

Improved Parton Showers at Large Transverse Momenta*

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

Several methods to improve the parton-shower description of hard processes by an injection of matrix-element-based information have been presented over the years. In this article we study (re)weighting schemes for the first/hardest emission. One objective is to provide a consistent matching of the POWHEG next-to-leading order generator to the PYTHIA shower algorithms. Another is to correct the default behaviour of these showers at large transverse momenta, based on a comparison with real-emission matrix elements.

*Work supported by the Marie Curie Early Stage Training program “HEP-EST” (contract number MEST-CT-2005-019626) and in part by the Marie Curie RTN “MCnet” (contract number MRTN-CT-2006-035606)

¹richard.corke@thep.lu.se

²torbjorn@thep.lu.se

1 Introduction

With the start of the LHC in mind, there has been a recent focus on improving the description of event topologies and cross sections, going beyond the Born level for many processes of interest. Firstly, it involves the effects of events with one or more extra jets in the final state, which may affect the impact of background processes and thereby the choice of analysis strategies. Secondly, it includes next-to-leading order (NLO) corrections to production cross sections, which are needed for precision tests of the Standard Model and, hopefully, of physics beyond it.

For the first point, real-emission matrix element (ME) calculations give a good description of hard and widely separated jets, while parton shower (PS) models give the correct behaviour in the soft and collinear regions of phase space. The goal is to find a way to combine these two methods, so that each is used in its region of validity, with a smooth transition between the two in all physical distributions. This is not a trivial task. One key issue is that ME calculations describe inclusive events, while the PS generates exclusive ones. The CKKW [1] method solves this by using ME's supplemented by analytical Sudakov form factors to go from an inclusive to an exclusive language in the hard region, and then switching to a PS below some ME cutoff scale. In CKKW-L [2] the Sudakovs are instead generated from fictitious showers using the same PS algorithm as in the soft region, to improve the consistency and continuity. These methods continue to evolve [3, 4] while other approaches include MLM [5] and Pseudo-showers [6]. Comparisons between the methods have been made [7, 8].

For the second point, the virtual correction terms required at NLO make the calculations more complicated. Analytically the cancellation of real and virtual ME divergences occur in the soft/collinear region, i.e. where we would rather want to use the PS description. The first approach to solve this issue for nontrivial cases was MC@NLO [9, 10]. Here the analytical expression is derived for the phase space population by the first shower emission, in the absence of a Sudakov form-factor correction. The difference between the real-emission ME and this analytical PS expression, which should be finite in the soft and collinear limits, defines the differential cross section for events with one real “ME-based” emission, from which the shower should start. The rest of the cross section, wherein the analytical PS divergences and the virtual divergences cancel, gives the events where the shower should start from the lowest-order process.

The MC@NLO approach has the disadvantage that the PS emission rate may well be higher than the ME one in some parts of phase space, in which case one is forced to introduce negative-weight events. Furthermore, the analytical shower expressions are specific to a particular PS algorithm, so lengthy work has to be redone not only between different generators but also for minor changes inside a given generator [11]. Both of these issues are solved with the POWHEG method [12, 13], where the ME's themselves are exponentiated to provide a process-dependent Sudakov form factor. Thereby, a positive-weight algorithm can be obtained wherein POWHEG (almost) always generates one emission, chosen to be the one with largest transverse momentum. It is then up to the subsequent shower to respect this constraint, but otherwise without the need for a tight connection between the ME and PS stages.

When POWHEG is used with showers that are not p_{\perp} -ordered, then further thought must be given to the interface. Specifically, in HERWIG, with its angular ordered showers

[14,15], the first emission is at largest angle but not necessarily at largest p_{\perp} . This leads to the idea of a “truncated shower” [12], where the shower is modified such that the hardest emission is generated with a modified Sudakov form factor. Subsequent emissions are then generated with the usual algorithm, but with a p_{\perp} veto so that they cannot be harder than the hardest emission.

The POWHEG approach is especially convenient, then, if the shower itself is p_{\perp} -ordered, as is now the case in the PYTHIA generator [16–18]. Nevertheless there are some subtleties that should be taken into account to optimise the interface. We discuss these issues in Section 2, and in particular interface to POWHEG-hvq [19], an event generator for heavy quark production in hadronic collisions at NLO QCD, where some simple comparisons are presented for top and bottom production. We also introduce a “poor man’s POWHEG” for cases where NLO calculations are available, but only in the traditional phase space slicing approach, where shower Sudakovs can be used to provide a smoother matching.

One should note that the more sophisticated the description aimed for, the higher the price in terms of manpower that goes into the detailed simulation of a specific process. The point of injecting ME information is precisely to move away from the universal shower behavior, which means that each new process must be considered “from scratch”. To provide a sensible first estimate, however, it is useful if the shower can be improved to get at least the qualitative behaviour right for a wide range of processes. In Section 3 we use the MadGraph/MadEvent [20] generator to look at further pair-production processes with a jet in the final state. By adding a Sudakov to the cross-section for real emission of a jet, to approximate the prescription used by the POWHEG method, we are able to make comparisons to the PYTHIA shower, and find modest changes that improve agreement significantly.

Specifically, we will address the issue of “power showers” vs. “wimpy showers”. In [21], the authors compare two extreme choices for the maximum emission scale of the parton shower; either the full CM energy of the incoming hadrons (power) or the transverse mass of the particles produced in the hard collision (wimpy), for both virtuality- and p_{\perp} -ordered showers. Their conclusion was that these options bracketed the matrix-element behaviour for the top and SUSY production processes studied, but also that the spread of “predictions” is large.

One must note that an ultimate goal would be to have a matching scheme that allows for both a matching to multiple real emissions and to virtual corrections. Some algorithms have recently been proposed in this direction [22–24], but are not yet at a stage to be used for serious LHC studies. We will not discuss such issues further here.

The outline of the article is to study the POWHEG approach and its relation to PYTHIA in Section 2, to use MadEvent to gain an improved understanding of sensible default shower behaviour in Section 3, and to provide a summary and outlook in Section 4.

2 POWHEG

2.1 The PYTHIA merging strategy

2.1.1 Evolution equations

The PYTHIA parton shower orders final-state radiation (FSR) emissions in terms of an evolution variable Q^2 , such as m^2 (previous PYTHIA versions) or p_\perp^2 (PYTHIA nowadays), with an additional energy-sharing variable z in the branching. For QCD emissions, introducing $t = \ln(Q^2/\Lambda_{QCD}^2)$, the DGLAP evolution equations lead to the probability for a splitting $a \rightarrow bc$

$$d\mathcal{P} = \sum_{b,c} \frac{\alpha_s}{2\pi} P_{a \rightarrow bc}(z) dt dz, \quad (1)$$

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, no emissions can have occurred at a larger Q^2 . The probability that a branching occurs at t is now given by

$$\frac{d\mathcal{P}}{dt} = \left(\sum_{b,c} \mathcal{I}_{a \rightarrow bc}(t) \right) \exp \left\{ - \int_t^{t_{max}} dt' \sum_{b,c} \mathcal{I}_{a \rightarrow bc}(t') \right\}, \quad (2)$$

with

$$\mathcal{I}_{a \rightarrow bc}(t) = \int_{z_{min}(t)}^{z_{max}(t)} dz \frac{\alpha_s(t)}{2\pi} P_{a \rightarrow bc}(z).$$

The introduction of this Sudakov form factor turns the unnormalised distribution into a normalised one, i.e. with unit integral over the full phase space. In practice, a lower cutoff, t_0 , is introduced to keep the shower away from soft/collinear regions, which leads to a fraction of events with no emissions inside the allowed region.

For initial-state radiation (ISR), the evolution is performed using backwards evolution [25], where 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. Such a change leads to a Sudakov with the form

$$S_b(x, t_{max}, t) = \exp \left(- \int_t^{t_{max}} dt' \sum_{a,c} \int dz \frac{\alpha_s(p_\perp^2)}{2\pi} P_{a \rightarrow bc}(z) \frac{x' f_a(x', t')}{x f_b(x, t')} \right). \quad (3)$$

One feature of the above equations is the running renormalisation and factorisation shower scales, i.e. the scales at which α_s and the PDF’s are evaluated. For both ISR and FSR, α_s is evaluated at the p_\perp scale of the emission (the definition of p_\perp in this context is discussed later in Section 2.3). For ISR, the flavour dependent ratio of PDF’s given in eq. (3) is evaluated at the selected t value which, for the current versions of PYTHIA, is p_\perp^2 . Thus, the renormalisation and factorisation scales are the same.

2.1.2 The merging strategy

Probably the first use of an explicit matching of PS to ME, the so-called merging strategy, was introduced to handle the case of three-jet events in e^+e^- annihilation [26]. An outline is given below, starting from the Born cross section σ_B for the lowest-order process $e^+e^- \rightarrow \gamma^*/Z^0 \rightarrow q\bar{q}$.

On the ME side, the real-emission cross section $e^+e^- \rightarrow \gamma^*/Z^0 \rightarrow q\bar{q}g$ is given, for massless quarks, by the well-known expression

$$\frac{1}{\sigma_B} \frac{d\sigma_R}{dx_1 dx_2} = W^{\text{ME}} = \frac{\alpha_s}{2\pi} \frac{4}{3} \frac{x_1^2 + x_2^2}{(1-x_1)(1-x_2)}, \quad (4)$$

with α_s typically evaluated at a scale of s , the invariant mass of the system.

The DGLAP inclusive $q \rightarrow qg$ emission probability in (Q^2, z) space can be mapped onto the (x_1, x_2) space

$$W_q^{\text{PS}} = \frac{\alpha_s(p_\perp^2)}{2\pi} \frac{4}{3} \frac{1}{Q^2} \frac{1+z^2}{1-z} \frac{d(Q^2, z)}{d(x_1, x_2)}. \quad (5)$$

The sum of emissions off the q and \bar{q} gives $W^{\text{PS}} = W_q^{\text{PS}} + W_{\bar{q}}^{\text{PS}}$. With the addition of a Sudakov form factor, as above, this becomes

$$W_{\text{corrected}}^{\text{PS}}(Q^2) = W^{\text{PS}}(Q^2) \exp\left(-\int_{Q^2}^{Q_{\text{max}}^2} dQ'^2 W^{\text{PS}}(Q'^2)\right). \quad (6)$$

For ease of notation we have here omitted the dependence of W^{PS} on z and the need of an integration dz' over a range $z'_{\text{min}}(Q'^2) < z' < z'_{\text{max}}(Q'^2)$ in the exponent.

It now so happens that the PYTHIA shower algorithm covers the full three-jet phase space and that $W^{\text{PS}} > W^{\text{ME}}$ everywhere (so long as the former is true, the latter can always be achieved by a suitable rescaling). One can therefore use the veto algorithm [17] to correct down the emission rate. That is, whenever a trial Q^2 has been selected according to eq. (6), the ratio $W^{\text{ME}}/W^{\text{PS}}$ in the chosen phase space point is the probability that this choice should be retained. If not, the evolution is continued downwards from the rejected Q^2 scale (*not* from Q_{max}^2). This gives a change from eq. (6) to

$$W_{\text{corrected}}^{\text{PS+ME}}(Q^2) = W^{\text{ME}}(Q^2) \exp\left(-\int_{Q^2}^{Q_{\text{max}}^2} dQ'^2 W^{\text{ME}}(Q'^2)\right). \quad (7)$$

Note that, while all explicit dependence on W^{PS} is gone in eq. (7), an implicit dependence on the shower remains in two respects. Firstly, the Sudakov-factor modification of the basic ME shape reflects the order in which the shower algorithm sweeps over phase space, i.e. the shower Q^2 definition. Secondly, if the α_s factors are omitted from the $W^{\text{ME}}/W^{\text{PS}}$ reweighting, the p_\perp^2 -dependent expression used in the shower is retained, instead of the fixed value normal for ME's.

Once the first emission has been considered, an uncorrected shower is allowed to continue downwards from the chosen Q^2 scale. Thus, in this algorithm there is no fixed scale for the transition from ME to PS, but a smooth merging of one into the other. This cannot give rise to discontinuities in the behaviour at any phase space point, except of course at the soft/collinear shower cutoff.

The formalism so far has only considered real emissions. For the e^+e^- case, however, the cancellation of real (R) and virtual (V) divergences results in a particularly simple expression (when integrated over the possible orientations of the event)

$$\sigma_{\text{NLO}} = \sigma_B + \int d\sigma_R + \sigma_V = \left(1 + \frac{\alpha_s}{\pi}\right) \sigma_B. \quad (8)$$

It is therefore trivial to retain eq. (7) as an NLO prescription for the distribution of events, just by raising the cross section associated with each event from σ_B to σ_{NLO} . Identifying $W^{\text{ME}} dQ^2 = d\sigma_{\text{R}}/\sigma_B$, we arrive at the final equation for the differential cross section

$$d\sigma = \sigma_{\text{NLO}} \frac{d\sigma_{\text{R}}}{\sigma_B} \exp\left(-\int_{Q^2}^{Q_{\text{max}}^2} \frac{d\sigma_{\text{R}}(Q'^2)}{\sigma_B}\right). \quad (9)$$

It is thus assumed that the $\mathcal{O}(\alpha_s)$ “new” part, $\sigma_{\text{NLO}} - \sigma_B$, of the total cross section should be associated with the same radiation function as σ_B . Alternative choices on this count would only show up in $\mathcal{O}(\alpha_s^2)$, which is beyond the certified accuracy of the algorithm. Physicswise it is the minimal assumption, relative to having a different radiation function for the pure NLO part of the cross section.

The merging formalism is easy to extend to the case of the emission of one extra gluon in any $1 \rightarrow 2$ decay, and has also been worked out and implemented for all such cases within the Standard Model and the Minimal Supersymmetric Extension to the Standard Model [27]. It can also be applied to ISR for $2 \rightarrow 1$ production of colour singlet particles, such as $q\bar{q} \rightarrow Z^0$ [28]. Here, however, NLO corrections are not so trivial. Experimental practice has been to rescale the total cross section to the NLO answer, as in eq. (9), thereby neglecting other potential kinematical differences between the LO and NLO answers.

The merging approach has also been applied to HERWIG [29–31], with two main differences. Firstly, since the HERWIG shower is ordered in angle rather than hardness, the ME/PS correction weight must be applied to any emission that is the hardest so far, rather than only to the first. Secondly, the HERWIG algorithm leaves holes in the phase space covered by the first shower emission, so it becomes necessary also to introduce a matching procedure, whereby such holes are filled directly by the ME rather than by the PS.

2.2 The POWHEG strategy

A general NLO cross section with hadronic incoming states will also contain remnant (counter) terms from the subtraction of initial-state collinear singularities, which have already been incorporated into the PDF’s. The complete differential cross section can therefore be written as the sum of contributions from leading order (Born), virtual, real and counter terms

$$d\sigma = d\sigma_B + d\sigma_V + d\sigma_R - d\sigma_C. \quad (10)$$

We note that the virtual term has the same n -body phase space as the Born term, while the real and counter terms have an $(n + 1)$ -body phase space. In the POWHEG method, the phase space kinematics are factorised in terms of Born (v) and radiation (r) kinematic variables, such that the overall cross section may now be written as

$$d\sigma = B(v) d\Phi_v + V(v) d\Phi_v + R(v, r) d\Phi_v d\Phi_r - C(v, r) d\Phi_v d\Phi_r. \quad (11)$$

Defining a function that integrates over the radiation variables

$$\bar{B}(v) = B(v) + V(v) + \int d\Phi_r [R(v, r) - C(v, r)], \quad (12)$$

once a Born event has been generated, distributed according to $\bar{B}(v) d\Phi_v$, the differential cross section for the hardest emission may be written as

$$d\sigma = \bar{B}(v) d\Phi_v \left[\frac{R(v, r)}{B(v)} \exp\left(-\int_{p_{\perp}} \frac{R(v, r')}{B(v)} d\Phi'_r\right) d\Phi_r \right]. \quad (13)$$

The evolution variable of the POWHEG “shower” is taken as the kinematical p_\perp of the emission with respect to the parton that branches, identifying this highest p_\perp branching as the hardest. For ISR this is the p_\perp with respect to the beam axis. In this way, the hardest radiation is generated according to exact NLO matrix elements, but in a probabilistic, exclusive language, which can be directly interfaced to a suitable shower program.

This expression shares many features with eq. (9). The integral is from some upper scale associated with the Born event, with a lower cutoff to avoid soft/collinear regions. In both cases it is possible for an event to evolve down to this lower cutoff without radiating. Here, however, the constant NLO prefactor, σ_{NLO} , is upgraded to $\bar{B}(v)d\Phi_v$, a fully differential quantity which will encapsulate all kinematical differences between the LO and NLO answers. It is thus in this term that all the sophistication and hard work of a full NLO calculation lies. As in the PYTHIA shower, the radiation is generated with a running α_s expression, evaluated at the p_\perp scale of the emission.

This formalism has been used for the hadronic production of vector bosons [32,33], heavy quark pairs [34], single tops [35] and Higgs bosons via gluon/vector-boson fusion [36,37]. Codes for generating all these processes, except for vector boson pairs, is publicly available. The latest development is the POWHEG BOX; a general framework for implementing NLO calculations with the POWHEG method [38].

2.3 The PYTHIA transverse-momentum-ordered showers

The PYTHIA 8.1 showers are ordered in transverse momentum [16], both for ISR and for FSR. Also, multiparton interactions (MPI) are ordered in p_\perp [39]. This allows a picture where MPI, ISR and FSR are interleaved in one common sequence of decreasing p_\perp values [18]. 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

$$\begin{aligned} \frac{d\mathcal{P}}{dp_\perp} &= \left(\frac{d\mathcal{P}_{\text{MPI}}}{dp_\perp} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp_\perp} + \sum \frac{d\mathcal{P}_{\text{FSR}}}{dp_\perp} \right) \\ &\times \exp \left(- \int_{p_\perp}^{p_{\perp\text{max}}} \left(\frac{d\mathcal{P}_{\text{MPI}}}{dp'_\perp} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp'_\perp} + \sum \frac{d\mathcal{P}_{\text{FSR}}}{dp'_\perp} \right) dp'_\perp \right) \end{aligned} \quad (14)$$

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. Starting from a single hard interaction, eq. (14) can be used repeatedly to construct a complete parton-level event of arbitrary complexity. Recently also rescattering has been included as a further (optional) component of the MPI framework [40].

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 ? It does not have a simple interpretation in absolute time; all the MPI occur essentially simultaneously (in a simpleminded picture where the protons have been Lorentz contracted to pancakes), while ISR stretches backwards in time (and is handled by backwards evolution

[25]) and FSR forwards in time. The closest would be to view eq. (14) as an evolution towards increasing formation time.

For the following, a relevant aspect is that the p_{\perp} definition is not exactly the same for MPI, ISR and FSR in PYTHIA. 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 problems arise 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 (LC) momentum along the a axis that b takes. Then

$$p_{\perp\text{LC}}^2 = z(1-z)m_a^2 - (1-z)m_b^2 - zm_c^2, \quad (15)$$

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

$$\text{ISR : } p_{\perp\text{evol}}^2 = (1-z)Q^2 \quad \text{with } m_b^2 = -Q^2 \text{ and } m_a^2 = m_c^2 = 0, \quad (16)$$

$$\text{FSR : } p_{\perp\text{evol}}^2 = z(1-z)Q^2 \quad \text{with } m_a^2 = Q^2 \text{ and } m_b^2 = m_c^2 = 0, \quad (17)$$

which are monotonous functions of the virtuality Q^2 . However, once a branching has been found and the kinematics is to be reconstructed, the Q^2 interpretation is retained, but z is now replaced by a Lorentz invariant definition. For ISR this is chosen to be $z = m_{br}^2/m_{ar}^2$, where r is the ‘‘recoiler’’, i.e. the incoming parton from the other side of the subcollision. The actual p_{\perp} of b and c then becomes

$$p_{\perp b,c}^2 = (1-z)Q^2 - \frac{Q^4}{m_{ar}^2} = p_{\perp\text{evol}}^2 - \frac{p_{\perp\text{evol}}^4}{p_{\perp\text{evol,max}}^2}, \quad (18)$$

where $p_{\perp\text{evol,max}}^2 = (1-z)Q_{\text{max}}^2 = (1-z)^2m_{ar}^2$. Instead, for FSR, $z = E_b/E_a$ in the rest frame of a and its colour-connected recoiler. The $p_{\perp b,c}^2$ expression in this case becomes somewhat lengthier than eq. (18), but shares the same physical properties; 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.

2.4 Interfacing POWHEG-hvq to PYTHIA

In this study, the interface from POWHEG-hvq to PYTHIA was performed using the Les Houches Event File (LHEF) [41] format. Both programs come with a number of PDF sets available for selection and PYTHIA gives access to external PDF sets through the LHAPDF library [42]. For consistency, the CTEQ6L [43] PDF set has been used throughout this study, as this set is available in all three programs used; no other differences in the PDF set itself have been taken into account. One caveat is the associated α_s running expression for a given PDF set. PYTHIA by default will use the first order expression, and allow $\alpha_s(M_Z^2)$

to be set independently for spacelike and timelike showers. Here, $\alpha_s(M_Z^2)$ was set to match the one used in the PDF set, but with first order running.

Historically the LHA conventions [44], used by LHEF, only encompass one event “scale”, specified to be the factorisation one. The α_s and α_{em} values at the renormalisation scale are also supplied, but not that scale itself, and there is no provision for shower-matching scales. A proposal for an extended standard that would store more scales has recently been presented [45]. When reading an LHEF into a shower program, the question then is: at what scale should the subsequent shower evolution begin?

In PYTHIA, such choices need to be made for ISR and FSR separately. The basic principle should be to avoid doublecounting to the largest extent possible. With LHEF input it would be natural to let the showers begin at the factorisation scale, the only known one, and then proceed downwards. This is an allowed choice both for ISR and FSR, but for ISR another possibility is default. Here, events are split into two kinds, based on the absence or presence of particles that the shower can produce, i.e. d, u, s, c, b, g and γ . If the LHEF final state contains no such particles then the shower can populate the full phase space without any risk of doublecounting, which should give more realistic event shapes. This would be the case for many W/Z, top, Higgs and New Physics processes. If, on the other hand, the final state does contain particles that could be produced in the shower, then doublecounting would be more likely than not, and the factorisation scale again becomes the only reasonable choice.

This more flexible attitude works well if the ME program does not mix different topologies, but breaks down if, say, both $t\bar{t}$ and $t\bar{t}g/t\bar{t}q$ events are supplied, with the latter intended to correspond to the fraction of $t\bar{t}$ events with an extra emission above the factorisation scale. When such a mixing is present, the showers off $t\bar{t}$ should not be allowed to populate the whole phase space. Thus, we conclude that interfacing to an LHEF cannot be done completely automatically, but must be made with some knowledge of which rules were used to produce the LHEF.

Therefore, to provide a consistent interface to POWHEG, we must both consider its generation strategy and the information it stores in the LHEF. In events with an emission above the lower cutoff scale $p_{\perp\min}$, POWHEG chooses the factorisation scale to be the p_{\perp} of the emission, $p_{\perp\text{POWHEG}} = p_{\perp}$ (defined with respect to the beam axis in the case of POWHEG-hvq).

When the POWHEG “shower” reaches $p_{\perp\min}$ without any emissions, the factorisation scale is instead set equal to this $p_{\perp\min}$, $p_{\perp\text{POWHEG}} = p_{\perp\min}$. Since $p_{\perp\min}$ normally is rather small, of the order of 1 GeV, the fraction of no-emission events is also small.

The LHEF choice of storing α_s and α_{em} rather than the renormalisation scale makes sense for traditional ME calculations, where typically one fixed scale is used. The POWHEG case is somewhat special, as it can have different renormalisation scales for the Born level process and for the subsequent radiative emission.

For the case where POWHEG has already generated an emission, the obvious choice is to begin the evolution at the factorisation scale $p_{\perp\text{POWHEG}}$, such that the shower will not generate harder emissions. However, there are then two potential complications: first, the mismatch between $p_{\perp\text{POWHEG}}$ and the lightcone-inspired $p_{\perp\text{evol}}$ scales of PYTHIA, and second, specifically for FSR, that p_{\perp} is defined with respect to the direction of the emitting parton rather than to the beam axis. For ISR, eq. (18) shows that $p_{\perp} < p_{\perp\text{evol}}$, such that starting the shower from $p_{\perp\text{evol}} = p_{\perp\text{POWHEG}}$ will lead to a small area of phase space not

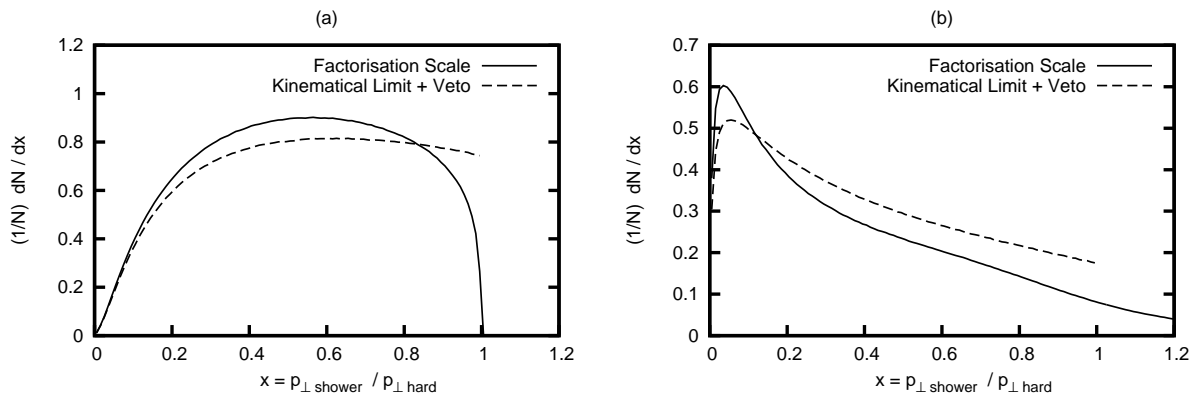


Figure 1: Ratio of the kinematic p_{\perp} of the first shower emission to the POWHEG emission, where the shower emission is (a) ISR or (b) FSR. In both cases, the results are shown when starting the shower at the factorisation scale and when starting the shower at the kinematical limit and vetoing above the POWHEG scale

being covered. An FSR emission, on the other hand, may have a small $p_{\perp\text{evol}}$ with respect to the emitting parton but still a $p_{\perp} > p_{\perp\text{POWHEG}}$ with respect to the beam axis.

A simple solution to both these problems is instead to begin the shower at the largest possible scale, and then veto any emissions with a kinematic $p_{\perp} > p_{\perp\text{POWHEG}}$. If we consider the first shower emission, the multiplicative nature of the no-emission probability ensures that the emission rates below $p_{\perp\text{POWHEG}}$ will be correct, i.e. unaffected by the vetoes above it. The picture is slightly less clear for subsequent emissions; having accepted one shower emission below $p_{\perp\text{POWHEG}}$, it is still possible for a later emission to be above it, since the first emission may well have had $p_{\perp\text{evol}} > p_{\perp\text{POWHEG}}$. The probability of such an occurrence is small, and effects formally of NNLO character, unenhanced by any large logarithms. They mainly show up for low- p_{\perp} first emissions, where their importance on the event as a whole is less, but still nonzero. Another NNLO issue is that recoil effects from one emission can shift the p_{\perp} of the previous ones, along with the hard process itself, either to lower or higher values.

The current POWHEG-hvq generator uses a second order running α_s expression, but with a Λ fixed at $n_f = 5$. Although slightly inconsistent, this only leads to changes beneath the Bottom and Charm scales. The Λ value is taken from a selected PDF set and is modified as in [46]. In the LHEF output file, all incoming and radiated partons are massless and the values of the couplings, α_s and α_{em} , are set to zero in all events.

To quantify how well the proposed interfacing works, we begin with top pair production ($m_t = 171\text{ GeV}$), where all results are generated at LHC energies (pp, $\sqrt{s} = 14\text{ TeV}$). In this case, the number of light flavours, which defines the content of the proton and the allowed radiation flavours, goes up to and includes the bottom quark ($n_l = 5$). To study the effect of the different shower starting scales, we examine the ratio of the first p_{\perp} in the shower to the p_{\perp} of the POWHEG emission (where the shower p_{\perp} value is taken directly after the emission). This is shown in Fig. 1, split into contributions from (a) ISR and (b) FSR. For ISR, we note that the ratios do not become larger than unity, but that when starting the shower at the factorisation scale, there is a region close to $p_{\perp\text{shower}}/p_{\perp\text{hard}} = 1$ where the phase space is not completely filled. This gap is filled when starting the shower at the

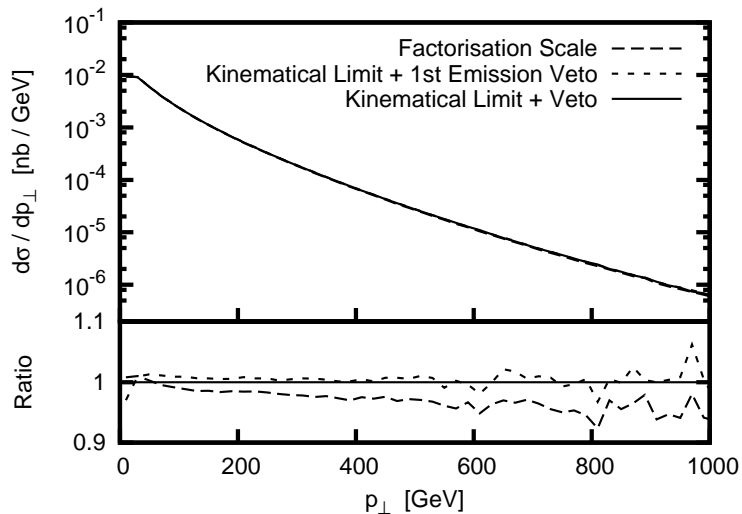


Figure 2: Final p_{\perp} of the top pair for three different cases (see text). In the upper plot, it is difficult to distinguish the different curves, but the results are clearer to see in the lower plot, which shows the ratio to the “Kinematical Limit + Veto” result

kinematical limit and vetoing emissions above $p_{\perp\text{POWHEG}} = p_{\perp\text{hard}}$; that is, the shower is running at “full steam” when it reaches the emission p_{\perp} threshold. That the curve happens to be so flat near the endpoint is a coincidence, related to a cancellation between the blowup of the naïve emission rate for smaller p_{\perp} , with a Sudakov damping in the same limit, helped along by many events having a small $p_{\perp\text{hard}}$ in the first place. It does not happen for FSR, where the large top mass reduces radiation to a lower level overall. For FSR, when starting at the factorisation scale, there is a tail beyond unity, as discussed previously, while this no longer happens with the veto scheme in place. Note that the FSR rate below $p_{\perp\text{POWHEG}}$ does come up, however, since emitted partons now are allowed at a larger separation from their mother parton so long as they are still at small p_{\perp} with respect to the beam. As already noted, the corrections are of higher order, but their inclusion is worthwhile for overall consistency.

We move on to study the effect of shower emissions beyond the first. As discussed previously, after a first allowed shower emission below $p_{\perp\text{POWHEG}}$, a subsequent emission may be generated above $p_{\perp\text{POWHEG}}$. To examine this effect, three different cases are considered: (1) the showers are started at the factorisation scale, (2) the showers are started at the kinematical limit and only the first shower emission is vetoed, and (3) the showers are started at the kinematical limit and all emissions above $p_{\perp\text{POWHEG}}$ are vetoed. As a measure of this effect, one may use the final p_{\perp} of the top pair, which is shown in Fig. 2 for the three different cases. At first glance, all three approaches appear to give similar results, but the ratio plot reveals that the difference between the factorisation scale and the veto schemes is around 10%, while there is little difference between the two different veto schemes. Instead, in Fig. 3, we study the smearing of the first shower emission p_{\perp} due to subsequent emissions. In this case only ISR is generated, so there is no ambiguity in picking out the first shower emission in the final event, and again, we show the ratio of the first shower emission to the POWHEG one, $p_{\perp\text{shower}}/p_{\perp\text{hard}}$, (a) immediately after the first emission and (b) after the shower evolution has finished. Fig. 3a shows the same features

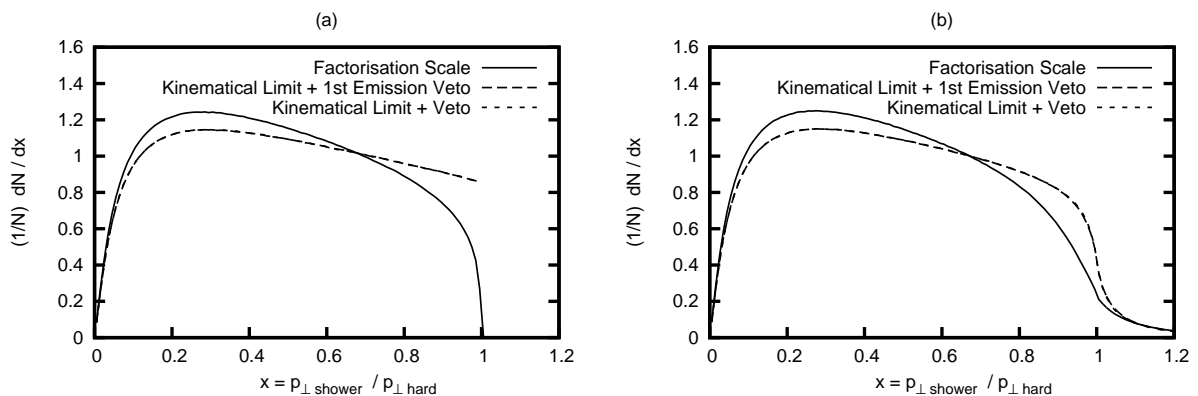


Figure 3: Ratio of the kinematic p_{\perp} of the first shower emission to the POWHEG emission for ISR only. (a) Shows the results immediately after the first emission, while (b) shows the results after the full shower evolution. In both cases, the “Kinematical Limit + 1st Emission Veto” and “Kinematical Limit + Veto” curves lie on top of each other

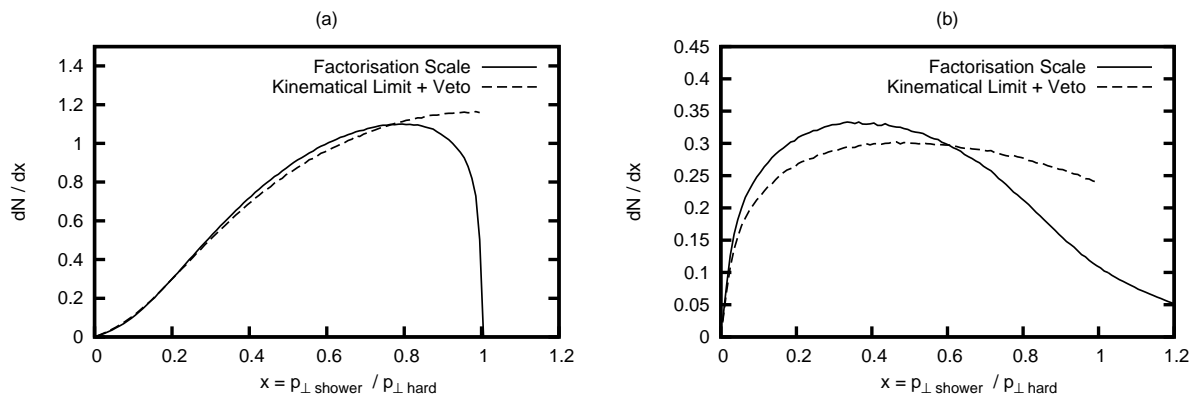


Figure 4: Ratio of the kinematic p_{\perp} of the first shower emission to the POWHEG emission for bottom pair production, where the shower emission is (a) ISR or (b) FSR

as Fig. 1, although with different normalisation due to the lack of FSR. In Fig. 3b, the first shower emission p_{\perp} is smeared by small amounts due to subsequent emissions, but again the difference between vetoing just the first emission and all emissions is negligible.

For bottom production ($m_b = 4.8 \text{ GeV}$), the number of light flavours is now set to four ($n_l = 4$), meaning that there will be no incoming or radiated b quarks. This is different from the default PYTHIA settings, where b quarks are both taken from the beam and allowed to be created in radiative emissions. The results are shown in Fig. 4, again split into contributions from (a) ISR and (b) FSR. Although the overall pattern of radiation is different, due to the smaller bottom mass, the details with respect to the interface with PYTHIA show the same features as Fig. 1.

2.5 Poor Man’s POWHEG

The POWHEG approach is a very powerful tool for NLO studies. Unfortunately, POWHEG-based implementations are not available for all relevant processes. Many older

NLO calculations are instead available in terms of phase space slicing results (see e.g. [47]). That is, the $(n + 1)$ -body cross section is based on the pure real-emission cross section $d\sigma_R$, without any Sudakov corrections. Unresolved emissions, virtual corrections and counterterms are lumped together with the Born cross section to provide the n -body cross section. The absence of a Sudakov factor means that soft/collinear divergences in $d\sigma_R$ are not dampened, thus requiring a reasonably large cutoff scale to stay away from the region where σ_n would turn negative. On the other hand, $d\sigma_R$ should be used where it can. Therefore σ_n and σ_{n+1} typically are chosen to be of the same order, while $\sigma_n \ll \sigma_{n+1}$ in POWHEG, where a much smaller cutoff can be used without any risk of inconsistencies.

In order to make these older NLO codes also useful to the experimental community they need to be interfaced to event generators, for further showering, MPI and hadronisation. It is then convenient to add Sudakovs to $d\sigma_R$, to bring it closer to the POWHEG approach. This would also smoothen out the transition between $(n + 1)$ - and n -body phase space, in the spirit of modern-day leading-order matching procedures, already mentioned in the introduction.

As a practical example we mention the WGAMMA_NLO package for $W\gamma$ production to NLO [48], with up to one additional quark or gluon in the final state, and allowing for anomalous $WW\gamma$ couplings. Here an implementation of the principles to be presented is already available [49].

Assume a lowest-order process that does not contain any coloured particles in the final state. The NLO processes thus include one quark or gluon extra, with divergences when its $p_\perp \rightarrow 0$. A scale $p_{\perp\min}$ is used to separate the $(n + 1)$ - and n -body phase space. An $(n + 1)$ -body state is characterised by the transverse momentum of the quark or gluon, p_\perp^{ME} , with $p_\perp^{\text{ME}} > p_{\perp\min}$. We now want to include the Sudakov to express that p_\perp^{ME} is the hardest jet of the event, i.e. that there are no jets at a larger scale.

In the CKKW-L [2] approach a fictitious trial shower is used to provide this. The better this shower attaches to the correct ME behaviour, the more accurate the Sudakov will be (thus, an advantage to having a more accurate default behaviour of the shower, as we strive for in this article). The point should not be overstressed, however; at large p_\perp the Sudakov suppression is negligible anyway, and at small p_\perp the universal behaviour should dominate.

The $(n + 1)$ topology has to be projected down to the core n process, without the extra parton, to provide the starting point for the trial shower. This requires a choice of which side the emission occurs on. Sometimes flavours allow only one possibility; for $u\bar{g} \rightarrow W^+\gamma d$ the hard subprocess must be $u\bar{d} \rightarrow W^+\gamma$ since the gluon cannot couple to the $W^+\gamma$ state. For $u\bar{d} \rightarrow W^+\gamma g$ the emission could be on either side, and the relative probability for a shower emission is related to the respective $1/p_{\perp\text{evol}}^2$. Since $z = m_n^2/m_{n+1}^2$ is independent of emission side, it follows that $p_{\perp\text{evol}}^2 \propto Q^2 \propto 1 \mp \cos\hat{\theta}$, where $\hat{\theta}$ is the angle of the gluon in the rest frame of the subcollision. The constancy of z also implies that the splitting kernel values are the same on both sides, and so give no net contribution. In another case, where the core process would be different depending on the side of emission, the weight of the respective splitting kernel would have to be taken into account, but still at the same z . In summary, the relative probability for a gluon emission from side 1 is $Q_1^{-2}/(Q_1^{-2} + Q_2^{-2}) = Q_2^2/(Q_1^2 + Q_2^2) = (1 + \cos\hat{\theta})/2$. The choice of side determines a new $x' = xz$ for it, while x on the other side is unchanged. A combined rotation and boost can bring the n particles to this new frame.

From there, the fictitious shower is allowed to generate the first/hardest emission at a

scale p_{\perp}^{PS} . The probability for $p_{\perp}^{\text{ME}} > p_{\perp}^{\text{PS}}$ is precisely the desired Sudakov, in the CKKW-L spirit. That is

- (i) if $p_{\perp}^{\text{ME}} > p_{\perp}^{\text{PS}}$ the original $n + 1$ -body topology is retained;
- (ii) else the projected n -body topology is selected.

In case (i), a normal shower can then be started up from p_{\perp}^{ME} and downwards; that is, further jets may be generated above the $p_{\perp\text{min}}$ scale, to give $(n + 2)$ -body topologies etc. The events in case (ii) are to be lumped together with the ones already originally classified as n -body, and allowed to shower from the $p_{\perp\text{min}}$ scale downwards.

In this approach, there is not guaranteed to be a smooth matching at $p_{\perp\text{min}}$, at least for the p_{\perp} scale of the hardest emission. The hope is that this step will be smeared out by subsequent showers and hadronization. Also note that, unlike POWHEG defined by eq.(13), but in line with MC@NLO, the high- p_{\perp} tail is defined by the LO $(n + 1)$ expression, without any $\bar{B}(v)/B(v)$ ‘‘K factor’’.

3 POWHEG without NLO

3.1 Correction to the PYTHIA shower

As will be shown in Section 3.3 for top pair production, an evolution with $Q_{\text{max}}^2 = s$ (power) overestimates the high- p_{\perp} tail while $Q_{\text{max}}^2 = m_{\perp t}^2$ (wimpy) underestimates it, cf. [21]. This is not so surprising, as follows. Let us recall that the matrix element for a QCD process such as $gg \rightarrow ggg$ scattering roughly behaves like [50, 51]

$$d\sigma \sim \frac{dp_{\perp 1}^2}{p_{\perp 1}^2} \frac{dp_{\perp 2}^2}{p_{\perp 2}^2} \frac{dp_{\perp 3}^2}{p_{\perp 3}^2} \delta^{(2)}(\mathbf{p}_{\perp 1} + \mathbf{p}_{\perp 2} + \mathbf{p}_{\perp 3}) , \quad (19)$$

where the $p_{\perp i}$ are the transverse momenta of the three outgoing gluons. In the limit $p_{\perp 3} \ll p_{\perp 1}$, where $|\mathbf{p}_{\perp 1}| \approx |\mathbf{p}_{\perp 2}|$, this reduces to $1/p_{\perp 1}^4 p_{\perp 3}^2$, i.e. the hard interaction behaves like $dp_{\perp 1}^2/p_{\perp 1}^4$ and the subsequent shower emission of an additional gluon like $dp_{\perp 3}^2/p_{\perp 3}^2$. Put another way, for a fixed $p_{\perp 1}$, you may distinguish a low- $p_{\perp 3}$ region with a fall-off like $dp_{\perp 3}^2/p_{\perp 3}^2$ and a high- $p_{\perp 3}$ region where the fall-off instead is like $dp_{\perp 3}^2/p_{\perp 1}^4$, with a smooth transition when $p_{\perp 3} \approx p_{\perp 1}$. In practice you would not want to simulate the process like this, of course, but reserve the hardest propagator to be described by matrix elements, which also would avoid doublecounting problems.

Obviously the picture is not equivalent for the $gg \rightarrow t\bar{t}g$, but let us apply a similar reasoning to the extra gluon in this case. For small $p_{\perp g}$ the picture of an ISR branching $g \rightarrow gg$ followed by a hard process $gg \rightarrow t\bar{t}$ ought to be a valid approximation, and so we expect a $dp_{\perp g}^2/p_{\perp g}^2$ falloff. At large $p_{\perp g}$, on the other hand, it would make more sense to think in terms of an ISR branching $g \rightarrow t\bar{t}$ followed by a hard process $gt \rightarrow gt$, and thus a shape more like $dp_{\perp g}^2/p_{\perp g}^4$. Now, we don’t simulate the latter kind of hard processes, e.g. because the top is so heavy that a top PDF is not a fruitful approximation, and so we would like to obtain these high- $p_{\perp g}$ configurations in the context of the simulation of the $gg \rightarrow t\bar{t}$ process. This leads to an ansatz of the form

$$\frac{d\mathcal{P}_{\text{ISR}}}{dp_{\perp}^2} \propto \frac{1}{p_{\perp}^2} \frac{k^2 M^2}{k^2 M^2 + p_{\perp}^2} , \quad (20)$$

where M^2 is a reasonable scale to associate with the hard $2 \rightarrow 2$ process (discussed further below) and k^2 is a fudge factor of order unity, parameterising at what scale the transition from a $1/p_\perp^2$ to a $1/p_\perp^4$ behaviour occurs.

How generic would such an ansatz be? First of all, for QCD processes involving light quarks and gluons, the showers should be cut off at (around) the scale of the hard process, or else one would doublecount, since in this case all possible $2 \rightarrow 2$ hard subprocesses one could construct out of a $2 \rightarrow 3$ process are already being simulated. Secondly, for $2 \rightarrow 1$ production of a colour singlet particle, the dp_\perp^2/p_\perp^2 ansatz works well up to the kinematical limit [28]. A reasonable assumption is that this generalises to processes where two or more colour singlet particles are produced in the core process, while a shape similar to eq. (20) should occur in processes that involve one or several coloured particles in the final state.

The argument for such a difference is one of colour coherence; with colour charge in both the initial and the final state one expects a destructive interference between ISR and FSR emissions that limits the radiation [52], while no such interference occurs with colours only in the initial state. One can then argue exactly what the scale M^2 appearing in eq. (20) should be; for pair production of coloured particles, the factorisation or renormalisation scale should be a reasonable choice, but when considering the production of a mixed coloured/non-coloured final state, by the coherence argument, it is primarily the coloured particles that should play a role in this scale. We note that the default choice for internal PYTHIA $2 \rightarrow 2$ processes is to set the renormalisation scale equal to the geometric mean of the squared transverse masses of the two outgoing particles, while the factorisation scale is set to the smaller of the squared transverse masses of the two outgoing particles. In the next sections we will check how well the ansatz of eq. (20) fares for a number of different processes.

3.2 MadEvent

3.2.1 Hard emission probability

In order to study processes for which an implementation of POWHEG is not yet available, we turn to the MadEvent matrix element generator. In eq. (13) we see that the POWHEG Sudakov, used to generate hard emissions, contains only Born and real terms; all NLO corrections are contained in a separate prefactor.

We therefore use MadEvent to generate events with an extra jet in the final state in order to extract the cross section for real jet emission, $d\sigma_R$. An approximate ‘‘POWHEG’’ style probability for emission is then formed by normalising to the overall lowest order (Born) cross section (generated by simulating the corresponding $2 \rightarrow 2$ process in MadEvent) and adding a Sudakov. This leads to a probability distribution described by

$$d\mathcal{P} = \frac{d\sigma_R}{\sigma_B} \exp\left(-\int_{p_\perp}^{p_\perp^{max}} \frac{d\sigma_R}{\sigma_B}\right), \quad (21)$$

where the p_\perp integral begins at the kinematical limit. As noted previously, the NLO prefactor can lead to kinematical differences in the resulting distribution. Under the assumption that such kinematical differences are small (to be addressed further in the case of top pair production), a qualitative comparison to the PYTHIA shower can be made. The effect of the Sudakov will be most visible in the low- p_\perp region, where the ME $d\sigma_R$ calculation is

divergent. Here the integrand will blow up and the Sudakov will make the distribution turn over, such that it will have a unit integral (up to cutoff effects). Given that the low- p_{\perp} region is where the parton shower should be most accurate, we expect the turnover in the Sudakov-modified MadEvent distribution to roughly correspond with the PYTHIA distribution.

All Standard Model masses in MadEvent are set equal to the default PYTHIA values. Variations in particle widths between MadEvent and PYTHIA are not expected to play a large role and are therefore neglected. The definition of both the proton and jets includes bottom quarks, as would be included in the PYTHIA default settings. In what follows, we will only consider heavy final states (top mass and higher), such that we can neglect FSR and only consider ISR. In this case, there are only soft and collinear singularities from the branching of incoming partons. A low- p_{\perp} cutoff of 2 GeV is introduced on real-jet emission, such that these divergences are avoided. No further cuts are applied.

3.2.2 Scale corrections

In generating events with MadEvent, there are again some different choices available relating to running α_s expressions and PDF's. The choice of PDF determines the running of α_s and the value $\alpha_s(M_z)$ and as before, the CTEQ6L PDF set was chosen. The ME calculations were generated with fixed renormalisation and factorisation scales, taken to be the geometric mean of the masses of the two heavy final-state particles. The results of these choices remain evident in the ME-derived jet distribution. In eqs. (2) and (3) it was shown that the PYTHIA shower algorithm instead picks scales related to the evolution variable when generating emissions. These differences can be non-negligible, especially in the low- p_{\perp} region, where α_s will become large and we expect a similar turnover to the PYTHIA distribution. To account for these effects, additional weights are applied to MadEvent distributions after generation, with the new scale taken to be the kinematic p_{\perp} of the jet.

A correction for α_s is simple to achieve, weighting events by the ratio of the new and old α_s values. With M defined as the fixed renormalisation/factorisation scale used in the ME calculations, we have a weight

$$\frac{\alpha_s(p_{\perp}^2)}{\alpha_s(M^2)}. \quad (22)$$

With ISR, correcting for differences in the factorisation scale is more difficult. The distribution for real emission from the ME calculation contains two PDF factors from the incoming partons, unlike ISR generated from the shower algorithm. This difference makes it unclear what correction should be applied, especially as in our ME calculation, we do not know which incoming parton has branched. Taking an exclusive point of view, there should have been no emissions between M^2 and p_{\perp}^2 on either side of the event, and so both should be reweighted, overall giving a correction factor

$$\frac{x_1 f_1(x_1, p_{\perp}^2)}{x_1 f_1(x_1, M^2)} \frac{x_2 f_2(x_2, p_{\perp}^2)}{x_2 f_2(x_2, M^2)}. \quad (23)$$

This may be a slight over-correction, however, and one would expect the results with just the α_s factor and those with this additional PDF factor to bracket the ‘‘correct’’ distribution. We examine this further in Sec. 3.3 for the case of top pair production.




Figure 5: Two possible graphs for \tilde{u}_L squark pair production with an additional jet which do not directly correspond to a $2 \rightarrow 2$ process and a PYTHIA shower splitting

3.2.3 Graph topologies

When moving to a final state containing an extra jet, MadEvent will correctly generate all possible topologies, but when comparing against PYTHIA, some of these graphs may not correspond to a shower history. Two examples are given in Fig. 5 for the case of up squark pair production in the MSSM.

In (a), the two t-channel squark propagators cannot directly be reproduced by a $2 \rightarrow 2$ process and a shower emission. In a traditional merging approach, these topologies can always be associated to some $2 \rightarrow 2$ hard process with the radiation associated to an incoming leg. As radiation off a massive leg is suppressed, topologies of this kind are not expected to dominate the overall behaviour of the real-emission matrix elements.

Instead, in (b), the squark and jet are created through the resonant decay of a gluino. This resonant mechanism leads to events with a different topology than we wish to consider; as PYTHIA works in the narrow width approximation, where the gluino production and decay are described as part of the separate $2 \rightarrow 2$ process $ug \rightarrow \tilde{u}\tilde{g}$, it would be doublecounting to include the same graph as a $2 \rightarrow 3$ process.

To remove double counting of this type, we use the method outlined in [53]; the idea here is to assign propagators to events statistically, based on the relative size of the relevant squared amplitudes. Such events then explicitly contain this propagator information in the LHEF output, meaning they can be easily vetoed when subsequently being processed. As a simple check that the results are reasonable, many processes were also generated with these resonant graphs manually removed, e.g. as in [21]. Although this procedure is not fully gauge invariant, as the widths of the resonances are small, we expect the results to be in line with the MadEvent veto method.

3.3 Standard Model processes

We begin by comparing the jet emission probability in top pair production of POWHEG, MadEvent and the first PYTHIA shower emission (ISR only). In all studies that follow, events are generated at LHC energies (pp, $\sqrt{s} = 14$ TeV) and, additionally, when the damping ansatz is used, the scale M^2 of eq. (20) is set to be the factorisation scale unless stated otherwise. For top production, we restrict both PYTHIA and MadEvent to QCD production only, as this is the dominant contribution to the cross section, and also allows a consistent comparison to POWHEG. As noted previously, the default PYTHIA ISR shower will begin evolution at the kinematical limit, given that a top cannot be produced in the

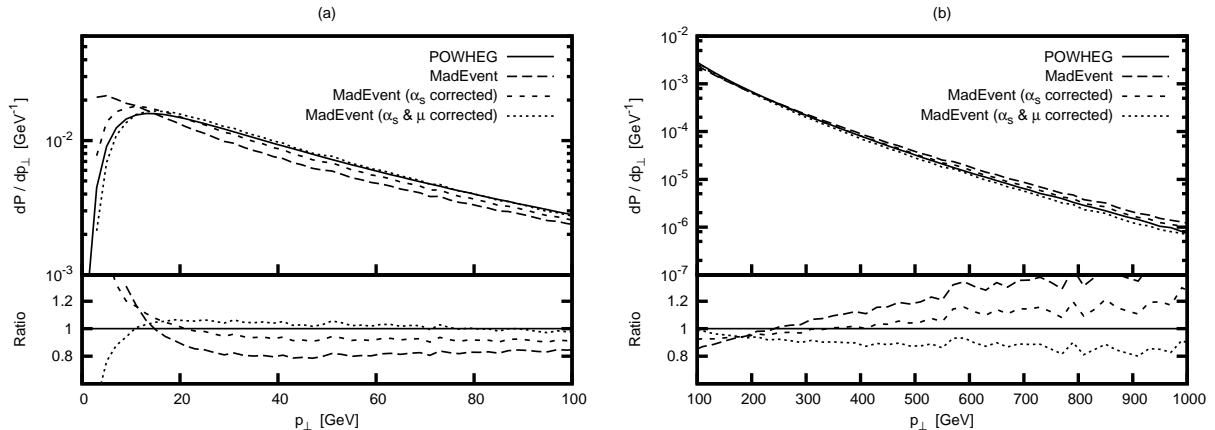


Figure 6: First emission probability as a function of p_{\perp} for top pair production in POWHEG and MadEvent. Ratio plots are normalised to the POWHEG result

shower. The POWHEG to MadEvent comparison allows us to check the validity of the MadEvent approximation method outlined in Sec. 3.2.1. The comparison to PYTHIA then lets us examine the effectiveness of the ansatz described in Sec. 3.1.

The probability distributions for jet emission in top pair production are expected to be roughly in agreement over the entire p_{\perp} range for POWHEG and MadEvent, as long as the kinematical differences that come from the full NLO prefactor of eq. (13) are small. This is shown in Fig. 6, separately for low and high p_{\perp} regions, for three different sets of the MadEvent results. All three have the Sudakov correction applied, but one set additionally has the α_s scale correction (“ α_s corrected”), and the other both the α_s and factorisation scale corrections (“ α_s & μ corrected”). These will all meet when the p_{\perp} scale matches that of the fixed factorisation/renormalisation scale used in the MadEvent generation, but begin to diverge as the p_{\perp} rises and falls away from this value. In the low- p_{\perp} region, the Sudakovs make the distributions turn over, while in the high p_{\perp} tail, there is good agreement between the distributions of POWHEG and the approximation from MadEvent; in the ratio plot the POWHEG curve sits between the two different corrected MadEvent curves, but overall closer to the α_s & μ -corrected one.

We now move on to compare POWHEG against the PYTHIA shower and the ansatz of Sec. 3.1. The expectation is that the default PYTHIA power shower will not fall off quickly enough and therefore overestimate the emission probability in the high p_{\perp} tail. In Fig. 7 we show, again separately for low and high p_{\perp} regions, the POWHEG results against four variations of the PYTHIA shower (wimpy, power, damped with $k = 1$ and damped with $k = 2$). Again, there is good agreement in the low- p_{\perp} region. In the high- p_{\perp} tail, the wimpy shower sits clearly below the POWHEG result, while the PYTHIA power shower does indeed overestimate the jet emission probability. The damping procedure, particularly with $k = 2$, then brings the PYTHIA distribution into closer agreement with POWHEG.

It is interesting to study the kinematical differences that come from the full NLO prefactor in POWHEG by considering the top pair rapidity both with and without such a prefactor. When full NLO corrections are included, we not only expect an overall change in cross section, but also that there may be kinematic differences in the distributions. In Fig. 8a, the ratio of two POWHEG pair rapidity distributions is shown. The first is from the default POWHEG-hvq generator, with a full NLO prefactor, while the second is gen-

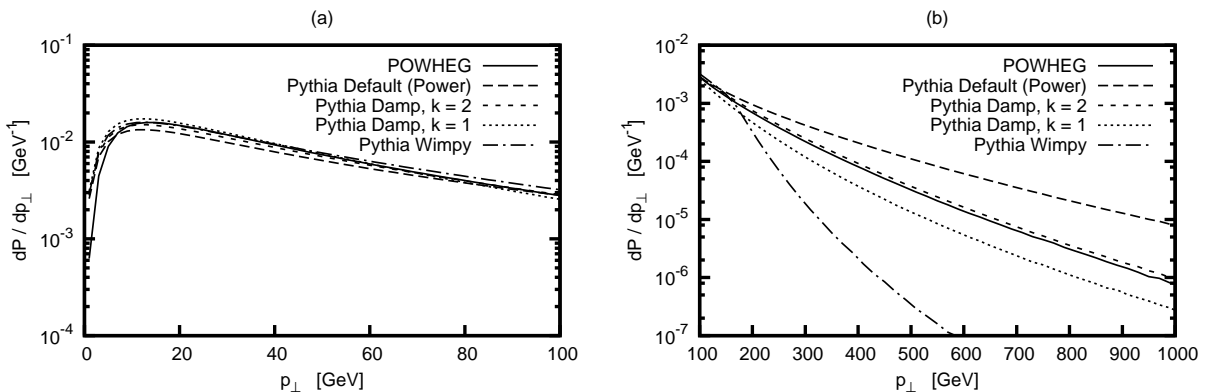


Figure 7: First emission probability as a function of p_{\perp} for top pair production in POWHEG and PYTHIA, split into (a) the low- p_{\perp} region where the Sudakov makes the distributions turn over and (b) the high- p_{\perp} tail

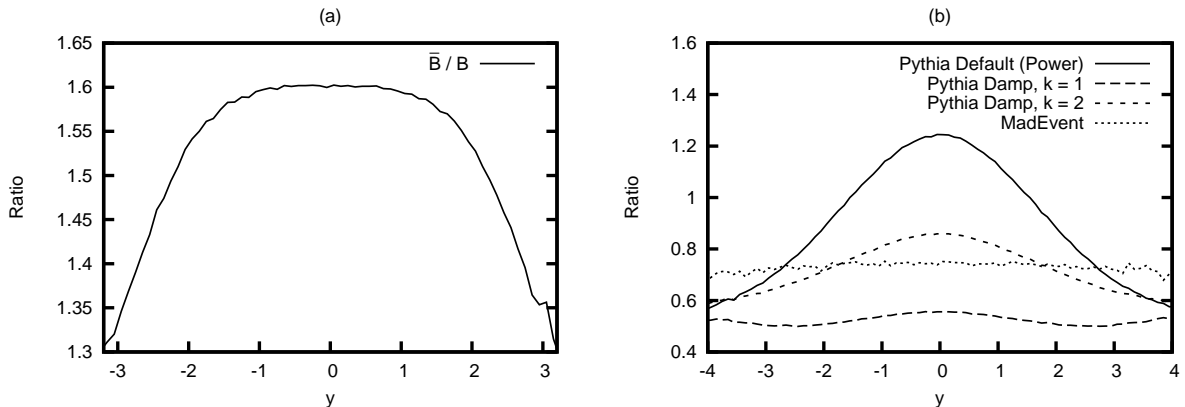


Figure 8: (a) Ratio of top pair rapidity between POWHEG with the full NLO prefactor, \bar{B} , of eq. (13) and B only. Note the suppressed zero on the y axis. (b) Jet rapidity distributions of PYTHIA and MadEvent for $p_{\perp jet} > 100\text{GeV}$, normalised to the POWHEG result

erated by replacing the \bar{B} term of eq. (13) with just the Born contribution B . The figure shows both features; an overall shift in cross section (a k -factor of around 1.5) and a small shift in pair rapidity to more central regions.

We finally compare the jet rapidity distributions from POWHEG, MadEvent and PYTHIA. In PYTHIA, the rapidity of the extra jet is taken immediately after the first shower emission (if present). For MadEvent, however, we can not apply the Sudakov correction in this case, and the direct output of the generator is taken. We restrict ourselves to $p_{\perp jet} > 100\text{GeV}$, such that the effect of the Sudakov and other scale corrections will be small. Here it is harder to know what to expect. For the damped PYTHIA shower, it is clear that by only reducing the high- p_{\perp} tail relative to the power shower, events will be taken out of the central region. These distributions are shown in Fig. 8b, normalised to the POWHEG result. MadEvent and POWHEG are in good agreement, although with an overall shift in cross section. The default PYTHIA shower is strongly peaked in the central region, but comes into better agreement when damped.

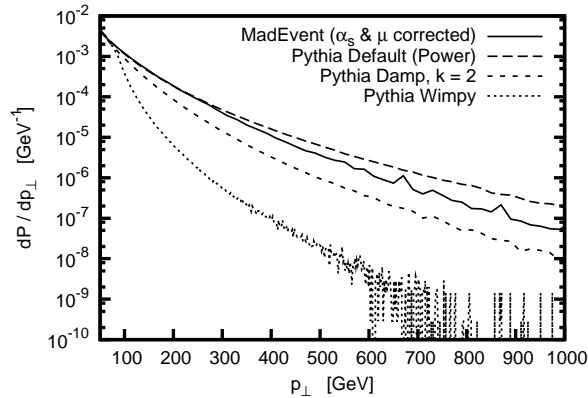


Figure 9: First emission probability as a function of p_{\perp} for Z pair production. Results show PYTHIA compared to the approximate MadEvent (α_s & μ corrected) prescription

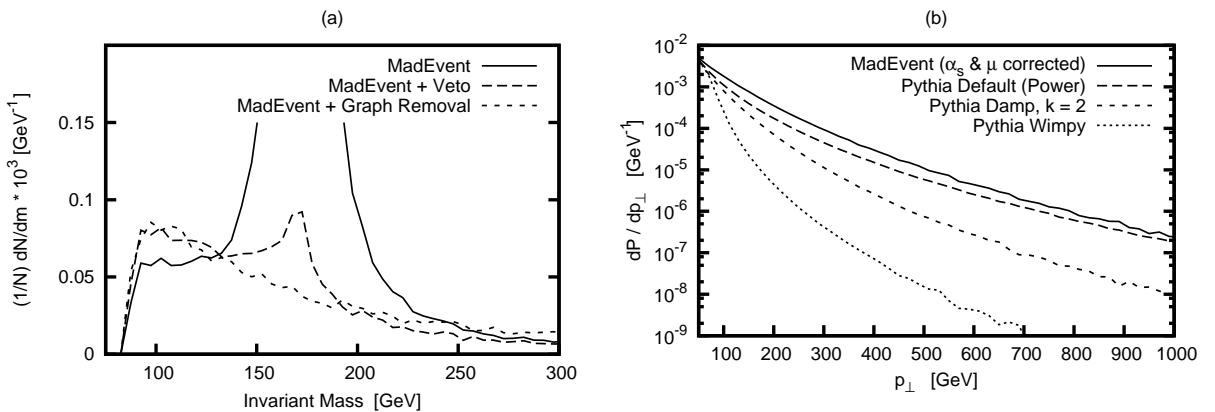


Figure 10: W pair production. (a) Shows the invariant mass of the jet and W boson when the jet is a bottom quark for three different MadEvent runs (see text). The large resonant peak in the “MadEvent” sample has been cut off above 0.15. (b) Shows the first emission probability as a function of p_{\perp}

We now study processes where there is not (yet) a “correct” POWHEG result to compare with, but instead rely on the MadEvent approximation. If the ansatz of Sec. 3.1 is correct, the damping of the PYTHIA shower is not expected to help improve the p_{\perp} distribution in Z and W pair production in a significant way. Here, there are no coloured final-state particles to use as a guide in setting the M^2 scale of eq. (20), so we continue to use the factorisation scale when generating the damped results. Again, in this case, the PYTHIA default is to start ISR shower at the kinematical limit. Fig. 9 shows the results for Z pair production for the default, damped ($k = 2$) and wimpy shower against MadEvent (α_s & μ corrected). The results show that the default shower, although giving a slightly too hard p_{\perp} tail, still does a reasonable job of reproducing the MadEvent curve without additional damping.

The case of W pair production is useful as a check on the MadEvent resonance veto method (see Sec. 3.2.3). When bottom quarks are allowed in both the incoming beams and in the definition of the extra jet, there is now a resonant contribution coming from $t \rightarrow bW$. In Fig. 10a, the invariant mass of the jet and the matching W are shown when the jet is a bottom quark. A clear peak is visible when those events with a resonance are

Particle	Mass (GeV)
\tilde{u}_L	561.1
\tilde{g}	607.7
$\tilde{\chi}_1^0$	96.7
$\tilde{\chi}_1^\pm$	181.7

Table 1: Subset of MSSM SPS1a masses used in process generation

not vetoed (“MadEvent”), while this peak is almost completely removed when these events are vetoed (“MadEvent + Veto”); there is only a small remaining trace of it from e.g. interference terms. As an additional check, also plotted is the result when resonant graphs are explicitly removed from the MadEvent generation (“MadEvent + Graph Removal”). Fig. 10b, shows the results for the first emission probability in the p_\perp tail for the MadEvent vetoed sample. As a final check, the radiation pattern was compared to a sample not including bottom quarks, and found to contain no large differences.

3.4 MSSM studies

To further study the production of heavy final states, MSSM processes were chosen split into three groups: coloured final states, non-coloured final states, and coloured/non-coloured final states. All events were generated with the SPS1a [54] set of parameters (the relevant final-state masses used are shown in Tab. 1). Here, also the lowest order $2 \rightarrow 2$ processes were generated using MadEvent, and fed into PYTHIA for showering, again with the factorisation and renormalisation scales fixed at the geometric mean of the masses of the two heavy final-state particles. Note that this is slightly different from the default PYTHIA internal $2 \rightarrow 2$ process scale choices and that this choice enters into the M^2 scale selection of the damping ansatz, eq. 20.

3.4.1 Coloured final states

For the fully coloured final states, all events were generated with QCD only; as with top pair production, this is where the dominant contribution to the cross section lies. As before, the expectation is that the coloured final states will benefit from a damping of the high- p_\perp tail. Figure 11 shows the tail of the p_\perp distributions, all compared to the MadEvent α_s & μ corrected data, for (a) $\tilde{u}_L \bar{\tilde{u}}_L$, (b) $\tilde{u}_L \tilde{g}$ and (c) $\tilde{g} \tilde{g}$. Similarly to top pair production, the power shower overestimates the high- p_\perp tail, while the damping ansatz brings the curves into better agreement. For $\tilde{u}_L \bar{\tilde{u}}_L$ production, $k = 2$ leads to the best agreement, while for $\tilde{u}_L \tilde{g}$ and $\tilde{g} \tilde{g}$ production, the MadEvent curve lies between the damped $k = 1$ and $k = 2$ curves.

3.4.2 Non-coloured final states

We move on to study the non-coloured final states, $\tilde{\chi}_1^0 \tilde{\chi}_1^0$ and $\tilde{\chi}_1^+ \tilde{\chi}_1^-$, where we expect the damping ansatz not to improve the p_\perp tail of the parton shower. Here, for both processes, at lowest order, there are large resonant H^0/A^0 contributions. These are relatively long-lived intermediate states which we expect to follow the rules for $2 \rightarrow 1$ singlet production, where

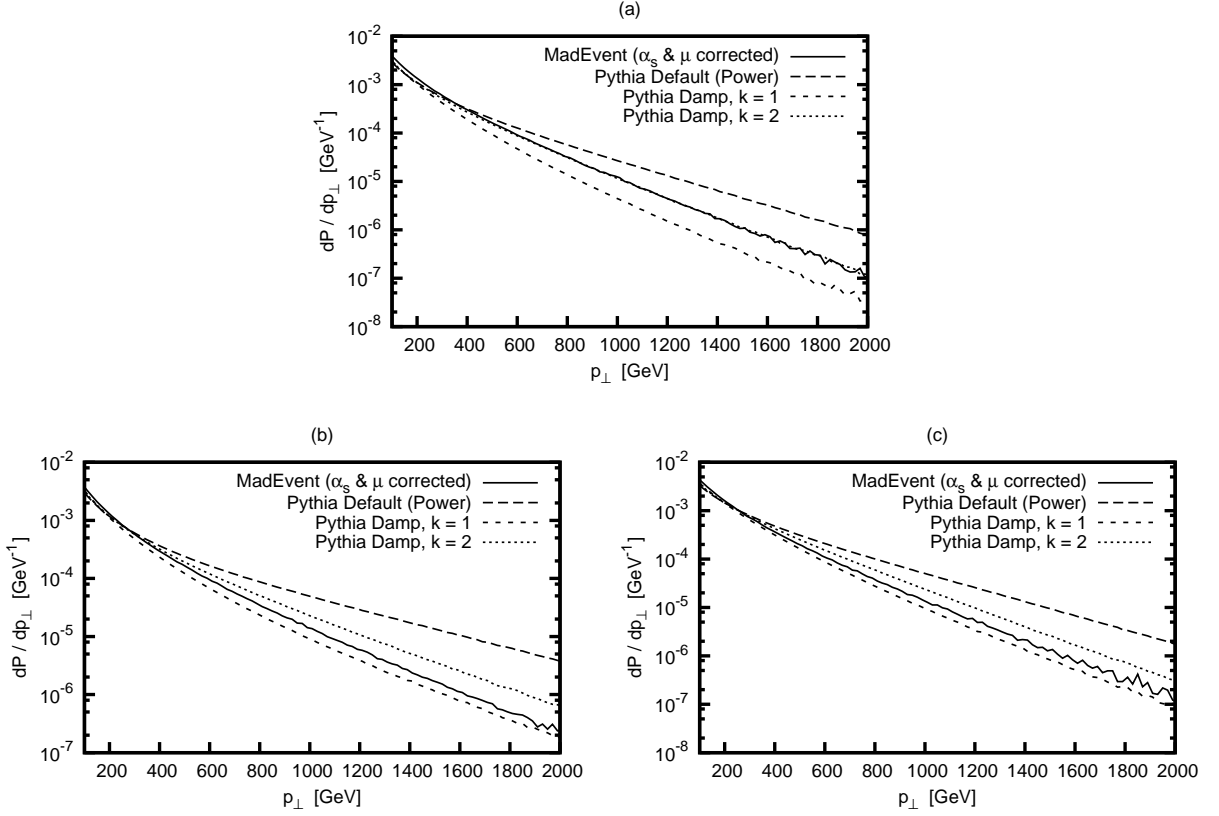


Figure 11: First emission probability as a function of p_{\perp} for (a) $\tilde{u}_L \bar{\tilde{u}}_L$, (b) $\tilde{u}_L \tilde{g}$ and (c) $\tilde{g} \tilde{g}$ production. In (a) the $k = 2$ and MadEvent curves lie on top of each other

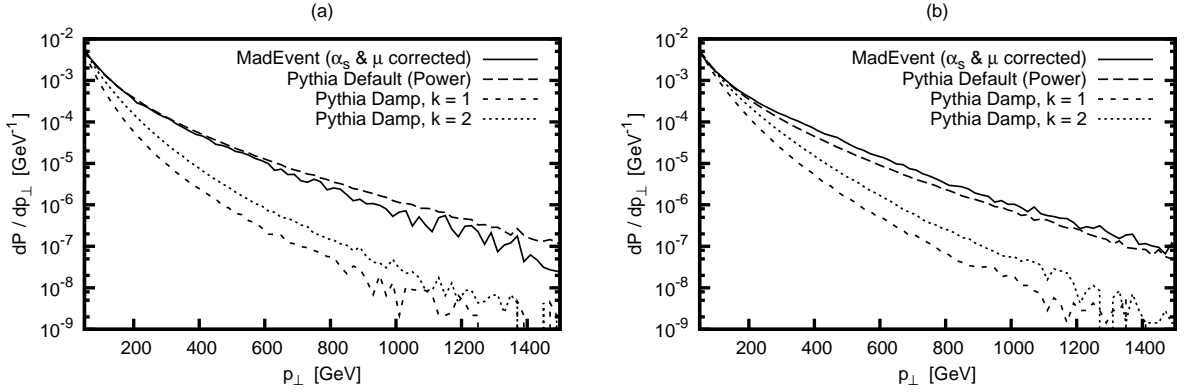


Figure 12: First emission probability as a function of p_{\perp} for (a) $\tilde{\chi}_1^0 \tilde{\chi}_1^0$ and (b) $\tilde{\chi}_1^+ \tilde{\chi}_1^-$ production

the shower already does a good job in covering the entire phase space. For both processes then, the MadEvent veto scheme is used to remove events of this type.

The results are shown in Fig. 12 for (a) $\tilde{\chi}_1^0 \tilde{\chi}_1^0$ and (b) $\tilde{\chi}_1^+ \tilde{\chi}_1^-$ production. The results follow the pattern for Z/W^{\pm} pair production; in both cases, the default power shower does a reasonable job in the high- p_{\perp} tail.

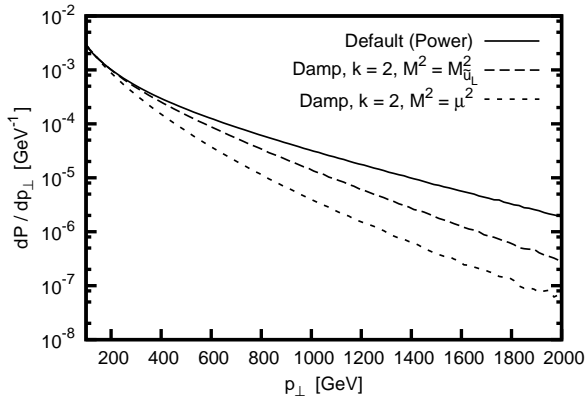


Figure 13: PYTHIA first emission probability for the damped shower ($k = 2$) in $\tilde{u}_L \tilde{\chi}_1^0$ production, where M^2 is set either to the factorisation scale or to the mass of the heavy outgoing squark. The default power shower curve is shown for comparison

3.4.3 Coloured/non-coloured final states

Finally, we study mixed coloured/non-coloured final states. Here, we would expect some measure of damping to improve the shower description due to the presence of colour in the final state. As discussed in Sec. 3.1, one issue here is the meaning of the M^2 scale of eq. (20). With the mixed final state, the difference between the factorisation scale and e.g. just the mass of the final-state coloured object is large enough to give noticeable differences in the damped shower tail. As an example, Fig. 13 shows the difference in the damped shower ($k = 2$) in $\tilde{u}_L \tilde{\chi}_1^0$ production, when M^2 is set to the factorisation scale and when it is instead set to the mass of the outgoing squark. In the remaining results of this section, we retain the choice of setting M^2 to the mass of the outgoing coloured state. Also note that all of the processes studied in this section have strong resonant contributions, where a squark decays into a neutralino/chargino and a jet, making it difficult to generate large statistics when using the MadEvent veto scheme.

Results are shown in Fig. 14 for (a) $\tilde{u}_L \tilde{\chi}_1^0$ and (b) $\tilde{u}_L \tilde{\chi}_1^-$. While the curve for $\tilde{u}_L \tilde{\chi}_1^-$ does show the expected behaviour, with the damping improving the shower description, for $\tilde{u}_L \tilde{\chi}_1^0$ we do not obtain the expected results. Instead, Fig. 15a shows the results for $\tilde{u}_L \tilde{\chi}_1^0$ when resonant graphs are manually set to zero in the MadEvent generation code. Here, beyond $p_\perp \sim 1$ TeV, the tail does appear to fall off with a damped behaviour. Comparing the real-emission Feynman diagrams for $\tilde{u}_L \tilde{\chi}_1^0$ against those of $\tilde{u}_L \tilde{\chi}_1^-$, a key difference is the appearance of right-handed intermediate squarks. In Fig. 15b, we show the results when the right-handed squark masses ($\tilde{u}_R, \tilde{d}_R, \tilde{s}_R$ and \tilde{c}_R) are set high. In this case, we do recover the behaviour of the $\tilde{u}_L \tilde{\chi}_1^-$ result.

Finally, in Fig. 16, we study the processes (a) $\tilde{g} \tilde{\chi}_1^0$ and (b) $\tilde{g} \tilde{\chi}_1^-$. Again the expected behaviour is not apparent; both sit closest to the default power shower curve.

In summary, for the mixed processes studied in this section, our ansatz is not an obvious improvement relative to the power shower. There is a nontrivial dependence of the emission pattern on the SUSY parameter choices that we do not understand, and do not go on to study further at this time.

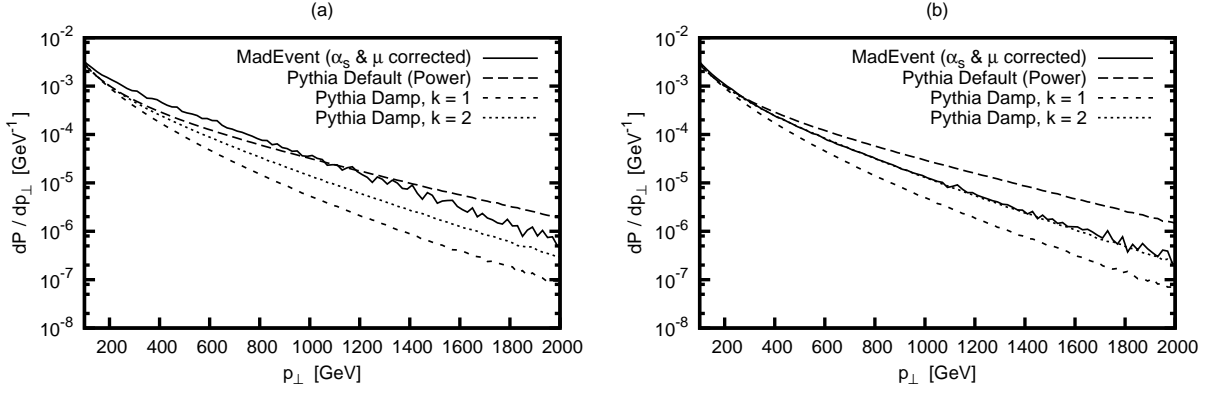


Figure 14: First emission probability as a function of p_{\perp} for (a) $\tilde{u}_L \tilde{\chi}_1^0$ and (b) $\tilde{u}_L \tilde{\chi}_1^-$

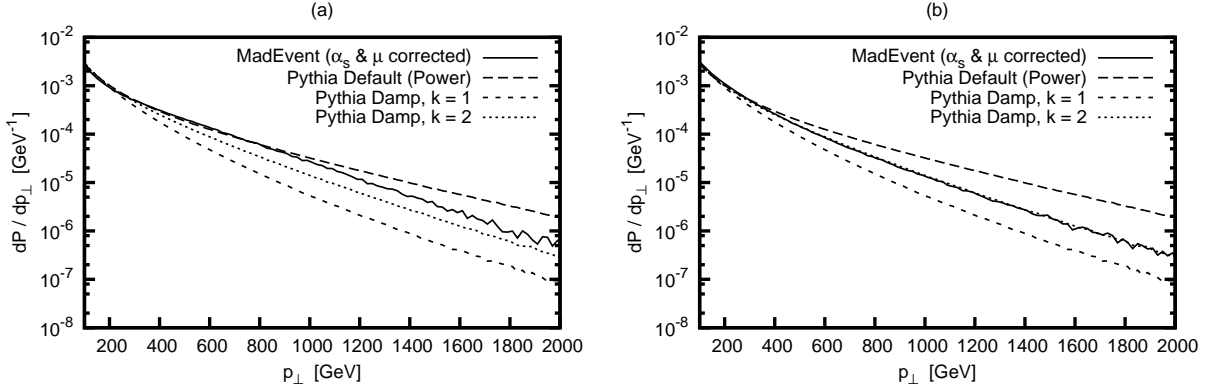


Figure 15: First emission probability as a function of p_{\perp} for $\tilde{u}_L \tilde{\chi}_1^0$, where in (a) resonant graphs have been removed and (b) the right handed squark masses have been set high

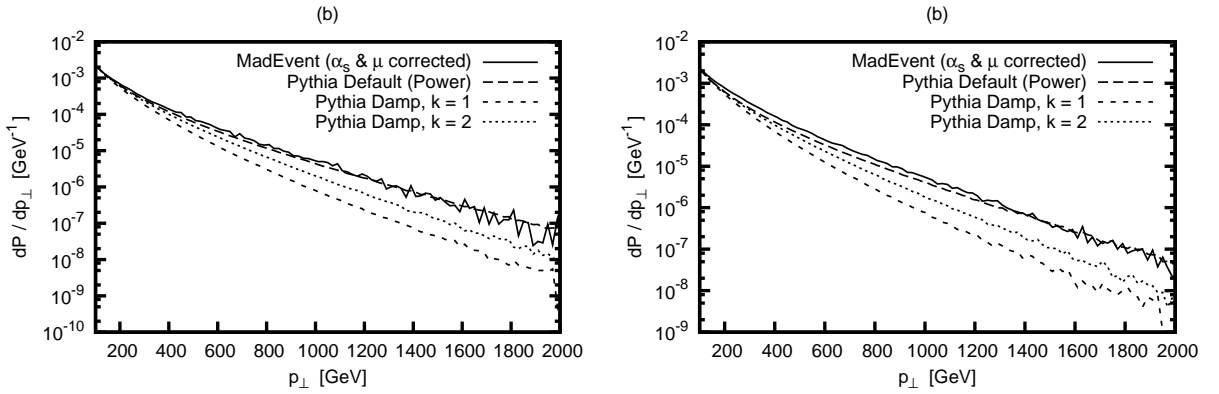


Figure 16: First emission probability as a function of p_{\perp} for (a) $\tilde{g} \tilde{\chi}_1^0$ and (b) $\tilde{g} \tilde{\chi}_1^-$

4 Summary and Outlook

One direct outcome of this study is an improved matching between the existing POWHEG program family and the PYTHIA p_{\perp} -ordered showers. The main point is that the PYTHIA p_{\perp} definition does not quite agree with the normal kinematical one, although the two do agree closely over most of the phase space populated by showers. In light of this, the numerical effect of these improvements are fairly modest, but, in view of the important role that we foresee for POWHEG-based studies at the LHC, there is every reason to have the interface to PYTHIA well understood.

The POWHEG approach, in complete or in simplified form, also allows us to test the default behaviour of the p_{\perp} -ordered showers. In particular we address the issue of “power” vs. “wimpy” showers, i.e. what starting scale to use for the downwards evolution, specifically for pair production of heavy particles. Here we show that many processes obey an intermediate behaviour, where the characteristic dp_{\perp}^2/p_{\perp}^2 shower fall-off is replaced by a steeper dp_{\perp}^2/p_{\perp}^4 fall-off for large p_{\perp} values. That is, emissions are allowed up to the kinematical limit, but at a dampened rate. The damping can be approximated by a factor

$$\mathcal{P}_{\text{dampen}} = \frac{k^2 M^2}{k^2 M^2 + p_{\perp}^2} \quad , \quad (24)$$

for the emission of a parton of transverse momentum p_{\perp} , where M^2 is a characteristic scale of the hard process and k is a fudge factor.

More specifically, we have seen that for coloured pair production, with M^2 set equal to the factorisation scale (noting the slightly different scale choice used in the PYTHIA internal processes and the external MSSM ones) and $k = 2$, we get a reasonable shower behaviour for the different final states studied here. Instead, when the particles in the final state of the hard process are colour singlets only, we have seen that there is no need to impose a damping, consistent with previous results for single W/Z production. The differences between these two cases can be understood based on the destructive interference between ISR and FSR that is expected for coloured particles but not for the uncoloured ones. In no case is the wimpy shower a good choice.

For the mixed coloured/non-coloured final states of Sec. 3.4.3, we have argued that M^2 should instead be related to the coloured partons only, due to this coherence argument. Unfortunately, for these processes, the results are not as expected. We do not currently understand the structure of the real-emission matrix elements that gives rise to the behaviour seen here and, specifically, whether it is an artifact of the complicated SUSY structure or an inherent property of the radiation pattern.

Not studied here is the case of the production of light coloured particles, i.e. of normal QCD jets. There, the hard process and the showers produce the same kind of partons, and doublecounting becomes a main concern. We intend to return to this class of event.

With the advances in computational tools, allowing automatised (Born-level) higher-order calculations, one may question the need for more accurate showers. We believe there are two main points in favour of improving showers, especially when this can be done with only a modest effort. One is that new physics scenarios are continuously being proposed, where higher-order calculations may be overkill for first studies, but nevertheless a realistic population over all possible event topologies is useful, even if off by a factor of two or so in some tails. The other is that trial showers as a means of obtaining Sudakov factors

is a crucial ingredient of CKKW-L matching schemes, so that the quality of matching is improved if the quality of the shower is also improved.

Acknowledgments

This work was supported in by the Marie Curie Early Stage Training program “HEP-EST” (contract number MEST-CT-2005-019626) and in part by the Marie Curie research training network “MCnet” (contract number MRTN-CT-2006-035606).

References

- [1] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, *JHEP* **11** (2001) 063, [arXiv:hep-ph/0109231](#).
- [2] L. Lönnblad, *JHEP* **05** (2002) 046, [arXiv:hep-ph/0112284](#).
- [3] S. Hoeche, F. Krauss, S. Schumann, and F. Siegert, *JHEP* **05** (2009) 053, [arXiv:0903.1219 \[hep-ph\]](#).
- [4] K. Hamilton, P. Richardson, and J. Tully, *JHEP* **11** (2009) 038, [arXiv:0905.3072 \[hep-ph\]](#).
- [5] M. Mangano. <http://www-cpd.fnal.gov/personal/mrenna/tuning/nov2002/mlm.pdf>.
- [6] S. Mrenna and P. Richardson, *JHEP* **05** (2004) 040, [arXiv:hep-ph/0312274](#).
- [7] J. Alwall *et al.*, *Eur. Phys. J.* **C53** (2008) 473–500, [arXiv:0706.2569 \[hep-ph\]](#).
- [8] N. Lavesson and L. Lönnblad, *JHEP* **04** (2008) 085, [arXiv:0712.2966 \[hep-ph\]](#).
- [9] S. Frixione and B. R. Webber, *JHEP* **06** (2002) 029, [arXiv:hep-ph/0204244](#).
- [10] S. Frixione, P. Nason, and B. R. Webber, *JHEP* **08** (2003) 007, [arXiv:hep-ph/0305252](#).
- [11] O. Latunde-Dada, *JHEP* **11** (2007) 040, [arXiv:0708.4390 \[hep-ph\]](#).
- [12] P. Nason, *JHEP* **11** (2004) 040, [arXiv:hep-ph/0409146](#).
- [13] S. Frixione, P. Nason, and C. Oleari, *JHEP* **11** (2007) 070, [arXiv:0709.2092 \[hep-ph\]](#).
- [14] G. Marchesini and B. R. Webber, *Nucl. Phys.* **B238** (1984) 1.
- [15] G. Marchesini and B. R. Webber, *Nucl. Phys.* **B310** (1988) 461.
- [16] T. Sjöstrand and P. Z. Skands, *Eur. Phys. J.* **C39** (2005) 129–154, [arXiv:hep-ph/0408302](#).
- [17] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *JHEP* **05** (2006) 026, [arXiv:hep-ph/0603175](#).

- [18] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178** (2008) 852–867, arXiv:0710.3820 [hep-ph].
- [19] S. Frixione, P. Nason, and G. Ridolfi, arXiv:0707.3081 [hep-ph].
- [20] J. Alwall *et al.*, *JHEP* **09** (2007) 028, arXiv:0706.2334 [hep-ph].
- [21] T. Plehn, D. Rainwater, and P. Z. Skands, *Phys. Lett.* **B645** (2007) 217–221, arXiv:hep-ph/0510144.
- [22] W. T. Giele, D. A. Kosower, and P. Z. Skands, *Phys. Rev.* **D78** (2008) 014026, arXiv:0707.3652 [hep-ph].
- [23] C. W. Bauer, F. J. Tackmann, and J. Thaler, *JHEP* **12** (2008) 010, arXiv:0801.4026 [hep-ph].
- [24] N. Lavesson and L. Lönnblad, *JHEP* **12** (2008) 070, arXiv:0811.2912 [hep-ph].
- [25] T. Sjöstrand, *Phys. Lett.* **B157** (1985) 321.
- [26] M. Bengtsson and T. Sjöstrand, *Phys. Lett.* **B185** (1987) 435.
- [27] E. Norrbin and T. Sjöstrand, *Nucl. Phys.* **B603** (2001) 297–342, arXiv:hep-ph/0010012.
- [28] G. Miu and T. Sjöstrand, *Phys. Lett.* **B449** (1999) 313–320, arXiv:hep-ph/9812455.
- [29] M. H. Seymour, *Comp. Phys. Commun.* **90** (1995) 95–101, arXiv:hep-ph/9410414.
- [30] G. Corcella and M. H. Seymour, *Phys. Lett.* **B442** (1998) 417–426, arXiv:hep-ph/9809451.
- [31] G. Corcella and M. H. Seymour, *Nucl. Phys.* **B565** (2000) 227–244, arXiv:hep-ph/9908388.
- [32] P. Nason and G. Ridolfi, *JHEP* **08** (2006) 077, arXiv:hep-ph/0606275.
- [33] S. Alioli, P. Nason, C. Oleari, and E. Re, *JHEP* **07** (2008) 060, arXiv:0805.4802 [hep-ph].
- [34] S. Frixione, P. Nason, and G. Ridolfi, *JHEP* **09** (2007) 126, arXiv:0707.3088 [hep-ph].
- [35] S. Alioli, P. Nason, C. Oleari, and E. Re, *JHEP* **09** (2009) 111, arXiv:0907.4076 [hep-ph].
- [36] S. Alioli, P. Nason, C. Oleari, and E. Re, *JHEP* **04** (2009) 002, arXiv:0812.0578 [hep-ph].
- [37] P. Nason and C. Oleari, arXiv:0911.5299 [hep-ph].
- [38] S. Alioli, P. Nason, C. Oleari, and E. Re, arXiv:1002.2581 [hep-ph].

- [39] T. Sjöstrand and M. van Zijl, *Phys. Rev.* **D36** (1987) 2019.
- [40] R. Corke and T. Sjöstrand, *JHEP* **01** (2010) 035, arXiv:0911.1909 [hep-ph].
- [41] J. Alwall *et al.*, *Comput. Phys. Commun.* **176** (2007) 300–304, arXiv:hep-ph/0609017.
- [42] M. R. Whalley, D. Bourilkov, and R. C. Group, arXiv:hep-ph/0508110.
- [43] J. Pumplin *et al.*, *JHEP* **07** (2002) 012, arXiv:hep-ph/0201195.
- [44] E. Boos *et al.*, arXiv:hep-ph/0109068.
- [45] J. M. Butterworth *et al.*, arXiv:1003.1643 [hep-ph].
- [46] S. Catani, B. R. Webber, and G. Marchesini, *Nucl. Phys.* **B349** (1991) 635–654.
- [47] W. T. Giele, E. W. N. Glover, and D. A. Kosower, *Nucl. Phys.* **B403** (1993) 633–670, arXiv:hep-ph/9302225.
- [48] U. Baur, T. Han, and J. Ohnemus, *Phys. Rev.* **D48** (1993) 5140–5161, arXiv:hep-ph/9305314.
- [49] D. Majumder, K. Mazumdar, and T. Sjostrand, arXiv:1002.4296 [Unknown].
- [50] F. A. Berends, R. Kleiss, P. De Causmaecker, R. Gastmans, and T. T. Wu, *Phys. Lett.* **B103** (1981) 124.
- [51] B. Andersson, G. Gustafson, H. Kharraziha, and J. Samuelsson, *Z. Phys.* **C71** (1996) 613–624.
- [52] R. K. Ellis, G. Marchesini, and B. R. Webber, *Nucl. Phys.* **B286** (1987) 643.
- [53] J. Alwall, S. de Visscher, and F. Maltoni, *JHEP* **02** (2009) 017, arXiv:0810.5350 [hep-ph].
- [54] B. C. Allanach *et al.*, *Eur. Phys. J.* **C25** (2002) 113–123, arXiv:hep-ph/0202233.