

Interleaved Parton Showers and Tuning Prospects

R. Corke¹ and T. Sjöstrand²

*Theoretical High Energy Physics,
Department of Astronomy and Theoretical Physics,
Lund University,
Sölvegatan 14A,
S-223 62 Lund, Sweden*

Abstract

General-purpose Monte Carlo event generators have become important tools in particle physics, allowing the simulation of exclusive hadronic final states. In this article we examine the PYTHIA 8 generator, in particular focusing on its parton-shower algorithms. Some relevant new additions to the code are introduced, that should allow for a better description of data. We also implement and compare with $2 \rightarrow 3$ real-emission QCD matrix elements, to check how well the shower algorithm fills the phase space away from the soft and collinear regions. A tuning of the generator to Tevatron data is performed for two PDF sets and the impact of first new LHC data is examined.

¹richard.corke@thep.lu.se

²torbjorn@thep.lu.se

1 Introduction

The production of exclusive hadronic final states involves many aspects, ranging from calculable perturbative to incalculable nonperturbative physics. General-purpose Monte Carlo (MC) event generators remain vital tools in the simulation of such events. Here, typically, leading-order (LO) matrix elements are combined with parton showers and hadronisation models. Additionally, when the incoming beams are comprised of hadrons, parton distribution functions (PDFs) and multiple parton interactions (MPIs) are required to account for the (multi)parton content of the hadrons. Alternatively, a large number of specialized generators concentrate on providing a more accurate perturbative description, e.g. by higher-order calculations. There need not be a contradiction between the two approaches; indeed a large amount of work has been done to combine them, in order to get the best possible overall description.

Two of the most widely used general-purpose MC tools have been PYTHIA 6 [1] and HERWIG [2, 3]. These two programs both strive to describe the same physics, but differ in the models used to do so. One key difference between the two programs is in their shower algorithms; both are based on DGLAP evolution [4–6], supplemented with Sudakov form factors [7], but while HERWIG uses angular ordering of emissions, PYTHIA 6 instead offers showers with either virtuality (Q^2) or transverse-momentum (p_\perp) ordering. Even with all the differences between the two programs, it is remarkable the range of physics which both are able to describe.

Despite the successes, both continue to evolve. They have been ported from Fortran to C++, but neither of these new versions is a simple rewrite. New features continue to be added, including changes to their shower algorithms. For example, Herwig++ [8] has a modified shower-ordering variable to be better suited for heavy particles. Other new shower developments include the release of the Sherpa event generator [9], featuring a Catani-Seymour dipole shower [10–12], and Vincia, based on the dipole-antenna picture [13].

The original PYTHIA 6 shower was virtuality-ordered. Combined with the MPI model [14], a large amount of Tevatron data can successfully be described. Later, a transverse-momentum-ordered shower was added in conjunction with an updated MPI model [15, 16]. An advantage to this p_\perp ordering is the implicit inclusion of coherence effects [17], also present in an angular-ordered shower, but only crudely implemented in the Q^2 -ordered shower. Another feature is a dipole-style approach to recoils, as pioneered by the ARIADNE program [17–19]. The use of a p_\perp evolution variable for both the shower and MPI frameworks also allows for interleaving, i.e. for a single downwards evolution in p_\perp where, at any stage, either a shower emission or an MPI can take place. In PYTHIA 6, initial-state radiation (ISR) and MPIs are interleaved, while final-state radiation (FSR) is only considered afterwards. The argument is that it is primarily MPI and ISR that compete for beam momentum. Also new is the inclusion of showers off all MPI subsystems, limited to just the hard scattering in the Q^2 -ordered showers.

In PYTHIA 8, the Q^2 -ordered showers are no longer available, while the p_\perp -ordered showers are updated [20]. The main change is the additional interleaving of FSR along with ISR and MPI. It has been possible to obtain a good agreement with much data also in this new model, but some time ago it was noted that problems exist in the description of underlying-event (UE) physics when the model has been tuned to minimum-bias (MB) data [21]. This is exemplified in Fig. 1, which shows the transverse region charged (a)

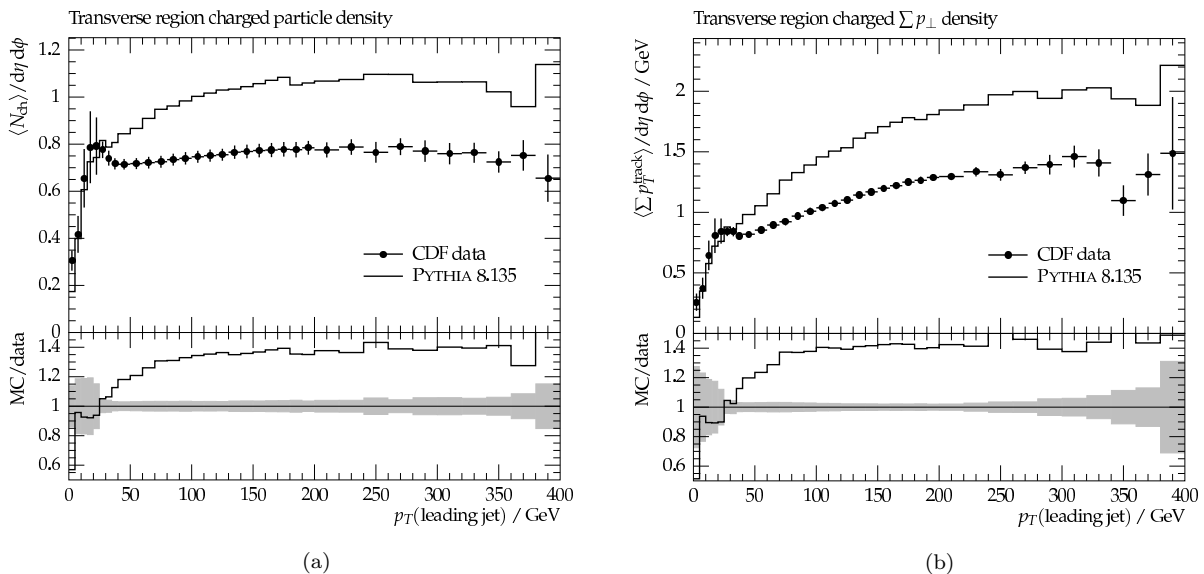


Figure 1: Rick Field’s leading-jet underlying-event analysis implemented in Rivet. Results for PYTHIA 8.135, default tune, for transverse region charged (a) number and (b) $\sum p_{\perp}$ density

number and (b) $\sum p_{\perp}$ density for PYTHIA 8.135 in Rick Field’s analysis of the UE activity as a function of the leading-jet p_{\perp} [22]. Although the default tune is able to describe minimum-bias multiplicities, the activity in the underlying event clearly rises too quickly with jet p_{\perp} . With this in mind, we study the new shower framework of interleaved FSR and ISR, and introduce modest changes to both that significantly improve agreement.

Although the shower description is an approximation of QCD primarily intended to be accurate in the soft and collinear limits, it can also be used to fill the phase space away from these regions. One direct way to examine how well the shower does away from these limits is through a comparison to $2 \rightarrow 3$ real-emission matrix elements. We will perform such a study in this article, for the improved shower framework we present.

This should be put in context. In preparing for LHC phenomenology, there has been a large focus on techniques to combine higher-order matrix elements with parton showers. One goal is to describe many hard and widely separated jets, achieved by algorithms such as MLM or CKKW(-L) [23–25], while another is to combine full NLO calculations with showers, with algorithms such as MC@NLO and POWHEG [26–28]. Attempts to combine both NLO calculations and higher-order Born-level matrix elements have also been made [13, 29–31]. PYTHIA 8 comes with matching built in for a few processes [32, 33], and can be interfaced to external matching programs [34], but QCD multijets has not been high on the list. An eventual goal would be to have at least a complete built-in matching of the shower to the QCD $2 \rightarrow 3$ matrix elements for the first shower emission, e.g. in the spirit of [33].

There are at least two good reasons why the default showers should still be fixed up to give as good a description of the three-jet (and higher-jet) region as possible. One is that the CKKW-L-matching approach is based on obtaining Sudakov form factors from trial showers, so the more accurate the shower, the more precise the matching. The other is that one can foresee that time-consuming matching approaches will primarily be applied to

the hardest interaction of an event, whereas further MPIs will be handled to the accuracy provided by the default showers.

The formulation of sound physics models is central to the development of a successful generator, but it is not enough. By the nonperturbative, and therefore incalculable, nature of much of the physics, and by the uncertainty that comes from shower approximations and unknown higher orders in the perturbative regime, generators inevitably come to contain a number of free parameters. Tuning of these parameters thus becomes a key part of obtaining an agreement with data. Assuming jet universality, LEP data can be used to tune FSR and hadronisation. The p_{\perp} spectrum of Z^0 provides a handle on ISR (and, at low p_{\perp} , on primordial k_{\perp}). PDFs are a topic unto themselves, addressed further in later sections, still leaving the MPI model, the crosstalk between ISR and FSR, beam remnants, colour flow issues, diffractive physics and more. The underlying-event analyses here serve as vital ingredients [22, 35, 36]. Many of the tunes for the Q^2 -ordered showers of PYTHIA 6 are based on Rick Field’s original Tune A [37].

The tuning process has now been systematised to a large extent, and Rivet [38] and Professor [39] have become the tools of choice for many MC programs. For PYTHIA 6, the Professor team have released both a Q^2 -shower tune, Pro-Q20, and a p_{\perp} -shower tune, Pro-pT0 [39, 40]. For the p_{\perp} -ordered showers, Peter Skands has also released the “Perugia” tunes, where a central parameter set is supplemented by 8 related variations that attempt to systematically explore different changes in the parameter sets [41]. The current default tune for PYTHIA 8 is based on a tuning to LEP data by Hendrik Hoeth, using the Rivet + Professor framework, supplemented by hadron collider data comparisons by Peter Skands. Other tunes are often developed within the experimental collaborations [42, 43] and tunes that incorporate (preliminary) LHC data have also started to appear [44]. While the new energies probed by the LHC offer an exciting opportunity to test and constrain models and model parameters inside generators, data from lower energy runs ($\sqrt{s} = 900$ GeV and 2.36 TeV) currently appears slightly incompatible with earlier Tevatron results [45]. It remains to be seen if a “global tune” that can encompass both LHC and Tevatron data is possible.

In Section 2, the shower algorithms of PYTHIA 8 are reviewed, and some developments made to allow an improved description of underlying-event and minimum-bias data. Other parts of the event generation framework are also outlined, especially those not previously documented and relevant for tuning exercises. In Section 3, a comparison of the shower is made to $2 \rightarrow 3$ QCD real-emission matrix elements. In Section 4, some initial tunes of the modified framework are made to Tevatron data, while Section 5 contains a first look at the new LHC data. Finally, in Section 6, a summary and outlook is given.

2 Event-generation framework

2.1 Interleaved evolution

The PYTHIA 8 showers are ordered in transverse momentum [15], both for ISR and for FSR. Also, multiparton interactions (MPI) are ordered in p_{\perp} [14]. This allows a picture where MPI, ISR and FSR are interleaved in one common sequence of decreasing p_{\perp} values [20]. This is most important for MPI and ISR, since they are in direct competition for momentum

from the beams, while FSR (mainly) redistributes momenta between already kicked-out partons. The interleaving implies that there is one combined evolution equation

$$\begin{aligned} \frac{d\mathcal{P}}{dp_{\perp}} &= \left(\frac{d\mathcal{P}_{\text{MPI}}}{dp_{\perp}} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp_{\perp}} + \sum \frac{d\mathcal{P}_{\text{FSR}}}{dp_{\perp}} \right) \\ &\times \exp \left(- \int_{p_{\perp}}^{p_{\perp\text{max}}} \left(\frac{d\mathcal{P}_{\text{MPI}}}{dp'_{\perp}} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp'_{\perp}} + \sum \frac{d\mathcal{P}_{\text{FSR}}}{dp'_{\perp}} \right) dp'_{\perp} \right), \end{aligned} \quad (1)$$

that probabilistically determines what the next step will be. Here the ISR sum runs over all incoming partons, two per already produced MPI, the FSR sum runs over all outgoing partons, and $p_{\perp\text{max}}$ is the p_{\perp} of the previous step. ISR is described by backwards evolution [46], wherein branchings are constructed from the hard process, back to the shower initiators. Starting from a single hard interaction, eq. (1) can be used repeatedly to construct a complete parton-level event of arbitrary complexity.

The decreasing p_{\perp} scale can be viewed as an evolution towards increasing resolution; given that the event has a particular structure when activity above some p_{\perp} scale is resolved, how might that picture change when the resolution cutoff is reduced by some infinitesimal dp_{\perp} ? That is, let the “harder” features of the event set the pattern to which “softer” features have to adapt. It does not have a simple interpretation in absolute time; all the MPIs occur essentially simultaneously (in a simpleminded picture where the protons have been Lorentz contracted to pancakes), while ISR stretches backwards in time and FSR forwards in time. The closest would then be to view eq. (1) as an evolution towards increasing formation time.

2.2 Dipoles vs. Feynman graphs

Before the detailed presentation of the initial- and final-state showers, it is useful to put them in the context of the overall philosophy. As a simple illustration, consider the process $qg \rightarrow qg$. One of the contributing graphs is a t -channel exchange of a gluon, with one colour-anticolour being annihilated in the process, another created, and one colour flowing through from the initial to final state, Fig. 2a. This can be transformed to Fig. 2b, illustrating how the partons move out from the interaction vertex. In this picture, the incoming particles are represented by fictitious outgoing antiparticles; from a colour point of view, these “holes” also represent the beam remnants left behind by the hard interactions. The three colour lines of the Feynman graph here transforms into the presence of three colour dipoles. By the nature of the endpoints, they can be classified as final-final (FF), final-initial (FI) or initial-initial (II). We begin by considering the first kind, but with many of the statements relevant for all three.

In planar QCD, i.e. in the limit of infinitely many colours, $N_C \rightarrow \infty$ [47], the radiation pattern of soft gluons becomes a simple sum of the radiation from each dipole separately. Interference between the dipoles are suppressed by a factor $\mathcal{O}(1/N_C^2)$ and so can often be neglected. When emitted gluons are not soft, however, recoil effects come into play; the endpoint partons of a radiating dipole obtain shifted momenta, affecting all dipoles they belong to. It is therefore not possible to consider emission from the different dipoles independently. Instead, one should impose a well-defined common order in which possible emissions are considered, such as p_{\perp} [17, 18]. Each new gluon emission also means that




Figure 2: (a) One possible Feynman graph and colour flow for the $qg \rightarrow qg$ process, and (b) the resulting colour dipoles between the scattered partons and the beam remnants

the radiating dipole is split into two, so the number of dipoles will increase as the shower evolves.

While radiation should be viewed as emanating from the dipole-as-a-whole, it is necessary to impose a strategy for how the recoil of an emission should be shared between the two endpoints of a dipole. There is no unique prescription, but obviously the original dipole direction should not change more than necessary. One possible choice is to view the emission as being associated with one of the two endpoint partons, the radiator, while the other parton acts as a recoiler. If viewed in the rest frame of the dipole, the recoiler does not change direction by the emission, but it obtains a reduced absolute momentum such that the overall four-momentum of the dipole is preserved during the emission. The radiator and the emitted parton obtain opposite p_{\perp} kicks by the branching, but the distribution in azimuthal angle is isotropic, at least to first approximation. All partons, both before and after the emission, are put on mass shell.

By the radiator-recoiler strategy, the radiation of a dipole is split into that of its two endpoints. This can be done in a smooth manner, such that radiation close to either endpoint is associated with that parton radiating, whereas large-angle radiation is split between the two in reasonable proportions. In this way, it becomes possible to describe gluon radiation from a quark and gluon end by the splitting kernels

$$P_{q \rightarrow qg} = C_F \frac{1+z^2}{1-z}, \quad (2)$$

$$\frac{1}{2} P_{g \rightarrow gg} = \frac{N_C}{2} \frac{1+z^3}{1-z}, \quad (3)$$

respectively, with a continuous transition between the two for a qg dipole.

The dipole picture can also be extended to apply to initial-state partons. Here the effect of radiation is that the incoming parton has to have a larger momentum than previously considered, in the spirit of backwards evolution. There are then four kinds of dipole ends, FF, FI, IF and II, classified by whether the radiator and recoiler are in the final or initial state.

For ISR, however, the dipole picture is known to have shortcomings. Specifically, it is in direct conflict with results from traditional Feynman-diagram calculations. As an example,




Figure 3: (a) Colour flow (dipole) topology for $q\bar{q} \rightarrow Z^0g$, and (b) in PYTHIA the emission of a second gluon would give a recoil to the whole existing Z^0g system rather than only to the dipole-connected first gluon

consider the production of a Z^0 . For the emissions of a first gluon the two approaches would agree, but this emission would lead to a picture with two IF dipoles, stretched from the incoming q and \bar{q} to the gluon, respectively, Fig. 3a. Any further radiation would therefore occur from these two dipoles, and their daughters, but never again have any impact on the kinematics of the Z^0 . With Feynman diagrams, on the other hand, the Z^0 takes a recoil that is modified as further gluon emissions are considered, and typically resummation techniques [48] are used to sum up the effects of infinitely many gluon emissions on the p_\perp spectrum of the Z^0 . This is a well-known problem e.g. with Catani-Seymour (CS) dipoles [10]. Therefore CS-based dipole showers do not necessarily apply the CS recoil strategy to the letter. Furthermore, when matching procedures are used, hard emissions are described by matrix elements and only softer radiation would be affected by the dipole picture.

Given these shortcomings, the ISR in PYTHIA is not based on the CS-style dipole picture. Instead, an ISR emission is allowed to give a recoil to all partons “downstream” of it, i.e. that have been created by the hard process or by previously considered emissions at a larger p_\perp scale, Fig. 3b. This applies both to the II and IF dipole ends. In the latter case one would still expect a memory of the colour flow topology to show up as an azimuthal anisotropy of radiation, studied further in later sections. For the FF and FI types, a more standard dipole picture is used, but we note and address an issue with the allowed emission region for the FI case. As it turns out, these moderate changes to the IF and FI dipole handling do help address the problem with tuning to Tevatron data.

2.3 Transverse-momentum-ordered showers

The current PYTHIA parton shower orders FSR emissions in terms of a p_\perp^2 evolution variable, with an additional energy-sharing variable z in the branching. For QCD emissions the DGLAP evolution equations lead to the probability for the splitting of parton $a \rightarrow bc$

$$d\mathcal{P}_a = \frac{dp_\perp^2}{p_\perp^2} \sum_{b,c} \frac{\alpha_s(p_\perp^2)}{2\pi} P_{a \rightarrow bc}(z) dz, \quad (4)$$

where $P_{a \rightarrow bc}$ are the DGLAP splitting kernels. This inclusive quantity can be turned into an exclusive one by requiring that, for the first (“hardest”) emission below the starting scale $p_{\perp \max}^2$, no emissions can have occurred at a scale larger than p_{\perp}^2 . The no-emission probability is the Sudakov form factor

$$S_a(p_{\perp \max}^2, p_{\perp}^2) = \exp \left(- \int_{p_{\perp}^2}^{p_{\perp \max}^2} \frac{dp_{\perp}'^2}{p_{\perp}'^2} \sum_{b,c} \int_{z_{\min}(p_{\perp}'^2)}^{z_{\max}(p_{\perp}'^2)} dz \frac{\alpha_s(p_{\perp}'^2)}{2\pi} P_{a \rightarrow bc}(z) \right). \quad (5)$$

Its introduction turns the unnormalised distribution into a normalised one,

$$d\mathcal{P}_a = \frac{dp_{\perp}^2}{p_{\perp}^2} \sum_{b,c} \frac{\alpha_s(p_{\perp}^2)}{2\pi} P_{a \rightarrow bc}(z) dz S_a(p_{\perp \max}^2, p_{\perp}^2), \quad (6)$$

i.e. with unit integral over the full phase space. In practice, a lower cutoff, $p_{\perp 0}^2$, is introduced to keep the shower away from soft and collinear regions, which leads to a fraction of events with no emissions inside the allowed region.

For initial-state radiation (ISR), using backwards evolution, a given parton b entering a hard scattering is unresolved into a parton a which preceded it. Here, the parton distribution functions, reflecting the contents of the incoming hadron, must be taken into account. This leads to a Sudakov with the form

$$S_b(x, p_{\perp \max}^2, p_{\perp}^2) = \exp \left(- \int_{p_{\perp}^2}^{p_{\perp \max}^2} dp_{\perp}'^2 \sum_{a,c} \int_{z_{\min}(p_{\perp}'^2)}^{z_{\max}(p_{\perp}'^2)} dz \frac{\alpha_s(p_{\perp}'^2)}{2\pi} P_{a \rightarrow bc}(z) \frac{x' f_a(x', p_{\perp}'^2)}{x f_b(x, p_{\perp}'^2)} \right), \quad (7)$$

where $z = x/x' = x_b/x_a$, and a corresponding normalised distribution

$$d\mathcal{P}_b = \frac{dp_{\perp}^2}{p_{\perp}^2} \sum_{a,c} \frac{\alpha_s(p_{\perp}^2)}{2\pi} P_{a \rightarrow bc}(z) \frac{x' f_a(x', p_{\perp}^2)}{x f_b(x, p_{\perp}^2)} dz S_b(x, p_{\perp \max}^2, p_{\perp}^2). \quad (8)$$

Currently both the running renormalisation and factorisation shower scales, i.e. the scales at which α_s and the PDFs are evaluated, are chosen to be p_{\perp}^2 .

For the following, a relevant aspect is that, while p_{\perp} ordering is used both for MPI, ISR and FSR in PYTHIA, the p_{\perp} definition is slightly different in the three components. For an MPI the p_{\perp} is the expected one; the transverse momentum of the two scattered partons in a $2 \rightarrow 2$ process, defined in a Lorentz frame where the two incoming beams are back-to-back. To understand why the same choice is not made for ISR or FSR, consider a $q \rightarrow qg$ branching, where the p_{\perp} of the emitted gluon is defined with respect to the direction of the initial quark. The p_{\perp} as a function of the gluon emission angle θ increases up till 90° , but then decreases again, $p_{\perp} \rightarrow 0$ for $\theta \rightarrow 180^\circ$. Thus, an ordering in such a p_{\perp} would classify a $\sim 180^\circ$ emission as collinear and occurring late in the evolution, although it would involve a more off-shell propagator than an emission at 90° . It could also erroneously associate a $1/p_{\perp}^2$ divergence with the $\theta \rightarrow 180^\circ$ limit. Therefore it is natural to choose an evolution variable that does not turn over at 90° .

To this end, consider a branching $a \rightarrow bc$ (e.g. $q \rightarrow qg$), where z is defined as the lightcone momentum along the axis of a that b takes. Then

$$p_{\perp \text{LC}}^2 = z(1-z)m_a^2 - (1-z)m_b^2 - zm_c^2, \quad (9)$$

and this equation can be used as inspiration to define evolution variables

$$\text{ISR} : p_{\perp\text{evol}}^2 = (1 - z)Q^2 \quad (10)$$

$$\text{with } m_b^2 = -Q^2 \text{ and } m_a^2 = m_c^2 = 0,$$

$$\text{FSR} : p_{\perp\text{evol}}^2 = z(1 - z)Q^2 \quad (11)$$

$$\text{with } m_a^2 = Q^2 \text{ and } m_b^2 = m_c^2 = 0,$$

which are monotonous functions of the virtuality Q^2 . In the actual branching kinematics, z is replaced by a Lorentz invariant definition while the Q^2 interpretation is retained.

2.3.1 Initial-state showers

For the initial-state shower, the radiating dipole is always chosen such that the recoiler is the incoming parton from the other side of the subcollision. Thus the whole enclosed systems share the recoil of a new emission, unlike e.g. a Catani-Seymour dipole scheme, as discussed previously.

The z definition in eq. (10) is chosen to be $z = m_{br}^2/m_{ar}^2$, where r is the recoiler, which allows for a straightforward bookkeeping of the x values to be used in PDFs. The actual p_{\perp} of b and c then becomes

$$p_{\perp b,c}^2 = (1 - z)Q^2 - \frac{Q^4}{m_{ar}^2} = p_{\perp\text{evol}}^2 - \frac{p_{\perp\text{evol}}^4}{p_{\perp\text{evol,max}}^2}, \quad (12)$$

where $p_{\perp\text{evol,max}}^2 = (1 - z)Q_{\text{max}}^2 = (1 - z)^2 m_{ar}^2$ is the kinematically possible upper limit, not to be confused with the starting scale for the downwards evolution in $p_{\perp\text{evol}}^2$. Usually $p_{\perp\text{evol}}^2 \ll p_{\perp\text{evol,max}}^2$, so that $p_{\perp b,c}^2 \approx p_{\perp\text{evol}}^2$.

To understand how the two dipole ends combine to produce the radiation pattern, again consider $q(1)\bar{q}(2) \rightarrow g(3)Z^0(4)$, cf. Fig. 3a. Defining $\hat{\theta}$ as the emission angle of the gluon in the rest frame of the process, and with the the Mandelstam variables $\hat{t}, \hat{u} = -\hat{s}(1 \mp \cos \hat{\theta})/2$, for the gluon being emitted by the incoming q , one has

$$\frac{dp_{\perp\text{evol}}^2}{p_{\perp\text{evol}}^2} = \frac{dQ^2}{Q^2} = \frac{d\hat{t}}{\hat{t}} = \frac{d(\cos \hat{\theta})}{1 - \cos \hat{\theta}} = (1 + \cos \hat{\theta}) \frac{d(\cos \hat{\theta})}{1 - \cos^2 \hat{\theta}}, \quad (13)$$

while an emission by the incoming \bar{q} is exactly the same, except for a sign flip from $\hat{t} \rightarrow \hat{u}$. It then follows that the quark is responsible for a fraction $(1 + \cos \hat{\theta})/2$ of the total amount of radiation at an angle $\hat{\theta}$, i.e. there is a smooth transition in the middle of the event. One simplifying factor is that $z = M_Z^2/\hat{s}$ is independent of the radiator side. Another is that we have considered the process for fixed incoming partons, so that PDF weights do not enter. If viewed instead as backwards evolution from $q\bar{q} \rightarrow Z^0$, a different PDF ratio would enter for emissions on the two sides, but the rest would be the same. The simple relations between the evolution and the Mandelstam variables makes this shower convenient for some matching tasks [32]. Further details of the kinematics, also for massive quarks, can be found in [15].

The direct relation between $p_{\perp\text{evol}}$ and emission angle (for fixed z) implies that typically central emissions will be considered first, and then successively emissions closer in angle or, equivalently, rapidity to the two beam directions. This ordering need not be strict however,

since z will also vary from one emission to the next. This implies that the resulting dipoles can zigzag back and forth in rapidity. We note that in CCFM-based [49,50] approaches [51] (as well as in angular-ordered showers, of course) emissions are ordered in rapidity, but then need not be ordered in p_{\perp} . Hard emissions tend to be ordered in both, however, and it is only in the soft region that a random walk in p_{\perp} sets in. While the amount of emitted partons need not be incorrect in our p_{\perp} -ordered approach, there is a worry that the zigzagging colour connections will give too large dipoles, which then gives more FSR and hadronisation activity than if colours were angularly ordered. Ultimately this is an issue that could be studied by the rate of forward jets and forward particle production in general. Different modifications to the existing algorithm could be envisioned. For now, however, we will content ourselves with an option where rapidity-unordered ISR emissions are vetoed, subsequent to the first emission off each dipole end. That is, while p_{\perp} is always ordered, the rapidity may or may not be.

Perturbation theory breaks down in the $p_{\perp} \rightarrow 0$ limit, and thus some cutoff procedure is necessary for showers. The simplest is a fixed lower $p_{\perp\text{min}}$ cutoff. However, a more attractive possibility is a smooth damping. Actually, the damping introduced within the MPI framework, see below in Sec. 2.4.2, could be viewed as a reduced “resolved” partonic content of the two incoming hadrons, and thus be the same for the ISR description. If so, the shower rate should be reduced roughly like

$$\frac{d\mathcal{P}}{dp_{\perp}^2} \propto \frac{\alpha_s(p_{\perp}^2)}{p_{\perp}^2} \rightarrow \frac{\alpha_s(p_{\perp 0}^2 + p_{\perp}^2)}{p_{\perp 0}^2 + p_{\perp}^2} . \quad (14)$$

We will allow for either possibility in the following.

2.3.2 Azimuthal asymmetries

For an ISR emission, as described previously, the dipole is always taken as stretching between the two incoming partons. In the colour-dipole picture, however, it will involve a colour line that is either stretched along the beam (collision) axis, or out to a final-state parton. This can lead to azimuthal asymmetries, as follows. In the first emission from a $q\bar{q} \rightarrow Z^0$ event, the colour line stretched between the two beams means there is no preferred azimuthal direction, and emissions are isotropic in φ . For a second emission, the colour line is instead now stretched from the radiating incoming parton to the final state. This colour dipole configuration leads to an enhancement of radiation near the azimuthal angle of this final-state parton. For the emission of a soft gluon from this dipole, and with a convenient separation of radiation between the two dipole ends, it is possible to derive what this distribution should look like [52]. To be specific, if the radiator is moving along the $+z$ direction and its colour partner along a (reasonably small) polar angle θ_{par} and azimuthal angle φ_{par} , then soft emissions should be distributed roughly like

$$\frac{d\mathcal{P}}{d\varphi} \propto \frac{1 - r \cos(\Delta\varphi)}{1 + r^2 - 2r \cos(\Delta\varphi)} \quad \text{with } \Delta\varphi = \varphi - \varphi_{\text{par}} , \quad (15)$$

where $r = \theta/\theta_{\text{par}}$. For small r the bias is tiny, but then increases as $r \rightarrow 1$. It turns partly negative for $r > 1$, and there integrates to 0 — the angular-ordering condition. It must be stressed, however, that the behaviour near $r \rightarrow 1$ depends on how the dipole is split into two ends, and also obtains corrections when the energy of the radiated parton is not

vanishingly small. As already discussed, we also optionally allow for cases where emissions are not rapidity-ordered, meaning $r > 1$, where azimuthal anisotropies again should fall off in the limit of a large r . Instead, we define an alternative

$$r = N \frac{2\theta_{\text{par}}}{\theta^2 + \theta_{\text{par}}^2}, \quad (16)$$

where N is a free parameter to be tuned later. It has a natural value of $1/2$ by the above small-angle behaviour, but we allow for somewhat different values, though never getting close to unity.

Another reason for anisotropies is the plane polarization of the gluon. That is, the planes defined by the production and by the decay of a gluon, respectively, tend to be correlated [52]. If the gluon is produced in a $g \rightarrow gg$ branching then a $g \rightarrow gg$ decay leads to a small tendency for the two planes to line up. By contrast, a $g \rightarrow q\bar{q}$ branching gives a strong tendency for the two planes to be at 90° to each other. The branching $g \rightarrow q\bar{q}$ is less frequent than $g \rightarrow gg$, so to first approximation it is expected that the net effect should be vanishingly small [53]. This is also what we observe, both when applied for ISR and for FSR.

2.3.3 Final-state showers

For FSR, the dipole picture is again used but, unlike ISR, the recoil of emissions is here assumed to be taken entirely by the other end of the colour dipole, as already discussed in Sec. 2.2. For the case of FF dipoles all relevant kinematics equations are already published [15], and we only sketch some relevant points here.

In the rest frame of a dipole ar , the branching $a \rightarrow bc$ at an evolution scale $p_{\perp\text{evol}}$ implies that parton a acquires a virtuality $Q^2 = p_{\perp\text{evol}}^2/(z(1-z))$. Its energy is thereby increased from the original $m_{ar}/2$ to $(m_{ar}^2 + Q^2)/(2m_{ar})$, and the recoiler energy reduced accordingly. Parton b then takes a fraction z of the increased radiator energy, and c takes $1-z$. This choice gives an exact match e.g. to the singularity structure of the well-known $q\bar{q} \rightarrow q(1)\bar{q}(2)g(3)$ branching in e^+e^- annihilation:

$$\frac{dp_{\perp\text{evol},q}^2}{p_{\perp\text{evol},q}^2} \frac{dz_q}{1-z_q} + \frac{dp_{\perp\text{evol},\bar{q}}^2}{p_{\perp\text{evol},\bar{q}}^2} \frac{dz_{\bar{q}}}{1-z_{\bar{q}}} = \frac{dx_1 dx_2}{(1-x_2)x_3} + \frac{dx_1 dx_2}{(1-x_1)x_3} = \frac{dx_1 dx_2}{(1-x_1)(1-x_2)} \quad (17)$$

with $x_i = 2E_i/\hat{E}_{\text{CM}}$. In the soft-gluon limit, $x_3 \rightarrow 0$, the kinematics simplifies to $1-x_2 \propto m_{\text{qg}}^2 \propto 1 - \cos \hat{\theta}$ and $1-x_1 \propto m_{\bar{\text{qg}}}^2 \propto 1 + \cos \hat{\theta}$, where $\hat{\theta}$ is the emission angle of the gluon with respect to the original $q\bar{q}$ axis. Thus one obtains exactly the same continuous transition from q - to \bar{q} -emitted gluons as for ISR, eq. (13). For nonvanishing gluon energy, the angular dependence becomes more complicated, but retains the same qualitative features.

The relation between the formal $p_{\perp\text{evol}}$ and the $p_{\perp b} = p_{\perp c}$ of the daughters with respect to the mother direction is somewhat lengthy. It shares the same physical properties as eq. (12) for ISR; at small values $p_{\perp\text{evol}}^2$ and $p_{\perp b,c}^2$ follow suit, and so correspond to identical $1/p_{\perp}^2$ singular behaviours, but the latter then turns around and vanishes when $p_{\perp\text{evol}}^2$ approaches the kinematical limit.

While FF dipoles are familiar from and constrained by e^+e^- physics, the FI ones are less well studied. In the PYTHIA 6 shower, FSR was deferred until after all ISR, such that

the FI dipoles at the end of the ISR evolution were typically characterised by a low p_{\perp} scale, meaning they did not radiate much. Indeed, it was shown that disregarding emission from them altogether had only a tiny influence on event properties. Instead, with ISR and FSR interleaved, typically each event contains one or a few FI dipole ends that should be evolved all the way from the p_{\perp} scale of the hard interaction.

The kinematics of an FI branching gives some differences relative to an FF one. In the dipole rest frame a fraction Q^2/m_{ar}^2 of the recoiler energy is given from the recoiler to the emitter, as above. That is, the emitter four-momentum is increased to

$$p'_a = p_a + \frac{Q^2}{m_{ar}^2} p_r , \quad (18)$$

which is valid in any Lorentz frame. But the recoiler is not a final-state particle, so the increase of a momentum is not compensated anywhere in the final state. Instead the incoming parton that the recoiler represents must have its momentum increased, not decreased, by the same amount as the emitter. That is, its x value needs to be scaled up as

$$x'_r = \left(1 + \frac{Q^2}{m_{ar}^2} \right) x_r . \quad (19)$$

The dipole mass m_{ar} and the squared subcollision mass \hat{s} are increased in the process, the latter by the same factor as x_r . As with ISR, the increased x value leads to an extra PDF weight

$$\frac{x'_r f_r(x'_r, p_{\perp}^2)}{x_r f_r(x_r, p_{\perp}^2)} \quad (20)$$

in the emission probability and Sudakov form factor. This ensures a proper damping of radiation in the $x'_r \rightarrow 1$ limit.

The above equations are valid in the massless limit. Incoming particles are always bookkept as exactly massless, so that p'_r remains parallel with the beam axis if p_r was. If the recoiler has an on-shell mass m_r^2 , then its energy is already from the onset bigger than the recoiler one in the dipole rest frame, and the momentum and x are not rescaled by as much

$$x'_r = \left(1 + \frac{Q^2 - m_a^2}{m_{ar}^2 - m_a^2} \right) x_r . \quad (21)$$

The issue that remains is how the emission off the FI dipole end should be damped as a function of location along the dipole. It has already been shown that there is a natural fall-off at around 90° in the rest frame of the dipole. For an FF and II dipole this offers a convenient subdivision, with the two ends together providing a smooth coverage. What must now be asked is how to subdivide FSR and ISR in an FI dipole. The natural choice is to view the dipole in a boosted+rotated frame where the radiator is along the $+z$ axis, with an energy equal to its original p_{\perp} , and the recoiler along $-z$ with energy $m_{ar}^2/(4p_{\perp})$. This is close to the Breit frame definition in Deeply Inelastic Scattering, with p_{\perp} as the hard scale instead of (the DIS variable) $Q/2$ and m_{ar}^2 instead of W^2 . The Breit frame offers a convenient separation between asymmetric current and target hemispheres.

A more precise motivation is offered by the fractal picture of dipole emission. While originally formulated for FSR [54], the same kinematics argument can be applied here, as follows. Consider a dipole stretched between two partons i and j along the $\pm z$ axis.

At some arbitrary (e.g. cutoff) scale m_0 the rapidity range open for emissions is then $\Delta y = \ln(m_{ij}^2/m_0^2)$. Assume a gluon k is emitted at an angle θ , sufficiently soft that recoil effects can be neglected. Then the new range(s) allowed subsequently is

$$\begin{aligned}\Delta y' &= \ln(m_{ik}^2/m_0^2) + \ln(m_{kj}^2/m_0^2) = \ln\left(\frac{2E_i E_k (1 - \cos\theta)}{m_0^2} \frac{2E_j E_k (1 + \cos\theta)}{m_0^2}\right) \\ &= \ln\left(\frac{4E_i E_j}{m_0^2} \frac{E_k^2 \sin^2\theta}{m_0^2}\right) = \Delta y + 2\ln(p_{\perp k}/m_0) .\end{aligned}\tag{22}$$

Here the unchanged Δy term can be associated with radiation essentially along the z axis. In our case this would correspond not only to II dipoles but also to the IF ends of FI dipoles, which are both considered as a continuous chain along the whole allowed rapidity range. The last term is the additional radiation to be associated with the two new FI dipole ends.

In summary, what we see is that the p_{\perp} scale of the hard emission serves a double purpose. On the one hand, it sets the maximum scale for subsequent emissions, on the other it delimits the phase space that emissions from an FI dipole end can populate without doublecounting.

This clarification especially concerns cases when the two relevant scales, the dipole mass m_{ar} and the radiator scale p_{\perp} are well separated. Actually, this is precisely what you expect to happen by t -channel gluon exchange; a (relatively) small change of momentum but a complete swap of colours. So in a process such as $ud \rightarrow ud$ the outgoing u will be hooked up with the beam remnant of the incoming d , which most of the time will be in the other hemisphere of the event, with $m_{ar} \gg p_{\perp}$. Such a dipole would have a large rapidity range to radiate inside, if hemispheres are separated in the dipole rest frame, and produce a doublecounting of FSR and ISR effects. As we will show, indeed it does appear to make a difference in event properties.

Finally, we need to specify a convenient expression for the damping at the border between the two uneven hemispheres. Again consider the ‘‘Breit’’-frame arrangement, with parton a having a momentum p_{\perp} along the $+z$ axis. In the branching $a \rightarrow bc$ parton b takes the fraction z , where we assume $z < 1/2$, so that b is the softer one. Both have transverse momentum $p_{\perp b,c}$. If z is interpreted as lightcone momentum then rapidity 0 for b is seen to correspond to $p_{\perp b,c} = p_{+b} = zp_{+a} = z2p_{\perp}$. The complete kinematics is more messy, but qualitatively it holds that $p_{\perp b,c} = z(1-z)p_{\perp}$ is a reasonable measure of rapidity 0, give or take factors of 2, and with $1-z$ inserted for symmetry. We therefore introduce a suppression weight

$$\mathcal{W} = \frac{z(1-z)p_{\perp}}{z(1-z)p_{\perp} + p_{\perp b,c}} = \frac{p_{\perp} p_{\perp b,c}}{p_{\perp} p_{\perp b,c} + p_{\perp b,c}/(z(1-z))} = \frac{p_{\perp} p_{\perp b,c}}{p_{\perp} p_{\perp b,c} + Q^2}\tag{23}$$

to smoothly kill the unwanted radiation in the backward hemisphere. Note that this damping comes on top of the ‘‘natural’’ subdivision of radiation in the dipole rest frame. Whichever one is more restrictive will dominate.

2.4 Other model details

In later sections, parameters relating to other parts of the event generation framework enter when tuning to data. In the rest of this section, then, these parts of the generator are outlined.

2.4.1 Total cross sections

The total cross section for pp collisions (excluding the Coulomb contribution) is subdivided into elastic and inelastic ones, with the latter further subdivided into diffractive and non-diffractive contributions. Many diffractive topologies are allowed, but the most common and the only ones currently implemented in PYTHIA are single and double diffraction. We therefore write

$$\sigma_{\text{tot}}(s) = \sigma_{\text{el}}(s) + \sigma_{\text{sd}}(s) + \sigma_{\text{dd}}(s) + \sigma_{\text{nd}}(s) , \quad (24)$$

with s the squared center-of-mass energy. While the class of elastic events is unambiguously defined, the other three are less easily distinguished. Single and double diffraction should be characterised by the presence of a rapidity gap, with either one or both of the incoming protons excited to a higher-mass system, whereas non-diffractive events should have no rapidity gaps. In practice, however, generated diffractive systems can fluctuate in mass up to the kinematical limit, in which case there is no gap, while generated non-diffractive events can still display gaps, through fluctuations in the hadronisation process. The experimental definition of diffractive events therefore does not have a one-to-one correspondence with the underlying picture in a generator. This opens up the possibility for alternative experimental definitions, but it is important that the exact one is clearly specified.

The “minimum bias” subset is also poorly defined as all events that are triggered within the context of some specific detector setup. In practice, elastic and low-mass diffractive events are excluded, leaving inelastic non-diffractive events with a modest high-mass diffractive contamination. Within the PYTHIA generator “minbias” is used as a convenient shorthand for the more clumsy and nondescriptive “inelastic non-diffractive” classification.

The different cross sections can be calculated from Regge theory, once a few free parameters have been extracted from data itself. This works fine at lower, fixed-target energies. It also appears to work for the total cross section up to higher energies, as e.g. in the DL parameterisation used by us [55]. In this approach the diffractive (and elastic) cross sections increase too fast with energy, however, and ultimately exceed the parameterised total cross section. This signals the need for higher-order corrections to dampen the growth. The parameterisations developed for PYTHIA [56, 57] introduce such corrections in an empirical fashion. In a program like PHOJET a more sophisticated eikonalisation approach is chosen instead, which leads to a slower growth of diffractive cross sections than in PYTHIA. Other arguments have also been raised why the diffractive cross section should cease to increase [58, 59]. A recent ATLAS study of the diffractive cross section at 7 TeV (although not currently fully corrected for detector effects) suggests that PYTHIA is on the high side, though not overly so, while PHOJET rather undershoots [60].

It seems likely that some tuning downwards of diffractive rates will be required. While it has been possible for a user to set the respective cross sections by hand, this would have to be done separately for each energy. We have now introduced an alternative simple mechanism to dampen the growth relative to the parameterisation of [56], without (significantly) affecting fixed-target phenomenology. For each of the diffractive rates it is thus possible to specify a maximum value that will be approached asymptotically

$$\sigma_i^{\text{mod}}(s) = \frac{\sigma_i^{\text{old}}(s) \sigma_i^{\text{max}}}{\sigma_i^{\text{old}}(s) + \sigma_i^{\text{max}}} . \quad (25)$$

We do not propose that $\sigma_i^{\text{mod}}(s)$ should approach asymptotia at current energies, only that

the increase should be slowed down. Since the total and elastic cross sections are not affected, a reduction in diffractive cross sections means an increase of the non-diffractive one.

2.4.2 Multiple interactions

Let us now study the inelastic non-diffractive event class, and focus on just the contribution from MPI. The probability for an interaction is then given by

$$\frac{d\mathcal{P}_{\text{MPI}}}{dp_{\perp}} = \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp_{\perp}} \exp\left(-\int_{p_{\perp}}^{p_{\perp i-1}} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp'_{\perp}} dp'_{\perp}\right). \quad (26)$$

Here, interactions are supposed to be generated in a decreasing sequence of p_{\perp} values, $p_{\perp 1} > p_{\perp 2} > p_{\perp 3} > \dots$, starting from some maximum scale, like the p_{\perp} scale of the hard interaction or the upper kinematical limit. The exponential is a standard ‘‘Sudakov’’ factor that restores unitarity. It encodes the probability that there is no interaction intermediate in p_{\perp} between the previous ($i - 1$) and the current (i) one.

The $d\sigma/dp_{\perp}$ is given by the perturbative QCD $2 \rightarrow 2$ cross section

$$\frac{d\sigma}{dp_{\perp}^2} = \sum_{i,j} \int dx_1 \int dx_2 f_i(x_1, p_{\perp}^2) f_j(x_2, p_{\perp}^2) \frac{d\hat{\sigma}}{dp_{\perp}^2}, \quad (27)$$

representing the convolution of the hard partonic cross section, $d\hat{\sigma}/dp_{\perp}^2$, with the two incoming parton densities, f_i and f_j . This cross section is dominated by t -channel gluon exchange, and diverges roughly as dp_{\perp}^2/p_{\perp}^4 (further steepened by the rise of the PDFs at small x). To avoid this divergence, the idea of colour screening is introduced. The concept of a perturbative cross section is based on the assumption of free incoming states, which is not the case when partons are confined in colour-singlet hadrons. One therefore expects a colour charge to be screened by the presence of nearby anti-charges; that is, if the typical charge separation is d , gluons with a transverse wavelength $\sim 1/p_{\perp} > d$ are no longer able to resolve charges individually, leading to a reduced effective coupling. This is introduced by reweighting the interaction cross section such that it is regularised according to

$$\frac{d\hat{\sigma}}{dp_{\perp}^2} \propto \frac{\alpha_s^2(p_{\perp}^2)}{p_{\perp}^4} \rightarrow \frac{\alpha_s^2(p_{\perp 0}^2 + p_{\perp}^2)}{(p_{\perp 0}^2 + p_{\perp}^2)^2}, \quad (28)$$

where $p_{\perp 0}$ (related to $1/d$ above) is now a free parameter in the model. To be more precise, it is the physical cross section $d\sigma/dp_{\perp}^2$ that needs to be regularised, i.e. the convolution of $d\hat{\sigma}/dp_{\perp}^2$ with the two parton densities, eq. (27). One is thus at liberty to associate the screening factor with the incoming hadrons, half for each of them, instead of with the interaction. Such an association also gives a recipe to regularise the ISR divergence, as already noted.

Not only $p_{\perp 0}$ itself, as determined e.g. from Tevatron data, comes with a large uncertainty, but so does the energy scaling of this parameter. The ansatz for the energy dependence of $p_{\perp 0}$ is that it scales in a similar manner to the total cross section, i.e. driven by an effective power related to the Pomeron intercept [55], which in turn could be related to the small- x behaviour of parton densities. This leads to a scaling

$$p_{\perp 0}(E_{\text{CM}}) = p_{\perp 0}^{\text{ref}} \times \left(\frac{E_{\text{CM}}}{E_{\text{CM}}^{\text{ref}}}\right)^{E_{\text{CM}}^{\text{pow}}}, \quad (29)$$

where $E_{\text{CM}}^{\text{ref}}$ is some convenient reference energy and $p_{\perp 0}^{\text{ref}}$ and $E_{\text{CM}}^{\text{pow}}$ are parameters to be tuned to data. This same scaling can also be adopted when regularising the ISR divergence, with the possibility to set the parameters used in the shower independently of those used in the MPI model.

The so-far Poissonian nature of the framework is then changed, first by requiring that there is at least one interaction, such that there is a physical event, and second by including an impact parameter, b . For a given matter distribution, $\rho(r)$, the time-integrated overlap of the incoming hadrons during collision is given by

$$\mathcal{O}(b) = \int dt \int d^3x \rho(x, y, z) \rho(x + b, y, z + t), \quad (30)$$

after a suitable scale transformation to compensate for the boosted nature of the incoming hadrons. This matter profile, $\rho(r)$ may then provides new model parameters within the MPI framework. In the original MPI model, accompanying the Q^2 -ordered showers of PYTHIA 6, the default selection was a double Gaussian

$$\rho(r) \propto \frac{1 - \beta}{a_1^3} \exp\left(-\frac{r^2}{a_1^2}\right) + \frac{\beta}{a_2^3} \exp\left(-\frac{r^2}{a_2^2}\right), \quad (31)$$

such that a fraction β of the matter is contained in a radius a_2 , which in turn is embedded in a radius a_1 containing the rest of the matter. Peter Skands subsequently noted that the inclusion of radiation off all scattering subsystems in an event provides a dynamical source of fluctuations, mirroring the effect of this double Gaussian matter distribution, and was able to achieve good agreement with data using only a single Gaussian profile, introducing no free parameters [41]. Both of these options are available in PYTHIA 8. A further option, intermediate in the sense that it has one free parameter, $E_{\text{exp}}^{\text{pow}}$, is to have an overlap function for the convolution of the two incoming matter distributions of the form

$$\exp\left(-b^{E_{\text{exp}}^{\text{pow}}}\right). \quad (32)$$

2.4.3 Parton densities

Both hard processes, showers and multiple interactions make use of parton densities. Since these are not specified from first principles, different parameterisations are on the market. The CTEQ5L [61] PDF set, dating from 1999, is currently the default choice for both PYTHIA 6 and PYTHIA 8. A large range of PDF sets has always been available by linking to the LHAPDF library [62].

Recently, a selection of newer PDF sets has been made available directly in PYTHIA 8 [63], to make the PYTHIA program more easy to install and run standalone, with a reproducible behaviour. One also saves some computer time by not having to go via interfaces. These new sets include members from the CTEQ6 and MSTW2008 families [64, 65], as well as specially modified sets designed for LO event generators, such as those from the CT09 and MRST* families [66, 67]. These modified sets have been adopted by experimental collaborations [43].

The reason for the modified sets is the following. Precision tests of QCD today normally involve comparisons with NLO matrix elements convoluted with NLO PDFs. NLO expressions are not guaranteed to be positive definite, and do not have a simple probabilistic interpretation. For PDFs, specifically, it is well-known that the gluon has a tendency

to start out negative at small x and Q^2 , and only turn positive by the QCD evolution towards larger Q^2 . In the MRST/MSTW sets this negativity is explicit, while it is masked in the CTEQ sets by picking an ansatz form that cannot be negative, but which still gives a very small gluon contribution. The combination of LO matrix elements and NLO PDFs thus gives an abnormally small rate of interactions at small x and p_\perp , relative to either a pure LO or a pure NLO combination. This region may not be important for New Physics searches, but it is crucial for the modeling of MPI, i.e. minimum-bias and underlying-event physics, and also plays a role for showers.

So if NLO PDFs are not convenient for use in LO generators, also LO PDFs have their problems. Usually NLO calculations give an enhanced rate relative to LO ones, the “K factor” is above unity, and offer a better description to data. This introduces a tension in LO PDF fits, where data sensitive to a specific x range attempts to pull in more of the total momentum, at the expense of other x ranges. The solution of the new modified PDFs is to allow the momentum sum rule to be broken (typically by 10 – 15%), so that the whole x range can be enhanced.

These new PDFs therefore offer the hope to allow improved descriptions of data within a LO framework for MPI and showers. They are not necessarily a panacea, however [63], and comparing results between the modified sets and more traditional ones will therefore be useful. PYTHIA 8 also allows one PDF set to be used for the hard process and another for MPI and showers, so that one could use NLO for the former while still having a sensible behaviour in the $p_\perp \rightarrow 0$ limit with the latter.

2.4.4 Primordial k_\perp

When the combined MPI/ISR/FSR p_\perp evolution has come to an end, the beam remnants will consist of the remaining valence content of the incoming hadrons, as well as any companion (anti)quarks to kicked-out sea quarks. These remnants must carry the remaining fraction of longitudinal momentum. PYTHIA will pick x values for each component of the beam remnants according to distributions such that the valence content is “harder” and will carry away more momentum. In the rare case that there is no remaining quark content in a beam, a gluon is assigned to take all the remaining momentum.

The event is then modified to add primordial k_\perp . Partons are expected to have a non-zero k_\perp value just from Fermi motion within the incoming hadrons. A rough estimate based on the size of the proton gives a value of ~ 0.3 GeV, but when comparing to data, for instance the p_\perp distribution of Z^0 at CDF, a value of ~ 2 GeV appears to be needed. The current solution is to decide a k_\perp value for each initiator parton taken from a hadron based on a Gaussian whose width is generated according to an interpolation

$$\sigma(Q, \widehat{m}) = \frac{Q_{\frac{1}{2}} \sigma_{\text{soft}} + Q \sigma_{\text{hard}}}{Q_{\frac{1}{2}} + Q} \frac{\widehat{m}}{\widehat{m}_{\frac{1}{2}} + \widehat{m}}, \quad (33)$$

where Q is the hardness of a sub-collision (p_\perp for a $2 \rightarrow 2$ QCD process) and \widehat{m} its invariant mass, σ_{soft} and σ_{hard} is a minimal and maximal value, and $Q_{\frac{1}{2}}$ and $\widehat{m}_{\frac{1}{2}}$ the respective scale giving a value halfway between the two extremes. Beam remnants are assigned a separate width σ_{remn} comparable with σ_{soft} . The independent random selection of primordial k_\perp values gives a net imbalance within each incoming beam, which is shared between all

initiator and remnant partons, with a reduction factor $\widehat{m}/(\widehat{m}_{\perp} + \widehat{m})$ for initiators of low-mass systems. With the k_{\perp} 's of the two initiators of a system known, all the outgoing partons of the system can be rotated and Lorentz boosted to the relevant frame. During this process, the invariant mass and rapidity of all systems is maintained by appropriately scaling the lightcone momenta of all initiator partons.

2.4.5 Colour reconnection

The final step at the parton level, before hadronisation, is colour reconnection. The idea of colour reconnection can be motivated by noting that MPI leads to many colour strings that will overlap in physical space, which makes the separate identity of these strings questionable. Alternatively, moving from the limit of $N_C \rightarrow \infty$ to $N_C = 3$, it is not unreasonable to allow these strings to be connected differently due to a coincidence of colour. Adapting either of these approaches, dynamics is likely to favour reconnections that reduce the total string length and thereby the potential energy.

In the old MPI framework, good agreement to CDF data is obtained if 90% of additional interactions produces two gluons with “nearest neighbour” colour connections [22]. More recently, an annealing algorithm has been used [68,69], again requiring a significant amount of reconnection to describe data. In PYTHIA 8, colour reconnection is currently performed by giving each system a probability to reconnect with a harder system

$$\mathcal{P} = \frac{p_{\perp\text{Rec}}^2}{(p_{\perp\text{Rec}}^2 + p_{\perp}^2)}, \quad p_{\perp\text{Rec}} = R \times p_{\perp 0}, \quad (34)$$

where R is a user-tunable parameter and $p_{\perp 0}$ is the same parameter as in eq. (28).

With the above probability for reconnection, it is easier to reconnect low- p_{\perp} systems, which can be viewed as them having a larger spatial extent, such that they are more likely to overlap with other colour strings. When a reconnection is allowed, the partons of the lower- p_{\perp} systems are attached to the existing higher- p_{\perp} colour dipoles in a way that minimizes the total string length.

Currently, however, all of this is only a convenient ansatz. More than that, given the lack of a firm theoretical basis, the need for colour reconnection has only been established within the context of specific models.

2.4.6 Diffraction

The diffractive treatment in PYTHIA 8 has been extended from the simple one implemented in PYTHIA 6, to share many more features with the sophisticated description of non-diffractive events [70]. This is possible by using the Ingelman–Schlein [71] picture, wherein single diffraction is viewed as the emission of a Pomeron pseudoparticle from one incoming proton, leaving that proton intact but with reduced momentum, followed by the subsequent collision between this Pomeron and the other proton. The Pomeron is to first approximation to be viewed as a glueball state with the quantum numbers of the vacuum, but by QCD interactions it will also have a quark content. The Pomeron–proton collision can then be handled as a normal hadron–hadron non-diffractive event, using the full machinery of MPI, ISR, FSR and other aspects as already described above.

The simple factorisation into a Pomeron flux times a Pomeron-proton collision should be viewed as an effective picture, i.e. as a first approximation to physics that is much more

complicated than that. Indeed the HERA studies show deviations from perfect factorisation [72], and diffractive W and jet production at the Tevatron is reduced by a significant factor relative to naive expectations [73]. Nevertheless, this factorisation does allow what hopefully is a realistic picture for the bulk of the diffractive cross section.

In the generation, the first step is to use the parameterised diffractive cross sections, Sec. 2.4.1. Once a single diffractive topology has been decided on, the second step is to use the Pomeron flux $f_{\mathbb{P}/p}(x_{\mathbb{P}}, t)$ to pick the fraction $x_{\mathbb{P}}$ that the Pomeron takes out of the proton momentum, and the squared momentum transfer t of the emitted Pomeron, which is related to the scattering angle of the proton. To first approximation

$$f_{\mathbb{P}/p}(x_{\mathbb{P}}, t) \approx \frac{1}{x_{\mathbb{P}}} \exp(Bt) , \quad (35)$$

where B is an energy-dependent slope parameter. The squared mass of the diffractive $\mathbb{P}p$ system X is given by $M_X^2 = x_{\mathbb{P}}s$, so it follows that $dx_{\mathbb{P}}/x_{\mathbb{P}} = dM_X^2/M_X^2 = d \ln M_X^2$. Since the rapidity range over which particle production can occur is proportional to $\ln M_X^2$, it follows that all kinematically possible ranges are about equally probable, and by complementarity, the same applies to the size of the rapidity gap. This is only a first approximation, however; in the program four alternative $f_{\mathbb{P}/p}(x_{\mathbb{P}}, t)$ shapes can be used, and all of them are somewhat more peaked towards lower masses than the ansatz above indicates.

Note that only the relative shape of $f_{\mathbb{P}/p}(x_{\mathbb{P}}, t)$ is of interest to us. If the absolute normalization were taken at face value it would lead to a too rapidly increasing $\sigma_{sd}(s)$, as already indicated in Sec. 2.4.1. Instead, we implicitly assume that the screening corrections appear as an energy-dependent but $x_{\mathbb{P}}$ - and t -independent reduction factor.

The third step is to set up the partonic state of the diffractive system. Here we make a distinction between low-mass and high-mass diffraction, not because physics is expected to be discontinuous, but because we need to consider such a wide range of diffractive masses M_X , where perturbation theory may be applied for the higher masses but not for the lower ones. A pragmatic dividing line is set at 10 GeV; below it everything is soft, above it the fraction of soft events smoothly drops towards zero.

The soft description is simulated as a mix of two components. In one, the Pomeron is assumed to kick out one of the valence quarks of the proton, so that the partonic final state consists of a single string stretched between the kicked-out quark and the remnant diquark. In the other, the Pomeron instead kicks out a gluon. This gives a hairpin topology, where the string is stretched from a quark in the remnant to the kicked-out gluon and then back to the diquark of the remnant. At the lower end of the range, extending to 1.2 GeV, i.e. around the Δ mass, the former is expected to dominate, and then gradually the latter takes over. Above 10 GeV, the perturbative description starts to be used.

The Pomeron PDFs are significantly worse known than those of the proton. Quite aside from the considerable measurement problems, the data usually probe a convolution of the Pomeron flux with its PDF. Therefore the two cannot be specified independently. It is not even guaranteed that the two factorise. Nevertheless a few alternatives are chosen, both toy-model ones and four H1 sets. Default in PYTHIA 8 is the H1 2006 Fit B LO distribution. Note that the H1 sets use an arbitrary normalisation of the flux, such that the PDFs do not come out normalised to unit momentum sum, but have a sum of roughly 1/2.

Strictly speaking the Pomeron PDF (convoluted with its flux) contributes to the overall

PDF of a proton. This implies that the diffractive component ought to be subtracted from the normal PDFs used for non-diffractive events. In reality the diffractive contribution comes out to be a tiny fraction of the event rate at medium-to-large p_{\perp} scales, and so this doublecounting correction has been neglected for now.

In the description of the MPI interaction probability, eq. (26), $\sigma_{\text{nd}} = \sigma_{\text{nd}}^{\text{PP}}(s)$ plays an important role as a normalisation factor. For the handling of a diffractive system it should be replaced by $\sigma_{\text{tot}}^{\text{IP}}(s)$, which unfortunately is not directly measurable. As with PDFs one reason is that the separation from the Pomeron flux factor is not unambiguous, and another is that we have allowed for a damping of $\sigma_{\text{nd}}^{\text{PP}}(s)$ away from the result of a naive convolution. Instead, an order-of-magnitude estimate is used, as follows. At low energies $\sigma_{\text{nd}}^{\text{PP}}(s)$ is in the ballpark 25–30 mb, and then slowly increases. By contrast $\sigma_{\text{nd}}^{\pi\text{P}}(s)$ is slightly below 20 mb, i.e. roughly 2/3, as suggested by quark counting rules. The same number is now assumed for $\sigma_{\text{tot}}^{\text{IP}}(s)$, by analogy between the two-gluon \mathbb{P} state and the two-quark π state. The larger colour charge of gluons could imply that the Pomeron couples more strongly to protons than pions do. By the same token the Pomeron wave function could be smaller, which would act in the opposite direction, so we will assume that such effects cancel. As a first guess, we therefore take 20 mb to be a reasonable number to use in eq. (26). This assumes, however, that the Pomeron PDFs are normalised to unit momentum sum. Since the H1 fits are normalised only to roughly half of this, σ_{nd} needs to be reduced accordingly, in order for the ratio to balance out. The actual number used therefore is 10 mb. As a sanity check we note that this gives an average charged multiplicity for diffractive systems at a given M_X that is comparable with non-diffractive pp events at the same energy.

Double diffraction is handled in the same spirit. Specifically, the final state will here consist of two Pomeron-proton collision, each with properties as already described. What needs some more consideration is the form of the combined two-side Pomeron flux $f_{\mathbb{P}/\text{p} \times \mathbb{P}/\text{p}}(x_{\mathbb{P}1}, x_{\mathbb{P}2}, t)$, that in particular determines the masses of the two diffractive systems. The old machinery [56] here gives a nontrivial correlation, while some of the alternative forms have only been formulated for single diffraction and therefore here are used by minimal extension, which is likely to be too simpleminded.

2.4.7 Hadronisation

Once the partonic configuration has been specified, including its colour flow, the normal Lund string fragmentation machinery can be used to turn it into a set of primary hadrons [74]. Many of these are unstable and subsequently decay further.

Here we assume that fragmentation parameters tuned to e^+e^- data, notably from LEP, can be carried over unchanged for application to pp collisions. Such a “jet universality” assumption is by no means obvious. In $e^+e^- \rightarrow \gamma^*/Z^0 \rightarrow q\bar{q}$ there is only one original dipole that may radiate further, typically giving rise to one string that does not bend over so much that it can overlap with itself. By contrast, the MPI structure of pp events ensures that many colour strings overlap in space and time during the hadronisation process. The assumption that they then can be treated as completely independent of each other (once colour reconnection has been taken into account) may be overly optimistic. All kinds of collective effects could be imagined, bordering on those associated with a quark–gluon plasma. This could affect e.g. the fraction of strangeness production, which is known to be higher in heavy-ion collisions than in e^+e^- ones [75]. For now, however, such potential

issues are left aside.

3 Comparison to 3-jet matrix elements

3.1 $2 \rightarrow 3$ real emission matrix elements

In this section, we compare the first shower emission against $2 \rightarrow 3$ real-emission matrix elements in QCD events. Eventually, we foresee a full matching of the first emission to these matrix elements, as is done e.g. for electroweak gauge boson production, but here we begin with some simple kinematic comparisons of the default shower. Even with such a matching, emissions after the first will still be handled by the standard shower machinery, and it is questionable if it would be used in all MPI subsystems rather than just the hardest process in an event. It is therefore important to have the best understanding of the shower as possible. There is quite some freedom in how to handle recoil effects, and when the shower has been designed to give sensible tails away from soft and collinear regions, the manner in which this is done grows in importance. Indeed, as we will show, in the soft and collinear regions, the showers give good agreement with the ME behaviour, while further out in phase space, the accuracy does degrade, but without exceedingly large deviations. In the rest of this section, some details of the $2 \rightarrow 3$ ME implementation are given, before comparisons are made, firstly to examine the azimuthal asymmetries introduced in Sec. 2.3.2, and secondly to study the rates and kinematics of shower emissions.

The formulae for the squared matrix elements in the massless QCD approximation, summed and averaged over spin, are taken from [76], with crossing applied where necessary. These expressions are compact, and so easy to input, but the tradeoff is the lack of information for generating the colour structure of events; for the kinematic distributions we consider here, this is not a serious limitation. In what follows, the outgoing partons are labeled 3, 4 and 5, and ordered such that $p_{\perp 3} > p_{\perp 4} > p_{\perp 5}$. The renormalisation scale is expected to relate to each vertex of the process, and by default is taken to be the geometric mean of the squared transverse masses of the three outgoing particles, $\mu_r^2 = \sqrt[3]{m_{\perp 3}^2 * m_{\perp 4}^2 * m_{\perp 5}^2}$. The factorisation scale, instead, in a picture comparable to backwards evolution in ISR, would be related to the smaller scales in the event, and is taken to be the geometric mean of the two smallest squared transverse masses of the three outgoing particles, $\mu_f^2 = \sqrt{m_{\perp 4}^2 * m_{\perp 5}^2}$.

A phase space generator has been written that is specially adapted for these processes, taking into account the important role of the soft and collinear singularities. The variables $p_{\perp 3}$, $p_{\perp 5}$, y_3 , y_4 and y_5 are used to generate the phase space. $p_{\perp 3}$ is first picked according to a $d^2 p_{\perp 3} / p_{\perp 3}^4$ distribution and then $p_{\perp 5}$ according to $d^2 p_{\perp 5} / p_{\perp 5}^2$. φ_3 and φ_5 are picked flat, before $p_{\perp 4}$ is finally reconstructed. All three rapidities are picked flat in their respective allowed range. Phase space cuts are placed on the two p_{\perp} variables, $p_{\perp 3}^{\min}$ and $p_{\perp 5}^{\min}$, and, additionally, there is an $R = \sqrt{(\Delta\eta)^2 + (\Delta\varphi)^2}$ separation cut, R_{sep}^{\min} , applied to all three possible pairs of partons. $p_{\perp 4}$ is not explicitly constrained, but must lie in the range specified by $p_{\perp 3}$ and $p_{\perp 5}$.

To validate these processes, the cross sections and kinematics have been compared to AlpGen [77–79] (excluding processes with no final state gluons) and MadEvent [80] (for processes with no final state gluons). For all numbers shown here, the CTEQ6L1 PDF set has been used, with a first order running of α_s , with $\alpha_s(M_W) = 0.13$. In both AlpGen and

| Cross section (μb) | | | | |
|---------------------------------|-------|--------|--------|--------|
| Tevatron | | | LHC | |
| Cuts | Ref | PYTHIA | Ref | PYTHIA |
| 1 | 1404 | 1409 | 23938 | 24164 |
| 2 | 2.786 | 2.798 | 171.58 | 172.97 |
| 3 | 0.191 | 0.191 | 17.87 | 17.92 |

Table 1: Overall cross sections for $2 \rightarrow 3$ partonic processes, for three sets of cuts (see text) at the Tevatron ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV) and LHC (pp , $\sqrt{s} = 14$ TeV). Results from PYTHIA are compared to summed cross sections from AlpGen and MadEvent

MadEvent, the factorisation and renormalisation scales have been set to follow the PYTHIA defaults. In all programs, the allowed incoming and outgoing flavours have been set to include the charm quark and below. In Tab. 1, overall cross sections are given for three sets of cuts at Tevatron ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV) and LHC (pp , $\sqrt{s} = 14$ TeV) energies:

- 1) $p_{\perp 3}^{\min} = 5.0$ GeV, $p_{\perp 5}^{\min} = 5.0$ GeV, $R_{\text{sep}}^{\min} = 0.1$
- 2) $p_{\perp 3}^{\min} = 50.0$ GeV, $p_{\perp 5}^{\min} = 5.0$ GeV, $R_{\text{sep}}^{\min} = 0.1$
- 3) $p_{\perp 3}^{\min} = 50.0$ GeV, $p_{\perp 5}^{\min} = 25.0$ GeV, $R_{\text{sep}}^{\min} = 1.0$

The difference between the reference and PYTHIA cross sections is less than 1% for all three sets of cuts. In the PS, the amount of ISR is roughly 70%, 60% and 90% respectively, and so dominates, even when using a rather small R_{sep}^{\min} . In Tab. 2, the different classes of $2 \rightarrow 3$ processes are given, with indicative cross sections for each process for the first set of cuts at Tevatron and LHC energies. Overall, the cross section is dominated by gg and qg scattering, and the effect grows at LHC energies, in particular driven by the large increase in small- x gluons.

3.2 Parton shower comparison

The comparison we wish to make is between the $2 \rightarrow 3$ matrix elements and $2 \rightarrow 2$ events with one additional PS emission (hereafter referred to as $2 \rightarrow 2 \otimes \text{PS}$). Overall, the generation must be such that the comparison is as fair as possible, but still allowing a reasonable efficiency, so the following settings have been applied:

- 1) The CTEQ6L1 PDF set is used everywhere.
- 2) The coupling strength for the hard process, ISR and FSR is set constant, $\alpha_s = 0.130$.
- 3) QED radiation is disabled.
- 4) The number of incoming and outgoing flavours possible is set equal for both the ME and PS (up to and including the charm quark).
- 5) All other stages of the event generation framework are switched off.

| Class | Processes | Cross section (μb) | |
|--------------------------|--|---------------------------------|----------|
| | | Tevatron | LHC |
| HardQCD:gg2ggg | $gg \rightarrow ggg$ | 717.46 | 15236.92 |
| HardQCD:qqbar2ggg | $q\bar{q} \rightarrow ggg$ | 0.51 | 2.47 |
| HardQCD:qg2qgg | $qg \rightarrow qgg$ | 525.65 | 6776.26 |
| HardQCD:qq2qqgDiff | $qq' \rightarrow qq'g, q\bar{q}' \rightarrow q\bar{q}'g$ | 45.47 | 404.31 |
| HardQCD:qq2qqgSame | $qq \rightarrow qqg, \bar{q}\bar{q} \rightarrow \bar{q}\bar{q}g$ | 8.92 | 94.51 |
| HardQCD:qqbar2qqbargDiff | $q\bar{q} \rightarrow q'\bar{q}'g$ | 0.55 | 2.68 |
| HardQCD:qqbar2qqbargSame | $q\bar{q} \rightarrow q\bar{q}g$ | 16.08 | 81.93 |
| HardQCD:gg2qqbarg | $gg \rightarrow q\bar{q}g$ | 71.22 | 1302.09 |
| HardQCD:qg2qqqbarDiff | $qg \rightarrow qq'\bar{q}', \bar{q}g \rightarrow \bar{q}\bar{q}'q'$ | 18.05 | 197.95 |
| HardQCD:qg2qqqbarSame | $qg \rightarrow qq\bar{q}, \bar{q}g \rightarrow \bar{q}q\bar{q}$ | 6.02 | 66.01 |

Table 2: The different classes of event implemented and their corresponding processes. Primes refer to different quark/anti-quark flavours and processes include all possible permutations of the initial or final state. Cross sections from PYTHIA are given for the first set of cuts as detailed in the text

- 6) To improve efficiency (without affecting the outcome), a cut $\hat{p}_{\perp\text{min}} = 0.4 * p_{\perp 3}^{\text{min}}$ is used in generating the underlying $2 \rightarrow 2$ events.

Even with these settings, a couple of issues remain. One is the different scale at which PDF factors are evaluated in ISR, eq. (8). Specifically, both the numerator and denominator of this factor are evaluated at a scale p_{\perp}^2 , giving an implicit running α_s when compared to the ME result. A quick check using typical scales and x values for the above cuts shows that this effect is not large, and we choose not to correct for this.

More important is the presence of Sudakov form factors in the PS results. To generate inclusive $2 \rightarrow 2 \otimes \text{PS}$ distributions, which can be compared against the ME results, a $2 \rightarrow 2$ hard process is selected, then a parton shower emission allowed to create a $2 \rightarrow 2 \otimes \text{PS}$ event. This event is then analysed to determine if it meets the required cuts, and if so, it is accepted and added to the statistics. The emission is then vetoed; PYTHIA will revert the event to its original $2 \rightarrow 2$ state, and then continue the shower evolution from the scale of the now-vetoed branching. Thus, a single $2 \rightarrow 2$ event, in the course of its evolution, can give rise to several $2 \rightarrow 3$ configurations. In fact, for an evolution variable t , ranging between 0 and 1, the distribution becomes a Poissonian with an average $\int_0^1 p(t)dt$. By the veto algorithm, the Sudakov form factor is exactly compensated by the Poissonian distribution

$$p(t) \exp\left(-\int_0^t p(t') dt'\right) \cdot \sum_{n=0}^{\infty} \frac{1}{n!} \left(\int_0^t p(t') dt'\right)^n = p(t), \quad (36)$$

where the sum runs over the n $2 \rightarrow 3$ events that may have been generated at $t' < t$.

3.3 The α angle

We begin by studying the azimuthal asymmetries introduced for ISR in Sec. 2.3.2. The strength of the asymmetry, tuned through the N parameter of eq. (16), enters in all following studies, and one goal here is to fix its value. The question to be asked is, how strongly does the plane of an initial-state emission correlate with its colour partner in the final state, when present. The distribution can be calculated in the soft limit, but obtains corrections away from this region; hard emissions, especially in the case when strict rapidity ordering is not enforced (as for the first emission and optionally for subsequent ones) are expected to be constrained to a much lower degree. In both the ME calculation and in data, there is no clean separation between ISR and FSR, so we instead use the α angle observable, used by the CDF collaboration in order to study colour coherence effects in hadron colliders [81].

At the hadron level, the idea is to search for a leading jet which is hard enough such that “soft” radiation is still hard enough to form secondary jets. The study is based on QCD events, where at least three jets must be present (after jet clustering with the CDF fixed-cone algorithm, $R_{\text{cone}} = 0.7$). Both the hardest and second hardest jets must lie in the region $|\eta| < 0.7$ and the leading jet must have $E_{\perp 1} > 110$ GeV. The two leading jets must be approximately back-to-back in the transverse plane, $||\varphi_1 - \varphi_2| - \pi| < \pi/9$, while the third jet must have $E_{\perp 3} > 10$ GeV. These cuts should isolate a region in which the shower approximation is still valid, and azimuthal asymmetries should appear as a correlation between the directions of the second and third jets.

Defining $\Delta\eta = \eta_3 - \eta_2$ and $\Delta\varphi = \varphi_3 - \varphi_2$, an $R_{23} = \sqrt{\Delta\eta^2 + \Delta\varphi^2}$ cut is introduced, $1.1 < R_{23} < \pi$. The lower limit is placed due to clustering cone-size effects, while the upper limit avoids the region where $|\Delta\varphi|$ approaches its maximal value of π . The α angle is then defined by

$$\alpha = \tan^{-1} \left(\frac{\text{sgn}(\eta_2)\Delta\eta}{|\Delta\varphi|} \right). \quad (37)$$

A full discussion of the details of the different α -angle regions is given in [81]. Here, we restrict ourselves to azimuthal correlations only. The primary effect should be a shift of events towards the endpoints of the distribution, $\alpha \rightarrow \pm\pi/2$, such that ISR is biased to sit in the plane of its colour-connected partner, $\Delta\varphi \rightarrow 0$. It should be noted that $\alpha < 0$ will be slightly favoured due to the third jet having access to central rapidity regions, where phase space is larger.

A similar study has also been done by the D0 collaboration, but instead using $\beta = \tan^{-1}(\text{sgn}(\eta_2)\Delta\varphi/\Delta\eta)$ as their observable [82]. The overall conclusions of both of these studies is that parton shower MCs that implement colour coherence, either through the choice of evolution variable or through an angular veto are able to describe data.

First, a comparison is made against the $2 \rightarrow 3$ ME, using the second set of cuts from the previous section. In this kinematic region, the softest jet is expected to come from radiation and should be relatively soft compared to the underlying $2 \rightarrow 2$ process, where the parton shower is expected to be most accurate. At the parton level, many of the cuts used in the full experimental analysis are no longer relevant, and so the only restriction placed is that $0.5 < R_{23} < 1.0$. The lower cut, although no longer necessary for clustering reasons, is included to enhance the effect of ISR by removing some contributions from FSR, which will be highest in the region of low R_{23} . The upper cut helps to isolate the region where the azimuthal weighting should have the biggest effect ($r \sim 1$, Sec. 2.3.2).

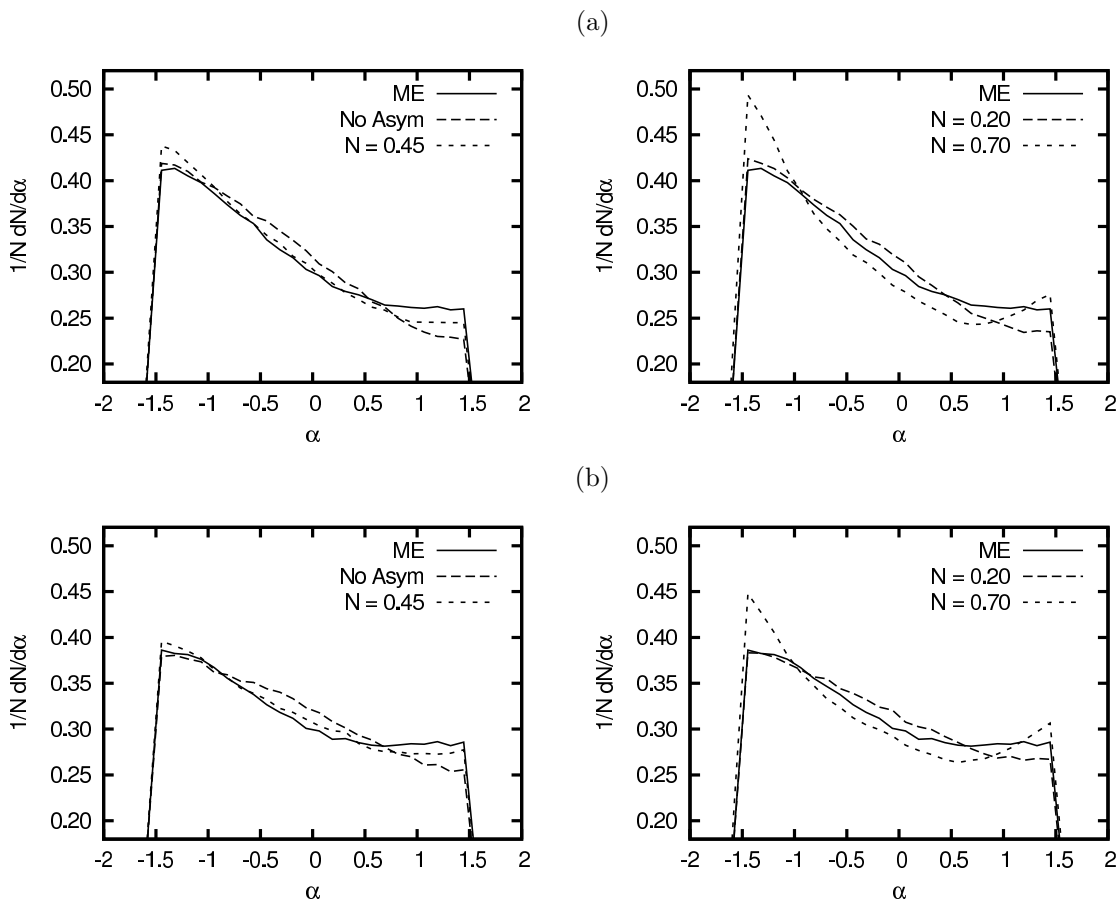


Figure 4: α -angle comparison for different strengths of asymmetry at (a) Tevatron ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV) and (b) LHC (pp , $\sqrt{s} = 14$ TeV) energies

Results are shown in Fig. 4 for different strengths of asymmetry, N , at (a) Tevatron ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV) and (b) LHC (pp , $\sqrt{s} = 14$ TeV) energies. For the PS, it is clear that when all initial-state azimuthal weighting has been switched off, there is an excess of events in the region of small $|\alpha|$. When the weighting is switched on, the results are as expected; events are shifted towards the endpoints of the distribution, with the effect growing as N is increased. One unexpected feature is the agreement of the PS and ME at $\alpha = -\pi/2$ when no azimuthal weighting is switched on. This suggests that ISR emitted toward the central region already has some natural bias towards small $\Delta\varphi$. The extra weighting, then, overshoots in this region, but does have the desired effect at $\alpha = \pi/2$, where the slope of the tail is brought up. In moving to higher energies, it should be noted that the same cuts are used, meaning phase space effects will affect the exact shape of the distribution and the balance between the positive and negative α regions. Considering only the first emission, a value of $N = 0.45$ gives a reasonable agreement across the whole α range and for both energies.

This is not the whole story however; before event generation is complete, there is still the rest of the shower evolution and additional non-perturbative effects to be added. Afterwards, some memory of the parton-level asymmetries should remain, which, in the α -angle observable, should show up in the same way as in the ME comparison. Additionally, at the hadron level we also study the the pseudorapidity distribution of the third jet (no explicit

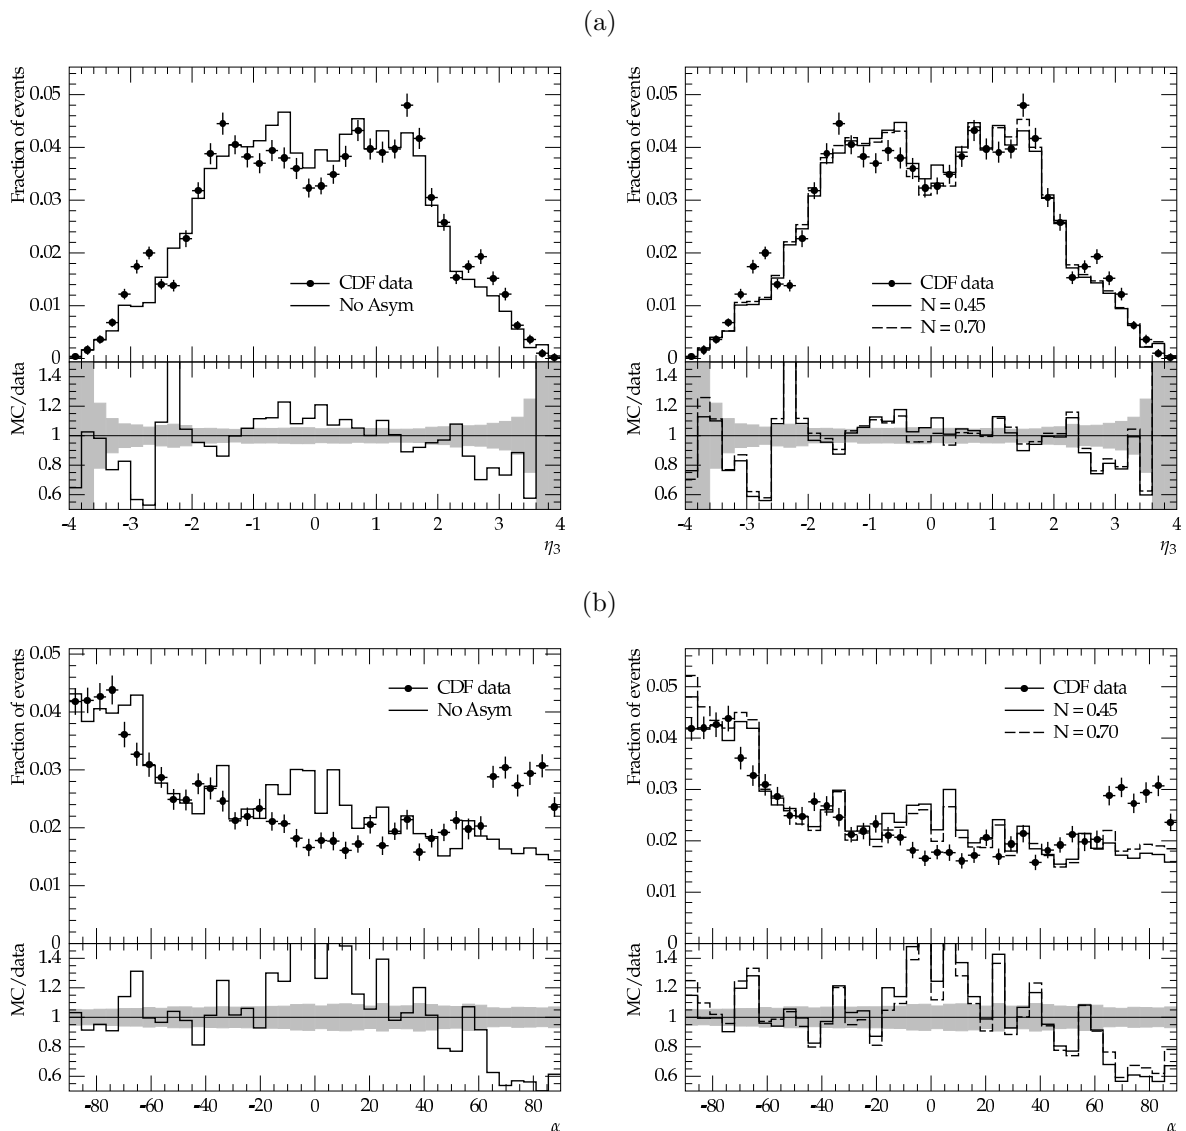


Figure 5: Rivet CDF α -angle analysis ($p\bar{p}$, $\sqrt{s} = 1.80$ TeV)

R_{23} separation cut is applied to this observable). It has been shown that colour coherence effects lead to a broadening in this distribution, as jets are more likely to be emitted at higher rapidity; a good agreement is also a sign that coherence effects are well modeled.

The default tune of PYTHIA 8.142 is used, and only those settings related to the azimuthal asymmetries in ISR changed. The output of the generator is fed through the CDF α -angle Rivet analysis and bin-by-bin detector corrections added to the final results. The results are given in Fig. 5, for (a) η_3 and (b) α . The width of the η_3 distribution is already well described, even without the extra azimuthal weighting. Increasing the strength of the asymmetry does lead to a very slight broadening of the distribution, but the largest difference is in the dip at $\eta_3 = 0$. In general, the high- p_\perp jets will prefer to sit at central rapidities. Biasing $\Delta\varphi$ to smaller values, in combination with the implicit R separation cut from the jet algorithm, will shift the third jet slightly further out in rapidity.

The results for the α -angle are similar to the ME comparison. At the smallest values,

there is a small overshoot, while at higher values, the agreement is improved. At $\alpha = \pi/2$, the tail does not come up to the levels of the data, but is no longer falling. Indeed, with the weighting that has been introduced, bringing the tail up to match the data is a hard task, indicating that additional correlations in rapidity may be involved in later stages of the event-generation process. Based on the CDF data, then, a stronger asymmetry is favoured, and a value $N = 0.70$ is adopted as the default value for all the following sections.

3.4 Kinematic distributions

The main freedom in adjusting the PS to better fill the phase space is given by the starting scale of the initial- and final-state showers. For those hard processes which do not contain any final-state particles which may be shower produced, the shower can begin at the kinematical limit, so as to populate the full phase space without risk of double counting. For the QCD processes considered here, where the hard process already contains particles which may be produced in the shower, double counting becomes an issue and the shower is instead started at the factorisation scale of the hard process. In the case of a $2 \rightarrow 2$ process, this defaults to the smaller of the squared transverse masses of the two outgoing particles. Neglecting parton masses, this implies that the p_{\perp} of the shower jets are below the $2 \rightarrow 2$ p_{\perp} scale. This appears to be a reasonable starting point, but note that the recoil of an emitted parton can boost one of the original two jets to have a smaller p_{\perp} than the emitted jet. Effects can also go in the opposite direction, where a parton emitted at the hard p_{\perp} scale boosts both original partons to a higher p_{\perp} , opening an unfilled region.

Results are shown using the default showers for the three sets of cuts defined previously, for the transverse momenta, rapidities and R separations of the three jets. The first set covers as large an area of phase space as possible, while the second isolates a region where the parton shower is expected to perform well. The final set should isolate a region where the three jets are all relatively hard, and well separated, exactly the region where the parton shower is not expected to perform well. In the plots that follow, we stress that the main aim is for a good qualitative agreement with the matrix elements.

For the second set of cuts, results are shown in Fig. 6 for Tevatron energies ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV) and in Fig. 7 for LHC energies (pp , $\sqrt{s} = 14$ TeV). While at the lower energy, the PS rate is slightly below that of the ME over most of the phase space, and at the higher, the ME is slightly above, all distributions are well reproduced. The small excesses in events at low $p_{\perp 4}$, $p_{\perp 5} > 25$ GeV and $R_{45} > \pi/2$ are consistent with a slightly too high contribution from hard- and widely-separated jets, which we examine further below.

In Fig. 8, the results for the third set of cuts at Tevatron energies are shown. In this region of hard, widely separated jets, the PS is not expected to perform so well. Indeed, here the PS rate is somewhat larger than the ME one. As previously, there is a noticeable excess in events at low $p_{\perp 4}$, suggesting configurations where a softer $2 \rightarrow 2$ process is shifted to meet the jet requirements through an ISR emission and its recoil. The rapidity distribution y_4 is also too broad. A simple check shows that there is no large contribution from an ISR emission becoming harder than the original two jets. That the shapes of $p_{\perp 5}$ and y_5 are relatively well described suggests that there are effects arising from the rotation and boosts employed in the recoil handling, also giving an excess in the tail of ΔR_{45} . We do not study this further at this time. Although the distributions are by no means perfect, we content ourselves with the fact that these events are in a disfavoured region of phase

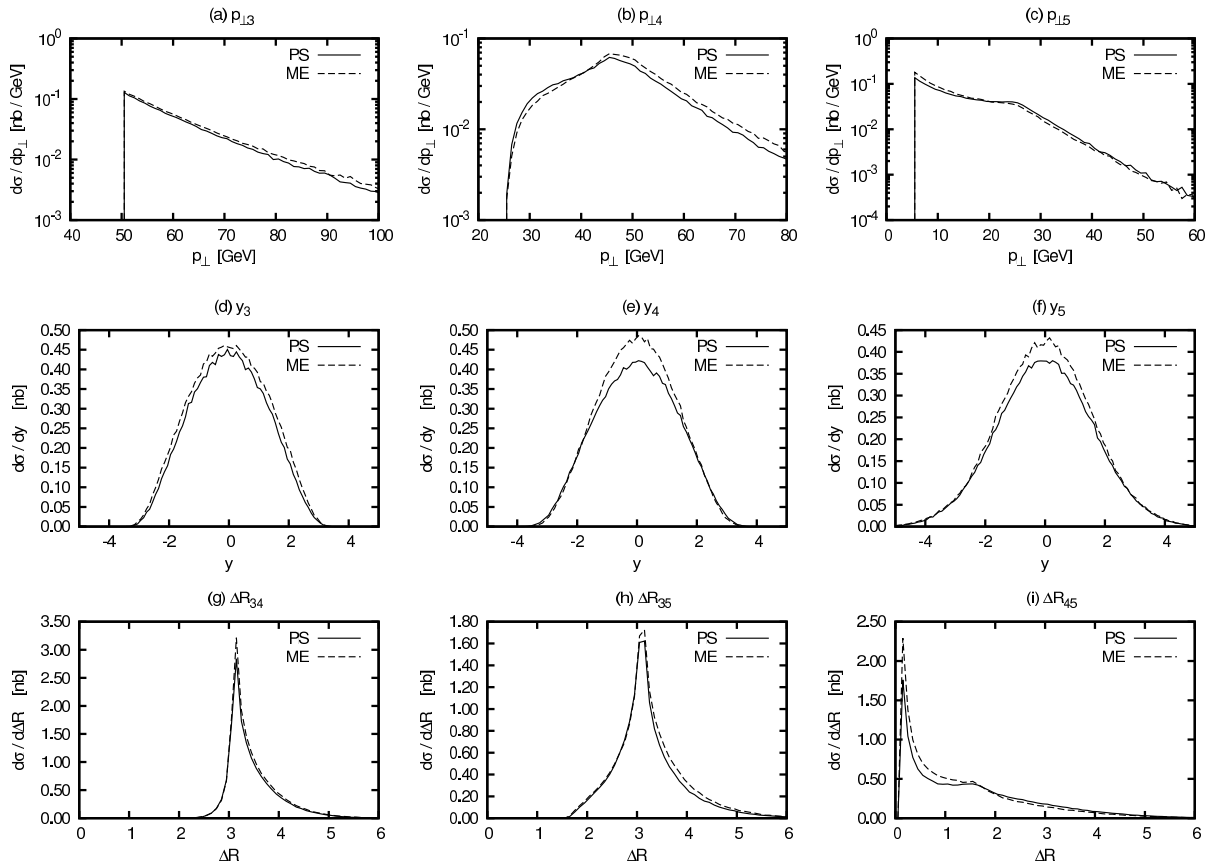


Figure 6: Kinematic distributions for cut set (2) at Tevatron energies ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV): $p_{\perp 3}^{\min} = 50.0$ GeV, $p_{\perp 5}^{\min} = 5.0$ GeV, $R_{\text{sep}}^{\min} = 0.1$

space and that the description gives the broad features of the shapes.

In Fig. 9, the results for the first set of cuts are shown at Tevatron energies. This cut set is designed to be as inclusive as possible, representative of the MPI that are important for MB and UE physics. The distributions show a mixture of the features of the second and third cut sets as shown previously. In general, both the rates and shapes are well reproduced, a sign of the dominance of the soft and collinear regions of phase space.

Some freedom to probe the ambiguities of the shower starting scale is offered by the possibility of introducing a constant factor, k , such that $p_{\perp \text{max}} = k * Q_{\text{fac}}$. This constant may be set separately for the initial- and final-state showers, and in all results shown so far has been set to $k = 1$. In Fig. 10, the p_{\perp} of the softest jet is shown for each set of cuts at Tevatron energies when $k = 0.8$ and $k = 1.2$, both for the initial- and final-state shower. Changes in k give variations both in rate and shape, and the agreement with the ME result is dependent on both the exact cuts used and the energy. Although scale variations away from $k = 1$ may give better agreement in certain regions of phase space and at certain energies, there is no indication that a systematic shift, higher or lower, would improve the overall PS description.

Finally, as an example of the impact of these variations on final observables, the jet-jet azimuthal-angle distribution, as studied by D0, is shown in Fig. 11. Other than the different factors of k , all other settings remain at their default value. As with many other

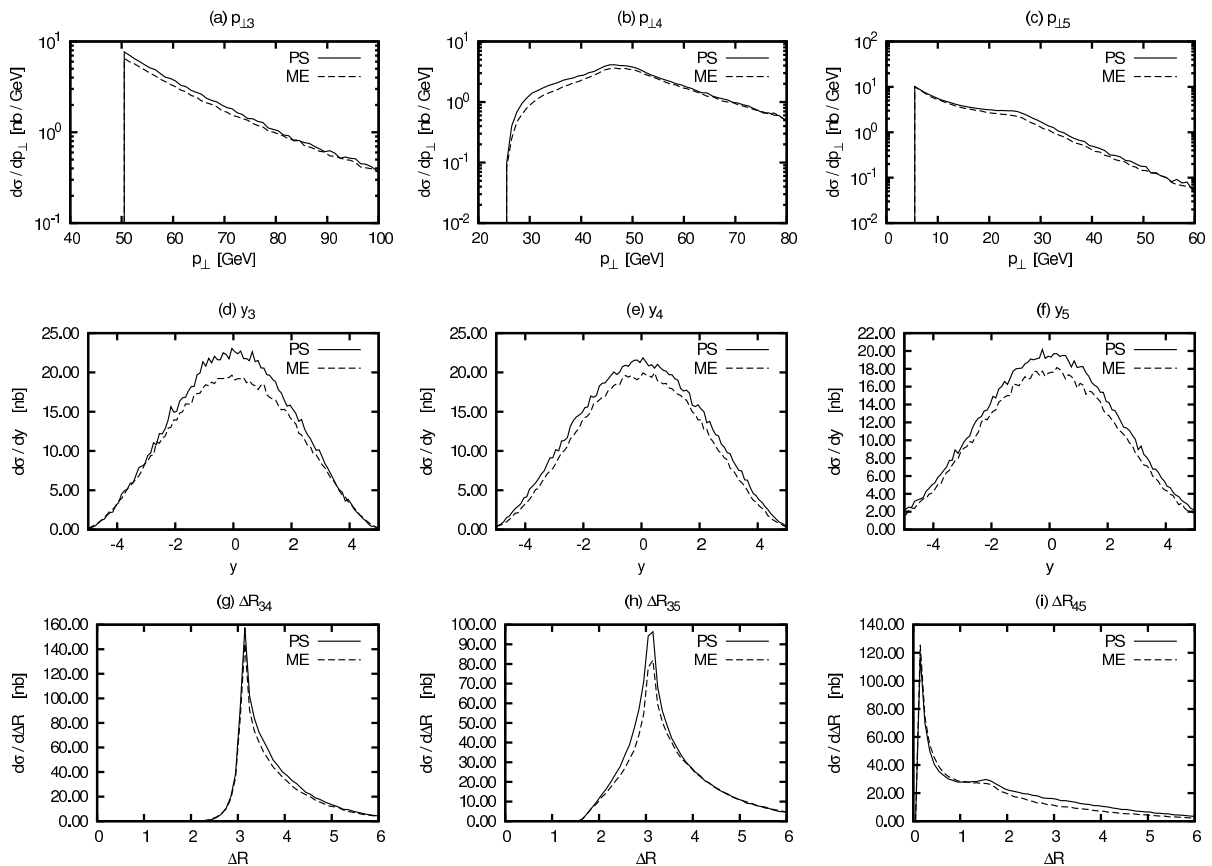


Figure 7: Kinematic distributions for cut set (2) at LHC energies (pp , $\sqrt{s} = 14$ TeV): $p_{\perp 3}^{\min} = 50.0$ GeV, $p_{\perp 5}^{\min} = 5.0$ GeV, $R_{\text{sep}}^{\min} = 0.1$

observables, many different parameters in the overall framework will play a role, and these must then be adjusted to give the best description of the data. The shower k factors could be varied in the context of a more complete tuning to data, at the cost of an increase in the size of the parameter space, but based on the evidence above we choose to keep $k = 1$ in all that follows.

4 Tunes to Tevatron data

In this section, having “validated” the now-modified PS framework in the previous section, we study the tuning prospects of PYTHIA 8 to Tevatron data. All changes to the framework outlined previously apply solely to the case of hadron collisions, so the existing LEP tune of the final-state shower and hadronisation can remain unchanged. However, even the subset of parameters relating only to hadron collisions is large, so the strategy is to pick a minimal set of key parameters, and try to find a region in this parameter space that describes key pieces of Tevatron data. One key goal is to show that a combined MB/UE tuning of Tevatron data is within reach for PYTHIA 8, as is possible in previous versions. It should be stressed that this is not designed to be a “complete” tuning of the generator. The search of this limited parameter space is guided by the principles we outline below, but there is still no guarantee that other regions of parameter space do not give a similar or better

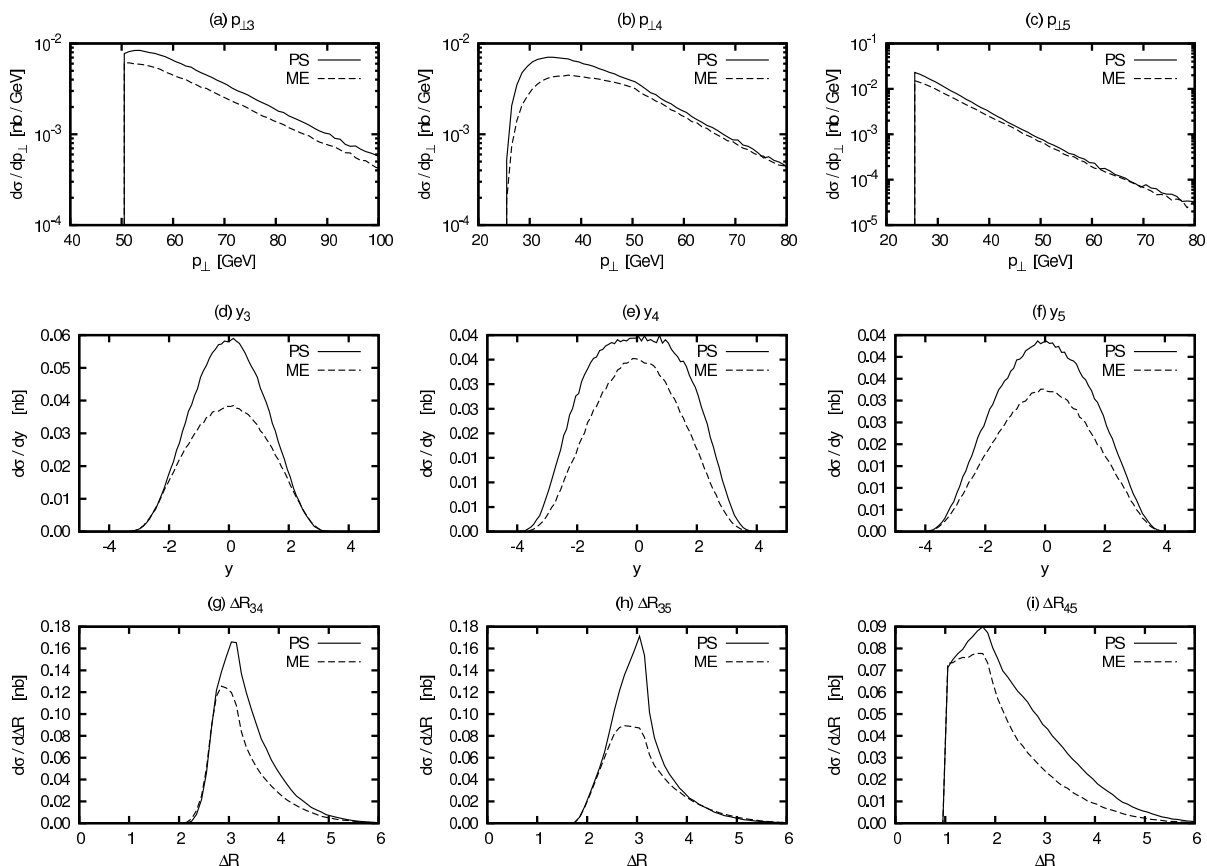


Figure 8: Kinematic distributions for cut set (3) at Tevatron energies ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV): $p_{\perp 3}^{\min} = 50.0$ GeV, $p_{\perp 5}^{\min} = 25.0$ GeV, $R_{\text{sep}}^{\min} = 1.0$

descriptions of the data. We hope that what we present here offers a good starting point for a more global Professor-type tuning, encompassing a wider set of data and parameters, as well as offering a general guide to those getting involved with such efforts.

We also take this opportunity to move away from the CTEQ5L PDF set. As explained in Sec. 2.4.3, a selection of newer sets is now available directly in PYTHIA 8, and we pick one standard LO set (CTEQ6L1) and one modified set (MRST LO**) with which to begin the tuning attempts.

The following Rivet analyses have been used in the tuning process:

CDF_2000_S4155203: Z pT measurement in Run I $Z \rightarrow e^+e^-$ events [83]

p_{\perp} and cross section measurement of e^+e^- pairs in the region of $66 < m_{ee} < 116$ GeV.

CDF_2001_S4751469: Field & Stuart Run I underlying event analysis [84]

The direction of the leading charged particle jet in an event is used to define three regions in $\eta - \varphi$ space. With $\Delta\varphi = \varphi - \varphi_{\text{LeadingJet}}$, the toward region is defined by $|\Delta\varphi| < 60^\circ$ (centered on the leading jet), the away region by $|\Delta\varphi| > 120^\circ$ and the transverse region by $60^\circ < |\Delta\varphi| < 120^\circ$. All regions are constrained such that $|\eta| < 1$. Two datasets are used; one using a minimum bias trigger, the other using a “JET20” trigger, where one calorimeter tower cluster must have $E_T > 20$ GeV. Results are given as a function of the transverse momentum of the leading jet, p_{\perp}^{lead} .

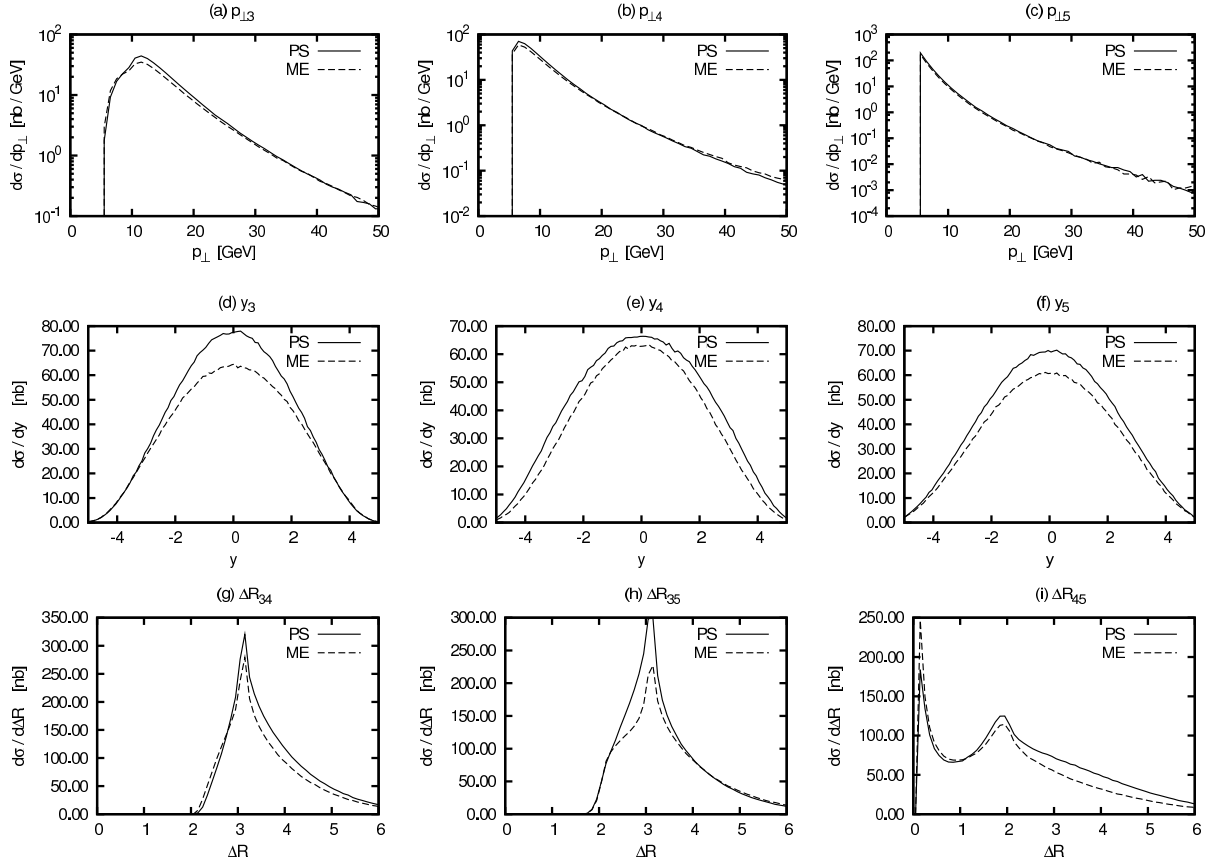


Figure 9: Kinematic distributions for cut set (1) at Tevatron energies ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV): $p_{\perp 3}^{\min} = 5.0$ GeV, $p_{\perp 5}^{\min} = 5.0$ GeV, $R_{\text{sep}}^{\min} = 0.1$




Figure 10: $p_{\perp 5}$ for each set of cuts at Tevatron energies ($p\bar{p}$, $\sqrt{s} = 1.96$ TeV) with $k = 0.8$ and $k = 1.2$

CDF_2002_S4796047: Run I charged multiplicity measurement [85]

Charged multiplicity measurements at $\sqrt{s} = 630$ & 1800 GeV. The $\langle p_{\perp} \rangle (N_{\text{ch}})$ measurements have largely been superseded by CDF_2009_S8233977, below.

D0_2004_S5992206: Run II jet azimuthal decorrelation analysis [86]

Distributions of $\Delta\varphi = \varphi_1 - \varphi_2$, where φ_1 is the azimuthal angle of the hardest jet and φ_2 that of the second hardest jet. Results are given in 4 different bins of $p_{\perp 1}$.

CDF_2008_LEADINGJETS: Run II underlying event in leading jet events [87]




Figure 11: Jet-jet azimuthal-angle distribution showing the changes when the shower starting scale is adjusted by a factor $k = 0.8$ and $k = 1.2$

Regions are defined as in CDF_2001_S4751469, but with the two transverse regions additionally separated into “transMIN” and “transMAX” regions, based on which contains the lowest/highest number/ $\sum p_{\perp}$ density, event-by-event. A “transDIF” region is defined by the difference between the transMAX and transMIN regions.

CDF_2008_NOTE_9351: Run II underlying event in Drell-Yan [88]

As CDF_2008_LEADINGJETS, but using a leading Z^0 reconstructed from a Drell-Yan lepton pair.

CDF_2009_S8233977: Run II min bias cross-section analysis [89]

Measurements of track p_{\perp} , $\langle p_{\perp} \rangle (N_{\text{ch}})$ and $\sum E_{\text{T}}$ for $\sqrt{s} = 1.96$ TeV.

Our limited parameter space consists of the following:

SigmaProcess:alphaSvalue

$\alpha_s(M_Z)$ for the hard process. Directly affects the single-particle p_{\perp} spectrum in the low- p_{\perp} region and jet cross sections at high p_{\perp} . For MRST LO**, a decrease of $\sim 15\%$ is expected due to similarly-sized increase in the PDF content (the QCD $2 \rightarrow 2$ cross section is quadratic both in PDFs and in α_s).

SpaceShower:alphaSvalue

$\alpha_s(M_Z)$ for the initial-state shower. Together with the regularisation parameters (see below), gives control over the amount of ISR activity. Although convenient to be set equal to α_s of the hard process, above, this is not a requirement.

SpaceShower:pt0Ref, SpaceShower:ecmRef, SpaceShower:ecmPow,

SpaceShower:samePTasMI, SpaceShower:pTmin

Parameters used to regularise the $p_{\perp} \rightarrow 0$ divergence in the initial-state shower (Sec. 2.3.1). While the smooth dampening of eq. (14) is physically preferable to a step-function regularisation, as provided by **SpaceShower:pTmin**, there is currently no real constraint on these options. For now, we choose to fix the $p_{\perp 0}^{\text{ref}}$ to an energy independent value, rather than use the same running as used in the MI framework. Together with $\alpha_s(M_Z)$ above, these values are tuned mainly based on the $p_{\perp}(Z^0)$ spectrum (although this also contains large contributions from primordial k_{\perp}) and the D0

jet azimuthal decorrelation analysis. While a change of $p_{\perp 0}$ will give more activity only at small p_{\perp} 's, a change to α_s will affect the whole spectrum, although still not evenly due to its running.

SpaceShower:rapidityOrder

If switched on, rapidity-unordered ISR emissions, subsequent to the first emission off each dipole, are vetoed (Sec. 2.3.1). Studies suggest that having this option activated helps to dampen the rise of the underlying event, leading to a better agreement with data.

MultipleInteractions:alphaSvalue

$\alpha_s(M_Z)$ for the MPI framework. This can be used to help control the amount of MPI activity, in conjunction with the regularisation parameters (see below). In the context of the MPI model there is no strong reason to vary this independently of **SigmaProcess:alphaSvalue**, and we choose to restrict them to be equal or very close to each other. As with ISR, a change to α_s will affect the whole spectrum of MPI as opposed to just the soft region. This balance can have effects on e.g. $\langle p_{\perp} \rangle(N_{\text{ch}})$ distributions, which in turn depends heavily on colour reconnection effects.

MultipleInteractions:pt0Ref, MultipleInteractions:ecmRef, MultipleInteractions:ecmPow

Parameters for the regularisation of the MPI framework (Sec. 2.4.2). $p_{\perp 0}$ is a key ingredient in getting both the charged multiplicity and the underlying-event distributions correct. As we only consider Tevatron data here, the energy dependence is only constrained by the CDF charged multiplicity at $\sqrt{s} = 630$ GeV. The reference energy is left unchanged at 1800 GeV.

MultipleInteractions:bProfile, MultipleInteractions:expPow

Parameters used to select the matter profile used in the MPI framework (Sec. 2.4.2). As noted previously, the introduction of showers off all scattering subsystems reduces the need for a double Gaussian profile; instead we choose the overlap function, eq. (32), giving one degree of freedom, $E_{\text{exp}}^{\text{pow}}$. This parameter is important in the same distributions as $p_{\perp 0}$ above; for charged multiplicities it affects high-multiplicity tail, while it also controls how quickly the underlying-event activity grows as a function of the transverse momentum of the leading jet in MB events. A value $E_{\text{exp}}^{\text{pow}} = 1$ gives results somewhat similar to the default double-Gaussian matter profile, giving more fluctuations out in the high-multiplicity tail and a faster rise of the underlying event, while $E_{\text{exp}}^{\text{pow}} = 2$ reduces back to the single Gaussian case. We examine values between these two limits to try to find the best balance.

BeamRemnants:reconnectRange

The R parameter of eq. (34). This is primarily tuned to the latest measurement of $\langle p_{\perp} \rangle(N_{\text{ch}})$ from Run II Tevatron data, but should also give reasonable results for the lower energy runs. It should be noted that this parameter is non-linear; at $R \sim 5$, the amount of reconnection has almost completely saturated.

Not included in the above list are those parameters related to primordial k_{\perp} (Sec. 2.4.4). For the hard component, early studies indicated that the $p_{\perp}(Z^0)$ peak was well described

| Parameter | Tune 2C | Tune 2M | Tune 4C |
|----------------------------------|---------|---------|---------|
| SigmaProcess:alphaSvalue | 0.135 | 0.1265 | 0.135 |
| SpaceShower:rapidityOrder | on | on | on |
| SpaceShower:alphaSvalue | 0.137 | 0.130 | 0.137 |
| SpaceShower:pT0Ref | 2.0 | 2.0 | 2.0 |
| MultipleInteractions:alphaSvalue | 0.135 | 0.127 | 0.135 |
| MultipleInteractions:pT0Ref | 2.320 | 2.455 | 2.085 |
| MultipleInteractions:ecmPow | 0.21 | 0.26 | 0.19 |
| MultipleInteractions:bProfile | 3 | 3 | 3 |
| MultipleInteractions:expPow | 1.60 | 1.15 | 2.00 |
| BeamRemnants:reconnectRange | 3.0 | 3.0 | 1.5 |
| SigmaDiffraction:dampen | off | off | on |
| SigmaDiffraction:maxXB | N/A | N/A | 65 |
| SigmaDiffraction:maxAX | N/A | N/A | 65 |
| SigmaDiffraction:maxXX | N/A | N/A | 65 |

Table 3: Parameters for Tunes 2C, 2M and 4C (described in Sec. 5). Any parameter not shown is set to the default value of PYTHIA 8.142

with its default setting ($\sigma_{\text{hard}} = 2 \text{ GeV}$), while for the soft component, variations showed little sensitivity in the mean p_{\perp} per parton. In Tab. 3, the parameters for two tunes, 2C (CTEQ6L1) and 2M (MRST LO**) are given. Any parameter not shown is set to the default value of PYTHIA 8.142.

Selected plots are shown in Figs. 12 & 13 for tunes 2C and 2M respectively. The tunes are compared against data and PYTHIA 6 tunes Pro-Q20 and Perugia 0 for (left-to-right, top-to-bottom):

- 1) CDF_2002_S4796047: Charged multiplicity at $\sqrt{s} = 630 \text{ GeV}$
- 2) CDF_2002_S4796047: Charged multiplicity at $\sqrt{s} = 1800 \text{ GeV}$
- 3) CDF_2008_LEADINGJETS: Transverse charged particle density
- 4) CDF_2008_LEADINGJETS: Transverse charged $\sum p_{\perp}$ density
- 5) CDF_2009_S8233977: $\langle p_{\perp} \rangle (N_{\text{ch}})$
- 6) D0_2004_S5992206: Jet-jet azimuthal angle for $100 < p_{\perp}^{\text{max}} < 130 \text{ GeV}$

It is notable that less colour reconnection is needed to match the $\langle p_{\perp} \rangle (N_{\text{ch}})$ data than in the old default. One contribution to this is the increased α_s in the MPI framework, giving more activity at all p_{\perp} values rather than just extra low- p_{\perp} activity from a change in $p_{\perp 0}$. The good agreement in the D0 jet-jet azimuthal distribution gives a good sign that the balance between contributions from ISR and MPI is well described, supplemented by $p_{\perp}(Z^0)$ measurements, not shown.

There is a large difference in the $E_{\text{exp}}^{\text{pow}}$ parameter between the two tunes. In fact, in Tune 2M, the multiplicities at both 630 & 1800 GeV do overshoot in the tails; the low value of $E_{\text{exp}}^{\text{pow}}$ is required to correctly describe the underlying event in the region $p_{\perp \text{lead}} < 200 \text{ GeV}$. One key reason is that the the MRST PDF set combined with the retuned α_s differs in

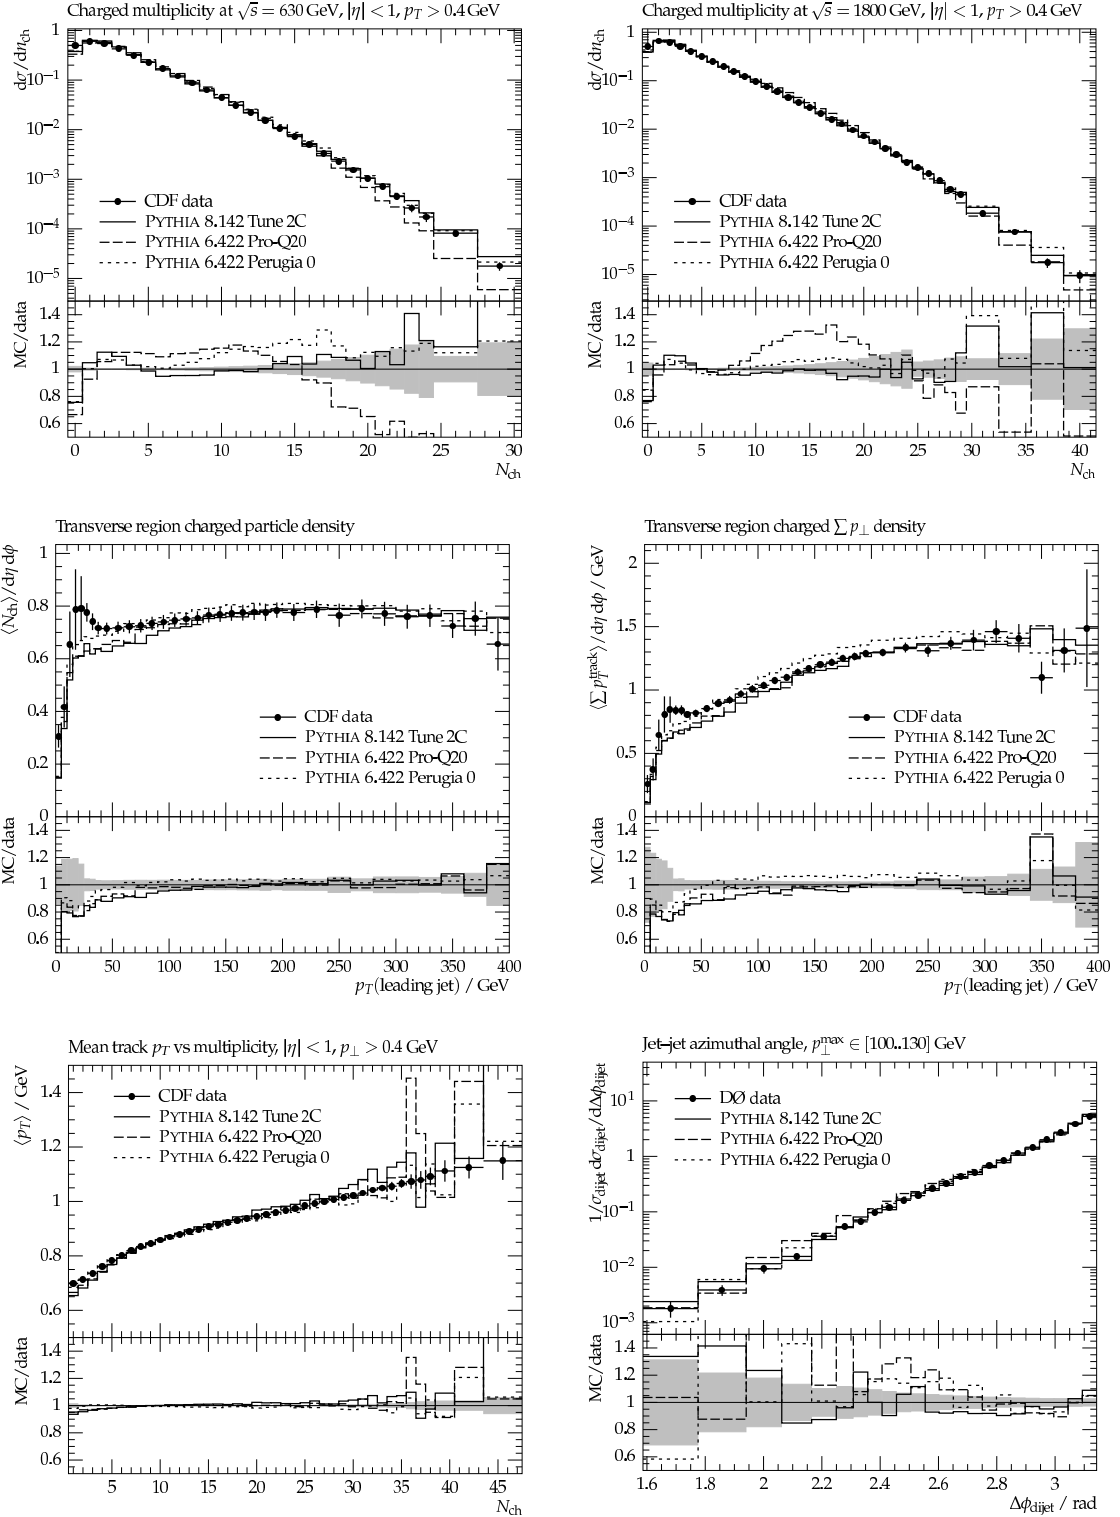


Figure 12: Selected Rivet plots for Tune 2C, compared against data, and tunes Pro-Q20 and Perugia 0

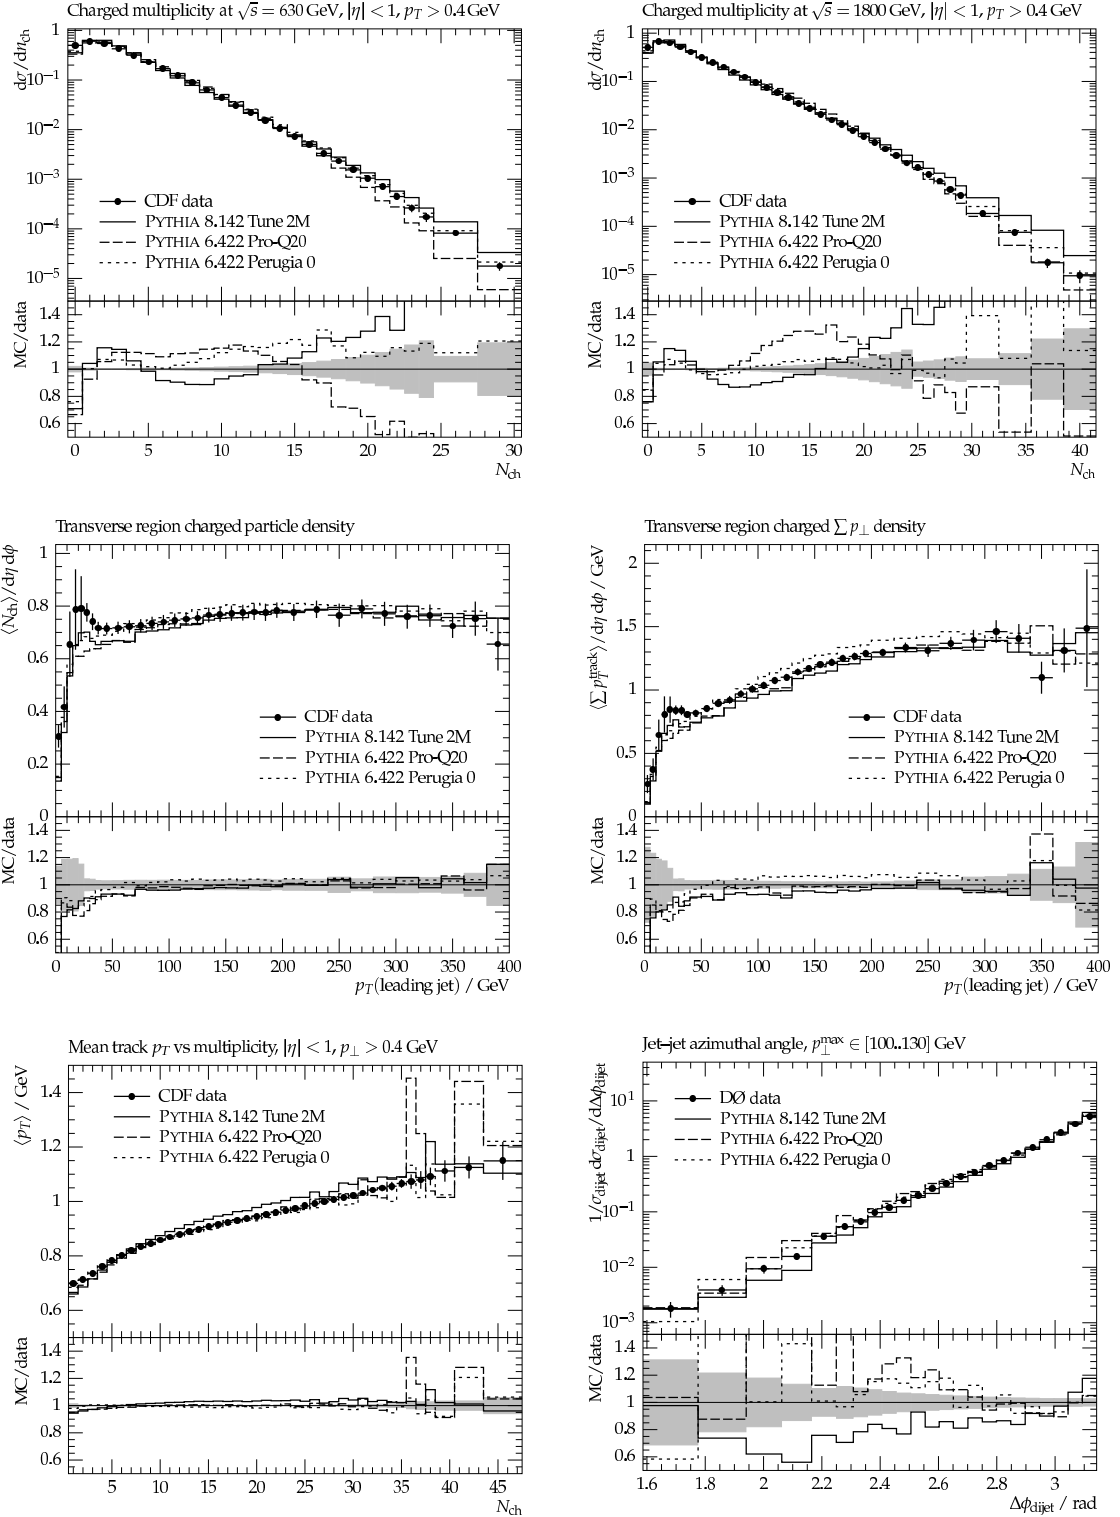


Figure 13: Selected Rivet plots for Tune 2M, compared against data, and tunes Pro-Q20 and Perugia 0

both shape and normalisation relative to CTEQ6L1, and when probing the amount of activity as a function of the leading jet, these differences become more apparent. A value $E_{\text{exp}}^{\text{pow}} \sim 1.4$ would give a better description of the overall multiplicity distributions, but with this setting, the average activity in this $p_{\perp\text{lead}}$ region is lower than for CTEQ6L1. In Tune 2M, we have compensated for this by allowing more impact-parameter fluctuations out to the high multiplicity tails, but expect regions in parameter space to exist which better address this balance.

For both tunes, the transverse activity in the underlying event study no longer shows the large rise noted in Fig. 1. Both tunes also show success in a combined MB/UE description at the Tevatron. Over all datasets, Tunes 2C and 2M never do significantly worse than Pro-Q20 and Perugia 0, and it is hoped that a full tuning, considering a larger parameter set, could improve agreement further.

5 First LHC data

5.1 Diffractive cross sections

For the LHC data, we begin by studying the diffractive cross section and the simple scheme to dampen its growth outlined in Sec. 2.4.1. On the experimental side, we turn to a recent ATLAS study designed to enhance the diffractive content of minimum bias events [60]. The study relies on the Minimum Bias Trigger Scintillators (MBTS), covering two rings in pseudorapidity, $2.09 < |\eta| < 2.82$ and $2.82 < |\eta| < 3.84$. By selecting events with a hit on only one side of the MBTS, diffractive events are preferentially chosen. While only total cross sections are examined here, the ATLAS study also shows that track distributions are better described with the new high-mass diffractive framework of Sec. 2.4.6. It is noted that this study is not corrected for detector/reconstruction effects, but that such corrections are not expected to change the conclusions of the study.

The particular quantity of interest is $R_{\text{ss}} = N_{\text{ss}}/N_{\text{any}}$, where N_{ss} is the number of events with a hit on exactly one side of the MBTS and N_{any} the number with a hit on either side. We follow the approach of the ATLAS study, in using fixed acceptance values for the different event classes while varying the contributions of the different diffractive modes.

From data, a value $R_{\text{ss}} = [4.52 \pm 0.02(\text{stat}) \pm 0.61(\text{syst})]\%$ is preferred, while the default settings from PYTHIA 8 gives a value $R_{\text{ss}} = 5.11\%$, lying just inside the error range. Assuming that the single- and double-diffractive components saturate at the same value, a damping $\sigma_{\text{SD}}^{\text{max}} = \sigma_{\text{DD}}^{\text{max}} = 65$ mb brings R_{ss} into closer agreement with the data. With this value, in this simpleminded model, saturation would then still be far away. Tab. 4 gives the non-, single- and double-diffractive cross sections for different energies in both the default and damped scenarios. R_{ss} values are also provided, but calculated with the same acceptance values as the 7 TeV ATLAS study, and the diffractive to inelastic cross section ratio. In Fig. 14, R_{ss} is shown as a function of the diffractive contribution to the total inelastic cross section. In the damped scenario, the generator R_{ss} now sits upon the central data value.

In Fig. 15, charged multiplicity distributions at $\sqrt{s} = 7$ TeV are shown, giving the breakdown of non-, single- and double-diffractive events, (a) over the whole rapidity range and (b) for $|y| < 2.5$. When the entire phase space is taken into account, the single- and

| | \sqrt{s} (TeV) | σ_{ND} (mb) | σ_{SD} (mb) | σ_{DD} (mb) | R_{ss} (%) | $\sigma_{\text{diff}}/\sigma_{\text{inel}}$ (%) |
|----------------|----------------------------|---------------------------|---------------------------|---------------------------|---------------------|---|
| Default | 0.90 (pp) | 34.4 | 11.7 | 6.4 | 5.5 | 34.5 |
| | 1.96 (p $\bar{\text{p}}$) | 39.0 | 12.5 | 7.5 | 5.4 | 33.8 |
| | 7.00 (pp) | 48.5 | 13.7 | 9.3 | 5.1 | 32.1 |
| Damped | 0.90 (pp) | 36.0 | 10.7 | 5.8 | 5.0 | 31.5 |
| | 1.96 (p $\bar{\text{p}}$) | 40.9 | 11.4 | 6.7 | 4.8 | 30.1 |
| | 7.00 (pp) | 50.9 | 12.4 | 8.1 | 4.5 | 28.7 |

Table 4: Non-, single- and double-diffractive cross sections at different energies for default and damped ($\sigma_{\text{SD}}^{\text{max}} = \sigma_{\text{DD}}^{\text{max}} = 65$ mb) settings. The percentage of same-side events (using the same acceptance values as the 7 TeV study) and ratio of diffractive to inelastic cross sections are also given

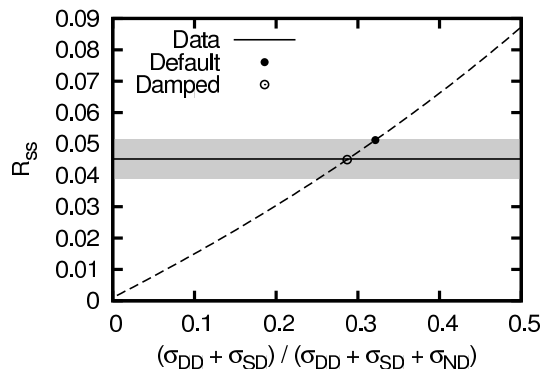


Figure 14: R_{ss} as a function of diffractive contribution to the total inelastic cross section. The dashed curve gives the generator prediction when the ratio is fixed at the default value, $\sigma_{\text{SD}}/\sigma_{\text{DD}} = 1.53$

double-diffractive events dominate in the lower bins, with the non-diffractive component reaching the same level only at $N_{\text{ch}} \sim 30$. In (b), just the rapidity cut is enough to allow the non-diffractive component to rise above the diffractive one down to much lower multiplicities.

It should be remembered that the non-diffractive cross section enters in the MPI framework, eq. (26). This is most clearly seen in the ratio plots; while the diffractive components show a more or less constant drop across the multiplicity range, in (b), the non-diffractive component begins to drop in the high-multiplicity tail. This drop will also occur in (a), but at higher multiplicities than shown. Any retuning to compensate for the effects of the damping must then also take into account the change of slope of the multiplicity distribution.

5.2 Tuning prospects

We now try a more complete tune to early LHC MB/UE data. Where possible, data has been taken from the online HEPDATA database, but certain key datasets are not presently

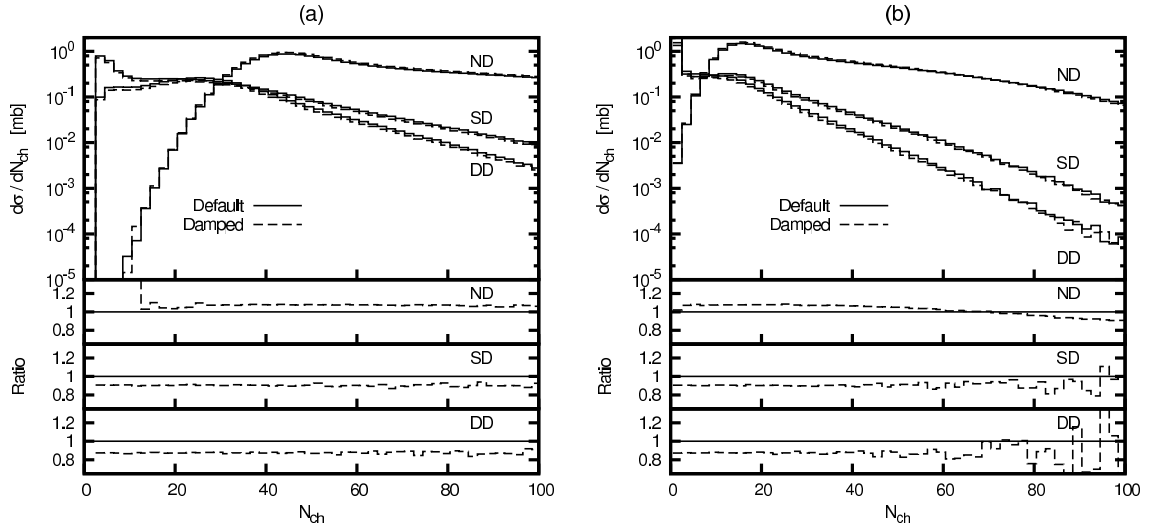


Figure 15: Charged multiplicity distributions at $\sqrt{s} = 7$ TeV showing the breakdown of non-, single- and double-diffractive events, (a) over the whole rapidity range and (b) for $|\eta| < 2.5$

available. For these, instead, data has been read off from the relevant plots available in ATLAS publications. In particular, from HEPDATA, the following datasets have been taken

- 1) ALICE ($|\eta| < 1$) charged multiplicity and rapidity distributions at $\sqrt{s} = 0.90, 2.36$ & 7.00 TeV [90, 91]. The charged multiplicity distributions at $\sqrt{s} = 0.90$ & 2.36 TeV are taken from the INEL dataset, but are shown with the zeroth bin removed. All others are taken from the INEL>0 dataset, where one track in the acceptance region is required to trigger.
- 2) ATLAS ($|\eta| < 2.5, p_{\perp} > 500$ MeV) charged multiplicity, track p_{\perp} , mean p_{\perp} as a function of charged multiplicity and rapidity distributions at $\sqrt{s} = 0.90$ TeV using the INEL>0 dataset [92].

The following data has been taken from publications:

- 1) As (2) above, but at $\sqrt{s} = 7.00$ TeV. Errors for the rapidity distribution are taken, but this was not possible for the remaining observables [93].
- 2) ATLAS ($|\eta| < 2.5, p_{\perp} > 500$ MeV) charged track based underlying event [45]. Charged particle number and sum- p_{\perp} density distributions are taken with errors for the toward, away and transverse regions. A charged track of $p_{\perp} > 1$ GeV in the η acceptance is required to trigger an event.

The results for Tune 2C, introduced in Sec. 4, are shown in Fig. 16, compared against these datasets. Where available, systematic and statistical errors are shown summed in quadrature. At all energies, the multiplicities lie below the data. In the underlying event, the charged number density at 7 TeV undershoots in the toward and away regions while at both energies there is too little activity in the transverse region. The same plots for Tune 2M show the same general features.




Figure 16: Tune 2C compared against early LHC data (see text). Where available, systematic and statistical errors are shown summed in quadrature as the grey band

In attempting to improve agreement with these datasets, we stay with the CTEQ6L1 PDF set, using Tune 2C as a starting point, while turning on the diffractive damping outlined previously. Additionally, we only vary those parameters relating to the MPI framework and colour reconnection (specifically the R parameter of Sec. 2.4.5). By doing this, we ignore the possibility of shifting contributions between ISR and MPI; this balance is perhaps better determined with observables such as $p_{\perp}(Z^0)$ and jet-jet azimuthal correlations. While this procedure is unlikely to give us the best tune possible, it should give a good indication of whether the generator is capable of describing the broad features of the data.

The primary tools available, then, are the $p_{\perp 10}^{\text{ref}}$ and $E_{\text{exp}}^{\text{pow}}$ parameters of the MPI framework, supplemented by $E_{\text{CM}}^{\text{pow}}$ to control the energy scaling (keeping $E_{\text{CM}}^{\text{ref}}$ fixed at 1800 GeV). As an example of the balance that needs to be found between them, Fig. 17 shows the ATLAS and ALICE charged multiplicity distributions and the transverse region charged




Figure 17: ATLAS and ALICE multiplicities and the transverse region charged particle density when tuned only to the ALICE data (see text)

particle density, when solely tuned to match the ALICE multiplicity data ($p_{\perp 0}^{\text{ref}} = 2.047$, $E_{\text{exp}}^{\text{pow}} = 1.5$ and $E_{\text{CM}}^{\text{pow}} = 0.19$). In order to match the ALICE multiplicities, especially all the way out in the tails (made very slightly more difficult by the diffractive damping, as shown previously), a lower $E_{\text{exp}}^{\text{pow}}$ is needed, but this leads to an overshoot in the ATLAS multiplicity tails and an underlying event that rises too quickly.

In Fig. 18, results are shown for a set of parameters which give a reasonable description to all sets of data, Tune 4C in Tab. 3. In this set, the colour reconnection parameter, R is adjusted to match the ATLAS $\langle p_{\perp} \rangle(N_{\text{ch}})$ data. This decrease, although preferable in terms of describing the data with less non-perturbative corrections, has direct effects on the other MB/UE data; it does help to improve the description of the high-multiplicity tails, but the extra multiplicity leads to a rise in the transverse charged number density, while there is not a similar rise in the $\sum p_{\perp}$ density. The $E_{\text{exp}}^{\text{pow}}$ parameter is pushed out such that it in fact reduces to the single-Gaussian case. With this setting, the rise in the toward and away regions is still slightly too high. While overall, at 7 TeV, the multiplicities are slightly on the high side, the level of activity in the transverse region in the underlying event is slightly too low. As shown in the previous example, getting out to the very high-multiplicity tails of the ALICE data is a hard task.

Fig. 19 shows the charged particle multiplicity at $\sqrt{s} = 1800$ GeV and the transverse region charged particle density when Tune 4C is run against the Tevatron analyses of the previous section. In general Tune 4C gives too much activity, also for the multiplicity at $\sqrt{s} = 630$ GeV. In both the transverse region (where the $\sum p_{\perp}$ density is in fact well described) and overall, the average p_{\perp} per particle sits below the data. This suggests a tension between Tevatron and LHC data, or at the very least a new kind of energy interpolation between the 900 GeV and 7 TeV data, as suggested by R. Field, based on his studies with the PYTHIA 6 generator [36].

6 Summary and outlook

In this article, modest changes to the parton-shower framework have been introduced, including an improved description of dipoles stretching from the final to the initial state in the final-state shower and additional azimuthal correlations in the initial-state shower. To accompany these changes, studies of the first shower emission in QCD events have been made against $2 \rightarrow 3$ matrix elements, and simple tunes to Tevatron and early LHC data have been produced.




Figure 18: Tune 4C compared to LHC data

The azimuthal weighting is primarily based on the expected angular distribution for the emission of soft gluons, but it is known that away from this limit it obtains corrections; hard emissions are free to “write their own rules”. These effects have been studied using the α angle, both in comparison to $2 \rightarrow 3$ matrix elements and Tevatron data. In the comparison against data, the width of the pseudorapidity distribution for the softest jet is well described, even without the additional weighting, although an improvement in the central “dip” is seen. For the α angle, good agreement with $2 \rightarrow 3$ ME’s is noted, but while the agreement to data is improved, it is still not as good as hoped for. Further sources of asymmetry thus have to be found.

While an eventual goal would be to have a matching of the first shower emission in QCD events to matrix elements, we have begun here with some simple kinematic comparisons. Even with such a matching, subsequent emissions would still be handled by the default shower, so it is important to understand its behaviour in as much detail as possible. In

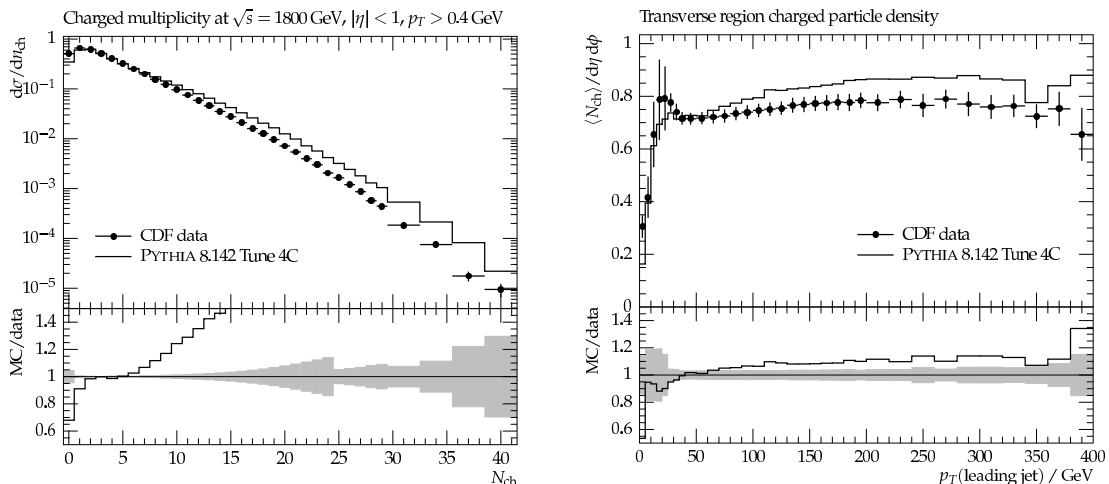


Figure 19: Charged particle multiplicity at $\sqrt{s} = 1800$ GeV and transverse region charged particle density at the Tevatron for Tune 4C

regions where the parton shower is expected to perform well, the description is good both in terms of rate and shape, across a wide range of energies. Away from the soft and collinear regions, there are non-trivial kinematic effects, but still a reasonable overall agreement, generally not worse than could be expected from nontrivial higher orders.

As shown in Sec. 1, the inclusion of FSR in the interleaved PS and MPI framework has not been without some issues. In particular, the simultaneous description of MB and UE data at the Tevatron has been problematic. The changed dipole handling in FSR has turned out to be a key factor, and we have performed a simple hand tuning to the Tevatron data to show that a combined MB/UE description is possible. This has been done for both the CTEQ6L1 and MRST LO** PDF sets. Across all datasets considered, these tunes are never significantly worse than the Pro-Q20 and Perugia 0 tunes of PYTHIA 6.4. These tunes are in no way intended to be the final answer; to simplify the task, we have limited ourselves to a few key parameters and datasets. It is hoped that what is presented here will serve both as a starting point for a more complete global tune to data, and as a general guide to some of the principles involved.

Finally, the first batches of LHC data have been released, including many new MB and UE studies. Starting with the CTEQ6L1 Tevatron tune, and limiting ourselves to changes in MPI and colour reconnection parameters, supplemented by a slight change in diffractive cross sections, a tune that broadly describes the features of much of the data has been produced. It is encouraging that this has been possible given the limited scope of the tuning. It is hoped that, with a more complete tuning effort, PYTHIA will be able to describe this data more accurately.

What is shown in this article is a very early look at this new data, and we look forward to more data from the LHC experiments. New studies to help separate diffractive contributions and their modeling will be welcome. $p_{\perp}(Z^0)$ and jet-jet azimuthal distributions will help to separate the contributions from ISR and MPI, while other MB/UE data may help to resolve some of the tensions seen so far. It remains an open question as to whether the current models will be able to simultaneously describe both Tevatron and LHC data. There may well be a region in the overall parameter space that will be able to do this to some

extent, but we do not rule out differences due to experimental effects or deficiencies in the models, e.g. related to the energy dependence of different parameters.

Acknowledgments

This work was supported by the Marie Curie Early Stage Training program “HEP-EST” (contract number MEST-CT-2005-019626), the Marie Curie research training network “MCnet” (contract number MRTN-CT-2006-035606), and the Swedish Research Council (contract numbers 621-2008-4252 and 621-2007-4157). The authors wish to thank Deepak Kar for his help in understanding the ATLAS UE analysis and Andy Buckley and Hendrick Hoeth for their help with the Rivet framework and analyses. For reading data from plots, the EasyNData tool was used (<http://puwer.web.cern.ch/puwer/EasyNData/>).

References

- [1] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *JHEP* **05** (2006) 026, [arXiv:hep-ph/0603175](#).
- [2] G. Marchesini and B. R. Webber, *Nucl. Phys.* **B238** (1984) 1.
- [3] G. Marchesini and B. R. Webber, *Nucl. Phys.* **B310** (1988) 461.
- [4] V. N. Gribov and L. N. Lipatov, *Sov. J. Nucl. Phys.* **15** (1972) 438–450.
- [5] G. Altarelli and G. Parisi, *Nucl. Phys.* **B126** (1977) 298.
- [6] Y. L. Dokshitzer, *Sov. Phys. JETP* **46** (1977) 641–653.
- [7] V. V. Sudakov, *Sov. Phys. JETP* **3** (1956) 65–71.
- [8] M. Bähr *et al.*, *Eur. Phys. J.* **C58** (2008) 639–707, [arXiv:0803.0883 \[hep-ph\]](#).
- [9] T. Gleisberg *et al.*, *JHEP* **02** (2009) 007, [arXiv:0811.4622 \[hep-ph\]](#).
- [10] S. Catani and M. H. Seymour, *Nucl. Phys.* **B485** (1997) 291–419, [arXiv:hep-ph/9605323](#).
- [11] Z. Nagy and D. E. Soper, [arXiv:hep-ph/0601021](#).
- [12] S. Schumann and F. Krauss, *JHEP* **03** (2008) 038, [arXiv:0709.1027 \[hep-ph\]](#).
- [13] W. T. Giele, D. A. Kosower, and P. Z. Skands, *Phys. Rev.* **D78** (2008) 014026, [arXiv:0707.3652 \[hep-ph\]](#).
- [14] T. Sjöstrand and M. van Zijl, *Phys. Rev.* **D36** (1987) 2019.
- [15] T. Sjöstrand and P. Z. Skands, *Eur. Phys. J.* **C39** (2005) 129–154, [arXiv:hep-ph/0408302](#).
- [16] T. Sjöstrand and P. Z. Skands, *JHEP* **03** (2004) 053, [arXiv:hep-ph/0402078](#).

- [17] G. Gustafson, *Phys. Lett.* **B175** (1986) 453.
- [18] G. Gustafson and U. Pettersson, *Nucl. Phys.* **B306** (1988) 746.
- [19] L. Lönnblad, *Comput. Phys. Commun.* **71** (1992) 15–31.
- [20] T. Sjöstrand, S. Mrenna, and P. Skands, *Comput. Phys. Commun.* **178** (2008) 852–867, [arXiv:0710.3820](#) [hep-ph].
- [21] H. Hoeth private communication, 2009.
- [22] **CDF** Collaboration, R. D. Field, [arXiv:hep-ph/0201192](#).
- [23] J. Alwall *et al.*, *Eur. Phys. J.* **C53** (2008) 473–500, [arXiv:0706.2569](#) [hep-ph].
- [24] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, *JHEP* **11** (2001) 063, [arXiv:hep-ph/0109231](#).
- [25] L. Lönnblad, *JHEP* **05** (2002) 046, [arXiv:hep-ph/0112284](#).
- [26] S. Frixione and B. R. Webber, *JHEP* **06** (2002) 029, [arXiv:hep-ph/0204244](#).
- [27] S. Frixione, P. Nason, and B. R. Webber, *JHEP* **08** (2003) 007, [arXiv:hep-ph/0305252](#).
- [28] S. Frixione, P. Nason, and C. Oleari, *JHEP* **11** (2007) 070, [arXiv:0709.2092](#) [hep-ph].
- [29] C. W. Bauer, F. J. Tackmann, and J. Thaler, *JHEP* **12** (2008) 010, [arXiv:0801.4026](#) [hep-ph].
- [30] N. Lavesson and L. Lönnblad, *JHEP* **12** (2008) 070, [arXiv:0811.2912](#) [hep-ph].
- [31] K. Hamilton and P. Nason, *JHEP* **06** (2010) 039, [arXiv:1004.1764](#) [hep-ph].
- [32] G. Miu and T. Sjöstrand, *Phys. Lett.* **B449** (1999) 313–320, [arXiv:hep-ph/9812455](#).
- [33] E. Norrbin and T. Sjöstrand, *Nucl. Phys.* **B603** (2001) 297–342, [arXiv:hep-ph/0010012](#).
- [34] R. Corke and T. Sjöstrand, *Eur. Phys. J.* **C69** (2010) 1–18, [arXiv:1003.2384](#) [hep-ph].
- [35] **CDF** Collaboration, D. Kar, [arXiv:0905.2323](#) [hep-ex].
- [36] Field, R.D., “recent talks available at http://www.phys.ufl.edu/~rfield/cdf/rdf_talks.html.”.
- [37] **CDF** Collaboration, R. Field and R. C. Group, [arXiv:hep-ph/0510198](#).
- [38] A. Buckley *et al.*, [arXiv:1003.0694](#) [hep-ph].

- [39] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and J. E. von Seggern, *Eur. Phys. J.* **C65** (2010) 331–357, [arXiv:0907.2973 \[hep-ph\]](#).
- [40] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and E. von Seggern, [arXiv:0906.0075 \[hep-ph\]](#).
- [41] P. Z. Skands, *Phys. Rev.* **D82** (2010) 074018, [arXiv:1005.3457 \[hep-ph\]](#).
- [42] A. Moraes, C. Buttar, and I. Dawson, *Eur. Phys. J.* **C50** (2007) 435–466.
- [43] **ATLAS** Collaboration. ATLAS-PHYS-PUB-2010-002 (2010).
- [44] **ATLAS** Collaboration. ATLAS-CONF-2010-031 (2010).
- [45] **ATLAS** Collaboration. ATLAS-CONF-2010-029 (2010).
- [46] T. Sjöstrand, *Phys. Lett.* **B157** (1985) 321.
- [47] G. 't Hooft, *Nucl. Phys.* **B72** (1974) 461.
- [48] Y. L. Dokshitzer, D. Diakonov, and S. I. Troian, *Phys. Rept.* **58** (1980) 269–395.
- [49] M. Ciafaloni, *Nucl. Phys.* **B296** (1988) 49.
- [50] S. Catani, F. Fiorani, and G. Marchesini, *Nucl. Phys.* **B336** (1990) 18.
- [51] B. Andersson, G. Gustafson, and J. Samuelsson, *Nucl. Phys.* **B467** (1996) 443–478.
- [52] B. R. Webber, *Ann. Rev. Nucl. Part. Sci.* **36** (1986) 253–286.
- [53] M. Bengtsson and P. M. Zerwas, *Phys. Lett.* **B208** (1988) 306.
- [54] G. Gustafson and A. Nilsson, *Nucl. Phys.* **B355** (1991) 106–122.
- [55] A. Donnachie and P. V. Landshoff, *Phys. Lett.* **B296** (1992) 227–232, [arXiv:hep-ph/9209205](#).
- [56] G. A. Schuler and T. Sjöstrand, *Phys. Rev.* **D49** (1994) 2257–2267.
- [57] G. A. Schuler and T. Sjöstrand, *Z. Phys.* **C73** (1997) 677–688, [arXiv:hep-ph/9605240](#).
- [58] K. A. Goulianos, *Phys. Lett.* **B358** (1995) 379–388, [arXiv:hep-ph/9502356](#).
- [59] K. A. Goulianos, *Phys. Rev.* **D80** (2009) 111901.
- [60] **ATLAS** Collaboration. ATLAS-CONF-2010-048 (2010).
- [61] **CTEQ** Collaboration, H. L. Lai *et al.*, *Eur. Phys. J.* **C12** (2000) 375–392, [arXiv:hep-ph/9903282](#).
- [62] M. R. Whalley, D. Bourilkov, and R. C. Group, [arXiv:hep-ph/0508110](#).
- [63] T. Kasemets and T. Sjöstrand, [arXiv:1007.0897 \[hep-ph\]](#).

- [64] J. Pumplin *et al.*, *JHEP* **07** (2002) 012, arXiv:hep-ph/0201195.
- [65] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, *Eur. Phys. J.* **C63** (2009) 189–285, arXiv:0901.0002 [hep-ph].
- [66] H.-L. Lai *et al.*, *JHEP* **04** (2010) 035, arXiv:0910.4183 [hep-ph].
- [67] A. Sherstnev and R. S. Thorne, *Eur. Phys. J.* **C55** (2008) 553–575, arXiv:0711.2473 [hep-ph].
- [68] M. Sandhoff and P. Skands. Presented at Les Houches Workshop on Physics at TeV Colliders, Les Houches, France, 2-20 May 2005.
- [69] P. Skands and D. Wick, *Eur. Phys. J.* **C52** (2007) 133–140, arXiv:hep-ph/0703081.
- [70] S. Navin, arXiv:1005.3894 [hep-ph].
- [71] G. Ingelman and P. E. Schlein, *Phys. Lett.* **B152** (1985) 256.
- [72] W. Slominski and A. Valkarova. DESY-PROC-2009-02.
- [73] **CDF** Collaboration, A. A. Affolder *et al.*, *Phys. Rev. Lett.* **84** (2000) 5043–5048.
- [74] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Phys. Rept.* **97** (1983) 31–145.
- [75] **STAR** Collaboration, B. I. Abelev *et al.*, *Phys. Rev.* **C77** (2008) 044908, arXiv:0705.2511 [nucl-ex].
- [76] F. A. Berends, R. Kleiss, P. De Causmaecker, R. Gastmans, and T. T. Wu, *Phys. Lett.* **B103** (1981) 124.
- [77] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, *JHEP* **07** (2003) 001, arXiv:hep-ph/0206293.
- [78] M. L. Mangano, M. Moretti, and R. Pittau, *Nucl. Phys.* **B632** (2002) 343–362, arXiv:hep-ph/0108069.
- [79] F. Caravaglios, M. L. Mangano, M. Moretti, and R. Pittau, *Nucl. Phys.* **B539** (1999) 215–232, arXiv:hep-ph/9807570.
- [80] J. Alwall *et al.*, *JHEP* **09** (2007) 028, arXiv:0706.2334 [hep-ph].
- [81] **CDF** Collaboration, F. Abe *et al.*, *Phys. Rev.* **D50** (1994) 5562–5579.
- [82] **D0** Collaboration, B. Abbott *et al.*, *Phys. Lett.* **B414** (1997) 419–427, arXiv:hep-ex/9706012.
- [83] **CDF** Collaboration, A. A. Affolder *et al.*, *Phys. Rev. Lett.* **84** (2000) 845–850, arXiv:hep-ex/0001021.
- [84] **CDF** Collaboration, A. A. Affolder *et al.*, *Phys. Rev.* **D65** (2002) 092002.

- [85] **CDF** Collaboration, D. E. Acosta *et al.*, *Phys. Rev.* **D65** (2002) 072005.
- [86] **D0** Collaboration, V. M. Abazov *et al.*, *Phys. Rev. Lett.* **94** (2005) 221801, arXiv:hep-ex/0409040.
- [87] **CDF** Collaboration, R. Field. The Underlying Event and Comparisons with MC, First International Workshop on Multiple Partonic Interactions at the LHC, 2008.
- [88] **CDF** Collaboration, D. Kar and R. Field. CDF Note 9351 (2008).
- [89] **CDF** Collaboration, T. Aaltonen *et al.*, *Phys. Rev.* **D79** (2009) 112005, arXiv:0904.1098 [hep-ex].
- [90] **ALICE** Collaboration, K. Aamodt *et al.*, *Eur. Phys. J.* **C68** (2010) 89–108, arXiv:1004.3034 [hep-ex].
- [91] **ALICE** Collaboration, K. Aamodt *et al.*, *Eur. Phys. J.* **C68** (2010) 345–354, arXiv:1004.3514 [hep-ex].
- [92] **ATLAS** Collaboration, G. Aad *et al.*, *Phys. Lett.* **B688** (2010) 21–42, arXiv:1003.3124 [hep-ex].
- [93] **ATLAS** Collaboration. ATLAS-CONF-2010-024.