Multiple Interactions in PYTHIA 8

Richard Corke[1] [†‡]
[1]Dept. of Theoretical Physics, Solvegatan 14A, S-223 62 Lund, Sweden

Abstract

Modelling multiple partonic interactions in hadronic events is vital for understanding minimum-bias physics, as well as the underlying event of hard processes. A brief overview of the current PYTHIA 8 multiple interactions (MI) model is given, before looking at two additional effects which can be included in the MI framework. With rescattering, a previously scattered parton is allowed to take part in another subsequent scattering, while with enhanced screening, the effects of varying initial-state fluctuations are modelled.

1 Introduction

The run-up to the start of the LHC has led to a greatly increased interest in the physics of multiple parton interactions in hadronic collisions. Existing models are used to try to get an insight into what can be expected at new experiments, extrapolating fits to Tevatron and other data to LHC energies [1]. Such extrapolations, however, come with a high level of uncertainty; within many models are parameters which scale with an uncertain energy dependence. There is, therefore, also the exciting prospect of new data, with which to further constrain and improve models.

In terms of theoretical understanding, MI is one of the least well understood areas. While current models, after tuning, are able to describe many distributions very well, there are still many others which are not fully described. This is a clear sign that new physical effects need to be modelled and it is therefore not enough to "sit still" while waiting for new data. It is with this in mind that we look at two new ideas in the context of MI and their potential effects.

With rescattering, an already scattered parton is able to undergo another subsequent scattering. Although, in general, such rescatterings may be relatively soft, even when compared to normal $2 \rightarrow 2$ MI scatterings, they can lead to non-trivial colour flows which change the structure of events. Another idea is to consider partonic fluctuations in the incoming hadrons before collision. In such a picture, it is possible to get varying amounts of colour screening on an event-by-event basis. The question then is, what effects such new ideas would have on multiple interactions and how can they be included in the PYTHIA framework?

In Section 2, a brief introduction to the existing MI model in PYTHIA 8 is given. For more comprehensive details about what is contained in the model, readers are directed to [2] and the references therein. In Sections 3 and 4, an initial look at rescattering and enhanced screening is given. A summary and outlook is given in Section 5.

2 Multiple Interactions in PYTHIA 8

The MI model in PYTHIA 8 [3] is a model for non-diffractive events. It is an evolution of the model introduced in PYTHIA 6.3 [2], which in turn is based on the model developed in earlier versions of PYTHIA. The earliest model [4] was built around the virtuality-ordered parton showers available at the time and introduced many key features which are still present in the later models, such as $p_\perp$ ordering, perturbative QCD cross sections dampened at small $p_\perp$, a variable impact parameter, PDF rescaling, and colour reconnection.

The next-generation model [5, 6] was developed after the introduction of transverse-momentum-ordered showers, opening the way to have a common $p_\perp$ evolution scale for initial-state radiation (ISR), final-state radiation (FSR) and MI emissions. The second key ingredient was the addition of junction fragmentation to the Lund String hadronisation model, allowing the handling of arbitrarily complicated beam remnants. This permitted the MI framework to be updated to include a more complete set of QCD $2 \to 2$ processes, with the inclusion of flavour effects in the PDF rescaling.

The PYTHIA 8 MI framework also contains additional new features which are not found in previous versions, such as

- a richer mix of underlying-event processes (γ, J/ψ, Drell-Yan, etc.),
- the possibility to select two hard interactions in the same event, and
- the possibility to use one PDF set for hard processes and another for other subsequent interactions.

2.1 Interleaved $p_\perp$ Ordering

Starting in PYTHIA 6.3, ISR and MI were interleaved with a common $p_\perp$ evolution scale. In PYTHIA 8, this is taken a step further, with FSR now also fully interleaved. The overall probability for the i^{th} interaction or shower branching to take place at $p_\perp = p_{\perp i}$ is given by

$$
\begin{aligned}
\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}p_\perp} &= \left(\frac{\mathrm{d}\mathcal{P}_{\mathrm{MI}}}{\mathrm{d}p_\perp} + \sum \frac{\mathrm{d}\mathcal{P}_{\mathrm{ISR}}}{\mathrm{d}p_\perp} + \sum \frac{\mathrm{d}\mathcal{P}_{\mathrm{FSR}}}{\mathrm{d}p_\perp} \right) \\
&\times \exp\left(-\int_{p_\perp}^{p_{\perp i-1}} \left(\frac{\mathrm{d}\mathcal{P}_{\mathrm{MI}}}{\mathrm{d}p'_\perp} + \sum \frac{\mathrm{d}\mathcal{P}_{\mathrm{ISR}}}{\mathrm{d}p'_\perp} + \sum \frac{\mathrm{d}\mathcal{P}_{\mathrm{FSR}}}{\mathrm{d}p'_\perp} \right) \mathrm{d}p'_\perp \right),
\end{aligned} \tag{1}
$$

with contributions from MI, ISR and FSR unitarised by a Sudakov-like exponential factor.

If we now focus on just the MI contribution, the probability for an interaction is given by

$$
\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}p_{\perp i}} = \frac{1}{\sigma_{nd}} \frac{\mathrm{d}\sigma}{\mathrm{d}p_\perp} \exp\left(-\int_{p_\perp}^{p_{\perp i-1}} \frac{1}{\sigma_{nd}} \frac{\mathrm{d}\sigma}{\mathrm{d}p'_\perp} \mathrm{d}p'_\perp \right), \tag{2}
$$

where $\mathrm{d}\sigma/\mathrm{d}p_\perp$ is given by the perturbative QCD $2 \to 2$ cross section. This cross section is dominated by t-channel gluon exchange, and diverges roughly as $\mathrm{d}p_\perp^2/p_\perp^4$. 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

$$\frac{\mathrm{d}\hat{\sigma}}{\mathrm{d}p_\perp^2} \propto \frac{\alpha_S^2(p_\perp^2)}{p_\perp^4} \to \frac{\alpha_S^2(p_{\perp0}^2 + p_\perp^2)}{(p_{\perp0}^2 + p_\perp^2)^2}, \tag{3}$$

where $p_{\perp0}$ (related to $1/d$ above) is now a free parameter in the model.

2.2 Impact Parameter

Up to this point, all parton-parton interactions have been assumed to be independent, such that the probability to have n interactions in an event, $\mathcal{P}_n$, is given by Poissonian statistics. This picture is now changed, first by requiring that there is at least one interaction, such that we have a physical event, and second by including an impact parameter, b. The default matter distribution in PYTHIA is a double Gaussian

$$\rho(r) \propto \frac{1-\beta}{a_1^3} \exp\left(-\frac{r^2}{a_1^2}\right) + \frac{\beta}{a_2^3} \exp\left(-\frac{r^2}{a_2^2}\right), \tag{4}$$

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. The time-integrated overlap of the incoming hadrons during collision is given by

$$\mathcal{O}(b) = \int \mathrm{d}t \int \mathrm{d}^3x \ \rho(x,y,z) \ \rho(x+b,y,z+t), \tag{5}$$

after a suitable scale transformation to compensate for the boosted nature of the incoming hadrons.

Such an impact parameter picture has central collisions being generally more active, with an average activity at a given impact parameter being proportional to the overlap, $\mathcal{O}(b)$. While requiring at least one interaction results in $\mathcal{P}_n$ being narrower than Poissonian, when the impact parameter dependence is added, the overall effect is that $\mathcal{P}_n$ is broader than Poissonian. The addition of an impact parameter also leads to a good description of the "Pedestal Effect", where events with a hard scale have a tendency to have more underlying activity; this is as central collisions have a higher chance both of a hard interaction and of more underlying activity. This centrality effect naturally saturates at $p_{\perp hard} \sim 10\,\mathrm{GeV}$.

2.3 PDF Rescaling

In the original model, PDFs were rescaled only such that overall momentum was conserved. This was done by evaluating PDFs at a modified x value

$$x_i' = \frac{x_i}{1 - \sum_{j=1}^{i-1} x_j}, \tag{6}$$

where the subscript i refers to the current interaction and the sum runs over all previous interactions. The original model was affected by a technical limitation in fragmentation; it was only possible to take one valence quark from an incoming hadron. This meant that the MI framework

was limited to q$\overline{\text{q}}$ and gg final states and that it was not possible to have ISR from secondary scatterings. By introducing junction fragmentation, where a central junction is connected to three quarks and carries baryon number, these limitations were removed. This allowed the next-generation model to include a more complete set of MI processes and flavour effects in PDF rescaling.

ISR, FSR and MI can all lead to changes in the incoming PDFs. In the case of FSR, a colour dipole can stretch from a radiating parton to a beam remnant, leading to (a modest amount of) momentum shuffling between the beam and the parton. Both ISR and MI can result in large x values being taken from the beams, as well as leading to flavour changes in the PDFs. If a valence quark is taken from one of the incoming hadrons, the valence PDF is rescaled to the remaining number. If, instead, a sea quark (q$_s$) is taken from a hadron, an anti-sea companion quark (q$_c$) is left behind. The x distribution for this companion quark is generated from a perturbative ansatz, where the sea/anti-sea quarks are assumed to have come from a gluon splitting, g $\rightarrow$ q$_s$q$_c$. Subsequent perturbative evolution of the q$_c$ distribution is neglected. Finally, there is the issue of overall momentum conservation. If a valence quark is removed from a PDF, momentum must be put back in, while if a companion quark is added, momentum must be taken from the PDF. This is done by allowing the normalisation of the sea and gluon PDFs to fluctuate such that overall momentum is conserved.

2.4 Beam Remnants, Primordial $k_\perp$ and Colour Reconnection

When the $p_\perp$ evolution has come to an end, the beam remnant will consist of the remaining valence content of the incoming hadrons as well as any companion 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 $\sim 0.3\,\text{GeV}$, but when comparing to data, for instance the $p_\perp$ distribution of Z^0 at CDF, a value of $\sim 2\,\text{GeV}$ 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

$$\sigma(Q) = \max\left(\sigma_{min}, \sigma_\infty \frac{1}{1 + Q_{\frac{1}{2}}/Q}\right), \tag{7}$$

where Q is the hardness of a sub-collision, σ_{min} is a minimal value ($\sim 0.3\,\text{GeV}$), σ_∞ is a maximal value that is approached asymptotically and $Q_{\frac{1}{2}}$ is the Q value at which $\sigma(Q)$ is equal to half σ_∞. The recoil is shared among all initiator and remnant partons from the incoming hadrons, and the $k_\perp$ given to all daughter partons through a Lorentz boost.

The final step is colour reconnection. In the old MI framework, Rick Field found a good agreement to CDF data if 90% of additional interactions produced two gluons with "nearest

<center>(a)   (b)</center>

Fig. 1: (a) Two $2 \to 2$ scatterings, (b) a $2 \to 2$ scattering followed by a rescattering

neighbour" colour connections [9]. In PYTHIA 8, with its more general MI framework, colour reconnection is performed by giving each system a probability to reconnect with a harder system

$$\mathcal{P} = \frac{p_{\perp Rec}^2}{(p_{\perp Rec}^2 + p_\perp^2)}, \qquad p_{\perp Rec} = RR * p_{\perp 0},\qquad (8)$$

where RR, ReconnectRange, is a user-tunable parameter and $p_{\perp 0}$ is the same parameter as in eq. (3). The idea of colour reconnection can be motivated by noting that MI leads to many colour strings that will overlap in physical space. Moving from the limit of $N_C \to \infty$ to $N_C = 3$, it is perhaps not unreasonable to consider these strings to be connected differently due to a coincidence of colour, so as to reduce the total string length and thereby the potential energy. 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. Currently, however, given the lack of a firm theoretical basis, the need for colour reconnection has only been established within the context of specific models.

3 Rescattering

A process with a rescattering occurs when an outgoing state from one scattering is allowed to become the incoming state in another scattering. This is illustrated schematically in Figure 1, where (a) shows two independent $2 \to 2$ processes while (b) shows a rescattering process. An estimate for the size of such rescattering effects is given by Paver and Treleani [7], where a factorised form is used for the double parton distribution, giving the probability of finding two partons of given x values inside an incoming hadron. Their results show that, at Tevatron energies, rescattering is expected to be a small effect when compared against the more dominant case of multiple disconnected scatterings.

If we accept MI as real, however, then we should also allow rescatterings to take place. They would show up in the collective effects of MI, manifesting themselves as changes to multiplicity, $p_\perp$ and other distributions. After a retuning of $p_{\perp 0}$ and other model parameters, it is likely that their impact is significantly reduced, so we should therefore ask whether there are more direct ways in which rescattering may show up. Is there perhaps a region of low $p_\perp$ jets,

	Tevatron		LHC	
	Min Bias	QCD Jets	Min Bias	QCD Jets
Scatterings	2.81	5.11	5.21	12.20
Single rescatterings	0.37	1.20	0.93	3.64
Double rescatterings	0.01	0.03	0.02	0.11

Table 1: Average number of scatterings, single rescatterings and double rescatterings in minimum bias and QCD jet events at Tevatron ($\sqrt{s} = 1.96\,\text{GeV}$, QCD jet $\hat{p}_{\perp min} = 20\,\text{GeV}$) and LHC ($\sqrt{s} = 14.0\,\text{TeV}$, QCD jet $\hat{p}_{\perp min} = 50\,\text{GeV}$) energies

where an event is not dominated by ISR/FSR, where this extra source of three-jet topologies will be visible? A further consideration is that such rescatterings will generate more $p_\perp$ in the perturbative region, which may overall mean it is possible to reduce the amount of primordial $k_\perp$ and colour reconnections necessary to match data, as discussed in Section 2.4.

3.1 Rescattering in PYTHIA 8

If we begin with the typical case of small-angle t-channel gluon scattering, we can imagine that a combination of a scattered parton and a hadron remnant will closely match one of the incoming hadrons. In such a picture, we can write the complete PDF for a hadron as

$$f(x, Q^2) \rightarrow f_{rescaled}(x, Q^2) + \sum_n \delta(x - x_n) = f_u(x, Q^2) + f_\delta(x, Q^2), \quad (9)$$

where the subscript u/δ is the unscattered/scattered component. That is, each time a scattering occurs, one parton is fixed to a specific x_n value, while the remainder is still a continuous probability distribution. In such a picture, the momentum sum should still approximately obey

$$\int_0^1 x \left[f_{rescaled}(x, Q^2) + \sum_n \delta(x - x_n) \right] dx = 1. \quad (10)$$

Of course, in general, it is not possible to uniquely identify a scattered parton with one hadron, so an approximate prescription must be used instead, such as rapidity based. If we consider the original MI probability given in eqs. (1) and (2), we can now generalise this to include the effects of rescattering

$$\frac{d\mathcal{P}_{\text{MI}}}{dp_\perp} \rightarrow \frac{d\mathcal{P}_{\text{uu}}}{dp_\perp} + \frac{d\mathcal{P}_{\text{u}\delta}}{dp_\perp} + \frac{d\mathcal{P}_{\delta\text{u}}}{dp_\perp} + \frac{d\mathcal{P}_{\delta\delta}}{dp_\perp}, \quad (11)$$

where the uu component now represents the original MI probability, the uδ and δu components a single rescattering and the $\delta\delta$ component a double rescattering, where both incoming states to an interaction are previously scattered partons.

Some indicative numbers are given in Table 1, which shows the average number of scatterings and rescatterings for different types of event at Tevatron and LHC energies. The average distribution of such scatterings per event is also shown in Figure 2 for Tevatron minimum bias
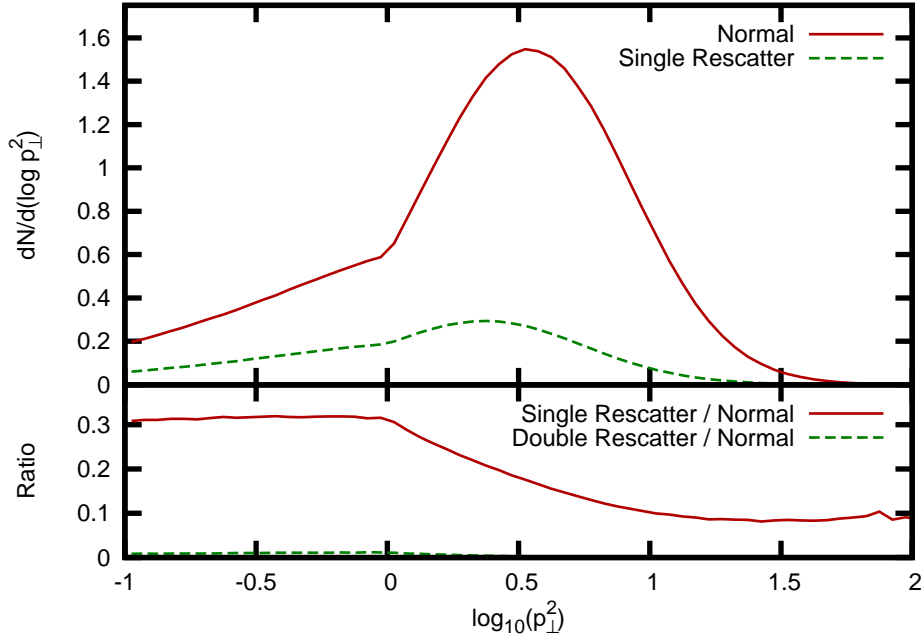
Fig. 2: Average distribution of scatterings, single rescatterings and double rescatterings per event ($\sqrt{s} = 1.96\,\mathrm{GeV}$, minimum bias). Double rescattering is not visible at this scale in the $\mathrm{d}N/\mathrm{d}(\log p_\perp^2)$ plot, but is visible in the ratio

events. In the upper plot of $\mathrm{d}N/\mathrm{d}(\log p_\perp^2)$, the suppression of the cross section at small $p_\perp^2$ is caused mainly by the regularisation outlined in eq. (3), but is also affected by the scaling violation in the PDFs. Below $p_\perp^2 \sim 1\,\mathrm{GeV}^2$, the PDFs are frozen, giving rise to an abrupt change in slope. Normal scatterings dominate, but there is a clear contribution from single rescatterings. In the upper plot, it is not possible to see the effects of double rescattering, but this is (barely) visible in the ratio plot below. Given the overall small contribution from double rescatterings, we neglect these in the following. As previously predicted, rescattering is a small effect at larger $p_\perp$ scales, but, when evolving downwards, its relative importance grows as more and more partons are scattered out of the incoming hadrons and become available to rescatter. Note that here, we classify the original scattering and the rescattering by $p_\perp$, but make no claims on the time ordering of the two.

3.2 Mean $p_\perp$ vs Charged Multiplicity

While a preliminary framework is in place which allows for hadronic final states, there are non-trivial recoil kinematics when considering the combination of rescattering, FSR and primordial $k_\perp$. With the dipole-style recoil used in the parton showers, a final-state radiating parton will usually shuffle momenta with its nearest colour neighbour. Without rescattering, colour dipoles are not spanned between systems, and individual systems will locally conserve momentum. With rescattering enabled, you instead have the possibility of colour dipoles spanning different scattering systems and therefore the possibility of an individual system no longer locally conserving

Fig. 3: Mean $p_\perp$ vs Charged Multiplicity, $|\eta| \leq 1$ and $p_\perp \geq 0.4\,\mathrm{GeV/c}$, CDF Run II data against Pythia 6.418 (Tune A) and Pythia 8.114 (default settings) with and without deferred FSR

momentum. When primordial $k_\perp$ is now added through a Lorentz boost, these local momentum imbalances can lead to global momentum non-conservation. In order to proceed and be able to take an initial look at the effects of rescattering on colour reconnection, a temporary solution of deferring FSR until after primordial $k_\perp$ is added has been used, as is done in PYTHIA 6.4.

We begin by studying the mean $p_\perp$ vs charged multiplicity distribution, $\langle p_\perp \rangle (n_{ch})$, from PYTHIA 6.418 (Tune A) and PYTHIA 8.114 (default settings), compared to the CDF Run II data ($|\eta| \leq 1$ and $p_\perp \geq 0.4\,\mathrm{GeV/c}$) [10]. For each run, the $p_{\perp 0}$ parameter of the MI framework is tuned so that the mean number of charged particles in the central region is maintained at the Tune A value. This is shown in Figure 3, where we can see that PYTHIA 6, using virtuality-ordered showers and the old MI framework, does a reasonable job of describing the data. PYTHIA 8 does not currently have a full tune to data, but does qualitatively reproduce the shape of the data when colour reconnection is turned on, up to an overall normalisation shift. It is clear that without colour reconnection, the slope of the curve is much too shallow and unlikely to describe the data, even given an overall shift. The same results with deferred FSR are also shown; the slope is marginally steeper, but still in the same region as without deferred FSR.

Figure 4 now shows the results when rescattering is enabled. Starting without any colour reconnection, we see that when rescattering is turned on, there is a rise in the mean $p_\perp$, but also that this is in no way a large gain. This is also the case when colour reconnection is turned on and tuned such that the curve qualitatively matches the shape of the Run II data. The amount of colour reconnection used is given in the form $RR * p_{\perp 0}$, as described in eq. (8). That a rise

Fig. 4: Mean $p_\perp$ vs Charged Multiplicity, PYTHIA 8.114 (deferred FSR), effects of rescattering

in the mean $p_\perp$ is there with rescattering, but small, is something that was observed already in an early toy model study. Now, when the full generation framework is almost there, it is clear that rescattering is not the answer to the colour reconnection problem. Other potential effects of rescattering remain to be studied.

4 Enhanced Screening

The idea of enhanced screening came from the modelling of initial states using dipoles in transverse space [11]. A model using an extended Mueller dipole formalism has recently been used to describe the total and diffractive cross sections in pp and γ^*p collisions and the elastic cross section in pp scattering [8]. In such a picture, initial-state dipoles are evolved forwards in rapidity, before two such incoming states are collided. In the model, as the evolution proceeds, the number of dipoles with small transverse extent grows faster than that of large dipoles. The dipole size, r, determines the screening length, which appears in the interaction cross section as a $p_\perp$ cutoff, $p_{\perp 0} \sim 1/r$. Smaller dipoles imply a larger effective cutoff, and an enhanced amount of screening. A rough calculation shows that this screening effect is expected to grow as the square root of the number of dipoles.

To model this in PYTHIA, we consider the $p_{\perp 0}$ parameter of the MI framework that encapsulates colour screening, as given in eq. (3). By scaling this value by an amount that grows as the amount of initial-state activity grows, this enhanced screening effect can be mimicked. Such
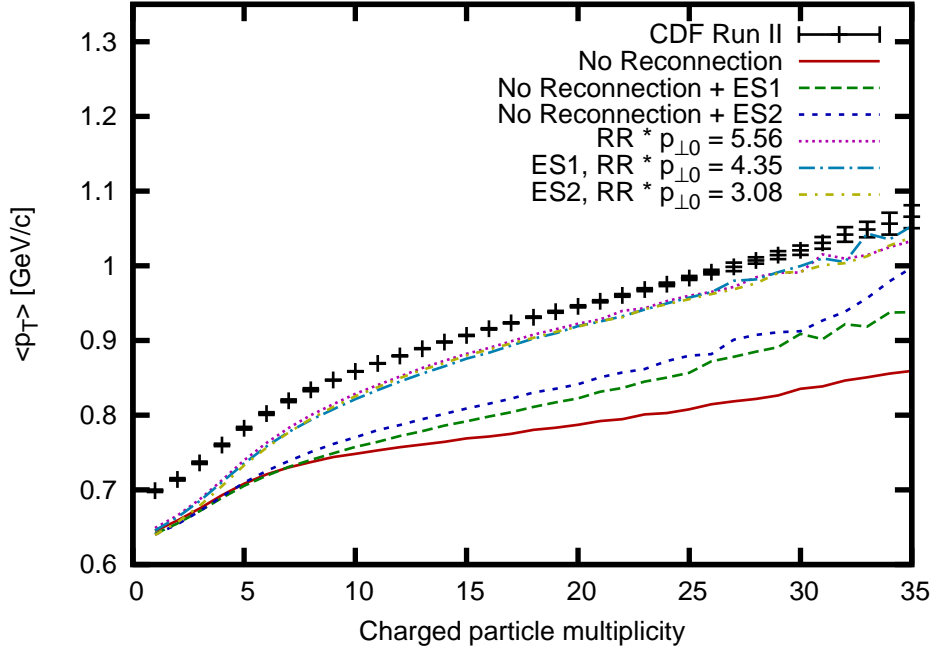
Fig. 5: Mean $p_\perp$ vs Charged Multiplicity, PYTHIA 8.114, effects of the enhanced screening ansatz

a change can be achieved by adjusting the weighting of the cross section according to

$$\frac{\mathrm{d}\hat{\sigma}}{\mathrm{d}p_\perp^2} \propto \frac{\alpha_S^2(p_{\perp 0}^2 + p_\perp^2)}{(p_{\perp 0}^2 + p_\perp^2)^2} \rightarrow \frac{\alpha_S^2(p_{\perp 0}^2 + p_\perp^2)}{(n\,p_{\perp 0}^2 + p_\perp^2)^2}, \tag{12}$$

where n takes a different meaning for two different scenarios. With the first scenario, ES1, n is set equal to the number of multiple interactions that have taken place in an event (including the current one). In the second, ES2, n is set equal to the number of MI+ISR interactions that have taken place in an event.

4.1 Mean $p_\perp$ vs Charged Multiplicity

We again study the $\langle p_\perp \rangle (n_{ch})$ distribution, this time with the enhanced screening ansatz. The results are given in Figure 5. Looking at the curves without colour reconnection, it is immediately apparent that both scenarios give a dramatic rise in the mean $p_\perp$, although not quite enough to explain data on their own. With colour reconnection now enabled and tuned, again so that the curves qualitatively match the shape of the Run II data, it is possible to noticeably reduce the amount of reconnection needed. With colour reconnection at these levels, there is still perhaps an uncomfortably large number of systems being reconnected, but the results are definitely encouraging. There are many more areas to study in relation to enhanced screening, but from these initial results, it is worth checking if it may play a role in reducing colour reconnections to a more comfortable level.

5 Conclusions

PYTHIA 8, the C++ rewrite of the PYTHIA event generator has now been released. It has been written with a focus on Tevatron and LHC applications, something that is evident given the sophisticated MI model present in the program. The original MI model, introduced in the early versions of PYTHIA, has been well proven when compared to experimental data. The new PYTHIA 8 MI framework, based on this original model, now generalises the physics processes available, as well as adding entirely new features.

We have also taken an early look at rescattering and enhanced screening, two new ideas for modifying the physics inside the MI framework. There is currently a preliminary framework for rescattering, although fully interleaved ISR, FSR and MI is still to come. It appears, at this early stage, that rescattering is not the answer to the colour reconnection problem, but there is still much more to investigate, such as three-jet multiplicities and other collective effects. The idea of enhanced screening leads to a simple ansatz that gives large changes when looking at the $\langle p_\perp \rangle (n_{ch})$ distribution. Again, there are still many questions to be asked, including how this modification affects other distributions.

References

[1] Field, R. D. and Hoeth, H. and Moraes, A. and Skands, P. presentation in these proceedings.

[2] Sjöstrand, Torbjorn and Mrenna, Stephen and Skands, Peter PYTHIA 6.4 physics and manual. *JHEP*, vol.05, pg.26, 2006.

[3] Sjöstrand, Torbjorn and Mrenna, Stephen and Skands, Peter A Brief Introduction to PYTHIA 8.1. *Comput. Phys. Commun.*, 178:852–867, 2008.

[4] Sjöstrand, Torbjorn and van Zijl, Maria A Multiple Interaction Model for the Event Structure in Hadron Collisions. *Phys. Rev.*, D36:2019, 1987.

[5] Sjöstrand, T. and Skands, P. Z. Transverse-momentum-ordered showers and interleaved multiple interactions. *Eur. Phys. J.*, C39:129–154, 2005.

[6] Sjöstrand, T. and Skands, P. Z. Multiple interactions and the structure of beam remnants. *JHEP*, vol.03, pg.53, 2004.

[7] Paver, N. and Treleani, D. MULTIPLE PARTON INTERACTIONS AND MULTI - JET EVENTS AT COLLIDER AND TEVATRON ENERGIES. *Phys. Lett.*, B156:252, 1984.

[8] Avsar, Emil and Gustafson, Gösta and Lönnblad, Leif Diffractive Excitation in DIS and pp Collisions. *JHEP*, vol.12, pg.12, 2007.
Flensburg, Christoffer and Gustafson, Gösta and Lönnblad, Leif Elastic and quasi-elastic pp and $\gamma^\star p$ scattering in the Dipole Model. *arXiv, hep-ph*, 0807.0325, 2008.

[9] Field, Rick D. The underlying event in hard scattering processes. *presentations at the 'Matrix Element and Monte Carlo Tuning Workshop', Fermilab, 4 October 2002 and 29-30 April 2003, talks available from webpage http://cepa.fnal.gov/psm/MCTuning/*

[10] Moggi, N and others Measurement of Inelastic $P\bar{P}$ Inclusive Cross Sections at $\sqrt{s} = 1.96$ TeV. *CDF Note 9337, http://www-cdf.fnal.gov/.*

[11] Gustafson, Gösta *private communication*.