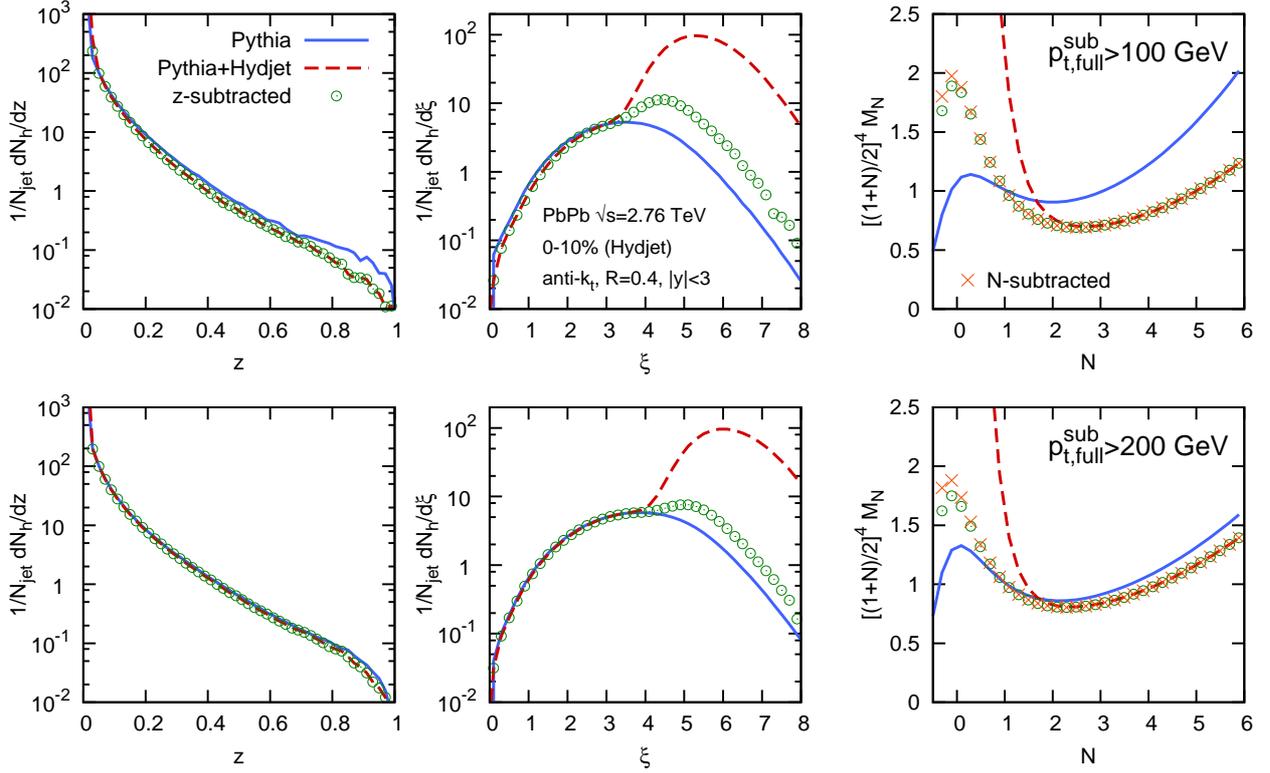


Figure 4.17: Jet fragmentation functions shown for plain Pythia, with the addition of the heavy-ion background (Pythia+Hydjet) and after subtraction of the heavy-ion background (z -subtracted and N -subtracted). For the results including the heavy-ion background, the jet p_t used to define z is always that after subtraction of the heavy-ion background. As in Fig. B.1, we show the results as a function of z (left), ξ (middle) and for the moments versus N (right). The upper (lower) row has a jet p_t threshold of 100 GeV (200 GeV).

one sees that there is a non-negligible *reduction* in the FF at large N . This is perhaps surprising given that the background adds particles. It is a feature related to the interplay between background fluctuations and the steeply falling jet spectrum. It is well known by the experiments and we will return to discuss it in section 2.6.

z -space subtraction. The traditional approach to background subtraction from a jet FF involves the construction of a distribution (in z or ξ space) intended to approximate that of background-only particles, and the subtraction of this distribution from the measured one. The way the background-only distribution is determined can vary, the simplest one probably being to consider a region of the event that is expected to be little affected by the hard jets, and measure it there.

To illustrate z -space subtraction here, we measure the distribution of hadrons in two regions transverse in azimuth with respect to the axis defined by the dijet event. Event by event, and jet by jet, we subtract those distributions (measured in a patch of phase space with an area equal to that of the jet that one is considering) from the jet fragmentation function. While exact experimental procedures differ in the details, most choices lead to similar results here. Perhaps the main distinction of the experimental procedures is that they sometimes address issues related to flow, which for simplicity we neglect here in our z -space subtraction.

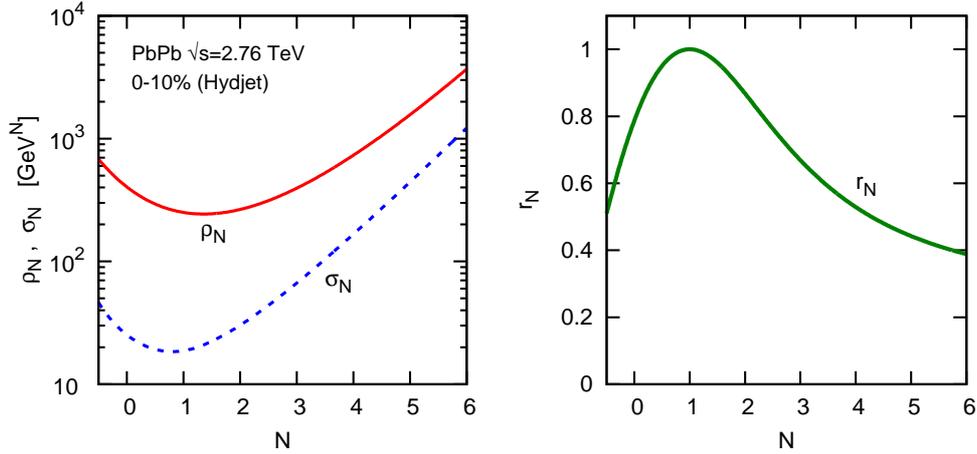


Figure 4.18: The quantities ρ_N , σ_N and the correlation coefficient r_N , shown as a function of N for 0 – 10% central Pb Pb collisions $\sqrt{s_{NN}} = 2.76$ TeV as obtained from simulations with Hydjet.

The results of applying a z -space subtraction are shown in figure 4.17 as green open circles, labelled ‘ z -subtracted’ (the moments of these z -subtracted results are also shown). One sees how these curves come closer to the pp results at small z and small N than do the unsubtracted (dashed red) ones. The agreement improves as the jet p_t threshold is increased. However, even with jets of $p_t \gtrsim 200$ GeV, this procedure still falls short of an accurate reconstruction of the hard FF. In fact, the region in z where the subtracted FF works well barely extends beyond that selected by simply truncating the unsubtracted FF at low $p_{t,h}$ so as to avoid the region dominated by the background.

Area–median moments subtraction. The results for M_N^{sub} are shown in figure 4.17 in the right-most plots, orange crosses labelled ‘ N -subtracted’. One can see that they are neither better nor worse than the corresponding z -subtracted ones, which have been translated to N -space and drawn in the same plots for direct comparison (green circles).

In order to provide a fragmentation function reconstruction that is better than that given by the standard z -space method, we introduce in the next section an improved background subtraction method that is most straightforwardly applied in moment space.

Area–median improved moments subtraction. The correction (2.55) can be applied jet by jet to correct for the effects of fluctuations. It requires the prior knowledge of the slope μ of the jet cross-section, which can be obtained from pp data, or from simulations.²⁸ All the other ingredients that enter this equation (σ , A , S_N , σ_N , r_N) can instead be determined event-by-event or jet-by-jet. In practice we determine σ , σ_N and r_N from the ensemble of jets contained in an annulus (or “doughnut”) of outer radius $3R$ and inner radius R , centred on the jet of interest. Typical values of ρ_N , σ_N and r_N are presented as a function of N in Fig. 4.18.²⁹

²⁸In practice, μ depends on p_t and should be taken at the scale $S_1 - q_t$ in the integrand. However, μ varies slowly with p_t and can easily be taken at the fixed scale p_t in our small-fluctuations limit. In our analysis μ ranged from ~ 9 GeV at $p_t \simeq 50$ GeV to ~ 28 GeV at $p_t \simeq 200$ GeV. At $p_t \simeq 100$ GeV we had $\mu \sim 16$ GeV.

²⁹The full specification of our procedure is as follows: ρ , ρ_N are evaluated using Eq. (2.44), where the “patches” correspond to the jets in the annulus; σ_N (and $\sigma \equiv \sigma_1$) is as returned by FastJet on the same set of jets, i.e. derived from the difference between the 16th percentile and the median of the set of $\frac{\sum_{i \in \text{jet}} p_{t,i}^N}{A_{\text{jet}}}$. For the last term in Eq. (2.55) (and also for Fig. 4.18), r_N is evaluated using Eq. (2.50) as determined on this set of jets, with Q_N for a given jet evaluated as $\sum_{i \in \text{jet}} p_{t,i}^N - \rho_N A_{\text{jet}}$ (and $q_t \equiv Q_1$), which corresponds to Eq. (2.49) in the approximation that the jets

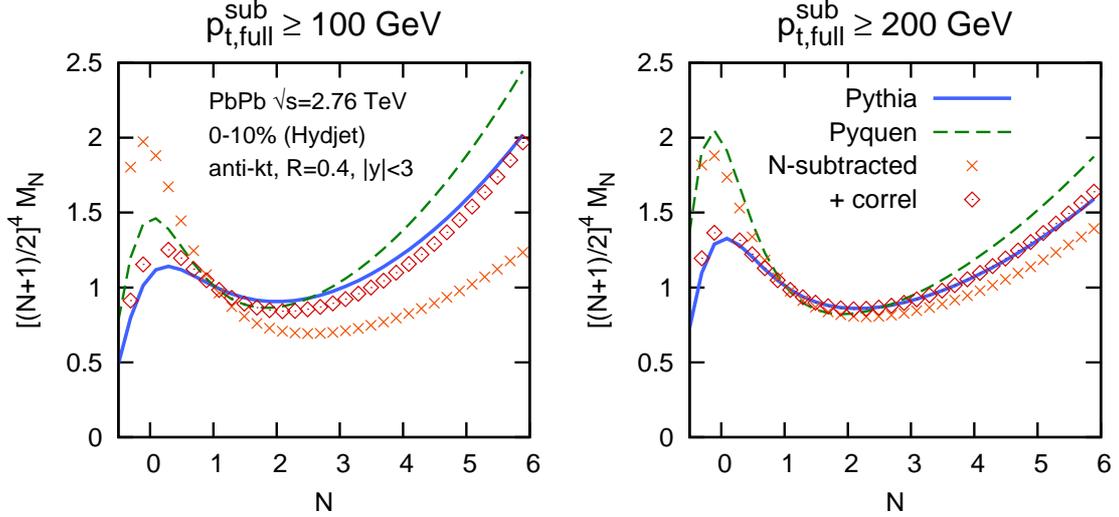


Figure 4.19: Jet fragmentation function moments, showing the plain Pythia result, the result after embedding in Hydjet and applying plain subtraction moment-space subtraction (“N-subtracted”) and after the additional improvement to account for correlations (“+ correl”), Eq. (2.55). A quenched result (“Pyquen”) is also shown, to help give an indication of the order of magnitude of quenching effects as compared to residual misreconstruction effects.

We show in figure 4.19 the result of applying Eq. (2.55) to our subtraction in moment space, for two jet p_t thresholds. The solid blue curve (pp reference) and orange crosses (N-subtracted) are identical to the results in the rightmost plots in figure 4.17. In addition, figure 4.19 also displays results obtained using Eq. (2.55), shown as red diamonds. One sees how the quality of the agreement with the ‘hard’ blue curve is markedly improved. At low N it is the last, additive, term in Eq. (2.55) that dominates this improvement, accounting for the correlation between a jet’s reconstructed p_t fluctuations and the fluctuations in the background’s contribution to the moment; at high N it is the multiplicative $N\sigma^2 A/S_1\mu$ term that dominates, correcting for the fact that the more common upwards fluctuations in the jet’s reconstructed p_t cause the fragmentation z value to be underestimated.

Figure 4.19 also shows (green dashed curve) the jet fragmentation function predicted by the quenching model used in Pyquen. One observes that the remaining deficiencies of the reconstruction are significantly smaller than the difference between the unquenched (solid blue) and the quenched (dashed green) FFs, pointing to a potential discriminating power. This is to be contrasted with subtraction without improvement (orange crosses), which especially in the soft region, $N < 1$, fails to describe the blue curve sufficiently well to tell whether quenching (as modelled in Pyquen) is present or simply if an imperfect reconstruction is taking place. This serves as an illustration that the improved subtraction may now be sufficiently good to allow one to discriminate a quenched FF from an unquenched one.

An implementation of the tools needed for the jet fragmentation function subtraction in moments space, as well as for the fluctuations unfolding improvement, is available in [fjcontrib](#).

in the annulus are predominantly composed of background particles. We also investigated an alternative procedure, in which the last term in Eq. (2.55) was evaluated as $\frac{\text{Cov}(q_t, Q_N)}{\text{Var}(q_t)} \times \frac{\sigma^2 A}{\mu S_1^N}$, i.e. directly combining Eqs. (2.51) and (2.54), with the ratio of covariance and variance evaluated using the usual set of jets in the annulus. In the region of moderate N , where the last term of Eq. (2.55) contributes significantly to the correction for M_N , this alternative procedure gives nearly identical results to our main procedure. However, for large N values, we have found that it gives a poor evaluation of the last term itself, though overall that term contributes negligibly there.

Chapter 5

Analytic insight

In this chapter we shall review a series of analytic properties related to the area–median pileup subtraction method. These are meant to serve two main purposes: first, jets being QCD objects, there is a clear interest to see if their properties are under analytic control in terms of perturbative QCD; second, we will see that simple analytic estimates allow to grasp a large fraction of the features observed in Monte Carlo simulations.

As far as we are willing to understand QCD aspects of jets and their properties, a large part of this chapter can be seen as a turf for theoretical QCD aficionados with little phenomenological consequences. Nevertheless, the results we obtain often allow for a deeper understanding of the dynamics of jets and can be beneficial in various respects.¹ In the following pages, we shall thus give an extended discussion of the two cornerstones of the area–median method by discussing the analytic properties of jet areas (Section 5.1) and an analytic understanding of the estimation of the pileup energy density ρ (see Section 5.2).

On the second front, there is a long list of cases where simple analytic understanding can help understanding the underlying physics and build a coherent picture of pileup mitigation, as we have already seen a few times in the previous chapters. These often rely on a Gaussian approximation for the pileup p_t distribution, simplified distributions for the underlying hard processes and basic understanding of jet areas. We shall provide strong bases for these arguments throughout this chapter.

Finally, as we shall see for example in Section 5.1.6, our analytic investigations open a series of doors towards future developments.

5.1 Properties of jet areas

In the context of pileup subtraction, we are mostly interested in the active jet area. However, passive areas are much simpler to handle since they do not involve the clustering of ghosts — it is tempting, here, to draw an analogy with QED v. QCD where, perturbatively, the latter is very often simpler to manage than the former because photons do not interact while gluons do. We shall therefore begin with a study of the properties of passive areas.

That said, our goal is to understand the properties of jet areas in perturbative QCD (pQCD). To do that, we start by considering jets made of a single particle, carry on with cases with 2 particles and try, when possible, to extrapolate to multi-particle configurations.

¹For example, the analytic understanding of jet areas has been a crucial element in devising the anti- k_t algorithm.

5.1.1 Passive areas

5.1.1.1 Areas for 1-particle configurations

We consider passive areas in the context of four different jet algorithms: (inclusive) anti- k_t [22], (inclusive) k_t [25], (inclusive) Cambridge/Aachen [28], and the seedless infrared-safe cone algorithm, SIScone [95]. Each of these jet algorithms contains a parameter R which essentially controls the reach of the jet algorithm in y and ϕ . Given an event made of a single particle p_1 , the passive area of the jet J_1 containing it is $a(J_1) = \pi R^2$ for all four algorithms.

5.1.1.2 Areas for 2-particle configurations

Let us now consider what happens to the passive jet areas in the presence of additional soft perturbative radiation. We add a particle p_2 such that the transverse momenta are strongly ordered,

$$p_{t1} \gg p_{t2} \gg \Lambda_{QCD} \gg g_t, \quad (5.1)$$

and p_1 and p_2 are separated by a geometrical distance $\Delta_{12}^2 = (y_1 - y_2)^2 + (\phi_1 - \phi_2)^2$ in the y - ϕ plane. Later, we shall integrate over Δ_{12} and p_{t2} . Note that g_t has been taken to be infinitesimal compared to all physical scales to ensure that the presence of the ghost particle does not affect the real jets.

For $\Delta_{12} = 0$ collinear safety ensures that the passive area is still equal to πR^2 for all three algorithms. However, as one increases Δ_{12} , each algorithm behaves differently.

Anti- k_t . Let us first consider the anti- k_t jet algorithm which has the 2-particle distance measure $d_{ij} = \min(k_{ti}^{-2}, k_{tj}^{-2}) \Delta_{ij}^2 / R^2$ and beam-particle distance $d_{iB} = k_{ti}^{-2}$. Taking $\Delta_{12} \sim \Delta_{1g} \sim \Delta_{2g} \sim R$ and exploiting the strong ordering (5.1) one has

$$d_{1B} \sim d_{12} \sim d_{1g} \ll d_{2B} \sim d_{2g} \ll d_{gB}. \quad (5.2)$$

From this ordering of the distances, one sees that the first step of the clustering will first combine p_1 with the closest of p_2 or the ghost, provided it is within a distance R of p_1 . p_1 will then cluster with the other of p_2 or the ghost, again provided it is within a distance R of p_1 . The ghost will therefore be clustered with p_1 if and only if they are within a distance R of p_1 , independently of the presence of the soft particle p_2 . The passive area of the jet which will thus remain πR^2 .

k_t . Let us now turn to the k_t jet algorithm which has the distance measures $d_{ij} = \min(k_{ti}^2, k_{tj}^2) \Delta_{ij}^2 / R^2$ and $d_{iB} = k_{ti}^2$. This time, we have

$$d_{1B} \gg d_{2B} \sim d_{12} \gg d_{g1} \sim d_{g2} \sim d_{gB}. \quad (5.3)$$

One sees that the ghost always undergoes its clustering before either of the perturbative particles p_1 and p_2 . Specifically, if at least one of Δ_{1g} and Δ_{2g} is smaller than R , the ghost clusters with the particle that is geometrically closer. If both Δ_{1g} and Δ_{2g} are greater than R the ghost clusters with the beam and will not belong to any of the perturbative jets.

Let us now consider various cases. If $\Delta_{12} < R$, (Fig. 5.1a,b) the particles p_1 and p_2 will eventually end up in the same jet. The ghost will therefore belong to the jet irrespectively of having been clustered first with p_1 or p_2 . The area of the jet will then be given by union of two circles of radius R , one centred on each of the two perturbative particles,

$$a_{k_t, R}(\Delta_{12}) = u \left(\frac{\Delta_{12}}{R} \right) \pi R^2, \quad \text{for } \Delta_{12} < R, \quad (5.4)$$

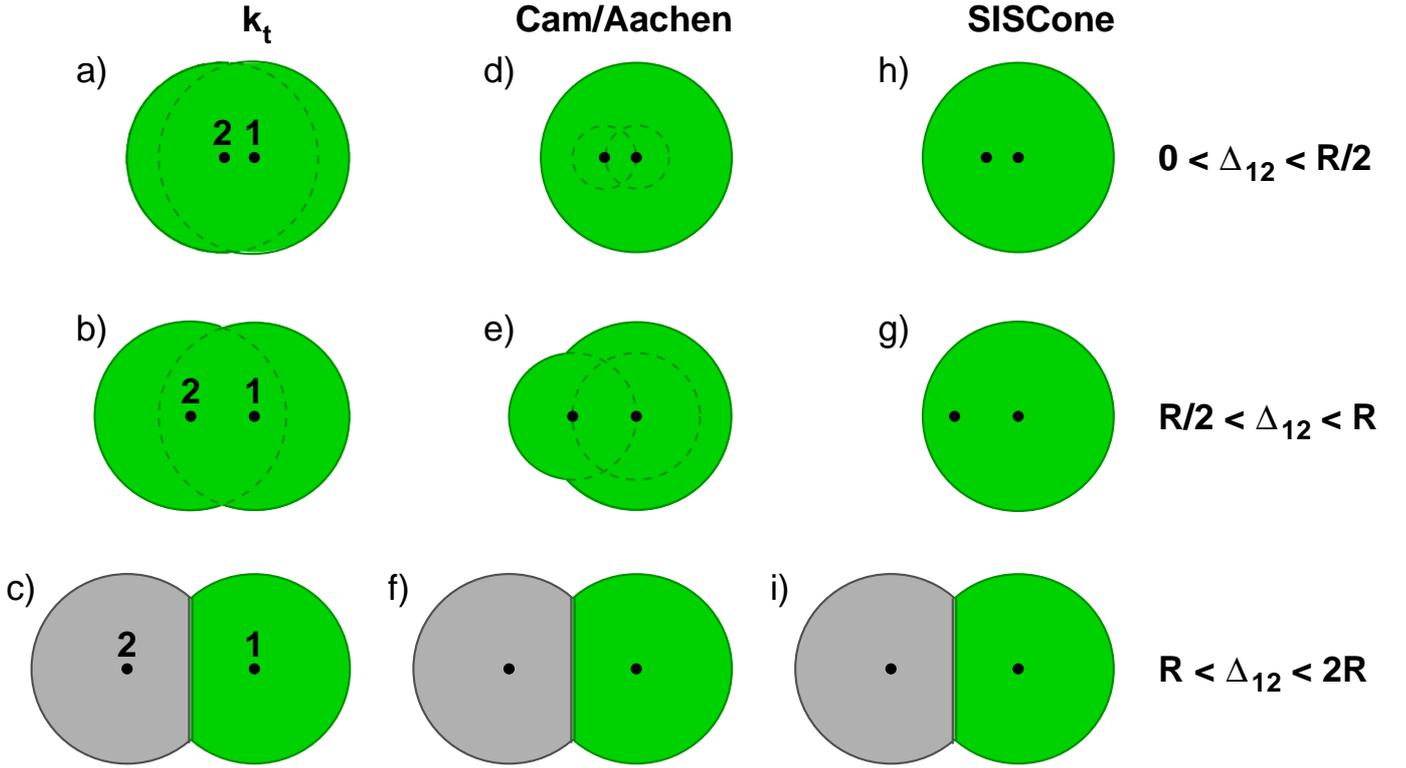


Figure 5.1: Schematic representation of the passive area of a jet containing one hard particle “1” and a softer one “2” for various separations between them and different jet algorithms. Different shadings represent distinct jets. We have not represented the anti- k_t algorithm for which the area remains πR^2 irrespective of the position of the soft particle p_2 .

where

$$u(x) = \frac{1}{\pi} \left[x \sqrt{1 - \frac{x^2}{4}} + 2 \left(\pi - \arccos \left(\frac{x}{2} \right) \right) \right], \quad (5.5)$$

represents the area, divided by π , of the union of two circles of radius one whose centres are separated by x .

The next case we consider is $R < \Delta_{12} < 2R$, Fig. 5.1c. In this case p_1 and p_2 will never be able to cluster together. Hence, they form different jets, and the ghost will belong to the jet of the closer of p_1 or p_2 . The two jets will each have area

$$a_{k_t, R}(\Delta_{12}) = \frac{u(\Delta_{12}/R)}{2} \pi R^2, \quad \text{for } R < \Delta_{12} < 2R. \quad (5.6)$$

Finally, for $\Delta_{12} > 2R$ the two jets formed by p_1 and p_2 each have area πR^2 .

The three cases derived above are summarised in table 5.1 and plotted in Fig. 5.2.

It is interesting to note that the above results also illustrate the fact that for the k_t algorithm, the passive area is equal to the Voronoi area. In this 2-particle configuration, the Voronoi diagram consists of a single line, equidistant between the two points, dividing the plane into two half-planes, each of which is the Voronoi cell of one of the particles (this is best seen in Fig. 5.1c). The intersection of the half-plane with the circle of radius R centred on the particle has area $\frac{1}{2} \pi R^2 u(\Delta_{12}/R)$, and this immediately gives us the results Eqs. (5.4), (5.6) according to whether the particles cluster into a single jet or not.

Cambridge/Aachen. Now we consider the behaviour of the Cambridge/Aachen jet algorithm [28], also a sequential recombination algorithm, defined by the 2-particle distance measure $d_{ij} = \Delta_{ij}^2/R^2$ and beam-particle distance $d_{iB} = 1$. Because in this case the distance measure does not involve the transverse momentum of the particles, the ghost only clusters first if $\min(\Delta_{1g}, \Delta_{2g}) < \Delta_{12}$. Otherwise, p_1 and p_2 cluster first into the jet J , and then J captures the ghost if $\Delta_{Jg} \simeq \Delta_{1g} < R$.

The region $\Delta_{12} < R$ now needs to be separated into two parts, $\Delta_{12} < R/2$ and $R/2 < \Delta_{12} < R$.

If the ghost clusters first, then it must have been contained in either of the dashed circles depicted in Figs. 5.1d,e. If $\Delta_{12} < R/2$ both these circles are contained in a circle of radius R centred on p_1 (Fig. 5.1d), and so the jet area is πR^2 .

If $R/2 < \Delta_{12} < R$ (Fig. 5.1e) the circle of radius Δ_{12} centred on p_2 protrudes, and adds to the area of the final jet, so that

$$a_{\text{Cam},R}(\Delta_{12}) = w\left(\frac{\Delta_{12}}{R}\right) \pi R^2, \quad \text{for } R/2 < \Delta_{12} < R, \quad (5.7)$$

where

$$w(x) = \frac{1}{\pi} \left[\pi - \arccos\left(\frac{1}{2x}\right) + \sqrt{x^2 - \frac{1}{4}} + x^2 \arccos\left(\frac{1}{2x^2} - 1\right) \right]. \quad (5.8)$$

For $\Delta_{12} > R$ a Cambridge/Aachen jet has the same area as the k_t jet, *cf.* Fig. 5.1f. As with the k_t algorithm, the above results are summarised in table 5.1 and Fig. 5.2. The latter in particular illustrates the significant difference between the k_t and Cambridge/Aachen areas for $\Delta_{12} \sim R/2$, caused by the different order of recombinations in the two algorithms.

SISCone. Modern cone jet algorithms like SISCone identify stable cones and then apply a split/merge procedure to overlapping stable cones.²

For $\Delta_{12} < R$ only a single stable cone is found, centred on p_1 . Any ghost within distance R of p_1 will therefore belong to this jet, so its area will be πR^2 , *cf.* Fig. 5.1h,g.

For $R < \Delta_{12} < 2R$ two stable cones are found, centred on p_1 and p_2 . The split/merge procedure will then always split them, because the fraction of overlapping energy is zero. Any ghost falling within either of the two cones will be assigned to the closer of p_1 and p_2 (see Fig. 5.1i). The area of the hard jet will therefore be the same as for the k_t and Cambridge algorithms.

Again these results are summarised in table 5.1 and Fig. 5.2 and one notices the large differences relative to the other algorithms at $R \lesssim 1$ and the striking feature that the cone algorithm only ever has negative corrections to the passive area for these energy-ordered two-particle configurations.

5.1.1.3 Area scaling violation

Given that jet passive areas are modified by the presence of a soft particle in the neighbourhood of a jet, one expects the average jet area to acquire a logarithmic dependence on the jet transverse momentum when the jets acquire a sub-structure as a consequence of radiative emission of gluons. To determine its coefficient we shall work in the approximation of small jet radii, motivated by the observation that corrections for finite R , proportional to powers of R , are usually accompanied by a small coefficient [97, 98].

²In the 2-particle case, the arguments that follow would be identical for the midpoint jet algorithm with a Tevatron run II type split-merge procedure [96]. For higher orders, or more realistic events, it is mandatory to use an infrared-safe seedless variant, like SISCone.

	$a_{\text{JA},R}(\Delta_{12})/\pi R^2$			
	anti- k_t	k_t	Cambridge/Aachen	SISCone
$0 < \Delta_{12} < R/2$	1	$u(\Delta_{12}/R)$	1	1
$R/2 < \Delta_{12} < R$	1	$u(\Delta_{12}/R)$	$w(\Delta_{12}/R)$	1
$R < \Delta_{12} < 2R$	1	$u(\Delta_{12}/R) / 2$	$u(\Delta_{12}/R) / 2$	$u(\Delta_{12}/R) / 2$
$\Delta_{12} > 2R$	1	1	1	1

Table 5.1: Summary of the passive areas for the three jet algorithms for the hard jet in an event containing one hard and one soft particle, separated by a y - ϕ distance Δ_{12} . The functions u and w are defined in Eqs. (5.5) and (5.8).

In the small-angle limit, the QCD matrix element for the emission of the perturbative soft gluon, p_2 , is

$$\frac{dP}{dp_{t2} d\Delta_{12}} = C_1 \frac{2\alpha_s(p_{t2}\Delta_{12})}{\pi} \frac{1}{\Delta_{12}} \frac{1}{p_{t2}}, \quad (5.9)$$

where C_1 is C_F or C_A according to whether the hard particle p_1 is a quark or a gluon. We make use of the fact that Δ_{12} is just the angle between the two particles (to within a longitudinal boost-dependent factor), and will assume that R is sufficiently small that the small-angle approximation is legitimate. The scale of the coupling is taken to be the transverse momentum of p_2 relative to p_1 .

At order α_s the mean jet area with a given jet algorithm (JA) can be written

$$\langle a_{\text{JA},R} \rangle = a_{\text{JA},R}(0) + \langle \Delta a_{\text{JA},R} \rangle = \pi R^2 + \langle \Delta a_{\text{JA},R} \rangle, \quad (5.10)$$

where we explicitly isolate the $\mathcal{O}(\alpha_s)$ higher-order contribution,

$$\langle \Delta a_{\text{JA},R} \rangle \simeq \int_0^{2R} d\Delta_{12} \int_{Q_0/\Delta_{12}}^{p_{t1}} dp_{t2} \frac{dP}{dp_{t2} d\Delta_{12}} (a_{\text{JA},R}(\Delta_{12}) - a_{\text{JA},R}(0)), \quad (5.11)$$

with the $-a_{\text{JA},R}(0)$ term accounting for virtual corrections. $\langle \dots \rangle$ represents therefore an average over perturbative emission. Note that because of the $1/p_{t2}$ soft divergence for the emission of p_2 , Eq. (5.11) diverges unless one specifies a lower limit on p_{t2} — accordingly we have to introduce an infrared cutoff Q_0/Δ_{12} on the p_t of the emitted gluon. This value results from requiring that the transverse momentum of p_2 relative to p_1 , i.e. $p_{t2}\Delta_{12}$, be larger than Q_0 . The fact that we need to place a lower limit on p_{t2} means that jet areas are infrared unsafe — they cannot be calculated order by order in perturbative QCD and for real-life jets they will depend on the details of non-perturbative effects (hadronisation). The case of the anti- k_t jet algorithm is an exception since the area remains πR^2 independently of Δ_{12} , hence the rhs of (5.11) vanishes independently of Q_0 .³ In all other cases, one can account for this to some extent by leaving Q_0 as a free parameter and examining the dependence of the perturbative result on Q_0 .⁴

³Another exception is for the SISCone algorithm with a cut, $p_{t,\text{min}}$, on the minimum transverse momentum of protojets that enter the split-merge procedure — in that situation Q_0 is effectively replaced by $p_{t,\text{min}}$.

⁴One should of course bear in mind that the multi-particle structure of the hadron level is such that a single-gluon result cannot contain all the relevant physics — this implies that one should, in future work, examine multi-soft gluon radiation as well, perhaps along the lines of the calculation of non-global logarithms [99, 100, 101, 102], though it is not currently clear how most meaningfully to carry out the matching with the non-perturbative regime. Despite these various issues, we shall see in Section 5.1.4 that the single-gluon results work remarkably well in comparisons to Monte Carlo predictions. In that section, we shall also argue that in cases with pileup, the pileup introduces a natural semi-hard (i.e. perturbative) cutoff scale that replaces Q_0 .

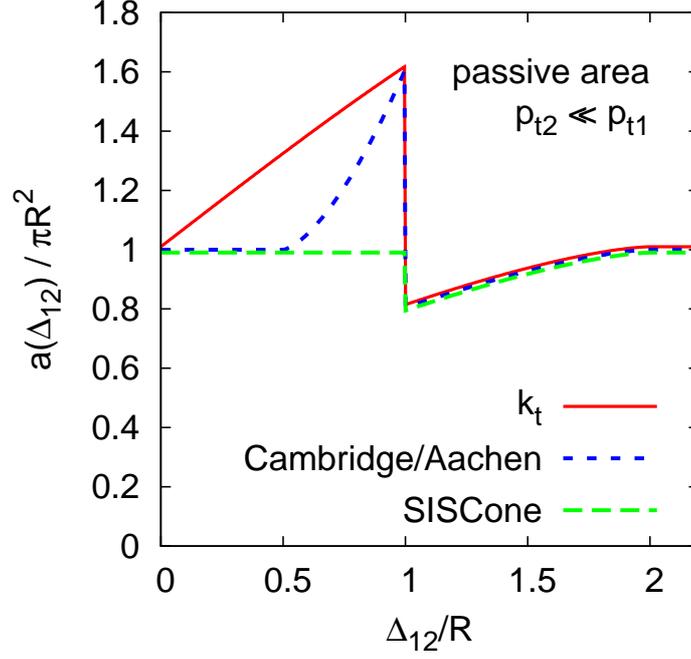


Figure 5.2: Plot of the passive areas of the hard jet as a function of the distance between the hard and the soft particle, as given in table 5.1. The anti- k_t case is not shown but corresponds to a constant line with $a(\Delta_{12})/(\pi R^2) = 1$.

As concerns the finiteness of the Δ_{12} integration, all the jet algorithms we consider have the property that

$$\lim_{\Delta_{12} \rightarrow 0} a_{\text{JA},R}(\Delta_{12}) = a_{\text{JA},R}(0) = \pi R^2; \quad (5.12a)$$

$$a_{\text{JA},R}(\Delta_{12}) = \pi R^2 \quad \text{for } \Delta_{12} > 2R \quad (5.12b)$$

so that the integral converges for $\Delta_{12} \rightarrow 0$, and we can place the upper limit at $\Delta_{12} = 2R$.

After evaluating Eq. (5.11), with the replacement $\Delta_{12} \rightarrow R$, both in the lower limit of the p_{t2} integral and the argument of the coupling, we obtain

$$\langle \Delta a_{\text{JA},R} \rangle = d_{\text{JA},R} \frac{2\alpha_s C_1}{\pi} \ln \frac{R p_{t1}}{Q_0}, \quad d_{\text{JA},R} = \int_0^{2R} \frac{d\theta}{\theta} (a_{\text{JA},R}(\theta) - \pi R^2), \quad (5.13)$$

in a fixed coupling approximation, and

$$\langle \Delta a_{\text{JA},R} \rangle = d_{\text{JA},R} \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(R p_{t1})}, \quad (5.14)$$

with a one-loop running coupling, where $b_0 = \frac{11C_A - 2n_f}{12\pi}$. The approximation $\Delta_{12} \sim R$ in the argument of the running coupling in the integrand corresponds to neglecting terms of $\mathcal{O}(\alpha_s)$ without any enhancements from logarithms of R or p_{t1}/Q_0 .

The results for $d_{\text{JA},R}$ are

$$d_{\text{anti-}k_t,R} = 0, \quad (5.15a)$$

$$d_{k_t,R} = \left(\frac{\sqrt{3}}{8} + \frac{\pi}{3} + \xi \right) R^2 \simeq 0.5638 \pi R^2, \quad (5.15b)$$

$$d_{\text{Cam},R} = \left(\frac{\sqrt{3}}{8} + \frac{\pi}{3} - 2\xi \right) R^2 \simeq 0.07918 \pi R^2, \quad (5.15c)$$

$$d_{\text{SISCone},R} = \left(-\frac{\sqrt{3}}{8} + \frac{\pi}{6} - \xi \right) R^2 \simeq -0.06378 \pi R^2, \quad (5.15d)$$

where

$$\xi \equiv \frac{\psi'(1/6) + \psi'(1/3) - \psi'(2/3) - \psi'(5/6)}{48\sqrt{3}} \simeq 0.507471, \quad (5.16)$$

with $\psi'(x) = d^2/dx^2(\ln \Gamma(x))$. One notes that the coefficient for the k_t algorithm is non-negligible, given that it is multiplied by the quantity $2\alpha_s C_1/\pi \ln Rp_{t1}/Q_0$ in Eq. (5.13) (or its running coupling analogue), which can be of order 1. In contrast the coefficients for Cambridge/Aachen and SISCone are much smaller and similar (the latter being however of opposite sign). Thus k_t areas will depend significantly more on the jet p_t than will those for the other algorithms. In contrast, the soft-resilience of the anti- k_t jet algorithm is manifest here with an area of πR^2 independently of p_t .

The fluctuation of the areas (over the emission p_2) can be calculated in a similar way. Let us define

$$\langle \sigma_{\text{JA},R}^2 \rangle = \langle a_{\text{JA},R}^2 \rangle - \langle a_{\text{JA},R} \rangle^2 = \sigma_{\text{JA},R}^2(0) + \langle \Delta \sigma_{\text{JA},R}^2 \rangle, \quad \sigma_{\text{JA},R}^2(0) = 0, \quad (5.17)$$

where we have introduced $\sigma_{\text{JA},R}^2(0)$, despite its being null, so as to facilitate comparison with later results. We then evaluate

$$\langle \Delta \sigma_{\text{JA},R}^2 \rangle = \langle \Delta a_{\text{JA},R}^2 \rangle - \langle \Delta a_{\text{JA},R} \rangle^2 \simeq \langle \Delta a_{\text{JA},R}^2 \rangle, \quad (5.18)$$

where we neglect $\langle \Delta a_{\text{JA},R} \rangle^2$ since it is of $\mathcal{O}(\alpha_s^2 \ln^2(Rp_{t1}/Q_0))$. The calculation of $\langle \Delta a_{\text{JA},R}^2 \rangle$ proceeds much as for $\langle \Delta a_{\text{JA},R} \rangle$ and gives

$$\langle \Delta a_{\text{JA},R}^2 \rangle = s_{\text{JA},R}^2 \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_{t1})}, \quad s_{\text{JA},R}^2 = \int_0^{2R} \frac{d\theta}{\theta} (a_{\text{JA},R}(\theta) - \pi R^2)^2 \quad (5.19)$$

for running coupling. The results are

$$s_{\text{anti-}k_t,R}^2 = 0, \quad (5.20a)$$

$$s_{k_t,R}^2 = \left(\frac{\sqrt{3}\pi}{4} - \frac{19}{64} - \frac{15\zeta(3)}{8} + 2\pi\xi \right) R^4 \simeq (0.4499 \pi R^2)^2, \quad (5.20b)$$

$$s_{\text{Cam},R}^2 = \left(\frac{\sqrt{3}\pi}{6} - \frac{3}{64} - \frac{\pi^2}{9} - \frac{13\zeta(3)}{12} + \frac{4\pi}{3}\xi \right) R^4 \simeq (0.2438 \pi R^2)^2, \quad (5.20c)$$

$$s_{\text{SISCone},R}^2 = \left(\frac{\sqrt{3}\pi}{12} - \frac{15}{64} - \frac{\pi^2}{18} - \frac{13\zeta(3)}{24} + \frac{2\pi}{3}\xi \right) R^4 \simeq (0.09142 \pi R^2)^2. \quad (5.20d)$$

We have a hierarchy between algorithms that is similar to that observed for the average area scaling violations, with anti- k_t showing no deviation from πR^2 . The coefficient for Cambridge/Aachen is now more intermediate between the k_t and SISCone values.

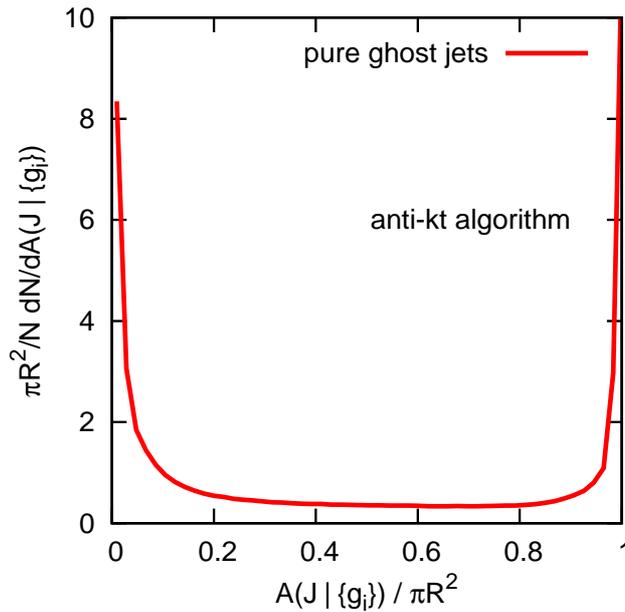


Figure 5.3: Distribution of the active area for pure-ghost jets obtained with the anti- k_t algorithm.

5.1.2 Active areas

5.1.2.1 Areas for 1-particle configurations and for ghost jets

This time, we consider a configuration with a single hard particle and an ensemble of ghosts, with the ghosts taking active part in the clustering.

Anti- k_t . The shortest anti- k_t distance will be between the hard particle and either a ghost or the beam, depending on the pure geometrical distance. The hard particle p_1 will therefore cluster with all the ghosts within a geometrical distance R from it before it becomes a jet (by clustering with the beam). Ghost-ghost distances are much larger, enhanced by a factor p_t^2/g_t^2 . The active area will therefore be πR^2 (with no fluctuations), once again a direct consequence of the soft resilience of the anti- k_t algorithm:

$$A_{\text{anti-}k_t, R}(\text{one-particle-jet}) = \pi R^2, \quad (5.21a)$$

$$\Sigma_{\text{anti-}k_t, R}(\text{one-particle-jet}) = 0. \quad (5.21b)$$

Beside the hard jets, there will also be pure-ghost jets. Their distribution is most readily studied numerically, by directly clustering large numbers of ghost particles, with or without additional hard particles. Typically we add ghosts with a density ν_g of ~ 100 per unit area,⁵ in the rapidity region $|y| < y_{\text{max}} = 6$, and study jets in the region $|y| < y_{\text{max}} - R$. This leads to about 7500 ghost particles. Each ghost is given a transverse momentum $\sim 10^{-100}$ GeV.

The resulting distribution of the anti- k_t active areas $A(J | \{g_i\})$ is presented in Fig. 5.3 for a large

⁵They are placed on a randomly scattered grid, in order to limit the impact of the finite density, i.e. one effectively carries out quasi Monte Carlo integration of the ghost ensembles, so that that finite density effects ought to vanish as $\nu_g^{-3/4}$, rather than $\nu_g^{-1/2}$ as would be obtained with completely random placement.

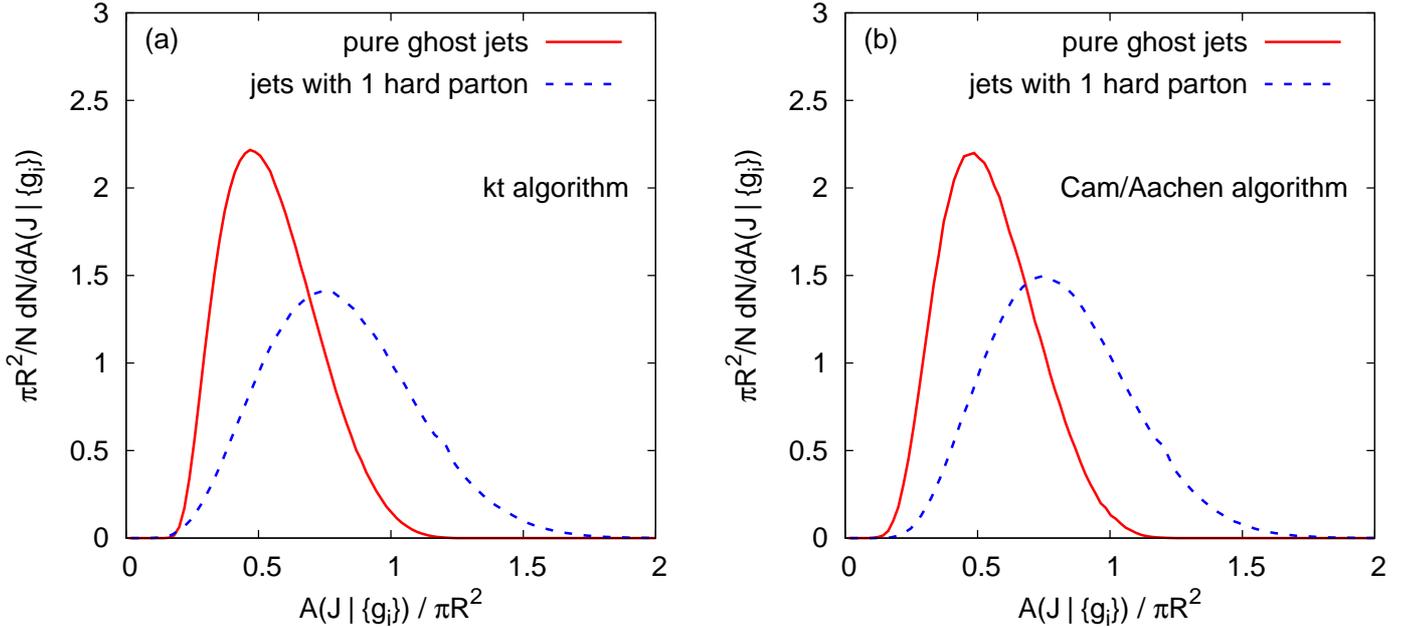


Figure 5.4: Distribution of active areas for pure ghost jets and jets with a single hard particle: (a) k_t algorithm, (b) Cambridge/Aachen algorithm.

ensemble of sets of ghosts.⁶ This distribution appears to be dominated by jets which are either almost circular and of area close to πR^2 and jets of much smaller area. We observe about 30% of the jets with an area smaller than $0.1\pi R^2$ and about 30% with an area larger than $0.95\pi R^2$. This distribution can be understood in the following way: pure-ghost jets will only be clustered once all the hard jets have been clustered. That clustering will start aggregating ghosts, giving them a larger p_t than the other ghosts, seeding a process of further aggregating the ghosts in the vicinity of that original cluster, in a way similar to the clustering of the hard jets. This would give a peak close to $A = \pi R^2$. As this goes on, the space between the jets becomes smaller and small pure-ghost jets, of area close to 0, fill the spaces between the jets of area close to πR^2 .

Note that this distribution justifies why anti- k_t jets are not suited for the area–median estimation of the pileup density ρ . The patches used for pileup estimation would follow a distribution close to that of the pure-ghost jets — potentially with an additional bias coming from pileup fluctuations — and the large number of small jets would lead to a large uncertainty on the estimation of ρ .

k_t and Cambridge/Aachen. The active area for the k_t and Cambridge/Aachen algorithms is again best studied numerically. On top of the ghosts distributed as discussed above, we added one hard particle with a transverse momentum of 100 GeV that we study. The results are insensitive to the transverse momentum as long as it is sufficiently larger than the ghost p_t . We briefly investigate in Appendix D.1 how the distribution of the “1-hard-parton” jet area gets modified when the transverse momentum of the parton is progressively reduced below the scale of a generic set of soft particles. Figure 5.4 shows the distribution of values of $A(J | \{g_i\})$ for pure ghost jets and jets with one hard particle. Let us concentrate initially on the case with a hard particle. Firstly, the average active area,

⁶In this particular case we have used about 10^7 separate random ghost sets, in order to obtain a smooth curve for the whole distribution. When calculating areas in physical events (or even at parton-shower level) the multiple real particles in the jet “fix” most of the area, and between 1 and 5 sets of ghosts particles are usually sufficient to obtain reliable area results (this is the case also for the other algorithms studied below).

Eq. (2.20) differs noticeably from the passive result of πR^2 :

$$A_{k_t,R}(\text{one-particle-jet}) \simeq 0.812 \pi R^2, \quad (5.22a)$$

$$A_{\text{Cam},R}(\text{one-particle-jet}) \simeq 0.814 \pi R^2. \quad (5.22b)$$

Secondly, the distributions of the area in Fig. 5.4 are rather broad. The randomness in the initial distribution of ghosts propagates all the way into the shape of the final jet and hence its area. This occurs because the k_t and Cambridge/Aachen algorithms flexibly adapt themselves to local structure (a good property when trying to reconstruct perturbative showering), and once a random perturbation has formed in the density of ghosts this seeds further growth of the soft part of the jet. The standard deviations of the resulting distributions are

$$\Sigma_{k_t,R}(\text{one-particle-jet}) \simeq 0.277 \pi R^2, \quad (5.23a)$$

$$\Sigma_{\text{Cam},R}(\text{one-particle-jet}) \simeq 0.261 \pi R^2. \quad (5.23b)$$

Figure 5.4 also shows the distribution of areas for pure ghost jets. One sees that pure ghost jets typically have a smaller area than hard-particle jets:⁷

$$A_{k_t,R}(\text{ghost-jet}) \simeq 0.554 \pi R^2, \quad (5.25a)$$

$$A_{\text{Cam},R}(\text{ghost-jet}) \simeq 0.551 \pi R^2, \quad (5.25b)$$

and the standard deviations are

$$\Sigma_{k_t,R}(\text{ghost-jet}) \simeq 0.174 \pi R^2, \quad (5.26a)$$

$$\Sigma_{\text{Cam},R}(\text{ghost-jet}) \simeq 0.176 \pi R^2. \quad (5.26b)$$

The fact that pure ghost jets are smaller than hard jets has an implication for certain physics studies: one expects jets made of soft ‘junk’ (minimum bias, pileup, thermal noise in heavy ions) to have area properties similar to ghost jets; since they are smaller on average than true hard jets, the hard jets will emerge from the junk somewhat more clearly than if both had the same area.

SISCone. The SISCone active area is amenable to analytical treatment, at least in some cases.

We recall that SISCone starts by finding all stable cones. A first stable cone is centred on the single hard particle. Additionally, there will be a large number of other stable cones, of order of the number of ghost particles added to the event [95]. In the limit of an infinite number of ghosts, all cones that can be drawn in the rapidity-azimuth plane and that do not overlap with the hard particle will be stable. Many of these stable cones will still overlap with the cone centred on the hard particle, as long as they do not contain the hard particle itself (see figure 5.5, left).

⁷Obtaining these values actually requires going beyond the ghost density and the rapidity range previously mentioned. In fact, when going to higher accuracy one notices the presence of small edge and finite-density effects, $\mathcal{O}(R/(y_{max} - R))$ and $\mathcal{O}(1/(\nu_g R^2))$ to some given power. Choosing the ghost area sufficiently small to ensure that finite-density effects are limited to the fourth decimal (in practice this means $1/(\nu_g R^2) < 0.01$) and extrapolating to infinite y_{max} one finds

$$A_{k_t}(\text{ghost-jet}) \simeq (0.5535 \pm 0.0005) \pi R^2, \quad (5.24a)$$

$$A_{\text{Cam}}(\text{ghost-jet}) \simeq (0.5505 \pm 0.0005) \pi R^2, \quad (5.24b)$$

with a conservative estimate of the residual uncertainty. This points to a small but statistically significant difference between the two algorithms.

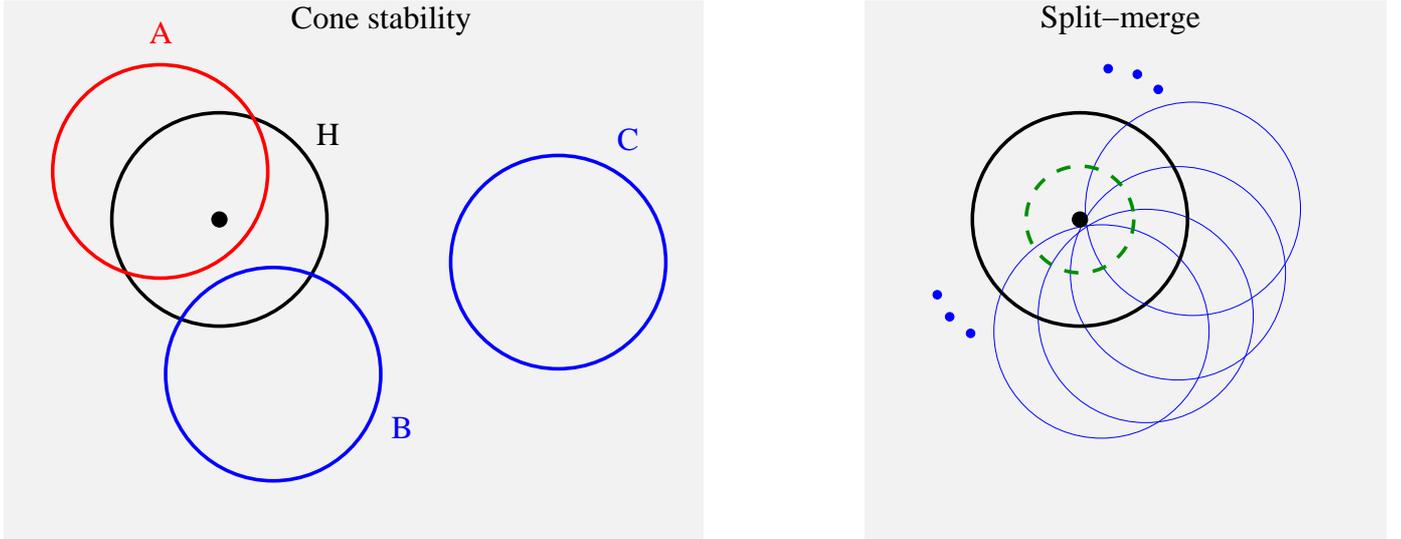


Figure 5.5: Left: the hard particle with the stable cone (H) centred on it, an example of a cone (A) that is unstable because it also contains the hard particle, and of two cones (B) and (C) that contain just ghost particles and are therefore stable. Right: some of the stable ghost cones (thin blue circles) that have the largest possible overlap with (H), together with the boundary of the hard jet after the split–merge procedure (dashed green line). In both diagrams, the grey background represents the uniform coverage of ghosts.

Next, the SIScone algorithm involves a split–merge procedure. One defines \tilde{p}_t for a jet as the scalar sum of the transverse momenta of its constituents. During the split–merge step, SIScone finds the stable cone with the highest \tilde{p}_t , and then the next hardest stable cone that overlaps with it. To decide whether these two cones (protojets) are to be merged or split, it determines the fraction of the softer cone’s \tilde{p}_t that is shared with the harder cone. If this fraction is smaller than some value f (the overlap parameter of the cone algorithm), the two protojets are split: particles that are shared between them are assigned to the protojet to whose centre⁸ they are closer. Otherwise they are merged into a single protojet. This procedure is repeated until the hard protojet no longer has any overlap with other protojets. At this point it is called a jet, and the split–merge procedure continues on the remaining protojets (without affecting the area of the hard jet).

The maximum possible overlap fraction,⁹ f_{\max} , between the hard protojet and a ghost protojet occurs in the situation depicted in figure 5.5 (right), i.e. when the ghost protojet’s centre is just outside the edge of the original hard stable cone (H). It is given by $f_{\max} = 2 - u(1) = \frac{2}{3} - \frac{\sqrt{3}}{2\pi} \approx 0.391$. This means that for a split–merge parameter $f > f_{\max}$ (commonly used values are $f = 0.5$ and $f = 0.75$) every overlap between the hard protojet and a pure-ghost stable cone will lead to a splitting. Since these pure-ghost stable cones are centred at distances $d > R$ from the hard particle, these splittings will reduce the hard jet to a circle of radius $R/2$ (the dashed green line in the right hand part of figure 5.5). The active area of the hard jet is thus

$$A_{\text{SIScone},R}(\text{one-particle-jet}) = \frac{\pi R^2}{4}. \quad (5.27)$$

This result has been verified numerically using the same technique employed above for k_t and Cam-

⁸This centre is given by the sum of momenta in the protojet before the split–merge operation.

⁹The fraction of momentum coincides with the fraction of area because the ghosts have uniform transverse momentum density.

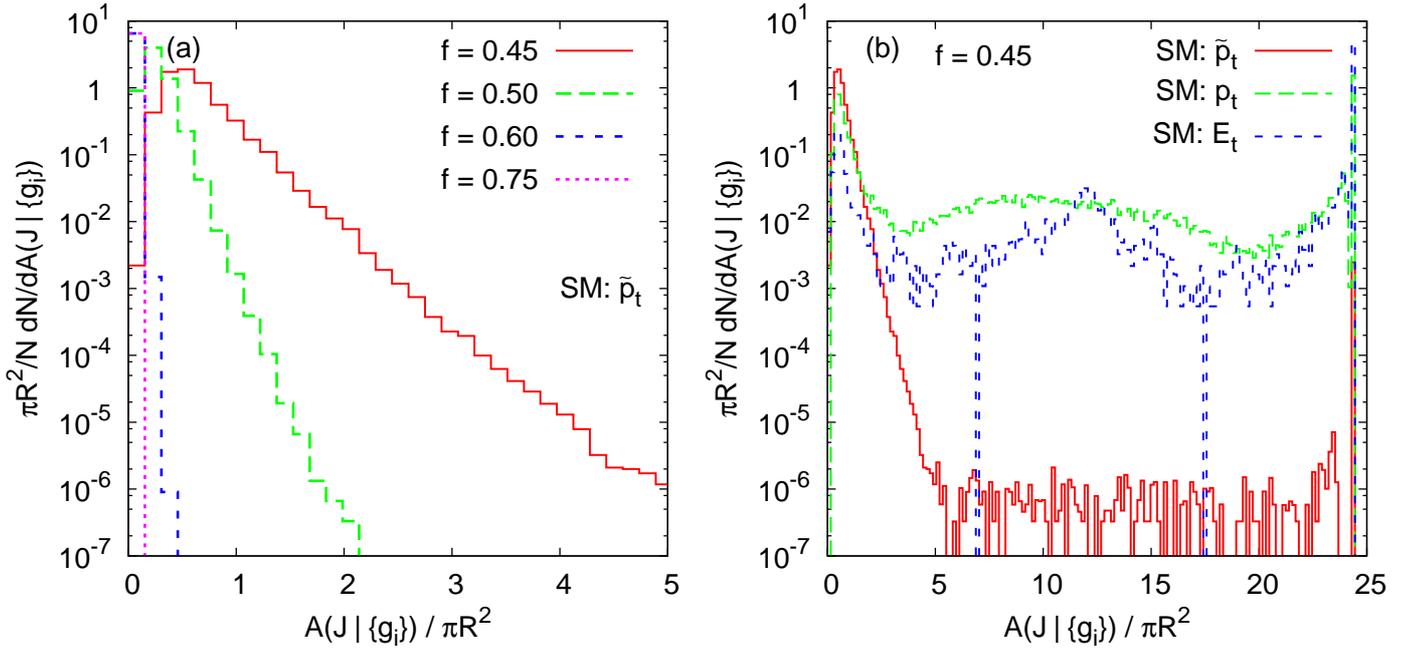


Figure 5.6: Distribution of pure-ghost jet areas for the SIS Cone algorithm, (a) with different values of the split–merge parameter f and (b) different choices for the scale used in the split–merge procedure. Ghosts are placed on a grid up to $|y| < 6$, with an average spacing of 0.2×0.2 in y, ϕ , and a random displacement of up to ± 0.1 in each direction, and transverse momentum values that are uniform modulo a $\pm 5\%$ random rescaling for each ghost. We consider all jets with $y < 5$. All jet definitions use $R = 1$ and multiple passes.

bridge/Aachen.

Note that this area differs considerably from the passive area, πR^2 . This shows that the cone area is very sensitive to the structure of the event, and it certainly does not always coincide with the naive geometrical expectation πR^2 , contrary to assumptions sometimes made in the literature (see for example [103]).

We further note that in contrast to k_t and Cambridge/Aachen algorithms, the SIS Cone algorithm always has the same active area for a single hard particle, independently of fluctuations of an infinitely dense set of ghosts, i.e.

$$\Sigma_{\text{SIS Cone}, R}(\text{one-particle-jet}) = 0, \quad (5.28)$$

a property that SIS Cone shares with the anti- k_t algorithm.

SIS Cone ghost-jet areas. While the area of a hard particle jet could be treated analytically, this is not the case for the pure-ghost jet area. Furthermore, numerical investigations reveal that the pure-ghost area distribution has a much more complicated behaviour than for k_t or Cambridge/Aachen. One aspect is that the distribution of pure-ghost jet areas is sensitive to the fine details of how the ghosts are distributed (density and transverse momentum fluctuations). Another is that it depends significantly on the details of the split–merge procedure. Figure 5.6a shows the distribution of areas of ghost jets for SIS Cone, for different values of the split–merge overlap threshold f . One sees, for example, that for smaller values of f there are occasional rather large ghost jets, whereas for $f \gtrsim 0.6$ nearly all ghost jets have very small areas.

One of the characteristics of SIS Cone that differs from previous cone algorithms is the specific

ordering and comparison variable used to determine splitting and merging. As explained above, the choice made in SISCone was \tilde{p}_t , the scalar sum of transverse momenta of all particles in a jet. Previous cone algorithms used either the vector sum of constituent transverse momenta, p_t (an infrared unsafe choice), or the transverse energy $E_t = E \sin \theta$ (in a 4-vector recombination scheme). With both of these choices of variable, split–merge thresholds $f \lesssim 0.55$ can lead to the formation of ‘monster’ ghost jets, which can even wrap around the whole cylindrical phase space. For $f = 0.45$ this is a quite frequent occurrence, as illustrated in figure 5.6b, where one sees a substantial number of jets occupying the whole of the phase space (i.e. an area $4\pi y_{\max} \simeq 24\pi R^2$). Monster jets can be formed also with the \tilde{p}_t choice, though it is a somewhat rarer occurrence — happening in ‘only’ $\sim 5\%$ of events.¹⁰

We have observed the formation of such monster jets also from normal pileup momenta simulated with Pythia [35], indicating that this disturbing characteristic is not merely an artefact related to our particular choice of ghosts. This indicates that a proper choice of the split–merge variable and threshold is critical in high-luminosity environments. The results from Figure 5.6a, suggest that if one wants to avoid monster jets, one has to choose a large enough value for f . Our recommendation is to adopt $f = 0.75$ as a default value for the split–merge threshold (together with the use of the \tilde{p}_t variable, already the default in SISCone, for reasons related to infrared safety and longitudinal boost invariance).

5.1.2.2 Areas for 2-particle configurations

In this section we study the same problem described in section 5.1.1.2, i.e. the area of jets containing two particles, a hard one and a softer (but still “perturbative”) one, Eq. (5.1), but now for active areas. As before, the results will then serve as an input in understanding the dependence of the active area on the jet’s transverse momentum when accounting for perturbative radiation.

Anti- k_t . With little surprise, the addition of the soft particle p_2 will not affect the area of the hard jet. The minimal anti- k_t distance will be between particle p_1 and the closest of p_2 and the (remaining) ghosts, until that geometric distance reaches R and the jet is clustered with the beam. The resulting area will be πR^2 ,

$$A_{\text{anti-}k_t,R}(\Delta_{12}) = \pi R^2, \quad (5.29)$$

with no fluctuations

$$\Sigma_{\text{anti-}k_t,R}(\Delta_{12}) = 0, \quad (5.30)$$

independently of the distance Δ_{12} between particles p_1 and p_2 .

k_t and Cambridge/Aachen. As was the case for the active area of a jet containing a single hard particle, we again have to resort to numerical analyses to study that of jets with two energy-ordered particles. We define $A_{\text{JA},R}(\Delta_{12})$ to be the active area for the energy-ordered two particle configuration already discussed in section 5.1.1.2.

Additionally since we have a distribution of areas for the single-particle active area case, it becomes of interest to study also $\Sigma_{\text{JA},R}(\Delta_{12})$ the standard deviation of the distribution of areas obtained for the two-particle configuration.

The results are shown in figure 5.7: the active areas can be seen to be consistently smaller than the passive ones, as was the case for the 1-particle area, but retain the same dependence on the

¹⁰This figure is not immediately deducible from Fig. 5.6b, which shows results normalised to the total number of ghost jets, rather than to the number of events.

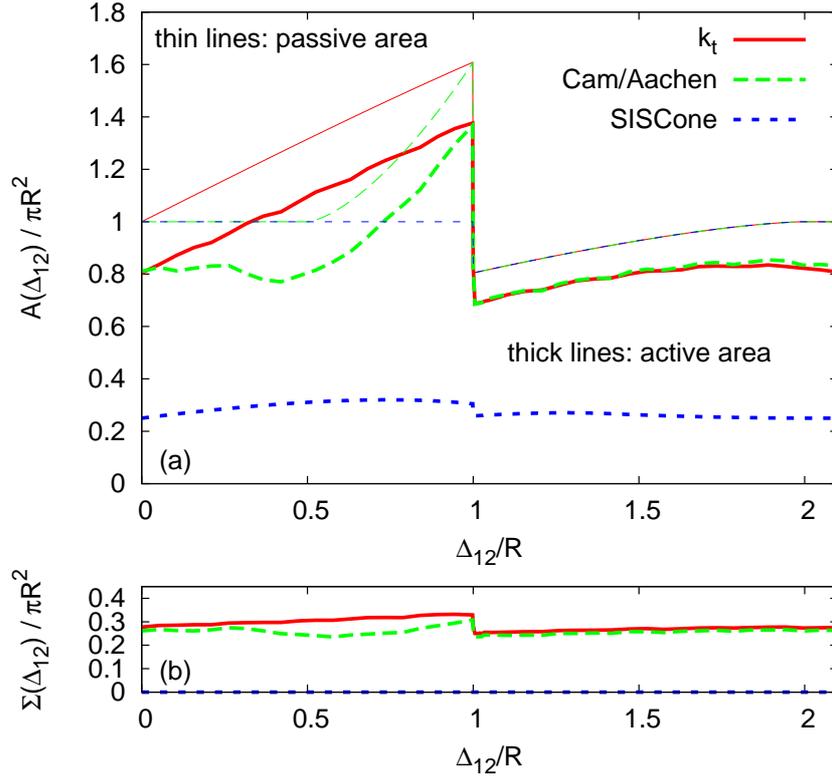


Figure 5.7: (a) Active areas (divided by πR^2) for the three jet algorithms as a function of the separation between a hard and a softer particle. For comparison we also include the passive areas, previously shown in Fig. 5.2. (b) The corresponding standard deviations.

angular separation between the two particles. Among the various features, one can also observe that the active area does not quite reach the single-particle value ($\simeq 0.81\pi R^2$) at $\Delta_{12} = 2R$, but only somewhat beyond $2R$. This contrasts with the behaviour of the passive area. The figure also shows results for the cone area, discussed in the following subsection.

SIScone In the case of the SIScone algorithm it is possible to find an analytical result for the two-particle active area, in an extension of what was done for one-particle case.¹¹

The stable-cone search will find one or two “hard” stable cones: the first centred on the hard particle and the second centred on the soft one, present only for $\Delta_{12} > R$. The pure-ghost stable cones will be centred at all positions distant by more than R from both p_1 and p_2 , i.e. outside the two circles centred on p_1 and p_2 .

We shall consider the active area of the jet centred on the hard particle p_1 . When $\Delta_{12} > R$, the jet centred on the soft particle has the same area.

As in section 5.1.2.1, the split–merge procedure first deals with the pure-ghost protojets overlapping with the hard stable cone. For $f > f_{\max}$, this again only leads to splittings. Depending of the value of Δ_{12} , different situations are found as shown on figure 5.8. If the y - ϕ coordinates of the particles are $p_1 \equiv (\Delta, 0)$ and $p_2 \equiv (0, 0)$, the geometrical objects that are present are: the circle centred on p_1 with radius $R/2$; the tangents to this circle at y - ϕ coordinates $(\Delta_{12}/4, \pm \frac{1}{2}\sqrt{R^2 - (\Delta_{12}/2)^2})$;

¹¹This is only possible for configurations with strong energy ordering between all particles — as soon as 2 or more particles have commensurate transverse momenta then the cone’s split–merge procedure will include ‘merge’ steps, whose effects on the active area are currently beyond analytical understanding.

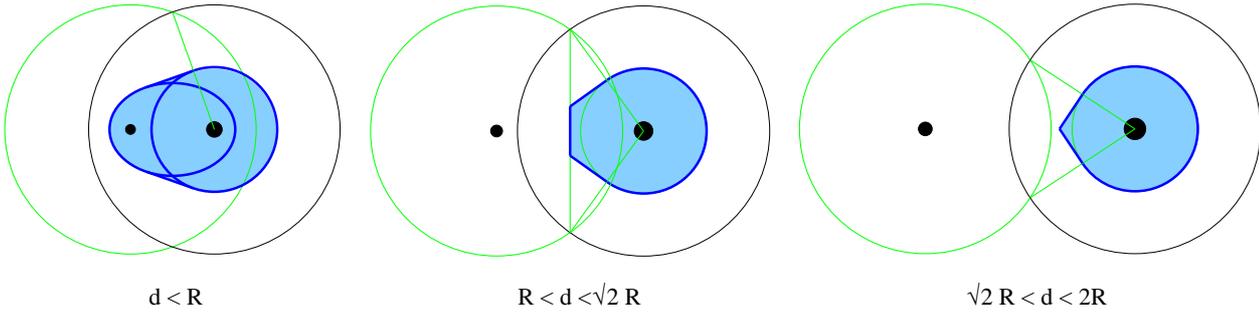


Figure 5.8: Picture of the active jet area for 2-particle configurations in the case of the SISCone jet algorithm. The black points represent the hard (big dot) and soft (small dot) particles, the black circle is the hard stable cone. The final hard jet is represented by the shaded area. The left (a) (centre (b), right (c)) plot corresponds to $\Delta_{12} < R$ ($R < \Delta_{12} < \sqrt{2}R$, $\sqrt{2}R < \Delta_{12} < 2R$).

and, for $\Delta_{12} < R$, the ellipse of eccentricity Δ_{12}/R whose foci are p_1 and p_2 (given by the equation, $\Delta_{a1} + \Delta_{a2} = R$, where a is a point on the ellipse). For $\Delta_{12} < R$ the area of the jet is given by the union of the ellipse, the circle and the regions between the ellipse, the circle and each of the tangents (Fig. 5.8a). For $R < \Delta_{12} < \sqrt{2}R$ it is given by the circle plus the region between the circle, the two tangents and the line equidistant from p_1 and p_2 (Fig. 5.8b). For $\sqrt{2}R < \Delta_{12} < 2R$ it is given by the circle plus the region between the circle and two tangents, up to the intersection of the two tangents (Fig. 5.8c). Finally, for $\Delta \geq 2R$, the area is $\pi R^2/4$.

An analytic computation of the active area gives

$$\frac{A_{\text{SISCone},R}(\Delta_{12})}{\pi R^2} = \frac{1}{4} \left[1 - \frac{1}{\pi} \arccos\left(\frac{x}{2}\right) \right] \quad (5.31)$$

$$+ \begin{cases} \frac{x}{2\pi} \sqrt{1 - \frac{x^2}{4}} + \frac{1}{4\pi} \sqrt{1 - x^2} \arccos\left(\frac{x}{2-x^2}\right) & \Delta_{12} < R \\ \frac{x}{2\pi} \sqrt{1 - \frac{x^2}{4}} - \frac{x}{8\pi} \frac{1}{\sqrt{1 - \frac{x^2}{4}}} & R < \Delta_{12} \leq \sqrt{2}R \\ \frac{1}{2\pi x} \sqrt{1 - \frac{x^2}{4}} & \sqrt{2}R < \Delta_{12} \leq 2R \end{cases} .$$

with $x = \Delta_{12}/R$, while for $\Delta_{12} > 2R$, one recovers the result $\pi R^2/4$. The SISCone active area is plotted in figure 5.9 and it is compared to the results for k_t and Cambridge/Aachen in figure 5.7. One notes that the SISCone result is both qualitatively and quantitatively much further from the passive result than was the case for k_t and Cambridge/Aachen.

The main reason why the 2-point active area is larger than the 1-point active area (whereas we saw the opposite behaviour for the passive areas) is that the presence of the 2-particle configuration causes a number of the pure-ghost cones that were present in the 1-particle case to now contain the second particle and therefore be unstable. Since these pure ghost cones are responsible for reducing the jet area relative to the passive result (during the split–merge step), their absence causes the active area to be ‘less-reduced’ than in the 1-particle case.

5.1.2.3 Area scaling violation

We can write the order α_s contribution to the average active area in a manner similar to the passive area case, Eq. (5.11):

$$\langle A_{\text{JA},R} \rangle = A_{\text{JA},R}(0) + \langle \Delta A_{\text{JA},R} \rangle, \quad (5.32)$$

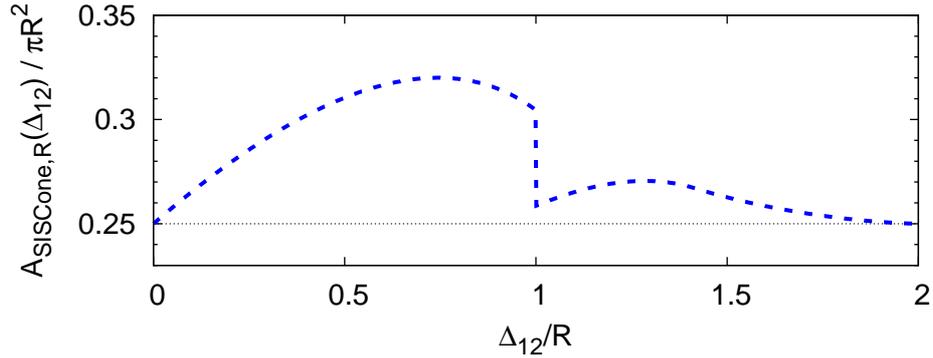


Figure 5.9: Active area of the hardest jet as a function of the distance between the hard and soft particle for the SISCone algorithm, *cf.* Eq. (5.31).

where we have used the relation $A_{JA,R}(\text{one-particle-jet}) \equiv A_{JA,R}(0)$ and where

$$\langle \Delta A_{JA,R} \rangle \simeq \int_0^{p_{t1}} d\Delta_{12} \int_{Q_0/\Delta_{12}}^{p_{t2}} dp_{t2} \frac{dP}{dp_{t2} d\Delta_{12}} (A_{JA,R}(\Delta_{12}) - A_{JA,R}(0)). \quad (5.33)$$

Note that compared to Eq. (5.11) we have removed the explicit upper limit at $2R$ on the Δ_{12} integral since, for active areas of some sequential recombination algorithms, $(A_{JA,R}(\Delta_{12}) - A_{JA,R}(0))$ may be non zero even for $\Delta_{12} > 2R$. Note also the notation for averages: we use $\langle \dots \rangle$ to refer to an average over perturbative emissions, while $\langle \dots \rangle_g$, implicitly contained in $A_{JA,R}$ (see Eq. (2.20)), refers to an average over ghost ensembles. We now proceed as in section 5.1.1.3, and write

$$\langle \Delta A_{JA,R} \rangle \simeq D_{JA,R} \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_{t1})}, \quad D_{JA,R} = \int_0 \frac{d\theta}{\theta} (A_{JA,R}(\theta) - A_{JA,R}(0)), \quad (5.34)$$

where for brevity we have given just the running-coupling result. One observes that the result continues to depend on Q_0 , an indication that the active area is an infrared-unsafe quantity, just like the passive area.¹² The coefficients for the anomalous dimension of the active area for the various algorithms are

$$D_{\text{anti-}k_t,R} \simeq 0, \quad (5.35a)$$

$$D_{k_t,R} \simeq 0.52 \pi R^2, \quad (5.35b)$$

$$D_{\text{Cam},R} \simeq 0.08 \pi R^2, \quad (5.35c)$$

$$D_{\text{SISCone},R} \simeq 0.1246 \pi R^2, \quad (5.35d)$$

where the anti- k_t and SISCone results has been obtained by integrating the analytical result, Eqs. (5.29) and (5.31) respectively, while the results for k_t and Cambridge/Aachen have been obtained both by integrating the 2-point active-area results shown in Fig. 5.7 and by a direct Monte Carlo evaluation of Eq. (5.34). Note that while the coefficients for k_t and Cambridge/Aachen are only slightly different from their passive counterparts, the one for SISCone has the opposite sign relative to the passive one.

The treatment of higher-order fluctuations of active areas is more complex than that for the passive ones, where the one-particle area was a constant. We can separate the fluctuations of active areas into two components, one (described above) that is the just the one-particle result and the other, $\langle \Delta \Sigma_{JA,R}^2 \rangle$, accounting for their modification in the presence of perturbative radiation:

$$\langle \Sigma_{JA,R}^2 \rangle = \Sigma_{JA,R}^2(0) + \langle \Delta \Sigma_{JA,R}^2 \rangle, \quad (5.36)$$

¹²We believe that any sensible (i.e. related to the jet's sensitivity to UE/pileup type contamination) definition of area would actually be infrared unsafe.

where $\Sigma_{\text{JA},R}(0)$ is given by Eqs. (5.23a,b) for the k_t and the Cambridge/Aachen algorithms respectively, and it is equal to zero for the anti- k_t and SISCone algorithms. The perturbative modification $\langle \Delta \Sigma_{\text{JA},R}^2 \rangle$ is itself now driven by two mechanisms: the fact that the second particle causes the average area to change, and that it also causes modifications of the fluctuations associated with the sampling over many ghost sets. We therefore write

$$\langle \Delta \Sigma_{\text{JA},R}^2 \rangle \simeq S_{\text{JA},R}^2 \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_{t1})}, \quad (5.37)$$

$$S_{\text{JA},R}^2 = \int_0^{\frac{d\theta}{\theta}} [(A_{\text{JA},R}(\theta) - A_{\text{JA},R}(0))^2 + \Sigma_{\text{JA},R}^2(\theta) - \Sigma_{\text{JA},R}^2(0)] \quad (5.38)$$

$$= \int_0^{\frac{d\theta}{\theta}} (A_{\text{JA},R}^2(\theta) - A_{\text{JA},R}^2(0)) - 2A_{\text{JA},R}(0)D_{\text{JA},R}, \quad (5.39)$$

where as usual we neglect contributions that are not enhanced by any logarithm or that are higher-order in α_s . The details of how to obtain these results are given in Appendix D.2.

The coefficient $S_{\text{JA},R}^2$ can be determined only numerically for the k_t and Cambridge/Aachen algorithms, while for the SISCone algorithm the result can be deduced from Eq. (5.31) together with the knowledge that $\Sigma_{\text{SISCone},R}(\theta) \equiv 0$:

$$S_{\text{anti-}k_t,R}^2 \simeq 0, \quad (5.40a)$$

$$S_{k_t,R}^2 \simeq (0.41 \pi R^2)^2, \quad (5.40b)$$

$$S_{\text{Cam},R}^2 \simeq (0.19 \pi R^2)^2, \quad (5.40c)$$

$$S_{\text{SISCone},R}^2 \simeq (0.0738 \pi R^2)^2. \quad (5.40d)$$

Again, both the values and their ordering are similar to what we have obtained for the passive areas (see Eq. (5.20)).

5.1.2.4 n -particle properties and large- n behaviour

Anti- k_t algorithm. If we consider a situation with a single hard particle p_1 , and multiple soft emissions p_i , $i \geq 2$, with $p_{t1} \gg p_{ti}$, the situation for the anti- k_t jet areas does not change compared to what we have seen so far, with both the passive and active areas of the jet containing p_1 being simply πR^2 . This is the main property of the anti- k_t algorithm, namely that the hard (isolated) jets are circular. Beside the fact that this soft-resilience makes it a nice algorithm in an experimental framework, it also has important consequences when studying soft gluon emissions analytically in QCD. For example, the calculation of non-global logarithm contributions will be considerably simplified and the Milan factor will take the “universal” value, 1.49 for 3 active non-perturbative flavours.

In the limit where multiple particles are hard and become close to each other, the anti- k_t areas can depart from πR^2 . We shall comment a bit more on this below, when we discuss back-reaction.

k_t algorithm. As for the passive area, the k_t algorithm’s active area has the property that it can be expressed as a sum of individual particle areas:

$$A_{k_t,R}(J) = \sum_{p_i \in J} A_{k_t,R}(p_i). \quad (5.41)$$

This is the case because the presence of the momentum scale k_t in the distance measure means that all ghosts cluster among themselves and with the hard particles, before any of the hard particles start clustering between themselves. However in contrast to the passive-area situation, there is no known simple geometrical construction for the individual particle area.

Equivalence of all areas for large n The existence of different area definitions is linked to the ambiguity in assigning ‘empty space’ to any particular jet. In the presence of a sufficiently large number of particles n , one expects this ambiguity to vanish because real particles fill up more and more of the empty space and thus help to sharpen the boundaries between jets. Thus in the limit of many particles, all sensible area definitions should give identical results for a given jet.

To quantify this statement, we examine (a bound on) the scaling law for the relation between the density of particles and the magnitude of the potential differences between different area definitions. We consider the limit of ‘dense’ coverage, defined as follows: take square tiles of some size and use them to cover the rapidity–azimuth cylinder (up to some maximal rapidity). Define λ as the smallest value of tile edge-length such that all tiles contain at least one particle. In an event consisting of uniformly distributed particles, λ is of the same order of magnitude as the typical inter-particle distance. The event is considered to be dense if $\lambda \ll R$.

Now let us define a boundary tile of a jet to be a tile that contains at least one particle of that jet and also contains a particle from another jet or has an adjacent tile containing one or more particle(s) belonging to a different jet. We expect that the difference between different jet-area definitions cannot be significantly larger than the total area of the boundary tiles for a jet.

The number of boundary tiles for jets produced by a given jet algorithm (of radius R) may scale in a non-trivial manner with the inter-particle spacing λ , since the boundary may well have a fractal structure. We therefore parametrise the average number of boundary tiles for a jet, $N_{b,JA,R}$ as

$$N_{b,JA,R} \sim \left(\frac{R}{\lambda}\right)^{\Upsilon_{JA}}, \quad (5.42)$$

where the fractal dimension $\Upsilon_{JA} = 1$ would correspond to a smooth boundary. The total area of these boundary tiles gives an upper limit on the ambiguity of the jet area

$$\langle |a_{JA,R} - A_{JA,R}| \rangle \lesssim N_{b,JA,R} \lambda^2 \sim R^{\Upsilon_{JA}} \lambda^{2-\Upsilon_{JA}}, \quad (5.43)$$

and similarly for the difference between active or passive and Voronoi areas. As long as $\Upsilon_{JA} < 2$ the differences between various area definitions are guaranteed to vanish in the infinitely dense limit, $\lambda \rightarrow 0$. We note that $\Upsilon_{JA} = 2$ corresponds to a situation in which the boundary itself behaves like an area, i.e. occupies the same order of magnitude of space as the jet itself. This would be visible in plots representing jet active areas (such as Fig. 2.2), in the form of finely intertwined jets. We have seen no evidence for this and therefore believe that $1 \leq \Upsilon_{JA} < 2$ for all three jet algorithms considered here.

In practice we expect the difference between any two area definitions to vanish much more rapidly than Eq. (5.43) as $\lambda \rightarrow 0$, since the upper bound will only be saturated if, for every tile, the difference between two area definitions has the same sign. This seems highly unlikely. If instead differences in the area for each tile are uncorrelated (but each of order λ^2) then one would expect to see

$$\langle |a_{JA,R} - A_{JA,R}| \rangle \sim \sqrt{N_{b,JA,R}} \lambda^2 \sim R^{\Upsilon_{JA}/2} \lambda^{2-\Upsilon_{JA}/2}. \quad (5.44)$$

We have measured the fractal dimension for the k_t and Cambridge/Aachen algorithms and find $\Upsilon_{k_t} \simeq \Upsilon_{\text{Cam}} \simeq 1.20 - 1.25$.¹³ Note that any measurement of the fractal dimension of jet algorithms in real data would be severely complicated by additional structure added to jets by QCD branching, itself also approximately fractal in nature.

¹³This has been measured on pure ghost jets, because their higher multiplicity facilitates the extraction of a reliable result, however we strongly suspect that it holds also for single-particle jets.

The fact that active and passive (or Voronoi) areas all give the same result in dense events has practical applications in real-life situations where an event is populated by a very large number of particles (heavy ion collisions being an example). In this case it will be possible to choose the area type which is fastest to compute (for instance the Voronoi area) and use the results in place of the active or passive one.

5.1.3 Back reaction

So far we have considered how a set of infinitely soft ghosts clusters with a hard jet, examining also cases where the jet has some finitely-soft substructure. This infinitely-soft approximation for the ghosts is not only adequate, but also necessary from the point of view of properly defining jet areas. However if we are to understand the impact of pileup and underlying-event radiation on jet finding — the original motivation for studying areas — then we should take into account the fact that these contributions provide a dense background of particles at some small but *finite* soft scale $\sim \Lambda_{QCD}$.

This has two main consequences. The first (more trivial) one is that pileup or underlying event can provide an alternative, dynamic infrared cutoff in the p_{t2} integration in equations such as Eq. (5.11): assuming that the density of PU transverse momentum per unit area is given by ρ then one can expect that when $p_{t2} \ll \pi R^2 \rho$, the presence of p_2 will no longer affect the clustering of the ghosts (i.e. the PU particles). In this case the p_{t2} integral will then acquire an effective infrared cutoff $\sim \pi R^2 \rho$ and in expressions such as Eqs. (5.13,5.14), Q_0 will be replaced by $\pi R^3 \rho$. Note that neither with infinitely nor finitely-soft ghosts do we claim control over the coefficient in front of this cutoff, though we do have confidence in the prediction of its R dependence for small R .

The second consequence of the finite softness of the PU contribution is that the addition of the PU particles can modify the set of non-PU particles that make it into a jet. This is the effect we called *back reaction* in our introductory discussion (see Section 2.2).

That back reaction should happen is quite intuitive since it concerns non-PU particles whose softness is commensurate with the PU scale. However the extent to which it occurs depends significantly on the jet algorithm. Furthermore for some algorithms it can also occur (rarely) even for non-PU particles that are much harder than the PU scale.

As with the studies of areas, there are two limits that can usefully be examined for back-reaction: the illustrative and mathematically simpler (but less physical) case of a pointlike minimum-bias background and the more realistic case with diffuse pileup radiation.

5.1.3.1 Back reaction from pointlike minimum-bias

Let us first calculate back reaction in the case of pointlike minimum bias. We will consider minimum-bias particles with transverse momentum p_{tm} distributed uniformly on the y - ϕ cylinder with density $\nu_m \ll 1$. We use a subscript m rather than g to differentiate them from ghost particles, the key distinction being that p_{tm} is small but finite, where p_{tg} is infinitesimal.

We thus consider the situation in which a particle p_1 , with large transverse momentum, $p_{t1} \gg p_{tm}$, has emitted a soft particle p_2 on a scale commensurate with the minimum-bias particles, $p_{t2} \sim p_{tm}$. We shall calculate the probability that p_2 was part of the jet in the absence of the minimum-bias particle, but is *lost* from it when the minimum-bias particle is added. This can be written

$$\frac{dP_{JA,R}^{(\text{loss})}}{dp_{t2}} = \int d\phi_m dy_m \nu_m \int d\Delta_{12} \frac{dP}{dp_{t2} d\Delta_{12}} H_{JA,R}(p_2 \in J_1) H_{JA,R}(p_2 \notin J_1 | p_m), \quad (5.45)$$

where $H_{JA,R}(p_2 \in J_1)$ is 1 (0) if, in the absence of p_m , p_2 is inside (outside) the jet that contains p_1 . Similarly, $H_{JA,R}(p_2 \notin J_1|p_m)$ is 1 (0) if, in the presence of p_m , p_2 is outside (inside) the jet that contains p_1 . One can also define the probability for p_2 to not be part of the jet in the absence of the minimum-bias particle, but to be *gained* by the jet when the minimum-bias particle is added,

$$\frac{dP_{JA,R}^{(\text{gain})}}{dp_{t2}} = \int d\phi_m dy_m \nu_m \int d\Delta_{12} \frac{dP}{dp_{t2} d\Delta_{12}} H_{JA,R}(p_2 \notin J_1) H_{JA,R}(p_2 \in J_1|p_m). \quad (5.46)$$

It is convenient to factor out the particle production probability as follows, in the small R limit,

$$\frac{dP_{JA,R}^{(\text{loss,gain})}}{dp_{t2}} = \Delta_{12} \frac{dP}{dp_{t2} d\Delta_{12}} \Big|_{\Delta_{12}=R} \nu_m b_{JA,R}^{(\text{loss,gain})}(p_{t2}/p_{tm}), \quad (5.47)$$

where $b_{JA,R}^{(\text{loss})}(p_{t2}/p_{tm})$ ($b_{JA,R}^{(\text{gain})}(p_{t2}/p_{tm})$) can be thought of as the effective ‘back-reaction area’ over which the minimum-bias particle causes a loss (gain) of jet contents, given a $d\Delta_{12}/\Delta_{12}$ angular distribution for the jet contents:¹⁴

$$b_{JA,R}^{(\text{loss})}(p_{t2}/p_{tm}) = \int d\phi_m dy_m \int \frac{d\Delta_{12}}{\Delta_{12}} H_{JA,R}(p_2 \in J_1) H_{JA,R}(p_2 \notin J_1|p_m), \quad (5.48)$$

$$b_{JA,R}^{(\text{gain})}(p_{t2}/p_{tm}) = \int d\phi_m dy_m \int \frac{d\Delta_{12}}{\Delta_{12}} H_{JA,R}(p_2 \notin J_1) H_{JA,R}(p_2 \in J_1|p_m). \quad (5.49)$$

Let us first consider the k_t and Cambridge/Aachen sequential-recombination algorithms for which the H functions in Eqs. (5.48,5.49) translate to a series of Θ -functions, e.g.

$$H_{k_t,R}(p_2 \in J_1) H_{k_t,R}(p_2 \notin J_1|p_m) = \Theta(R - \Delta_{12}) \Theta(R - \Delta_{2m}) \Theta(\Delta_{1(2+m)} - R) \times \\ \times \Theta(\Delta_{1m} - \min(1, p_{t2}/p_{tm}) \Delta_{2m}) \Theta(\Delta_{12} - \min(1, p_{tm}/p_{t2}) \Delta_{2m}), \quad (5.50)$$

for the k_t algorithm and

$$H_{\text{Cam},R}(p_2 \in J_1) H_{\text{Cam},R}(p_2 \notin J_1|p_m) = \Theta(R - \Delta_{12}) \Theta(R - \Delta_{2m}) \Theta(\Delta_{1(2+m)} - R) \times \\ \times \Theta(\Delta_{1m} - \Delta_{2m}) \Theta(\Delta_{12} - \Delta_{2m}), \quad (5.51)$$

for Cambridge/Aachen, where $\Delta_{1(2+m)}$ is the distance between p_1 and the recombined $p_2 + p_m$. Evaluating integrals with the above Θ -functions is rather tedious, but one can usefully consider the limit $p_{tm} \ll p_{t2} \ll p_{t1}$. This is of physical interest because it relates to the probability that the minimum-bias particle induces changes in jet momentum that are much larger than p_{tm} , and a number of simplifications occur in this limit. Since

$$\Delta_{1(2+m)} = \left| \vec{\Delta}_{12} + \frac{p_{tm}}{p_{t2} + p_{tm}} \vec{\Delta}_{2m} \right| = \Delta_{12} + \frac{p_{tm}}{p_{t2}} \frac{\vec{\Delta}_{12} \cdot \vec{\Delta}_{2m}}{\Delta_{12}} + \mathcal{O}\left(\frac{p_{tm}^2}{p_{t2}^2} R\right), \quad (5.52)$$

p_2 must be close to the edge of the jet in order for it to be pulled out by p_m , $|\Delta_{12} - R| \ll 1$. Without loss of generality, we can set $y_1 = \phi_1 = \phi_2 = 0$, so that $\Delta_{12} = y_2 \simeq R$, and

$$\Delta_{1(2+m)} = y_2 + \frac{p_{tm}}{p_{t2}} (y_m - R) + \mathcal{O}\left(\frac{p_{tm}^2}{p_{t2}^2} R\right). \quad (5.53)$$

¹⁴Strictly speaking it is the integral over area of the probability of causing a loss (or gain) of jet contents.

We can then carry out the integrations over ϕ_m and y_2 straightforwardly, leading to following the result for loss at high p_{t2} ,

$$b_{k_t,R}^{(\text{loss})}(p_{t2}/p_{tm} \gg 1) \simeq b_{\text{Cam},R}^{(\text{loss})}(p_{t2}/p_{tm} \gg 1) \simeq \int_R^{2R} dy_m \frac{(y_m - R) p_{tm}}{R p_{t2}} 2\sqrt{R^2 - (y_m - R)^2} = \frac{2}{3} \frac{p_{tm}}{p_{t2}} R^2. \quad (5.54)$$

For the anti- k_t algorithm it is tempting to follow the same path, i.e. to impose that the loss term is obtained by clustering p_m with p_2 in order to pull p_2 out of the jet, and write the H functions (for the loss term and assuming $p_{t2} \gg p_{tm}$) as

$$\Theta(R - \Delta_{12}) \Theta(\Delta_{1(2+m)} - R) \Theta(\Delta_{1m} - (p_{t1}/p_{t2})\Delta_{2m}) \Theta(\Delta_{12} - (p_{t1}/p_{t2})\Delta_{2m}) \Theta(R - \Delta_{2m}). \quad (5.55)$$

In the limit of interest, $p_{tm} \ll p_{t2} \ll p_{t1}$, the constraint that p_2 must be in the jet and that $p_2 + p_m$ must be outside the jet imposes, as before (see (5.53)), that Δ_{12} must be between $R - \frac{p_{tm}}{p_{t2}}(y_m - R)$ and R . But now, the condition that p_2 clusters with p_m and not p_1 (the fourth term in (5.55)) imposes that p_m lies within a half-circle of radius $R p_{t2}/p_{t1}$ centred on p_2 . Ultimately, this gives a contribution proportional to $\frac{p_{t2}^2 p_{tm}}{p_{t1}^3} R^2$ which, as we shall show below, is subleading for the case of interest, $p_{t2} \ll p_{t1}$.

Instead, the dominant contribution comes from the situation where p_m clusters with p_1 and the recoil causes p_2 to fall outside the jet, giving

$$H_{\text{anti-}k_t,R}(p_2 \in J_1) H_{\text{anti-}k_t,R}(p_2 \notin J_1 | p_m) = \Theta(R - \Delta_{12}) \Theta(\Delta_{(1+m)2} - R) \times \\ \times \Theta(\Delta_{1m} - \Delta_{12}) \Theta(\Delta_{12} - \min(p_{t1}/p_{t2}, p_{t1}/p_{tm})\Delta_{2m}) \Theta(R - \Delta_{1m}). \quad (5.56)$$

In the limit $p_{t1} \gg p_{t2}, p_{tm}$, the recoil (along the direction $\vec{\Delta}_{12}$) is found to be $p_{tm}/p_{t1} y_m$ and this ultimately leads to

$$b_{\text{anti-}k_t,R}^{(\text{loss})} \simeq \frac{2}{3} \frac{p_{tm}}{p_{t1}} R^2, \quad (5.57)$$

i.e. suppressed by a factor p_{t2}/p_{t1} compared to the result obtained in (5.54) for the k_t and Cambridge/Aachen algorithms.¹⁵ Note also that, contrary to the results for the k_t and Cambridge/Aachen algorithms, (5.57) is valid independently of the ordering between p_{t2} and p_{tm} , as long as they are both small compared to p_{t1} .

In a similar manner we obtain for the gain,

$$b_{k_t,R}^{(\text{gain})}(p_{t2}/p_{tm} \gg 1) \simeq b_{\text{Cam},R}^{(\text{gain})}(p_{t2}/p_{tm} \gg 1) \simeq \int_{R/2}^R dy_m \frac{(R - y_m) p_{tm}}{R p_{t2}} 2\sqrt{R^2 - (R - y_m)^2} = \\ = \left(\frac{2}{3} - \frac{\sqrt{3}}{4} \right) \frac{p_{tm}}{p_{t2}} R^2. \quad (5.58)$$

and for the anti- k_t algorithm,

$$b_{\text{anti-}k_t,R}^{(\text{gain})} \simeq b_{\text{anti-}k_t,R}^{(\text{loss})} \simeq \frac{2}{3} \frac{p_{tm}}{p_{t1}} R^2, \quad (5.59)$$

¹⁵A similar contribution, with the same p_{tm}/p_{t1} scaling, would also be present for the k_t and Cambridge/Aachen algorithms but would be subleading compared to the result of (5.54).

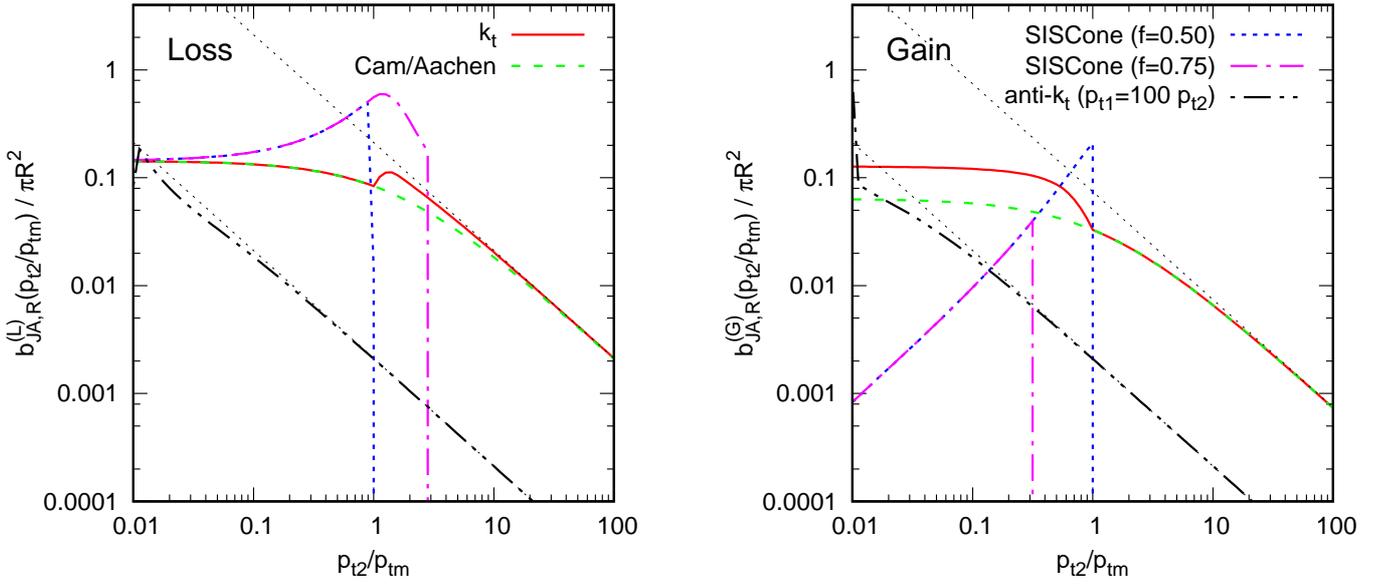


Figure 5.10: The effective area for back-reaction as a function of the ratio of the soft perturbative scale p_{t2} and the point-like minimum-bias scale p_{tm} , showing separately the loss (left) and gain (right) components for four jet definitions. For the anti- k_t results, we have set $p_{t1} = 100 p_{t2}$. The thin black dotted lines represent our analytic results valid for $p_{t2} \gg p_{tm}$.

The SIScone results are included and have the property that

$$b_{\text{SIScone},R}^{(\text{loss})}(p_{t2}/p_{tm}) = 0 \quad \text{for} \quad \frac{p_{t2}}{p_{tm}} > \frac{f}{1-f}, \quad (5.60a)$$

$$b_{\text{SIScone},R}^{(\text{gain})}(p_{t2}/p_{tm}) = 0 \quad \text{for} \quad \frac{p_{t2}}{p_{tm}} > \frac{1-f}{f}. \quad (5.60b)$$

i.e. at high p_{t2} , point-like minimum bias never induces back-reaction in the cone algorithm.¹⁶

The results for general p_{t2}/p_{tm} , determined numerically, are shown in figure 5.10. One sees that our analytic estimates work very well in the region where $p_{t2} \gg p_{tm}$ and that indeed, the anti- k_t algorithm has a suppressed back reaction and SIScone has zero back-reaction at large p_{t2} as expected. On the other hand, for $p_{t2} \sim p_{tm}$ back-reaction is more likely with the cone algorithm — the effective area over which the MB particle can cause a change in jet contents is $\sim 0.5\pi R^2$, to be compared to $\sim 0.1\pi R^2$ for the k_t and Cambridge/Aachen algorithms.

One may use the results Eqs. (5.54)–(5.60) to determine the average change in jet-momentum due to back reaction. Because of the logarithmic spectrum of emissions $dP/(dp_{t2} d\Delta_{12})$, one finds that it receives contributions from the whole logarithmic region $p_{tm} < p_{t2} < p_{t1}$,

$$\langle \Delta p_{t,JA,R}^{(\text{gain}-\text{loss})} \rangle \simeq \int_{p_{tm}}^{p_{t1}} dp_{t2} p_{t2} \left[\frac{dP_{JA,R}^{(\text{gain})}}{dp_{t2}} - \frac{dP_{JA,R}^{(\text{loss})}}{dp_{t2}} \right] = \beta_{JA,R} \rho \cdot \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(p_{tm} R)}{\alpha_s(p_{t1} R)}, \quad (5.61)$$

(evaluated for running coupling), where $\rho = \nu_m p_{tm}$ corresponds to the average transverse momentum

¹⁶For large but finite p_{t1} , back-reaction can actually occur beyond the above limits, but only with probability $\sim p_{tm}/p_{t1}$.

of minimum-bias radiation per unit area and

$$\beta_{\text{JA},R} = \lim_{p_{t2} \rightarrow \infty} \frac{p_{t2}}{p_{tm}} \left(b_{\text{JA},R}^{(\text{gain})}(p_{t2}/p_{tm}) - b_{\text{JA},R}^{(\text{loss})}(p_{t2}/p_{tm}) \right). \quad (5.62)$$

The structure of the correction in Eq. (5.61) is very similar to that for the actual contamination from minimum bias, $\rho \langle \Delta a_{\text{JA},R} \rangle$, with $\langle \Delta a_{\text{JA},R} \rangle$ as determined in section 5.1.1.3: notably, for fixed coupling, the average back-reaction scales with the logarithm of the jet p_t . The coefficients $\beta_{\text{JA},R}$,

$$\beta_{k_t,R} = \beta_{\text{Cam},R} = -\frac{\sqrt{3}}{4} R^2 \simeq -0.1378\pi R^2, \quad (5.63a)$$

$$\beta_{\text{anti-}k_t,R} = \beta_{\text{SISCone},R} = 0. \quad (5.63b)$$

can be directly compared to the results for the $d_{\text{JA},R}$ there. The values are relatively small, similar in particular to what one observes for the Cambridge/Aachen algorithm, though of opposite sign. Note also that for the SISCone and anti- k_t algorithms, one would expect a contribution similar to (5.61), proportional to ρ , but without the logarithmic enhancement, although, in the case of the anti- k_t algorithm, further cancellations could happen since $b_{\text{anti-}k_t,R}^{(\text{gain})} \simeq b_{\text{anti-}k_t,R}^{(\text{loss})}$.

Though the average change in jet momentum, both from scaling violations of the area and from back-reaction, have a similar analytical structure, it is worth bearing in mind that these similar analytical structures come about quite differently in the two cases. Regarding area scaling violations, a significant fraction of jets, $\sim \alpha_s \ln p_{t1}/p_{tm}$, are subject to a change in area $\sim R^2$ (cf. section 5.1.1.3), and a consequent modification of the minimum-bias contamination by a modest amount $\sim p_{tm}$. In contrast the average back reaction effect $\sim \alpha_s p_{tm} \ln p_{t1}/p_{tm}$, is due to large modifications of the jet momentum $\sim p_{t2}$ with $p_{tm} \ll p_{t2} \ll p_{t1}$ occurring rarely, with a differential probability that depends on p_{t2} as $\sim \alpha_s dp_{t2} p_{tm} / p_{t2}^2$. One consequence of this is that the mean square change in transverse momentum due to back-reaction is dominated by very rare modifications of jets in which $p_{t2} \sim p_{t1}$, giving

$$\left\langle \left(\Delta p_{t,\text{JA},R}^{(\text{gain,loss})} \right)^2 \right\rangle \sim \alpha_s p_{t1} p_{tm} \nu_m. \quad (5.64)$$

Note that the coefficient of this dispersion is non-zero even for the SISCone algorithm, due to the residual probability $\sim p_{tm}/p_{t1}$ that it has, for finite p_{t1} , for any change in structure with $p_{tm} \ll p_{t2} \lesssim p_{t1}$. For the anti- k_t algorithm, one would also obtain the behaviour given by Eq. (5.64) since, for $p_{t2} \sim p_{t1}$, the gain and loss probabilities also behave as $\sim \alpha_s dp_{t2} p_{tm} / p_{t2}^2$.

5.1.3.2 Back reaction from diffuse pileup

Let us now examine back reaction in the case where the minimum-bias radiation has a uniform diffuse structure, consisting of a high density of pileup particles, $\nu_m \gg 1$. The relation between back reaction in the point-like and diffuse cases is rather similar to the relation between passive and active areas: key analytical features remain unaffected by the change in the structure of the soft background, but certain specific coefficients change.

We define the probability for loss in the presence of diffuse PU as

$$\frac{dP_{\text{JA},R}^{(\text{loss})}}{dp_{t2}} = \lim_{\substack{\nu_m \rightarrow \infty \\ \rho = \nu_m \langle p_{tm} \rangle \text{ fixed}}} \left\langle \int d\Delta_{12} \frac{dP}{dp_{t2} d\Delta_{12}} H_{\text{JA},R}(p_2 \in J_1) H_{\text{JA},R}(p_2 \notin J_1 | \rho) \right\rangle_{\text{PU}}, \quad (5.65)$$

where the average is performed over the ensemble of PU configurations, and now $H_{JA,R}(p_2 \notin J_1 | \rho)$ is 1 (0) if p_2 is outside (inside) the jet containing p_1 in the presence of the specific pileup instance. A similar equation holds for the gain probability.

Then, as with Eq. (5.47) we factorise this,

$$\frac{dP_{JA,R}^{(\text{loss})}}{dp_{t2}} = \Delta_{12} \frac{dP}{dp_{t2} d\Delta_{12}} \Big|_{\Delta_{12}=R} B_{JA,R}^{(\text{loss})}(p_{t2}/\rho), \quad (5.66)$$

into one piece related to the probability for perturbative emission and a second piece $B_{JA,R}^{(\text{loss})}$ that is the diffuse analogue of the effective back-reaction area $b_{JA,R}^{(\text{loss})}$ in the point-like case. Note however that it is not so obvious exactly what geometrical area $B_{JA,R}^{(\text{loss})}$ actually corresponds to.¹⁷ Similarly, we introduce $dP_{JA,R}^{(\text{gain})}/dp_{t2}$ and $B_{JA,R}^{(\text{gain})}$ corresponding to the gain in the presence of a diffuse PU.

As for the point-like case, we can obtain analytical results for $B_{JA,R}^{(\text{loss,gain})}(p_{t2}/\rho)$ in the case of the SISCone algorithm, for which (with $f > f_{\text{max}} \simeq 0.391$)

$$B_{\text{SISCone},R}^{(\text{loss})}(p_{t2}/\rho) = B_{\text{SISCone},R}^{(\text{gain})}(p_{t2}/\rho) = 0. \quad (5.67)$$

This is a consequence of the facts (a) that the addition of a uniform background of PU particles has no effect on the stability (or instability) of a specific cone, (b) that for $p_{t2} \ll p_{t1}$ the split-merge step is immaterial to the jet finding if p_2 is within the cone around p_1 , and (c) that for p_2 outside the cone around p_1 , the maximal possible overlap is of p_2 's stable cone with that of p_1 is f_{max} and if $f > f_{\text{max}}$ then the two cones will always be split, ensuring that p_2 remains in a jet distinct from p_1 . We believe that real-life corrections to the zero in Eq. (5.67) are proportional to the standard deviation, σ , of PU transverse-momentum density from point to point in the event.

Numerical results for the $B_{JA,R}^{(\text{loss,gain})}(p_{t2}/\rho)$ for the k_t and Cambridge/Aachen algorithms are given in figure 5.11 and compared to the results in the point-like case. One sees that the general functional form is rather similar though the normalisations are somewhat smaller (by a factor of 2 for loss, a factor ~ 10 for gain). The asymptotic large- p_{t2} behaviours are observed to be

$$B_{k_t,R}^{(\text{loss})}(p_{t2}/\rho) \simeq B_{\text{Cam},R}^{(\text{loss})}(p_{t2}/\rho) \simeq 0.11 \pi R^2 \frac{\rho}{p_{t2}} \quad (5.68a)$$

$$B_{k_t,R}^{(\text{gain})}(p_{t2}/\rho) \simeq B_{\text{Cam},R}^{(\text{gain})}(p_{t2}/\rho) \simeq 0.013 \pi R^2 \frac{\rho}{p_{t2}} \quad (5.68b)$$

The case of the anti- k_t algorithm is slightly more delicate. With a uniform background, the momentum imbalance, due to the background within a distance R of p_1 clustered instead with particle p_2 , will be of order $\pi/2 R^2 \rho (p_{t2}/p_{t1})^2$, with a corresponding recoil of order $\pi/2 R^3 (\rho/p_{t1})(p_{t2}/p_{t1})^2$. This means that one would obtain

$$B_{\text{anti-}k_t,R}^{(\text{loss})}(p_{t2}/\rho) \simeq B_{\text{anti-}k_t,R}^{(\text{gain})}(p_{t2}/\rho) \simeq \frac{1}{2} \frac{\pi R^2 \rho p_{t2}^2}{p_{t1} p_{t1}^2}. \quad (5.69)$$

One should however also notice that, as for SISCone, one should expect corrections proportional to the PU fluctuations σ which, as for the point-like case, should be proportional to $\pi R^2 \sigma / p_{t1}$. For $p_{t2} \ll p_{t1}$ the latter is expected to be the dominant behaviour.

¹⁷A related issue is that the precise choice of normalisation of $B_{JA,R}^{(\text{loss})}$ is somewhat arbitrary — our specific choice is intended to provide a meaningful connection with $b_{JA,R}^{(\text{loss})}$ in the large p_{t2} limit.

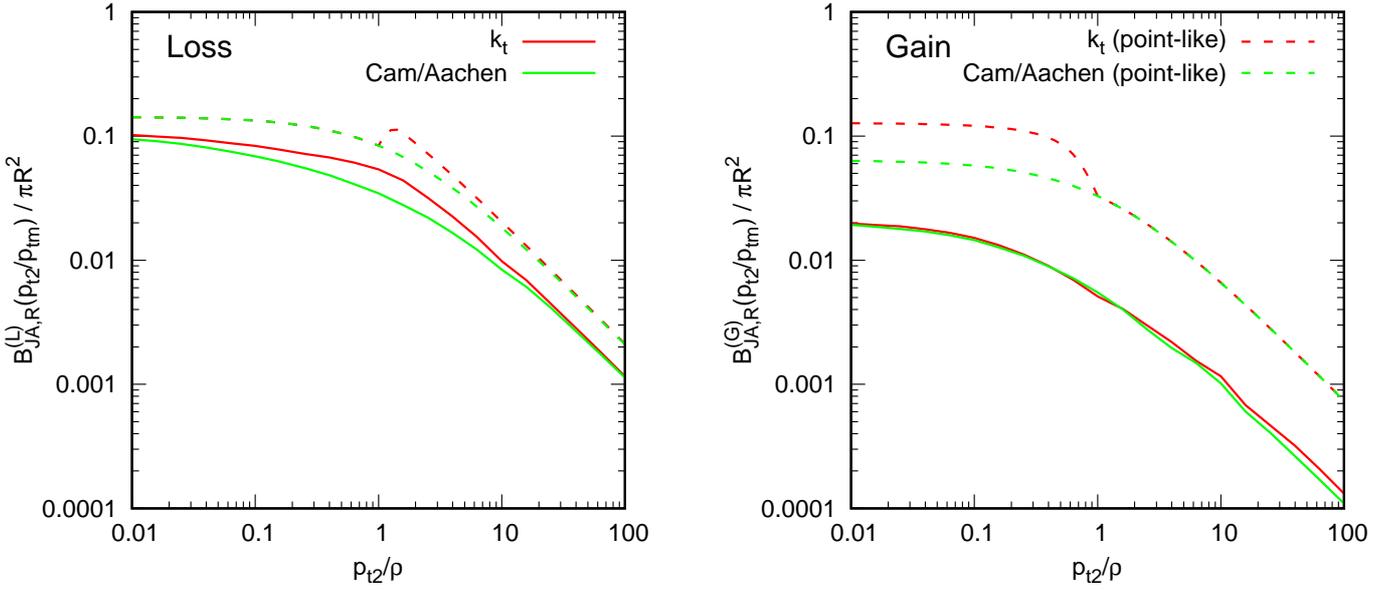


Figure 5.11: Numerical results for the diffuse effective back-reaction ‘area’, $B_{JA,R}^{(\text{loss,gain})}(p_{t2}/\rho)$, for the k_t and Cambridge/Aachen algorithms, with the point-like results $b_{JA,1}^{(\text{loss,gain})}(p_{t2}/p_{tm})$ shown for comparison also. Results obtained for $R = 1$ and verified also for $R = 0.7$.

As in the point-like case we can calculate the mean change in jet transverse momentum due to back reaction and we obtain

$$\langle \Delta p_{t,JA,R}^{(\text{gain-loss})} \rangle \simeq \int_{p_{tm}}^{p_{t1}} dp_{t2} p_{t2} \left[\frac{dP_{JA,R}^{(\text{gain})}}{dp_{t2}} - \frac{dP_{JA,R}^{(\text{loss})}}{dp_{t2}} \right] = \mathcal{B}_{JA,R} \rho \cdot \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(\rho R^3)}{\alpha_s(p_{t1} R)}, \quad (5.70)$$

with

$$\mathcal{B}_{JA,R} = \lim_{p_{t2} \rightarrow \infty} \frac{p_{t2}}{\rho} \left(B_{JA,R}^{(\text{gain})}(p_{t2}/\rho) - B_{JA,R}^{(\text{loss})}(p_{t2}/\rho) \right). \quad (5.71)$$

Even though $b_{JA,R}^{(\text{loss,gain})}$ had p_{t2}/p_{tm} as its argument and $B_{JA,R}^{(\text{loss,gain})}$ has p_{t2}/ρ , the final expressions for the average back-reaction in the point-like and diffuse cases, Eqs. (5.61), (5.70), have almost identical forms — in particular the overall scale appearing in each is ρ and the only difference appears in the denominator for the argument of the logarithm. The coefficients are slightly smaller,

$$\mathcal{B}_{k_t,R} = \mathcal{B}_{\text{Cam},R} \simeq -0.10\pi R^2, \quad (5.72a)$$

$$\mathcal{B}_{\text{SIScone},R} = \mathcal{B}_{\text{anti-}k_t,R} = 0, \quad (5.72b)$$

and will again translate into modest effects compared to the overall pileup contamination in the jets. Note however that in general the scaling with R of $b_{k_t,R}^{(\text{loss,gain})}(p_{t2}/p_{tm})$ and $B_{k_t,R}^{(\text{loss,gain})}(p_{t2}/\rho)$ is subtly different. The former truly behaves like an area, in that $b_{k_t,R}^{(\text{loss,gain})}(p_{t2}/p_{tm})/R^2$ is R -independent; the latter instead has the property that it is $B_{k_t,R}^{(\text{loss,gain})}(R^2 p_{t2}/\rho)$ that is R -independent.

Finally, as for the case of the point-like background, the diffuse PU will lead to a mean square change of the jet due to back-reaction which will be dominated by rare modifications with $p_{t2} \sim p_{t1}$ such that

$$\left\langle \left(\Delta p_{t,JA,R}^{(\text{gain,loss})} \right)^2 \right\rangle \sim \alpha_s p_{t1} \rho. \quad (5.73)$$

This is also true for the anti- k_t algorithm since the extra suppression factor become finite for $p_{t2} \sim p_{t1}$.

5.1.4 Comparison to Monte-Carlo simulations

In this section we examine the properties of jet areas in the context of realistic events, as simulated with Herwig [41] and Pythia [35]. There are two purposes to doing so. Firstly we wish to illustrate the extent to which the simple single-gluon emission arguments of the previous section hold once one includes full parton-showering and hadronisation. Secondly, jet areas play an important role in the estimation and subtraction of underlying event and pileup contamination. The study of jet areas in realistic events can help to highlight some of the issues that arise in such a procedure.

Let us start with an investigation of the distribution of the areas of hard jets with $p_t \gtrsim 1$ TeV in simulated LHC dijet events, and first concentrate on the k_t , Cambridge/Aachen and SISCone algorithms. The area distributions are shown on Fig. 5.12 for various ‘levels’ in the Monte Carlo: just after parton showering, after hadronisation, both with and without an UE, and finally with a pileup contribution corresponding to a luminosity of 0.25 mb^{-1} per bunch crossing (i.e. $\mu \sim 25$). We examine both passive and active areas.

The passive areas distributions at parton-shower level, Fig. 5.12a, are those that are most amenable to understanding in terms of our analytical results. Firstly one notes that the SISCone and Cambridge/Aachen algorithms have a clear peak at $a = \pi R^2$. These two algorithms both have the property (cf. section 5.1.1.2) that the area is not affected by moderately collinear ($\Delta_{12} < R$ for SISCone, $\Delta_{12} < R/2$ for Cambridge/Aachen) soft particle emission. Thus it is possible, even in the presence of parton showering (which is mostly collinear), for the passive area to remain πR^2 . For the cone algorithm, the other main structure is a “shoulder” region $0.8 \lesssim a/(\pi R^2) \lesssim 1$, which coincides nicely with the range of values that can be induced by 2-particle configurations (cf. Fig. 5.2). A similar shoulder exists for the Cambridge algorithm, which however additionally has jets with $a > \pi R^2$ — again the range of values spanned, up to $a \simeq 1.6\pi R^2$, is rather similar to what would be expected from the two-particle calculations. Further broadening of these distributions at the edges is attributable to parton-level states with more than two particles. In contrast the parton-level passive area distribution for the k_t algorithm seems less directly related to the 2-particle calculations of section 5.1.1.2. This can be understood by observing that the k_t passive area is modified even by rather collinear emissions, and the multiple collinear emissions present in parton showers add together to cause substantial broadening of the area distribution.

At parton shower level, there are relatively few particles in the event and there is no obvious boundary to the jet — the ghosts that we add provide a way of assigning that empty area. It is therefore not surprising to see significant differences between the two ways of adding ghosts, i.e. the passive and active area distributions. This is most marked for SISCone, as is to be expected from the results of section 5.1.2.1, which showed that for a 1-particle jet the active area is $\pi R^2/4$. There is a trace of this result in Fig. 5.12b, where one sees that the SISCone distribution now extends down to $A = \pi R^2/4$. There is however no peak structure, presumably because even a highly collinear emission gives a slight modification of the area, cf. Fig. 5.9 (the same argument given for the absence of a peak for the passive k_t area). For the Cambridge/Aachen and k_t algorithms, there is less difference between active and passive area distributions, again as expected.

As one moves to events with more particles, for example hadron level, Fig. 5.12c and d, the particles themselves start to give a clearer outline to the jets. Thus the passive and active distributions are far more similar. This is less so for SISCone than the others, as is to be expected, for which one still sees a trace of the peak at $a = \pi R^2$ for the passive area.

For events with many particles, for example with the UE (Fig. 5.12e) and then pileup (Fig. 5.12f) added, the difference between passive and active area distributions is so small that we show only the active area result. These last two plots are the most relevant to phenomenology. A feature

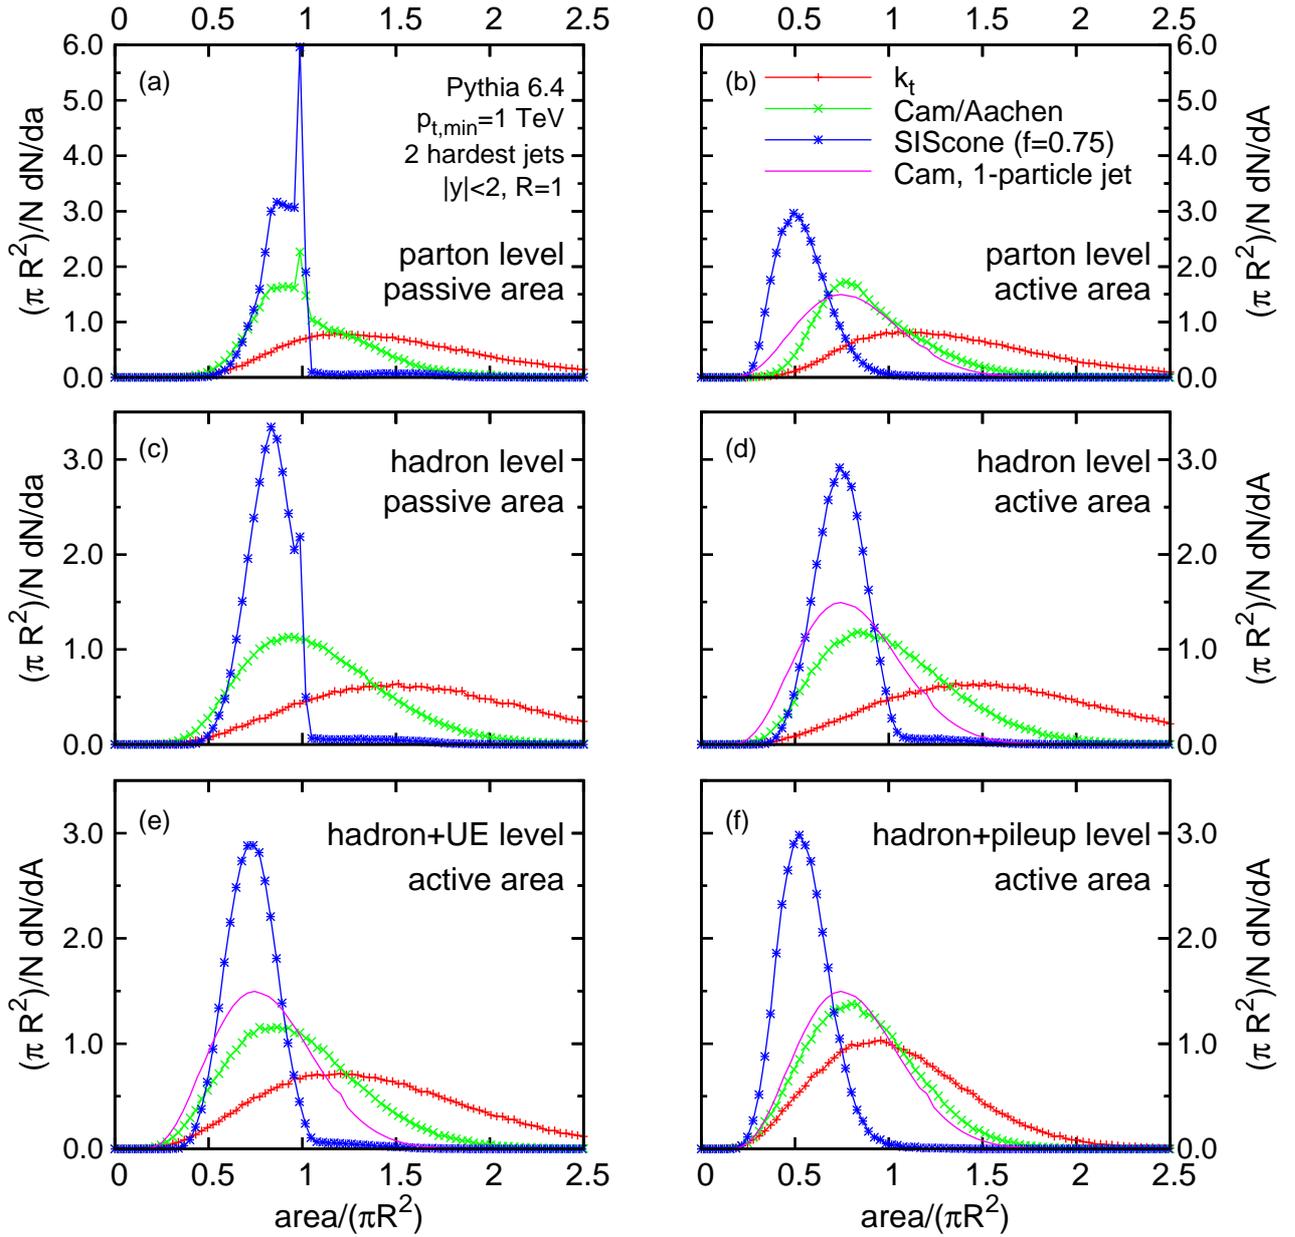


Figure 5.12: Distribution of active and passive areas of the two hardest jets in a range of simulated LHC dijet events, with a minimum p_t of 1 TeV in the Pythia event generation (v6.4, default tune). Only jets with $|y| < 2$ have been included in the histogram; ‘parton’ indicates parton-level, ‘hadron’ indicates hadron-level with the UE switched off, ‘UE’ corresponds to hadron-level with UE switched on and in the ‘pileup’ case the UE-level event is supplemented with additional pileup events corresponding to 0.25 mb^{-1} per bunch crossing (~ 25 simultaneous interactions). For all jet algorithms we use $R = 1$. A ghost area $\simeq 0.02$ was used throughout except for SIScone, where the ghost area was roughly 0.1 in the active area cases. Note that “area” in these plots corresponds to $a_{JA,R}(J_i)$ and $A_{JA,R}(J_i)$ respectively for the passive and active areas. As such the latter has been averaged over ghost ensembles (see Eq. (2.20)) and the dispersion is a consequence of the event and jet structure. We have found that 5 ghost ensembles were sufficient for k_t and Cambridge/Aachen algorithms, and 3 for SIScone (with pileup, 1 ensemble would actually have been enough). In contrast, the Cam/Aachen 1-particle jet result corresponds to $\pi R^2/N dN/dA(1\text{-particle-jet}|\{g_i\})$ and serves to illustrate how the impact of variability in the event structure in real events has consequences rather similar to that of the variability of ghost-particle ensembles in the theoretical arguments of the previous sections.

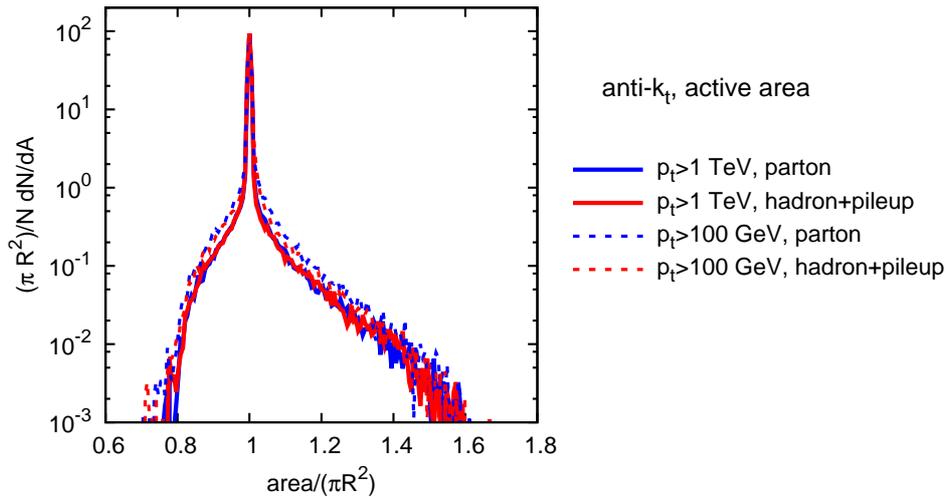


Figure 5.13: Distributions obtained in the same context as for Fig. 5.12, this time for the anti- k_t algorithm. We just show the two extreme situations, corresponding to parton-level (blue lines) and hadron level with pileup (red lines). We show the distributions for two different jet minimum transverse momentum: 100 GeV (solid lines) and 1 TeV (dotted lines).

of both is that the dispersion is smallest for SISCone (but significantly different from zero) and largest for k_t , precisely as one would expect from Eqs. (5.36) and (5.40). Another feature is how for Cambridge/Aachen the dispersion is essentially that associated with $\Sigma(0)$, the distribution being very similar to the 1-parton active area result, shown as the solid line. This similarity is strongest when there is pileup: the logarithmic enhancement that enters in Eq. (5.37) is reduced because the pileup introduces a large value for $Q_0 \sim 20$ GeV. In this case even the k_t algorithm starts to have a distribution of areas that resembles the 1-parton active area result, and for SISCone one once again sees signs of the lower limit at $A = \pi R^2/4$, the one-particle active area result.

Let us now move to area distributions for the anti- k_t algorithm. A few representative distributions are plotted in Fig. 5.13. Only the active area distribution is plotted for the two extreme particle multiplicities — the “parton” level and the “hadron+pileup” level — and for two jet transverse momenta. Most often, the jet area comes out as πR^2 as expected from our analytic studies so far. The use of a logarithmic scale on the vertical axis on Fig. 5.13 allows one to highlight the rare tails away from πR^2 , occurring for about 1% of the jets. These features are stable when we add pileup to the event or vary the jet p_t , showing that jet areas are constrained by the hard structure of the event.

While the plots of Figs. 5.12 and 5.13 provide an illustration of many of the features that we have discussed in the earlier sections of this Chapter, the fact that we have not systematically calculated area distributions means that the discussion necessarily remains qualitative. For quantitative checks, one should instead examine the mean active area and its dispersion and compare them to the results of section 5.1.2. This is done in Fig. 5.14, separately for $qq \rightarrow qq$ and for $gg \rightarrow gg$ scattering, as a function of the jet transverse momentum. The horizontal scale has been chosen uniform in $\ln \ln p_t/\Lambda$ so that our predictions correspond to straight lines.

The predictions of section 5.1.2 set the slope of those lines for each algorithm, and the agreement is reasonable in all cases. The infrared cutoff scale Q_0 is not predicted, and may differ both between quark and gluon jets and between algorithms. The values for Q_0 have therefore been fitted, using the results for the mean, and are consistent with a general non-perturbative origin (modulo issues

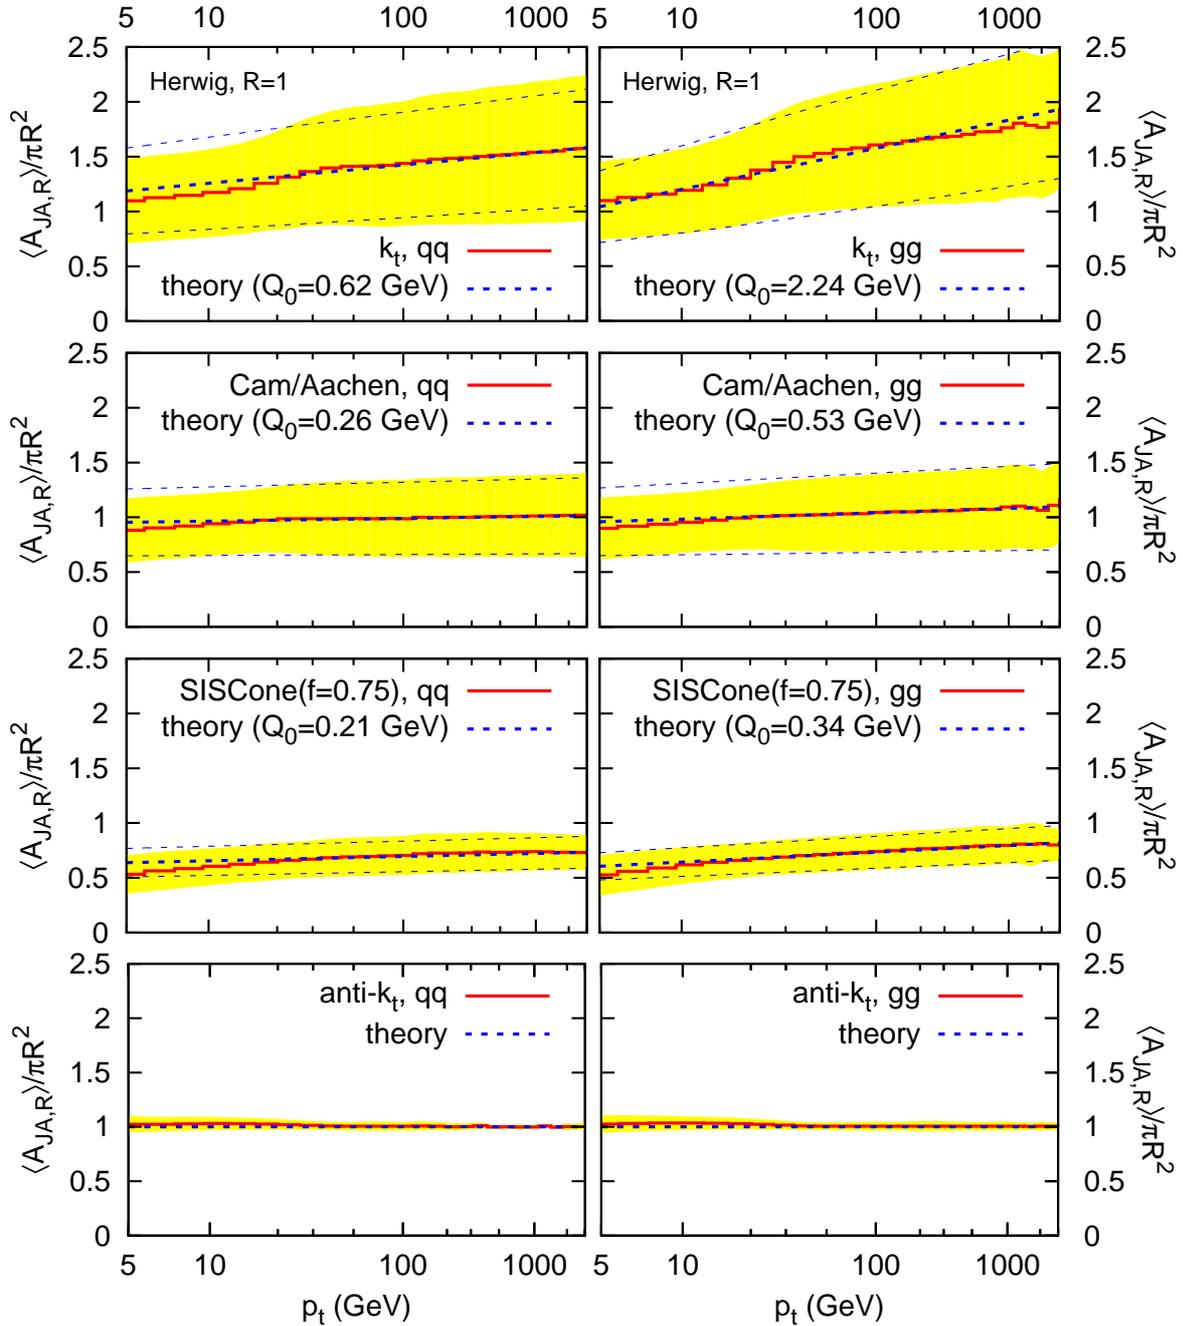


Figure 5.14: The mean (solid line) and standard deviation (band) of the active area for each of the two hardest jets in $qq \rightarrow qq$ and $gg \rightarrow gg$ events simulated with Herwig 6.5, as a function of the jet transverse momentum. The theory curves (thick dashed for mean, thin dashed for standard deviation) correspond to Eqs. (5.32)–(5.40) with a 1-loop coupling ($\Lambda_{\text{QCD}} = 0.2 \text{ GeV}$, $n_f = 5$), and Q_0 fitted (for the mean values). The areas have been obtained with ghosts of size 0.02 for anti- k_t , k_t and Cambridge/Aachen, and 0.1 for SISCone. For all algorithms, we used $R = 1$. Note that the horizontal scale is uniform in $\ln \ln p_t / \Lambda$. We considered pp collisions with $\sqrt{s} = 14 \text{ TeV}$.

in SISCone discussed below). The standard deviation (indicated by the band for the Monte Carlo simulation and the thin dashed lines for the theory result) is then entirely predicted, and also agrees remarkably well. Thus, overall, our simple analytical calculations provide a surprisingly successful picture of the mean and dispersions for various algorithms, over a range of jet transverse momenta. This is all the more striking considering that the calculation is based on just the first term in a series $\alpha_s^n \ln^n p_t$, and is in part based on a small angle approximation.

Some remarks are due regarding the Q_0 values. Clear patterns emerge: it is largest for k_t , smallest for SISCone, and systematically larger for gluon jets than for quark jets. In the case of SISCone with quark jets, the value is uncomfortably close to the value of $\Lambda_{QCD} = 0.2 \text{ GeV}$ used in the one-loop coupling. It may be that this low value is an artefact whose origin lies in the finite density both of actual event particles and of ghosts (the latter due to speed limitations in SISCone): when one has a limited density of particles and/or ghosts, the measured area may be intermediate between the passive and “ideal” active areas; because SISCone has such a large difference between passive and active areas (a factor of 4 for the 1-particle results), the finite density effects can be significant, and so Q_0 may be taking an extreme value in order to compensate for this.

Compared to k_t , Cambridge/Aachen and SISCone, the anti- k_t algorithm shows an area very close to πR^2 with almost no dependence on p_t and a very small dispersion.

A final comment concerns the choice of event generator. Here we have shown results from Herwig (v. 6.5). Pythia with the original shower (the default choice in v. 6.4) is known to have difficulties reproducing anomalous dimensions associated with soft large angle radiation [104, 105], whereas the new shower [106] mostly resolves these issues [105]. Similar concerns are potentially relevant here too. In our original studies [13], we actually found that both Pythia showers give results similar to Herwig’s (and ours) in all cases except the k_t algorithm, for which both Pythia showers give a slope a factor of two smaller than expected. This suggests that an experimental measure of the p_t dependence of jet areas might provide some non-trivial constraints on parton-showering dynamics.

5.1.5 Summary

The concept of a jet area is one that initially seems rather intuitive. Yet, as we have seen both in Chapter 2 (in particular in Section 2.3) and in this Chapter, there is a wealth of physics associated with it, both in terms of how one defines it and as concerns the interplay between a jet’s internal structure and its area. This is reflected in the range of quantities that can be studied to characterise the behaviour of jet areas, summarised in table 5.2.

Our guiding principle in defining jet areas has been that they should provide a measure of a jet’s susceptibility to additional underlying event and pileup radiation, in the limit in which these are infinitely soft. Two opposite hypotheses for the structure of such radiation, pointlike or diffuse, lead to two definitions of the area, respectively passive or active, both calculated in practice with the help of infinitely soft “ghost” particles. The two definitions can be used with any infrared safe jet algorithm, and we have studied them both with sequential recombination algorithms (anti- k_t , k_t and Cambridge/Aachen) and with a stable-cone with split–merge algorithm, SISCone.

The area of a jet may depend on its substructure. For the simplest case of a jet made of a single hard particle, the passive area coincides with one’s naive expectation of πR^2 for all jet algorithms considered here. The active area — that related to the more realistic, diffuse picture for UE and pileup — instead differs from πR^2 , except for the anti- k_t algorithm. The most surprising result is that for SISCone, whose active area is $\pi R^2/4$, a consequence of the split–merge dynamics (the midpoint algorithm widely used at the Tevatron behaves similarly). Thus the widespread assumption that cone-based algorithms automatically have an area of πR^2 is, in many cases, unjustified.

Real jets of course consist of more than a single particle. The first level of substructure involves the addition of a soft particle to the neighbourhood of the jet. For *all* except the anti- k_t algorithm, we have seen that this modifies the jet area, and that the average jet area then becomes an infrared unsafe quantity. Even if it looks like a bad feature at first sight, the infrared unsafety of jet areas is actually desirable and tightly related to the fact that areas are meant to precisely measure the jet's susceptibility to soft radiation. In that respect, the soft resilience of the anti- k_t algorithm corresponds to an area remaining πR^2 when we add any number of soft particles to the jet. Also, the small pileup contamination often observed with the SISCone algorithm can be related to its smaller catchment area than that of recombination algorithms.

From our analytic viewpoint, it emerges that the effects of gluon emission can usefully be summarised in terms of an anomalous dimension, which encodes how the jet's average area depends on its transverse momentum. We have calculated this to leading order and seen that it agrees remarkably well with the p_t dependence of measures of the jet areas in hadron-level Monte Carlo simulations.

A summary of our main results for active and passive areas, their fluctuations and their anomalous dimensions is given in table 5.3. As is visible also in figure 5.14 for a broad range of p_t , there is a hierarchy in the areas of the jet algorithms, $\langle A_{\text{SISCone},R} \rangle \lesssim \langle A_{\text{anti-}k_t,R} \rangle \sim \langle A_{\text{Cam},R} \rangle \lesssim \langle A_{k_t,R} \rangle$. A likely consequence is that SISCone jets will be the least affected by UE contamination, though the extent to which this holds depends on the precise balance between pointlike and diffuse components of the UE. The above hierarchy might also suggest an explanation for the opposite hierarchy in the size of hadronisation corrections, observed in Monte Carlo studies for the different jet algorithms in [98]. There the k_t algorithm was seen to be least affected, which now appears natural, since a larger jet is less likely to lose momentum through hadronisation.

In the context of pileup subtraction, the significant fluctuations of the area from one jet to the next (for algorithms other than anti- k_t) mean that it is important to take into account the area of each individual jet, rather than assume some typical mean value. Among results of relevance to the subtraction procedure, we highlight the demonstration that all areas (passive, active, Voronoi) are identical for highly populated events, which is important in ensuring that the subtraction procedure is free of significant ambiguities. We also remark on the calculation of back-reaction of pileup on jet structure, which is found to be a small effect, especially for the anti- k_t algorithm.

5.1.6 Future investigations

There are many avenues for potential further study of jet areas. The calculations presented here have extracted only the leading logarithmic part of the first non-trivial order in α_s , and usually we have concentrated on properties of the mean area. The Monte Carlo results in Section 5.1.4 also suggest interesting structures in the distributions of areas, and these merit further investigation.

In the specific case of the widely used anti- k_t algorithm, studying the deviations of the jet areas from πR^2 would also be of interest, especially in the context of jet calibration at the LHC. In particular, one would want to study the jet area when two hard jets become close to one another. A possibly related question is the one of the reach of the anti- k_t algorithm, where one may ask if there is a maximal area for an anti- k_t jet, or a maximal distance at which two particles can be clustered.

Stepping away from the anti- k_t algorithm, another question for future work is that of the transition between passive and active areas if one considers a continuous transformation of the ghosts from pointlike to diffuse. Since real UE and pileup contamination is neither fully pointlike, nor fully diffuse, this question is of particular relevance for the stable-cone type algorithms, for which there is a large difference between the passive and active areas.

quantity	discussed in	description
Passive area • average area $a(1PJ)$ $\langle a \rangle$ d • area fluctuations $\sigma^2(1PJ)$ $\langle \sigma^2 \rangle$ s^2	Section 5.1.1 Sect. 5.1.1.1 Eq. (5.10) Eq. (5.15) Eq. (5.17) Eq. (5.20)	single-ghost area passive area for a 1-particle jet average passive area of jets with QCD branching: $\langle a \rangle = a(1PJ) + d \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_t)} + \dots$ coefficient of leading scaling violations of $\langle a \rangle$ variance of one-particle passive area (0 by definition) variance of passive area of jets with QCD branching: $\langle \sigma^2 \rangle = \sigma^2(1PJ) + s^2 \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_t)} + \dots$ coefficient of leading scaling violations of $\langle \sigma^2 \rangle$
Active area • average area $A(1PJ)$ $\langle A \rangle$ D $A(GJ)$ • area fluctuations $\Sigma^2(1PJ)$ $\langle \Sigma^2 \rangle$ S^2 $\Sigma^2(GJ)$	Section 5.1.2 (5.21,5.22,5.27) Eq. (5.32) Eq. (5.35) Eq. (5.25) (5.21,5.23,5.28) Eq. (5.36) Eq. (5.40) Eq. (5.26)	many-ghost area (ghosts also cluster among themselves) active area for a 1-particle jet average active area of jets with QCD branching: $\langle A \rangle = A(1PJ) + D \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_t)} + \dots$ coefficient of leading scaling violations of $\langle A \rangle$ average active area for pure-ghost jets variance of one-particle active area variance of active area of jets with QCD branching: $\langle \Sigma^2 \rangle = \Sigma^2(1PJ) + S^2 \frac{C_1}{\pi b_0} \ln \frac{\alpha_s(Q_0)}{\alpha_s(Rp_t)} + \dots$ coefficient of leading scaling violations of $\langle \Sigma^2 \rangle$ variance of active area for pure-ghost jets
Back reaction • pointlike MB $dP^{(L,G)}/dp_{t2}$ $b^{(L,G)}(p_{t2}/p_{tm})$ β • diffuse MB $dP^{(L,G)}/dp_{t2}$ $B^{(L,G)}(p_{t2}/\rho)$ \mathcal{B}	Section 5.1.3 Eqs. (5.45, 5.46) Eqs. (5.48, 5.49) Eq. (5.62) Eq. (5.65) Eq. (5.66) Eq. (5.71)	action of finite-momentum min-bias (MB) on the clustering of the non-MB particles probability of jet losing (L) or gaining (G) p_{t2} worth of non-MB particles eff. area for loss, gain (MB particle has \perp mom. p_{tm}) coeff. of net high- p_{t2} gain–loss: $\lim_{p_{t2} \rightarrow \infty} \frac{p_{t2}}{p_{tm}} b^{(G-L)} \left(\frac{p_{t2}}{p_{tm}} \right)$ probability of jet losing (L) or gaining (G) p_{t2} worth of non-MB particles eff. area for loss, gain (MB has \perp -mom. density ρ) coeff. of net high- p_{t2} gain–loss: $\lim_{p_{t2} \rightarrow \infty} \frac{p_{t2}}{\rho} B^{(G-L)} \left(\frac{p_{t2}}{\rho} \right)$

Table 5.2: A summary of the mathematical quantities defined throughout this article, together with descriptions of the associated physical concepts.

	$a(1PJ)$	$A(1PJ)$	$\sigma(1PJ)$	$\Sigma(1PJ)$	d	D	s	S	$A(GJ)$	$\Sigma(GJ)$
k_t	1	0.81	0	0.28	0.56	0.52	0.45	0.41	0.55	0.17
Cam/Aachen	1	0.81	0	0.26	0.08	0.08	0.24	0.19	0.55	0.18
anti- k_t	1	0	0	0	0	0	0	0	—	—
SISCone	1	1/4	0	0	-0.06	0.12	0.09	0.07	—	—

Table 5.3: A summary of the numerical results for the main quantities worked out in the article and described in table 5.2. All results are normalised to πR^2 , and rounded to the second decimal figure. Anomalous dimensions multiply powers of $\alpha_s^n \ln^n p_t/Q_0$ that for typical jet transverse momenta sum to something of order 1. Active-area and anomalous-dimension results hold only in the small- R limit, though finite- R corrections are small. Back reaction is not easily summarised by a single number, and the reader is therefore referred to Figs. 5.10, 5.11.

Finally, an understanding of the behaviour of jet areas can play a key role in the choice of parameters for jet algorithms, as well as in the design of new algorithms. One example of this concerns SISCone, for which we saw in section 5.1.2.1 that a split–merge overlap threshold $f \simeq 0.5$ can lead to the formation of “monster jets,” whereas a choice of $f \simeq 0.75$, eliminates the problem, suggesting that the latter is a more appropriate default value.

5.2 Estimation of the pileup properties

We show in this Section that it is also possible to obtain analytic results about the estimation of ρ using from the area–median approach. The natural question in this respect is the precision that one should expect on the estimated value of ρ .

Many of the results included in this Section are adapted from Ref [19], where the authors considered several methods used to estimate event-by-event the Underlying Event characteristics at the LHC. Here, we basically follow the same steps but focus on the specific case of the area–median approach and put the derivation in the context of pileup subtraction.

We shall also discuss several other aspects like estimations of the impact of the size of the range used to estimate ρ , i.e. address the question of the minimal range size required to estimate ρ reliably.

In the following pages, we will concentrate on grid-based estimates. If, instead, we consider a jet-based determination of ρ , most of the results derived below hold with the replacement of the grid-cell size a by $\sqrt{c_J}R$ with R the jet radius and $c_J \approx 0.65\pi \approx 2.04$ (see Ref. [19]).

5.2.1 A toy model for pileup

To discuss the estimation of ρ , it is helpful to start by introducing a simple toy-model to describe the non-perturbative pileup particle distributions.

We will assume that pileup is producing, on average, ν particles per unit area and that the probability to find n particles in a patch of area A follows a Poisson distribution, i.e.

$$P_n = \frac{(\nu A)^n}{n!} e^{-\nu A}. \quad (5.74)$$

We further assume that the particles have independent transverse momenta and that the single-particle

transverse-momentum probability distribution decreases exponentially¹⁸

$$\frac{1}{P_1} \frac{dP_1}{dp_t} = \frac{1}{\mu} e^{-p_t/\mu}. \quad (5.75)$$

Both the average transverse momentum of a particle and its dispersion are therefore given by μ .

If one takes a patch of area A , containing n particles, its total p_t probability distribution is given by

$$\frac{1}{P_n} \frac{dP_n}{dp_t} = \frac{1}{(n-1)!} \frac{p_t^{n-1}}{\mu^n} e^{-p_t/\mu}. \quad (5.76)$$

From (5.74) and (5.76), it is possible to compute the overall probability distribution for the transverse momentum in the patch. One finds

$$\frac{dP}{dp_t}(A) = \sum_{n=0}^{\infty} \frac{dP_n}{dp_t} = \delta(p_t) e^{-\nu A} + e^{-\nu A - p_t/\mu} \sqrt{\frac{A\nu}{\mu p_t}} I_1 \left(2\sqrt{\frac{A\nu p_t}{\mu}} \right), \quad (5.77)$$

with $I_1(x)$ the modified Bessel function of the first kind. This distribution has a mean of $\nu A \mu$ and a standard deviation of $\sqrt{2\nu A} \mu$, where we recognise the expected scaling in μ and νA .

If we want to match this toy model to our description of pileup in terms of a deposit ρ per unit area and fluctuations σ , we have

$$\rho = \mu\nu \quad (5.78)$$

$$\sigma = \sqrt{2\nu} \mu \quad (5.79)$$

so that, in the end,

$$p_t(A) = \rho A \pm \sigma \sqrt{A}. \quad (5.80)$$

Note that, in the high-density limit, i.e. in the large pileup limit $\nu A \gg 1$, the distribution (5.77) will tend to a Gaussian.

In what follows, we shall consider typical values obtained in Section 3.1, say $\rho = \rho_1 N_{\text{PU}}$, with $\rho_1 \approx 700$ MeV, and $\sigma = \sigma_1 \sqrt{N_{\text{PU}}}$, with $\sigma_1 \approx 550$ MeV. These would correspond to $\nu = 2\rho^2/\sigma^2 \approx 3.25 N_{\text{PU}}$ and $\mu = \sigma^2/(2\rho) \approx 215$ MeV. If, instead, we use $1/P_1 dP_1/dp_t = 4p_t \mu^2 e^{-p_t/\mu}$, for which $\sigma/\rho = \sqrt{3/(2\nu)}$, one obtains $\nu = (3/2)\rho^2/\sigma^2 \approx 2.40 N_{\text{PU}}$ and $\mu = 2\sigma^2/(3\rho) \approx 290$ MeV. In both cases, we should expect a few particles per unit area for each pileup interaction.

5.2.2 Pileup distribution impact on ρ estimation

A first source of potential bias for the estimation of ρ using the area–median approach is coming from the fact that there might be a mismatch between the average and the median of the distribution (5.77). To estimate this effect, let us consider pure-pileup events, i.e. event which are just the superposition of N_{PU} minimum bias event without any hard interaction. Assuming that we use the grid-based approach with cells of area A_{cell} , the estimated ρ_{est} median would correspond to the solution of

$$\int_0^{A_{\text{cell}} \rho_{\text{est}}} dp_t \frac{dP}{dp_t}(A_{\text{cell}}) = \frac{1}{2}. \quad (5.81)$$

¹⁸Assuming a power-law behaviour would be slightly more physical but it would also complicate the calculations without changing significantly the features discussed here.

The integration above is delicate to handle analytically but one can obtain an approximate solution that is numerically close to the exact solution and keeps the right asymptotic behaviour:

$$\rho_{\text{est}}^{(\text{PU})} = \rho \frac{\nu A_{\text{cell}} - \log(2)}{\nu A_{\text{cell}} - \log(2) + \frac{1}{2}} \Theta(\nu A_{\text{cell}} > \log(2)). \quad (5.82)$$

This shows a turn-on point for $\nu A_{\text{cell}} = \log(2)$ and has the property that for asymptotic νA_{cell} one tends to the expected value for ρ according to

$$\frac{\rho_{\text{est}}^{(\text{PU})} - \rho}{\rho} = -\frac{1}{2\nu A_{\text{cell}}} + \mathcal{O}((\nu A_{\text{cell}})^{-2}). \quad (5.83)$$

This parametric behaviour is actually fairly independent of the details of the underlying 1-particle transverse momentum distribution. Although its coefficient changes in (5.83), other distribution, including those behaving like $p_t^{-m} e^{-(m+1)p_t/\mu}$ would all show a $1/(\nu A_{\text{cell}})$ bias.

It is also helpful to rewrite Eq. (5.83) in terms of the pileup properties ρ and σ . If a is the size of the grid cells, i.e. $A_{\text{cell}} = a^2$, we find

$$\frac{\rho_{\text{est}}^{(\text{PU})} - \rho}{\rho} = -\frac{\sigma^2}{4\rho^2 a^2} + \mathcal{O}\left(\frac{\sigma^4}{\rho^4 a^4}\right). \quad (5.84)$$

For the alternative model, $1/P_1 dP_1/dp_t = 4p_t\mu^2 e^{-p_t/\mu}$, the above results would become (see [19] for details)

$$\rho_{\text{est}}^{(\text{PU})} = \rho \sqrt{\frac{\nu A_{\text{cell}} - \log(2)}{\nu A_{\text{cell}} - \log(2) + \frac{2}{3}}} \Theta(\nu A_{\text{cell}} > \log(2)). \quad (5.85)$$

and

$$\frac{\rho_{\text{est}}^{(\text{PU})} - \rho}{\rho} = -\frac{2\sigma^2}{9\rho^2 a^2} + \mathcal{O}\left(\frac{\sigma^4}{\rho^4 a^4}\right). \quad (5.86)$$

5.2.3 Perturbative impact of the hard event

Let us now take the opposite situation and consider events without pileup or Underlying Event but with just a hard interaction. If we assume that we have selected events with at least 2 hard jets, extra jets can arise from perturbative radiation at higher orders. These ‘‘hard’’ jets will contribute to the estimation of the median, potentially biasing the estimated value of ρ towards larger values. We estimate these effects below.

To a decent approximation, we can consider that these emissions are soft, that they are distributed uniformly in azimuth and rapidity, and with a p_t spectrum proportional to $C_i \alpha_s(p_t)/p_t$, with C_i a colour factor. This means

$$\frac{dn}{dp_t dy d\phi} \approx \frac{C_i \alpha_s(p_t)}{\pi^2 p_t}. \quad (5.87)$$

We will consider emissions with p_t between an infra-red scale Q_0 , that we shall set to 1 GeV from now on, and a hard scale $Q = p_{t,\text{hard}}/2$. Assuming independent emissions, the probability distribution for the number of emissions will be Poissonian with a mean given by the integral of (5.87) over the whole phasespace available for radiation.

The average total number of “hard” jets $\langle n_h \rangle$ can be written as $\langle n_h \rangle = n_b + \langle n_p \rangle$ with n_b the number of original, Born-level, jets and $\langle n_p \rangle$ the number of perturbative emissions in the total area A_{tot} , given by

$$\langle n_p \rangle \approx \frac{C_i}{\pi^2} \int_{A_{\text{tot}}} dy d\phi \int_{Q_0}^Q \frac{dp_t}{p_t} \alpha_s(p_t) = \frac{C_i A_{\text{tot}}}{2\pi^2 \beta_0} \log \left(\frac{\alpha_s(Q_0)}{\alpha_s(Q)} \right), \quad (5.88)$$

with β_0 the 1-loop QCD β function.

For a given cell of area $A_{\text{cell}} = a^2$, the probability to have at least an emission in the cell is

$$P_{\text{em}} \approx 1 - \exp \left[\frac{C_i}{\pi^2} \int_{A_{\text{cell}}} dy d\phi \int_{Q_0}^Q \frac{dp_t}{p_t} \alpha_s(p_t) \right] \quad (5.89)$$

$$= 1 - \exp \left[\frac{C_i A_{\text{cell}}}{2\pi^2 \beta_0} \log \left(\frac{\alpha_s(Q_0)}{\alpha_s(Q)} \right) \right]. \quad (5.90)$$

Then, for the background density estimation to be non-zero, we need at least half of the cells non-empty. If there is a number $N = A_{\text{tot}}/A_{\text{cell}}$ cells, this probability is given by (recalling that we already have n_b initial hard jets)

$$P_{\text{nonzero}} = \sum_{k=N/2-n_b}^{N-n_b} \frac{(N-n_b)!}{k!(N-n_b-k)!} P_{\text{em}}^k \approx \frac{(N-n_b)!}{(N/2-n_b)!(N/2)!} P_{\text{em}}^{N/2-n_b}, \quad (5.91)$$

where we have assumed that in the limit $P_{\text{em}} \ll 1$, the sum is dominated by the first term. Given P_{nonzero} , one can estimate ρ_{est} by observing that the $(N/2)^{\text{th}}$ jet will be the softest of all and therefore have $p_t \approx Q_0$, hence

$$\rho_{\text{est}}^{(\text{pert})} \approx \frac{Q_0}{A_{\text{cell}}} P_{\text{nonzero}}. \quad (5.92)$$

5.2.4 Perturbative radiation and soft pileup

In practice, real events will have both soft radiation coming from pileup and perturbative radiation. The combination of the two effects is more delicate to handle analytically, unless we work in the dense pileup region, $\nu A_{\text{cell}} \gg 1$, far from the turn-on point $\nu A_{\text{cell}} = \log(2)$. In this region, relevant for most phenomenological considerations, the pileup distribution in the grid cells can be considered approximately Gaussian.

Out of the $N = A_{\text{tot}}/A_{\text{cell}}$ cells, an average of $n_b + NP_{\text{em}}$, with P_{em} given by Eq. (5.90), will be contaminated by hard perturbative radiation. The $N(1 - P_{\text{em}}) - n_b \approx N - \langle n_h \rangle + \langle n_h \rangle^2 / (2N)$ remaining jets will have an approximately Gaussian $\rho_{\text{cell}} = p_{t,\text{cell}}/A_{\text{cell}}$ distribution of the form

$$\frac{dn^{\text{PU}}}{d\rho_{\text{cell}}} \approx \frac{1}{\sigma} \sqrt{\frac{A_{\text{cell}}}{2\pi}} \exp \left[-\frac{A_{\text{cell}}}{2\sigma^2} (\rho_{\text{cell}} - \rho)^2 \right] (N(1 - P_{\text{em}}) - n_b). \quad (5.93)$$

Assuming $n_b/N + P_{\text{em}} < 1/2$, the median procedure will find ρ_{est} such that $N/2$ of the $N(1 - P_{\text{em}}) - n_b$ remaining Gaussian-distributed pileup jets have $\rho_{\text{jet}} < \rho_{\text{est}}$. This means that ρ_{est} will be determined by

$$\int_{-\infty}^{\rho_{\text{est}}} d\rho_{\text{jet}} \frac{dn^{\text{pileup}}}{d\rho_{\text{jet}}} = \frac{N}{2}. \quad (5.94)$$

In the small $\langle n_h \rangle$ limit, this can be solved analytically and we find

$$\rho_{\text{est}} \approx \rho_{\text{est}}^{(\text{PU})} + \sqrt{\frac{\pi}{2A_{\text{cell}}}} \sigma \left(\frac{\langle n_h \rangle}{N} + \frac{\langle n_h \rangle^2}{2N^2} \right), \quad (5.95)$$

$$\approx \rho - \frac{\sigma^2}{4\rho A_{\text{cell}}} + \sqrt{\frac{\pi A_{\text{cell}}}{2}} \sigma \left(\frac{\langle n_h \rangle}{A_{\text{tot}}} + A_{\text{cell}} \frac{\langle n_h \rangle^2}{2A_{\text{tot}}^2} \right), \quad (5.96)$$

where $\rho_{\text{est}}^{(\text{PU})}$ in the first equality is the result obtained in Eq. (5.83) in the pure-pileup case.¹⁹ For large-enough pileup, the scale Q_0 in the definition of $\langle n_h \rangle$ should be replaced by a scale of the order of the pileup fluctuations $\sigma\sqrt{A_{\text{cell}}}$, so that we have

$$\frac{\langle n_h \rangle}{A_{\text{tot}}} \approx \frac{n_b}{A_{\text{tot}}} + \frac{C_i}{2\pi^2\beta_0} \log \left(\frac{\alpha_s(\max(Q_0, \sigma\sqrt{A_{\text{cell}}}))}{\alpha_s(Q)} \right). \quad (5.97)$$

At the end of the day, the bias in the estimation of the pileup density ρ receives two main contributions: the non-Gaussianity of the pileup distributions results in a negative bias decreasing with the grid size a , the second term in the above equation, and the contamination from the hard radiation yields a positive bias increasing with a . We should therefore expect that there is an intermediate range in a where these two biases cancel approximately.

If we ignore the n_b contribution to $\langle n_h \rangle$ and the term proportional to $\langle n_h \rangle^2$, and set the logarithm in $\langle n_h \rangle$ to a typical value of 1, we find that the two biases cancel for

$$a \approx \sqrt{\frac{\pi}{2^{1/3}}} \left(\frac{\beta_0 \sigma}{C_i \rho} \right)^{1/3} \approx 0.9 \left(\frac{C_A \sigma}{C_i \rho} \right)^{1/3}. \quad (5.98)$$

5.2.5 Fluctuations in the estimated ρ

The estimation of ρ will fluctuate depending on the exact distribution of the pileup particles and the hard radiation. This is due to the fact that the pileup distribution is only sampled in a finite area and with a finite number of objects. The pileup fluctuations therefore also imply some fluctuations in the estimated value of ρ . In that sense, the results provided in the previous Section are really averaged values for the estimated ρ . In this Section, we concentrate on the fluctuations of the estimated ρ around this averaged value. As for the average estimate of ρ above, we shall first consider the contribution from the soft pileup itself and then the contribution from the hard radiation.

Pure-pileup case. For our purpose, it is sufficient to work in the dense limit where $\nu A_{\text{cell}} \gg 1$ so that we can approximate the $\rho_{\text{cell}} = p_{t,\text{cell}}/A_{\text{cell}}$ distribution (see Eq. (5.77)) by a Gaussian of average ρ and dispersion $\sigma/\sqrt{A_{\text{cell}}}$

$$\frac{dP}{d\rho_{\text{cell}}} = \sqrt{\frac{A_{\text{cell}}}{2\pi}} \frac{1}{\sigma} \exp \left[-\frac{A_{\text{cell}}}{2\sigma^2} (\rho_{\text{cell}} - \rho)^2 \right]. \quad (5.99)$$

From that we can easily obtain the cumulative probability

$$P(< \rho_{\text{cell}}) = \int_{-\infty}^{\rho_{\text{cell}}} d\rho_{\text{cell}} \frac{dP}{d\rho_{\text{cell}}} = \frac{1}{2} + \frac{1}{2} \text{erf} \left(\sqrt{\frac{A_{\text{cell}}}{2}} \frac{\rho_{\text{cell}} - \rho}{\sigma} \right) \approx \frac{1}{2} + \sqrt{\frac{A_{\text{cell}}}{2\pi}} \frac{\rho_{\text{cell}} - \rho}{\sigma}, \quad (5.100)$$

¹⁹In the second line, we have dropped the Θ function in $\rho_{\text{est}}^{(\text{PU})}$. In practical applications in Section 5.2.6, we will simply require that the final estimation be positive.

where the last expression is valid when $\sqrt{A_{\text{cell}}}(\rho_{\text{cell}} - \rho)/\sigma$ is small.

Now, say we have $N = 2m + 1$ such cells, where we have taken N odd for simplicity. In that case, the median will be ρ_{est} if m cells have a smaller ρ_{cell} , one cell has $\rho_{\text{cell}} = \rho_{\text{est}}$ and the remaining m cells a larger ρ_{cell} . This corresponds to the following probability distribution

$$\frac{dP_{\text{median}}}{d\rho_{\text{est}}} = (2m + 1) \frac{(2m)!}{(m!)^2} \frac{dP}{d\rho_{\text{est}}} [P(< \rho_{\text{est}})]^m [1 - P(< \rho_{\text{est}})]^m. \quad (5.101)$$

For large m , we can use the Stirling formula, $m! \approx \sqrt{2\pi m}(m/e)^m$, so that, working in the small $\sqrt{A_{\text{cell}}}(\rho_{\text{cell}} - \rho)/\sigma$ limit, we get

$$\frac{dP_{\text{median}}}{d\rho_{\text{est}}} \approx \frac{\sqrt{2mA_{\text{cell}}}}{\pi\sigma} \left[1 - \frac{2A_{\text{cell}}(\rho_{\text{cell}} - \rho)^2}{\pi\sigma^2} \right]^m \approx \frac{\sqrt{A_{\text{tot}}}}{\pi\sigma} \exp \left[-\frac{A_{\text{tot}}(\rho_{\text{cell}} - \rho)^2}{\pi\sigma^2} \right], \quad (5.102)$$

where, to obtain the last equality, we have used $N \approx 2m$, $NA_{\text{cell}} = A_{\text{tot}}$ and taken the large N limit. This means that the dispersion in the estimated ρ will be

$$S_{\text{est}}^{(\text{PU})} \approx \sqrt{\frac{\pi}{2A_{\text{tot}}}} \sigma. \quad (5.103)$$

This result has actually already been used in the previous Chapter when estimating the minimal size of a range required to obtain a good estimate of ρ (see Appendix 4.4.4).

Note that if we had computed ρ from the average of the ρ_{cell} , we would have obtained a dispersion $\sigma/\sqrt{NA_{\text{cell}}} = \sigma/\sqrt{A_{\text{tot}}}$. The median estimate is larger by a factor $\sqrt{\pi/2} \approx 1.25$. This moderate increase in resolution is the price to pay to have a method which is robust against hard perturbative radiation.

Hard perturbative radiation. To compute this effect, we take a slightly different approach than the one we have adopted before to calculate the average effect.

We start with an event with pileup and investigate the effect of adding hard perturbative radiation. For each emission, we then have two possible situations: either the hard emission contributes to a cell which was already above the median, leaving the median unchanged, or it contributes to a cell which was below the median, putting it above. If there are k emissions belonging to the second category, the median will be shifted by (still working in the large number of cell N limit)

$$\langle \delta\rho_{\text{est}} \rangle = k\delta_1, \quad \text{with } \delta_1 = \sqrt{\frac{2\pi}{A_{\text{cell}}}} \frac{\sigma}{N}. \quad (5.104)$$

For $k = \langle n_h \rangle / 2$, we correctly recover the hard contamination from (5.96).

Around that average value, we can have two sources of fluctuations: the first is related to the fact that, for a given k , the exact shift of the median will depend on how the ρ_{cell} are distributed, and the second comes from the fluctuations in k itself.

Let us start with the first source. For $k = 1$, the median will shift up by one jet, and the distribution of $\delta\rho$ will be simply be given by the distribution of the difference in ρ cell between two neighbouring cells in the sorted sequence of cells (at position in the sequence that is near the median). That distribution is an exponential distribution with mean δ_1 :

$$\frac{dP_{\text{shift}}}{d\delta\rho}(k = 1) = \frac{1}{\delta_1} e^{-\delta\rho/\delta_1}. \quad (5.105)$$

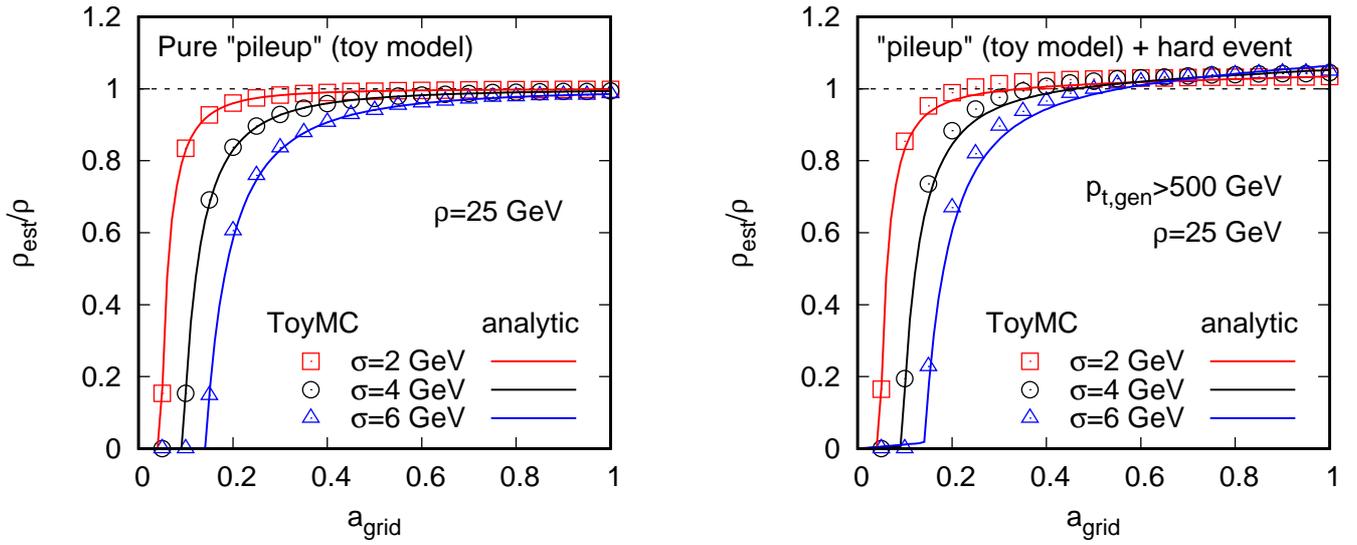


Figure 5.15: Analytic predictions obtained in this Section for the estimated value of ρ compared to numerical simulations of our toy-model pileup. For the hard events, we have generated dijet events with a generator \hat{p}_t of 500 GeV.

For an arbitrary k shifted jets, this therefore gives

$$\frac{dP_{\text{shift}}}{d\delta\rho}(k) = \frac{1}{k!} \frac{(\delta\rho)^{k-1}}{\delta_1^k} e^{-\delta\rho/\delta_1}, \quad (5.106)$$

which has a mean $k\delta_1$ as expected and standard deviation $\sqrt{k}\delta_1 \simeq \sqrt{\langle n_h \rangle / 2} \delta_1$.

Now, for the second effect, k will follow a Poisson distribution with mean $\langle n_h \rangle / 2$ and standard deviation $\sqrt{\langle n_h \rangle / 2}$.²⁰ This corresponds to a fluctuation $\sqrt{\langle n_h \rangle / 2} \delta_1$ in the estimated ρ .

Putting these two results together, we obtain

$$S_{\text{est}}^{(\text{hard})} = \sqrt{\langle n_h \rangle} \delta_1 = \sqrt{2\pi \frac{A_{\text{cell}} \langle n_h \rangle}{A_{\text{tot}} A_{\text{tot}}}} \sigma \quad (5.107)$$

with $\langle n_h \rangle / A_{\text{tot}}$ given by Eq. (5.97).

5.2.6 Comparison to Monte-Carlo studies

In order to cross-check our calculations, we implemented the toy pileup model in a Monte-Carlo event generator and compared results of numerical simulations to our analytic expectations.

Figure 5.15 shows the comparison for pure-pileup events on the left, and, on the right, for pure-pileup events in which we have embedded dijet events simulated with Pythia8 with $p_{t,\text{gen}} \geq 500$ GeV (right). In all cases, we have kept the particles up to $|y| = 5$ and varied the size of the grid used to estimate ρ between 0.05 and 1 by steps of 0.05.²¹ We have fixed $\rho = 25$ GeV, a value in the ballpark

²⁰For simplicity, we have treated Born partons and other emissions on equal footing. A more precise calculation would only include the hard perturbative emissions in the Poisson distribution.

²¹Note that we have used the `GridMedianBackgroundEstimator` class provided with `FastJet`. This will adapt the grid-size in both rapidity and azimuthal angle to fit an integer number of cells in the available interval. This should have a negligible influence on our results.

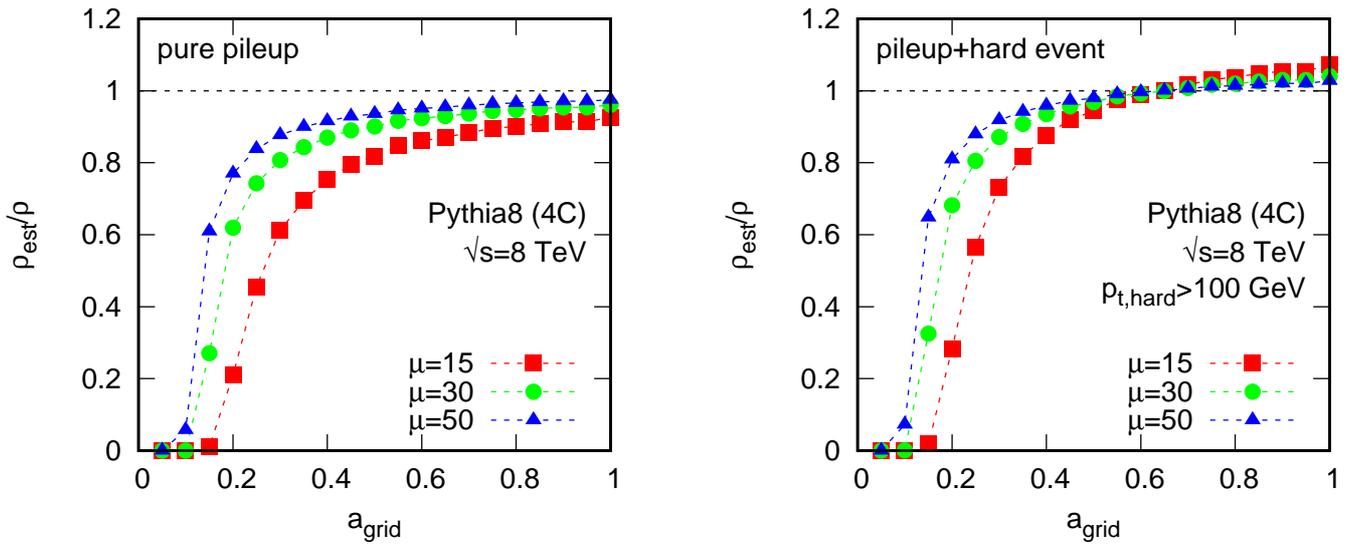


Figure 5.16: Similar results as in Fig. 5.15 for pure pileup and pileup with an additional hard dijet, now with pileup also obtained from Pythia8 simulations. For the hard events, we have generated dijet events with a generator \hat{p}_t of 100 GeV.

of what one could expect at the LHC and varied σ to probe the scaling of our analytic results. For the analytic curves, we have used $n_b = 2$, $C_i = C_A$ and a one-loop running coupling with $n_f = 5$ and $\Lambda_{\text{QCD}} = 100$ MeV.²² We see that, globally speaking, the analytic expressions do provide a very good qualitative and quantitative description of the features observed in the Monte Carlo simulations.

Finally, we also show on Fig. 5.16 simulations obtained with Pythia8 for both the hard events and the pileup. Although the quantitative details are different, especially for the transition from 0 to the region with $\rho_{\text{est}} \sim 1$. We know however that this region is sensitive to the details of the pileup model and we should therefore expect that our simplified toy model should at most capture the qualitative features of the full Pythia simulations. The plots on Figs. 5.15 and 5.16 show that it is indeed the case.

²²We have checked that varying Λ_{QCD} does not lead to any visible changes in our results.

Part II

Grooming techniques for fat jets

Chapter 6

Main methods

For the second part of this document, we shall switch gear and consider another approach that can be used to mitigate pileup, namely *grooming techniques*. In this Chapter, we discuss what is meant by grooming and describe the most common techniques that have recently been introduced and used at the LHC. As for the area–median subtraction method discussed in the previous part, we will then carry on with some Monte-Carlo studies, presented in Chapter 7, and finally show that groomers are amenable to some level of analytic understanding in Chapter 8.

Today, grooming techniques are mostly used for specific applications, usually in the context of boosted jets and jet substructure. Compared to the area–median techniques, groomers are not as widely used as generic pileup subtraction tools at the LHC. We shall therefore remain a bit more concise in our presentation than what we have done in the previous part of this document. However, I believe that grooming techniques have the potential to play a crucial role in the development of more advanced pileup mitigation techniques. They are for example an element of the cleansing pileup mitigation technique [10] recently introduced. In the third part of this document, which discusses new pileup mitigation techniques, we shall therefore take an excursion into currently unpublished, preliminary, results trying to use grooming techniques as a generic pileup subtraction tool (see Chapter 10).

6.1 Grooming and tagging fat jets

Let us put aside for a moment the “standard” use of jets at the LHC with relatively small jet radii and work with jets defined with larger R , in the range 0.8-1.5. These *fat jets* are severely affected by pileup and even without any pileup, are still substantially affected by the Underlying Event. One therefore often tries to exclude from the jet a fraction of its constituents which is deemed to contribute essentially to pileup or the UE. For example, one tends to remove soft contamination away from the jet centre, a region indeed expected to be dominated by UE or pileup. This is typically what *groomers* do: given an input jet, they clean it from soft and large angle radiation and return a groomed jet.

To give a slightly more concrete perspective to this discussion, let us illustrate our arguments with the specific example of *filtering*. It works as follows: given a jet, one re-clusters its constituents into subjets (with a smaller radius) and keeps only a (pre-decided) fixed number of the hardest subjets (i.e. with the largest p_t). The union of the subjets that have been kept is the filtered jet. This filtered jet is therefore made of a subset of the original jet’s constituents. Over a bit less than a decade, many grooming techniques have been introduced and we shall briefly review them in the next Section.

We should notice that there is an overlap between the grooming techniques and *boosted objects tagging*. In a nutshell, the latter refers to the fact that massive objects ($W/Z/H$ bosons or top

quarks), when produced with a p_t much larger than their mass and when decaying into hadrons, will be reconstructed as single jets. To separate these boosted fat jets, originating from the hadronic decay of a massive object, from the large QCD background, one also exploits the internal properties of jets, i.e. *jet substructure*. Such techniques are often referred to as *taggers*.

Taggers and groomers share many similarities, going beyond the simple observation that they are both used on fat jets. Physically, one approach to tag boosted jets is to rely on the fact that massive objects share their p_t across their decay products more symmetrically than QCD jets which tend to be dominated by soft radiation. In essence, this approach of finding hard prongs inside a jet bears resemblance to eliminating soft contamination from the Underlying Event or pileup. Hence taggers bear resemblance with groomers. In practice, many tools introduced as groomers can be used as taggers and vice versa. Furthermore, combinations of the two are not uncommon, *c.f.* the use of filtering after applying a mass-drop tagger in [59].

In my opinion (see also the discussion in [107]), the main difference between a tagger and a groomer lies in what we expect it to do: when we use a tagger, we understand that we want to identify multiple hard prongs in a fat jet with the ultimate goal to maximise signal-to-background efficiency in boosted objects searches; when we use a groomer, we mean that we want to eliminate soft contamination and improve the resolution on the jet kinematics and properties. In other words, a tagger either succeeds or fails, typically according to whether or not it finds hard substructure in the jet, while a groomer always returns a (groomed) jet.

An additional point which is worth keeping in mind in the context of pileup mitigation is that taggers and groomers can be seen as acting on different physical scales (although, in a practical context, the separation may not be as clear): improving the resolution on a jet, the work of groomers, is expected to be done by imposing cuts around the soft UE/pileup scale, while, for taggers, identifying hard prongs in a jet is likely to involve cuts at the perturbative scale(s) of the jet. In the end, this is an area of jet substructure which is still in full development and we can hope for exciting improvements in the next years.

Let us now come back to our original purpose of mitigating pileup. One should understand that grooming technique adopt a different approach to pileup mitigation than the area–median subtraction method used so far. The main conceptual difference is that “standard” pileup subtraction techniques like the area–median method try to correct the jet contaminated by soft radiation back to a jet as it was without that contamination, while grooming techniques would most often affect the hard jets event in the absence of pileup.

This can be understood in the context of our qualitative discussion of pileup effects in Section 2.1. The main goal of the area–median subtraction method is to eliminate the shift due to pileup independently of the hard process one is considering. This means, eliminate, on average, the ρA_{jet} contamination from the jet. In doing so, the smearing effects are reduced to $\sigma\sqrt{A_{\text{jet}}}$. Grooming techniques, on the other hand, target an improvement in p_t (or mass) resolution of the jet, i.e. want to reduce as much as possible the $\sigma\sqrt{A_{\text{jet}}}$ without guaranteeing an unbiased average contamination. By excluding soft parts of the jet, grooming typically reduces its area from A_{jet} to, say, $A_{\text{groomed}} < A_{\text{jet}}$. This is not bias-free since it comes at the expense of removing some of the constituents of the original hard jet (a negative bias) and leaving a positive ρA_{groomed} contribution from pileup. However, it significantly lowers the resolution degradation by reducing it to $\sigma\sqrt{A_{\text{groomed}}}$. Adjusting the free parameters of the groomer we could try to minimise the bias while keeping the resolution improvements but this has no guarantee to work on a wide range of applications and contexts.

Both approaches have their advantage. For a broad and robust applications of pileup mitigation to jets with not-so-large R , it is crucial to control the average bias and this is what the area–median

subtraction method is good at. On the other hand, if one considers a search based on fat jets, appropriate grooming techniques are perfectly suited to obtain a narrow signal peak on a reduced QCD background. Note that, in this case, we also wish to supplement the groomer with area–median subtraction in order to also correct for the positive bias due to ρA_{groomed} and further improve the resolution when averaging over the full jet sample.

Before digging into the grooming techniques themselves, it is interesting to illustrate some of the arguments above with a concrete example. Let us consider the measurement of the mass of boosted tops reconstructed using fat jets. The area–median subtraction alone would, on average reconstruct the correct average jet mass for both top quarks and pure-QCD jets, with a potentially large smearing due to pileup fluctuations (and the UE). A grooming technique, like filtering briefly introduced above where we would keep the 3 or 4 hardest subjets, would have a different behaviour. For signal jets coming from the hadronic decay of top quarks, one would keep the main structure of the 3-prong decays and hence reconstruct a mass with little average bias, and, at the same time, benefit from the grooming to obtain a good resolution on the mass i.e. a narrower mass peak than without filtering. For standard QCD jets, where most often the mass comes from soft radiation at larger angles, grooming would strongly reduce their mass, already at the perturbative level. This can be a large bias which, in the case of our search for top quarks, plays in our favour since it further suppresses the QCD jet background. Combining grooming and area–median subtraction would further improve the top signal and further reduce the QCD background. We shall see this explicitly in Chapter 7.

6.2 Mass-drop, filtering, trimming, pruning, soft drop

Here, we briefly introduce the main grooming techniques that have been used over the last decade. Some were initially introduced either as taggers, or as both taggers and groomers, but at the end of the day, they all eliminate soft and large-angle radiation and, in that sense, can be seen as groomers.

Filtering. Given a fat jet, filtering [59] re-clusters its constituents with the Cambridge/Aachen algorithm with a small radius R_{filt} . It then only keeps the n_{filt} hardest subjets (in terms of their p_t). The filtered jet is the union of the subjets that have been kept. This has two adjustable parameters: R_{filt} and n_{filt} . It is typically used to reduce soft contamination in situations where we have a prior knowledge of the number of hard prongs in a jet. For a jet with n_{prong} hard prongs — $n_{\text{prong}} = 2$ for a $W/Z/H$ bosons and $n_{\text{prong}} = 3$ for a top — we would typically use $n_{\text{filt}} = n_{\text{prong}} + 1$ which would also keep the (perturbative) radiation of an extra gluon. Filtering is implemented directly in **FastJet** where it is available as the `Filter` tool (using `SelectorNHardest` for the selection criterion).

Trimming. Trimming [108] shares some similarities with filtering. It also starts with re-clustering the jet with a smaller radius, R_{trim} , using either the k_t or the Cambridge/Aachen algorithm. It then keeps all the subjets with a transverse momentum larger than a fraction f_{trim} of the initial jet transverse momentum. On top of the choice of algorithm, this also has two parameters: R_{trim} and f_{trim} . It is often used as a generic groomer and tagger in boosted-jet studies. Trimming can be obtained from the `Filter` class in **FastJet** (using `SelectorPtFractionMin` for the selection criterion).

Pruning. Pruning [109] is similar in spirit to trimming except that it adopts a bottom-up approach (with trimming seen as a top-down approach). Given a jet, pruning reclusters its constituents using a user-specified jet definition (based on pairwise recombinations) with the additional constraint that at

each step of the clustering, objects i and j are only recombined if they satisfy at least one of these two criteria: (i) the geometric distance $\Delta R_{ij} = \sqrt{\Delta y_{ij}^2 + \Delta \phi_{ij}^2}$ is smaller than $R_{\text{prune}} = f_{\text{prune}} \times 2m_{\text{jet}}/p_{t,\text{jet}}$, with $p_{t,\text{jet}}$ and m_{jet} the original jet transverse momentum and jet mass, (ii) the splitting between i and j is sufficiently symmetric, i.e. $\min(p_{t,i}, p_{t,j}) \geq z_{\text{prune}} p_{t,(i+j)}$. If none of these criteria are met, only the hardest of i and j (in terms of their p_t) is kept and the other is rejected. On top of the jet definition used for the re-clustering, which is usually taken to be either k_t or Cambridge/Aachen with a radius much larger than the one of original jet, this has two parameters: f_{prune} and z_{prune} . z_{prune} plays the same role as f_{trim} for trimming and f_{prune} plays a role similar to R_{trim} . Note however that the R_{prune} is defined dynamically with the jet kinematics, compared to R_{trim} which is kept fixed. This can have important consequences both analytically and phenomenologically.¹ Pruning can also be considered as a general-purpose groomer and tagger and is often used in situations similar to trimming. It is available in `FastJet` as the `Pruner` class.

I and Y-Pruning. When pruning a jet, there might be situations where a soft emission at large angle dominates the mass of the jet but gets pruned away because it does not satisfy the symmetry condition. The mass of the pruned jet is then determined by radiation at smaller angle, typically within the pruning radius. This situation where the jet mass and the pruning radius are determined by different emissions within the jet would result in a jet with a single prong and are therefore called I-Pruning. For these situations, the pruning radius has no relation to the hard substructure of the jet. More precisely, I-Pruning is defined [107] as the subclass of pruned jets for which, during the sequential clustering, there was never a recombination with $\Delta R_{ij} > R_{\text{prune}}$ and $\min(p_{t,i}, p_{t,j}) > z_{\text{prune}} p_{t,(i+j)}$. The other situation, i.e. a pruned jet for which there was at least one recombination for which $\Delta R_{ij} > R_{\text{prune}}$ and $\min(p_{t,i}, p_{t,j}) > z_{\text{prune}} p_{t,(i+j)}$, corresponds to a genuine two-prong structure and is called Y-Pruning. This distinction between I- and Y-Pruning is mostly relevant for boosted jet tagging. However, it has been shown to have an impact on the analytical behaviour of Pruning, with Y-Pruning being under better control and than I-Pruning, the latter adding quite a lot of complexity to the calculation. If one wishes to reach some level of analytic control over groomed jets, Y-Pruning appears as a more natural choice than Pruning which also includes the I-Pruning class.

Mass-drop. Originally proposed [59] as a tool to isolate boosted Higgs bosons, decaying to $b\bar{b}$ pairs, from the QCD background, the mass-drop tagger assumes one works with a fat jet clustered with the Cambridge/Aachen algorithm — if it is not the case, one is always at liberty to re-cluster the jet. One then iteratively undoes the last step of the clustering $p_{i+j} \rightarrow p_i + p_j$ until both the following criteria are met: (i) there is a “mass drop” i.e. $\max(m_i, m_j) < \mu_{\text{cut}} m_{i+j}$, (ii) the splitting is sufficiently symmetric i.e. $\min(p_{t,i}^2, p_{t,j}^2) \Delta R_{ij}^2 > y_{\text{cut}} m_j^2$. When both criteria are met, we keep “ $i + j$ ” as the result of the mass-drop tagger, otherwise the procedure is repeated iteratively keeping only the most massive of i and j . If the procedure fails to find two subjets satisfying the conditions, i.e. end up recursing until it reaches a single constituent which can not be further de-clustered, it is considered as having failed and returns an empty jet. The mass-drop tagger has two parameters: μ_{cut} , the mass-drop parameter itself, and y_{cut} , the symmetry cut. The two conditions imposed by the mass-drop tagger make full sense in the physical situation where we want to tag the decay of 2-pronged boosted objects: the symmetry cut requires that one indeed finds two hard prongs and the mass-drop condition imposes that one goes from a massive boson jet to two jets originated from massless QCD partons. Although it was originally introduced as a tagger, the mass-drop tagger also acts as a groomer since, via the

¹It would be interesting to study the properties of trimming with a dynamical trimming radius.

Cambridge/Aachen declustering, it would iteratively remove the soft radiation at the outskirts of the jet, hence limiting the pileup/UE contamination. This is available as the `MassDropTagger` tool in `FastJet`.

Modified Mass-drop (mMDT). It has been shown recently [107] that substructure tools can be understood analytically. When trying to understand the behaviour of the mass-drop tagger on QCD jets, it was realised that the fact that the iterative de-clustering was following the most massive of the two branches could lead to pathological situations. It was therefore suggested to adapt the procedure so that it instead follows the hardest branch (in terms of p_t). This modification has the advantage that it makes the analytical calculation much easier and more robust without affecting the performance of the method (even improving it slightly). The same study also added two more minor modifications. First, it was realised that the symmetry condition could be replaced by $\min(p_{t,i}, p_{t,j}) > z_{\text{cut}}(p_{t,i} + p_{t,j})$ and that the latter would have a slightly reduced sensitivity to non-perturbative corrections. Second, perturbatively, the mass-drop condition would only enter as a subleading correction (in the strong coupling constant α_s) compared to the symmetry condition and so could usually be ignored.² The `ModifiedMassDropTagger` is available as part of the `RecursiveTools` contribution in `fjcontrib`. By default, it uses a z_{cut} symmetry condition, does not impose a mass-drop condition and follows the highest- p_t branch, but variants are allowed in the code.

Soft Drop. Soft Drop [21] can be seen a generalisation of the modified Mass-drop tagger. It also proceeds by iteratively de-clustering a jet obtained — either directly or re-clustered — with the Cambridge/Aachen algorithm until the Soft Drop condition is met. If we decluster p_{i+j} into p_i and p_j , we stop the de-clustering when

$$\frac{\min(p_{t,i}, p_{t,j})}{p_{t,i} + p_{t,j}} > z_{\text{cut}} \left(\frac{\Delta R_{ij}}{R} \right)^\beta, \quad (6.1)$$

with R the jet radius. In the limit $\beta \rightarrow 0$, this condition reduces to the mMDT. Hence, Soft Drop can be seen as an extension of the mMDT with the extra parameter β allowing for more freedom in constraining the phase space.³ Soft Drop has two parameters. The z_{cut} parameter plays the same role as in the (m)MDT of keeping the hard structure and exclude soft emissions, starting from large angles due to the Cambridge/Aachen de-clustering. The β parameter controls how aggressive the groomer is. It goes from no grooming at all for $\beta \rightarrow \infty$ to a quite aggressive grooming for $\beta = 0$ (mMDT) and carries on with a 2-prong tagging region for $\beta \leq 0$. Used as a groomer, we believe that values of β in the range between 1 and 2 are appropriate as they retain some of the large-angle structure of the jet, although larger values could also be used.⁴ It is also interesting to realise that, for $\beta = 2$, the Soft

²This argument was made and checked in the context of boosted object tagging [107]. It would still be interesting to see if the μ_{cut} parameter in the mass-drop condition has an impact on grooming properties.

³The Soft Drop procedure returns by default a single particle if it fails to find two subjets satisfying the Soft Drop condition. This “grooming mode” is different from the default “tagging mode” of the mMDT which would fail, i.e. return an empty jet, if no substructure are found.

⁴Note that in the context of boosted object tagging, when used alone, negative values of β , typically the mMDT or Soft Drop with $\beta = -1$, give perfectly adequate and efficient taggers (see the original Soft Drop paper [21]). However, if one wishes to combine Soft Drop with other tagging techniques, like N -subjettiness or Energy-Correlation Functions, which constrain radiation patterns in boosted jets, we wish to retain some of the large-angle structure of the jets, where the radiation pattern between a colourless $W/Z/H$ and a QCD jet is most different. In that case, Soft Drop with positive β would be a good compromise between retaining some of this structure while excluding soft and non-perturbative UE and pileup effects. When combining two-prong taggers with radiation constraints, it would then appear natural to use

Drop condition imposes a momentum cut proportional to the square of the angular separation of the two subjects, i.e. proportional to the area of the subjects. One might therefore expect that the choice $\beta = 2$ plays a special role when it comes to discussing pileup effects. For example, for $\beta = 2$, one could expect the pileup contamination to correspond to a shift in z_{cut} . In practice Soft Drop is made available as a `SoftDrop` class in the `RecursiveTools` contribution to `fjcontrib`, where, similarly to the `mMDT`, it also offers the option to impose an additional mass-drop condition.

Combination with area–median subtraction. As discussed earlier, area–median subtraction and groomers target different aspects of pileup mitigation. In many cases it is however possible to combine both approaches in order to benefit both from the correction of the bias implemented in the area–median subtraction and the resolution improvements brought by grooming. We briefly comment on how this can be achieved for the various grooming procedures defined above.

- For the case of **filtering** and **trimming**, the idea is to use the area–median subtraction to correct the momentum of the subjects before selecting which ones have to be kept. In the end, the groomed jet would be the union of the *subtracted* subjects that have been kept. Note that in the case of trimming there is an ambiguity as to whether the condition on the (subtracted) subject p_t uses the subtracted full jet as a reference or the unsubtracted jet. Both options are acceptable and have their respective advantage. Basically, the former tends to be more aggressive than the latter. Depending on the application and the choice of parameters, both options show little practical difference. In terms of the implementation, the combination of filtering/trimming with area–median subtraction can be straightforwardly done by passing a `Subtractor` to the `Filter` object. For trimming, it is the end-user’s responsibility to decide whether to pass a subtracted or an unsubtracted jet.
- For **pruning**, we could also decide, at least on paper, to apply an area–median subtraction at each step of the clustering and impose the pruning criteria on the subtracted jets. This has not yet been implemented in practice. This is due to a purely technical reason: jet areas in `FastJet` are assigned to the jets a posteriori, hence are not available when one would need them to veto or not a recombination during the re-clustering.
- For the case of the **(modified) Mass-drop Tagger** and **Soft Drop**, it is again sufficient to apply the kinematic conditions on subtracted jets. The groomed jet at the end of the procedure would then be the union of the two subtracted subjects found by the iterative de-clustering. In practice, a `Subtractor` can be passed to both the `ModifiedMassDropTagger` and to the `SoftDrop` classes. Note that by default, the original jet is assumed to be unsubtracted and is subtracted internally.⁵ Note that the `MassDropTagger` class in `FastJet` has no support for internal subtraction so the `ModifiedMassDropTagger` should be preferred, knowing that the latter could, if needed, be parametrised so as to reproduce the behaviour of the former.

Further discussion, similarities and differences. Before carrying on with some Monte-Carlo studies of groomers and an example of their analytic properties, there are still a few generic remarks that we find of generic use.

in parallel negative, or zero, β to identify the 2-prong structure and positive β with a jet shape, to impose a cut on radiation.

⁵This default behaviour can be changed.

- all the groomers introduced above cut soft particles (or subjets) at large angles. This is either achieved explicitly by removing soft subjets (which are automatically away from the hard subjets by, at least, the subjet radius), or, when using de-clustering, relying on the specific choice of the Cambridge/Aachen jet algorithm, which reflects an angular ordering and therefore throws away soft radiation at large angles until some hard structure is met.
- Since they all implement mainly the same physical ideas, the physical behaviour of all these tools share many similarities, something that can be seen both in Monte Carlo studies and analytic calculations.
- At some point, when the jet boost is increased, important differences between the various groomers start to appear. For example, the fact that the Mass Drop tagger can recurse all the way to arbitrary small opening angle, compared to trimming which has a built-in subjet radius, or to Pruning which has a dynamical minimum angular scale, translates into different behaviour of these three tools. These can sometimes be seen in Monte-Carlo simulations but are best understood when assessing the analytic properties of the groomers.
- Soft Drop differs from the other groomers in that the symmetry cut is angle-dependent. In the end, I believe that this gives Soft Drop some extra flexibility which could prove useful in a broad range of applications. It makes Soft Drop a powerful groomer with nice analytic properties (see also next Chapters).
- In terms of the contamination from pileup (or the UE) at an angular scale ΔR , one would expect an average p_t shift of order $\rho \pi \Delta R^2$. To some extent, this is reminiscent of Soft Drop with $\beta = 2$ where the condition would read $z > z_{\text{cut}}(\Delta R/R)^2$, suggesting $z_{\text{cut}} \gtrsim A_{\text{jet}}\rho/p_{t,\text{jet}}$.⁶

⁶A similar argument can be made for pileup fluctuations and $\beta = 1$ Soft Drop.

Chapter 7

Monte-Carlo studies

So far, most of the practical applications of grooming techniques have concentrated on boosted jets studies. We shall therefore investigate here the performance of grooming as a tool to mitigate pileup effects on boosted jets. This is mostly based on a study that we initially conducted in the context of the proceedings of the Boost 2012 workshop [20], although the new version presented here has been updated to include a broader set of groomers, especially more modern ones. We will come back to grooming techniques as a way to mitigate pileup from generic jets in Chapter 10.

Compared to Ref. [20], this study includes more recent groomers like the modified Mass Drop Tagger and Soft Drop. Similar to the initial study, the events are obtained by combining the hard (signal) events embedded in a series of minimum-bias events to simulate pileup. We consider collisions at a centre-of-mass energy of 13 TeV. Minimum-bias events are generated with the Pythia8 Monte Carlo (MC) generator [35, 36], with its 4C tune [37]. We considered different pileup conditions, Poisson distributed with $\mu = \{30, 60, 100, 140\}$. This set of choices covers several situations of interest ranging from typical pileup conditions for Run I and II of the LHC¹ up to the exploration of future high intensity scenarii at the LHC.

We have selected two example signals for the Monte Carlo simulation studies presented in this Section. The first is the decay of a possible heavy Z' boson with a chosen $M_{Z'} = 1.5$ TeV to a (boosted) top quark pair, as included in the proceedings of the Boost 2012 workshop. The top and anti-top quarks then decay fully hadronically ($t \rightarrow Wb \rightarrow jjb$). The second also consists of a Z' boson with $M_{Z'} = 1.5$ TeV, now decaying into WW with both W decaying hadronically. These two samples are meant as sources of 3-pronged and 2-pronged jets, respectively. The signal samples were also generated using Pythia8 with tune 4C.

The analysis uses the tools either directly available in the `FastJet` package or in `fjcontrib` for the jet finding and the jet substructure analysis. The initial jet finding is done with the anti- k_t algorithm with $R = 1$, to ensure that most of the final state W and top-quark decays can be collected into one jet. For the latter, this corresponds to top-quarks generated with $p_t \gtrsim 400$ GeV.

In practice, we will apply several grooming procedures, listed in the next Section, to the original anti- k_t jets. For the studies presented in this report we require that the jet p_t before grooming and pileup subtraction is larger than 100 GeV. We keep jets with $|y| < 3$. We only consider the two hardest p_t jets in the event. We further require that the rapidity difference between the two jets $|y_1 - y_2|$ is less than one.

¹The change in centre-of-mass energy should not bring any sizeable difference compared to the present study.

7.1 Jet grooming configurations

Several jet grooming techniques, in use at the LHC, are included in our study (see Chapter 6 for a more complete description):

Trimming. In this study, we used the k_t algorithm for the reclustering, with $R_{\text{trim}} = 0.2$, and we used $f_{\text{trim}} = 0.03$.

Filtering. Here, we reclustered the jets with the Cambridge-Aachen algorithm with $R_{\text{filt}} = 0.2$, and kept the n_{filt} hardest subjects. n_{filt} has been varied according to the process under consideration and we found good results using $n_{\text{filt}} = 2$ for the $Z' \rightarrow WW$ process and $n_{\text{filt}} = 4$ for the $Z' \rightarrow t\bar{t}$ case.

modified Mass Drop (mMDT) and Soft Drop (SD). We have fixed $z_{\text{cut}} = 0.1$ and studied $\beta = 0$ (mMDT), 1 and 2 (SD). Following the original prescription of [59], we have optionally applied to the result of the mMDT/SD procedure, an extra step of filtering, using $R_{\text{filt}} = \max(R_{\text{SD}}, 0.2)$ with R_{SD} the distance between the two prongs found by the mMDT/SD, and n_{filt} as for the plain filtering case described above. In a nutshell, once the mMDT/SD procedure has found a 2-pronged structure, thereby excluding radiation at angles larger than R_{SD} , this extra filtering step further cleans the contamination at smaller angles.

A specific point we want to study is the interplay between jet grooming and area-based pileup correction. The subtraction is applied directly on the 4-momentum of the jet using the area–median approach, Eq. (2.30), with the active area of the jet computed using **FastJet**. We included the ρ_m term which corrects for contamination to the total jet mass due to the mass of the pileup particles. The estimation of ρ and ρ_m is performed with **FastJet** using a grid-based estimate with a grid of cell-size 0.55 extending to $|y| = 4$. Corrections for the rapidity dependence of the pileup density ρ are applied using a rapidity rescaling (see Sections 2.4.2 and 3.1). Ghosts are placed up to $y_{\text{max}} = 4$ and explicit ghosts are enabled, ensuring the availability of the area information to grooming tools. When applying this subtraction procedure, we discard jets with negative transverse momentum or (squared) mass of the jet.

When we apply this background subtraction together with grooming, the subtraction is performed directly on the subjects, before applying any kinematic constraint, i.e. before deciding which subjects should be kept, in the case of trimming and filtering, or before checking if the symmetry criterion is met in the case of mMDT and Soft Drop. This is meant to limit the potential effects of pileup on kinematic tests. Note that in the case of trimming, the unsubtracted jet p_t is used as the reference to compute the p_t threshold for subjects. It was found that this gives slightly better results than using the subtracted jet.

7.2 Jet substructure performance

As a first test, we look at the jet mass distribution obtained for our $Z' \rightarrow t\bar{t}$ event sample. This is plotted in Figs. 7.1 and 7.2. The immediate expectation for the reconstructed jet mass is the top mass, i.e. $m \approx 171$ GeV, and no residual dependence on the pileup activity μ , after the pileup subtraction. The two plots in the upper row of Figure 7.1 show the distributions of the reconstructed jet masses: without pileup subtraction on the left (7.1(a)) and with area–median pileup subtraction on the right (7.1(b)). Concentrating on the sample without any pileup added, we see that a large fraction of the events are reconstructed close to the top mass and about 30% are compatible with the W mass.

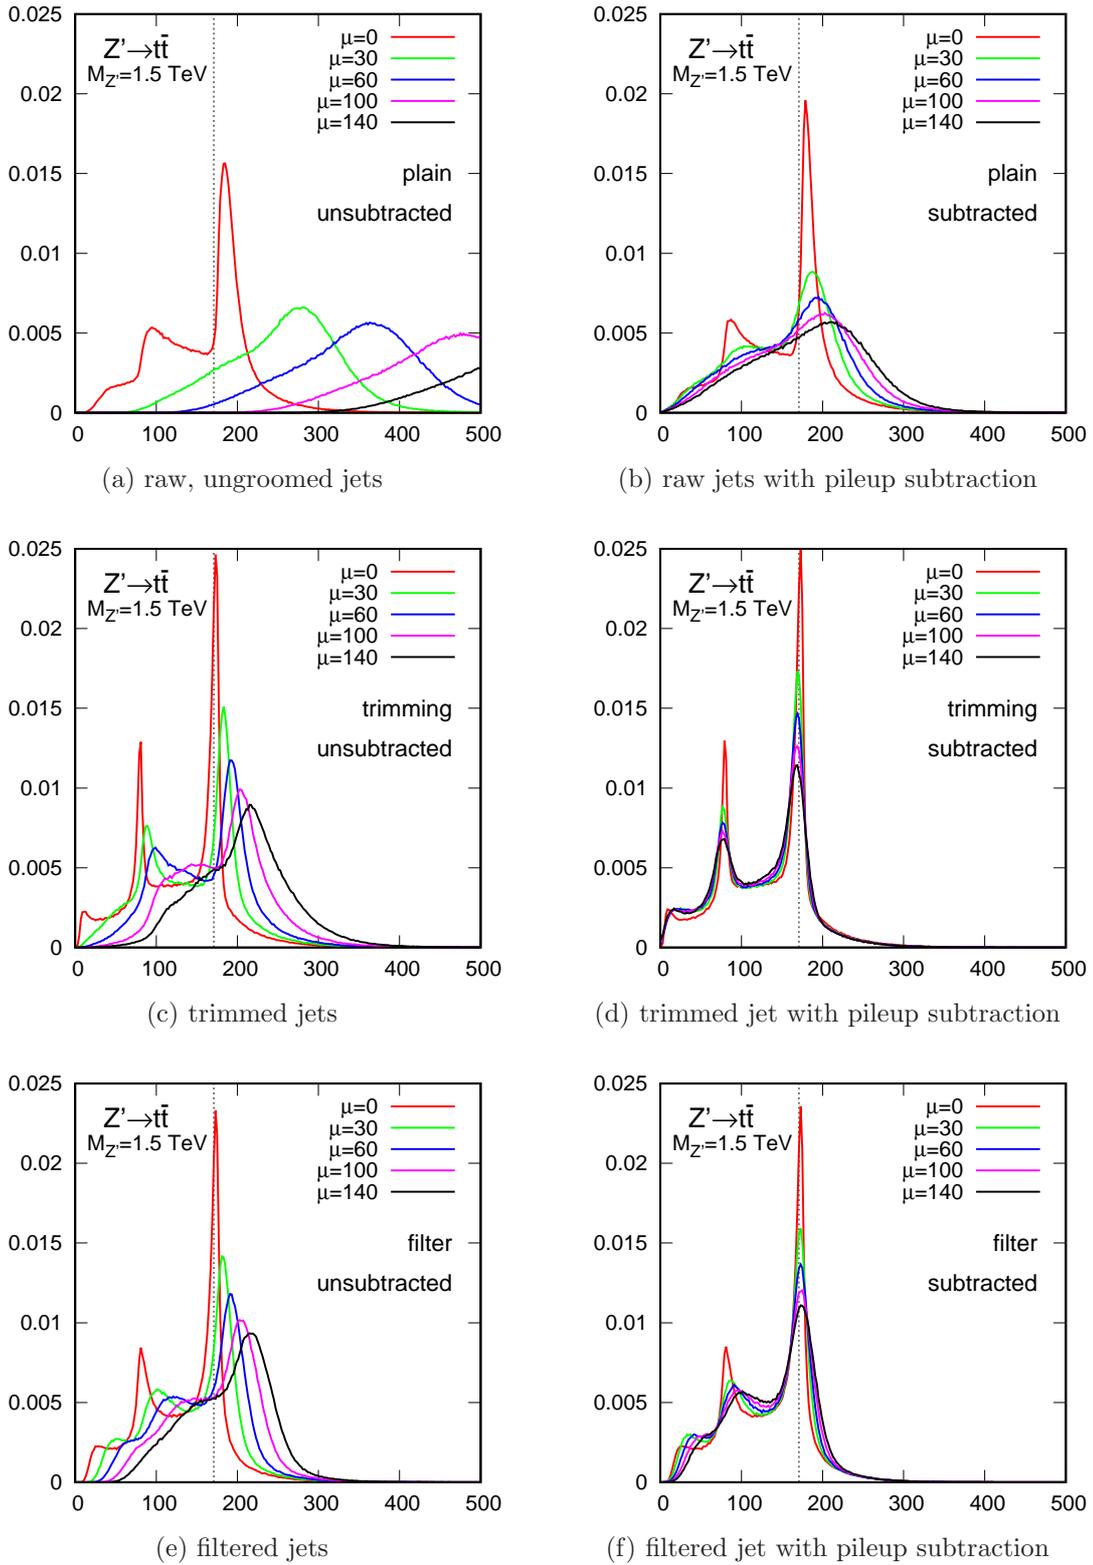


Figure 7.1: The impact of N_{PU} on the jet mass distribution. Top row: the raw jet mass distribution for $Z' \rightarrow t\bar{t}$ final states with $m_{Z'} = 1.5$ TeV, in the presence of pileup with $\mu = 30, 60, 100,$ and 140 , before (left) and after (right) area–median pileup subtraction. The second and third rows show the same results after trimming (middle row) and filtering (lower row). vertical size of the graphs

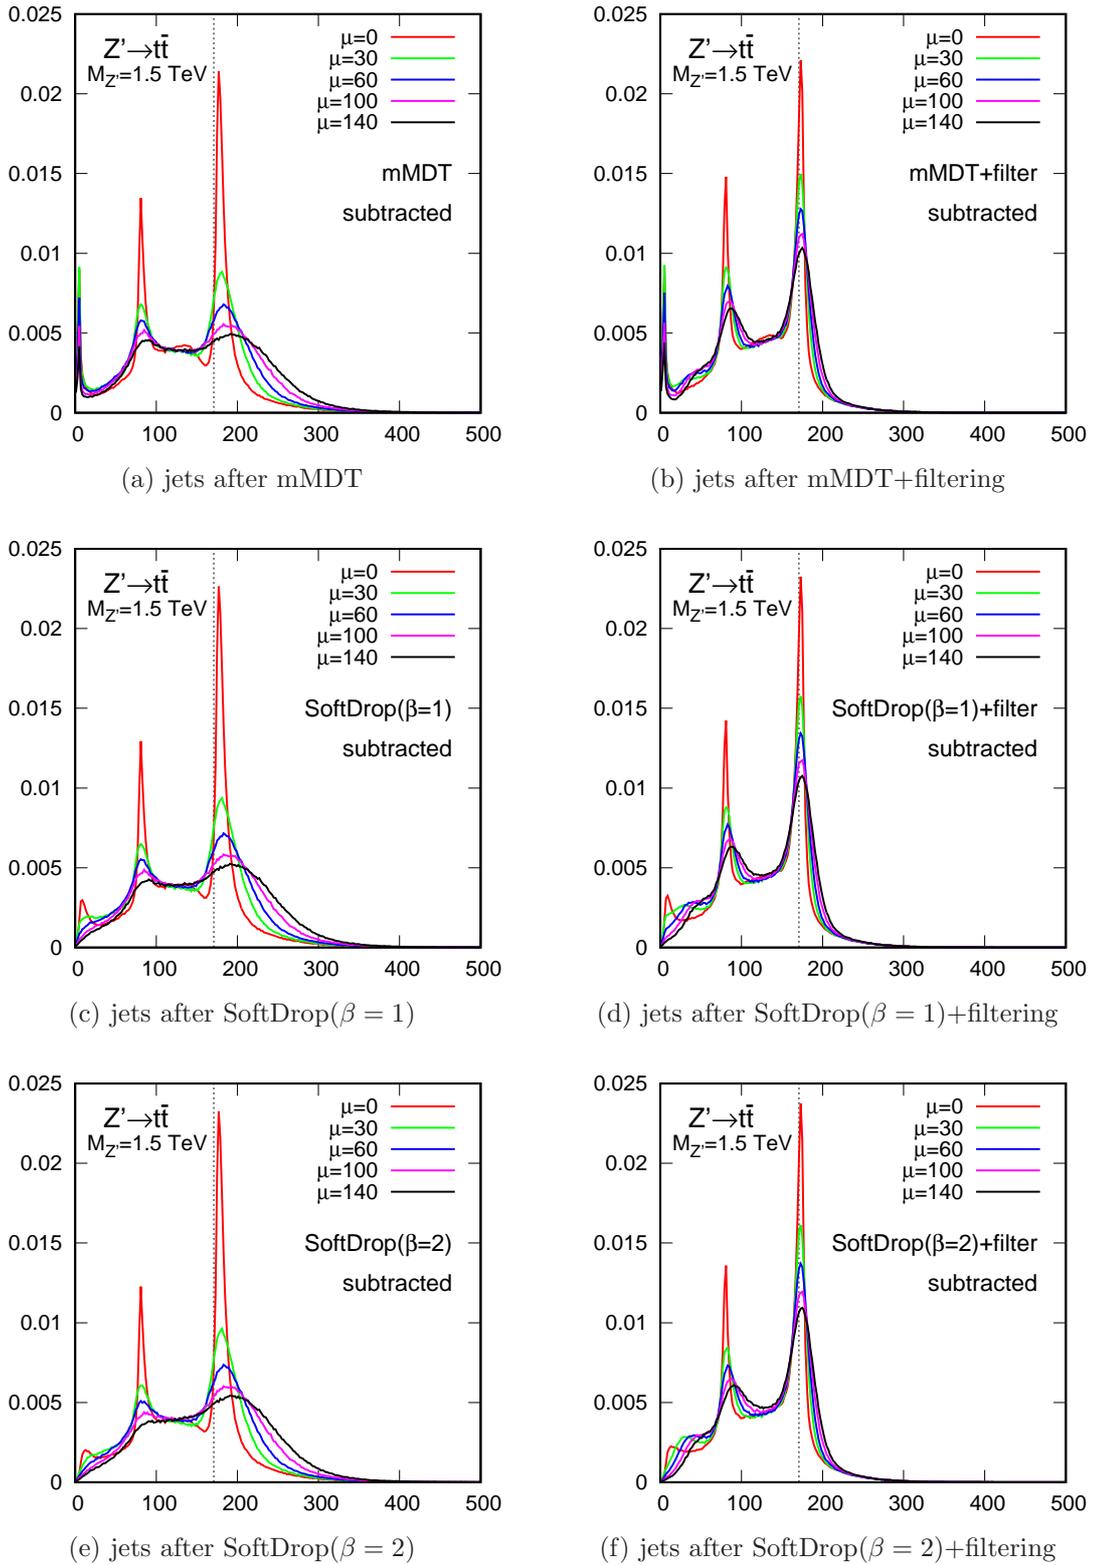


Figure 7.2: Same as Fig. 7.1 now showing the jet mass distribution after applying the modified Mass Drop (top row) or Soft Drop with $\beta = 1$ (middle row) or $\beta = 2$ (lower row). All distributions are shown after area–median pileup subtraction. Plots on the left column show the results of the initial grooming procedure alone, and plots on the right column show the results with an extra filtering step (see text for details).

The effect of pileup on the mass scale and resolution is clearly visible, increasingly shifting the peak towards much larger jet masses and smearing it.

Applying only the pileup subtraction, without applying any grooming to the jets, already improves the mass reconstruction significantly. The pileup shift is mostly removed from the jet mass spectrum, as shown in Fig. 7.1(b) and the position of the mass peak is recovered, independently of μ . With increasing pileup, the mass peak gets more and more smeared, an effect that we have already seen in the first part of this document, due to the fact that the pileup is not perfectly uniform. We have already stressed that these point-to-point fluctuations in an event lead to a smearing by $\pm\sigma\sqrt{A}$ in the subtracted jet p_t . For very large pileup, this smearing extends all the way to $m = 0$ as seen in Fig. 7.1(b) and we lose the clean separation between the W and the top mass peaks.²

The effect of the grooming techniques on the reconstructed jet mass distributions is summarised in the remaining plots of Figs. 7.1 and 7.2. For filtering and trimming, Figs. 7.1(c)-(f), we show the mass distributions with and without subtraction, while for the plot of Fig. 7.2 we show the result of the mMDT/SD procedure with and without the extra filtering step. The spectra show that grooming clearly improves the mass reconstruction. Already without applying area–median pileup subtraction, one sees that the peak is much less shifted and smeared than in the ungroomed case. The application of the pileup subtraction in addition to the grooming procedure further improves the mass reconstruction performance, fully correcting for the peak position and further improving the resolution. From Fig. 7.2, we also see that applying filtering on the result of the mMDT/SD grooming further improves the quality of the reconstructed spectrum.

The fact that the combination of grooming and area–median pileup subtraction appears as the optimal combination should not be a surprise. As discussed many times already, pileup has mainly two effects on the jet: a constant shift proportional to ρA and a smearing effect proportional to $\sigma\sqrt{A}$, with σ a measure of the fluctuations of the pileup within an event. In that language, the subtraction corrects for the shift leaving the smearing term untouched. Grooming, on the contrary, since it selects only part of the subjets, acts as if it was reducing the area of the jet³. This reduces both the shift and the dispersion. Combining grooming with subtraction thus allows to correct for the shift leftover by grooming and reduce the smearing effects at the same time. All these effects are observed on the plots.

The effect of the filtering step applied after the mMDT/SD grooming can also be understood from fairly simple arguments: the declustering procedure applied by mMDT/SD discards the soft activity at large angles in the jet until two hard prongs are found. This effectively reduces the jet area with the immediate consequence of reducing the shift and smearing of the mass peak. However, this results in a groomed jet with an effective radius dictated by the hard-prong structure of the jet. Say we have a 500 GeV W boson, this corresponds to a grooming radius around $R_{\text{SD}} \approx 0.4$ which gives an area πR_{SD}^2 still larger than the typical $2\pi R_{\text{filt,trim}}^2$ obtained with filtering or trimming. Applying an extra filtering step to the result of the mMDT/SD, as already suggested in the original work [59], would further groom the jet at angular scales smaller than the original grooming radius, further reducing the effective jet area and hence further reducing the peak shift and smearing.

To summarise in a more compact way the above findings and quantify more precisely the per-

²The remaining apparent small shift of the peak position still remaining after pileup subtraction is likely due to a selection bias coming from our selection of the hardest jets before grooming or subtraction is applied. Applying the p_t cut on the subtracted jets would likely remove that remaining bias.

³Note that grooming techniques do more than reducing the catchment area of a jet. Noticeably, the selection of the hardest subjets introduces a bias towards including upwards fluctuations of the background. This positive bias is balanced by a negative one related to the perturbative radiation discarded by the grooming. These effects go beyond the generic features explained here.

formance of the various grooming techniques analysed in this study, we have measured the position and width of the peak in all our different scenarii. We define this by looking, for each individual mass distribution, at the narrowest mass window that contains 50% of the reconstructed jets.⁴ We then define the peak position as the median of that window and the peak width as the width of the mass window. The findings from the spectra are quantitatively summarised in Fig. 7.3 for the peak position (left) and width (right), for both the top (top) and the W (bottom) samples. Since we have already seen that the combination of grooming with area–median subtraction was performing better than grooming alone, and that mMDT/SD was working best when supplemented with a filtering step, we focused only on these optimal cases.

We see that all methods succeed at maintaining the jet mass independent of the pileup conditions and around the expected mass value of 171 GeV, although the reconstructed average top mass tends to be slightly underestimated. This works well for all groomers even at large pileup multiplicity, although some differences are observed between the various groomers. All the groomers studied here improve the mass resolution to different degrees, but in any case perform better than pileup subtraction alone, as expected. Combining area-subtraction and filtering yields the least sensitivity to pileup in terms of the peak position and width, especially for the $t\bar{t}$ sample. For the WW sample, all groomers tend to perform similarly. The observed differences are anyway beyond the scope of this study and depend also on the specific choice of parameters for the different groomers.⁵

7.3 Summary

A brief Monte Carlo study of the effect of jet grooming techniques on the jet mass reconstruction in $Z' \rightarrow t\bar{t}$ and $Z' \rightarrow WW$ final states has been conducted. Jet trimming, filtering, the modified Mass-Drop tagger and Soft Drop are used by themselves, or in combination with the pileup subtraction using the four-vector area, to reconstruct the single jet mass and evaluate the stability of the mass scale and resolution at pileup levels of 30, 60, 100, and 140 extra pp collisions, in addition to the signal event. We find that for these particular final states all groomers work well for maintaining the mass scale and resolution, provided they are applied together with pileup subtraction so as to benefit both from the average shift correction from subtraction and noise reduction from the grooming. The results are stable up to large pileup multiplicities. In the case of the mMDT/SD procedures, we find useful to supplement them with an extra filtering step as suggested in the original work from Ref. [59].

Finally, note that a more thorough comparison of the performance of these groomers would require varying their parameters. This goes beyond the scope of this study.

⁴The original Boost report used 67%. In the case of the top mass, where a certain fraction of the events are initially compatible with a W boson, we find that a lower fraction, more often including only the top peak, is a more representative quality figure.

⁵A perhaps surprising feature is that in the top case, filtering alone tends to outperform mMDT/SD+filtering in terms of dispersion. We suspect that it might be related to situations with extra radiation at large angle and that radiation, kept by filtering alone but groomed away by the mMDT/SD de-clustering procedure. This would agree with the absence of that phenomena in the W case but seems to contradict the presumption that large-angle radiation in the case of the top should happen before the on-shell top decays. In that respect, it would be helpful to investigate more carefully the dependence on R_{filt} and n_{filt} .

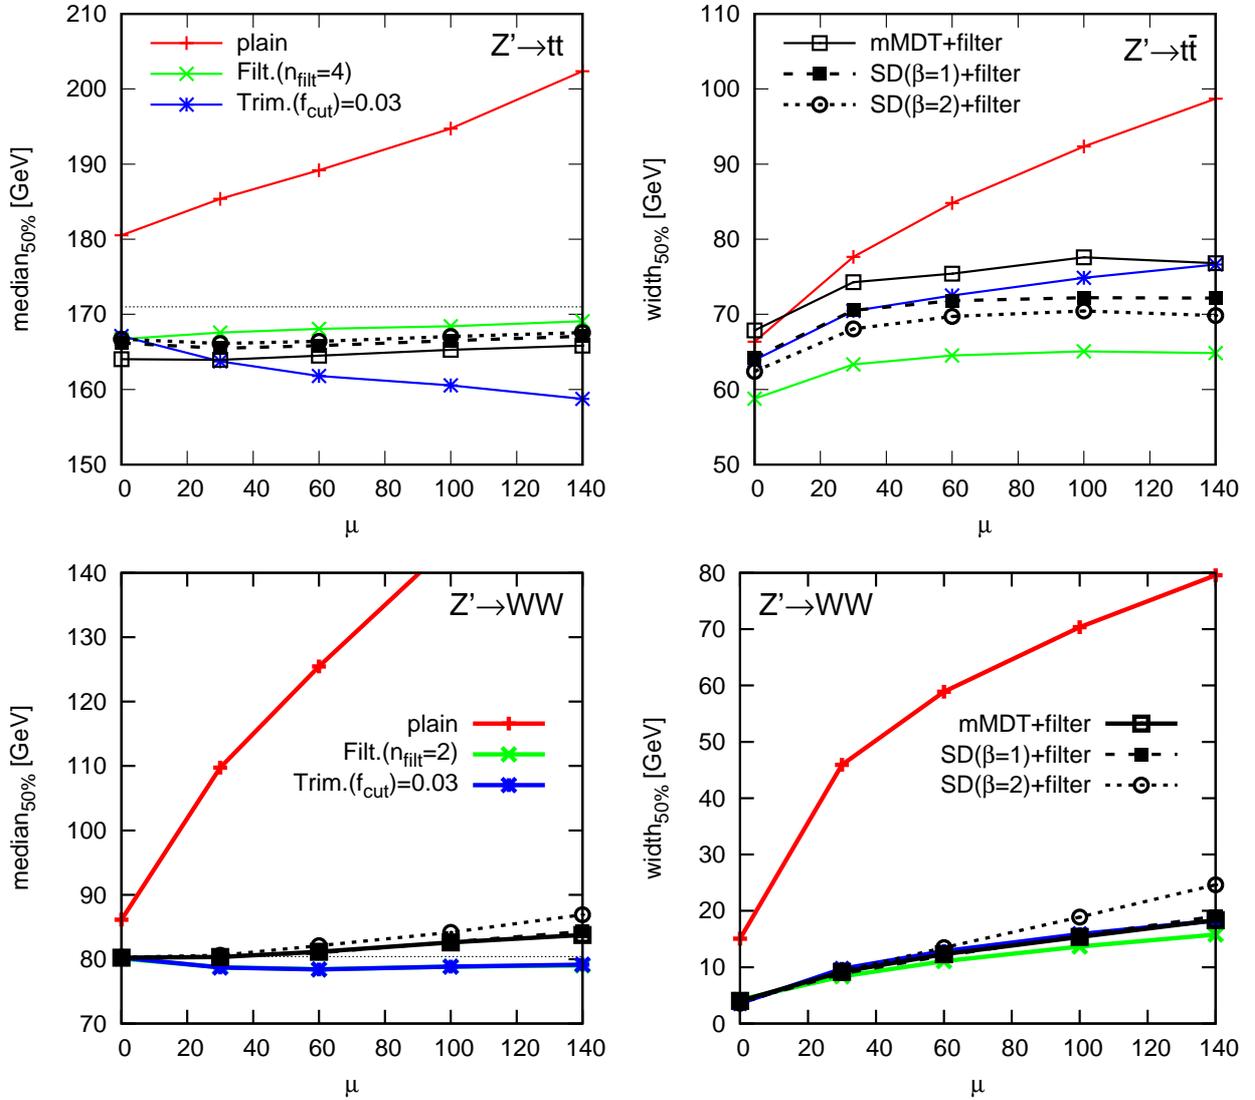


Figure 7.3: Position (leftmost figure) and width (rightmost figure) of the reconstructed jet mass distribution in $Z' \rightarrow t\bar{t}$ final states (top row) and WW events (bottom row), as a function of the pileup activity μ , for various jet grooming techniques. Area–median pileup subtraction is applied in all cases.

Chapter 8

Analytical insight

8.1 Foreword

We have seen that groomers are an efficient way to reduce pileup contamination from fat jets. Another aspect of groomers is that, even in the absence of pileup, they affect the structure of the hard jets in the event. This is usually a desired feature since it implies a reduction of the QCD background compared to the heavy-object signal when tagging boosted objects.

In this chapter we illustrate that the perturbative behaviour of groomers is amenable to an analytic treatment. After the very first results on the MassDrop Tagger, trimming and pruning in Refs. [107, 110], we have shown [21] that an analytic treatment is also possible for SoftDrop. To keep the discussion relatively contained and to avoid departing too much from the main theme of this document, we will just concentrate on a single observable, namely the groomed jet radius after one applies SoftDrop to a boosted jet. We will show that it can be computed in QCD in for boosted objects. This also covers the case of the modified MassDropTagger [107] which can be obtained from the SoftDrop results by taking the limit $\beta \rightarrow 0$. We refer directly to [107], [21] and [52] for additional analytic results for substructure techniques.

At first sight, this does not seem directly related to pileup mitigation. However, we believe that an analytic understanding of common tools in the field of jet substructure is crucial-enough that we devote a few pages to that topic. Furthermore, one might ultimately want to use insight from QCD first-principle calculations to develop optimal pileup subtraction techniques, e.g. to control their biases and sensitivities to different dynamical properties of the jet. The analytic control over groomers described in the next few pages can therefore also be seen as a step towards this broader goal. In particular, the groomed jet radius is one observable that is directly related to the residual sensitivity of the groomed jet to pileup. Gaining analytic control over it could be helpful in devising better pileup mitigation techniques.

We will also provide a brief Monte-Carlo study of non-perturbative effects for SoftDropped jets. In practice, it is partially possible to address the question of non-perturbative effects from an analytic perspective. Unfortunately, this work is still too preliminary to deserve a place here, although we can say that the underlying idea follows the non-perturbative calculation done for the modified Mass-Drop Tagger in Ref. [107].

8.2 Groomed Jet Radius

Because the soft drop procedure is defined through declustering a Cambridge/Aachen branching tree, there is a well-defined and IRC-safe meaning to the groomed jet radius. Concretely, the groomed jet radius R_g is the angle between the branches that first satisfy the SoftDrop condition (6.1). This is sensible since for a Cambridge/Aachen tree, all subsequent branches are separated by an angle less than R_g . From a practical perspective, R_g is particularly interesting, since the groomed jet area is approximately πR_g^2 . Thus, R_g serves as a proxy for the sensitivity of the groomed jet to possible contamination from pileup.

8.2.1 Analytic calculation

In what follows, we will provide an analytic calculation of the R_g distribution valid in the small-angle limit, $R_g \ll R \ll 1$, pertinent e.g. for the boosted jet regime, and focus on the leading logarithmic terms.

At this accuracy it is sufficient to assume that the jet is composed of the ‘‘leading parton’’ accompanied by independent soft-and-collinear emissions strongly ordered in angle and energy. Assume we have n such emissions with strongly-ordered angles to the jet axis, $\theta_1 \gg \theta_2 \gg \dots \gg \theta_n$. We also denote by z_1, \dots, z_n their respective transverse momentum fraction. The grooming radius will be θ_i if emission i is kept by Soft Drop and all previous emissions $1, \dots, i-1$ are groomed away. The grooming radius condition for that set of emissions can thus be written as

$$\delta(R_g(\{\theta_i, z_i\}) - R_g) = \sum_{i=1}^n \Theta_i^{\text{in}} \prod_{j=1}^{i-1} \Theta_j^{\text{out}} \delta(\theta_i - R_g) + \prod_{i=1}^n \Theta_i^{\text{out}} \delta(R_g), \quad (8.1)$$

where Θ_i^{in} indicates that emission i is kept by Soft Drop while Θ_i^{out} indicates that it is groomed away. The last term in the above expression corresponds to the situation where all emissions are groomed and where we have decided to set $R_g = 0$ in this case for simplicity. Recall also that if emission i is kept by grooming, all subsequent emissions at smaller angles will be kept as well.

We then have to sum over any number of such emissions and obtain the following expression for the groomed jet radius

$$\frac{R_g}{\sigma} \frac{d\sigma}{dR_g} = \lim_{\epsilon \rightarrow 0} \sum_{n=0}^{\infty} \prod_{i=0}^n \int_{\epsilon}^R \frac{d\theta_i}{\theta_i} \int_{\epsilon}^1 dz_i p(z_i) \frac{\alpha_s(z_i \theta_i p_t)}{\pi} \Theta(\theta_i > \theta_{i+1}) \quad (8.2)$$

$$R_g \delta(R_g(\{\theta_i, z_i\}) - R_g) \exp \left[- \int_{\epsilon}^R \frac{d\theta}{\theta} \int_{\epsilon}^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \right].$$

To write (8.2), we have introduced ϵ to regulate the soft and collinear divergences. The last, exponential, term corresponds to virtual corrections and will cancel the soft and collinear divergences in the end so that one can safely take the limit $\epsilon \rightarrow 0$. The argument of the strong coupling is taken as the relative transverse momentum of the emitted gluon w.r.t. the parent parton. Finally, $p(z)$ is the appropriate splitting function different for quark-initiated jets and gluon-initiated jets.

To evaluate (8.2), we first notice that the $\delta(R_g)$ in (8.1) will vanish in the limit $\epsilon \rightarrow 0$. Indeed, if all the emissions are outside of the grooming radius, the sum over real emissions factorises into

$$\sum_{n=0}^{\infty} \prod_{i=0}^n \int_{\epsilon}^R \frac{d\theta_i}{\theta_i} \int_{\epsilon}^1 dz_i p(z_i) \frac{\alpha_s(z_i \theta_i p_t)}{\pi} \Theta(\theta_i > \theta_{i+1}) \Theta_i^{\text{out}} = \exp \left[\int_{\epsilon}^R \frac{d\theta}{\theta} \int_{\epsilon}^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \Theta^{\text{out}} \right]. \quad (8.3)$$

Combined with virtual corrections, this gives

$$\exp \left[- \int_{\epsilon}^R \frac{d\theta}{\theta} \int_{\epsilon}^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \Theta^{\text{in}} \right] \quad (8.4)$$

which vanishes in the limit $\epsilon \rightarrow 0$.

Then, for the remaining terms in (8.1), it is helpful to reorganise the emissions where we pull out in front the integration over θ_i such that $\theta_i = R_g$, and rename it θ_0 . We then separate the remaining emissions in two groups: those at larger angles for which we have to impose the condition that they are outside the Soft Drop region (see (8.1)), and those at smaller angles which are automatically kept by Soft Drop. The first series can set to have n emissions with momenta and angles $(z_1, \theta_1), \dots, (z_n, \theta_n)$ and the second series \bar{n} emissions with momenta and angles $(\bar{z}_1, \bar{\theta}_1), \dots, (\bar{z}_{\bar{n}}, \bar{\theta}_{\bar{n}})$, where we have to sum over n and \bar{n} . This gives

$$\begin{aligned} \frac{R_g}{\sigma} \frac{d\sigma}{dR_g} &= \lim_{\epsilon \rightarrow 0} \int_{\epsilon}^R \frac{d\theta_0}{\theta_0} \int_{\epsilon}^1 dz_0 p(z_0) \frac{\alpha_s(z_0\theta_0 p_t)}{\pi} R_g \delta(R_g(\theta_0 - R_g)) \Theta(z_0 > z_{\text{cut}}(R_g/R)^\beta) \\ &\quad \sum_{n=0}^{\infty} \frac{1}{n!} \prod_{i=0}^n \int_{\theta_0}^R \frac{d\theta_i}{\theta_i} \int_{\epsilon}^1 dz_i p(z_i) \frac{\alpha_s(z_i\theta_i p_t)}{\pi} \Theta(z_i < z_{\text{cut}}(\theta_i/R)^\beta) \\ &\quad \sum_{\bar{n}=0}^{\infty} \frac{1}{\bar{n}!} \prod_{j=0}^{\bar{n}} \int_{\epsilon}^{\theta_0} \frac{d\bar{\theta}_j}{\bar{\theta}_j} \int_{\epsilon}^1 d\bar{z}_j p(\bar{z}_j) \frac{\alpha_s(\bar{z}_j\bar{\theta}_j p_t)}{\pi} \exp \left[- \int_{\epsilon}^R \frac{d\theta}{\theta} \int_{\epsilon}^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \right]. \end{aligned} \quad (8.5)$$

Once again, the sums over n and \bar{n} lead to exponentials:

$$\begin{aligned} \sum_{n=0}^{\infty} \frac{1}{n!} \prod_{i=0}^n \int_{\theta_0}^R \frac{d\theta_i}{\theta_i} \int_{\epsilon}^1 dz_i p(z_i) \frac{\alpha_s(z_i\theta_i p_t)}{\pi} \Theta(z_i < z_{\text{cut}}(R_g/R)^\beta) \\ = \exp \left[- \int_{\theta_0}^R \frac{d\theta}{\theta} \int_{\epsilon}^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \Theta(z < z_{\text{cut}}(\theta/R)^\beta) \right], \end{aligned} \quad (8.6)$$

and

$$\sum_{\bar{n}=0}^{\infty} \frac{1}{\bar{n}!} \prod_{j=0}^{\bar{n}} \int_{\epsilon}^{\theta_0} \frac{d\bar{\theta}_j}{\bar{\theta}_j} \int_{\epsilon}^1 d\bar{z}_j p(\bar{z}_j) \frac{\alpha_s(\bar{z}_j\bar{\theta}_j p_t)}{\pi} = \exp \left[- \int_{\epsilon}^{\theta_0} \frac{d\theta}{\theta} \int_{\epsilon}^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \right]. \quad (8.7)$$

These exponentials combine with the virtual corrections in particular to cancel the divergence at $\epsilon \rightarrow 0$ as expected and give, after integration over θ_0

$$\frac{R_g}{\sigma} \frac{d\sigma}{dR_g} = \int_0^1 dz_0 p(z_0) \frac{\alpha_s(z_0 R_g p_t)}{\pi} \Theta\left(z_0 > z_{\text{cut}} \frac{R_g^\beta}{R^\beta}\right) \exp \left[- \int_{R_g}^R \frac{d\theta}{\theta} \int_0^1 dz p(z) \frac{\alpha_s(z\theta p_t)}{\pi} \Theta\left(z > z_{\text{cut}} \frac{\theta^\beta}{R^\beta}\right) \right]. \quad (8.8)$$

The above expression can easily be evaluated for quark and gluon jets, using a 1-loop running-coupling approximation for α_s . This is what we have done for the comparison to Monte-Carlo simulations presented in the next Section. It is nevertheless illustrative to consider the fixed-coupling approximation which gives (up to subleading power corrections)

$$\int_0^1 dz_0 p(z_0) \frac{\alpha_s}{\pi} = \frac{\alpha_s C_R}{\pi} \left[\log \left(\frac{1}{z_{\text{cut}} R_g^\beta} \right) + B_i \right], \quad (8.9)$$

$$\int_{R_g}^R \frac{d\theta}{\theta} \int_0^1 dz p(z) \frac{\alpha_s}{\pi} \Theta\left(z > z_{\text{cut}} \frac{\theta^\beta}{R^\beta}\right) = \frac{\alpha_s C_R}{\pi} \left[\log \left(\frac{R^{\beta/2}}{z_{\text{cut}} R_g^{\beta/2}} \right) + B_i \right] \log \left(\frac{R}{R_g} \right), \quad (8.10)$$

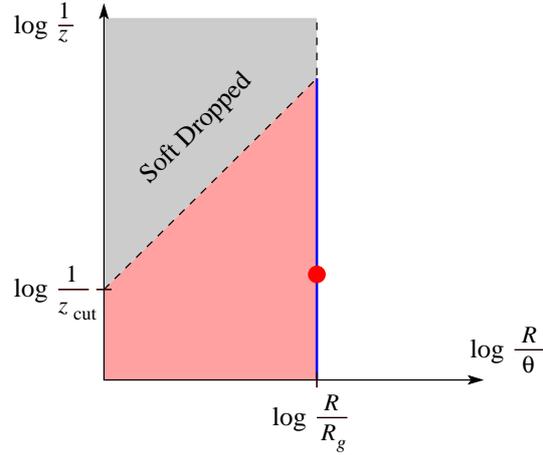


Figure 8.1: Phase space for emissions relevant for groomed jet radius R_g in the $(\log \frac{1}{z}, \log \frac{R}{\theta})$ plane. The region groomed away by Soft Drop is grey and the first emission satisfying the soft drop criteria is illustrated by the red dot which can be anywhere on the thick blue line. The forbidden emission region (the Sudakov exponent) is shaded in pink.

with $C_R = C_F$ and $B_i = -3/4$ for quark-initiated jets, and $C_R = C_A$ and $B_i = (11C_A - 4n_f T_R)/(12C_A)$ for gluon-initiated jets. This corresponds to a double-logarithmic suppression in R/R_g which becomes single-logarithmic for $\beta = 0$.

It is also helpful to give a physical interpretation of (8.8). This is illustrated in Fig. 8.1 where we represent the phase-space available for emission in the $(\log \frac{1}{z}, \log \frac{R}{\theta})$ plane. If one wants a given grooming radius R_g , we need an emission at an angle R_g passing the Soft Drop condition. This corresponds to the red dot in the Figure which can lie anywhere on the thick blue line. It leads to the prefactor in (8.8) where the integration over z_0 corresponds to the red dot spanning the blue line. We then need to guarantee that the Soft Drop procedure does not stop at a larger radius. We should therefore veto all real emissions at larger angles and that would pass the Soft Drop criterion. This is represented by the pink shaded area in the figure and to the Sudakov suppression factor in Eq. (8.8). Since real emissions are vetoed in that region, only virtual corrections remain. In all the rest of the phasespace, real emissions are allowed and cancel against virtual corrections.

On a more technical basis, it is worth noting that the groomed jet radius R_g is also simple in terms of multiple-emission corrections. Indeed, once one emission satisfies the soft drop criteria, the jet radius is set, so multiple emissions do not contribute to this observable. We can also verify that non-global contributions are suppressed by R_g for $\beta < \infty$. For these reasons, we believe that the expression in Eq. (8.8) is fully accurate to single-logarithmic level.¹

8.2.2 Comparison to Monte Carlo

There are two different ways one can define the groomed jet radius in Monte Carlo. The first method is to simply measure the R_g value of the Cambridge/Aachen branching that satisfies the Soft Drop condition. A second approach, more directly connected with pileup mitigation, is to determine the effective radius of the groomed jet from its active area (see Section 2.3.3). An effective jet radius R_{eff}

¹Strictly speaking, NLL accuracy would require evaluating the strong coupling at two loops in the CMW scheme [111], i.e. including two extra terms proportional to β_1 and K , suppressed by one power of α_s compared to the leading β_0 term.

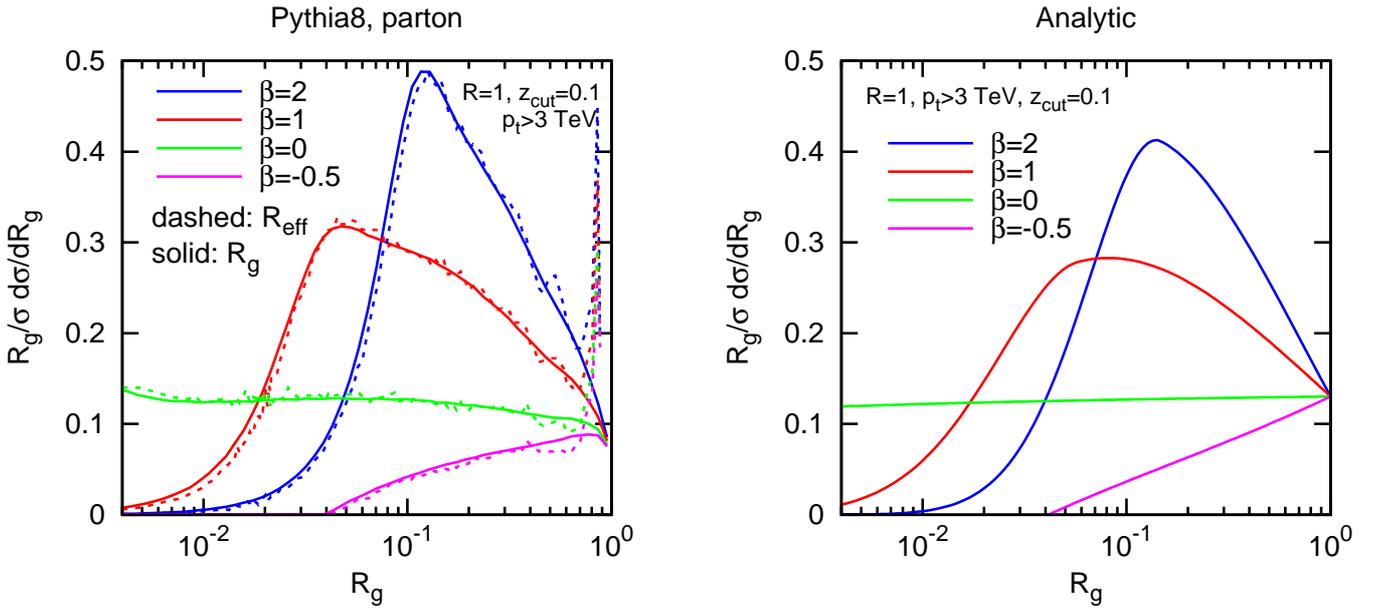


Figure 8.2: Comparison of the jet radius R_g distribution extracted from Pythia 8 (left, solid) and inferred from the active area in Pythia 8 using $R_{\text{eff}} = \sqrt{A_{\text{active}}/\pi\xi}$ (left, dashed), and computed to modified leading logarithmic accuracy (right). The original jet radius is set to be $R = 1$ and the jets have an ungroomed energy of 3 TeV. The Soft Drop parameter is $z_{\text{cut}} = 0.1$, while β is varied.

can then be defined from the groomed jet active area using:

$$R_{\text{eff}} \equiv \left(\frac{A_{\text{active}}}{\pi\xi} \right)^{1/2}, \quad (8.11)$$

where A_{active} is the active jet area, and $\xi \simeq (1.16)^2$ accounts for the fact that a typical Cambridge/Aachen jet of radius R has an average active area $\xi\pi R^2$.²

To compare R_g and R_{eff} and to validate our analytic calculation (8.8), we have generated with Pythia8 (v8.176) 14 TeV $qq \rightarrow qq$ events and selected anti- $k_t(R = 1)$ jets with $p_t \geq 3$ TeV. We have applied the Soft Drop procedure with $z_{\text{cut}} = 0.1$ and different values of β and computed R_g and R_{eff} . To obtain R_{eff} in practice, we have computed the groomed jet area using active areas as implemented in `FastJet` (v3), and we used a ghost area of 0.0005 and 10 repetitions in order to reach sufficiently small values of R_{eff} .

The resulting distributions are plotted in the left plot of Fig. 8.2. With the ξ offset factor, the two techniques give remarkably similar results, giving strong evidence that the groomed jet radius R_g is an effective measure of pileup sensitivity. The main difference is the spike at $R_{\text{eff}} = 1/\sqrt{\xi}$, corresponding to cases where the first Cambridge/Aachen branching already satisfies the soft drop condition, yet typically with $R_g < 1$. The nice reduction of the jet area even with mild grooming (e.g. $\beta = 2$) suggests that Soft Drop should work well for pileup mitigation.

In the right plot of Fig. 8.2, we show the distribution obtained from our analytic calculation, Eq. (8.8). There is good qualitative agreement with Pythia for a range of angular exponents β ,

²The numerical value for ξ can be read from Fig. 5.7. Strictly speaking, this result is only valid for a jet made of two particles separated by R_g with one of them much softer than the other. However, for Cambridge/Aachen jets, one expects that this would not vary much for more symmetric two-particle configurations (see e.g. Ref. [112]).

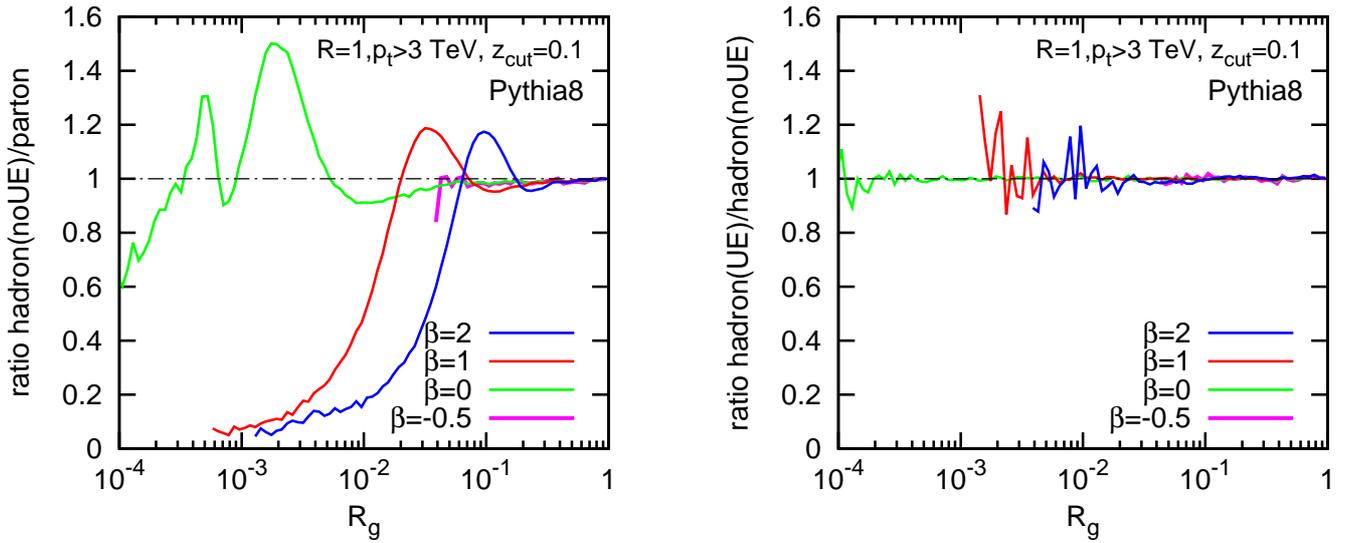


Figure 8.3: Effect of non-perturbative corrections on R_g . The left plot shows the ratio between hadron level and parton level predictions obtained with Pythia 8 (without UE). The right plot instead shows the ratio of hadron-level results with and without UE.

suggesting that our modified leading-logarithmic calculation for R_g captures much of the relevant physics effects present in the Monte Carlo simulation and that analytic calculations can be helpful in understanding how grooming techniques mitigate pileup.

Note that when increasing the jet p_t , one would effectively reduce the coupling $\alpha_s(p_t R)$ arising naturally from (8.8). Both the pre-factor and Sudakov exponent would then be reduced, shifting the peak in the R_g distribution to smaller values.

8.3 Non-Perturbative Contributions

In the above analytic calculations and Monte-Carlo simulations, we only considered the distributions generated at parton level. In this section, we present a brief Monte Carlo study to estimate the impact that non-perturbative effects from hadronisation and UE can have on the groomed jet radius distribution.

In Fig. 8.3, we show the effect of hadronisation (left) and UE (right) on the R_g distribution for different values of β . In the case of hadronisation, we plot the ratio between the hadronic and partonic distributions obtained from Pythia 8. In the case of UE, we plot the ratio between the distributions with and without UE. Apart from including non-perturbative effects, the details of the analysis are the same as for the previous Monte Carlo studies.

The left plot in Fig. 8.3 shows that Soft Drop declustering pushes the onset of hadronisation corrections to smaller values of R_g with decreasing β . As shown in the right plot of Fig. 8.3, Soft Drop has the remarkable ability to reduce the UE contribution to almost zero down to very small values of R_g .

Part III

Towards more efficient techniques

Chapter 9

Challenges beyond the area–median subtraction

9.1 Motivation

Let us put aside the specific case of fat jets and their grooming, and come back to our original task: providing a generic-purpose pileup mitigation technique that can be used for standard jet reconstruction at the LHC. This covers the majority of cases where jets are used at the LHC today. In the same kind of spirit, in most of the analysis where jets are used, the fundamental quantity that we use is the jet transverse momentum. In this last part, we shall therefore focus on pileup mitigation for “standard” jet clustering and mainly consider the most widely used jet characteristic: its transverse momentum. We shall only briefly investigate the case of other jet properties like the jet mass or even other jet shapes, and some applications to boosted jets. One should however keep in mind that even if these jet substructure observables are of growing importance, they usually have to be considered together with grooming techniques for practical applications.

In the first part of this document we have seen that the area–median pileup subtraction technique does a very good job at correcting, event by event, for the transverse momentum shift induced by pileup. A benefit of the event-by-event nature of the subtraction is that it also reduces the associated resolution degradation. The area–median is therefore an efficient pileup-mitigation tool that has proven perfectly adequate for the Run I of the LHC and at least the beginning of Run II, with values of μ slightly above 40 as of October 2016. However, as the LHC luminosity increases, so will μ , the average number of pileup interactions per bunch crossing. One might expect values of μ around 60 at maximal luminosity already in Run II, and in perspectives for the High-Luminosity LHC, one quotes values of μ that can reach the 140–200 range. While the area–median subtraction method would still be efficient at correcting for the average p_t shift, with higher pileup multiplicities, the effects of fluctuations would increase substantially. From the 6 GeV observed for 30 pileup interactions in Fig. (3.8), one should expect a resolution smearing of about 7.5 GeV for 60 pileup interactions and more than 10 GeV above 100, in good agreement with a scaling proportional to $\sqrt{N_{\text{PU}}}$. This increase should not be too dramatic for very high- p_t jets where detector resolution effects dominate over the pileup smearing effects, but become an issue for low- p_t jets.

One could simply decide to increase the minimal jet p_t threshold. However, the feasibility of key measurements at the LHC — like double-Higgs production, jet vetoes and the study of the properties of the Higgs boson — depend on the experiments’ ability to reconstruct jets down to transverse momenta as low as 20 GeV. This situation is therefore far from ideal and has lead people to think

about alternative pileup mitigation techniques.

In what follows, we shall discuss new ideas going along two different main avenues. The first one follows a logic similar to the area–median subtraction technique: it focuses on correcting for the right average pileup contamination but further uses additional information in the event. The typical example here is the case of charged tracks. With current technology, the experiments can associate charged tracks to a vertex of origin and discard those coming from pileup vertices. One is then left with pileup mitigation on the neutral part of the event. Although we do not want to enter into the experimental details, this is a core idea in the Charged-Hadron Subtraction (CHS) method in CMS. In an ideal situation, this would correct exactly for the charged part of pileup. Applying pileup subtraction on the remaining neutral part, e.g. using the area–median subtraction, would mean that, in the end, the resolution degradation is only affected by the neutral fluctuations of pileup. Since about 60% of the energy goes into charged particles, this would automatically reduce the resolution degradation by a factor around $\sqrt{0.4} \approx 0.63$. With this technique, the pileup smearing effects for $N_{\text{PU}} = 100$ would only be slightly worse than what one sees without using charged-track information at $N_{\text{PU}} = 30$. We shall discuss at length the use of charged tracks to mitigate pileup in Chapter 11 below, including the adaptation of the area–median method as well as other alternatives.

The second approach is to explicitly try to reduce the effect of pileup fluctuations. Thinking about the LHC as a (very fast) camera, this can be seen as a noise reduction algorithm.¹ This is a potentially vast domain of investigation and a few methods going in that direction have already been proposed. We will discuss a few such methods below: in Chapter 10, we will discuss potential use of grooming techniques as generic pileup mitigation tools (i.e. beyond their applications to fat jets); in Chapter 12 we shall report on the SoftKiller method [11] that we have introduced recently, and in Chapter 13 we will briefly comment on possible alternatives that have not yet been fully investigated but are of potential interest for the development of future pileup mitigation techniques. Note that around the same time as we introduced the SoftKiller method, another noise-reduction technique called PUPPI has also been proposed [113] and is currently increasingly used in the CMS collaboration. We will not discuss the method here but will briefly come back to it in our final discussion at the end of this document.

9.2 Challenges of “noise-reduction” techniques

Before going into the investigation of the techniques themselves, we would like to have a generic discussion about one faces when designing pileup-mitigation techniques that try to reduce the effects of fluctuations. We will argue that the main issue is that these techniques come with a high level of fine-tuning needed to keep the method unbiased.

In what follows, we will consider two kinds of methods. The first family starts by clustering the event and then, for a given jet decides on a subset of its constituents (or subjets) to keep, similar to what grooming procedures do. The other approach is to first decide on a set of particles to keep, deeming the others as pileup and removing them from the event prior to the clustering. Let us therefore consider a generic method that decides for each particle or subjet whether they are pileup particles or not, keeping the latter and rejecting the former. The discussion below trivially extends to the more generic case where one instead associates a weight to each particle or subjet.

The procedure of keeping or rejecting particles comes with two sources of bias. First some of the particles originally coming from pileup will be kept, yielding an under-subtraction, i.e. a positive

¹With the important difference that the boundary of the hard event are not as well defined as boundaries between objects on a standard picture.

bias. Then, some of the particles from the hard event will be deemed as pileup and removed, giving a negative bias. One will therefore only have an overall unbiased method if these two sources of bias compensate each other on average. This does require some level of fine-tuning of the free parameters of all the methods considered below. In other words, the positive tail of upwards pileup fluctuations, kept in the event, will have to balance the negative bias coming from, essentially soft, particles in the hard event that have been rejected by the pileup mitigation technique. The resulting fine-tuning has to be contrasted with the area–median method for which the estimation of the pileup density ρ remains acceptable in a relatively large range of the grid-size parameter used for the estimation. However, if the bias of our new method can be kept small, one can expect an improvement of the jet resolution if the fluctuations of the positive and negative biases discussed above are themselves small. In that respect, it is reasonable that the set of parameters for which the bias remains small should also roughly correspond to the values of the parameters for which the resolution degradation is at its minimum.

An important point to be aware of is that the parameters for which the bias remains small are likely to vary with the process under consideration, the energy scale of the hard event and the pileup conditions. It is therefore important to scan different hard processes, jet p_t and pileup multiplicities when validating any new method.

The last potential issue to be aware of concerns methods applying an event-wide pileup subtraction prior to the clustering. Pileup is roughly uniformly spread across the event while the hard jets are localised, peaked, energy deposits. One must therefore be careful, when balancing the positive pileup fluctuations left uniformly in the event, to avoid leaving some of these positive fluctuations at the expense of over-subtracting the hard jets. Ideally, an efficient pileup mitigation technique should therefore cut pileup more aggressively when moving away from the hard jets. This is somehow encoded in the PUPPI approach [113] but is not taken into account in a method like the SoftKiller (see Chapter 12). This is not a concern if one wishes to subtract pileup from jets of a given radius but requires a re-tuning for different radii (see also Section 13.2.1).

In the long term, it would be very interesting to understand if a first-principles QCD understanding of jet substructure could help controlling the biases mentioned above, in particular their variation wrt the jet p_t , the jet radius and the pileup conditions, as well as how the subtraction should vary when progressively moving away from the hard jets.

Chapter 10

Grooming as a generic tool for pileup subtraction

10.1 Foreword and motivation

Given the success of grooming techniques in the context of fat jets and boosted objects reconstruction observed in Chapter 7, one naturally wonders if they could also be useful for pileup mitigation in a more generic context. This is what we discuss in this Chapter. This work was initiated in the context of the “Les Houches, Physics at TeV colliders” workshop in 2013 but never really finalised or, at least, never pushed to a point where I thought it contained solid-enough a message to deserve a place in the workshop proceedings or as a standalone publication. The study below is therefore unpublished and should really be considered as preliminary. I find useful to include it in this review, not only because it directly pertains to the topic of pileup mitigation, but also because I want this document to keep an opened door towards future developments and I believe that grooming techniques can play a key role in that context.

That said, the study carried over the next few pages is also similar in spirit to the tests we have carried with the “Cambridge/Aachen+filtering” algorithm in our heavy-ion studies in Chapter 4 up to one crucial difference. In our heavy-ion studies, we considered the “Cambridge/Aachen+filtering” for both the hard reference and the full event including Underlying Event, i.e. as a jet algorithm, used in combination with the area–median subtraction. Here we only want to consider grooming on the full jet and compare that to the ungroomed hard jets. In a sense, the idea is really to consider grooming as a tool for pileup mitigation, which implies that, in the absence of pileup, we work with ungroomed anti- k_t jets.¹ Despite this sizeable difference, the reason for which we want to use grooming remains the same: by effectively reducing the area of the jet, we expect the groomed jet to be less sensitive to pileup fluctuations i.e. have a better energy resolution. This obviously comes with a price: grooming also affects the hard contents of jets and so, compared to the robust area–median approach, there is no guarantee that the average bias will remain close to 0. We refer the reader to Chapter 9 for a discussion of biases and reduction of fluctuation effects.

¹It might also be interesting to consider the case where the grooming is included as part of the jet definition, as done in Chapter 4.

10.2 Grooming configurations and workflow

Our study includes the three types of jet grooming techniques listed below. In each case, the subjects are subtracted using the area–median method before making the selection of which subjects have to be kept.

Filtering. We will consider both $n_{\text{filt}} = 2$ and $n_{\text{filt}} = 3$ and scan over a range of R_{filt} values, with subjects clustered with the Cambridge/Aachen algorithm.

Trimming. We use the Cambridge/Aachen algorithm with $R_{\text{trim}} = 0.2$ and scanned over a range of f_{trim} values. The reference p_t is taken to be that of the unsubtracted jet.

Area-trimming. This has not been introduced before. Area-trimming starts similarly to trimming by re-clustering the jet into subjects using the Cambridge/Aachen algorithm with $R_{\text{sub}} = 0.2$. We then keep the subjects for which the (subtracted) p_t is larger than $n_\sigma \sigma \sqrt{A_{\text{subject}}}$, with σ the pileup fluctuations as estimated with **FastJet**. The motivation for this choice is that, after the pileup subtraction, one has removed the negative tail of the pileup fluctuations but one is left with the positive tail which has $\sigma \sqrt{A_{\text{subject}}}$ as a characteristic scale. We have tested a range of n_σ values in our studies.

Note that in a more complete study it would be interesting to add Soft Drop to the above list and to study different subject radii for trimming and area-trimming.

The rest of the simulation procedure follows the same pattern as what we have done in Section 3.2 for our basic assessment of the area–median approach. We simulate hard events at $\sqrt{s} = 13$ TeV with Pythia 8.186 (tune 4C), with the Underlying Event switched off. We have only considered dijet events although tests with different event samples would be desirable. The hard events will be embedded in pileup, also generated with the same Pythia8 version and tune. Following the same terminology as in Chapter 3.2, the latter will be referred to as the “full event”. The estimation of the pileup properties are obtained using a grid-based area–median estimation with a cell size of 0.55 and rapidity rescaling.

In practice, we select jets from the hard event with $|y| < 4$ and a given $p_{t,\text{cut}}$, then match the jets in the full event to the hard jets.² Once we have a pair of matching jets, we apply one of the grooming techniques listed above and compute the p_t difference, Δp_t , between the original hard jet and the subtracted full jet. As before, we study the average shift $\langle \Delta p_t \rangle$, which should be close to 0 for an unbiased method, and the dispersion $\sigma_{\Delta p_t}$ which measures the residual resolution degradation.

10.3 Results

A summary of our findings is presented in Figs. 10.1 and 10.2. In both cases, we show the average shift and the corresponding dispersion for the different grooming techniques under scrutiny, varying their free parameter. In Fig. 10.1, we have fixed the pileup conditions to $\mu = 60$ and varied the cut on the p_t of the jet, while in fig. 10.2 we have worked with a fixed p_t cut of 50 GeV and varied the pileup conditions.

The first properties we expect from an efficient pileup mitigation technique is that it remains unbiased and robust, i.e. that the average shift is close to 0 regardless of the pileup conditions and of the details of the hard process. In all cases, we see that this requirement is not trivially satisfied. At the very least, the free parameter of each grooming technique needs to be carefully tuned for the bias

²As before, we do that by requiring that the p_t of the common constituents is at least 50% of the p_t of the hard jet.

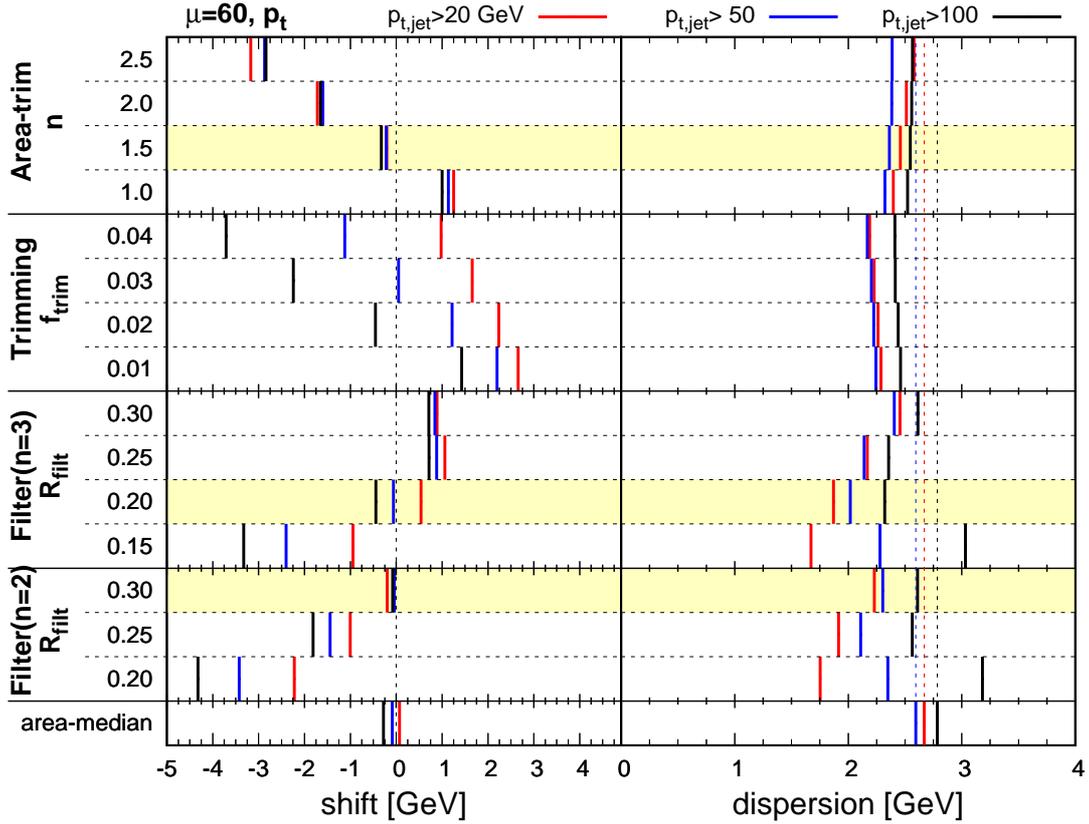


Figure 10.1: Summary plot of the average shift and bias for the different grooming methods we have tested and different values of their respective parameters. We have used $\mu = 60$ and different cuts on the jet p_t : 20 GeV (red), 50 GeV (blue) and 100 GeV (black). The “standard” area–median results are quoted at the bottom of the plot. On the average shift plot (left), the vertical dashed line corresponds to a zero bias, and on the dispersion plot (right), the vertical dashed lines correspond to the area–median results.

to remain small. This contrasts with the area–median approach where the average bias is very close to 0 in a wide variety of cases and where the free parameter, controlling the size of the patches, can be varied around the optimal scale without affecting significantly the performance.

Furthermore, we already see differences between the grooming methods at the level of the average bias. Filtering with $n_{\text{filt}} = 2$ and $R_{\text{filt}} = 0.3$, and area–trimming with $n_{\sigma} = 1.5$ show small biases, comparable with the area–median approach, for all the cases we have studied. Then filtering with $n_{\text{filt}} = 3$ and $R_{\text{filt}} = 0.2$ shows a reasonably small bias albeit with a larger dependence on the pileup multiplicity. For trimming, the smallest bias is obtained for $f_{\text{trim}} = 0.03$ but this choice appears less robust than what we observe with other grooming techniques. Physically, this seems to indicate that if we want to use a method that decides which subjects are kept based on a p_t cut, one obtains better results if the cut is set based on the pileup properties than based on the jet p_t . If some dependence on p_t is to be included — e.g. to account for the residual biases of the area–trimming method — it should at least be milder than linear.³

Turning to the dispersion, we see that the resolution degradation is, as expected, smaller than what one obtains with the area–median approach. We also see that the gain becomes larger as the

³This is also a question that one should be able to address analytically.

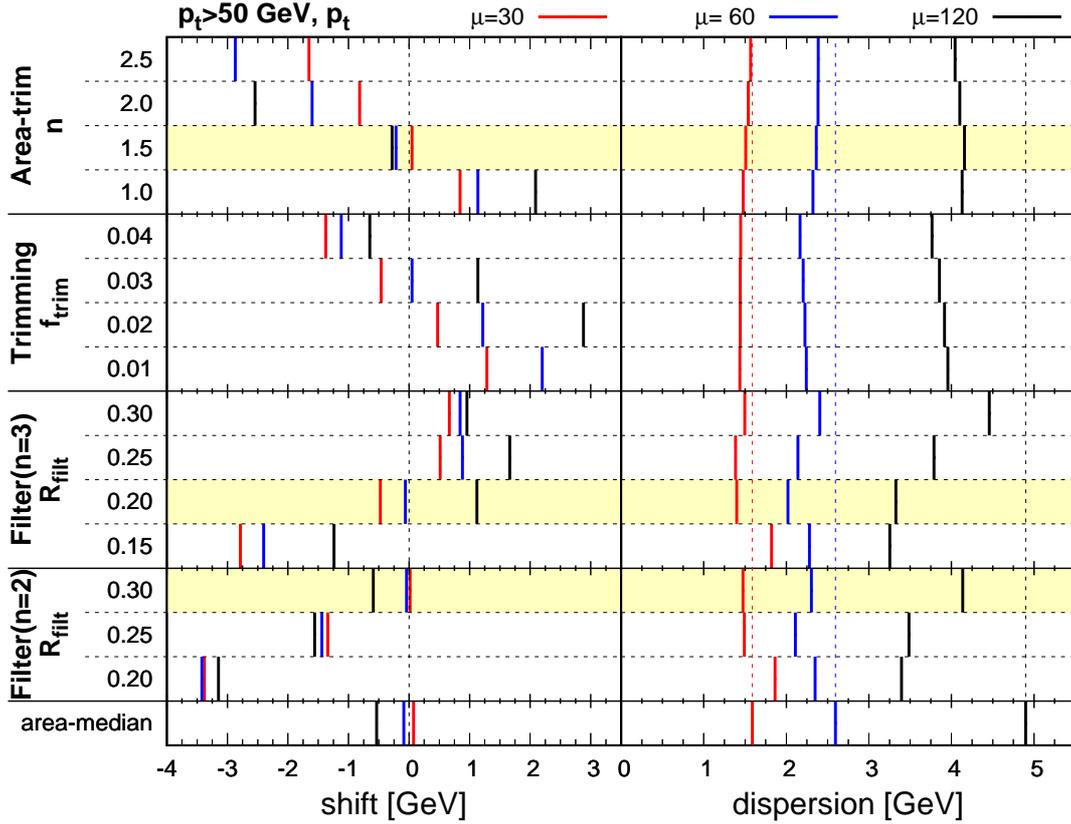


Figure 10.2: Same as Fig. 10.1 by with a fixed cut of 50 GeV on the jet p_t and, instead, varying the pileup average multiplicity $\mu = 30$ (red), 60 (blue) or 120 (black).

pileup increases.

To investigate these results a bit further, we have selected for each of the 4 groomer families the one showing the best performance and we have looked more closely at the average shift and dispersion as a function of the jet p_t , Fig. 10.3, and of the pileup conditions, Fig. 10.4. This shows clearly that filtering with $n_{\text{filt}} = 2$ and $R_{\text{filt}} = 0.3$, and area-trimming with $n_{\sigma} = 1.5$ have a bias similar to the area-median method, at least for the dijet processes that we have studied. Except for the filter at large p_t , these two options show an improvement over the area-median approach in terms of resolution degradation. At small p_t (and $\mu = 60$), we see an improvement around 5% for area-trimming and 15% for filtering. While this improvement persists across the whole p_t range for the case of area-trimming, it disappears for $p_t \geq 200$ GeV in the case of filtering. This might not be too much of an issue since pileup is most problematic at low p_t . Nonetheless, it suggests that improvements are possible at large p_t . At high pileup multiplicities, both methods show a resolution improvement slightly above 10%.

If one wishes to tolerate a slightly larger bias, filtering with $n_{\text{filt}} = 3$ and $R_{\text{filt}} = 0.2$ can also be considered and gives a more sizeable improvement in terms of dispersion, although once again, this is mostly seen at small p_t .

10.4 Concluding remarks and future prospects

The study we have presented here is fairly minimal and definitely requires more tests if one wishes to draw a firm conclusion. We see however some patterns emerging:

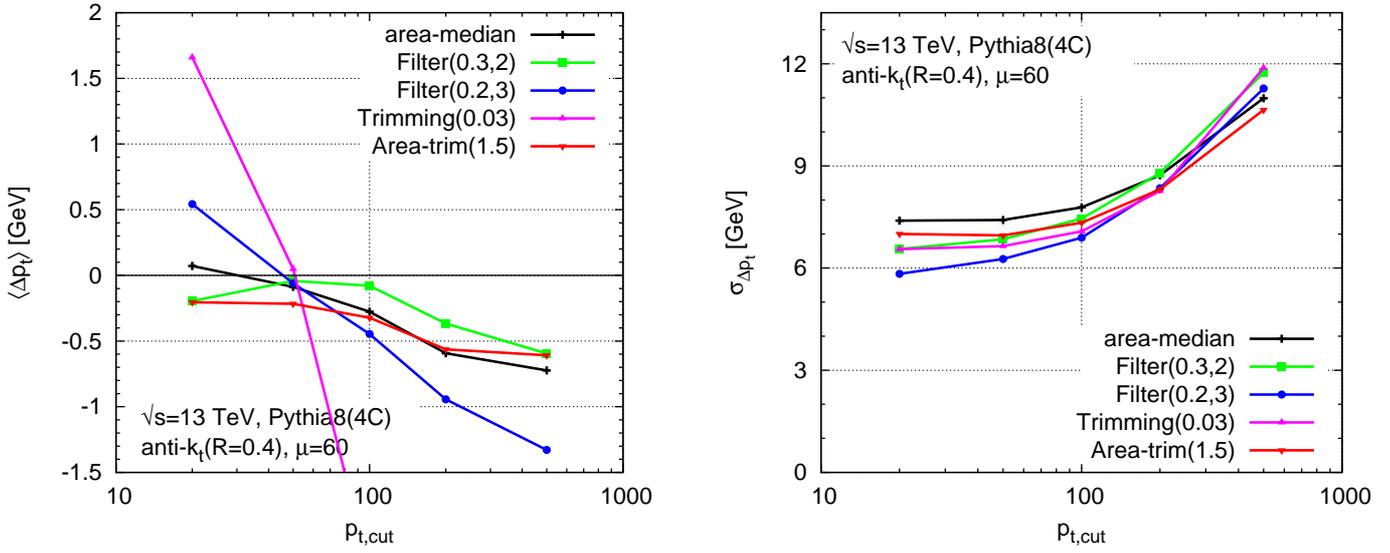


Figure 10.3: Plot of the average bias (left) and of the corresponding dispersion (right), as a function of the jet p_t cut, for a few specific grooming techniques and for $\mu = 60$,

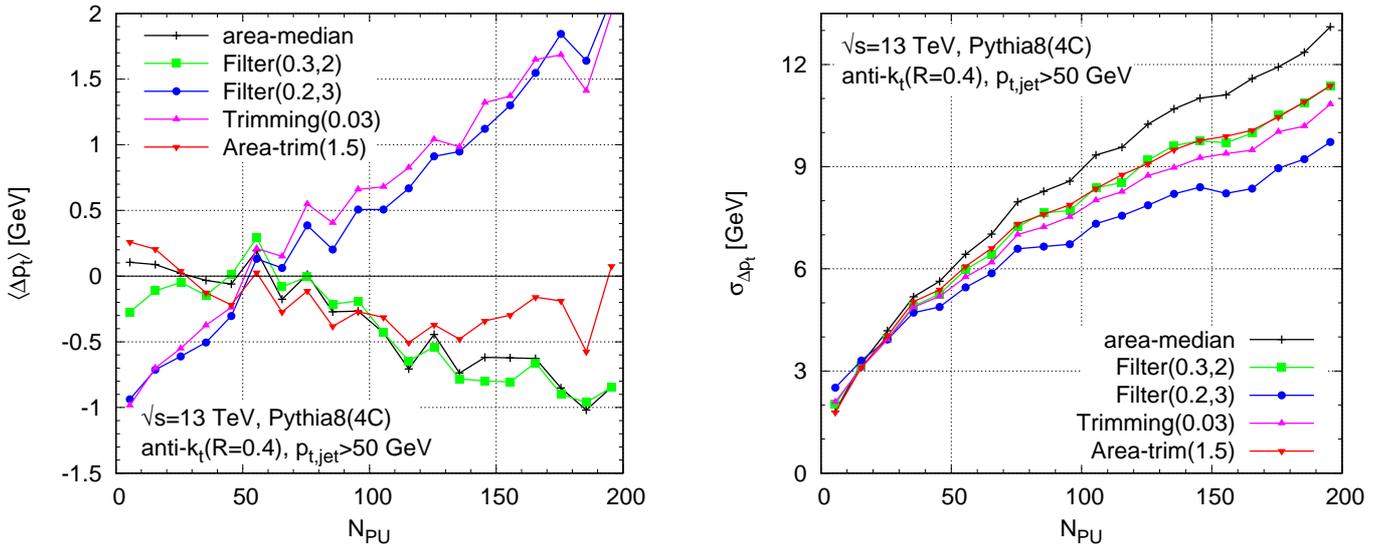


Figure 10.4: Plot of the average bias (left) and of the corresponding dispersion (right), as a function of the average number of pileup interactions, for jets with $p_t > 50$ GeV.

- using grooming techniques as a tool for pileup subtraction requires some level of fine-tuning if one wishes to keep the average bias small.
- In the region where the bias is small, one sees, as initially foreseen, a reduction of the sensitivity to pileup fluctuations.
- Filtering with $n_{\text{filt}} = 2$ and $R_{\text{filt}} = 0.3$ or with $n_{\text{filt}} = 3$ and $R_{\text{filt}} = 0.2$ as well as the, new, area-trimming with $n_{\sigma} = 1.5$ show overall good performances. Trimming seems to be disfavoured, in particular because it shows a larger dependence of the bias on the jet p_t .

The main message to remember from this preliminary pileup-mitigation study is that grooming can be a valuable tool in the context of pileup mitigation. It would also be interesting to carry further studies of the robustness of the methods that are seen to perform well here. This involves more detailed tests, including the rapidity dependence, a more complete scan over the jet p_t and N_{PU} and the study of other processes, but also tests of detector effects or applications to CHS events (see also Chapter 11 and our final discussion in Chapter 14).

Finally, one may also consider dynamical adjustments of the parameters, e.g. R_{filt} or n_{σ} , with the jet p_t and the pileup conditions. For example, we could groom more aggressively for low- p_t jets where we see larger gains in terms of resolution and where these gains are the most important for practical applications at the LHC. Note that these parameters are expected to vary relatively slowly, typically only logarithmically, with the jet p_t .

Chapter 11

Using information from charged tracks

In this chapter, we consider the situation where we can use vertex information to tell whether a charged track comes from the hard interaction vertex or from a pileup vertex. We shall assume perfect identification and measurement of charged tracks and all vertices. This can be seen as an idealised Charge-Hadron-Subtracted (CHS) approach.

In this approach, the particles in each event can therefore be split in three main categories: neutral particles, charged particles from the hard interaction vertex, and charged particles from any of the pileup vertices. In that context, we can perform an exact subtraction of pileup for the charged part of the event, keeping the charged tracks associated with the leading vertex and discarding the ones associated with a pileup vertex. However, if one wishes to retain information from the charged particles associated with a pileup vertex, it is always possible to keep them as ghosts in the event, i.e. to keep them with transverse momenta scaled down to infinitesimal values. This is what we shall consider for most of our CHS-like events in what follows.

We are therefore left with the subtraction of the neutral part of the event. In this chapter, we investigate several methods that concentrate on pileup mitigation for CHS-type of events and compare their performance.

11.1 Pileup subtraction methods for CHS events

Area–median with charged-tracks information. The generalisation of the area–median method to CHS events is straightforward: we simply compute ρ (and ρ_m) as described in Section 2.4, this time using the particles from the CHS events. The rapidity profile can also be obtained from the neutral component of minimum bias. Note that, in practice, we could either use only the neutral particles or all CHS particles for the estimation of ρ . Since this makes very little difference in the end, we decided to compute the pileup density from all particles in the CHS event.

Furthermore, we still use (2.24) to subtract pileup from the jets. In the context of CHS events, we can also slightly improve our positivity constraints for the final jet transverse momentum and mass. If, after subtraction, the transverse momentum of the jet is less than the transverse momentum carried by charged tracks from the leading vertex, we set the subtracted 4-momentum to the sum of the charged tracks from the leading vertex. If it has a larger momentum but a smaller mass, we set the transverse momentum and azimuthal angle according to the result of the pileup subtraction and set the rapidity and mass as the one obtained from summing the charged tracks from the leading vertex. We shall study later in this chapter the effect of this “safe” subtraction on the reconstruction of the jet mass.

Note finally that, in `FastJet`, one can still use the `Subtractor` tool as for the standard usage. One can even use the `Subtractor` directly on full events using `Subtractor::set_known_selectors(...)` which takes two `Selectors` as arguments: the first one isolates particles which can be associated to a known vertex, while the second one selects the particles from the leading vertex among the particles with a known vertex association.

If we want to impose positivity constraints on the mass, we can still use `Subtractor::set_safe_mass`. With a previous call to `Subtractor::set_known_selectors(...)`, the more constraining conditions, involving information from the particles associated with the leading vertex, will be used (see Section 2.4.4).

Neutral-proportional-to-charged (NpC). An alternative approach to pileup subtraction in a context where we know for each charged track whether it comes from the leading vertex or from a pileup interaction is to deduce the neutral momentum of the jet based on its charged information. The NpC method 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 (untouched) as part of the jet during clustering, then the corrected jet momentum is [114]

$$p_{\mu}^{\text{jet,sub}} = p_{\mu}^{\text{jet}} - \frac{1}{\gamma_0} p_{\mu}^{\text{jet,chg-PU}}, \quad (11.1)$$

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

$$\gamma_0 \equiv \left\langle \frac{\sum_{i \in \text{charged particles}} p_{ti}}{\sum_{i \in \text{all particles}} p_{ti}} \right\rangle_{\text{events}}, \quad (11.2)$$

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. With the Pythia8 simulations we have used in this Chapter — see Section 11.2 below and Appendix E — we have $\gamma_0 \approx 0.612$.

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 needs to make sure that pileup vertices are still included as ghost tracks, with their momenta rescaled by an infinitesimal factor $\epsilon \lll 1$. In this case the correction becomes

$$p_{\mu}^{\text{jet,sub}} = p_{\mu}^{\text{jet,CHS}} - \frac{1 - \gamma_0}{\gamma_0 \epsilon} p_{\mu}^{\text{jet,rescaled-chg-PU}}, \quad (11.3)$$

where $p_{\mu}^{\text{jet,CHS}}$ is the momentum of the jet as obtained from the 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 back-reaction associated with the charged pileup (this is useful also for the area–median 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 pseudo-rapidities. Additionally there are subtleties with hadron masses: in effect, γ_0 is different for transverse components and for longitudinal components. In the

following we will avoid this problem by treating all particles as massless.¹ The importance of the above limitations can only be fully evaluated in an experimental context.

NpC also comes with a few potential (experimental) advantages compared to the area–median subtraction method. First, the area–median method includes questions of non-trivial rapidity dependence and detector non-linearities (the latter is relevant also for NpC). These have, to a reasonable extent, been successfully overcome by the experiments [115, 116]. One respect in which NpC may have advantages of the area–median method is that the latter fails to correctly account for the fact pileup fluctuates from point to point within the event, a feature that cannot be encoded within the global pileup estimate ρ . 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 v. 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.

Jet cleansing. Jet cleansing [10] essentially puts several ingredients together into a pileup subtraction technique. It first re-clusters the jet into subjets as done for filtering and trimming, see Sec. 6.2, (cf. also the early work by Seymour [117]). It then applies an NpC-like technique to each of the subjets, before summing the resulting 4-momenta into the final subtracted jet. It also applies an additional constraint that we decided to call *zeroing* and that we shall discuss separately later in this 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).²³

To be a little more precise, the implementation of (linear) Jet Cleansing, available in `fjcontrib`, slightly differs from our definition of NpC. According to the description in Ref. [10], one additional characteristic of linear cleansing relative to NpC is that it switches to jet-vertex-fraction (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. In our Monte-Carlo studies presented below in Section 11.2, 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.

Furthermore, Eq. (11.3) for NpC corrects for the jet 4-momentum, while cleansing scales the 4-momentum $p_\mu^{(i)}$ of each neutral particle i in the jet as follows

$$p_\mu^{(i)} \rightarrow p_\mu^{(i)} \times \left(1 - \frac{1 - \gamma_0 p_t^{\text{jet, rescaled-chg-PU}}}{\gamma_0 \epsilon} \frac{p_t^{\text{jet, ntr}}}{p_t^{\text{jet, ntr}}} \right), \quad (11.4)$$

where the notation is as in Eq. (11.3) and additionally $p_t^{\text{jet, ntr}}$ is the transverse momentum of the neutral part of the jet (including pileup). Eq. (11.4) is correct for the jet p_t , but would be only be correct for the jet mass if the pileup and hard components had identical angular distributions.

¹Particle momenta are modified so as to become massless while conserving p_t , rapidity and azimuth.

²In v1 of arXiv:1404.7353 [118] in April 2014, and also in a version circulated to the authors of Ref. [10] several months before the arXiv submission, we used $f_{\text{cut}} = 0.05$ for cleansing, reflecting our understanding of the choices made in v1 of Ref. [10], 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 v1 of our article, the authors of Ref. [10] clarified that the results in their Fig. 6 (in arXiv v2, Fig. 4 in arXiv v1) 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.

³They 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 [115, 119, 120].

However in practice, the hard component is usually collimated, whilst, to a first approximation, the pileup component is uniformly distributed across the jet. Accordingly, the scaling gives incorrect jet mass results. The issue is less severe when cleansing is used as originally intended, i.e. including trimming.

A last small difference is that cleansing assumes that all the particles in the event are clustered, in a spirit similar to Eq. (11.1). As already mentioned in our description of NpC, this can bring an additional source of back-reaction compared to Eq. (11.3), where the charged tracks from pileup vertices are only included as ghosts.

Note finally that besides the “linear” version of cleansing presented above, Ref. [10] also introduces a variant called *Gaussian cleansing* which is particularly interesting in that it effectively carries out a χ^2 minimisation across different hypotheses for the ratio of charged to neutral energy flow, separately for the pileup and the hard event. We shall see below that its performance is usually marginally better than the much simpler linear cleansing.

Charged-tracks-based trimming, zeroing and protected zeroing. In the context of comparison studies between the area–median approach, NpC and Jet Cleansing (see below), we studied the public code for jet cleansing.⁴ This revealed an additional condition being applied to subjets: if a subjet contains no charged particles from the leading vertex (LV), then its momentum is set to zero.⁵ Since we will be discussing it extensively, it is 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 [120], whereby a JVF-style cut is applied to reject subjets whose charged-momentum fraction from the leading vertex is too low. As we will show in the Monte-Carlo studies presented in Section 11.2 zeroing turns out to be crucial in explaining the differences observed between CHS area-subtraction (or with NpC subtraction) with $f_{\text{cut}} = 0$ trimming and cleansing.

This approach may seem a little dangerous since there are rare cases where a hard jet — or a subjet carrying a large fraction of the jet p_t — can be purely neutral and zeroing would simply discard it. This can correspond to the loss of subjets with tens of GeV, while 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). We have found good results, and improvements compared to plain zeroing, taking $n = 10$ (as we shall use later in our Monte-Carlo studies) or a fixed $p_{t,\text{min}} = 20$ GeV.

Before starting an in-depth comparison of the methods described above, we also want to point out that similarly to the combination of a neutral-proportional-to-charged method with trimming, done by jet cleansing, several other similar approaches have been proposed and studied in the past in the context of the area–median approach[121]. We have already discussed in Chapter 4 (see also Ref.[5]) and in Chapter 7 (see also [20]) the possibility to combine the area–median subtraction together with grooming techniques and this combination can trivially be extended to the case of CHS events. The preliminary studies presented in Chapter 10 also relate to similar ideas.

⁴Version 1.0.1 from [fjcontrib](#).

⁵This crucial step appears not to have been mentioned in Ref. [10].

11.2 Monte-Carlo studies and physical discussions

Since our testing framework remains close to the one we have used for the studies in Chapter 3, modulo a few adaptations meant to be closer to the setup used in Ref. [10], we will directly concentrate on the results of our studies and their physical interpretation. We refer to Appendix E for a more detailed description of our simulations.

Since the main goal of this Chapter is to discuss the use of charged tracks in pileup mitigation, the first and most important point we shall discuss is a comparison between the area–median method and the NpC approach. It will show that the former gives a slightly better performance than the latter and we will explain why this could be expected. We shall then discuss extensively the comparison of our findings with the, apparently contradictory, findings from jet cleansing in Ref. [10].

11.2.1 Performance of NpC v. area–median

Let us therefore proceed with an investigation of NpC’s performance compared to the now-familiar area–median method, 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 centred at $y = \phi = 0$ and examine the correlation coefficient 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, ρ_{ntr} . Fig. 11.1 (left) shows these two correlation coefficients, “ntr v. chg” and “ntr v. $\rho_{\text{ntr}}A$ ”, 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} .

As already pointed out before, quantitative interpretations of correlation coefficients can sometimes be delicate (see also Appendix C). We instead find that it is more robust to investigate $\sigma_{\Delta p_t^{\text{ntr}}}$, defined (similarly to our previous studies in this document) as the standard deviation of

$$\Delta p_t^{\text{ntr}} = p_t^{\text{ntr}} - p_t^{\text{ntr,estimated}}, \quad (11.5)$$

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 plot of Fig. 11.1 shows $\sigma_{\Delta p_t^{\text{ntr}}}$ for the two methods, again as a function of R , for two levels of pileup. It is normalised 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 figure 11.1 suggests that one will be disappointed: the local estimate of the neutral component of pileup from NpC is marginally worse than that given by the global ρ_{ntr} from the area–median method.

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 non-perturbative effects. Inside an energetic jet, the non-perturbative 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 flow mostly tend to go in the same direction.

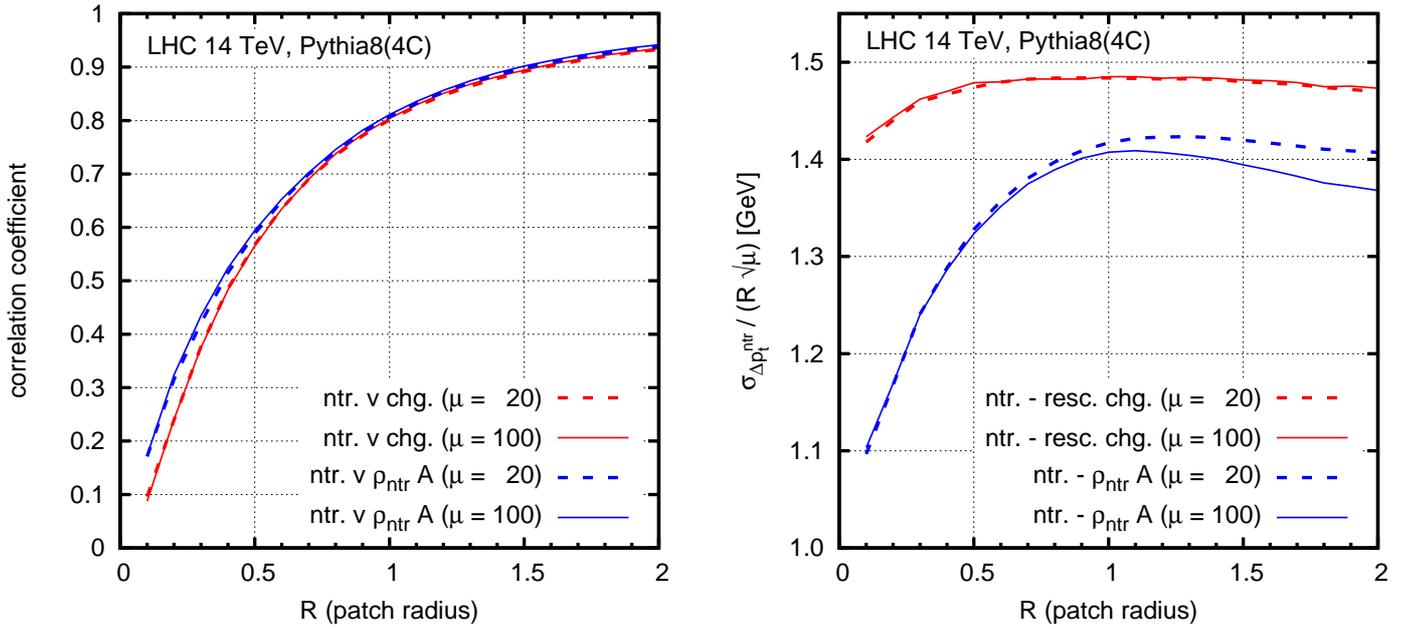


Figure 11.1: Left: the correlation coefficient between the neutral transverse momentum in a central patch and either the charged transverse momentum in that patch (rescaling this component would not change the correlation coefficient) or the prediction using the area–median method, i.e. $\rho_{\text{ntr}}A$. Right: 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 Pythia 8, tune 4C, and the number of collisions per event is Poisson distributed with average μ .

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 hadronisation, 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 , non-perturbative effects specifically tend to wash out the charged-neutral angular correlation.

This point is illustrated in Fig. 11.2. We consider zero-bias events and examine a circular patch of radius $R = 0.4$ centred at $y = \phi = 0$. The figure shows the distribution of the charged p_t fraction, r ,

$$r = \frac{p_t^{\text{chg}}}{p_t^{\text{chg+ntr}}}, \quad (11.6)$$

in the patch (filled histogram, broken into contributions where the patch contains 1, 2 or more particles). The same plot 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 about 60% of events the patch is either just charged or just neutral, quite often consisting of just a single particle. This is probably part of the

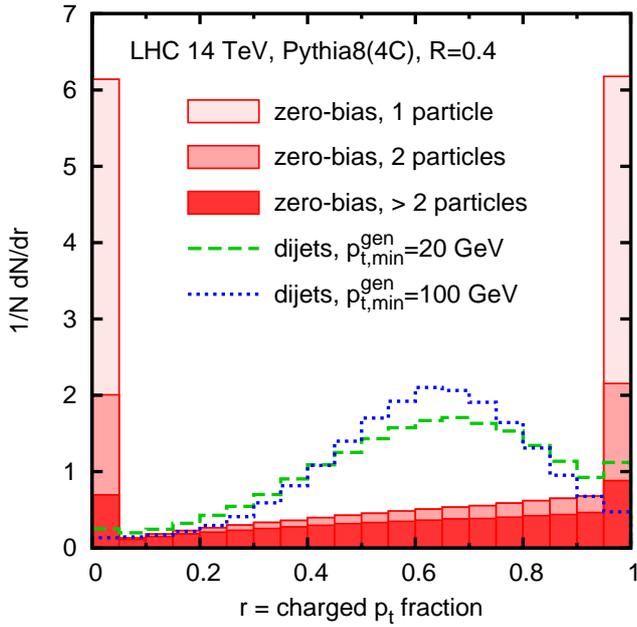


Figure 11.2: The filled histogram shows the distribution, for simulated zero-bias collisions, of the fraction, r , of the transverse momentum in a central circular patch of radius $R = 0.4$ that is due to charged particles. It is separated into components according to the multiplicity of particles in 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 $p_{t,\min}^{\text{gen}}$.

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, as implemented in FastJet⁶, with a jet radius parameter of $R = 0.4$. We study dijet and pileup events generated with Pythia 8.176, in tune 4C; we assume idealised 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 event⁷ and subtract them using either NpC, Eq. (11.3), or the area–median method, Eq. (2.24), with ρ estimated from the CHS event. The hard events are generated with the underlying event turned off so as to avoid subtleties related to the simultaneous subtraction of the Underlying Event.

Figure 11.3 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 plot 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 plot shows the standard deviation of the difference between the hard and 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

⁶A development snapshot of version 3.1, equivalent to the present v3.2 series for the purpose at hand.

⁷For the matching, we introduce a quantity $p_t^{\text{shared}}(j_i^{\text{hard}}, j_j^{\text{full}})$, the scalar sum 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 back-reaction effects. However this is exceedingly rare, since we consider only the two hardest jets in the hard event, which are nearly always back-to-back.

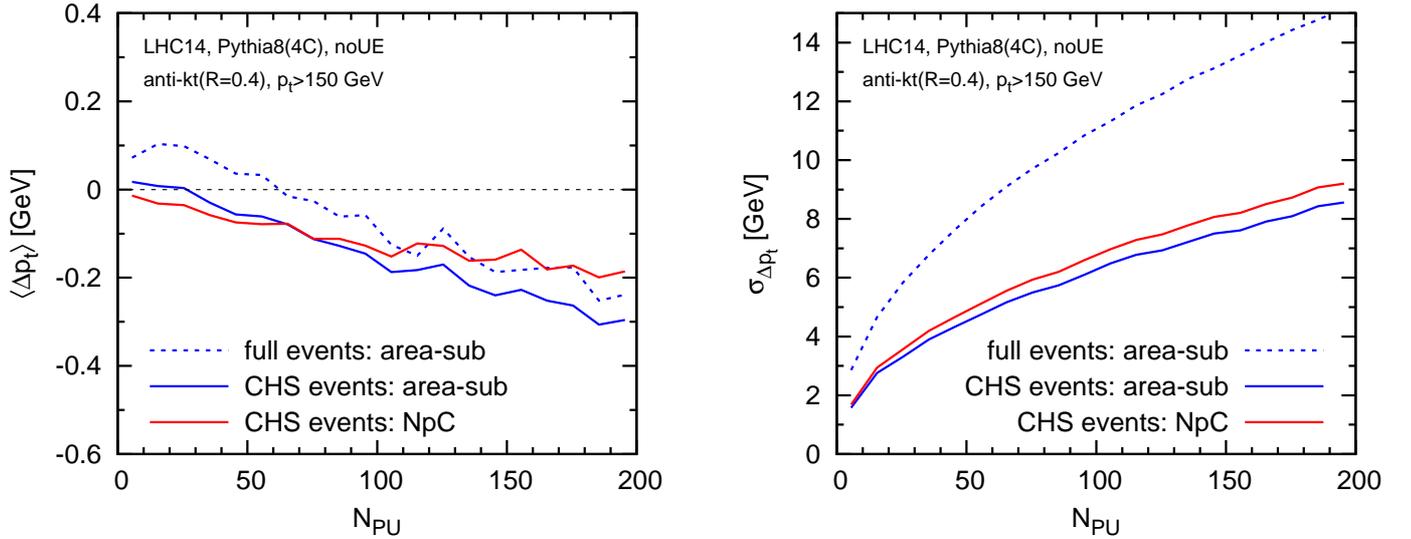


Figure 11.3: A comparison of the performance of the NpC and area–median subtraction methods. The left-hand plot shows, as a function of the number of pileup vertices N_{PU} , the average difference in p_t between a jet after pileup addition and subtraction and the corresponding matched jet in the hard sample, $\Delta p_t = p_t^{\text{jet,sub}} - p_t^{\text{jet,hard}}$. The right-hand plot 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”).

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 hadronisation 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.

11.2.2 Performance of cleansing v. area–median

We now turn to a broader discussion, comparing our discussion about the NpC and area–median approaches with the results found using Jet Cleansing.

The top left-hand plot of Fig. 11.4 shows the dispersion $\Delta m_{jj} = m_{jj}^{\text{sub}} - m_{jj}^{\text{hard}}$ for the dijet mass for several pileup mitigation methods.⁸ 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 plots 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 plot, for the jet mass, we see linear (and Gaussian) cleansing performing better than area subtraction, albeit not to the extent found in Ref. [10]. These (and, unless explicitly stated, our other Z' results) have been generated with the Z' decaying to all flavours except $t\bar{t}$, and B -hadrons have been kept

⁸Our original comparison [9] as well as Ref.[10] showed instead correlation coefficients between the hard and subtracted results. For more coherence through this document, we prefer to show results for the dispersion. The ordering between the different methods and physical interpretation of the results is unchanged.

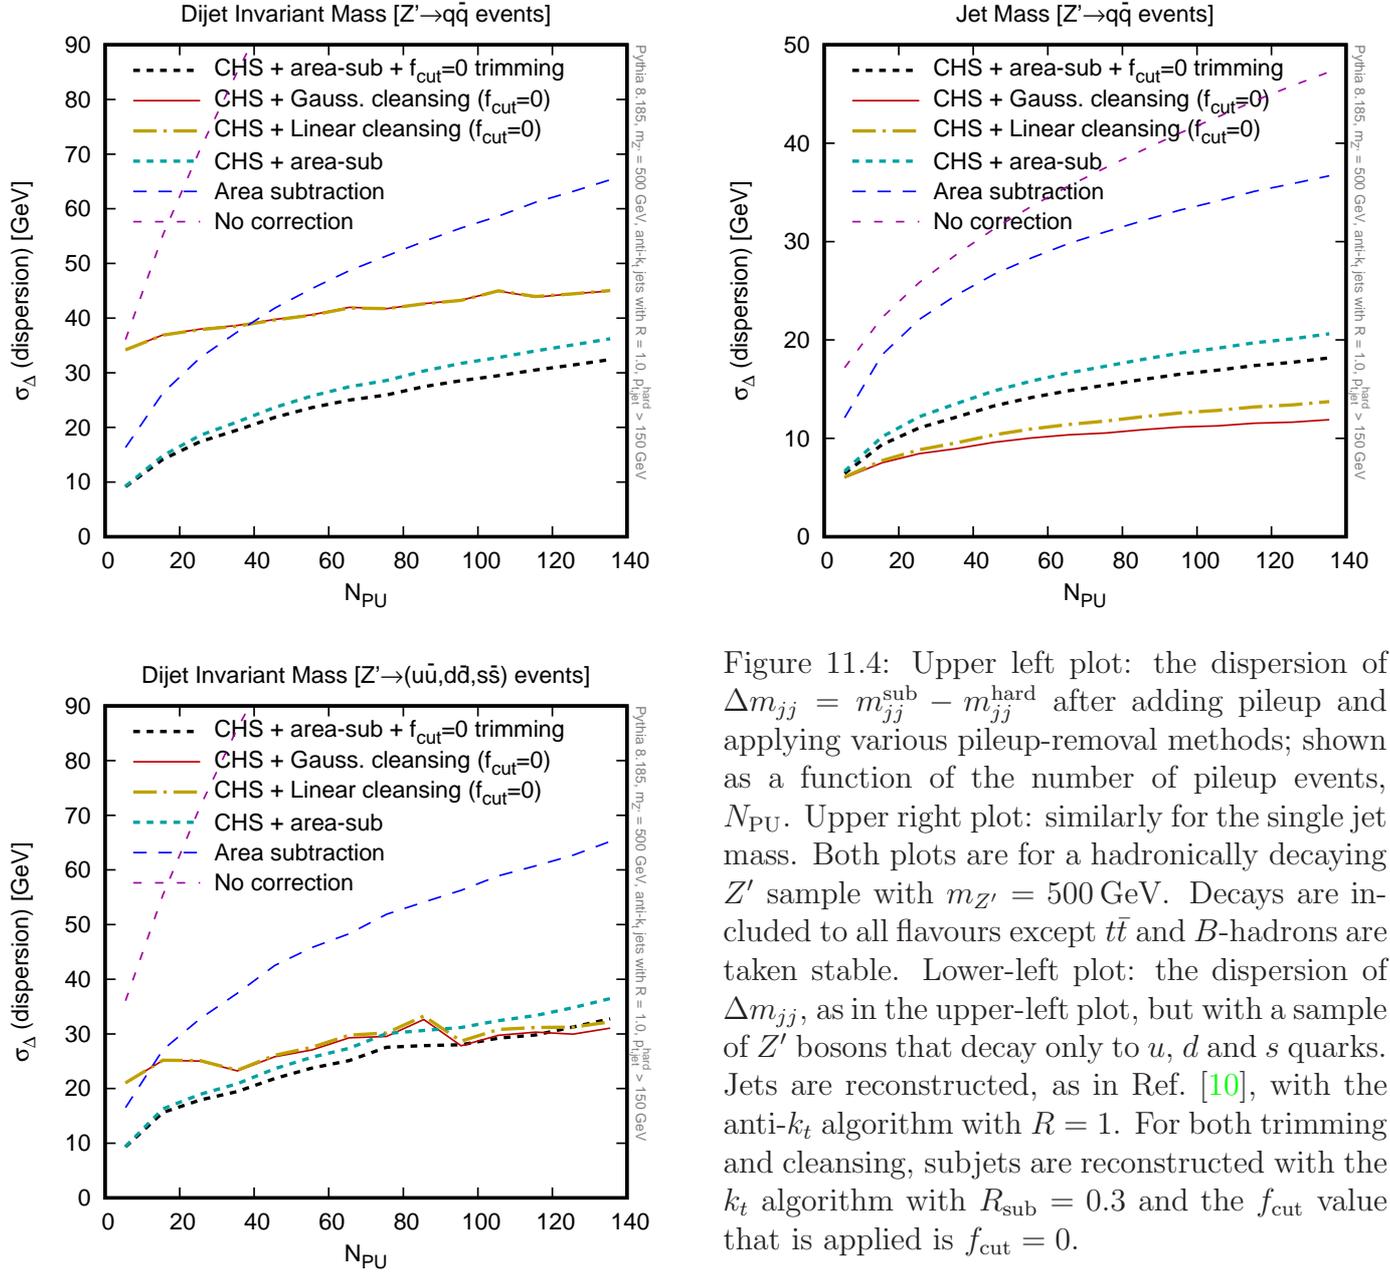


Figure 11.4: Upper left plot: the dispersion of $\Delta m_{jj} = m_{jj}^{\text{sub}} - m_{jj}^{\text{hard}}$ after adding pileup and applying various pileup-removal methods; shown as a function of the number of pileup events, N_{PU} . Upper right plot: similarly for the single jet mass. Both plots are for a hadronically decaying Z' sample with $m_{Z'} = 500$ GeV. Decays are included to all flavours except $t\bar{t}$ and B -hadrons are taken stable. Lower-left plot: the dispersion of Δm_{jj} , as in the upper-left plot, but with a sample of Z' bosons that decay only to u , d and s quarks. Jets are reconstructed, as in Ref. [10], with the anti- k_t algorithm with $R = 1$. For both trimming and cleansing, subjets are reconstructed with the k_t algorithm with $R_{\text{sub}} = 0.3$ and the f_{cut} value that is applied is $f_{\text{cut}} = 0$.

stable.⁹ The lower plot 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 plots in our Fig. 11.4 differ noticeably from the conclusion drawn from Fig. 6 (top) of Ref. [10] (arXiv v2) and in particular they are not consistent with KLSW's observation of much improved performance for the dijet mass with cleansing relative to area+CHS subtraction.¹⁰

In light of our results on NpC in section 11.2.1, the difference between the performance of area-subtraction plus trimming versus that of cleansing are puzzling: our expectation is that their performances should be similar.¹¹ The strong sample-dependence of the cleansing performance also calls for an explanation. We thus pursue with an in-depth study of the question.

As mentioned in our description of cleansing in Section 11.1, one additional characteristic of linear cleansing relative to area-subtraction is that it switches to jet-vertex-fraction (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 observed in Fig. 11.4.

The next difference that we will discuss at length is the zeroing step that is found when inspecting the details of the public code for jet cleansing.¹² 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. 11.5 shows a ‘‘fingerprint’’ for each of several pileup-removal methods, for both the jet p_t (left) and mass (right). 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 the 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

⁹We often find this to be useful for particle-level b -tagging studies. Experimentally, in the future, one might even imagine an ‘‘idealised’’ form of particle flow that attempts to reconstruct B -hadrons (or at least their charged part) from displaced tracks before jet clustering.

¹⁰We remain puzzled also by the relative pattern of area+CHS v. area-subtracted results in Fig. 6 (top) of Ref. [10], since the area+CHS curves appears to tend towards the area curves at large pileup, whereas the use of CHS should significantly reduce the impact of pileup.

¹¹KLSW state in [10] that fluctuations around a ‘best’ charged fraction $\bar{\gamma}_0$ decrease with increasing N_{PU} and suggest (see also [122], 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 $\sqrt{N_{\text{PU}}}$.

¹²Version 1.0.1 from fjcontrib [123]. This step appears not to have been mentioned in Ref. [10].

(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 fingerprints, but again NpC, area and cleansing are very similar.¹³

When used with $f_{\text{cut}} = 0$ trimming, and when examining quality measures such as the dispersion (in red, or the closely related correlation coefficient, cf. Appendix C), subjet 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 subjets 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. 11.5. On the other hand, zeroing sometimes incorrectly rejects subjets that have no charged tracks from the LV but do have significant neutral energy flow from the LV. This can lead to long tails for the Δp_t or Δm distributions, adversely affecting the dispersion.¹⁴ 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 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. 11.5 shows that for cleansing without zeroing, the sample dependence (dashed versus solid lines) vanishes, while it is substantial with zeroing. 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 can be avoided by using the alternative *protected zeroing* described in Section 11.1. 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. 11.5. Protected zeroing also eliminates the sample dependence.¹⁵

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. 11.5. 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

¹³One exception is that for the jet mass, Gaussian cleansing does differ in the $f_{\text{cut}} = 0$ case with zeroing, and shows an advantage from its combinations of different constraints on subjet momenta. Unfortunately, as we shall see this later, this advantage does not seem to carry over to jet masses with $f_{\text{cut}} \neq 0$ trimming, which are phenomenologically more relevant than the full jet mass.

¹⁴A discrepancy between dispersion and peak-width measures is to be seen in Fig. 15 of Ref. [124] for jet masses. Our “fingerprint” plot is in part inspired by the representation provided there, though our choice of peak-width measure differs.

¹⁵We learned while finalising this work that for the identification of pileup v. non-pileup 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 exchanges on this point.

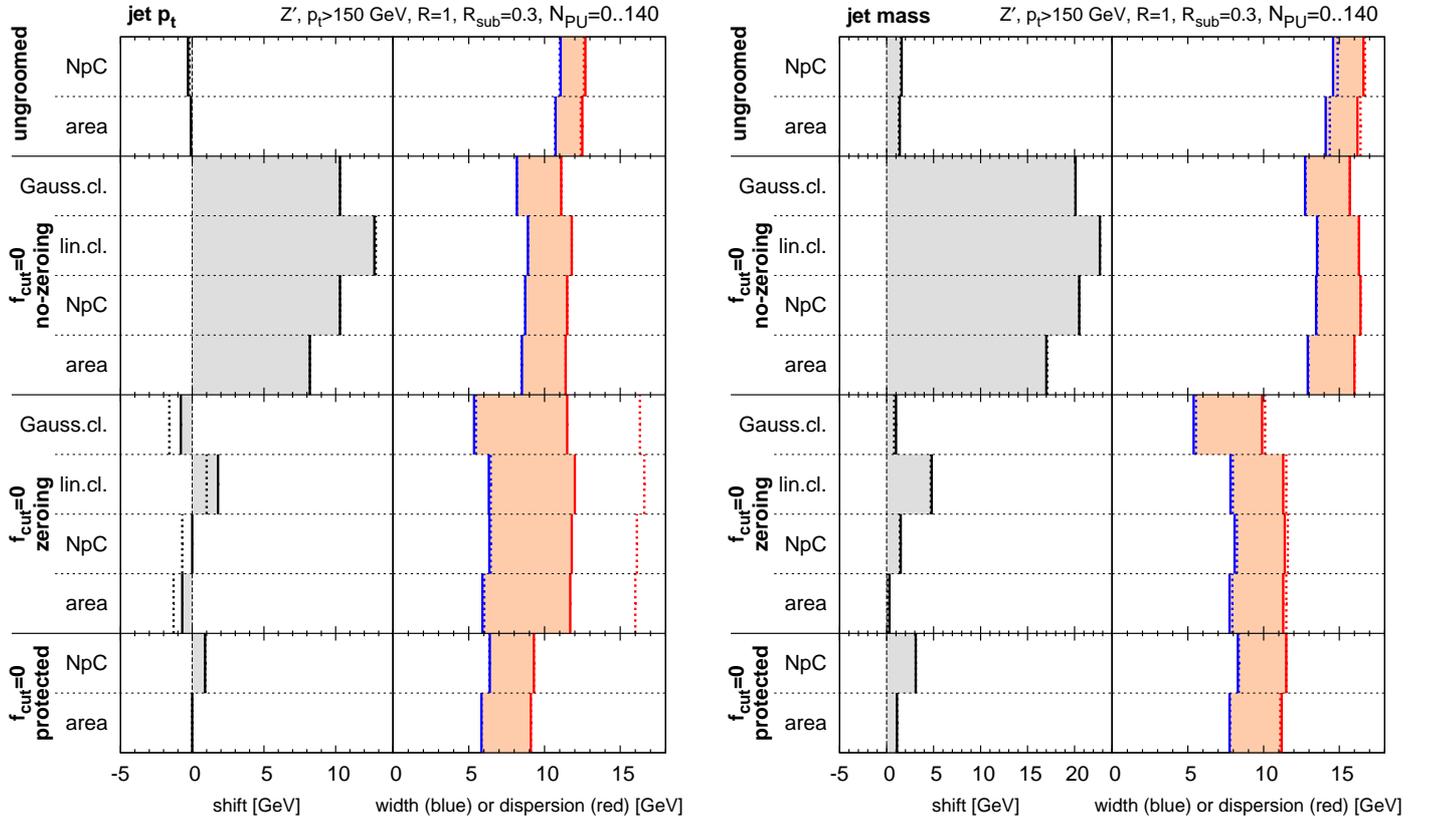


Figure 11.5: The left-hand plot illustrates various characteristics of the change (Δp_t) in the jet p_t after addition of pileup and removal by a range of methods. It shows the average shift $\langle \Delta p_t \rangle$ (in black) and the peak width (in blue) and dispersion (in red) of the Δp_t distribution. The peak width is defined as the smallest window of Δp_t that contains 90% of the Δp_t distribution, scaled by a factor $\simeq 0.304$ such that in the case of a Gaussian distribution the result agrees with the dispersion. The right-hand plot 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 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_t 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. 11.4, is “ $f_{\text{cut}} = 0$ zeroing”.

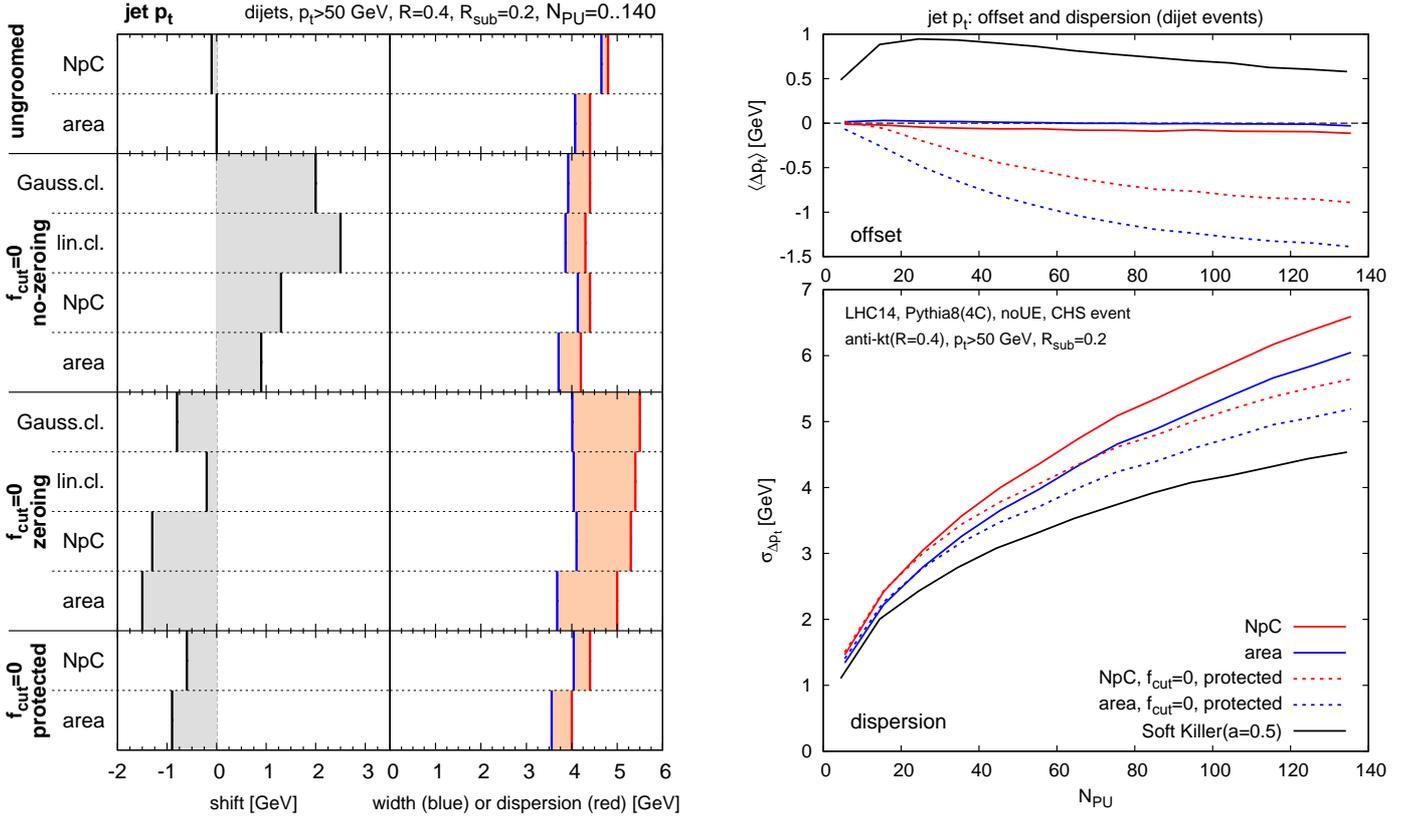


Figure 11.6: Left, Analogue of Fig. 11.5 (left) for a jet radius of $R = 0.4$, subjet radius (where relevant) of $R_{\text{sub}} = 0.2$ and a QCD continuum dijet sample generated with Pythia 8. The underlying event is turned off in the sample and B hadrons decay. We consider only jets that in the hard sample have $p_t > 50$ GeV and $|y| < 2.5$. Right: the dispersions for a subset of the methods, shown as a function of the number of pileup events.

fluctuations, those that happened not to have any charged tracks from the LV. It also introduces further negative fluctuations for subjets that happened to have some neutral energy flow but no charged tracks. Overall, one sees that the final net bias comes out to be relatively small. This kind of cancellation between different biases is common in noise-reducing pileup-reduction approaches [125, 11, 12].

Most of the studies so far in this section have been carried out with a setup that is similar to that of Ref. [10], 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 Fig. 11.6 (left) we show the analogue of Fig. 11.5'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. 11.5, with the area–median approach performing slightly better than NpC-based techniques as expected from our discussion in Section 11.2.1.¹⁶ Quantitatively, the difference between the various choices is much smaller, with about a 10% reduction in dispersion (or width) in going from ungroomed 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 plot of Fig. 11.6.

¹⁶Note, 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.

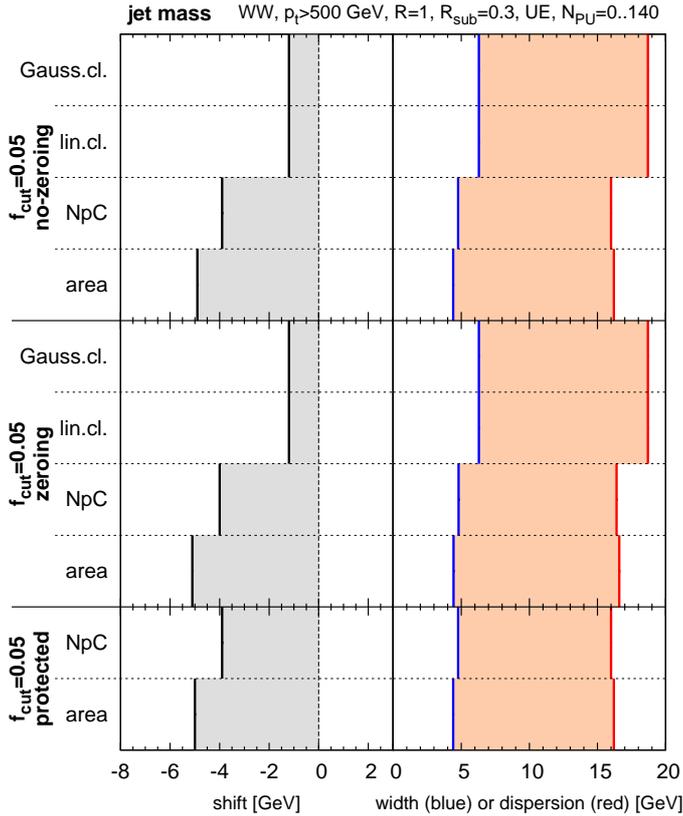


Figure 11.7: Analogue of Fig. 11.5 (right), showing the performance for the jet mass, but now with $f_{\text{cut}} = 0.05$ applied to both trimming and cleansing and in a sample of hadronically-decaying boosted W bosons ($pp \rightarrow W^+W^-$). The jets reconstructed after addition and subtraction 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 B hadrons decay.

That plot also includes results from the SoftKiller method discussed in Chapter 12 (see also [11]) 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 plots 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 behaviour across a range of transverse momenta, determine optimal choices for the protection of the zeroing and investigate also how best to combine it with particle-level subtraction methods such as SoftKiller.¹⁷ We will come back to this combination in Section 13.2.

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 non-zero 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 [107, 110] even in the absence of pileup.¹⁸ Accordingly in Fig. 11.7 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

¹⁷An interesting feature of protected zeroing, SoftKiller and another recently introduced method, PUPPI [12], 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.

¹⁸In contrast, for $f_{\text{cut}} = 0$ trimming, the jet structure is unchanged in the absence of pileup.

summary plot, protection of zeroing appears to have little impact for the trimmed jet mass.¹⁹ 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 19). In that context, it would also be interesting to consider the combination of zeroing with the grooming techniques suggested as pileup mitigation methods in Chapter 10.

Further systematic comparisons of the different methods, with and without zeroing can be found in the Appendices of Ref. [9].

We close this section with a summary of our findings. Based on its description in Ref. [10] and our findings about NpC v. 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. Ref. [10] 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 behaviour 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. [10] 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. [120]. 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 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 non-zero 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.

11.3 Digression: combining different methods

It is interesting to further probe the relation between NpC and the area–median method and 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

¹⁹Cleansing appears to perform slightly worse than trimming with NpC or area subtraction. One difference in behaviour 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.

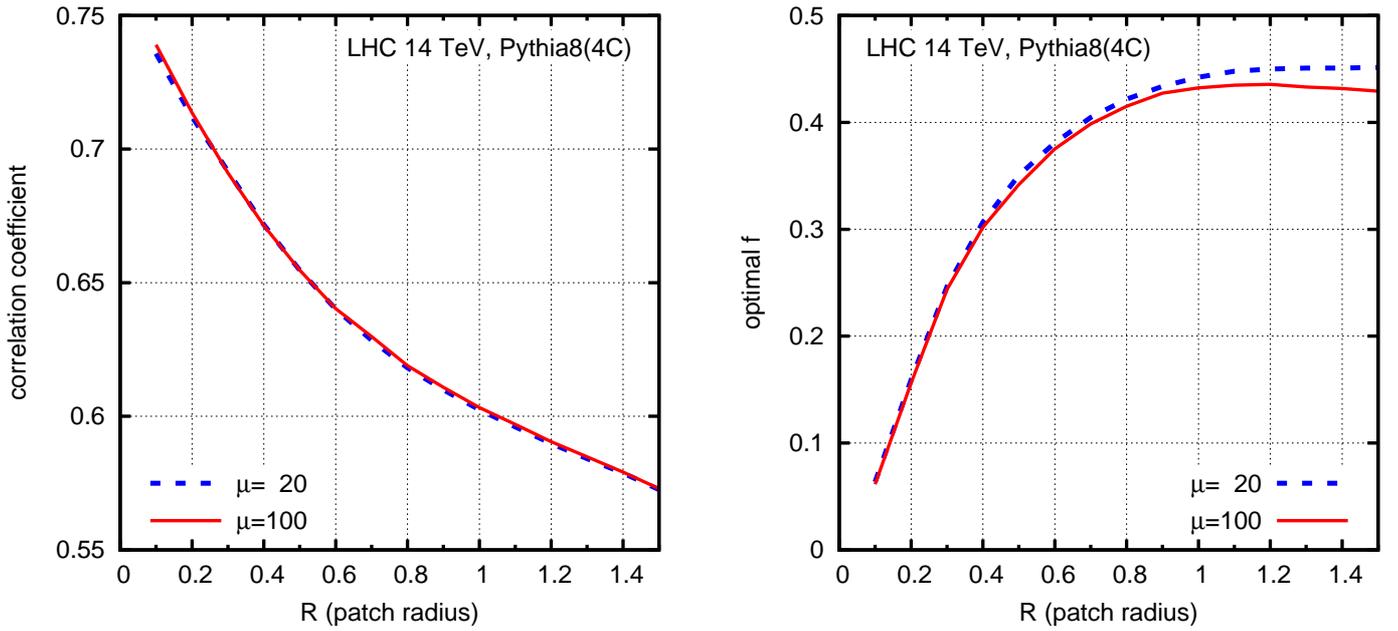


Figure 11.8: Left: correlation between $p_t^{\text{nr}} - \hat{p}_\mu^{\text{nr(NpC)}}$ and $p_t^{\text{nr}} - \hat{p}_\mu^{\text{nr}(\rho A)}$, shown as a function of R . Right: optimal weight f for combining NpC and area pileup subtraction, Eq. (11.9), as a function of R .

performance.

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

$$\hat{p}_\mu^{\text{nr(NpC)}} = \frac{1 - \gamma_0}{\gamma_0 \epsilon} p_\mu^{\text{jet, rescaled-chg-PU}}, \quad \hat{p}_\mu^{\text{nr}(\rho A)} = \rho A, \quad (11.7)$$

are, respectively, the estimates for the neutral pileup based on the local charged p_t flow and on ρA . Note that we for our analysis in this section we use CHS events and in particular ρ is as determined from the CHS event, i.e. $\rho \equiv \rho_{\text{CHS}}$. This means that it is essentially equal to ρ_{nr} , 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_{nr} instead of ρ_{CHS} to estimate $\hat{p}_\mu^{\text{nr}(\rho A)}$, and this would in some respects be a more coherent approach, though the numerical differences will be small.

Concentrating on the transverse components, the extent to which the two estimates in Eq. (11.7) provide complementary information can be quantified in terms of (one minus) the correlation coefficient, r , between $p_t^{\text{nr}} - \hat{p}_\mu^{\text{nr(NpC)}}$ and $p_t^{\text{nr}} - \hat{p}_\mu^{\text{nr}(\rho A)}$. That correlation is shown as a function of R in Fig. 11.8 (left), and it is quite high, in the range 0.6–0.7 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

$$\hat{p}_\mu^{\text{nr}} = f \hat{p}_\mu^{\text{nr(NpC)}} + (1 - f) \hat{p}_\mu^{\text{nr}(\rho A)}, \quad (11.8)$$

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

$$f = \frac{\sigma_{\rho A}^2 - r \sigma_{\text{NpC}} \sigma_{\rho A}}{\sigma_{\text{NpC}}^2 + \sigma_{\rho A}^2 - 2r \sigma_{\text{NpC}} \sigma_{\rho A}}, \quad (11.9)$$

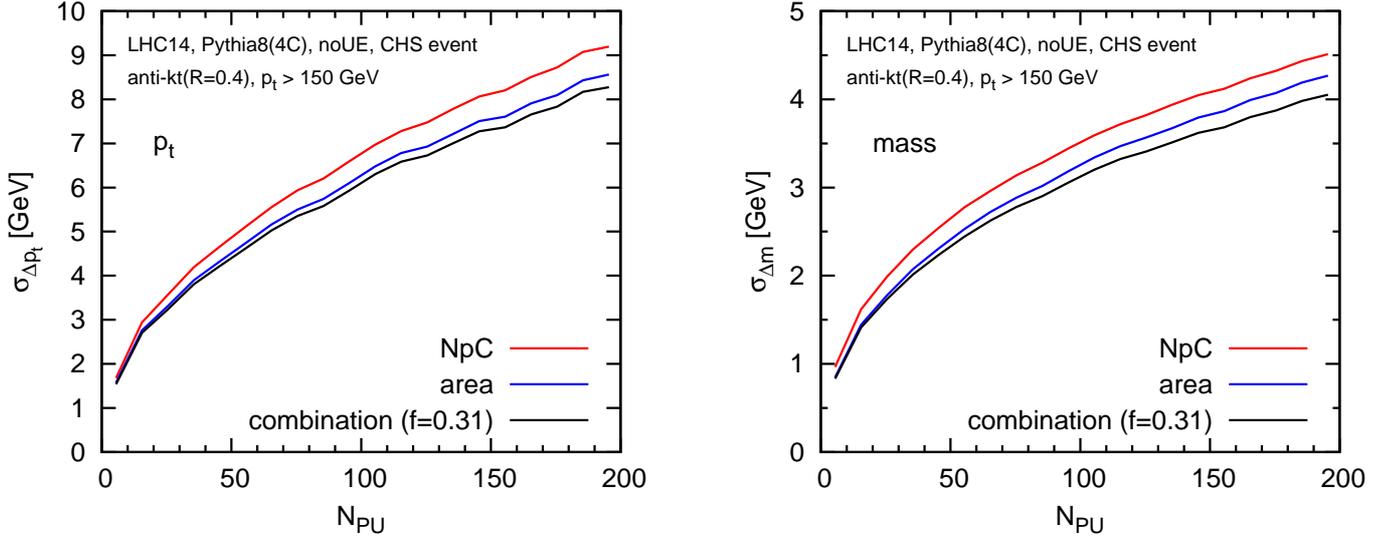


Figure 11.9: Comparison of the performance of NpC, area–median and combined subtraction, as a function of the number of pileup vertices. The left-hand plot is for the jet p_t and the right-hand one for the jet mass.

which is plotted as a function of R in Fig. 11.8 (right). The resulting squared dispersion for $p_t^{\text{tr}} - \hat{p}_t^{\text{tr}}$ is

$$\sigma^2 = \frac{(1 - r^2) \sigma_{\text{NpC}}^2 \sigma_{\rho A}^2}{\sigma_{\text{NpC}}^2 + \sigma_{\rho A}^2 - 2r \sigma_{\text{NpC}} \sigma_{\rho A}}. \quad (11.10)$$

Reading $r = 0.67$ from Fig. 11.8 (left) for $R = 0.4$, and $\sigma_{\text{NpC}} \simeq 1.14 \sigma_{\rho A}$ from Fig. 11.1 (right), one finds $f \simeq 0.31$ and $\sigma \simeq 0.96 \sigma_{\rho A}$. Because of the substantial correlation between the two methods, one expects only a modest gain from their linear combination.

In Fig. 11.9 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. 11.8 (right) 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.²⁰ 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.

Note finally that the approach described above can be applied to combine a different pair of pileup mitigation techniques. It can also easily be extended to more than two methods.

11.4 Concluding remarks and discussions

One natural approach to pileup subtraction 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 chapter, 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

²⁰We 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. (11.9). However the dependence on f around its minimum is very weak, rendering the details of its exact choice somewhat immaterial.

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 non-perturbative effects that characterise typical inelastic proton-proton collisions act to destroy local charged-neutral correlation.

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. [10], 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. [10], 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 behaviour of cleansing and area subtraction (or pure NpC) to the use in the cleansing code of a step that was not documented in Ref. [10] and that discards subjects that contain no tracks from the leading vertex. This “zeroing” step, similar to the charged-track based trimming introduced by ATLAS [120], can indeed be of benefit. It has a drawback of introducing tails in some distributions due to subjects with a substantial neutral p_t from the leading vertex, but no charged tracks. As a result, different quality measures (and different event samples) lead to different conclusions as to the benefits of zeroing. The tails can be alleviated by a variant of zeroing that we introduce in Ref. [9], “protected zeroing”, whereby subjects 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.

Code for our implementation of area subtraction with positive-definite mass is available as part of FastJet versions 3.1.0 and higher (see Section 2.7 and, in particular, 2.7.2.4). Public code and samples for carrying out a subset of the comparisons with cleansing described in section 11.2.2, including also the NpC subtraction tools, are available from Ref. [126].

Chapter 12

Noise reduction with the SoftKiller

In this Chapter, we describe a new method that we introduced in Ref. [11] and that we called SoftKiller (SK). This method acts as a noise-reduction method which has shown nice resolution improvements compared to the area–median method in Monte-Carlo simulations.

We will first describe the method itself in Section 12.1 followed by some validation in Monte-Carlo studies (Section 12.2) as well as possible adaptations to specific situations like CHS events and detector effects (Section 12.3).

One of the interesting features of the SoftKiller method is that it shows remarkable performance in terms of computing time. We will discuss that in Section 12.4.

The work presented in this Chapter is a minor adaptation of the original SoftKiller paper [11]. We will discuss in the next Chapter, devoted to preliminary ideas and results, some possible ways to further improve the SoftKiller method.

12.1 The SoftKiller method

The SoftKiller method involves eliminating particles below some p_t cutoff, p_t^{cut} , chosen to be the minimal value that ensures that ρ is zero. Here, ρ is the event-wide estimate of transverse-momentum flow density in the area–median approach, i.e. given by Eq. (2.23). In practice we will use a grid-base estimate for ρ .¹

Choosing the minimal transverse momentum threshold, p_t^{cut} , that results in $\rho = 0$ is equivalent to gradually raising the p_t threshold until exactly half of the patches contain no particles, which ensures that the median is zero. This is illustrated in Fig. 12.1. Computationally, p_t^{cut} is straightforward to evaluate: one determines, for each patch i , the p_t of the hardest particle in that patch, p_{ti}^{max} and then p_t^{cut} is given by the median of p_{ti}^{max} values:

$$p_t^{\text{cut}} = \text{median}_{i \in \text{patches}} \{p_{ti}^{\text{max}}\} . \quad (12.1)$$

With this choice, half the patches will contain only particles that have $p_t < p_t^{\text{cut}}$. These patches will be empty after application of the p_t threshold, leading to a zero result for ρ as defined in Eq. (2.23).² The computational time to evaluate p_t^{cut} as in Eq. (12.1) scales linearly in the number of particles and the method should be amenable to parallel implementation.

¹Note also that the estimation of ρ does not use any rapidity rescaling. Such a rapidity rescaling would actually have no effect on the determination of the SoftKiller p_t cut. This is further discussed at the end of Section 12.2.

²Applying a p_t threshold to individual particles is not collinear safe; in the specific context of pileup removal, we believe that this is not a significant issue, as we discuss in more detail in Appendix F.

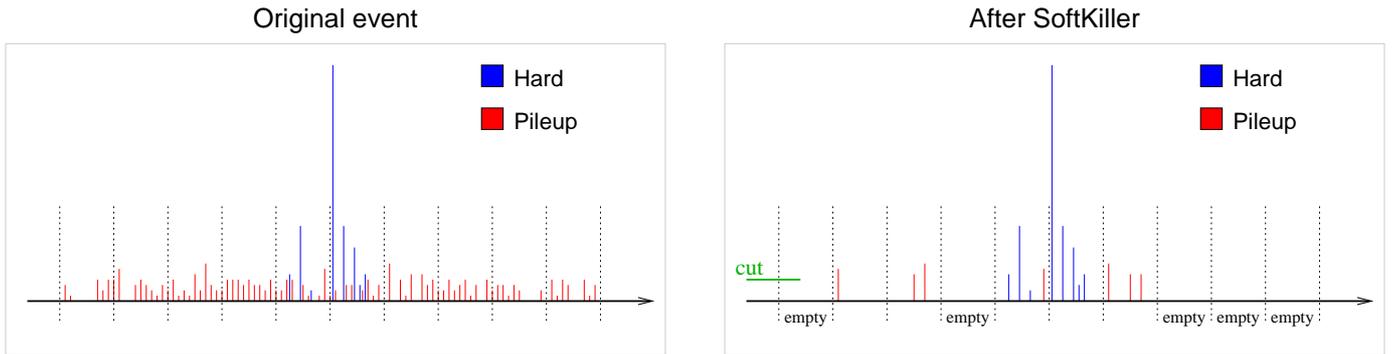


Figure 12.1: Illustration of the SoftKiller method. The left plot depicts particles in an event, with the hard event particles shown in blue and the pileup particles shown in red. On the right, the same event after applying the SoftKiller. The vertical dotted lines represent the edges of the patches used to estimate the pileup density ρ .

Imposing a cut on particles' transverse momenta eliminates most of the pileup particles, and so might reduce the fluctuations in residual pileup contamination from one point to the next within the event. However, as with other event-wide noise-reducing pileup and underlying-event mitigation approaches, notably the CMS heavy-ion method [125] (cf. the analysis in Appendix A.4 of Ref. [60]; see also our introductory discussion in Chapter 9), the price that one pays for noise reduction is the introduction of biases. Specifically, some particles from pileup will be above p_t^{cut} and so remain to contaminate the jets, inducing a net positive bias in the jet momenta. Furthermore some particles in genuine hard jets will be lost, because they are below the p_t^{cut} , inducing a negative bias in the jet momenta. The jet energy scale will only be correctly reproduced if these two kinds of bias are of similar size,³ so that they largely cancel. There will be an improvement in the jet resolution if the fluctuations in these biases are modest.

Figure 12.2 shows, on the left, the average p_t^{cut} value, together with its standard deviation (dashed lines), as a function of the number of pileup interactions, N_{PU} . The event sample consists of a superposition of N_{PU} zero-bias on one hard dijet event, in 14 TeV proton–proton collisions, all simulated with Pythia 8 (tune 4C) [36].⁴ The underlying event in the hard event has been switched off, and all particles have been made massless, maintaining their p_t , rapidity and azimuth.⁵ These are our default choices throughout this Chapter. The grid used to determine p_t^{cut} has a spacing of $a \simeq 0.4$ and extends up to $|y| < 5$. One sees that p_t^{cut} remains moderate, below 2 GeV, even for pileup at the level foreseen for the high-luminosity upgrade of the LHC (HL-LHC), which is expected to reach an average (Poisson-distributed) number of pileup interactions of $\mu \simeq 140$. The right-hand plot shows the two sources of bias: the lower (solid) curves, illustrate the bias on the hard jets induced by the loss of genuine hard-event particles below p_t^{cut} . Jet clustering is performed as usual with the anti- k_t jet algorithm with $R = 0.4$ (as implemented in FastJet 3.1). The three line colours correspond to

³For patch areas that are similar to the typical jet area, this can be expected to happen because half the patches will contain residual pileup of order p_t^{cut} , and since jets tend to have only a few low- p_t particles from the hard scatter, the loss will also be order of p_t^{cut} .

⁴For a study of the tune dependence, see Appendix C of Ref. [11] where we briefly examined the Pythia 6 [35] Z2 tune [127], and find very similar results.

⁵If one keeps the underlying event in the hard event, much of it (about 1 GeV for both the area–median approach and the SoftKiller) is subtracted together with the pileup correction, affecting slightly the observed shifts. Keeping massive particles does not affect the SK performance but requires an extra correction for the area–median subtraction, see Section 2.4.3. We therefore use massless particles for simplicity.

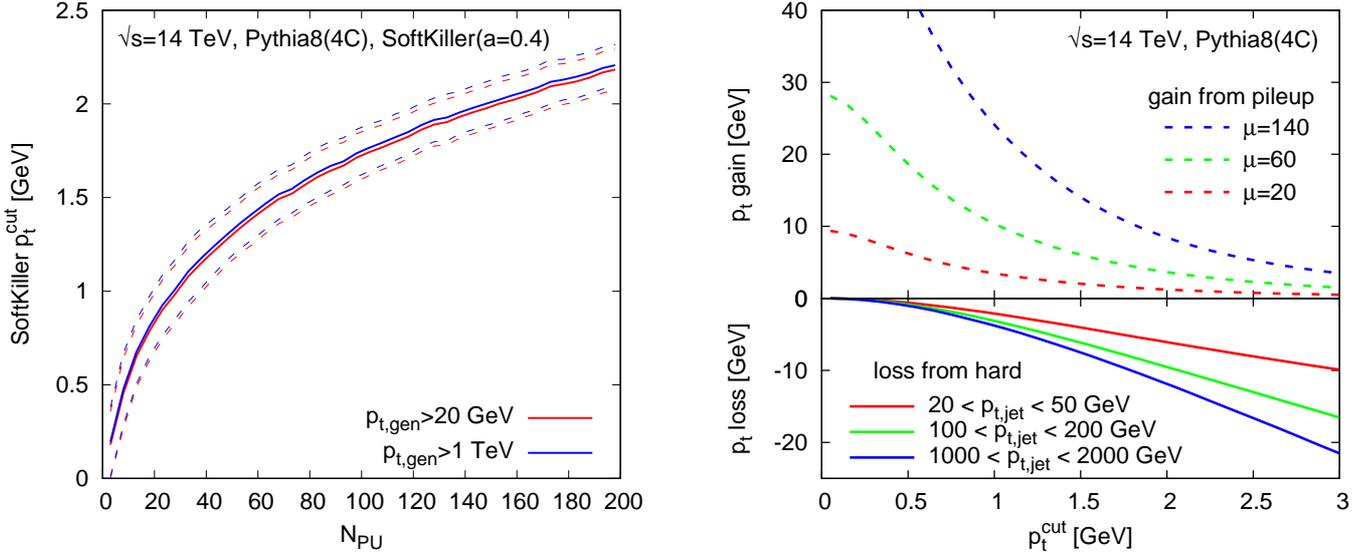


Figure 12.2: Left: value of the p_t cut applied by the SoftKiller, displayed as a function of the number of pileup events. We show results for two different values of the generator minimal p_t for the hard event, $p_{t,\text{gen}}$. The solid line is the average p_t^{cut} value while the dashed lines indicate the one-standard-deviation band. Right: plot of the p_t that is lost when applying a given p_t cut (the x axis) to the constituents of jets clustered (anti- k_t , $R = 0.4$) from the hard event (solid lines) and the residual pileup p_t that remains after applying that same cut to the constituents of circular patches of radius 0.4 in pure-pileup events (dashed lines).

different jet p_t ranges. The loss has some dependence on the jet p_t itself, notably for higher values of p_t^{cut} .⁶ In particular it grows in absolute terms for larger jet p_t 's, though it decreases relative to the jet p_t . The positive bias from residual pileup particles (calculated in circular patches of radius 0.4 at rapidity $y = 0$) is shown as dashed curves, for three different pileup levels. To estimate the net bias, one should choose a value for N_{PU} , read the average p_t^{cut} from the left-hand plot, and for that p_t^{cut} compare the solid curve with the dashed curve that corresponds to the given N_{PU} . Performing this exercise reveals that there is indeed a reasonable degree of cancellation between the positive and negative biases. Based on this observation, we can move forward with a more detailed study of the performance of the method.⁷

12.2 SoftKiller performance

Choice of grid size. For a detailed study of the SoftKiller method, the first step is to choose the grid spacing a so as to break the event into patches. The spacing a is the one main free parameter of the method. The patch-size parameter is present also for area–median pileup subtraction. There the exact choice of this parameter is not too critical because the median is quite stable when pileup levels are high: all grid cells are filled, and nearly all are dominated by pileup. However the SoftKiller method chooses the p_t^{cut} so as to obtain a nearly empty event. In this limit, the median operation becomes somewhat more sensitive to the grid spacing (see also the discussion in Section 5.2).

⁶In a local parton-hadron duality type approach to calculate hadron spectra, the spectrum of very low p_t particles in a jet of a given flavour is actually independent of the jet's p_t [128].

⁷A study of fixed p_t cutoffs, rather than dynamically determined ones, is performed in Appendix D of Ref.[11].

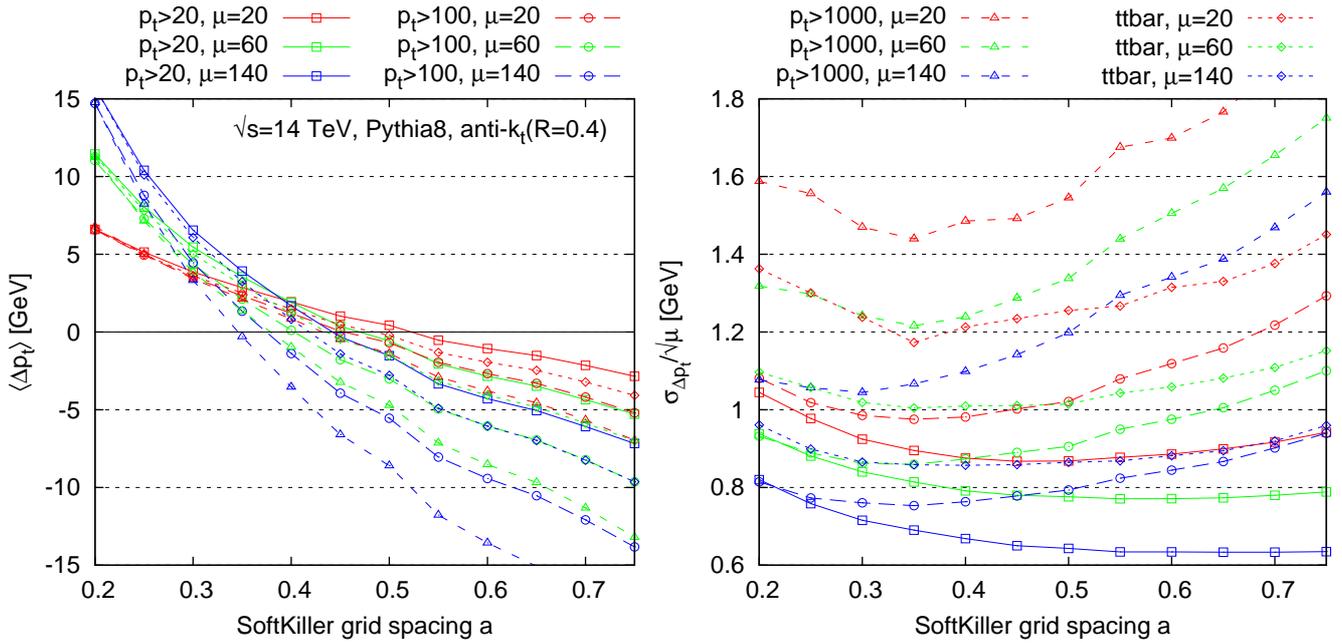


Figure 12.3: Scan of the SoftKiller performances as a function of the grid-spacing parameter a for different hard-event samples and three different pileup levels (Poisson-distributed with average pileup event multiplicities of $\mu = 20, 60, 140$). Left: average p_t shift; right: p_t shift dispersion, normalised to $\sqrt{\mu}$ for better readability. Results are given for a variety of hard processes and pileup conditions to illustrate robustness. Curves labelled $p_t > X$ correspond to dijet events, in which one studies only those jets that in the hard event have a transverse momentum greater than X . For the $t\bar{t}$ sample, restricted to fully hadronic decays of the top quarks, the study is limited to jets that have $p_t > 50$ GeV in the hard event.

Fig. 12.3 considers a range of hard event samples (different line styles) and pileup levels (different colours). For each, as a function of the grid spacing a , the left-hand plot shows the average, $\langle \Delta p_t \rangle$, of the net shift in the jet transverse momentum, defined as before⁸

$$\Delta p_t = p_t^{\text{corrected}} - p_t^{\text{hard}}, \quad (12.2)$$

while the right-hand plot shows the dispersion, $\sigma_{\Delta p_t}$, of that shift from one jet to the next, here normalised to $\sqrt{\mu}$ (right).

One sees that the average jet p_t shift has significant dependence on the grid spacing a . However, there exists a grid spacing, in this case $a \simeq 0.4$, for which the shift is not too far from zero and not too dependent either on the hard sample choice or on the level of pileup. In most cases the absolute value of the shift is within about 2 GeV, the only exception being for the $p_t > 1000$ GeV dijet sample, for which the bias can reach up to 4 GeV for $\mu = 140$. This shift is, however, still less than the typical best experimental systematic error on the jet energy scale, today of the order of 1% or slightly better [129, 130].

It is not trivial that there should be a single grid spacing that is effective across all samples and pileup levels: the fact that there is can be considered phenomenologically fortuitous. The value of the

⁸The matching between a hard jet and a “corrected” jet is done essentially as before: for each hard jet passing our selection cut, we select the jet in the subtracted event (i.e. after applying the SoftKiller) for which the sum of the transverse momentum of the constituents it has in common with the original hard jet is maximal.

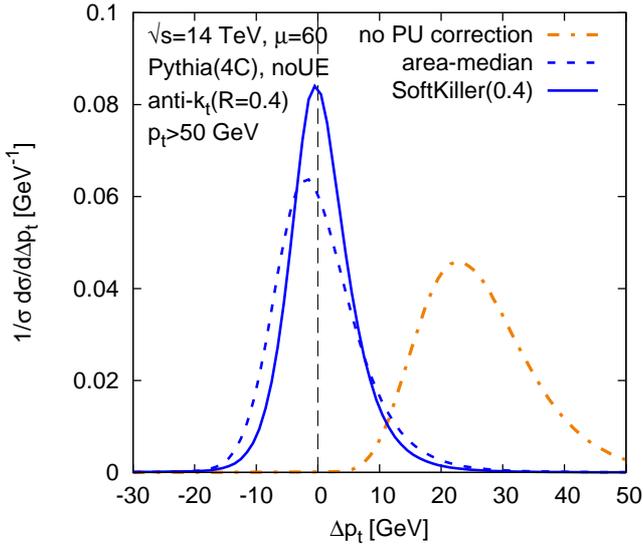


Figure 12.4: Performance of SoftKiller for 50 GeV jets and $\mu = 60$ Poisson-distributed pileup events. We plot the distribution of the shift Δp_t between the jet p_t after pileup removal and the jet p_t in the hard event alone. Results are compared between the area-median approach and the SoftKiller. For comparison, the (orange) dash-dotted line corresponds to the situation where no pileup corrections are applied.

grid spacing a that minimises the typical shifts is also close to the value that minimises the dispersion in the shifts.⁹ That optimal value of a isn't identical across event samples, and can also depend on the level of pileup. However the dispersion at $a = 0.4$ is always close to the actual minimal attainable dispersion for a given sample. Accordingly, for most of the rest of this article, we will work with a grid spacing of $a = 0.4$.¹⁰

Performance compared to the area–median method. We now compare the performance of the SoftKiller to that of area–median subtraction. Figure 12.4 shows the distribution of shift in p_t , for (hard) jets with $p_t > 50$ GeV in a dijet sample. The average number of pileup events is $\mu = 60$, with a Poisson distribution. One sees that in the SoftKiller approach, the peak is about 30% higher than what is obtained with the area–median approach and the distribution correspondingly narrower. The peak, in this specific case, is well centred on $\Delta p_t = 0$.

Figure 12.5 shows the shift (left) and dispersion (right) as a function of N_{PU} for two different samples: the $p_t > 50$ GeV dijet sample (in blue), as used in Fig. 12.4, and a hadronic $t\bar{t}$ sample, with a 50 GeV p_t cut on jets (in green). Again, the figure compares the area–median (dashed) and SoftKiller results (solid). One immediately sees that the area–median approach gives a bias that is more stable as a function of N_{PU} . Nevertheless, the bias in the SoftKiller approach remains between about -0.5 and 1.5 GeV, which is still reasonable when one considers that, experimentally, some degree of recalibration is anyway needed after area–median subtraction. As concerns sample dependence of the shift, comparing $t\bar{t}$ v. dijet, the area–median and SoftKiller methods appear to have similar systematic differences. In the case of SoftKiller, there are two main causes for the sample dependence: firstly the higher multiplicity of jets has a small effect on the choice of p_t^{cut} and secondly the dijet sample is mostly composed of gluon-induced jets, whereas the $t\bar{t}$ sample is mostly composed of quark-induced jets (which have fewer soft particles and so lose less momentum when imposing a particle p_t threshold). Turning to the right-hand plot, with the dispersions, one sees that the SoftKiller brings a significant

⁹In a context where the net shift is the sum of two opposite-sign sources of bias, this is perhaps not too surprising: the two contributions to the dispersion are each likely to be of the same order of magnitude as the individual biases, and their sum probably minimised when neither bias is too large.

¹⁰A single value of a is adequate as long as jet finding is carried out mostly with jets of radius $R \simeq 0.4$. Later in this section we will supplement our $R = 0.4$ studies with a discussion of larger jet radii.

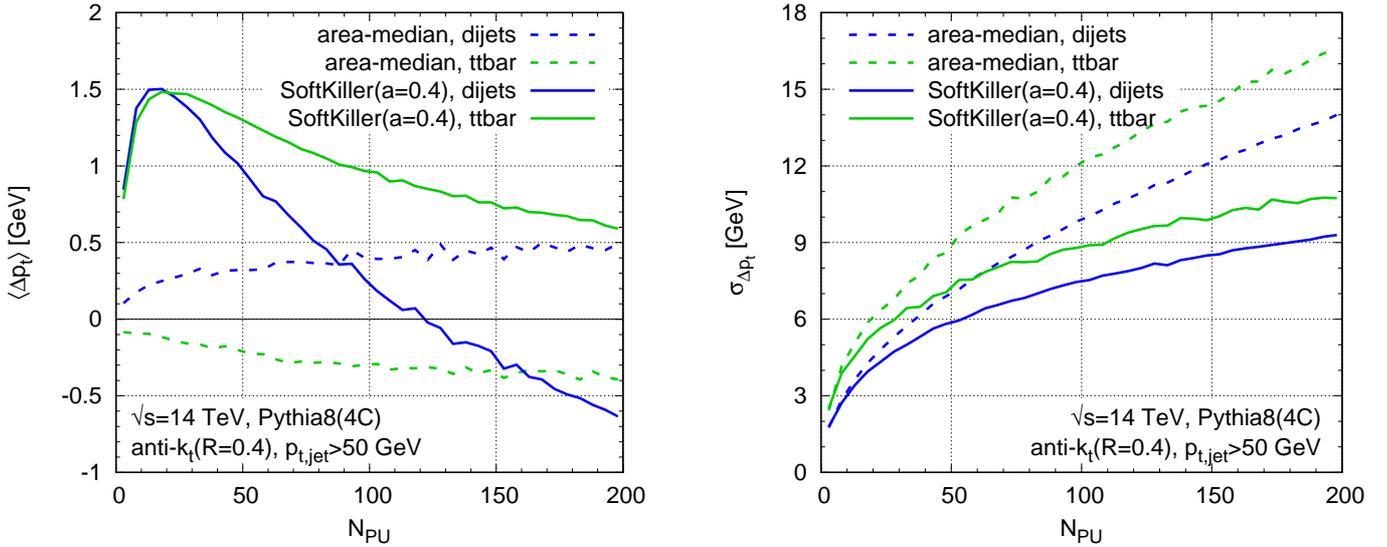


Figure 12.5: Performance of SoftKiller shown as a function of the pileup multiplicity and compared to the area-median subtraction method. Left: the average p_t shift after subtraction, compared to the original jets in the hard event. Right: the corresponding dispersion.

improvement compared to area-median subtraction for $N_{\text{PU}} \gtrsim 20$. The relative improvement is greatest at high pileup levels, where there is a reduction in dispersion of 30 – 35%, beating the $\sqrt{N_{\text{PU}}}$ scaling that is characteristic of the area-median method. While the actual values of the dispersion depend a little on the sample, the benefit of the SoftKiller approach is clearly visible for both.

Figure 12.6 shows the shift (left) and dispersion (right) for jet p_t 's and jet masses, now as a function of the hard jet minimum p_t . Again, dashed curves correspond to area-median subtraction, while solid ones correspond to the SoftKiller results. All curves correspond to an average of 60 pileup interactions. For the jet p_t (blue curves) one sees that the area-median shift ranges from 0.5 to 0 GeV as p_t increases from 20 GeV to 1 TeV, while for SK the dependence is stronger, from about 2 to -1 GeV, but still reasonable. For the jet mass (green curves), again the area-median method¹¹ is more stable than SK, but overall the biases are under control, at the 1 to 2 GeV level. Considering the dispersions (right), one sees that SK gives a systematic improvement, across the whole range of jet p_t 's. In relative terms, the improvement is somewhat larger for the jet mass ($\sim 30\%$) than for the jet p_t ($\sim 20\%$).

Fig. 12.7 shows the actual mass spectra of $R = 0.4$ jets, for two samples: a QCD dijet sample and a boosted $t\bar{t}$ sample. For both samples, we only considered jets with $p_t > 500$ GeV in the hard event. One sees that SK gives slightly improved mass peaks relative to the area-median method and also avoids area-median's spurious peak at $m = 0$, which is due to events in which the squared jet mass came out negative after four-vector area-subtraction and so was reset to zero. The plot also shows results from the recently proposed Constituent Subtractor method [131], using v. 1.0.0 of the corresponding code from FastJet Contrib [123]. It too performs better than area-median subtraction for the jet mass, though the improvement is not quite as large as for SK.¹²

¹¹using a “safe” subtraction procedure that replaces negative-mass jets with zero-mass jets, as described in Section 2.4.4.

¹²A further option is to use an “intrajet killer” that removes soft particles inside a given jet until a total p_t of ρA_{jet} has been subtracted. This shows performance similar to that of the Constituent Subtractor.

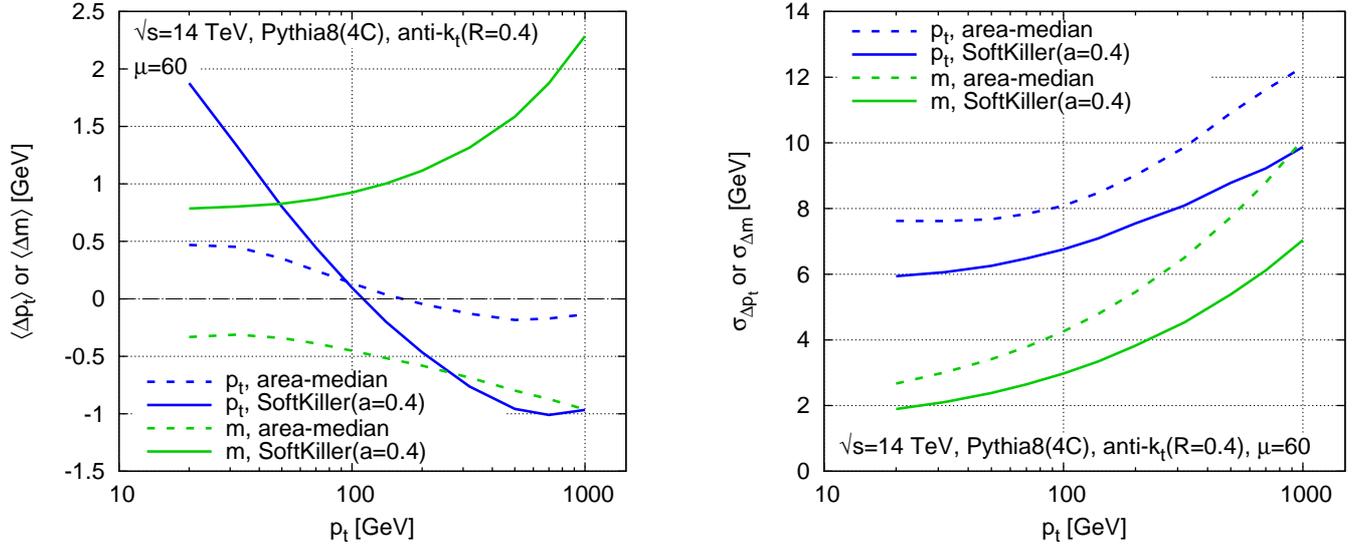


Figure 12.6: Performance of SoftKiller shown as a function of the hard jet p_t (solid lines), with a comparison to the area-median subtraction method (dashed lines). Left: the average shifts for the jet p_t (blue) and mass (green), after subtraction. Right: the corresponding dispersions.

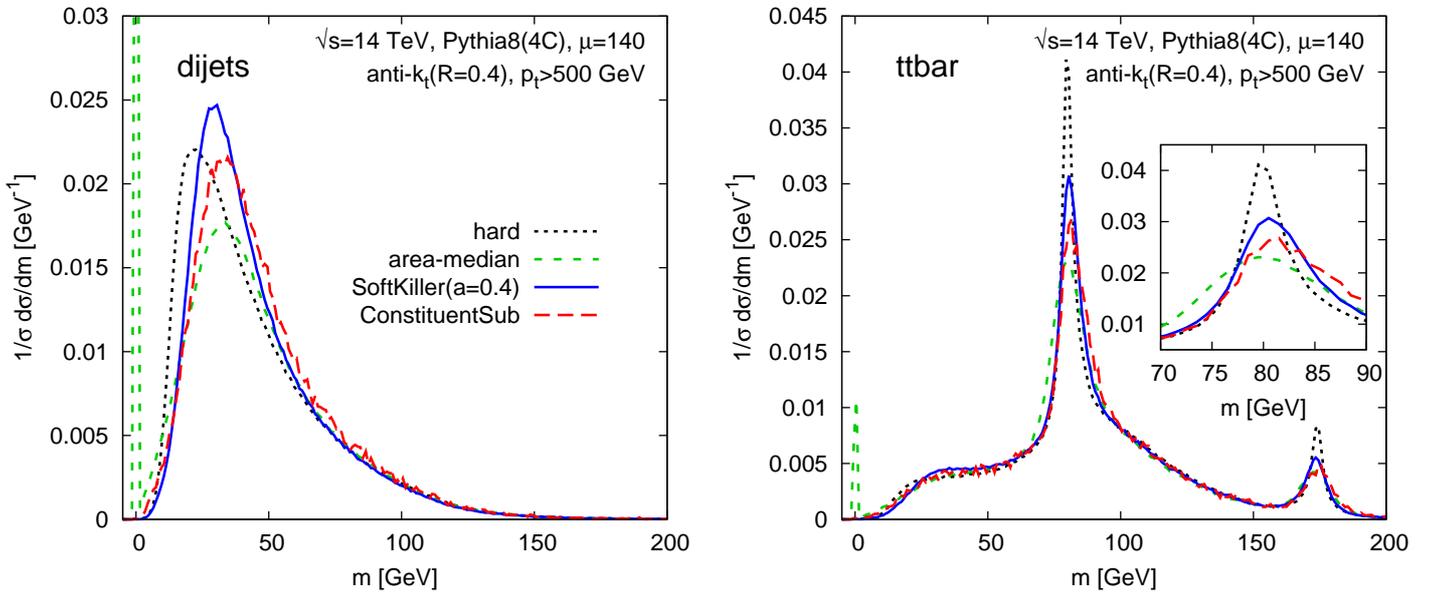


Figure 12.7: Performance of SoftKiller on jet mass reconstruction. We show the jet mass distribution after applying different pileup mitigation techniques. Left: dijet events, right: hadronically-decaying $t\bar{t}$ events.

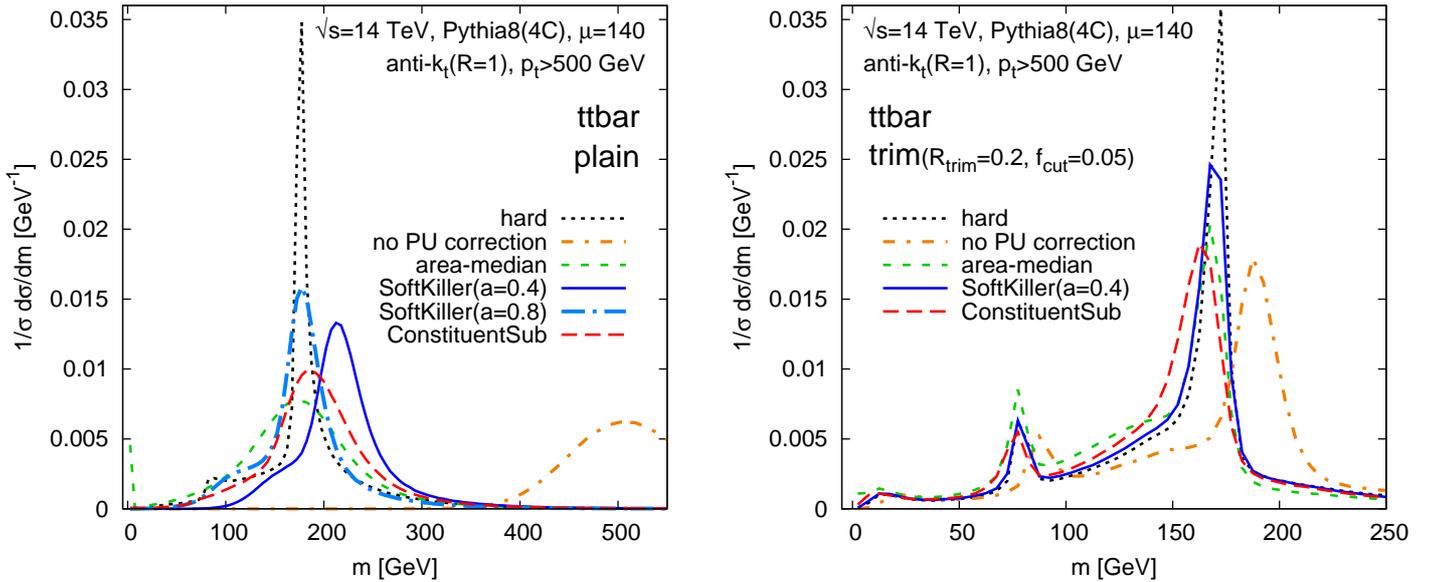


Figure 12.8: The results of the SoftKiller and other methods applied to large- R jets. The left-hand plot is without grooming, the right-hand plot with grooming (trimming with $R_{\text{trim}} = 0.2$ and $f_{\text{cut}} = 0.05$).

Large- R jets and grooming. One might ask why we concentrated on $R = 0.4$ jets here, given that jet-mass studies often use large- R jets. The reason is that large- R jets are nearly always used in conjunction with some form of grooming, for example trimming, pruning or filtering (see also Chapter 6). Grooming reduces the large-radius jet to a collection of small-radius jets and so the large-radius groomed-jet mass is effectively a combination of the p_t 's and masses of one or more small-radius jets.

For the sake of completeness, let us briefly also study the SoftKiller performance for large- R jets. Figure 12.8 shows jet-mass results for the same $t\bar{t}$ sample as in Fig. 12.7 (right), now clustered with the anti- k_t algorithm with $R = 1$. The left-hand plot is without grooming: one sees that SK with our default spacing of $a = 0.4$ gives a jet mass that has better resolution than area-median subtraction (or the ConstituentSubtractor), but a noticeable shift, albeit one that is small compared to the effect of uncorrected pileup. That shift is associated with some residual contamination from pileup particles: in an $R = 0.4$ jet, there are typically a handful of particles left from pileup, which compensate low- p_t particles lost from near the core of the jet. If one substantially increases the jet radius without applying grooming, then that balance is upset, with substantially more pileup entering the jet, while there is only slight further loss of genuine jet p_t . To some extent this can be addressed by using the SoftKiller with a larger grid spacing (cf. the $a = 0.8$ result), which effectively increases the particle p_t^{cut} . This comes at the expense of performance on small- R jets (cf. Fig. 12.3). An interesting, open problem is to find a simple way to remove pileup from an event such that, for a single configuration of the pileup removal procedure, one simultaneously obtains good performance on small- R and large- R jets. We will come back to candidate solutions to this open question in Chapter 13 (Section 13.2.1).

As we said above, however, large- R jet masses are nearly always used in conjunction with some form of grooming. Fig. 12.8 (right) shows that when used together with trimming, SoftKiller with our default $a = 0.4$ choice performs well both in terms of resolution and shift.

Jet shapes. Returning to $R = 0.4$ jets, the final figure of this section, Fig. 12.9, shows average shifts (left) and dispersions (right) as a function of N_{PU} for several different jet “shapes”: jet masses, k_t

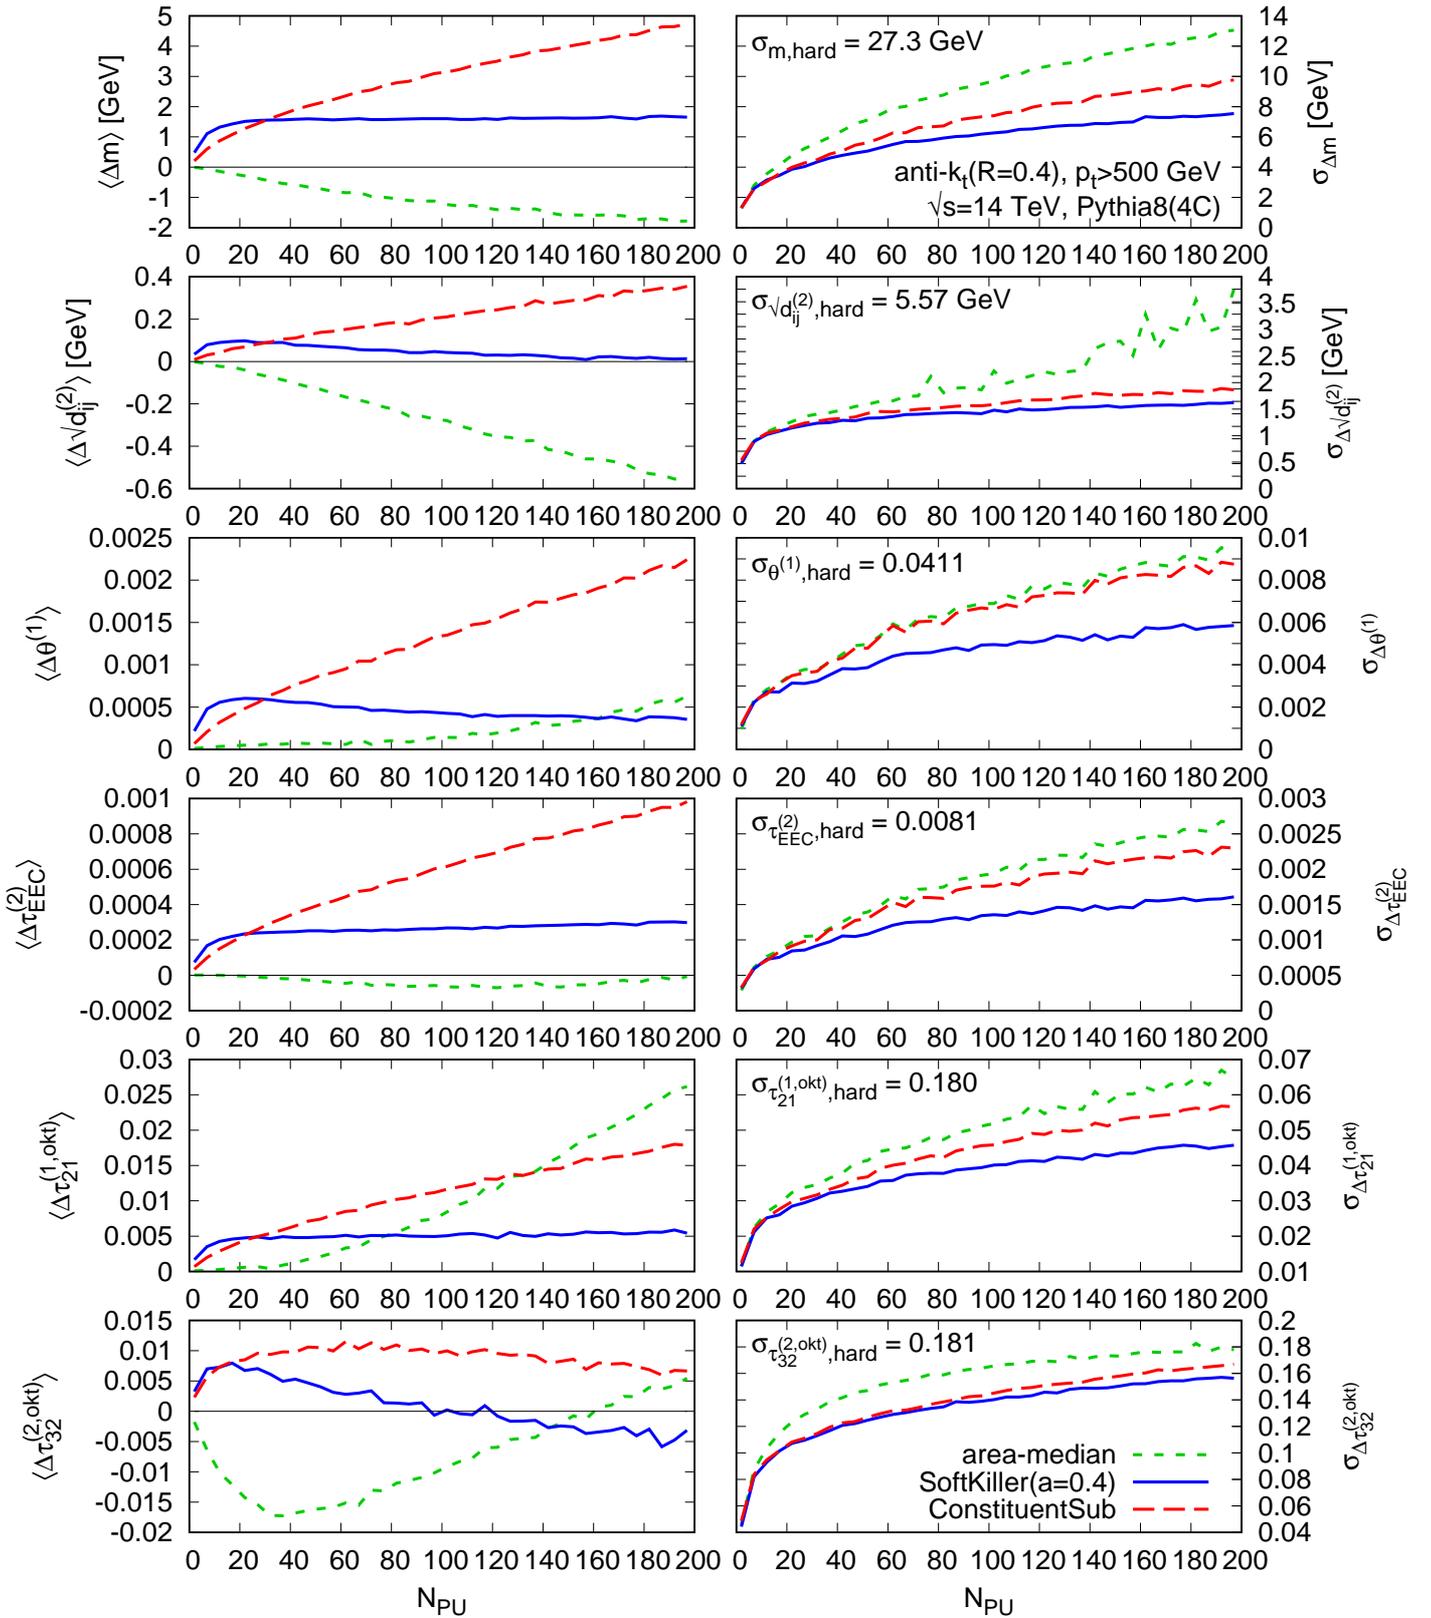


Figure 12.9: Performance of the SoftKiller on jet shapes, compared to area-median subtraction and the recently proposed Constituent Subtractor method [131]. All results are shown for dijet events with a 500 GeV p_t cut on anti- k_t , $R = 0.4$ jets. For comparison of the subtraction performance we also quote, for each observable X , $\sigma_{X,\text{hard}}$, the dispersion of the distribution of the observable in the hard event. For τ_{21} there is the additional requirement (in the hard event) that the jet mass is above 30 GeV and for τ_{32} we further impose that $\tau_{21} \geq 0.1$ (again in the hard event), so as to ensure infrared safety.

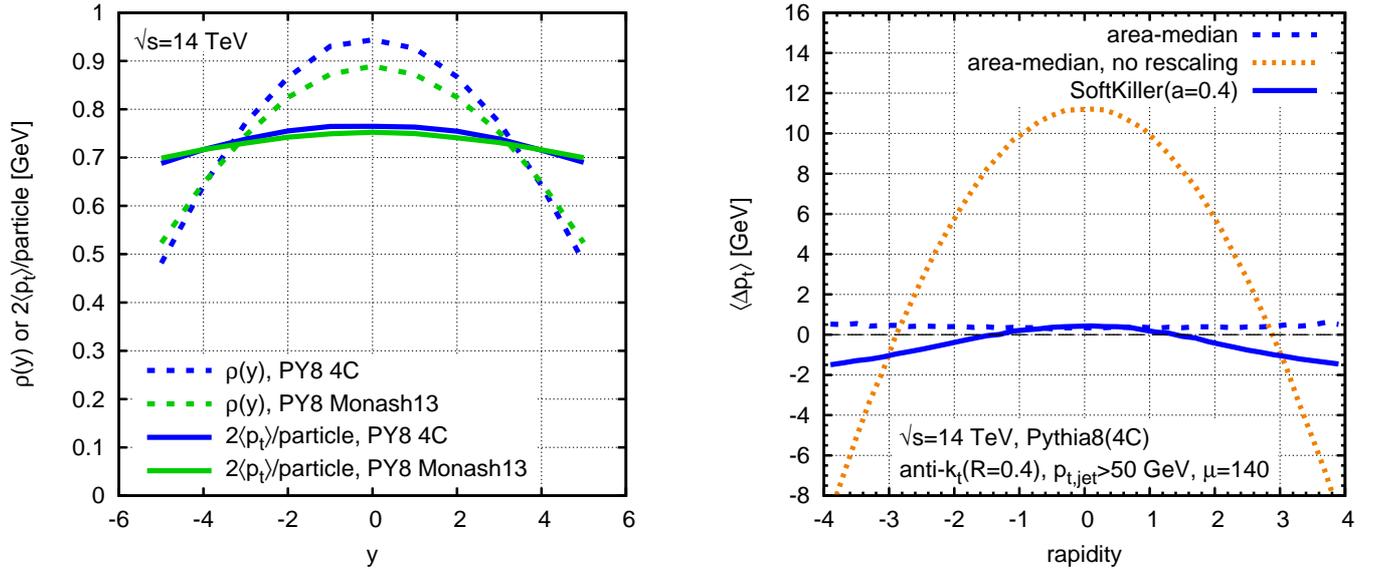


Figure 12.10: Left: rapidity dependence, in simulated zero-bias events, of the p_t density per unit area ($\rho(y)$), and of the average p_t per particle (scaled by a factor of 2 for readability), comparing Pythia8’s 4C and Monash 2013 [133] tunes. Right: the average shift in jet p_t as a function of rapidity for the area–median method with rapidity rescaling (our default), without rapidity rescaling, and for the SoftKiller.

clustering scales [25], the jet width (or broadening or girth [132, 45, 46]), an energy–energy correlation moment [51] and the $\tau_{21}^{(\beta=1)}$ and $\tau_{32}^{(\beta=2)}$ N-subjettiness ratios [54, 55], using the exclusive k_t axes with one-pass of minimisation. Except in the case of the jet mass (which uses “safe” area subtraction, as mentioned above), the area–median results have been obtained using the shape subtraction technique introduced in Section 2.5, as implemented in v. 1.2.0 of the GenericSubtractor in FastJet Contrib.

As regards the shifts, the SK approach is sometimes the best, other times second best. Which method fares worst depends on the precise observable. In all cases, when considering the dispersions, it is the SK that performs best, though the extent of the improvement relative to other methods depends strongly on the particular observable. Overall this figure gives us confidence that one can use the SoftKiller approach for a range of properties of small-radius jets.

Rapidity dependence. Given that the SoftKiller imposes a fixed p_t cut, one might wonder if it correctly account for the rapidity dependence of pileup. In the case of the area–median subtraction, that is handled by the rapidity rescaling used in Eq. (2.27) which we have used throughout this Chapter.¹³ The rapidity dependence of ρ is shown as the dashed lines in Fig. 12.10 (left) for two different tunes (4C and Monash 2013), illustrating the fact that they have somewhat different rapidity dependence.

The SoftKiller method acts not on the average energy flow, but instead on the particle p_t ’s. The solid lines in Fig. 12.10 (left) show that the average particle p_t is nearly independent of rapidity. This suggests that there may not be a need to explicitly account for rapidity in the SK method, at least at particle level (detector effects introduce further non-trivial rapidity dependence).

This is confirmed in the right-hand plot of Fig. 12.10, which shows the rapidity-dependence of the shift in the jet p_t with the area–median and SK methods. Our default area–median curve, which

¹³For our 14 TeV Pythia8(4C) simulations, we use $f(y) = 1.1685397 - 0.0246807 y^2 + 5.94119 \cdot 10^{-5} y^4$.

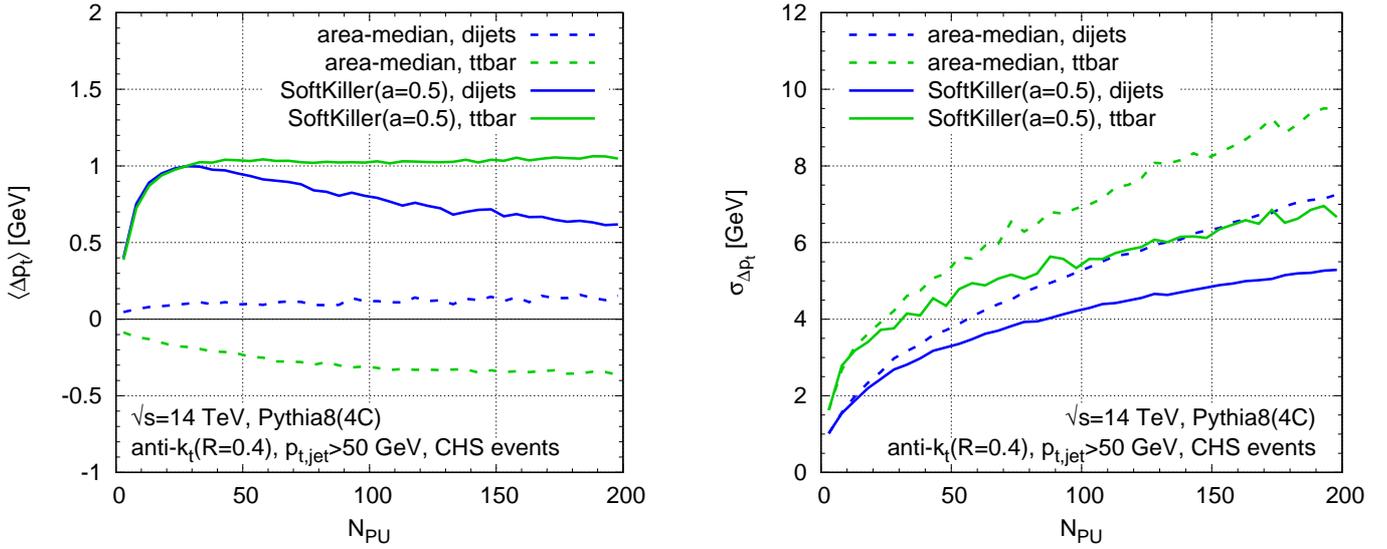


Figure 12.11: Same as Fig. 12.5, for events with charged-hadron subtraction (CHS). Note that the grid size used for the SoftKiller curves has been set to $a = 0.5$.

includes rapidity rescaling, leads to a nearly rapidity-independent shift. However, without the rapidity rescaling, there are instead rapidity-dependent shifts of up to 10 GeV at high pileup. In contrast, the SK method, which in our implementation does not involve any parametrisation of rapidity dependence, automatically gives a jet p_t shift that is fairly independent of rapidity, to within about 2 GeV. We interpret this as a consequence of the fact (cf. the left-hand plot of Fig. 12.10) that the average particle p_t appears to be far less dependent on rapidity than the average p_t flow. In a similar spirit to (2.27), one could also imagine introducing a rapidity-dependent rescaling of the particle p_t 's before applying SoftKiller, and then inverting the rescaling afterwards. Our initial tests of this approach suggest that it does largely correct for the residual SK rapidity dependence.

12.3 Adaptation to CHS events and calorimetric events

It is important to verify that the SoftKiller (or any other new pileup mitigation method with noise reduction) works not just at particle level, but also at detector level, essentially to make sure that the balance between the positive and negative biases discussed in Section 12.1 is not spoiled. There are numerous subtleties in carrying out detector-level simulation, from the difficulty of correctly treating the detector response to low- p_t particles, to the reproduction of actual detector reconstruction methods and calibrations, and even the determination of which observables to use as performance indicators. Here we will consider two cases: idealised charged-hadron-subtraction, which simply examines the effect of discarding charged pileup particles; and simple calorimeter towers.

Performance on CHS events. As in the previous Chapter, for events with particle flow [134] and charged-hadron subtraction (CHS), we imagine a situation in which all charged particles can be unambiguously associated either with the leading vertex or with a pileup vertex. We then apply the SK exclusively to the neutral particles, which we assume to have been measured exactly. This is almost certainly a crude approximation, however it helps to illustrate some general features.

One important change that arises from applying SK just to the neutral particles is that there is a reduced contribution of low- p_t hard-event particles. This means that for a given actual amount of pileup contamination (in terms of visible transverse momentum per unit area), one can afford to cut more aggressively, i.e. raise the p_t^{cut} as compared to the full particle-level case, because for a given p_t^{cut} there will be a reduced loss of hard-event particles. This can be achieved through a moderate increase in the grid spacing, to $a = 0.5$. Figure 12.11 shows the results, with the shift (left) and dispersion (right) for the jet p_t in dijet and $t\bar{t}$ samples. The SK method continues to bring an improvement, though that improvement is slightly more limited than in the particle-level case. We attribute this reduced improvement to the fact that SK's greatest impact is at very high pileup, and for a given N_{PU} , SK with CHS is effectively operating at lower pileup levels than without CHS. In Ref. [11], Appendix E, we also carried on a further study with our toy CHS simulation concerning lepton isolation.

Performance with calorimetric events. Next let us turn to events where the particles enter calorimeter towers. Here we encounter the issue (discussed also in Appendix F) that SK is not collinear safe. While we argue there that this is not a fundamental drawback from the point of view of particle-level studies, there are issues at calorimeter level: on one hand a single particle may be divided between two calorimeter towers (we won't attempt to simulate this, as it is very sensitive to detector details); on the other, within a given tower (say 0.1×0.1) it is quite likely that for high pileup the tower may receive contributions from multiple particles. In particular, if a tower receives contributions from a hard particle with a substantial p_t and additionally from pileup particles, the tower will always be above threshold, and the pileup contribution will never be removed. There are also related effects due to the fact that two pileup particles may enter the same tower. To account for the fact that towers have finite area, we therefore adapt the SK as follows. In a first step we subtract each tower:

$$p_t^{\text{tower,sub}} = \max(0, p_t^{\text{tower}} - \rho A^{\text{tower}}), \quad (12.3)$$

where ρ is as determined on the event prior to any correction.¹⁴ This in itself eliminates a significant fraction of pileup, but there remains a residual contribution from the roughly 50% of towers whose p_t was larger than ρA^{tower} . We then apply the SoftKiller to the subtracted towers, i.e. with a cut given by

$$p_t^{\text{cut,sub}} = \text{median}_{i \in \text{patches}} \left\{ p_{ti}^{\text{tower,sub,max}} \right\}, \quad (12.4)$$

where $p_{ti}^{\text{tower,sub,max}}$ is the p_t , after subtraction, of the hardest tower in patch i , in analogy with Eq. (12.1). In the limit of infinite granularity, a limit similar to particle level, $A^{\text{tower}} = 0$, the step in Eq. (12.3) then has no effect and one recovers the standard SoftKiller procedure applied to particle level.

Results are shown in Fig. 12.12. The energy E in each 0.1×0.1 tower is taken to have Gaussian fluctuations with relative standard deviation $1/\sqrt{E/\text{GeV}}$. A p_t threshold of 0.5 GeV is applied to each tower after fluctuations. The SK grid spacing is set to $a = 0.6$. Interestingly, with a calorimeter, the area–median method starts to have significant biases, of a couple of GeV, which can be attributed to the calorimeter's non-linear response to soft energy. The SK biases are similar in magnitude to those in Fig. 12.5 at particle level (note, however, the need for a different choice of grid spacing a). The presence of a calorimeter worsens the resolution both for area–median subtraction and SK, however SK

¹⁴We use our standard choices for determining ρ , namely the grid version of the area–median method, with a grid spacing of 0.55 and rapidity scaling as discussed in Section 3.2. One could equally well use the same grid spacing for the ρ determination as for the SoftKiller.

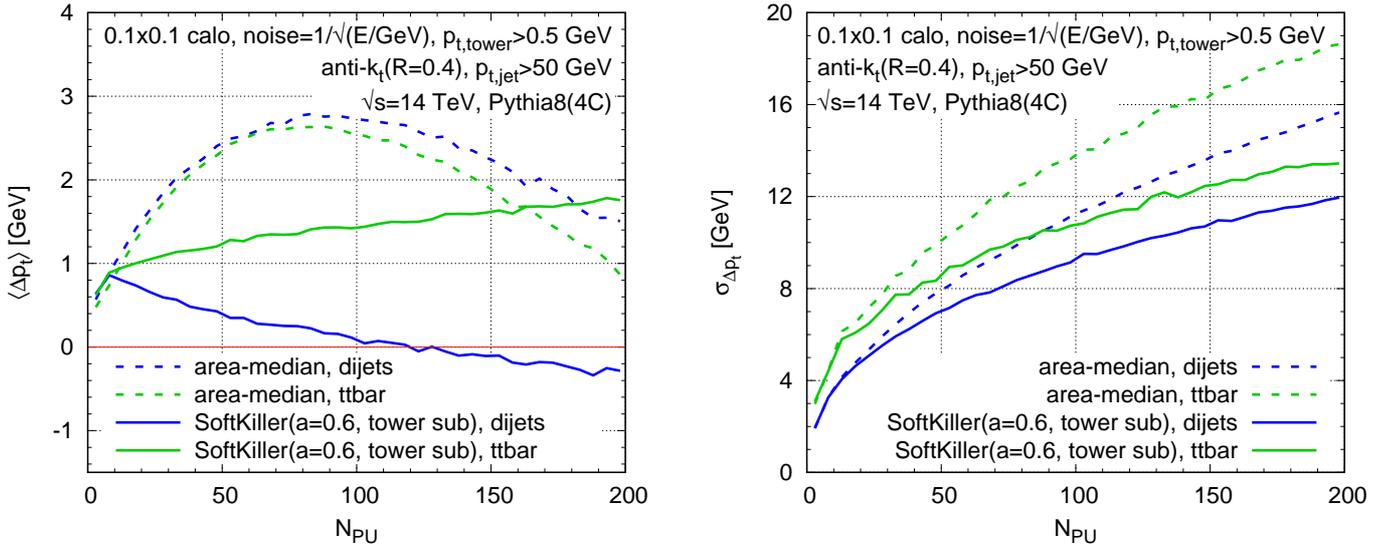


Figure 12.12: Same as Fig. 12.5 for events with a simple calorimeter simulation. The SoftKiller was used here with a grid spacing of $a = 0.6$ and includes the tower subtraction of Eq. (12.3).

continues to perform better, even if the improvement relative to area–median subtraction is slightly smaller than for the particle-level results.

We have also investigated a direct application of the particle-level SoftKiller approach to calorimeter towers, i.e. without the subtraction in Eq. (12.3). We find that the biases were larger but still under some degree of control with an appropriate tuning of a , while the performance on dispersion tends to be intermediate between that of area–median subtraction and the version of SoftKiller with tower subtraction.

The above results are not intended to provide an exhaustive study of detector effects. For example, particle flow and CHS are affected by detector fluctuations, which we have ignored; purely calorimetric jet measurements are affected by the fact that calorimeter towers are of different sizes in different regions of the detector and furthermore may be combined non-trivially through topoclustering. Nevertheless, our results help illustrate that it is at least plausible that the SoftKiller approach could be adapted to a full detector environment while retaining much of its performance advantage relative to the area–median method.

12.4 Computing time

The computation time for the SoftKiller procedure has two components: the assignment of particles to patches, which is $\mathcal{O}(N)$, i.e. linear in the total number of particles N and the determination of the median, which is $\mathcal{O}(P \ln P)$ where P is the number of patches. The subsequent clustering is performed with a reduced number of particles, M , which, at high pileup is almost independent of the number of pileup particles in the original event. In this limit, the procedure is therefore expected to be dominated by the time to assign particles to patches, which is linear in N . This assignment is almost certainly amenable to being parallelised.

In studying the timing, we restrict our attention to particle-level events for simplicity. We believe that calorimeter-type extensions as described in section 12.3 can be coded in such a way as to obtain similar (or perhaps even better) performance.

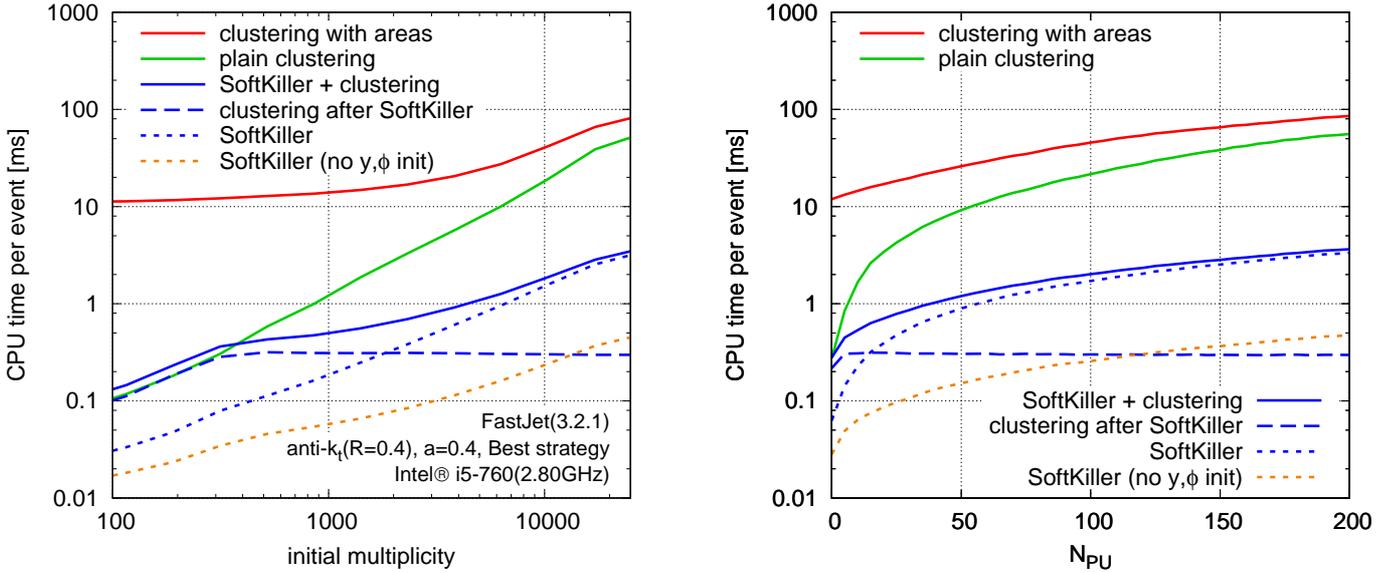


Figure 12.13: Timings of the SoftKiller compared to standard clustering shown as a function of the number of particles in the event (left) or as a function of the number of pileup vertices (right). We compare the SoftKiller timings to the time to cluster the full event with (red) or without (green) jet area information. For the SoftKiller timings (in blue), we show individually the time spent to apply the SoftKiller to the event (dotted line), the time spent to cluster the resulting event (dashed line) and their sum (solid line). The orange dotted line corresponds to the SoftKiller timing when the particle’s rapidity and azimuth have been precomputed.

Timings are shown in Fig. 12.13 versus initial multiplicity (left) and versus the number of pileup vertices (right).¹⁵ Each plot shows the time needed to cluster the full event and the time to cluster the full event together with ghosts (as needed for area-based subtraction). It also shows the time to run the SoftKiller procedure, the time to cluster the resulting event, and the total time for SK plus clustering.

Overall, one sees more than one order of magnitude improvement in speed from the SK procedure, with run times per event ranging from 0.15 to 4 ms as compared to 10 to 90 ms for clustering with area information. At low multiplicities, the time to run SK is small compared to that needed for the subsequent clustering. As the event multiplicity increases, SK has the effect of limiting the event multiplicity to about 300 particles, nearly independently of the level of pileup and so the clustering time saturates. However the time to run SK grows and comes to dominate over the clustering time. Asymptotically, the total event processing time then grows linearly with the level of pileup. A significant part of that time (about 130 ns per particle, 75% of the run-time at high multiplicity) is taken by the determination of the particles’ rapidity and azimuth in order to assign them to a grid cell. If the particles’ rapidity and azimuth are known before applying the SoftKiller to an event (as it would be the case e.g. for calorimeter towers), the computing time to apply the SoftKiller and cluster the resulting event would be yet faster, staying below 0.5 ms, as indicated by the dotted orange line on Fig. 12.13.

Because of its large speed improvement, the SoftKiller method has significant potential for pileup

¹⁵ These timings have been obtained on an Intel processor, i5-760 (2.80 GHz), using FastJet 3.2.1, with the “Best” clustering strategy. Note that this figure has been updated compared to what was originally published in [11], so as to include the speed improvements brought with the release of Fastjet 3.1.

removal at the trigger level. Since SoftKiller returns an event with fewer particles, it will have a speed performance edge also in situations where little or no time is spent in jet-area calculations (either because Voronoi areas or fast approximate implementations are used). This can be seen in Fig. 12.13 by comparing the green and the solid blue curves which stays well below 1 ms even at high multiplicities.

12.5 Discussion and conclusions

The SoftKiller method appears to bring significant improvements in pileup mitigation performance, in particular as concerns the jet energy resolution, whose degradation due to pileup is reduced by 20–30% relative to the area–median based methods. As an example, the performance that is obtained with area–median subtraction for 70 pileup events can be extended to 140 pileup events when using SoftKiller. This sometimes comes at the price of an increase in the biases on the jet p_t , however these biases still remain under control.

Since the method acts directly on an event’s particles, it automatically provides a correction for jet masses and jet shapes, and in all cases that we have studied brings a non-negligible improvement in resolution relative to the shape subtraction method, and also (albeit to a lesser extent) relative to the recently proposed Constituent Subtractor approach.

The method is also extremely fast, bringing nearly two orders of magnitude speed improvement over the area–median method for jet p_t ’s. This can be advantageous both in time-critical applications, for example at trigger level, and in the context of fast detector simulations.

There remain a number of open questions. It would be of interest to understand, more quantitatively, why such a simple method works so well and what dictates the optimal choice of the underlying grid spacing. This might also bring insight into how to further improve the method. In particular, the method is known to have deficiencies when applied to large- R ungroomed jets, which would benefit from additional study (see Section 13.2.1 for a possible solution). Finally, we have illustrated that in simple detector simulations it is possible to reproduce much of the performance improvement seen at particle level, albeit at the price of a slight adaption of the method to take into account the finite angular resolution of calorimeters. These simple studies should merely be taken as indicative, and we look forward to proper validation (and possible further adaptation) taking into account full detector effects. It would also be interesting to see a performance comparison of the SoftKiller with the recently-proposed PUPPI pileup mitigation technique [113]. We will come back to this point in our conclusions.

Chapter 13

A (rich) forest of ideas to explore

13.1 Alternative noise-reduction methods based on particle areas

Besides the SoftKiller, we have considered a series of alternative simple methods. These have only been put through a very minimal set of tests and should be considered as preliminary. Nonetheless, we find helpful to document our findings since some of the corresponding ideas can prove useful for future pileup mitigation techniques. Note that, since we were after a simple pileup-mitigation technique, all the methods presented in this Section have essentially a single free parameter. Introducing additional parameters could, at the expense of adding complexity, bring some extra flexibility and help finding improvements.¹

Since our goal is to provide potential alternatives for generic pileup-mitigation techniques, we shall focus mainly on the performance for the reconstruction of the jet transverse momentum. We will also consider briefly performance for the reconstruction of the jet mass. As we have already argued in the last two Chapters, jet properties like the jet mass or jet shapes are often used in the context of boosted jets, where one starts from a large- R jets and apply a grooming technique. The study of the jet mass presented below should therefore merely be seen as a first basic check that our methods do not spoil critically measurements of the internal properties of jets. Additional studies like jet shapes (the equivalent of Fig. 12.9 in our SoftKiller studies) or boosted jet reconstruction (the equivalent of Fig. 12.8 in our SoftKiller studies) are left for future studies.

One avenue that we have investigated is to try to extend our area-based approach beyond jets and to apply it directly to the particles in the event. This first requires that we extend the concept of jet area to individual particles. The definition we have adopted can be seen as a *particle Voronoi area*: we build the Voronoi graph from all the particles in the event, withing a given rapidity acceptance, and the area associated to each particle is simply the area of its Voronoi cell, restricted to the particle's rapidity acceptance so as to avoid infinite cells at the edges of the acceptance. The 4-vector area is then defined as the massless 4-vector with transverse momentum given by the Voronoi (scalar) area, and rapidity and azimuthal angle of the particle.

With that definition at hand, there is an apparent straightforward way to apply the area–median subtraction directly at the level of the particles: if A_{particle}^μ is the area of a given particle, one subtracts $\rho A_{\text{particle}}^\mu$ from its momentum, discarding the particle altogether if its transverse momentum p_t is smaller than ρA_{particle} . This is unfortunately over-simplified since many soft particles, corresponding

¹Like with the area–median subtraction (Section 5.2) and the substructure-based methods discussed in Chapter 10, it would be very interesting to see if analytic QCD calculations could help constraining the free parameters.

to local downwards fluctuations of the pileup background would be discarded (i.e. set to a zero p_t instead of a negative one), cutting the negative tail of pileup fluctuations while keeping the positive tail, ultimately leading to an under-subtraction of pileup.

In our preliminary studies, we have considered several particle-based pileup mitigation methods trying to avoid this under-subtraction issue. We have compared them to the area–median approach as well as to the SoftKiller approach presented in the previous Chapter.

13.1.1 List of methods

Here is the list of methods we shall discuss in this Section:

- The **Area–median** method serves as a robust unbiased baseline for comparison. Following our recommendations from Section 2.8 we have used a grid-based estimation of ρ with a grid-size parameter of 0.55 and rapidity rescaling. We have imposed positivity constraints on the jet p_t and mass.
- The **SoftKiller** method discussed at length in the previous Chapter, that removes the soft particles in the event until the estimated ρ vanishes.
- A somewhat similar idea, that we shall call the **Soft-Removal** method in what follows, is to remove the soft particles in the event until the total (scalar) p_t that has been eliminated reaches $f_{\text{SR}}\rho A_{\text{tot}}$, where ρ is the pileup density (if rescaling is active, it is estimated at $y = 0$), A_{tot} is the total detector area and f_{SR} an overall factor, taken as a free parameter. Taking f_{SR} to 1, we should expect to over-subtract the transverse momentum of the hard jets since the unsubtracted pileup fluctuations left across the whole event would be partially compensated by the subtraction of particles from the hard jets. We shall therefore consider values of f_{SR} smaller than 1 in what follows.
- A **Voronoi Killer** method which works, in a similar spirit as the SoftKiller, by removing the softest particles in the event until, a fraction f_{VK} of the total event area, computed from the particle Voronoi area, is empty.
- Let us come back to the idea of subtracting ρA_{particle} from each particle, setting to 0 the particles with $p_t < \rho A_{\text{particle}}$. As argued above, this would leave a potentially large positive bias (i.e. an under-subtraction). This bias is expected to be proportional to the internal fluctuations of the pileup in the event, i.e. to σ . A possible solution would be to subtract, on top of ρA_{particle} from each particle, an additional amount proportional to $\sigma\sqrt{A_{\text{particle}}}$. We have tried 3 approaches based on this idea:
 - **Voronoi Subtraction**: we apply a subtraction to each of the original particle. For a particle with original transverse momentum p_t and Voronoi area A , we set its subtracted transverse momentum to $\max(p_t - \rho A - f_{\text{VS}}\sigma\sqrt{A}, 0)$ where f_{VS} is an adjustable parameter.
 - **Voronoi Cut**: the subtracted event is built by keeping the particles with $p_t > \rho A - f_{\text{VC}}\sigma\sqrt{A}$ (untouched) and discarding the others. f_{VC} is again an adjustable parameter.
 - **Voronoi Subtract&Cut**: we first subtract $\rho\sqrt{A}$ from each particle in the event, discarding particles with $p_t < \rho\sqrt{A}$. We then remove all the particles for which the subtracted p_t is below $f_{\text{VSC}}\sigma\sqrt{A}$, with f_{VSC} an adjustable parameter. Note that this last option is somewhat reminiscent of the adaptation of the SoftKiller method to calorimeters — see Section 12.3 — and can therefore likely be applied straightforwardly in an experimental context.

There are also two other methods that we briefly investigated but were relatively rapidly abandoned since our initial findings did not give very promising results. For the sake of completeness, we introduce them here, discuss why we believe they are not efficient and quote possible options to improve them. Both methods subtract from the event a total transverse momentum of ρA_{tot} , with A_{tot} the total area of the detector. They both first subtract ρA_{particle} from each particle, discarding the particles for which this would give a negative transverse momentum. Since we have put some of the particles to 0 transverse momentum, the total subtracted scalar p_t is less than ρA_{tot} and an extra subtraction is needed. We then consider two possible options:

- the excess, between ρA_{tot} and what has been subtracted from the initial step is subtracted uniformly over the remaining particles. We again subtract an amount proportional to their area, setting to 0 the particles for which this would give a negative p_t . The operation is repeated until a total of ρA_{tot} has been subtracted. We dubbed this method *Voronoi subtraction with Uniform Balance*.
- alternatively, after the initial subtraction step, we successively remove the softest particle in the event until a total p_t of ρA_{tot} has been subtracted. We call this the *Voronoi subtraction with Soft Balance*.

In both cases, the idea here is to balance the negative part of the fluctuations, cut by the initial subtraction, with the positive part, hopefully discarded in subsequent subtractions or particle removal.

In both cases, this actually lead to a rather larger negative bias on the reconstruction of the jet p_t . We believe this is due to a reason similar to our argument in to use $f_{\text{SR}} < 1$ for the Soft-Removal method. After the initial subtraction, we are essentially left with two kinds of objects: hard jets located around a few directions in the event and pileup upwards fluctuations distributed uniformly across the event. Ideally, only the pileup component need to be subtracted but in practice, both these components will be affected by the subsequent steps of the subtraction. This will lead to some upwards pileup fluctuations left uniformly across the event and an oversubtraction of the hard jets to compensate for that effect. Similarly to the Soft-Removal method, a tentative solution to that issue would be to adapt the method and only subtract a fraction f of the total ρA_{tot} , with f an adjustable parameter. We have not (yet) investigated that option.

Note also that we could have considered a third option where the amount left to subtract after the initial subtraction is balanced using the particle's neighbours (there are actually several ways to achieve that). In practice, the hard particles in the jet would be used to compensate the negative fluctuations in the jet's vicinity, yielding once more an over-subtraction, possibly even larger than with the other two approaches.²

Finally, in a more complete study, it would be interesting to include a comparison to the ConstituentSubtractor, as well as other intra-jet methods, like an *intrajet killer* which would iteratively remove the softest particle in a given jet until a total p_t of ρA_{jet} has been subtracted (see footnote 12 in Chapter 12). These methods are expected to give results very close to the area-median for the jet transverse momentum and a small improvement for the jet mass.

13.1.2 Monte-Carlo studies

Description. As usual, we work by embedding hard events in a superposition of minimum bias events to simulate pileup. Both the signal events and the minimum bias events have been generated

²The Constituent Subtractor has likely a similar issue when applied to the whole event.

with Pythia 8186 with the 4C tune and, for simplicity, we have set the mass of the initial particles to zero, preserving their transverse momenta, rapidities and azimuthal angles.³ We keep particles up to $|y| = 5$. We cluster the jets using the anti- k_t algorithm with $R = 0.4$. In the hard event, we select the jets passing a given p_t cut that we shall vary to probe the robustness of the pileup mitigation methods, and with $|y| < 4$.

We apply the various event-wide pileup subtraction techniques introduced in the previous Section to the full event — including the hard event and pileup — and for each of the original hard jets, find the most overlapping jet in that subtracted event. We then focus on the shift $\langle \Delta p_t \rangle$ and dispersion $\sigma_{\Delta p_t}$ quality figures already used several times so far, in order to assess the performance of the methods introduced above. As an additional check, we also look at the reconstruction of the jet mass by studying $\langle \Delta m \rangle$ and $\sigma_{\Delta m}$.

The parameters of our candidate pileup mitigation techniques are fixed/varied as follows. For the area–median subtraction, acting as our baseline, we use a grid-based estimation of ρ with the grid-size parameter set to 0.55 and rescaling to handle the rapidity dependence. The SoftKiller size parameter is varied between 0.3 and 0.7.⁴ For the Soft Removal technique, we use the pileup density ρ coming from the area–median method, calculated at $y = 0$ and vary f_{SR} between 0.66 and 0.76. Moving to the methods based on Voronoi areas, we calculate the particle areas up to $|y| = 5$. For the Voronoi Killer, we vary f_{VK} between 0.94 and 0.98. Finally, for the last three Voronoi-based methods, we compute ρ and σ independently for each particle, using the area–median technique, including the rapidity dependence. We vary f_{VC} between 1.5 and 3.5, f_{VSC} between 0.5 and 3, and f_{VS} between 0.2 and 0.7.

Since a key element common to all these new techniques is their potentially uncontrolled biases, we want to study different processes and pileup conditions to assess the robustness of the methods. We do this by selecting a few working points. First, we work with fixed pileup conditions, taking $\mu = 60$ as a default, and study a few representative cuts on the jet p_t , namely 20, 100 and 500 GeV. Conversely, for a fixed jet p_t cut of 50 GeV, we vary the pileup conditions, choosing $\mu = 30, 60$ or 140. This should be sufficient to give a first overview of the performance of our candidate methods when varying the two main scales in the problem: the scale of the hard jet and that of pileup.

Results. Summary plots are presented in Fig. 13.1 for the p_t variations at fixed μ and in Fig. 13.2 for the μ dependence at fixed p_t . Overall, the message to take from these figures is that, as foreseen, one can obtain a much better resolution than with the area–median approach but at the expense of potentially large biases, strongly dependent on the adjustable parameter of the method.⁵ Below, we make more precise comments individually for each method.

Note that while the discussion focuses on the reconstruction of the most fundamental property of jets, their p_t , similar conclusions can be reached for the jet mass for which summary plots are presented in Figs. 13.3 and 13.4.

SoftKiller. The SoftKiller with $a \approx 0.4 - 0.5$ shows a nice stability against both p_t and μ variations, as well as a significant resolution improvement compared to the area–median approach. Here, we

³Compared to the studies done in the previous Chapter, we have kept the π^0 stable. As we shall see, this slightly affects the optimal parameter for the SoftKiller.

⁴This was already studied in the previous Chapter. We included it here in order to provide an easier comparison between the different pileup mitigation methods.

⁵It would be interesting to introduce a measure of the fine-tuning associated with these free parameters.

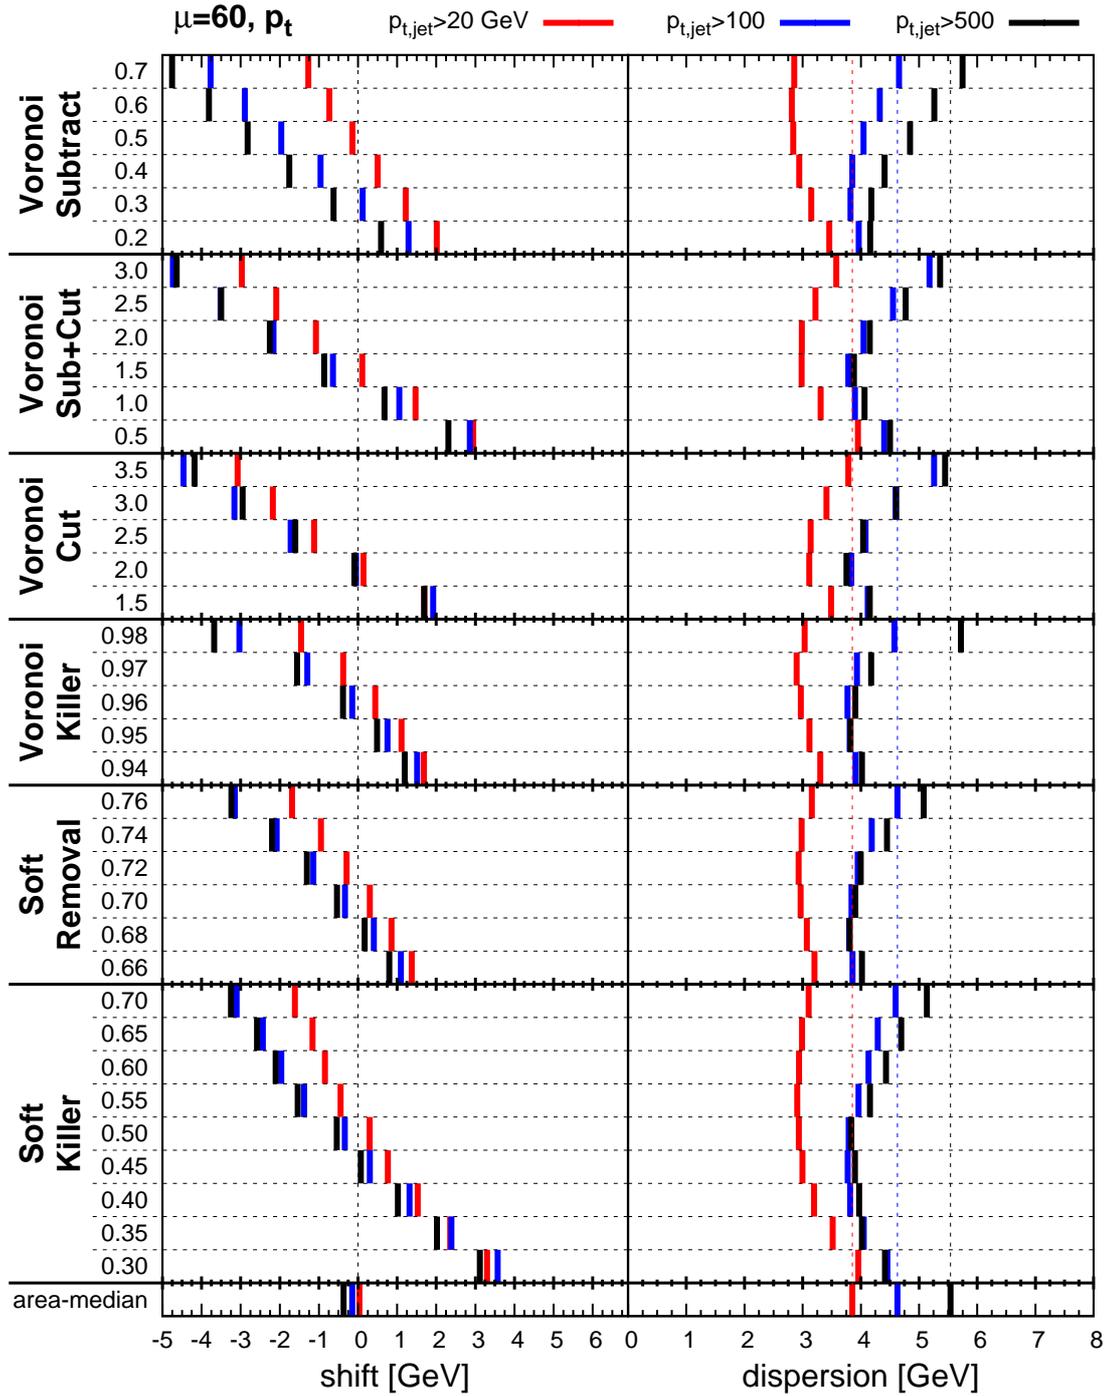


Figure 13.1: Summary of the quality measures for various pileup mitigation methods. The average p_t shift $\langle \Delta p_t \rangle$ (left panel) and the associated dispersion $\sigma_{\Delta p_t}$ (right panel) are plotted for $\mu = 60$ and different jet p_t cuts (20, 100 and 500 GeV respectively in red, blue and black). From bottom to top, we show the results for the area–median method, for the SoftKiller for various grid size a , for the SoftRemoval for different f_{SR} , for the Voronoi Cut, the Voronoi Subtract&Cut and the Voronoi Subtraction methods, for different f_{VC} , f_{VSC} and f_{VS} , respectively. For the average shift, the dashed vertical line corresponds to the ideal situation of a zero bias. For the dispersion plot, the vertical lines correspond to the baseline obtained with the area–median subtraction method.

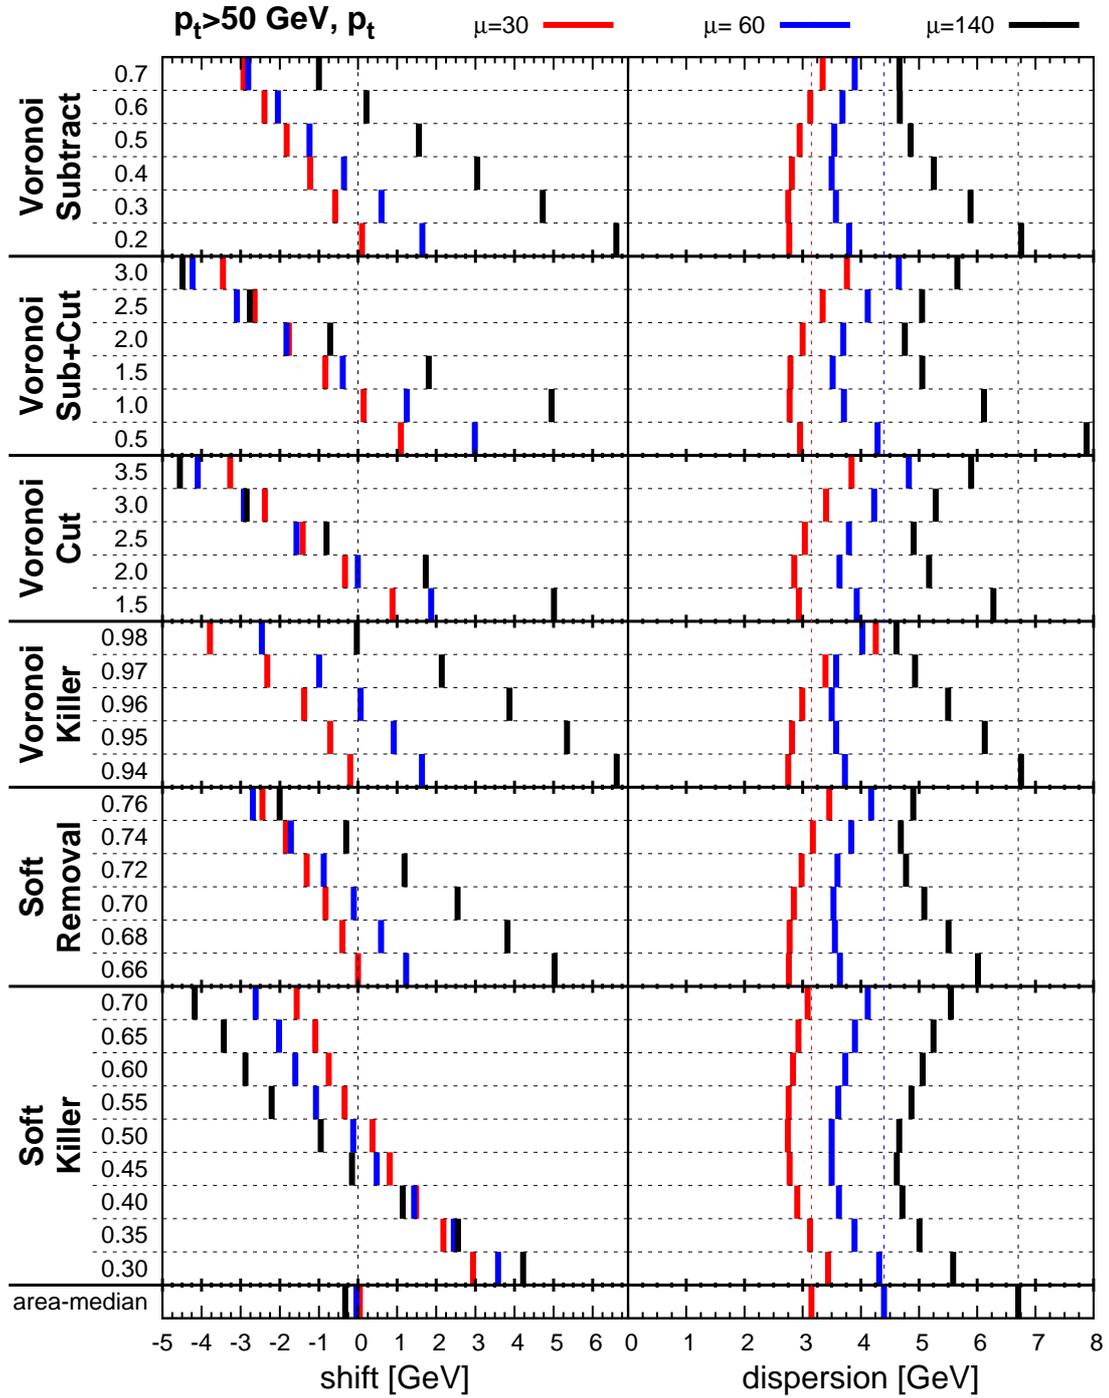


Figure 13.2: Same as Fig. 13.1 but this time for a sample of hard jets with $p_t \geq 50 \text{ GeV}$ and studying the sensitivity of our results w.r.t. to variations of the pileup multiplicity. Different curves correspond to $\mu = 30, 60$ and 140 (respectively in red, blue and black).

recover the results presented in details in the previous Chapter.⁶

Soft Removal and Voronoi Killer. These two methods show similar characteristics. For fixed pileup conditions, the subtraction bias is reasonably stable against changes in the jet p_t with a region where it remains close to zero, but a fixed value of the adjustable parameter is unable to stay unbiased when varying the pileup conditions.

This could actually have been anticipated. If one takes, for example, the case of the Voronoi Killer, as pileup increases, the area occupied by pileup particles (resp. hard particles) increases (resp. decreases). Removing a fixed area from the event will then result in a bias increasing with pileup. A possible improvement of the method would be to adjust f_{SR} or f_{VK} with the pileup conditions.⁷

It is also interesting to note that, while values of f_{VK} are close to one, f_{SR} is much smaller. This is likely due to the fact that we have used the value of ρ estimated at $y = 0$. If we take into account the rapidity profile of ρ , introduced in Chapter 3, we find that, for $\sqrt{s} = 13$ TeV, the averaged ρ for $|y| < 5$ is about 83% of the $\rho(y = 0)$, a value only slightly above our range for f_{SR} .

That said, if we consider a value of their parameter where the bias remains small, both the Soft Removal and the Voronoi Killer methods show a resolution similar to that obtained with the SoftKiller. This means that, modulo the dependence of f_{SR} and f_{VK} on the pileup multiplicity, which might likely be parametrised easily, these methods could give results similar to that of the SoftKiller.

Voronoi Cut, Subtract&Cut and Subtract. Comparing the average bias observed with these three methods, one sees that, in the region where the bias is small, the Voronoi Subtract Method shows a larger dependence on p_t and μ than the Voronoi Cut, with the mixed Voronoi Subtract&Cut method being intermediate. This indicates that a method that either keeps a particle untouched or discards it shows more stability than a method that applies a more democratic subtraction. The Voronoi Cut technique actually shows a behaviour comparable to the Soft Removal and Voronoi Killer methods discussed above with an even slightly smaller dependence on the jet p_t and μ . As long as we do not go to very large pileup multiplicities, the Voronoi Cut method with $f_{\text{VC}} \approx 2$ shows similar performance and robustness than the SoftKiller, both in terms of the average bias than in terms of dispersion, with a small advantage for the latter in terms of dispersion. Going to larger pileup multiplicities would probably necessitate a small increase of f_{VC} to maintain a similar level of performance.

Conclusions. In the end, all the methods show relatively similar potential resolution gains compared to the original area–median approach. Amongst the list of methods we have considered, the SoftKiller is the one showing the best stability when varying p_t and μ and so is the one we ultimately decided to select for a deeper study.⁸ It is nevertheless interesting to notice that the Soft Removal, the Voronoi Killer and, especially, the Voronoi Cut methods, for fixed pileup conditions, all have a working point for which the average bias remains small independently of the jet p_t .⁹ Adjusting their respective free parameter with the pileup conditions would deserve further study and can potentially give performances similar to what is reached with the SoftKiller. Since the studies here remain fairly

⁶Up to a small variation of the dependence on the SoftKiller parameter a related to our different event simulation with stable π^0 .

⁷For example, a behaviour like $f_{\text{VK}} = N_{\text{PU}}/(a + N_{\text{PU}})$ — with $a \sim 2$ according to Fig. 13.2 — would probably give the right asymptotic behaviour at large pileup and nicely go to 0 in the absence of pileup.

⁸Speed arguments also play in favour of the SoftKiller, although we have not put any effort in optimising the speed of the other methods.

⁹This test should also be carried with busier events like hadronically decaying $t\bar{t}$ events.

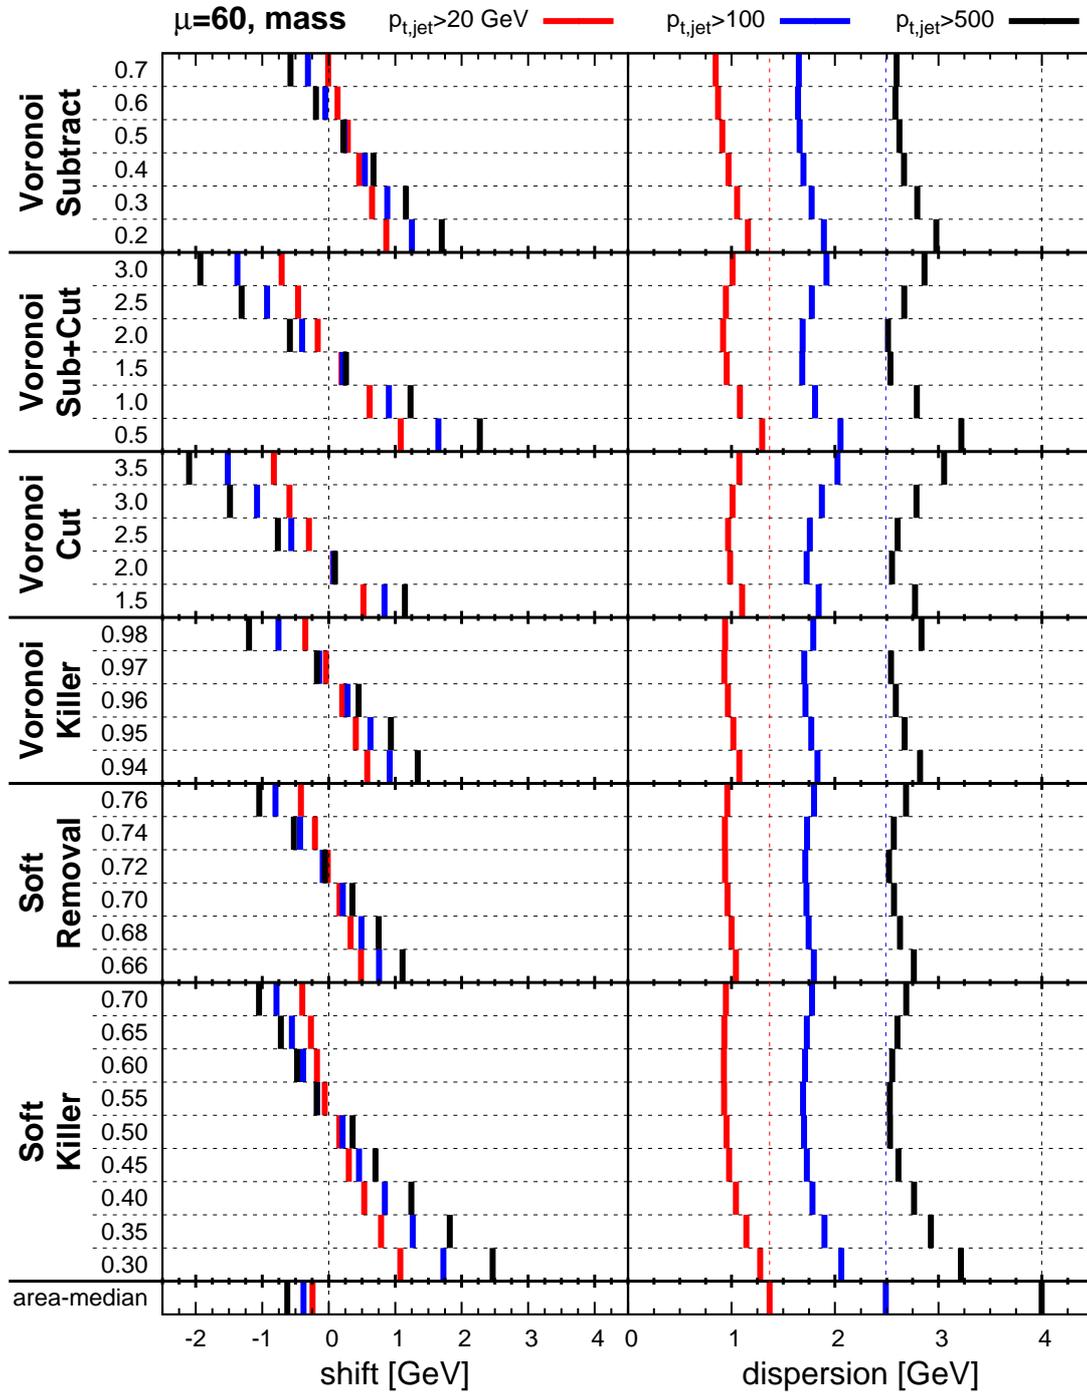


Figure 13.3: Same as Fig. 13.1 but this time for the jet mass rather than for the jet p_t . The left plot shows the average shift and the right plot shows the dispersion.

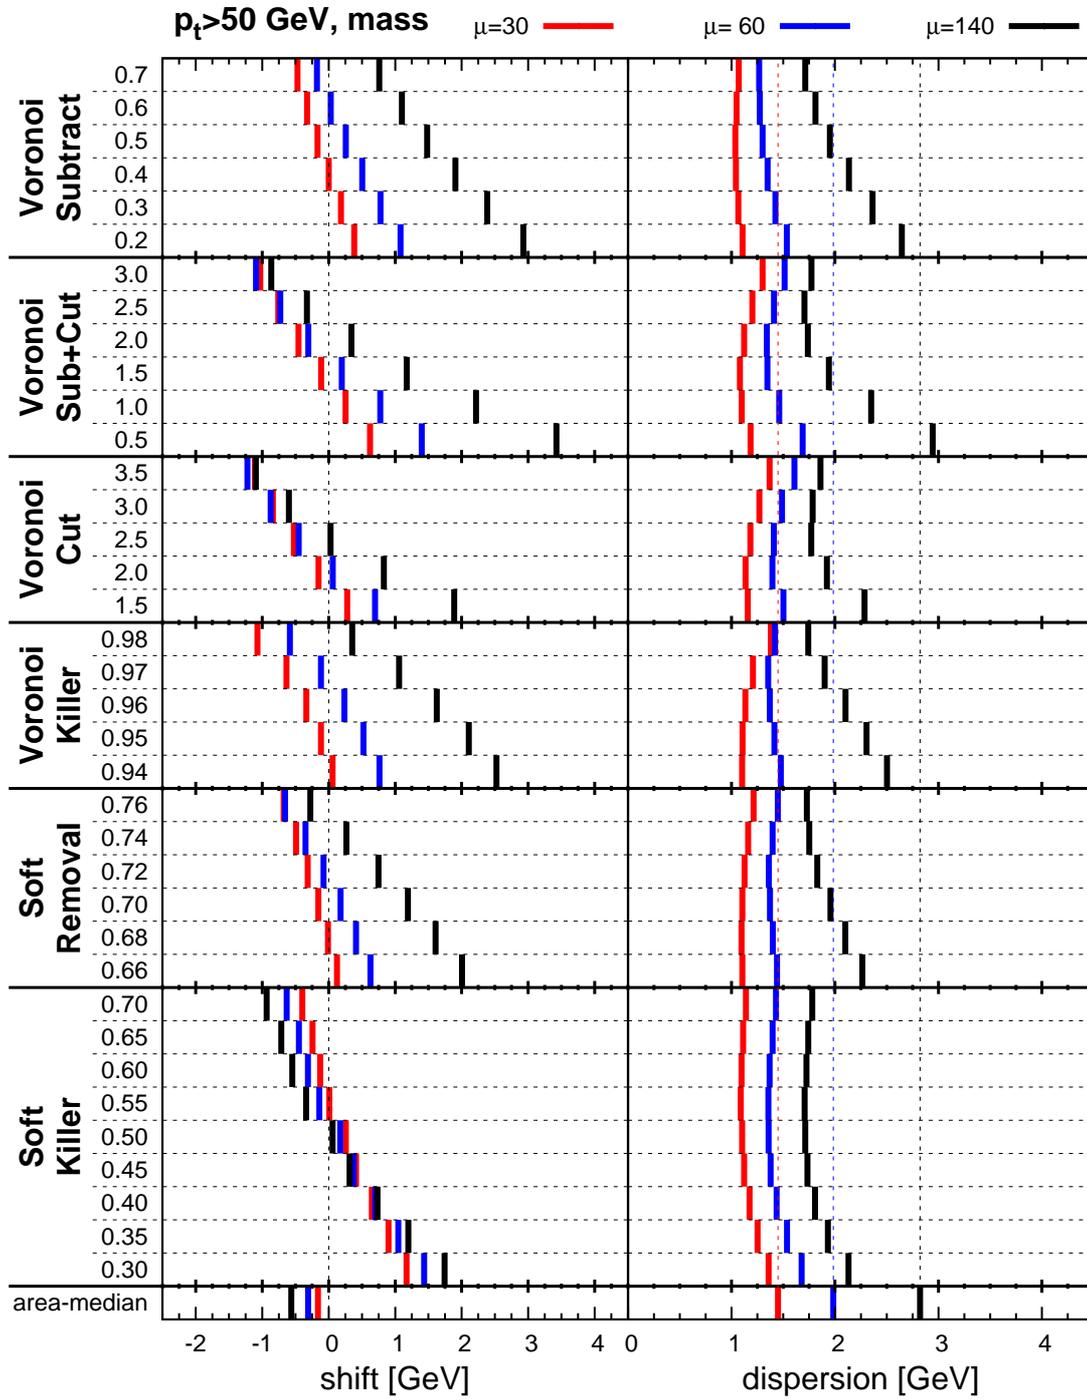


Figure 13.4: Same as Fig. 13.2 but this time for the jet mass rather than for the jet p_t . The left plot shows the average shift and the right plot shows the dispersion.

basic, e.g. they do not discuss CHS events or detector effects, the alternative approaches presented above, and the ideas they are based on, should be kept in mind for future pileup mitigation studies.

13.2 SoftKiller improvements

Studies in the previous Chapter and in the previous Section have shown that the SoftKiller is a very promising noise-reduction pileup mitigation technique. In this Section, we investigate possible improvements of the method. All the results presented here are to be considered as preliminary. We nevertheless believe that they carry useful information and might prove useful in the longer run.

In practice, we want to study two specific points. The first one is the question of the angular dependence: as we have seen in the previous Chapter, the parameter a of the SoftKiller tends to depend on the jet radius R used to cluster the final jets and we want to propose a possible workaround that would work independently of R . Note that fat jets are only one possible application of this study. The other application, potentially of broader impact, could be an improvement of the SoftKiller already at $R = 0.4$ through an improved description of the angular dependence at smaller R .

The second question is related to our observation in Chapter 11 that “(protected) zeroing” can be useful in reducing the impact of pileup fluctuations at the expense of potentially biasing the reconstructed averaged jet properties. We want to see if zeroing can be combined with the SoftKiller for further performance gains.

13.2.1 Improved R dependence

Motivation and description. One piece of physical information that is not encoded in the SoftKiller technique is the fact that the QCD radiation from a hard jet is collinear to the jet, peaking at small angles, while pileup tends to be mostly uniform in rapidity and azimuthal angle.

In practice, this means that if one goes further away from the jet axis, particles (of a given p_t) are more likely to come from pileup. In the context of the SoftKiller, the consequence is that if we want to reconstruct jets with a larger radius, one has to increase the p_t cut on the particles which is done in practice by increasing the size parameter a of the SoftKiller.

In that context, it is natural to wonder if one can encode some form of angular dependence directly in the SoftKiller recipe to automatically allow for jet clustering with any radius. A side benefit of that approach might also be that we could also obtain a performance gain at $R = 0.4$.

Implementing the physical idea that hard jets have collinear radiation while pileup is more uniform has one major difficulty: it requires some form of knowledge of where the hard jets are.

The approach we have tested is the following. For a given event, we first apply the SoftKiller with an initial grid-size parameter a . This corresponds to a p_t cut $p_{t,0}$. We then cluster the event with a jet radius R_0 for which we expect the initial SoftKiller to give good performances. In our studies, we chose the anti- k_t algorithm but this choice should not matter too much. We then keep the jets above a certain $p_{t,\text{cut}}$ as our “reference hard jets”. Finally, for each particle in the event, we apply the following cut: we compute its distance to the nearest reference hard jet, Δ , and discard that particle if its transverse momentum is below

$$p_{t,0} \left(\frac{\max(\Delta, R_0)}{R_0} \right)^\kappa, \quad (13.1)$$

where κ is a parameter controlling the angular dependence of the cut.

Compared to the plain SoftKiller, this has three extra parameters: R_0 and $p_{t,\text{cut}}$ defining the reference jets, and κ . Adding parameters allows for much more fine tuning but, at the same time, leads to a much larger phase-space to explore. A complete and thorough exploration of this phase-space goes beyond the reach of this preliminary analysis. In what follows, we shall just present a few results based on a first, very basic, exploration of the ideas we just introduced.

Testing framework. In practice, we have used the same framework as for our SoftKiller studies. We have worked with full events (i.e. not assuming anything about charged track or neutral particles), and let the π^0 decay. As a first attempt to test a R -dependent version of the SoftKiller, we have focused on a single hard process. We have considered dijet events with a p_t cut of 50 GeV. We have assumed $\mu = 60$ in all our studies. In order to explore the parameter phasespace of the method, we have varied a between 0.2 and 0.6 by steps of 0.1, R_0 between 0.1 and 0.4 by steps of 0.1, tested $\kappa = 0.5, 1$ or 2 , and $p_{t,\text{cut}} = 5, 10$ or 20 GeV. Since our main goal is to test the R dependence of the method, we have studied independently $R = 0.4, 0.7$ and 1 . With no surprise, we have looked at the average shift Δp_t and the associated dispersion $\sigma_{\Delta p_t}$ with the same matching procedure as the one used for our SoftKiller studies.

Results. We present a summary of our findings in Fig. 13.5. For simplicity we show only results for $p_{t,\text{cut}} = 10$ GeV and $\kappa = 0.5$. Smaller values of $p_{t,\text{cut}}$ tend to give worse results and larger values only give a small improvement and could be risky for the reconstruction of jets of lower p_t . Larger κ values show similar performance but a slightly larger dependence on the jet radius R .¹⁰

In order to visualise the performance in a compact form, we have shown on each plot both the average shift and the dispersion, as a function of the remaining parameters, a and R_0 . The black lines indicate contours of equal $\langle \Delta p_t \rangle$ and we show them for $\langle \Delta p_t \rangle = 0, \pm 1$ or ± 2 GeV (respectively the solid, dashed and dotted lines). Ideally, we want to stay as close as possible to the solid black line which corresponds to an unbiased method. The other lines indicate contours of equal $\sigma_{\Delta p_t}$. For those we have first computed the optimal $\sigma_{\Delta p_t}$ for the plain SoftKiller (with a kept within the values mentioned above, namely 0.2 – 0.6) and show contours where the dispersion is larger than this minimum by 2, 5, 10 or 20%. The results for the SoftKiller alone correspond to the values reported at the top of the plot, as a function of the parameter a only.

In the end, it seems that a choice $a \approx 0.35$, $R_0 \approx 0.2 - 0.25$ shows good performance across the whole range of R values. However, this comes with a series of remarks and caveats.

First, we have only tested one process, one jet p_t and one value of μ . Checking the robustness of our findings for a range of processes and pileup conditions is mandatory before claiming any solid conclusion. In a similar spirit, it would be interesting to repeat the study with CHS events and, optionally, with (protected) zeroing (see the next Section). Performance for other observables like missing E_T would also be interesting to study.

Moreover, it appears that the SoftKiller alone, with $a \approx R$ seems to give small biases and a dispersion comparable or even better than a fixed “ R -dependent Killer”. Since larger values of R are mostly used for specific purpose, one might therefore also consider applying the SoftKiller as described in Chapter 12, guaranteeing good performance for $R = 0.4$ and then apply (or re-apply) the SoftKiller with a larger a whenever one needs a larger R . In that context, it is important to note that the use of grooming techniques, inherent to the majority of the analyses involving larger R values, are likely

¹⁰One should however also keep in mind that a too small value for the p_t cut could increase sensitivity to detector effects as well as to pileup jets.

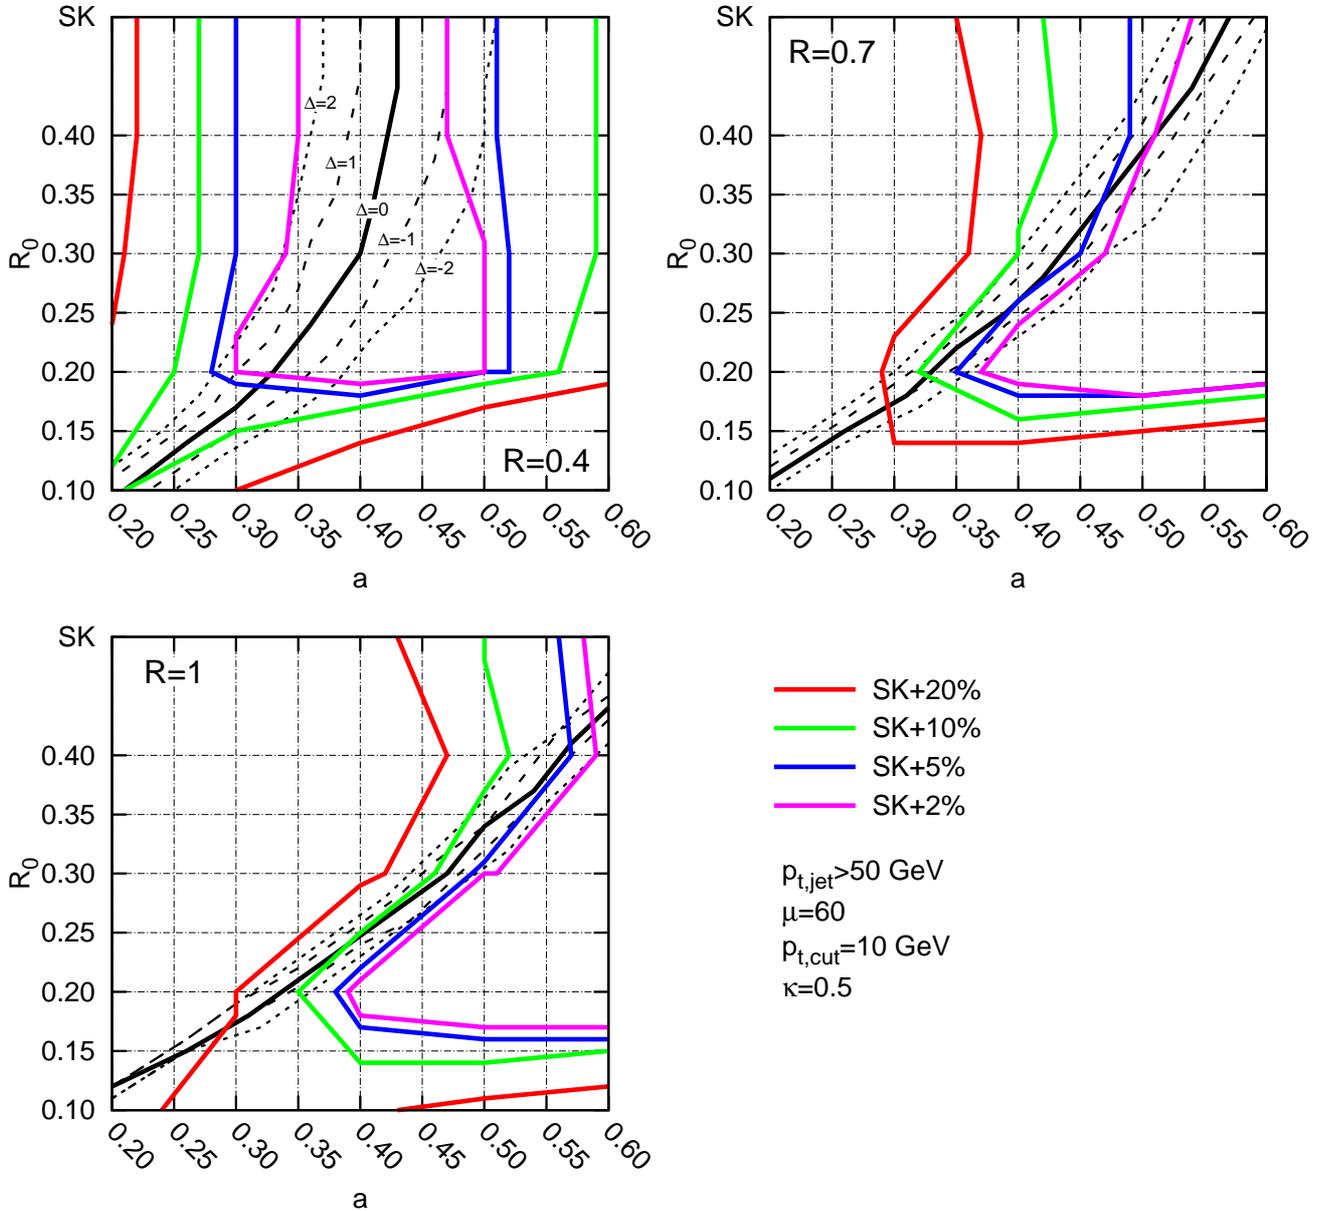


Figure 13.5: Summary of our results for the R -dependent extension of the SoftKiller. Each plot corresponds to different radius used to cluster jets: $R = 0.4$ on the top left plot, 0.7 on the top right plot and $R = 1$ on the bottom plot. On each plot, we have plotted contours of equal shift and dispersion in the plane of the a and R_0 parameters of the method. The top of the plot corresponds to the SoftKiller which only depends on a . The black lines indicate contours where the average bias $\Delta \equiv \langle \Delta p_t \rangle$ is $0, \pm 1$ or ± 2 GeV. The other lines indicate contours of equal dispersion $\sigma_{\Delta p_t}$ where we have used the optimal dispersion for the plain SoftKiller (with $0.2 \leq 1 \leq 0.6$) as a reference.

to affect the picture. This was already discussed in Section 12.2 where we noticed that a value of a valid for $R = 0.4$ gives good results for the groomed mass distribution of large- R jets.

Finally, focusing on the results for $R = 0.4$, there is a small region around $a \approx 0.35 - 0.4$, $R_0 \approx 0.2 - 0.25$,¹¹ where the bias is small and the dispersion marginally better than the SoftKiller alone. This means that this approach has a potential to bring a marginal improvement over the SoftKiller method for the standard $R = 0.4$ use-case. A more definitive statement would involve a more complete study of the other parameters of the model and, in particular, an in-depth study of the robustness of the method across a range of pileup conditions and jet p_t .

Conclusions. In the end, we find that the R -dependent extension of the SoftKiller method proposed here brings at best minor improvements compared to the SoftKiller alone. However, it indicates that there might be a region of our parameter space where a single working point could allow clustering with any jet radius, where a plain SoftKiller would require adapting the size parameter a .

There are also some signs indicating that a small improvement could be waiting around the corner, although this requires more systematic tests. In particular, we have found improvements of a few percents for $R = 0.4$, compared to the SoftKiller alone. Further tests could also involve or suggest modifications of our simple central formula, Eq. (13.1). This points again to the need for studies with different processes and different pileup conditions as well as for performance tests with CHS events. Also, given similarities between the performance obtained with the SoftKiller and PUPPI (see also the summary study presented in our conclusions, Chapter 14), with the latter involving some form of angular dependence, one might wonder about the size of the gains that could be obtained from such an approach.

Finally, it would be interesting to investigate whether a toy-model approach like the one discussed in Section 5.2 could help understanding from an analytic perspective the angular dependence of the SoftKiller or its extensions.

13.2.2 Combination with (protected) zeroing

The next extension of the SoftKiller approach that we want to investigate is its combination with (protected) zeroing.

Description. We have tested this using the following method: given a CHS event, we first apply the SoftKiller, with a size parameter a , on it; we then apply protected zeroing i.e. discard each of the remaining neutral particle which have no charged tracks coming from the leading vertex within a distance R_{zero} , unless they have a transverse momentum larger than $p_{t,\text{prot}}$.

We have followed the same testing scheme as for the SoftKiller in the previous Chapter and so refer to Section 12.2 for details of our simulations. The approach described above has three parameters: the grid size a for the SoftKiller that we varied between 0.2 and 0.6, the zeroing radius R_{zero} that we varied between 0.1 and 0.3, and the protection cut $p_{t,\text{prot}}$ that we fixed to 10 GeV.

Results for the jet p_t . On Figs. 13.6 and 13.7 we show the usual performance indicators: the average bias $\langle \Delta p_t \rangle$ of the reconstructed jets after pileup mitigation compared to the original hard jets, Fig. 13.6, and the associated dispersion $\sigma_{\Delta p_t}$, Fig. 13.7. In both cases, the results are shown as a function of the p_t of the hard jet and as a function of the pileup multiplicity. We have selected a few representative values of a and R_{zero} which show good performance. As expected, lowering the zeroing

¹¹That region will also depend on the precise value chosen for $p_{t,\text{cut}}$ and κ .

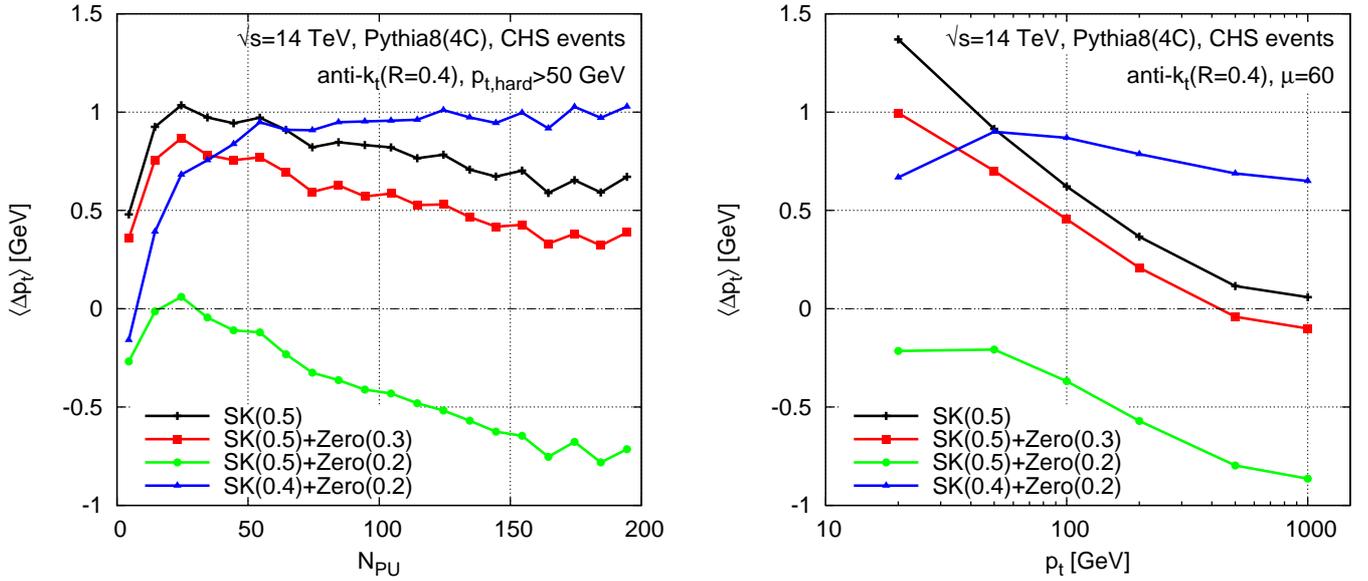


Figure 13.6: Average shift, $\langle \Delta p_t \rangle$, obtained after applying the SoftKiller either alone (black line) or followed by a zeroing step. The shift is shown as a function of the number of pileup vertices for hard jets above 50 GeV on the left plot and as a function of the hard jet p_t for $\mu = 60$ Poisson-distributed pileup events on the right plot.

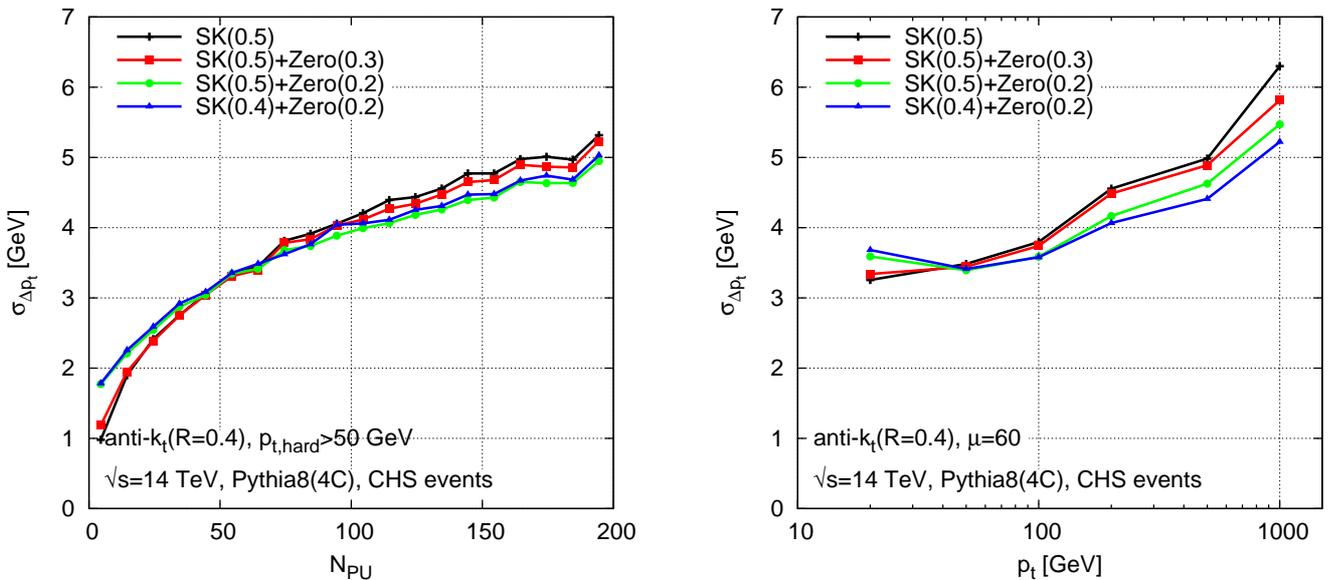


Figure 13.7: Dispersion $\sigma_{\Delta p_t}$ corresponding to the bias from Fig. 13.6.

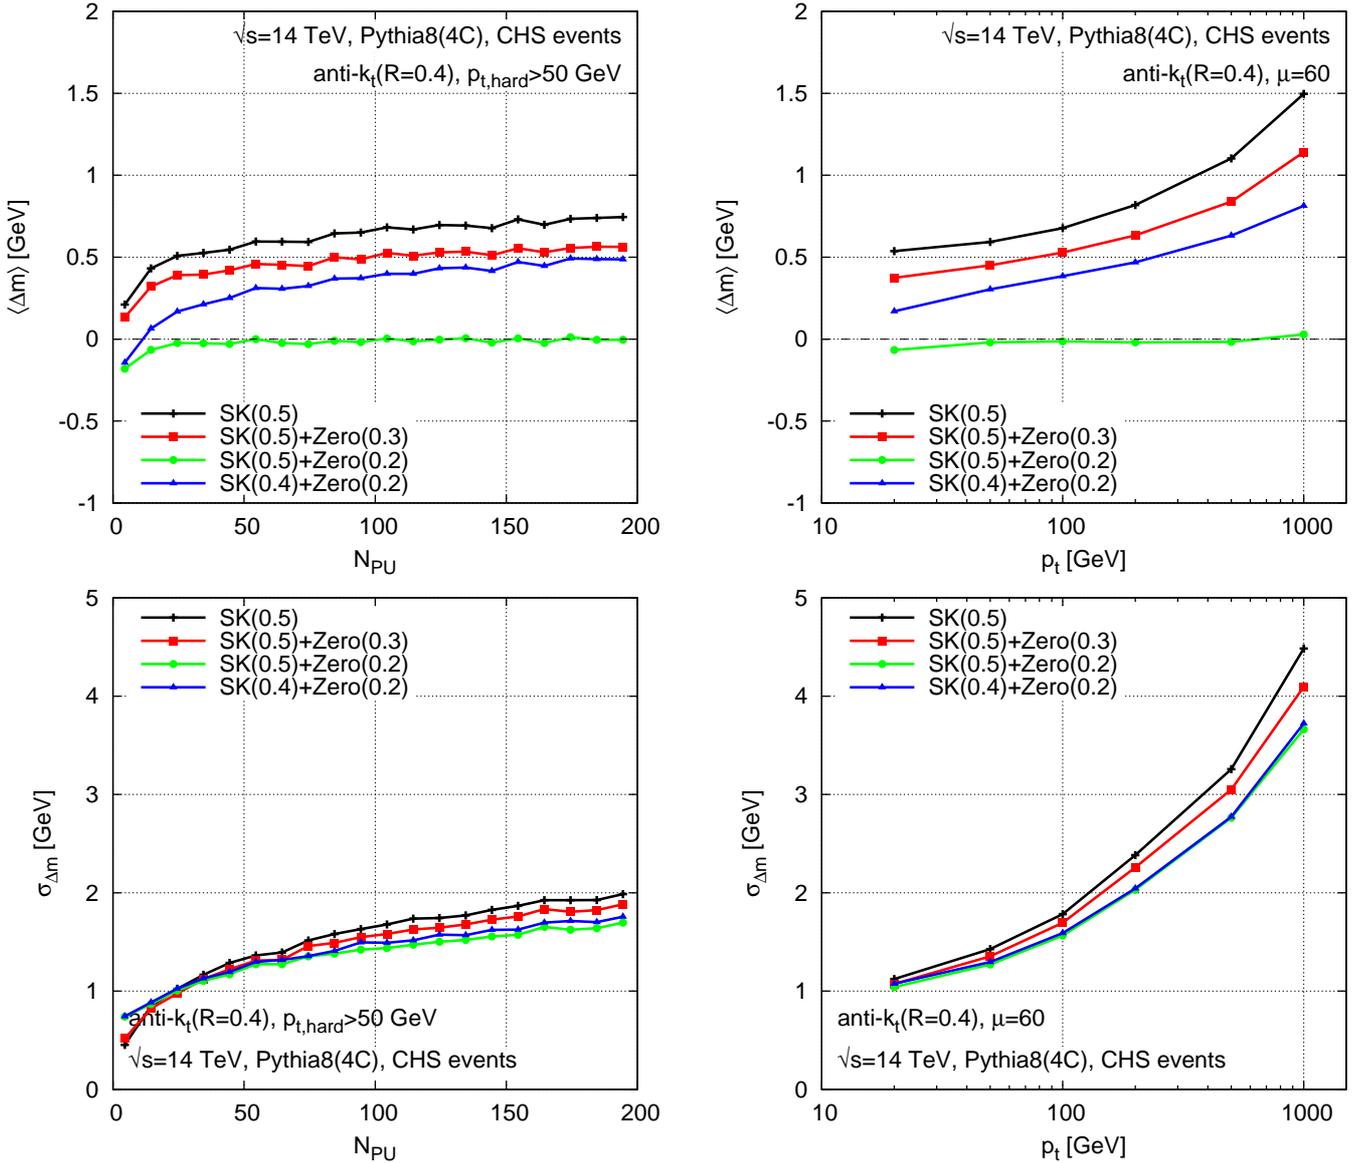


Figure 13.8: Same as Figs. 13.6 and 13.7 now for the jet mass instead of the jet p_t .

radius shifts the bias downwards and a first inspection seems to indicate that values of R_{zero} in the 0.2-0.3 range give the optimal performance. In terms of the average bias, all the combinations shown in the Figure give results of similar quality to, or marginally better than, the SoftKiller alone.

Let us now turn to the dispersion. We see that zeroing gives an improvement in extreme situations: for $a = 0.4$ or 0.5 and $R_{zero} = 0.2$, we gain about 15% in dispersion at large jet p_t and around 8% at large pileup multiplicities. Unfortunately, this improvement is not present uniformly, especially at low p_t or low pileup multiplicity where our results indicate a preference for a larger zeroing radius and only bring marginal improvement, if any, compared to the SoftKiller alone. A more careful tuning of the parameters is probably needed to reach a performance gain over the whole kinematic spectrum.

Results for the jet mass. If one now considers the reconstruction of the jet mass instead of the jet transverse momentum, see Fig. 13.8, one sees a more interesting behaviour. First, the performance gain brought by the extra zeroing step is more manifest. Then, the combination of the SoftKiller

with $a = 0.5$ and zeroing with $R_{\text{zero}} = 0.2$ shows remarkable performance: on average the mass is reconstructed with very little bias over the whole kinematic range we have tested, and the dispersion comes out 15-20% lower than with the SoftKiller alone, without a sign of degradation at low p_t . This is coherent with our earlier observations from Chapter 11.

Conclusions. In the end, this preliminary study indicates that (protected) zeroing can indeed improve the performance of the SoftKiller but that some degree of tuning is needed to fully optimise the method. For example, our studies show that a better performance would be obtained with a larger zeroing radius at low pileup multiplicity and low jet p_t . After that tuning is carefully done, one could expect a performance gain over the SoftKiller alone that reaches up to 15% in terms of dispersion with no significant impact on the average bias. Finally, if one focuses on the reconstruction of the jet mass instead of the jet transverse momentum, the combination of the SoftKiller with $a = 0.5$ and zeroing with $R_{\text{zero}} = 0.2$ shows remarkable performance with nearly negligible bias and nice resolution gains. This would therefore certainly deserve more extensive studies, including tests with more busy environments like $t\bar{t}$ events and tests with a more realistic, experimental, setup.

Chapter 14

Conclusion and perspectives

Pileup, the superposition of simultaneous proton-proton collision, contaminates almost all measurements at the LHC. Being hadronic objects, jets are particularly affected by pileup and one needs to design procedures to subtract/mitigate pileup effects if one wants to reach optimal precise for measurements involving jets. This is the topic of this document. The main quantity of concerns is the widely used jet transverse momentum but we have also discussed other jet properties.

In this concluding Chapter, we will first give a brief overview of the methods that have been covered with their respective advantages and potential issues. To put that overview into perspective, we will provide a benchmark study comparing their performance in a couple of summary plots. We will then spend the last few paragraphs discussing future challenges and perspectives.

A significant fraction (the first part) of this review has concentrated on the area–median pileup subtraction method which has been used in Run I of the LHC and is still extensively in use during Run II. In Chapters 2-5, we have described how the method can be applied to a wide range of jet observables: the jet transverse momentum, the jet mass, jet shapes and even the collinear-unsafe jet fragmentation function; we have then tested its performance and studied some of its analytical properties. A summary of our recommendations for the usage of the area–median method at the LHC can be found in Section 2.8. The main benchmark of the area–median method is that it provides a robust pileup subtraction, with a very small residual bias — of order of a few 100 MeV — and a large reduction of the pileup resolution degradation effects. By “robust”, we mean that the method depends essentially on a single free parameter — the patch size used for the estimation of the pileup density — which can be chosen over a relatively wide range without affecting the overall performance of the method, independently of the process and transverse momentum scale under consideration.

One aspect that we have not discussed at all in this document is detector effects which will certainly require additional jet-calibration corrections.¹ However, the stability of the area–median approach wrt the process and kinematic dependence and wrt the choice of its free parameter, indicates that the residual detector effects should also be small.

Furthermore, as we have shown in Chapter 11, the area–median method can be straightforwardly extended to Charged-Hadron-Subtracted (CHS) events, and is expected to perform slightly better than the Neutral-proportional-to-Charged (NpC) method which uses the (known) charged pileup contamination to the jet to infer their neutral contamination. This comparison could however be affected by several detector effects like the effectiveness of the CHS event reconstruction and out-of-time pileup.

¹One specific issue we have not discussed is the non-linearity coming from the fact that pileup contamination can push a calorimeter tower from below to above the noise threshold.

In the end, the area–median pileup subtraction method can be seen as a robust and efficient reference method for pileup subtraction.

The main limitation of the area–median method is that it leaves some residual resolution degradation after the subtraction has been applied. This smearing remains small with current pileup conditions ($\mu \sim 30 - 50$) — see below for more quantitative information — but will increase significantly in future runs of the LHC ($\mu \sim 60 - 140$). This suggests that one investigates alternative pileup mitigation techniques which are less sensitive to pileup fluctuations between different points within an event. This has been investigated in the third part of our review. We have proposed a series of new methods following two main directions of thought: jet grooming techniques and event-wide pileup subtraction.

Jet grooming techniques are substructure tools extensively used in the context of boosted fat jets at the LHC, and we have briefly studied this in the second part of this document. Our preliminary results for the use of jet grooming as a generic pileup mitigation tool have been presented in Chapter 10 and a discussion of the jet cleansing technique, which uses jet substructure as well, is included in Chapter 11. As far as event-wide techniques are concerned, we have made an extensive study of the SoftKiller in Chapter 13. We have also provided a brief Monte-Carlo study of a series of alternative (and still preliminary) event-wide pileup-mitigation methods, in Chapter 13, together with a few possible (also preliminary) extensions of the SoftKiller method.

These substructure-based or event-wide techniques all share similar generic patterns. As expected, they come with sometimes sizeable gains in jet resolution, especially at large pileup multiplicities. In particular, they often show a dependence on N_{PU} which grows slower than the naive $\sqrt{N_{PU}}$ statistical behaviour. However, this gain in resolution does not come for free: in order to keep the average transverse momentum bias small, one usually needs to fine-tune the free parameters of the method, sometimes including variations of these parameters with the pileup conditions or jet energy. This contrasts with the robustness of the area–median method. Among the methods we have studied, the most stable is the SoftKiller for which a single value of the one free parameter gives acceptable biases independently of the pileup conditions, process and jet energy under consideration. The remaining average bias is around 1-2 GeV, larger than what we obtain from the area–median method but of the order of, or smaller than, the jet-energy-scale corrections applied in an experimental context.

To summarise our main findings and the many pileup mitigation techniques introduced throughout this document, we have selected a few representative methods (see below) and studied them with the framework developed for the “*PileUp WorkShop*” held at CERN in May 2014 [135]. Since the goal of the workshop was precisely to discuss and compare different pileup mitigation techniques using a common open-source software framework, this is perfectly suited for our summary. The code used to obtain the results presented hereafter is publicly available in the `comparisons/review` folder of the <https://github.com/PileupWorkshop/2014PileupWorkshop> GitHub [136] repository.

To avoid a proliferation of curves, we have limited the study presented here to the following representative methods:

- *Area–median*: this is the main method described in this review (see Part I) and should be considered as a baseline for comparisons.
- *Filtering*: this is meant to illustrate how jet grooming techniques can be used as generic pileup subtraction methods (see Chapter 10 for more details and other groomer options).
- *Linear cleansing*: this is the technique introduced in Ref. [10] and discussed in Chapter 11. It also illustrates the use of the Neutral-proportional-to-Charged method, combined with subjet techniques (and zeroing).

- *SoftKiller*: this method discussed in Chapter 12 is an event-wide, noise-reduction, pileup mitigation method.
- *SoftKiller+Zeroing*: this is the extension of the SoftKiller method proposed in Section 13.2.2, supplemented with (protected) zeroing.
- *PUPPI*: this is the recent method proposed in Ref. [113] and used by the CMS collaboration in Run II of the LHC.

The first 3 methods are applied individually on each jet in the full event, the last three methods are applied on the full event prior to the clustering. For all cases, we have assumed massless particles and perfect CHS events where the charged particles can be exactly reconstructed and associated either to the leading vertex or to a pileup vertex. Note that the software available from the GitHub repository includes additional options — typically, filtering with alternative parameters, area-trimming, the ConstituentSubtractor and the subtractions based on the Voronoi particle area introduced in Section 13.1 — which we will not discuss here.

The details of how the simulation and analysis are carried on are described in Appendix G, together with the parameters adopted for each of the pileup mitigation techniques. In the end, we study our usual quality measures: the average bias for the jet transverse momentum and the associated resolution. We also consider the jet mass but this should merely be considered as a cross-check that the pileup mitigation does not severely affects jet substructure measurements.

Results are plotted in Fig. 14.1 for the jet p_t and in Fig. 14.2 for the jet mass. We have used the format advocated during the workshop which shows the dispersion versus the average bias. Each curve corresponds to a different method with the 4 points corresponding to the 4 pileup multiplicities, $N_{\text{PU}} = 30, 60, 100$ and 140 from bottom to top. Each panel of the plot corresponds to a different cut on the jet p_t in the initial sample without pileup. We first discuss the case of the jet p_t and move to the jet mass later on.

First and above all, we clearly see the stability of the area-median subtraction method (the solid red curves with empty circles): the p_t bias remains well below 1 GeV over a wide range of pileup conditions and jet p_t . We also observe many of the other features highlighted in the first part of this document, namely that the dispersion increases like $\sqrt{N_{\text{PU}}}$ (we have, roughly, $\sigma_{\Delta p_t} \approx 0.55\sqrt{N_{\text{PU}}}$ GeV) and that the dispersion starts increasing at large p_t (an effect associated with back-reaction).

Let us now examine to the two other jet-by-jet techniques: filtering, combined with the area-median subtraction for the subjets (the dot-dot-dashed blue line with diamond symbols), and jet cleansing, using the neutral-proportional-to-charged approach (the dotted magenta line with crosses). We see that they typically bring gains in resolution but come with a few caveats: the average bias increases and depends on either the pileup multiplicity (for cleansing) or the jet p_t (for filtering). Furthermore, while the gain in resolution is clear at small transverse momentum and large pileup, arguably where it matters the most, the situation is less clear for large p_t where the more robust area-median approach also gives a better resolution, or at smaller pileup multiplicity and small p_t , where only filtering shows a good performance. This illustrates the fact that, compared to the area-median method, these new techniques require an additional fine-tuning to keep the average bias under control. Some level of fine-tuning (e.g. of the subjet radius used in both filtering and jet cleansing) with N_{PU} and p_t could also lead to a more systematic improvement in jet resolution. Note also that it would be interesting to study a combination of filtering with (protected) zeroing since this has shown performance gains with cleansing.

Finally, let us consider the 3 event-wide techniques we have included in our study: the SoftKiller alone (the long-dashed black line with filled squares) or supplemented with (protected) zeroing (the

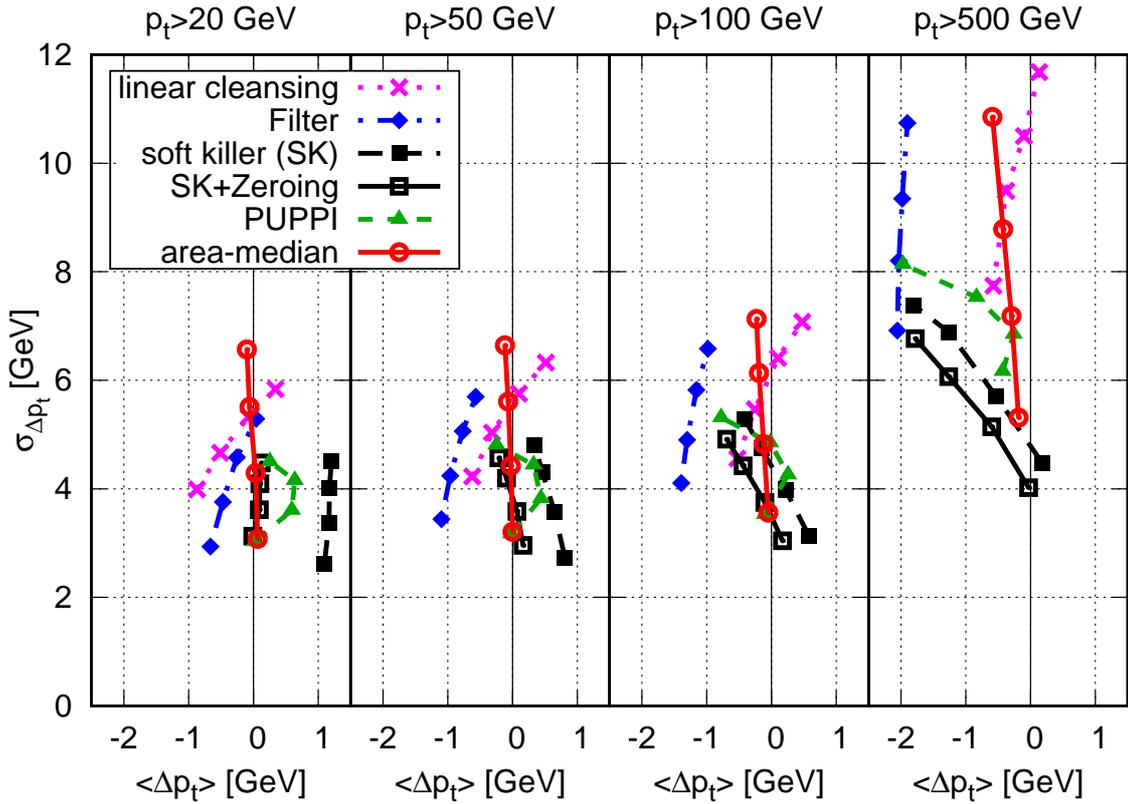


Figure 14.1: Performance summary for 6 representative pileup mitigation techniques. The p_t resolution degradation due to pileup is plotted vs. the average p_t bias for cuts on the jet p_t ranging from 20 GeV (left panel) to 500 GeV (right panel). The 4 points on each curve corresponds to $N_{PU} = 30, 60, 100$ and 140 from bottom to top. All methods are applied on idealistic CHS events.

solid black line with empty squares) and PUPPI (the short-dashed green line with triangles). In terms of the average bias, we see the same trend as with the subjet techniques: a fine-tuning of the free parameters of the method is necessary to keep the bias under reasonable control, reasonable here meaning around 1-2 GeV with some residual dependence on the jet transverse momentum and the pileup conditions. Such a bias can be considered as acceptable since it remains comparable to jet energy scale corrections one would anyway have to apply in an experimental context. Furthermore, the three event-wide approaches show substantial resolution gains, especially at large pileup multiplicities, with only small differences between the methods. At low p_t , PUPPI and SoftKiller combined with zeroing behave very similarly with the SoftKiller alone showing a larger bias but a slightly better resolution for $N_{PU} \sim 30 - 60$. Moving to intermediate and large p_t , the SoftKiller supplemented with zeroing seems to perform slightly better than the other two methods. It would be definitely interesting to see how this comparison behaves in a more realistic experimental context.

Note that it is worth comparing our numbers obtained for the resolution, to the CMS detector performance for the jet energy resolution [137] using particle flow [134] and CHS events. With almost no pileup ($\mu < 10$), CMS reports a jet energy resolution around 3.2 GeV for 20 GeV jets, growing to ~ 5.5 GeV for 50 GeV jets and ~ 10 GeV for 100 GeV jets. This indicates, as expected, that pileup affects more low and intermediate- p_t jets than high- p_t jets. For 20 GeV jets, a pileup resolution degradation of 4 GeV would give about a 50% degradation in jet energy resolution, which is about what is obtained with current (end of 2016) pileup conditions. For larger pileup multiplicities, fig. 14.1

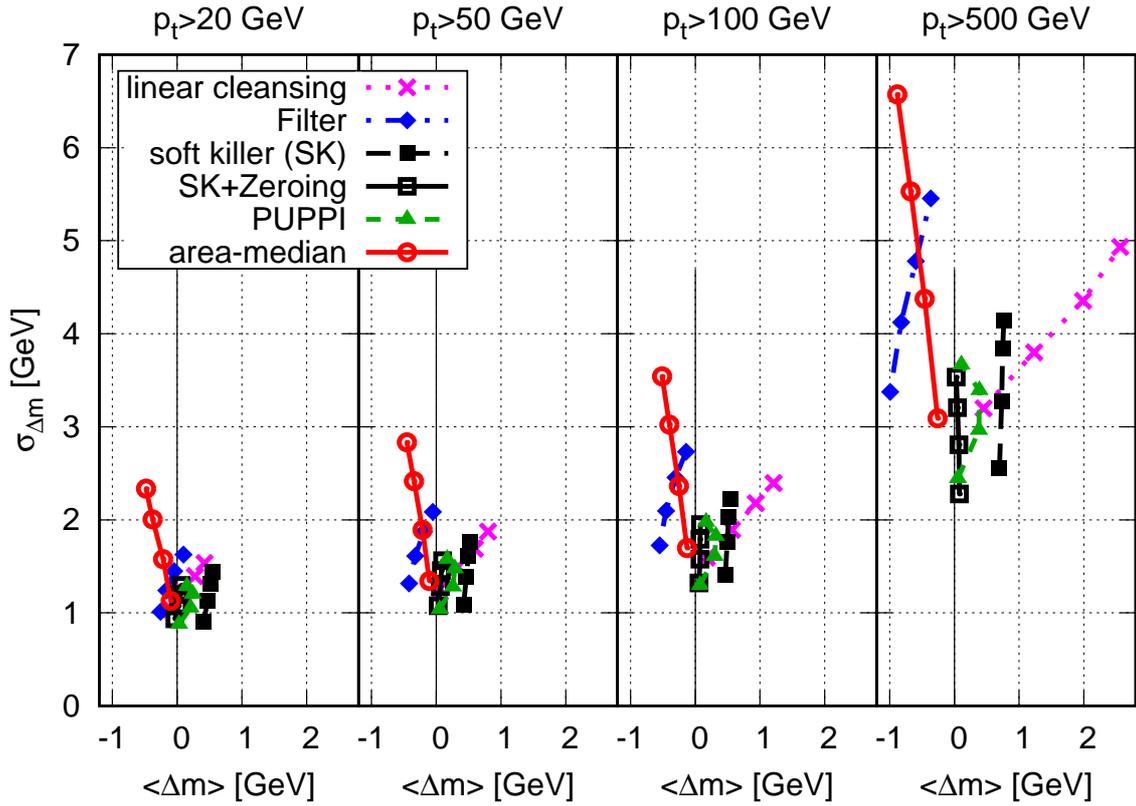


Figure 14.2: Same as Fig. 14.1, now for the jet mass instead of the jet p_t .

clearly points in the direction of using the new event-wide mitigation techniques, like the SoftKiller and its possible improvements, or PUPPI. For those techniques the resolution degradation is much less dramatic. Whether this can be further improved or not remains however an open question.

Turning briefly to the jet mass, Fig. 14.2, we see roughly the same pattern as for the jet p_t , although the resolution gain is larger in the case of the jet mass. In particular, it can already be seen for filtering and cleansing which use subjets, and the combination of SoftKiller and protected zeroing shows a very good performance in terms of resolution with a remarkably small average bias.

This summary study illustrates the main points developed in this document: we see that the area–median does an excellent job at providing a robust method with very small bias and process dependence, and new techniques, especially those based on event-wide subtractions, have a potential for significant resolution improvements in the coming runs of the LHC, at the expense of requiring a fine-tuning of their free parameters and a potentially more involved jet calibration.

We will now conclude with a few last generic remarks and open questions for future investigations.

First, a brief comment about speed: jet clustering (including pileup mitigation) is one of the most computer-expensive tasks in the reconstruction of an event, together with the reconstruction of charged tracks. The fast implementation of sequential algorithms provided by FastJet was clearly a step forwards in that context and recent releases of FastJet brought valuable additional gains in speed. With the increasing LHC luminosity, one might wonder if further gains can be found. In that context, one has to highlight that the SoftKiller can be implemented so as to run extremely fast (see Section 12.4). Given the results of our summary study above, one might wonder if the extension of SoftKiller with protected zeroing is also amenable to such a fast implementation?

Additionally, the many discussions throughout this report have focused essentially on the impact

of pileup on jets. This is the most common application but is far from the only one: pileup affects the reconstruction of all objects in a hadronic collider. The typical other examples are lepton/photon isolation and missing transverse energy (MET) reconstruction. These observables are however more dependent on the details of the detectors and hence more delicate to address using Monte Carlo studies. A few elements in this document point in the direction that a good reconstruction of MET should be more aggressive than what is used for jets. For example, in the case of an area–median-like approach, one should typically consider an extra pileup subtraction to remove the potentially large positive bias from positive fluctuations of the pileup density. Similarly, our SoftKiller studies require a larger grid-size parameter a (and hence a larger associated p_t cut) when the jet radius increases. However, other indices suggest a less aggressive pileup mitigation for MET reconstruction. For example, with an ideal, fully hermetic detector, energy-momentum conservation would guarantee a good reconstruction of MET even without any pileup mitigation.

More generally, we have (almost) not addressed detector effects and have instead tried to concentrate on more analytic aspects of pileup mitigation (see the remarks below). Most definitely, any pileup mitigation technique has to be validated in an experimental context but I believe that it is preferable to leave that delicate task directly to members of the experimental collaborations which are more aware of the associated subtle technical details. In that viewpoint, tests based on Monte Carlo simulations should serve as a solid indication that a method is worth of further investigations. Active discussions between the LHC phenomenologists and the experimental collaborations can only be beneficial to address potential issues or technical details and to ultimately aim at efficient pileup mitigation techniques.

Let me finish this final Chapter with a few remarks on how analytic insight and calculations have been, and can still be, useful in the context of pileup mitigation, a topic that might seem purely experimental at first sight. First, we have seen repeatedly in this document that simple analytic descriptions of the pileup energy deposit often allow one to understand the main features observed in Monte Carlo simulations. Already using a Gaussian approximation is sufficient to grasp many aspects of how pileup affects jets and how the area–median subtraction behaves. It is also interesting to point out that, as pileup multiplicity increases, the Gaussian approximation becomes valid over an increasingly wide range of scales. This fundamental understanding is key to gain confidence in a pileup mitigation method. Identifying physical properties of existing pileup mitigation techniques through a simple analytic understanding can also be of potential help towards developing new, more efficient, methods.

On a different level, analytic developments presented in this review are intimately connected to other aspects of perturbative QCD. The area–median method introduces a new property of jets, namely their area, and we have seen in Chapter 5 that many analytic results can be derived for that quantity.² Our results show rich structures which can be of potential interest for tests of Monte Carlo event generators. Since jet areas are intrinsically non-perturbative quantities, sensitive to soft particles, this is naturally true for the tuning of non-perturbative models of the Underlying Event. Additionally, we have seen that the dependence of the jet area on the jet p_t can be understood from perturbative QCD. This could provide constraints on parton showers, and we have indeed seen some differences between different generators and shower models included in our studies. Even on a more technical ground, jet areas show fundamental properties that are still to be understood. This includes for example the distribution of pure-ghost jets (see Fig. 5.4) or the appearance of fractal jet boundaries (see Section 5.1.2.4).

Another situation where analytic calculations are important is the case of jet grooming, discussed

²Or, to be more precise, “quantities”, since there are several ways to define the area of a jet.

briefly in Chapter 8. This is intimately related to the understanding of jet substructure from first-principles in QCD, an area that is still fairly young and has received quite a lot of interest over the past couple of years.

More generally, I believe that a combination of perturbative QCD understanding of jets and modelling of the properties of pileup — either using a simple Gaussian approximation, more refined toy models like those described in Section 5.2, or directly measurements from data — is one of the most promising way towards developing and controlling new pileup mitigation techniques.

As for jet areas and jet substructure, a first-principles understanding comes with two advantages. First, it can provide a simple picture to explain the main features of a given pileup mitigation method. More importantly, it could help understanding from basic principles how the average bias and resolution behaves when varying pileup conditions, the jet transverse momentum, the jet radius, or the parameters of a given pileup mitigation technique. This is somehow similar to the calculation of the anomalous dimension for jet areas and to our analytic discussions of jet grooming. In a context where the tendency is to develop event-wide pileup mitigation methods with noise reduction, such a first-principles understanding might be the key to keep under control the potentially large biases inherent to these methods. It could therefore largely reduce the need for fine-tuning and improve both the effectiveness and the robustness of future pileup mitigation techniques.

Appendix A

Jet algorithms and definitions

Throughout this document we have mostly considered four jet algorithms: anti- k_t [22], k_t [25], Cambridge/Aachen [28] and SISCone [95]. We provide here a very brief description of how they work.

A.1 Recombination algorithms

The anti- k_t , k_t , and Cambridge/Aachen algorithms cluster the jets via successive recombinations. They can all be seen as special cases of the generalised- k_t algorithm [22].

One first introduces a distance d_{ij} between each pair of particles and a beam distance d_{iB} for each particle:

$$d_{ij} = \min(k_{ti}^{2p}, k_{tj}^{2p}) \frac{\Delta y_{ij}^2 + \Delta \phi_{ij}^2}{R^2}, \quad d_{iB} = k_{ti}^{2p}. \quad (\text{A.1})$$

The free parameters are R , the jet radius, and p . The anti- k_t , k_t , and Cambridge/Aachen algorithms respectively correspond to $p = -1$, 1 and 0 .

One then iteratively repeats the following procedure until no particles are left: at each step, the smallest distance is computed. If it involves a pair of particles, they are recombined using a E -scheme sum of four-momenta (i.e. direct addition of the four-momenta), otherwise the particle is “clustered with the beam”. i.e. called a jet.

A.2 The SISCone algorithm

The SISCone jet algorithm is an infrared- and collinear-safe implementation of a modern cone algorithm. It first finds all stable cones of radius R , a stable cone being a circle in the (y, ϕ) plane such that the E -scheme sum of the momenta inside the cone points in the same direction as the centre of the cone. It then runs a Tevatron run II type [96] split–merge procedure to deal with overlapping stable cones. The stable cones are ordered in \tilde{p}_t , the scalar sum of the p_t of a cone’s constituents, to produce the initial list of protojets. One takes the hardest protojet and finds the next hardest one that overlaps with it (its scaled transverse momentum being $\tilde{p}_{t,j}$). If the shared \tilde{p}_t is larger than $f\tilde{p}_{t,j}$ (f is the split–merge overlap threshold parameter for the algorithm), they are merged, otherwise, the common particles are attributed to the protojet with the closer centre. If no overlaps are found, the protojet is added to the list of jets and removed from the list of protojets. This is repeated until no protojets remain.

Appendix B

Representations of the jet fragmentation

A jet fragmentation function can be defined as the distribution dN_h/dz of the momentum fraction $z = p_{t,h}/(\sum_{i \in \text{jet}} p_{t,i})$, of hadrons in the jet, where $p_{t,h}$ is the transverse momentum of the hadron.

The left-most plot of Figure B.1 shows the typical shape of the z distribution (normalised to the total number of jets N_{jet}), given for anti- k_t ($R = 0.4$) jets with $p_t > 100$ GeV in pp with $\sqrt{s} = 2.76$ TeV. The middle plot shows identical information, represented as a differential distribution in $\xi = \ln 1/z$. The z representation helps visualise the hard region of the FF, while ξ devotes more space to the soft part. Two curves are shown: the solid blue, labelled ‘‘Pythia 6’’, is a pp reference curve obtained from Pythia 6.425 with its virtuality-ordered shower; the dashed green was obtained with the Pyquen program [61, 62], which modifies Pythia showering so as to simulate quenching. We have used it with settings corresponding to 0 – 10% centrality. Its effect on the FF is of the same order of magnitude as the effects seen experimentally [138, 139]. Independently of any question of whether it correctly models the underlying physics, Pyquen provides therefore a useful reference when establishing whether FFs are being reconstructed with sufficient accuracy by some given procedure.

The right-most plot of Fig. B.1 shows moments of the fragmentation functions. The N^{th} moment, M_N , of the fragmentation function is given by the integral

$$M_N = \frac{1}{N_{\text{jet}}} \int_0^1 z^N \frac{dN_h}{dz} dz = \frac{1}{N_{\text{jet}}} \int_0^\infty e^{-N\xi} \frac{dN_h}{d\xi} d\xi. \quad (\text{B.1})$$

In practice, the moments for a single jet can be calculated as

$$M_N^{\text{jet}} = \frac{\sum_{i \in \text{jet}} p_{t,i}^N}{(\sum_{i \in \text{jet}} p_{t,i})^N}, \quad (\text{B.2})$$

where the sum runs over all the jet’s constituents.¹ The results can then be averaged over many jets, so that $M_N = \langle M_N^{\text{jet}} \rangle_{\text{jets}}$. Obviously, M_0 represents the average particle multiplicity in a jet, and M_1 is equal to one by virtue of momentum conservation, provided one measures all hadrons, as we assume here (taking π_0 ’s to be stable). If instead only charged tracks are used in the numerator, then it is clear that M_1 will be significantly below 1. There is another value of N that is of special interest: given a jet spectrum $d\sigma_{\text{jet}}/dp_t$ that falls as p_t^{-n} , the ratio of the inclusive hadron spectrum and inclusive jet spectrum is given by M_{n-1} . Thus, $M_{n-1}^{\text{AA}}/M_{n-1}^{\text{PP}}$ corresponds to the ratio of (charged-)hadron and jet

¹We have chosen to normalise the moments using the scalar sum of the transverse momenta of the constituents, so that $M_1 = 1$. Using instead the transverse component of the jet momentum rather than the scalar sum leads to small violations (of the order of a fraction of one percent, for the jet radius $R = 0.4$ used in this work) of the $M_1 = 1$ relation.

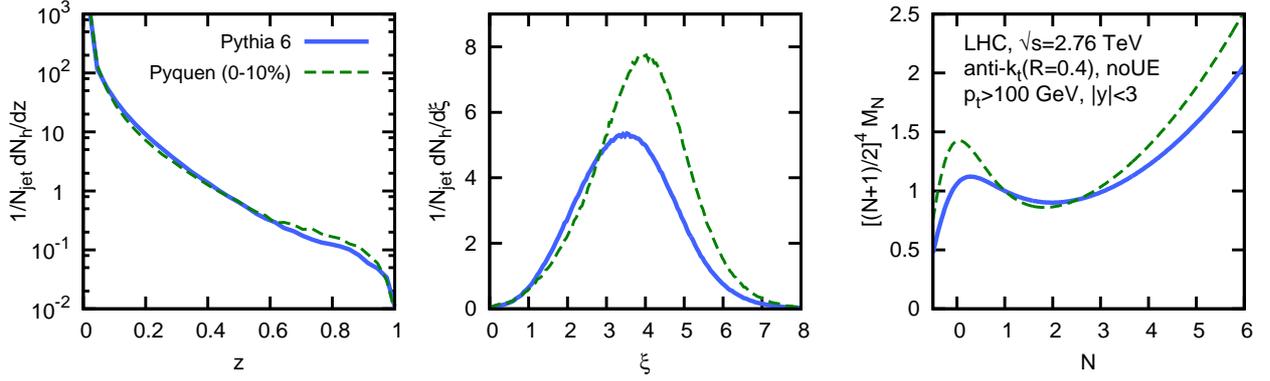


Figure B.1: Jet fragmentation functions (versus z and ξ) and moments (versus N) in pp collisions at the LHC ($\sqrt{s_{NN}} = 2.76$ TeV), obtained without quenching (Pythia 6), and with quenching (Pyquen).

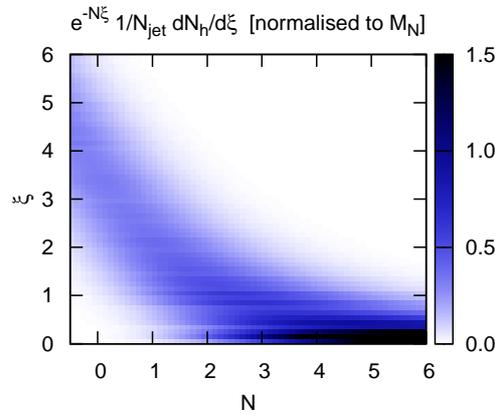


Figure B.2: Representation of the ξ values that contribute dominantly to the M_N integral for a given N , shown as a function of N . Shown for the Pythia 6 results and cuts of Fig. B.1.

R_{AA} values (in the approximation that n is exactly independent of p_t),

$$\frac{M_{n-1}^{AA}}{M_{n-1}^{PP}} = \frac{R_{AA}^h}{R_{AA}^{\text{jet}}}. \quad (\text{B.3})$$

For p_t in the range 100 – 200 GeV, at $\sqrt{s_{NN}} = 2.76$ TeV, the relevant n value has some dependence on p_t and is in the range $n = 6 - 7$, corresponding to $N = 5 - 6$.

In representing the moments in figure B.1(right), we include a factor $((N+1)/2)^4$, which allows a broad range of N values to be shown on a linear vertical scale. The same features visible in the plots versus z and ξ are visible versus N too, for example that Pyquen leads to higher multiplicities than Pythia at small N (corresponding to $z < 0.05$ or $\xi > 3$) and large N ($z > 0.5$, $\xi < 0.7$), and a slightly reduced multiplicity at intermediate N ($z \sim 0.2$, $\xi \sim 1.5$).

To help understand the quantitative relationship between N and ξ , one may examine figure B.2, a colour-map that shows as a function of N the contribution to the M_N moment from each ξ value. It shows clearly how large ξ values dominate for low N (and vice-versa). This ξ, N relationship depends to some extent on the shape of the fragmentation function and it is given for the same Pythia 6 fragmentation function that was used in Fig. B.1.

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

$$\Delta v = v^{\text{sub}} - v^{\text{hard}}, \quad (\text{C.1})$$

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 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 on the offset. In a search for a resonance peak, then one aims for the narrowest peak, and so the smallest possible standard deviation.¹

Another quality measure, advocated for example in [10], is 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. C.1. On the left, one has a scatter plot of the dijet mass in PU-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 plot). Now the correlation coefficient is 0.59, i.e. much worse. This does not reflect a much 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

¹This 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. Other measures of the peak width that are less affected by long tails can be considered, see e.g. the discussion in Section 11.2.2.

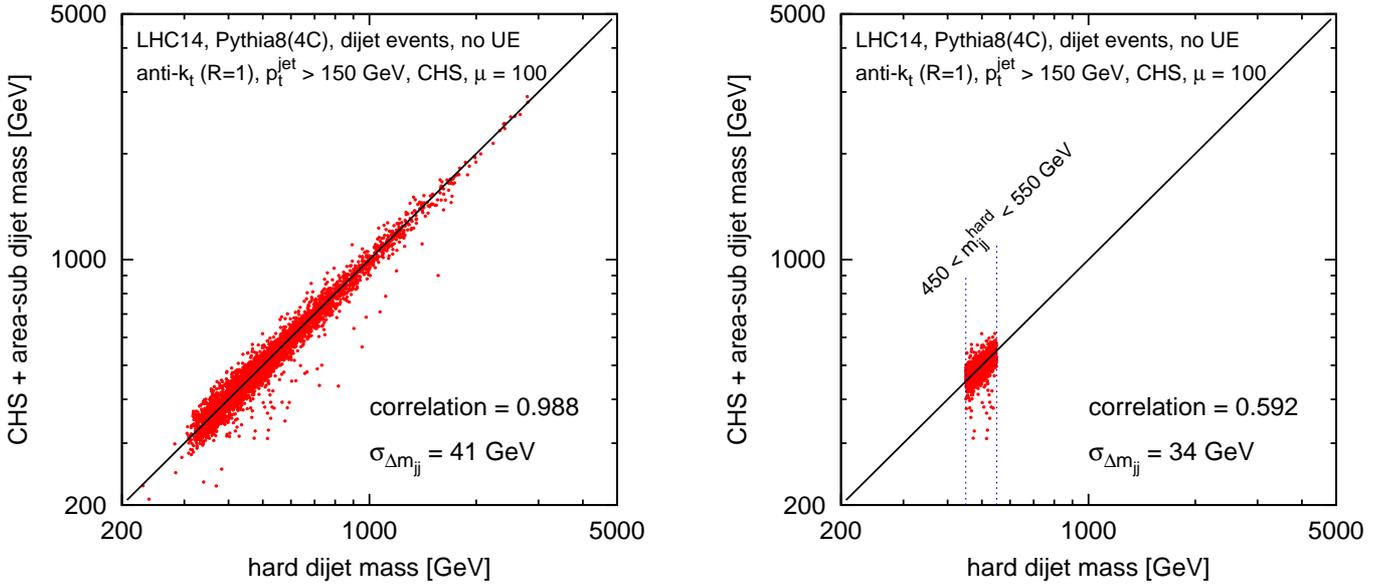


Figure C.1: Left: scatter plot of the dijet mass after addition of an average of 100 pileup events and area–median subtraction (in CHS events) versus the dijet mass in the original hard event. The hard dijet sample and the analysis are as described in appendix E, with a jet radius of $R = 1$. The right-hand plot is identical except for the following additional condition on 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.

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

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

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 behaviour seen in Fig. C.1. 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} (as is actually the case for several practical applications, like dijet mass studies).

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 non-trivial 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.

Appendix D

More analytics results for jet areas

D.1 Transition from one-particle jet to soft jet.

Having observed in Chapter 5 the different properties of the area of pure ghost jets and jets containing a hard particle, one may wonder what happens to the area of a jet containing a “trigger” particle whose transverse momentum p_t is progressively reduced, until it becomes much smaller (ultrasoft) than the momentum scale of a soft background, $\pi R^2 \rho$.¹

Figure D.1 show our findings for the areas obtained with the Cambridge/Aachen algorithm. As one expects, that there are two asymptotic regions. At large p_t , the area tends to that of a jet with an infinitely hard particle, Eq. (5.22). At small p_t , it tends not to the pure ghost jet area, Eq. (5.25), as one might naively expect, but rather to a value which can be predicted as

$$A(\text{ultrasoft-trigger-jet}) = \frac{\int dA A^2 dN/dA}{\int dA A dN/dA}, \quad (\text{D.1})$$

where dN/dA is the distribution of the number of pure ghost jets with a given area, and corresponds to the solid curves depicted in figure 5.4. This equation can be understood in the following way: when the momentum of the trigger particle becomes negligible compared to that of the soft background, it does not influence the size of its own jet. However, the likelihood of the trigger particle being found in a soft-background jet is proportional to that soft-background jet’s area and one gets an extra factor of A in the integrand of the numerator of Eq. (D.1). One is able to use the pure-ghost-jet area distribution in Eq. (D.1), because it coincides with that of soft background jets.

Figure D.1 helps to illustrate the point that a diffuse background of particles such as pileup (represented here by the ghosts) provides an effective cutoff scale, below which a particle will not have any effect on the jet’s area. Numerically, the effective cutoff does indeed coincide with the soft-transverse momentum scale $\pi R^2 \rho$.

D.2 Fluctuations of the active area

Finally, we derive the results Eqs. (5.37)–(5.39) for the fluctuation coefficient $S_{JA,R}^2$ in the case of active areas. This is slightly more technical than for passive areas as we also have to deal with averages over ghosts distributions. Let us briefly recall our notation: $\langle \dots \rangle$ represents an average over perturbative

¹Imagine the trigger particle being a B meson, which you could tag through a secondary vertex — you could then recognise its presence in any jet, regardless of its transverse momentum relative to the background.

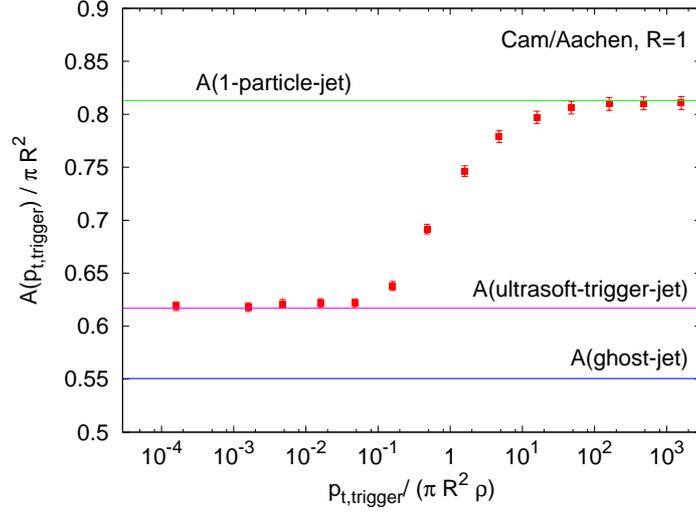


Figure D.1: Area of the jet containing a trigger particle when it is immersed in a bath of soft particles with transverse momentum momentum density ρ ; shown as a function of the trigger particle transverse momentum.

emission, while $\langle \dots \rangle_g$ is an average over ghosts ensembles. When a quantity is evaluated for a specific ghost ensemble $\{g_i\}$, we will explicitly state so with notation of the form $A_{JA,R}(\dots | \{g_i\})$. Otherwise, we have the implicit notation that a quantity specified without any mention of ghosts, such as $A_{JA,R}(0)$ is already averaged over ghost ensembles.

At order α_s , the mean active area can then be written (using $A_{JA,R}(\Delta = 0) = A_{JA,R}(\text{one particle})$)

$$\langle A_{JA,R} \rangle \equiv \langle \langle A_{JA,R}(\dots | \{g_i\}) \rangle \rangle_g \quad (\text{D.2})$$

$$\simeq \left\langle A_{JA,R}(0 | \{g_i\}) + \int dP [A_{JA,R}(\Delta | \{g_i\}) - A_{JA,R}(0 | \{g_i\})] \right\rangle_g \quad (\text{D.3})$$

$$\simeq A_{JA,R}(0) + D_{JA,R} \frac{C_1}{\pi b_0} \log \left(\frac{\alpha_s(Q_0)}{\alpha_s(Rp_{t,1})} \right), \quad (\text{D.4})$$

with

$$D_{JA,R} = \int \frac{d\Delta}{\Delta} (\langle A_{JA,R}(\Delta | \{g_i\}) \rangle_g - \langle A_{JA,R}(0 | \{g_i\}) \rangle_g). \quad (\text{D.5})$$

Here, we have taken into account both the corrections due to the radiation of a soft particle and the average over the distribution of the ghosts. Note that we have used the shorthand dP to represent Eq. (5.9), $dP/(dp_{t2}d\Delta)$, times $dp_{t2}d\Delta$, and that “ \dots ” in Eq. (D.2) represents all possible perturbative states.

For the corresponding fluctuations, the derivation goes along the same line

$$\langle \Sigma_{JA,R}^2 \rangle = \langle \langle A_{JA,R}^2(\dots | \{g_i\}) \rangle \rangle_g - \langle A_{JA,R} \rangle^2 \quad (\text{D.6})$$

$$\begin{aligned} &= \left\langle A_{JA,R}^2(0 | \{g_i\}) + \int dP [A_{JA,R}^2(\Delta | \{g_i\}) - A_{JA,R}^2(0 | \{g_i\})] \right\rangle_g \\ &\quad - \left\langle A_{JA,R}(0 | \{g_i\}) + \int dP [A_{JA,R}(\Delta | \{g_i\}) - A_{JA,R}(0 | \{g_i\})] \right\rangle_g^2. \end{aligned} \quad (\text{D.7})$$

Neglecting the terms proportional to α_s^2 , we can write $\langle \Sigma_{JA,R}^2 \rangle = \Sigma_{JA,R}^2(0) + \langle \Delta \Sigma_{JA,R}^2 \rangle$, where

$$\Sigma_{JA,R}^2(\Delta) = \langle A_{JA,R}^2(\Delta | \{g_i\}) \rangle_g - \langle A_{JA,R}(\Delta | \{g_i\}) \rangle_g^2, \quad (\text{D.8})$$

which for $\Delta = 0$ is the leading order result. We can also write

$$\langle \Delta \Sigma_{\text{JA},R}^2 \rangle = S_{\text{JA},R}^2 \frac{C_1}{\pi b_0} \log \left(\frac{\alpha_s(Q_0)}{\alpha_s(Rp_{t,1})} \right). \quad (\text{D.9})$$

Using straightforward algebra one obtains

$$S_{\text{JA},R}^2 = \int \frac{d\Delta}{\Delta} \left\{ \langle [A_{\text{JA},R}^2(\Delta | \{g_i\}) - A_{\text{JA},R}^2(0 | \{g_i\})] \rangle_g \right. \\ \left. - 2 \langle A_{\text{JA},R}(0 | \{g_i\}) \rangle_g \langle [A_{\text{JA},R}(\Delta | \{g_i\}) - A_{\text{JA},R}(0 | \{g_i\})] \rangle_g \right\} \quad (\text{D.10})$$

$$= \int \frac{d\Delta}{\Delta} \langle [A_{\text{JA},R}^2(\Delta | \{g_i\}) - A_{\text{JA},R}^2(0 | \{g_i\})] \rangle_g - 2 \langle A_{\text{JA},R}(0 | \{g_i\}) \rangle_g D_{\text{JA},R} \quad (\text{D.11})$$

$$= \int \frac{d\Delta}{\Delta} \left[\langle A_{\text{JA},R}(\Delta | \{g_i\}) - A_{\text{JA},R}(0 | \{g_i\}) \rangle_g^2 + \Sigma_{\text{JA},R}^2(\Delta) - \Sigma_{\text{JA},R}^2(0) \right]. \quad (\text{D.12})$$

The second equality is a direct rewriting of the first. One gets to the last line by rearranging the different terms of the first one. The last two lines correspond exactly to Eqs. (5.39) and (5.38), respectively.

Appendix E

Details of our NpC and cleansing study

Let us first fully specify what we have done in our study and then comment on (possible) differences relative to the KLSW study from Ref. [10].

Our hard event sample consists of dijet events from pp collisions at $\sqrt{s} = 14$ TeV, simulated with Pythia 8.176 [36], 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 plots presented in Figs. 11.4 and 11.5, where we use Z' events with $m_{Z'} = 500$ GeV. Jets are reconstructed with the anti- k_t algorithm [22] after making all particles massless (preserving their rapidity) and keeping only particles with $|y| < 4$. We have $R = 0.4$, except for the some of the results presented in Section 11.2.2, where we use $R = 1$ as in Ref. [10].

Given a hard event, we select all the jets with $p_t > 150$ GeV and $|y| < 2.5$. We then add pileup and cluster the resulting full event without imposing any p_t or rapidity cut on the resulting jets. For each jet selected in the hard event, 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. 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. 11.4) 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.¹ 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.²

Most of the studies shown in this paper use idealised 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 backreaction (see Section 5.1.3). The jet selection and matching procedures are independent of the use of CHS or full events. As throughout this document, results plotted as a

¹In the case of the Z' events used for Fig. 11.4, 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 $|\Delta y| < 1.2$ [140, 141]. This ensures that high dijet-mass events are not dominated by low p_t forward-backward jet pairs, which are usually enhanced in QCD v. 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 $R = 0.6$ [140], 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”) [141]. This too can affect conclusions about pileup.

²One alternative would have been to impose the cuts on the jets in the full event (after 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. [10]. This can give overly optimistic results because it neglects backreaction.

function of N_{PU} use a fixed number of zero-bias events, while results shown for a fixed value of μ use Poisson-distributed zero-bias events with an average μ . Clustering and area determination are performed with a development version of FastJet 3.1 and with FastJet 3.1.0 and 3.1.1 for the results in section 11.2.2. They all behave identically to the current 3.2 series for the features used here.

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 , 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 rapidity dependence of ρ using a rescaling as discussed in Section 2.4.2. We carry out 4-vector subtraction, Eq. (2.24).

We additionally avoid unphysical situations as described in Sections 2.4.4, 2.4.5 and 2.7.2.4, with *known selectors* specified to identify the charged particles and their vertex of origin. In particular, we enabled the `safe_mass` option for both the area–median and NpC subtractions.

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 towers. In our original (v1) 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.³

Cleansing has two options: one can give it jets clustered from the full event, and then it uses an analogue of Eq. (11.1): 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. (11.3), 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 back-reaction. 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, 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.⁴ For comparisons with cleansing we generally use $f_{\text{cut}} = 0$ unless explicitly indicated otherwise.

³We choose to show particle-level results here because of the difficulty of correctly simulating a full detector, especially given the non-linearities of responses to pileup and the subtleties of particle flow reconstruction in the presence of real detector fluctuations.

⁴And is the default behaviour 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.

Appendix F

SoftKiller collinear safety issued

Collinear safety is normally essential in order to get reliable results from perturbation theory. One reaction to the SoftKiller proposal is that it is not collinear safe, because it relies only on information about individual particles' transverse momenta. There are at least two perspectives on why this is not a severe issue.

The first relates to the intrinsic low- p_t nature of the p_t^{cut} , which is typically of order 1 – 2 GeV. At these scales, non-perturbative dynamics effectively regulates the collinear divergence. Consider one element of the hadronisation process, namely resonance decay, specifically $\rho \rightarrow \pi\pi$: if the ρ has a p_t of order 2 GeV, the rapidity-azimuth separation of the two pions is of the order of 0.7 – 1 (see e.g. Section 11.2.1). Alternatively, consider the emission from a high-energy parton of a gluon with a p_t of the order of 1 GeV: this gluon can only be considered perturbative if its transverse momentum relative to the emitter is at least of order a GeV, i.e. if it has an angle relative to the emitter of order 1. Both these examples illustrate that the collinear divergence that is of concern at parton level is smeared by non-perturbative effects when considering low- p_t particles. Furthermore, the impact of these effects on the jet p_t will remain of the order of p_t^{cut} , i.e. power-suppressed with respect to the scale of hard physics.

The second perspective is from the strict point of view of perturbative calculations. One would not normally apply a pileup reduction mechanism in such a context. But it is conceivable that one might wish to *define* the final state such that it always includes a pileup and underlying event (UE) removal procedure.¹ Then one should understand the consequences of applying the method at parton level. Considering patches of size $0.5 \times \pi/6$ and particles with $|y| < 2.5$, there are a total of 120 patches; only when the perturbative calculation has at least 60 particles, i.e. attains order α_s^{60} , can p_t^{cut} be non-zero; so the collinear safety issue would enter at an inconceivably high order, and all practical fixed-order parton-level calculations would give results that are unaffected by the procedure.

Collinear safety, as well as being important from a theoretical point of view, also has experimental relevance: for example, depending on its exact position, a particle may shower predominantly into one calorimeter tower or into two. Collinear safety helps ensure that results are independent of these details. While we carried out basic detector simulations in section 12.3, a complete study of the impact of this type of effect would require full simulation and actual experimental reconstruction methods (e.g. particle flow or topoclustering).

¹For example, so as to reduce prediction and reconstruction uncertainties related to the modelling of the UE (we are grateful to Leif Lönnblad for discussions on this subject). This might, just, be feasible with area–median subtraction, with its small biases, but for the larger biases of SK does not seem phenomenologically compelling. Still, it is interesting to explore the principle of the question.

Appendix G

Details of the summary simulations

We give here the details of the analysis presented as a summary study in Chapter 14.

Framework. All studies have been performed using the framework developed for the Pileup Workshop [135] held at CERN in May 2014. The framework is publicly available from the GitHub repository <https://github.com/PileupWorkshop/2014PileupWorkshop> which includes the code used to obtain the results presented in this document in the `comparisons/review` folder (as of December 15 2016).

Hard event samples. We have used the “dijetset1” set of event samples available from <http://puws2014.web.cern.ch/puws2014/events>, corresponding to dijet events with at least one reconstructed anti- k_t ($R = 0.4$) jet satisfying $p_t \geq 20, 50, 100$ or 500 GeV. Each event sample contains 100 000 events. The events have been simulated with Pythia8 (v8.185) with tune 4C. B hadrons have been kept stable and the Underlying Event has been switched off. The last point avoids bringing into the discussion extra complications related to the subtraction of the Underlying Event.

Pileup simulation. Pileup is simulated here as the superposition of a fixed number N_{PU} of minimum bias events simulated using the Pythia8 event generator (v8.185) with tune 4C, at $\sqrt{s} = 14$ TeV. In practice, we have used the sample of 10 million zero-bias events also available from <http://puws2014.web.cern.ch/puws2014/events>.

We have studied four pileup multiplicities: $N_{\text{PU}} = 30, 60, 100$ or 140 , ranging between conditions obtained at the end of Run I and the beginning of Run II, to multiplicities expected for HL-LHC.

Event and jet selections. To avoid potential difficulties related to the inclusion of the ρ_m term, we take particles to be massless, preserving rapidity and azimuthal angle, We also work with idealised CHS events, where charged tracks can be exactly reconstructed and their vertex of origin exactly identified. Charged tracks from pileup vertices are kept in the analysed event, with their transverse momentum scaled down by a factor 10^{-60} . We consider the hard event alone, as well as “full” events which include both the hard event and pileup. In both cases, we keep particles with $|y| < 4$.

We then build a set of reference jets by clustering the hard event using the anti- k_t algorithm with $R = 0.4$ and selecting jets with $p_t \geq p_{t,\text{min}}$ and $|y| < 2.5$.¹ The $p_{t,\text{min}}$ cut is taken to match that of the

¹We could also apply the selection cuts on the jets obtained from the full event, after pileup mitigation. This would however come with additional complications related to the selection bias inherent to this approach.

initial jet sample being studied, i.e. 20, 50, 100 or 500 GeV. For each pileup mitigation technique, we then reconstruct and subtract the jets: for methods which work on each individual jets, we first cluster the jets and then subtract each jet individually, while for event-wide techniques, we first subtract the whole event and then cluster the resulting subtracted event. In both cases, the clustering is also done using the anti- k_t algorithm with $R = 0.4$ and no selection cuts are applied.

For each jet in the full event, we then search for the closest selected jet in the hard event, requiring that both are separated by $\Delta R = \sqrt{\Delta y^2 + \Delta \phi^2} < 0.3$. We have checked that the resulting matching efficiencies are close to one, generally well above 99.9%, with some methods showing a small decrease to $\sim 99.5\%$ for $p_{t,\min} = 20$ GeV and $N_{\text{PU}} = 140$. For matching pairs of jets, $(j_{\text{hard}}, j_{\text{full,sub}})$, we measure the bias in transverse momentum, $\Delta p_t = p_{t,\text{full,sub}} - p_{t,\text{hard}}$ and mass, $\Delta m = p_{t,\text{full,sub}} - p_{t,\text{hard}}$. We study the our usual quality measures, namely the average bias, $\langle \Delta p_t \rangle$ and the associated dispersion, $\sigma_{\Delta p_t}$, computed over the full jets sample.²

Finally, unless explicitly mentioned otherwise, all jet manipulations are done using FastJet (v3.2.1).

Pileup mitigation methods. For the methods under consideration in this study, we use the following configurations and parameters:

- *Area–median*: we use a grid-based estimation of ρ with a grid-size parameter of 0.55 and rapidity rescaling according to the profile $f(y) = 0.9084412 - 0.0189777 y^2 + 4.05981 10^{-5} y^4$. We have enabled the protection against negative masses after subtraction. We use the full CHS event for the estimation of ρ . For the computation of the jet area, the ghosts are placed up to the edge of the particle rapidity acceptance, $|y| < 4$, with a ghost area of 0.01.
- *Filtering*: for each jet in the full event, we recluster the jet constituents into subjets using the Cambridge/Aachen algorithm with $R_{\text{filt}} = 0.2$. We then apply the area–median subtraction described above and keep the $n_{\text{filt}} = 3$ hardest resulting subjets.
- *Linear cleansing*: we use the linear cleansing variant with the `input_nc_separate` mode. We take $\gamma_0 = 0.612$ and applied trimming with $R_{\text{trim}} = 0.2$ and $f_{\text{trim}} = 0$, as suggested by the authors of Ref. [10] during the PileUp WorkShop [135].
- *SoftKiller*: we have used a grid-size parameter of 0.5, with the grid extending to the particle acceptance, $|y| < 4$. We apply the SoftKiller only on the neutral particles in the event, leaving the charged particles from the leading vertex untouched.
- *SoftKiller+Zeroing*: here, we first apply a SoftKiller (as described above but now with a grid-size parameter of 0.45) to the neutral particles in the event. We then apply protected zeroing: we remove neutral particles with $p_t < 10$ GeV for which there are no charged tracks within a radius $R_{\text{zero}} = 0.2$ of the neutral particle.
- *PUPPI*: we have used the implementation provided by the PUPPI authors in the context of the PileUp WorkShop [135]. The code has initially been taken from revision 158 of the GitHub repository (`puppiContainer` from the `example-puppi` directory) with the following adaptation meant to retain the user-defined information of the jet constituents without affecting any of the PUPPI results: for a particle `PseudoJet` \mathbf{p} with 4-momentum (p_x, p_y, p_z, E) and PUPPI weight w , we set the resulting subtracted `PseudoJet` to $w*\mathbf{p}$ instead of (wp_x, wp_y, wp_z, wE) .

²In practice, the code also computes the correlation coefficient between p_t (or the mass) of the hard and full jets. Since we have argued in Chapter 11 and Appendix C that correlation coefficients are more delicate to interpret, we will not use them in the results presented here.

Bibliography

- [1] See e.g. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>
- [2] O. Brning and L. Rossi, *The High Luminosity Large Hadron Collider : the new machine for illuminating the mysteries of Universe*, doi:10.1142/9581#t=toc
- [3] L. Rossi and O. Brning, *Introduction to the HL-LHC Project*, doi:10.1142/9789814675475_0001
- [4] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, Phys. Lett. B **659** (2008) 119 [[arXiv:0707.1378](https://arxiv.org/abs/0707.1378)] [hep-ph].
- [5] M. Cacciari, J. Rojo, G. P. Salam and G. Soyez, *Jet Reconstruction in Heavy Ion Collisions*, Eur. Phys. J. C **71** (2011) 1539; [[arXiv:1010.1759](https://arxiv.org/abs/1010.1759)] [hep-ph].
- [6] J. Alcaraz Maestre *et al.* [SM and NLO MULTILEG Working Group and SM MC Working Group Collaborations], *The SM and NLO Multileg and SM MC Working Groups: Summary Report*, [arXiv:1203.6803](https://arxiv.org/abs/1203.6803) [hep-ph].
- [7] M. Cacciari, P. Quiroga-Arias, G. P. Salam and G. Soyez, *Jet Fragmentation Function Moments in Heavy Ion Collisions*, Eur. Phys. J. C **73** (2013) no.3, 2319 [[arXiv:1209.6086](https://arxiv.org/abs/1209.6086)] [hep-ph].
- [8] G. Soyez, G. P. Salam, J. Kim, S. Dutta and M. Cacciari, *Pileup subtraction for jet shapes*, Phys. Rev. Lett. **110** (2013) no.16, 162001 [[arXiv:1211.2811](https://arxiv.org/abs/1211.2811)] [hep-ph].
- [9] M. Cacciari, G. P. Salam and G. Soyez, *Use of charged-track information to subtract neutral pileup*, Phys. Rev. D **92** (2015) no.1, 014003 [[arXiv:1404.7353](https://arxiv.org/abs/1404.7353)] [hep-ph].
- [10] D. Krohn, M. D. Schwartz, M. Low and L. T. Wang, *Jet Cleansing: Pileup Removal at High Luminosity*, Phys. Rev. D **90** (2014) no.6, 065020 [[arXiv:1309.4777](https://arxiv.org/abs/1309.4777)] [hep-ph].
- [11] M. Cacciari, G. P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, Eur. Phys. J. C **75** (2015) no.2, 59 [[arXiv:1407.0408](https://arxiv.org/abs/1407.0408)] [hep-ph].
- [12] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup Per Particle Identification*, JHEP **1410** (2014) 059 [[arXiv:1407.6013](https://arxiv.org/abs/1407.6013)] [hep-ph].
- [13] M. Cacciari, G. P. Salam and G. Soyez, *The Catchment Area of Jets*, JHEP **0804** (2008) 005 [[arXiv:0802.1188](https://arxiv.org/abs/0802.1188)] [hep-ph].
- [14] M. Cacciari and G. P. Salam, *Dispelling the N^{*3} myth for the $k(t)$ jet-finder*, Phys. Lett. B **641** (2006) 57 [[hep-ph/0512210](https://arxiv.org/abs/hep-ph/0512210)].

- [15] M. Cacciari, G. P. Salam and G. Soyez, *FastJet user manual*, Eur. Phys. J. C **72** (2012) 1896 [[arXiv:1111.6097](#) [hep-ph]], <http://www.fastjet.fr>.
- [16] G. Aad *et al.* [ATLAS Collaboration], *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, [arXiv:1510.03823](#) [hep-ex].
- [17] S. Chatrchyan *et al.* [CMS Collaboration], *Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS*, JINST **6** (2011) P11002 [[arXiv:1107.4277](#) [physics.ins-det]]; CMS Collaboration, CMS-DP-2012-012, <http://cdsweb.cern.ch/record/1460989>.
- [18] G. P. Salam, *Towards Jetography*, Eur. Phys. J. C **67** (2010) 637 [[arXiv:0906.1833](#) [hep-ph]].
- [19] M. Cacciari, G. P. Salam and S. Sapeta, *On the characterisation of the underlying event*, JHEP **1004** (2010) 065 [[arXiv:0912.4926](#) [hep-ph]].
- [20] A. Altheimer *et al.*, *Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd-27th of July 2012*, Eur. Phys. J. C **74** (2014) no.3, 2792 [[arXiv:1311.2708](#) [hep-ex]].
- [21] A. J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, JHEP **1405** (2014) 146 [[arXiv:1402.2657](#) [hep-ph]].
- [22] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, JHEP **0804** (2008) 063 [[arXiv:0802.1189](#) [hep-ph]].
- [23] D. Y. Grigoriev, E. Jankowski and F. V. Tkachov, *Towards a standard jet definition*, Phys. Rev. Lett. **91**, 061801 (2003) [[hep-ph/0301185](#)]; *Optimal jet finder*, Comput. Phys. Commun. **155**, 42 (2003) [[hep-ph/0301226](#)].
- [24] L. Lonnblad, *Arclus: A New Jet Clustering Algorithm Inspired By The Color Dipole Model*, Z. Phys. C **58** (1993) 471.
- [25] S. Catani, Y. L. Dokshitzer, M. H. Seymour and B. R. Webber, *Longitudinally invariant k_t clustering algorithms for hadron hadron collisions*, Nucl. Phys. B **406** (1993) 187 and refs. therein; S. D. Ellis and D. E. Soper, *Successive Combination Jet Algorithm For Hadron Collisions*, Phys. Rev. D **48** (1993) 3160 [[hep-ph/9305266](#)].
- [26] S. J. Fortune, *Algorithmica* (1987) 153–174.
- [27] G. Aad *et al.* [Atlas Collaboration], *Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{S(NN)} = 2.76$ TeV with the ATLAS Detector at the LHC*, Phys. Rev. Lett. **105** (2010) 252303 [[arXiv:1011.6182](#) [hep-ex]].
- [28] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better jet clustering algorithms*, JHEP **9708**, 001 (1997) [[hep-ph/9707323](#)]; M. Wobisch and T. Wengler, *Hadronization corrections to jet cross sections in deep-inelastic scattering*, [hep-ph/9907280](#).
- [29] [CMS Collaboration], *Measurement of the Underlying Event Activity with the Jet Area/Median Approach at 0.9 TeV*, [CMS-PAS-QCD-10-005](#).

- [30] S. S. Adler *et al.* [PHENIX Collaboration], *Elliptic flow of identified hadrons in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. Lett. **91** (2003) 182301 [[nucl-ex/0305013](#)]; J. Adams *et al.* [STAR Collaboration], *Particle dependence of azimuthal anisotropy and nuclear modification of particle production at moderate p_t in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. Lett. **92** (2004) 052302 [[nucl-ex/0306007](#)]; C. Alt *et al.* [NA49 Collaboration], *Directed and elliptic flow of charged pions and protons in Pb + Pb collisions at 40-A-GeV and 158-A-GeV*, Phys. Rev. C **68** (2003) 034903 [[nucl-ex/0303001](#)]; B. B. Back *et al.* [PHOBOS Collaboration], *Centrality and pseudorapidity dependence of elliptic flow for charged hadrons in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **72** (2005) 051901 [[nucl-ex/0407012](#)].
- [31] J. Y. Ollitrault, *Anisotropy As A Signature Of Transverse Collective Flow*, Phys. Rev. D **46** (1992) 229; A. M. Poskanzer and S. A. Voloshin, *Methods for analyzing anisotropic flow in relativistic nuclear collisions*, Phys. Rev. C **58** (1998) 1671 [[nucl-ex/9805001](#)].
- [32] S. Chatrchyan *et al.* [CMS Collaboration], *Measurement of jet fragmentation into charged particles in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, [arXiv:1205.5872](#) [[nucl-ex](#)].
- [33] G. Aad *et al.* [ATLAS Collaboration], *Measurement of inclusive jet charged-particle fragmentation functions in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector*, Phys. Lett. B **739** (2014) 320 [[arXiv:1406.2979](#)] [[hep-ex](#)].
- [34] S. Chatrchyan *et al.* [CMS Collaboration], *Observation and studies of jet quenching in PbPb collisions at nucleon-nucleon center-of-mass energy = 2.76 TeV*, Phys. Rev. C **84** (2011) 024906 [[arXiv:1102.1957](#)] [[nucl-ex](#)].
- [35] T. Sjostrand *et al.*, *High-energy-physics event generation with PYTHIA 6.1*, Comput. Phys. Commun. **135** (2001) 238 [[hep-ph/0010017](#)]; T. Sjostrand *et al.*, *PYTHIA 6.3: Physics and manual*, [hep-ph/0308153](#); T. Sjostrand, S. Mrenna and P. Skands, *PYTHIA 6.4 Physics and Manual*, JHEP **0605** (2006) 026 [[hep-ph/0603175](#)].
- [36] T. Sjostrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178** (2008) 852 [[arXiv:0710.3820](#)] [[hep-ph](#)].
- [37] R. Corke and T. Sjostrand, *Interleaved Parton Showers and Tuning Prospects*, JHEP **1103** (2011) 032 [[arXiv:1011.1759](#)] [[hep-ph](#)].
- [38] A. Karneyeu, L. Mijovic, S. Prestel and P. Z. Skands, *MCPLOTS: a particle physics resource based on volunteer computing*, Eur. Phys. J. C **74** (2014) 2714 [[arXiv:1306.3436](#)] [[hep-ph](#)]; see also <http://mcplots.cern.ch/>, as well as an alternative comparison, <http://rivet.hepforge.org/tunecmp/index.html>.
- [39] P. M. Jacobs [STAR Collaboration], *Background Fluctuations in Heavy Ion Jet Reconstruction*, Nucl. Phys. A **855** (2011) 299 [[arXiv:1012.2406](#)] [[nucl-ex](#)].
- [40] P. Z. Skands, *Tuning Monte Carlo Generators: The Perugia Tunes*, Phys. Rev. D **82** (2010) 074018 [[arXiv:1005.3457](#)] [[hep-ph](#)].
- [41] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour and B. R. Webber, JHEP **0101** (2001) 010 [[hep-ph/0011363](#)]; Herwig 6.5 release note, [hep-ph/0210213](#).

- [42] M. Bahr *et al.*, Eur. Phys. J. C **58** (2008) 639 [[arXiv:0803.0883](#) [hep-ph]]; S. Gieseke *et al.*, *Herwig++ 2.5 Release Note*, [arXiv:1102.1672](#) [hep-ph].
- [43] For the ATLAS collaboration, the jet resolution can be extracted from the first plot from <http://twiki.cern.ch/twiki/pub/AtlasPublic/JetResolutionPreliminaryResults>, as extracted on 13 April 2011.
- [44] For the CMS collaboration, see e.g. [CMS PAS JME-10-014](#).
- [45] C. F. Berger, T. Kucs and G. F. Sterman, *Event shape / energy flow correlations*, Phys. Rev. D **68** (2003) 014012 [[hep-ph/0303051](#)].
- [46] L. G. Almeida, S. J. Lee, G. Perez, G. F. Sterman, I. Sung and J. Virzi, *Substructure of high- p_t Jets at the LHC*, Phys. Rev. D **79** (2009) 074017 [[arXiv:0807.0234](#) [hep-ph]].
- [47] J. Gallicchio and M. D. Schwartz, *Quark and Gluon Tagging at the LHC*, Phys. Rev. Lett. **107** (2011) 172001 [[arXiv:1106.3076](#) [hep-ph]].
- [48] M. L. Gonzalez Silva (for the ATLAS Collaboration), presentation at Boost 2012, Valencia, Spain.
- [49] A. Banfi, G. P. Salam and G. Zanderighi, *Principles of general final-state resummation and automated implementation*, JHEP **0503** (2005) 073 [[hep-ph/0407286](#)].
- [50] M. Jankowiak and A. J. Larkoski, *Jet Substructure Without Trees*, JHEP **1106** (2011) 057 [[arXiv:1104.1646](#) [hep-ph]]; *Angular Scaling in Jets*, JHEP **1204** (2012) 039 [[arXiv:1201.2688](#) [hep-ph]].
- [51] A. J. Larkoski, G. P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, JHEP **1306** (2013) 108 [[arXiv:1305.0007](#) [hep-ph]].
- [52] M. Dasgupta, L. Schunk and G. Soyez, *Jet shapes for boosted jet two-prong decays from first-principles*, [arXiv:1512.00516](#) [hep-ph].
- [53] J. -H. Kim, *Rest Frame Subjet Algorithm With SISCone Jet For Fully Hadronic Decaying Higgs Search*, Phys. Rev. D **83** (2011) 011502 [[arXiv:1011.1493](#) [hep-ph]].
- [54] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N -subjettiness*, JHEP **1103** (2011) 015 [[arXiv:1011.2268](#) [hep-ph]].
- [55] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing N -subjettiness*, JHEP **1202** (2012) 093 [[arXiv:1108.2701](#) [hep-ph]].
- [56] G. Aad *et al.* [ATLAS Collaboration], *Search for pair production of massive particles decaying into three quarks with the ATLAS detector in $\sqrt{s} = 7$ TeV pp collisions at the LHC*, [arXiv:1210.4813](#) [hep-ex].
- [57] J. Thaler and L. -T. Wang, *Strategies to Identify Boosted Tops*, JHEP **0807** (2008) 092 [[arXiv:0806.0023](#) [hep-ph]].
- [58] A. J. Larkoski and J. Thaler, *Unsafe but Calculable: Ratios of Angularities in Perturbative QCD*, JHEP **1309** (2013) 137 [[arXiv:1307.1699](#) [hep-ph]].

- [59] 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** (2008) 242001 [[arXiv:0802.2470](#) [hep-ph]].
- [60] M. Cacciari, G. P. Salam and G. Soyez, *Fluctuations and asymmetric jet events in PbPb collisions at the LHC*, Eur. Phys. J. C **71** (2011) 1692 [[arXiv:1101.2878](#) [hep-ph]].
- [61] I. P. Lokhtin and A. M. Snigirev, *A model of jet quenching in ultrarelativistic heavy ion collisions and high- p_t hadron spectra at RHIC*, Eur. Phys. J. C **45** (2006) 211 [[hep-ph/0506189](#)].
- [62] I. P. Lokhtin, A. V. Belyaev, L. V. Malinina, S. V. Petrushanko, E. P. Rogochaya and A. M. Snigirev, *Hadron spectra, flow and correlations in PbPb collisions at the LHC: interplay between soft and hard physics*, [arXiv:1204.4820](#) [hep-ph].
- [63] I. P. Lokhtin, L. V. Malinina, S. V. Petrushanko, A. M. Snigirev, I. Arsene and K. Tywoniuk, *Heavy ion event generator HYDJET++ (HYDrodynamics plus JETs)*, Comput. Phys. Commun. **180** (2009) 779 [[arXiv:0809.2708](#)].
- [64] R. Baier, Y. L. Dokshitzer, A. H. Mueller, S. Peigne and D. Schiff, *Radiative energy loss and p_t -broadening of high energy partons in nuclei*, Nucl. Phys. B **484** (1997) 265 [[hep-ph/9608322](#)]; B. G. Zakharov, *Radiative energy loss of high energy quarks in finite-size nuclear matter and quark-gluon plasma*, JETP Lett. **65** (1997) 615 [[hep-ph/9704255](#)].
- [65] J. Casalderrey-Solana and C. A. Salgado, *Introductory lectures on jet quenching in heavy ion collisions*, Acta Phys. Polon. B **38** (2007) 3731 [[arXiv:0712.3443](#)]; D. d'Enterria, *Jet quenching*, [arXiv:0902.2011](#); U. A. Wiedemann, *Jet Quenching in Heavy Ion Collisions*, [arXiv:0908.2306](#); A. Majumder and M. Van Leeuwen, *The theory and phenomenology of perturbative QCD based jet quenching*, [arXiv:1002.2206](#); D. d'Enterria and B. Betz, *High- p_t Hadron Suppression And Jet Quenching*,
- [66] N. Armesto, L. Cunqueiro and C. A. Salgado, *Q-PYTHIA: a medium-modified implementation of final state radiation*, Eur. Phys. J. C **63** (2009) 679 [[arXiv:0907.1014](#)]; N. Armesto, L. Cunqueiro, C. A. Salgado and W. C. Xiang, *Medium-evolved fragmentation functions*, JHEP **0802** (2008) 048 [[arXiv:0710.3073](#)].
- [67] Y. -S. Lai, B. A. Cole, *Jet reconstruction in hadronic collisions by Gaussian filtering*, [[arXiv:0806.1499](#)].
- [68] S. Salur, *Full Jet Reconstruction in Heavy Ion Collisions*, Nucl. Phys. A **830** (2009) 139C [[arXiv:0907.4536](#)]; M. Ploskon [STAR Collaboration], *Inclusive cross section and correlations of fully reconstructed jets in 200 GEV Au+Au and p+p collisions*, Nucl. Phys. A **830** (2009) 255C [[arXiv:0908.1799](#)].
- [69] J. Adams *et al.* [STAR Collaboration], *Measurements of transverse energy distributions in Au + Au collisions at $\sqrt{s_{nn}} = 200$ GeV*, Phys. Rev. C **70** (2004) 054907 [[nucl-ex/0407003](#)].
- [70] S. S. Adler *et al.* [PHENIX Collaboration], *Systematic studies of the centrality and $\sqrt{s_{NN}}$ dependence of $dE_T/d\mu$ and $dN_{ch}/d\mu$ in heavy ion collisions at mid-rapidity*, Phys. Rev. C **71** (2005) 034908 [Erratum-ibid. C **71** (2005) 049901] [[nucl-ex/0409015](#)].
- [71] B. B. Back *et al.*, *The significance of the fragmentation region in ultrarelativistic heavy ion collisions*, Phys. Rev. Lett. **91** (2003) 052303 [[nucl-ex/0210015](#)].

- [72] B. B. Back *et al.* [PHOBOS Collaboration], *Charged hadron transverse momentum distributions in au + au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Lett. **B578** (2004) 297-303. [[nucl-ex/0302015](#)].
- [73] A. Adare *et al.* [PHENIX Collaboration], *Suppression pattern of neutral pions at high transverse momentum in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV and constraints on medium transport coefficients*, Phys. Rev. Lett. **101** (2008) 232301. [[arXiv:0801.4020](#)].
- [74] N. Armesto, *Predictions for the heavy-ion programme at the Large Hadron Collider*, [arXiv:0903.1330](#) [hep-ph].
- [75] K. Aamodt *et al.* [The ALICE Collaboration], *Charged-particle multiplicity density at mid-rapidity in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, [arXiv:1011.3916](#) [nucl-ex].
- [76] S. L. Blyth *et al.*, *A cone jet-finding algorithm for heavy-ion collisions at LHC energies*, J. Phys. G **34** (2007) 271 [[nucl-ex/0609023](#)].
- [77] M. Estienne [ALICE Collaboration], *Jet reconstruction in heavy ion collisions (emphasis on Underlying Event background subtraction)*, [arXiv:0910.2482](#).
- [78] I. P. Lokhtin and A. M. Snigirev, *Fast simulation of flow effects in central and semi-central heavy ion collisions at LHC*, [hep-ph/0312204](#).
- [79] K. Zapp, G. Ingelman, J. Rathsman, J. Stachel and U. A. Wiedemann, *A Monte Carlo Model for 'Jet Quenching'*, Eur. Phys. J. C **60** (2009) 617 [[arXiv:0804.3568](#)].
- [80] B. Schenke, C. Gale and S. Jeon, *MARTINI: An event generator for relativistic heavy-ion collisions*, Phys. Rev. C **80** (2009) 054913 [[arXiv:0909.2037](#)].
- [81] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, Phys. Lett. B **659** (2008) 119 [[arXiv:0707.1378](#)].
- [82] P. Jacobs, presentation at "Jets in p+p and heavy-ion collisions, Prague, August 12-14, 2010, http://bazille.fjfi.cvut.cz/~jetworkshop/talkdir/peter_jacobs.pdf
- [83] N. Grau, B. A. Cole, W. G. Holzmann, M. Spousta and P. Steinberg [ATLAS Collaboration], *Identification and Rejection of Fake Reconstructed Jets From a Fluctuating Heavy Ion Background in ATLAS*, Eur. Phys. J. C **62** (2009) 191 [[arXiv:0810.1219](#)].
- [84] Y. S. Lai for the PHENIX collaboration, *Direct jet reconstruction in p + p and Cu + Cu at PHENIX*, [arXiv:0911.3399](#).
- [85] E. Bruna for the STAR Collaboration, presentation at "Jets in p+p and heavy-ion collisions", Prague, August 12-14, 2010, http://bazille.fjfi.cvut.cz/~jetworkshop/talkdir/elena_bruna.ppt.
- [86] N. Borghini and U. A. Wiedemann, *Distorting the hump-backed plateau of jets with dense QCD matter*, [hep-ph/0506218](#).
- [87] X.-F. Guo and X.-N. Wang, *Multiple scattering, parton energy loss and modified fragmentation functions in deeply inelastic e A scattering*, Phys. Rev. Lett. **85**, 3591 (2000) [[hep-ph/0005044](#)].

- [88] N. Armesto, L. Cunqueiro, C. A. Salgado and W. -C. Xiang, *Medium-evolved fragmentation functions*, JHEP **0802**, 048 (2008) [[arXiv:0710.3073](#) [hep-ph]].
- [89] F. Arleo, *(Medium-modified) Fragmentation Functions*, Eur. Phys. J. C **61**, 603 (2009) [[arXiv:0810.1193](#) [hep-ph]].
- [90] S. Sapeta and U. A. Wiedemann, *Jet hadrochemistry as a characteristics of jet quenching*, Eur. Phys. J. C **55**, 293 (2008) [[arXiv:0707.3494](#) [hep-ph]].
- [91] A. Majumder, *The In-medium scale evolution in jet modification*, [arXiv:0901.4516](#) [nucl-th].
- [92] A. Beraudo, J. G. Milhano and U. A. Wiedemann, *Medium-induced color flow softens hadronization*, Phys. Rev. C **85** (2012) 031901 [[arXiv:1109.5025](#) [hep-ph]]; A. Beraudo, J. G. Milhano and U. A. Wiedemann, *The contribution of medium-modified color flow to jet quenching*, JHEP **1207** (2012) 144 [[arXiv:1204.4342](#) [hep-ph]].
- [93] M. G. Albrow *et al.* [TeV4LHC QCD Working Group], *Tevatron-for-LHC Report of the QCD Working Group*, [hep-ph/0610012](#).
- [94] B. Abelev *et al.* [ALICE Collaboration], *Measurement of Event Background Fluctuations for Charged Particle Jet Reconstruction in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, JHEP **1203** (2012) 053 [[arXiv:1201.2423](#) [hep-ex]].
- [95] G. P. Salam and G. Soyez, *JHEP* **05** (2007) 086, [[arXiv:0704.0292](#)]; see also, <http://projects.hepforge.org/siscone>.
- [96] G. C. Blazey *et al.*, *Run II jet physics*, [hep-ex/0005012](#).
- [97] D. de Florian and W. Vogelsang, *Resummed cross section for jet production at hadron colliders*, Phys. Rev. D **76** (2007) 074031 [[arXiv:0704.1677](#) [hep-ph]].
- [98] M. Dasgupta, L. Magnea and G. P. Salam, *Non-perturbative QCD effects in jets at hadron colliders*, [arXiv:0712.3014](#) [hep-ph].
- [99] M. Dasgupta and G. P. Salam, *Resummation of non-global QCD observables*, Phys. Lett. B **512** (2001) 323 [[hep-ph/0104277](#)]; *Accounting for coherence in interjet $E(t)$ flow: A case study*, JHEP **0203** (2002) 017 [[hep-ph/0203009](#)].
- [100] A. Banfi, G. Marchesini and G. Smye, *Away-from-jet energy flow*, JHEP **0208**, 006 (2002) [[hep-ph/0206076](#)].
- [101] R. B. Appleby and M. H. Seymour, *Non-global logarithms in inter-jet energy flow with kt clustering requirement*, JHEP **0212** (2002) 063 [[hep-ph/0211426](#)]; *The resummation of inter-jet energy flow for gaps-between-jets processes at HERA*, JHEP **0309** (2003) 056 [[hep-ph/0308086](#)].
- [102] A. Banfi and M. Dasgupta, *Problems in resumming interjet energy flows with $k(t)$ clustering*, Phys. Lett. B **628** (2005) 49 [[hep-ph/0508159](#)].
- [103] A. Bhatti *et al.*, *Determination of the jet energy scale at the Collider Detector at Fermilab*, Nucl. Instrum. Meth. A **566** (2006) 375 [[hep-ex/0510047](#)].

- [104] G. P. Salam and D. Wicke, *Hadron masses and power corrections to event shapes*, JHEP **0105**, 061 (2001) [[hep-ph/0102343](#)].
- [105] A. Banfi, G. Corcella and M. Dasgupta, *Angular ordering and parton showers for non-global QCD observables*, JHEP **0703**, 050 (2007) [[hep-ph/0612282](#)].
- [106] T. Sjostrand and P. Z. Skands, *Transverse-momentum-ordered showers and interleaved multiple interactions*, Eur. Phys. J. C **39** (2005) 129 [[hep-ph/0408302](#)].
- [107] M. Dasgupta, A. Fregoso, S. Marzani and G. P. Salam, *Towards an understanding of jet substructure*, JHEP **1309** (2013) 029 [[arXiv:1307.0007](#) [hep-ph]].
- [108] D. Krohn, J. Thaler and L. T. Wang, *Jet Trimming*, JHEP **1002** (2010) 084 [[arXiv:0912.1342](#) [hep-ph]].
- [109] S. D. Ellis, C. K. Vermilion and J. R. Walsh, *Techniques for improved heavy particle searches with jet substructure*, Phys. Rev. D **80** (2009) 051501 [[arXiv:0903.5081](#) [hep-ph]].
- [110] M. Dasgupta, A. Fregoso, S. Marzani and A. Powling, *Jet substructure with analytical methods*, Eur. Phys. J. C **73** (2013) no.11, 2623 [[arXiv:1307.0013](#) [hep-ph]].
- [111] S. Catani, B. R. Webber and G. Marchesini, *QCD coherent branching and semiinclusive processes at large x* , Nucl. Phys. B **349** (1991) 635.
- [112] S. Sapeta and Q. C. Zhang, *The mass area of jets*, JHEP **1106** (2011) 038 [[arXiv:1009.1143](#) [hep-ph]].
- [113] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup Per Particle Identification*, JHEP **1410** (2014) 059 [[arXiv:1407.6013](#) [hep-ph]].
- [114] M. Cacciari, G. P. Salam and G. Soyez, unpublished, presented at CMS Week, CERN, Geneva, Switzerland, March 2011, available publicly since then at <http://www.lpthe.jussieu.fr/~salam/talks/repo/2011-CMS-week.pdf>.
- [115] The ATLAS collaboration, *Pile-up subtraction and suppression for jets in ATLAS*, ATLAS-CONF-2013-083.
- [116] CMS Collaboration, *Jet Energy Scale performance in 2011*, CMS-DP-2012-006.
- [117] M. H. Seymour, *Searches for new particles using cone and cluster jet algorithms: A Comparative study*, Z. Phys. C **62** (1994) 127.
- [118] M. Cacciari, G. P. Salam and G. Soyez, [arXiv:1404.7353v1](#).
- [119] CMS Collaboration, *Pileup Jet Identification*, CMS-PAS-JME-13-005.
- [120] ATLAS Collaboration, *Tagging and suppression of pileup jets*, ATL-PHYS-PUB-2014-001.
- [121] M. Cacciari, J. Rojo, G. P. Salam and G. Soyez, *Quantifying the performance of jet definitions for kinematic reconstruction at the LHC*, JHEP **0812** (2008) 032 [[arXiv:0810.1304](#) [hep-ph]].
- [122] M. D. Schwartz, Talk at Workshop on Mitigation of Pileup Effects at the LHC, CERN, May 2014, <https://indico.cern.ch/event/306155/session/2/contribution/1/material/slides/0.pdf>

- [123] M. Cacciari, G. P. Salam and G. Soyez, *FastJet contrib*, <http://fastjet.hepforge.org/contrib/>
- [124] CMS Collaboration, *Pileup Removal Algorithms*, CMS-PAS-JME-14-001.
- [125] O.L. Kodolova *et al*, *Study of γ +Jet Channel in Heavy ion Collisions with CMS*, CMS-NOTE-1998-063; V. Gavrilov, A. Oulianov, O. Kodolova and I. Vardanian, *Jet Reconstruction with Pileup Subtraction*, CMS-RN-2003-004; 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** (2007) 117.
- [126] 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>.
- [127] R. Field, *Min-Bias and the Underlying Event at the LHC*, *Acta Phys. Polon. B* **42** (2011) 2631 [[arXiv:1110.5530](https://arxiv.org/abs/1110.5530) [hep-ph]].
- [128] V. A. Khoze, S. Lupia and W. Ochs, *Perturbative universality in soft particle production*, *Eur. Phys. J. C* **5** (1998) 77 [[hep-ph/9711392](https://arxiv.org/abs/hep-ph/9711392)].
- [129] H. Kirschenmann [CMS Collaboration], *Determination of the jet energy scale in CMS*, *J. Phys. Conf. Ser.* **404** (2012) 012013.
- [130] G. Aad *et al.* [ATLAS Collaboration], *Jet energy measurement and its systematic uncertainty in proton–proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, [arXiv:1406.0076](https://arxiv.org/abs/1406.0076) [hep-ex].
- [131] P. Berta, M. Spousta, D. W. Miller and R. Leitner, *Particle-level pileup subtraction for jets and jet shapes*, *JHEP* **1406** (2014) 092 [[arXiv:1403.3108](https://arxiv.org/abs/1403.3108) [hep-ex]].
- [132] S. Catani, G. Turnock and B. R. Webber, *Jet broadening measures in e^+e^- annihilation*, *Phys. Lett. B* **295** (1992) 269.
- [133] P. Skands, S. Carrazza and J. Rojo, *Tuning PYTHIA 8.1: the Monash 2013 Tune*, [arXiv:1404.5630](https://arxiv.org/abs/1404.5630) [hep-ph].
- [134] CMS Collaboration, *Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taups, and MET*, CMS-PAS-PFT-09-001.
- [135] *Mitigation of Pileup Effects at the LHC*, <https://indico.cern.ch/event/306155/>.
- [136] <https://github.com>.
- [137] [CMS Collaboration], *Jet energy scale and resolution performances with 13TeV data*, CMS-DP-2016-020.
- [138] CMS Collaboration [CMS Collaboration], *Detailed Characterization of Jets in Heavy Ion Collisions Using Jet Shapes and Jet Fragmentation Functions*, CMS-PAS-HIN-12-013.
- [139] [ATLAS Collaboration], *Measurement of inclusive jet charged particle fragmentation functions in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector*, ATLAS-CONF-2012-115.
- [140] ATLAS Collaboration, *Search for New Phenomena in the Dijet Mass Distribution updated using 13fb^{-1} of pp Collisions at $\sqrt{s} = 8$ TeV collected by the ATLAS Detector*, ATLAS-CONF-2012-148.

- [141] CMS Collaboration, *Search for Narrow Resonances using the Dijet Mass Spectrum with 19.6 fb⁻¹ of pp Collisions at $\sqrt{s} = 8$ TeV*, [CMS-PAS-EXO-12-059](#).