Hard Diffraction
with Dynamic Gap Survival

Christine O. Rasmussen and Torbjörn Sjöstrand

Theoretical Particle Physics,

Department of Astronomy and Theoretical Physics,

Lund University,

Sölvegatan 14A,

SE-223 62 Lund, Sweden

Abstract

We present a new framework for the modelling of hard diffraction in pp and p$\overline{\text{p}}$ collisions. It starts from the the approach pioneered by Ingelman and Schlein, wherein the single diffractive cross section is factorized into a Pomeron flux and a Pomeron PDF. To this it adds a dynamically calculated rapidity gap survival factor, derived from the modelling of multiparton interactions. This factor is not relevant for diffraction in ep collisions, giving non-universality between HERA and Tevatron diffractive event rates. The model has been implemented in PYTHIA 8 and provides a complete description of the hadronic state associated with any hard single diffractive process. Comparisons with p$\overline{\text{p}}$ and pp data reveal improvement in the description of single diffractive events.

1 Introduction

The nature of diffractive excitation in hadron-hadron collisions remains a bit of a mystery. We may motivate *why* it happens, e.g. based on the optical analogy that lies behind its name, or in the related Good-Walker formalism [1]. But to explain *how* diffractive events are produced, and with what properties, is a longer story. In a first step the single diffractive cross section should be describable as a function of the diffractive mass M and the squared momentum transfer t. In a second step the generic properties of a diffractive system of a given mass should be explained: multiplicity distributions, rapidity and transverse momentum spectra and other event characteristics. In a third step the existence and character of exclusive diffractive processes and the underlying events associated therewith should be understood.

Over the years much data has accumulated, and many models have been presented, but so far without any model that explains all aspects of the data, and without any consensus which models are the most relevant ones. It is beyond the scope of the current article to review all the data and models; for a selection of relevant textbooks and reviews see [2, 3, 4, 5, 6, 7, 8].

For the path we will follow in this article, Regge theory provides the basic mathematical framework. In it, poles in the plane of complex spin α may be viewed as the manifestations of hadronic resonances in the crossed channels. A linear trajectory of poles $\alpha(t) = \alpha(0) + \alpha' t$ corresponds to a $\sigma_{\text{tot}} \sim s^{\alpha(0)-1}$. Several trajectories appear to exist, but for high-energy applications the most important is the Pomeron ($\mathbb{P}$) one, which with its $\alpha(0) > 1$ is deemed responsible for the observed rise of the total cross section, and in modern terminology would correspond to a set of glueball states. With single-Pomeron exchange as the starting point, higher orders involve multiple Pomeron exchanges, and also interactions between the Pomerons being exchanged, driven by a triple-Pomeron vertex. Out of this framework the cross section for various diffractive topologies can be derived, differentially in mass and t, given a set of numbers that must be extracted from data.

Such models do not address the structure of the diffractive system. The fireball models of older times implied isotropic decay in the rest frame of the diffractive system, or possibly elongated along the collision axis, but without internal structure. The Ingelman-Schlein (IS) model [9] made the bold assumption that the exchanged Pomeron could be viewed as a hadronic state, and that therefore a diffractive system could be described as a hadron-hadron collision at a reduced energy. This implied the existence of Parton Distribution Functions (PDFs) for the Pomeron. Thereby also hard processes became available, confirmed by the observation of jet production in diffractive systems [10]. The PomPyt program [11] combined Pomeron fluxes and PDFs, largely determined from HERA data, with the PYTHIA event generator of the time [12] to produce complete hadronic final states, and PomWig [13] did similarly for Herwig [14].

One limitation of these models is that they are restricted to the exchange of one Pomeron per hadron-hadron collision, not the multiple ones expected in Regge theory. Translated into a QCD-based, more modern view of such collisions, Multiple Partonic Interactions (MPIs) occur between the incoming hadrons [15]. That is, since hadrons are composite objects, there is the possibility for several partons from a hadron to collide, predominantly by semisoft $2 \to 2$ QCD interactions. These together create colour flows (strings [16]) criss-crossing the event, typically filling up the whole rapidity range between the two beam particles with hadron production. Thereby a "basic" process containing a rapidity gap can lose that by MPIs. (MPIs and soft colour exchanges could also be sources of gaps [17, 18], a possibility we will not study further in this article, so as to keep the discussion focussed.)

A spectacular example is Higgs production by gauge-boson fusion, $W^+W^- \to H^0$ and $Z^0Z^0 \to H^0$, where the naive process should result in a large central gap only populated by the Higgs decay products, since no colour exchange is involved. Including MPIs, this gap largely fills up

[19], although a fraction of the events should contain no further MPIs [20], a fraction denoted as the Rapidity Gap Survival Probability (RGSP). Such a picture has been given credence by the observation of "factorization breaking" between HERA and the Tevatron: the Pomeron flux and PDFs determined at HERA predicts about an order of magnitude more QCD jet production than observed at the Tevatron, e.g. [21].

In this article the intention is to provide a dynamical description of such factorization breaking, as a function of the hard process studied and its kinematics, and to predict the resulting event structure for hard diffraction in hadronic collisions. This is done in three steps. Firstly, given a hard process selected based on the inclusive PDFs, the fraction of a PDF that should be associated with diffraction is calculated, as a convolution of the Pomeron flux and its PDFs. Secondly, the full MPI framework of PYTHIA, including also the effects of initial- and final-state radiation, is applied to find the fraction of events without any further MPIs. Those events that survive these two steps define the diffractive event fraction, while the rest remain as regular nondiffractive events. Thirdly, diffractive events may still have MPIs within the $\mathbb{P}$p subsystem, and therefore the full hadron-hadron underlying-event generation machinery is repeated for this subsystem. The nondiffractive events are kept as they are in this step.

One should not expect perfect agreement with data in this approach; there are too many uncertainties that enter in the description. Nevertheless a qualitative description can be helpful, not only to understand the trend of existing data, but also to pave the way for future studies. The new framework we present here has been implemented as an integrated part of the PYTHIA 8.2 event generator [22], and can be switched on for any standard hard process. It thereby complements the already existing modelling of soft diffraction, i.e. of diffractive events with no discernible hard process. The dividing line between these two descriptions is not sharp, and in the future we will explore tensions between the two.

As should be clear from this introduction, our model is "just" a combination of the existing IS and RGSP ideas, and thus not anything fundamentally new. The devil lies in the details, however, and to the best of our knowledge nobody has previously worked out a complete model of this character.

The article is structured as follows. In Section 2 we introduce the new model framework, which then is validated in Section 3. Some tentative comparisons with data are presented in Section 4. The article concludes with a summary and outlook in Section 5.

2   The model

In this article we study hard diffraction, so this means we assume the presence of some hard process in the events of interest. Standard examples would be jet, Z^0 and $W^\pm$ production. By factorization a cross section involving partons i, j from incoming beams A, B can be written as

$$\sigma = \sum_{i,j} \iint \mathrm{d}x_1 \, \mathrm{d}x_2 \, f_{i/A}(x_1, Q^2) \, f_{j/B}(x_2, Q^2) \, \hat{\sigma}_{ij}(\hat{s} = x_1 x_2 s, Q^2) \ , \tag{1}$$

where $\hat{\sigma}$ is the parton-level cross section, integrated over relevant further degrees of freedom, like a $p_\perp$ range for jets.

Assuming Pomerons to have some kind of existence inside the proton, in the Ingelman-Schlein spirit, we introduce a Pomeron flux $f_{\mathbb{P}/\mathrm{p}}(x_{\mathbb{P}}, t)$, where $x_{\mathbb{P}}$ is the $\mathbb{P}$ momentum fraction and t its (spacelike) virtuality. The $\mathbb{P}$ has a partonic substructure, just like a hadron, and thus we can define PDFs $f_{i/\mathbb{P}}(x, Q^2)$. The PDF could also depend on the t scale, just like the photon has a PDF strongly dependent on its virtuality. For lack of a model for such a dependence we assume the $\mathbb{P}$ PDF is a suitable average over the t range probed. As a consequence we will not need t for most of the studies, and so it can be integrated out of the flux, $f_{\mathbb{P}/\mathrm{p}}(x_{\mathbb{P}}) = \int f_{\mathbb{P}/\mathrm{p}}(x_{\mathbb{P}}, t) \, \mathrm{d}t$.

Given the ansatz with Pomeron flux and PDF, the PDF of a proton can be split into one regular nondiffractive (ND) and one $\mathbb{P}$-induced diffractive (D) part,

$$f_{i/\mathrm{p}}(x, Q^2) = f_{i/\mathrm{p}}^{\mathrm{ND}}(x, Q^2) + f_{i/\mathrm{p}}^{\mathrm{D}}(x, Q^2) \ , \tag{2}$$

where

$$\begin{aligned} f_{i/\mathrm{p}}^{\mathrm{D}}(x, Q^2) &= \int_0^1 \mathrm{d}x_{\mathbb{P}} \, f_{\mathbb{P}/\mathrm{p}}(x_{\mathbb{P}}) \int_0^1 \mathrm{d}x' \, f_{i/\mathbb{P}}(x', Q^2) \, \delta(x - x_{\mathbb{P}} x') \\ &= \int_x^1 \frac{\mathrm{d}x_{\mathbb{P}}}{x_{\mathbb{P}}} \, f_{\mathbb{P}/\mathrm{p}}(x_{\mathbb{P}}) \, f_{i/\mathbb{P}}\left(\frac{x}{x_{\mathbb{P}}}, Q^2\right) \ . \end{aligned} \tag{3}$$

The assumption that the diffractive part $f_{i/\mathrm{p}}^{\mathrm{D}}(x, Q^2; x_{\mathbb{P}}, t)$ of the full PDF can be decomposed in this way is in approximate agreement with the HERA data [23].

For two incoming protons (or antiprotons, or other hadrons) A and B, an initial probability for diffraction $\mathcal{P}^{\mathrm{D}} \approx \mathcal{P}_A^{\mathrm{D}} + \mathcal{P}_B^{\mathrm{D}}$ is obtained from the ratio of diffractive to inclusive PDFs,

$$\begin{aligned} \mathcal{P}_A^{\mathrm{D}} &= \frac{f_{i/B}^{\mathrm{D}}(x_B, Q^2)}{f_{i/B}(x_B, Q^2)} \quad \text{for} \quad AB \to XB \ , \\ \mathcal{P}_B^{\mathrm{D}} &= \frac{f_{i/A}^{\mathrm{D}}(x_A, Q^2)}{f_{i/A}(x_A, Q^2)} \quad \text{for} \quad AB \to AX \ , \end{aligned} \tag{4}$$

where $\mathcal{P}_{A/B}^{\mathrm{D}}$ is the probability for side A/B to be the diffractive system, thus being dependent on the variables of the opposite side.

This probability is used to determine, on an event-by-event basis, the nature of the selected hard scattering, whether diffractive or not. Currently we concentrate on single diffraction. A natural extension would be to associate the product $\mathcal{P}_A^{\mathrm{D}} \mathcal{P}_B^{\mathrm{D}}$ with central diffraction (CD), where two Pomerons collide and one parton is extracted from each $\mathbb{P}$. It would also be possible to extend the formalism such that part of the SD rate is reassigned as double diffraction (DD), where the hard collision happens inside one of the two diffractive systems. Neither CD nor DD are considered in this first study, however. Instead, for the fraction $\mathcal{P}_A^{\mathrm{D}} \mathcal{P}_B^{\mathrm{D}}$ of events, which normally is small anyway, a random choice is made between $AB \to AX$ and $AB \to XB$.

The key aspect of the model is now that it contains a dynamical gap survival. This means that we do not allow any further MPIs to occur between the two incoming hadrons, so as to ensure that the gap survives. In practise the tentative classification as diffractive, based on eq. (4), initially has no consequences: all events are handled as if they were nondiffractive hadron-hadron collisions.

Only if no additional MPIs occur does a diffractive classification survive and only then is the $\mathbb{P}$p subsystem set up. Specifically the $x_{\mathbb{P}}$ value is selected according to the distribution implied by eq. (4), and also a t value is selected for the outgoing proton. Technically, it is only at this stage that "pure" samples of diffractive events can be selected, should one wish to single out such events.

Once the $\mathbb{P}$p system has been set up, it is allowed to develop a partonic structure just like any hadron-hadron collision. Both initial-state radiation (ISR) and final-state radiation (FSR) thereby dress the original hard process by the emission of further softer partons. Also further MPIs inside this system are allowed, based on the $f_{i/\mathbb{P}}(x, Q^2)$ PDFs, successively modified to take into account the momentum and flavours already carried away by the MPI, ISR and FSR activity at $p_\perp$ scales above the currently considerd one, just like for nondiffractive systems.

The ISR/FSR/MPI description is based on the perturbative parton picture. Nonperturbative aspects have to be added to this. Beam remnants carry the momentum not kicked out of the

incoming $\mathbb{P}$ and p. For the former a fictitious "valence quark" content of either d$\bar{\text{d}}$ or u$\bar{\text{u}}$ is chosen at random for each new event. It is essentially equivalent to having a gluon as remnant, but is slightly more convenient. All outgoing partons are colour-connected by colour flux lines - strings - that fragment to produce the primary hadrons of the final state. The colour flow in an event is not unambiguously determined, however, and data suggest that colours tend to be more correlated than naively comes out of the perturbative picture, a phenomenon known as Colour Reconnection (CR).

We can by combining these two simple ideas give an explanation of the discrepancies between Tevatron and HERA. The dynamical gap survival introduces an additional suppression factor, reducing the number of diffractive events without any additional parameters.

2.1 Pomeron fluxes and PDFs

For numerical studies it is necessary to specify Pomeron flux and PDF parametrizations. There are currently seven parametrizations/models for the former and five for the latter available in PYTHIA.

The parametrizations for the Pomeron flux $f_{\mathbb{P}/\text{p}}(x_{\mathbb{P}}, t)$ are

- Schuler-Sjöstrand model (SaS) [24],
- the Bruni-Ingelman model [25],
- the Streng-Berger model [26],
- the Donnachie-Landshoff model [27],
- the Minimum Bias Rockefeller model (MBR) [28] with an option to renormalize the flux, and
- the H1 models Fit A and B [29, 30].

All have to obey an approximate form $f_{\mathbb{P}/\text{p}}(x_{\mathbb{P}}) \sim 1/x_{\mathbb{P}}$ in order to obtain an approximate diffractive mass spectrum $\sim \mathrm{d}M_X^2/M_X^2$, as required by Regge theory and by data. Just like the rise of the total cross section requires a supercritical Pomeron $\alpha(0) = 1 + \epsilon > 1$, with $\epsilon \approx 0.08$, several of the fluxes have adapted this steeper slope $f_{\mathbb{P}/\text{p}}(x_{\mathbb{P}}) \sim 1/x_{\mathbb{P}}^{1+2\epsilon}$ (where the factor of 2 in front of ϵ comes from the optical theorem). There are also some attempts to account for an excess in the low-mass resonance region. The t dependence is typically parametrized as a single exponential $f_{\mathbb{P}/\text{p}}(x_{\mathbb{P}}, t) \sim \exp(B_{\text{SD}}t)$, but also as a sum of two exponentials, or as a (power-like) dipole form factor. The MBR model differs from the others, since the model renormalizes the flux to unity. This renormalization suppresses the flux, thus making the dynamical gap survival obsolete. In order to make direct comparisons to the other available flux-models, we have implemented the renormalization as an option with the default being the non-renormalized flux.

The parametrizations for the Pomeron PDFs $f_{i/\mathbb{P}}(x, Q^2)$ are

- PomFix, a simple (toy) Q^2-independent parametrization,
- the H1 Fit A and B NLO PDFs [29],
- the H1 Jets NLO PDF [31], and
- the H1 Fit B LO PDF [29].

The first of these has a momentum sum of unity, whereas the latter four are not normalized to any specific value, the argument used being that the Pomeron is not a real particle and so does not obey that kind of constraints [32, 33]. (Technically H1 chose to normalize the $\mathbb{P}$ flux to unity at $x_{\mathbb{P}} = 0.003$, and then let the PDF normalization float.) Pragmatically it could be argued that what is measured is the convolution of the $\mathbb{P}$ flux and the $\mathbb{P}$ PDF, so that is is feasible to shuffle any constant number between the two. Unfortunately this makes it less trivial to mix freely, and makes it almost a necessity to combine H1 PDFs with H1 fluxes.

This is not the end of the story from an event-generator point of view, however. In the H1

parametrizations the momentum sums to approximately 0.5, but this does not mean that half of the $\mathbb{P}$ momentum in the $\mathbb{P}$p collision can just be thought away. At the very least this other half has to be considered as an inert component that sails through without interacting, but is present in the beam remnant. A further complication arises when MPIs are introduced. Normally these are generated in a sequence of decreasing $p_\perp$, with the PDFs for an MPI adjusted to take into account the momentum and flavours carried away by the preceding MPIs. So if 0.4 of the $\mathbb{P}$ momentum has already been taken, does it mean that 0.1 or 0.6 of it remains? This is an issue that did not exist at HERA, where MPIs are negligible outside of the photoproduction region. The choice made in PYTHIA is to assume that the full $\mathbb{P}$ momentum is available for MPIs. Furthermore we allow the option to rescale the H1 PDFs by a constant factor so as to change the momentum, notably by a factor of two to restore (approximately) the momentum sum rule. This should then be compensated by a corresponding rescaling of the $\mathbb{P}$ flux in the opposite direction. That way the $\mathbb{P}$ can be brought closer to an ordinary hadron, and more $\mathbb{P}$ fluxes can be used.

Another problem is that most PDF fits are NLO ones. Given the sparsity of data, it should be clear that "NLO accuracy" does not mean the same thing as it does for the inclusive proton PDF. Since PYTHIA only contains LO matrix elements (MEs) for QCD processes there is no extra bonus for using NLO PDFs. Worse, it is well known that the gluon PDF (of the proton) is much smaller in NLO than in LO for small x and Q^2; in principle it can even become negative. This behaviour compensates for the NLO MEs being larger than the LO ones in this region, but the compensation is nontrivial. Therefore an all-LO description, for all its weaknesses, is more robust in the small-$p_\perp$ region, which is where the MPI machinery largely operates. The default choice thus is H1 Fit B LO.

Finally also the inclusive proton PDF $f_{i/\mathrm{p}}(x, Q^2)$ should be chosen. Here several options come with PYTHIA, and many more can be obtained via the interfaces to LHAPDF5 and LHAPDF6 [34, 35]. The current default set is the NNPDF 2.3 QCD+QED LO one with $\alpha_\mathrm{s}(M_\mathrm{Z}) = 0.130$ [36]. The argument for using LO has already been outlined above. Since the proton PDF is much better constrained than that of the $\mathbb{P}$, there is less of a point in varying it between different options consistent with current p data. Note that, for diffractive events, the dependence on the original choice of proton PDF is largely removed on the $\mathbb{P}$ side by applying eq. (4). It does remain on the proton side, and in the dynamical calculation of rapidity gap survival, however.

2.2 MPI phenomenology

The QCD $2 \to 2$ processes are dominated by t-channel gluon exchange, which gives a perturbative cross section $\mathrm{d}\hat{\sigma}/\mathrm{d}p_\perp^2 \sim \alpha_\mathrm{s}^2(p_\perp^2)/p_\perp^4$ that diverges in the $p_\perp \to 0$ limit. Two modifications are needed to make sense out of this divergence.

Firstly a divergent integrated QCD cross section should not necessarily be construed as a divergent total pp cross section. Rather a $\mu = \sigma_{2\to2}^\mathrm{tot}/\sigma_\mathrm{pp}^\mathrm{tot} > 1$ for $p_\perp > p_{\perp\mathrm{min}}$ should be interpreted as implying an average of μ such partonic interactions per pp collision. Overall energy-momentum conservation will reduce the naively calculated rate, but would still kick out essentially all beam momentum if we allow $p_{\perp\mathrm{min}} \to 0$, in contradiction with the presence of well-defined beam jets wherein a single particle can carry an appreciable fraction of the incoming beam momentum.

Secondly, therefore, it is important to note the presence of a screening mechanism: whereas standard perturbation theory is based on asymptotically free incoming states, reality is that partons are confined inside colour singlet states. This introduces a nonperturbative scale of the size of a hadron, or rather of the average distance d between two opposite-colour charges. In this spirit we introduce a free parameter $p_{\perp 0} \sim 1/d$ that is used to dampen the cross section

$$\frac{\mathrm{d}\sigma}{\mathrm{d}p_\perp^2} \propto \frac{\alpha_\mathrm{s}^2(p_\perp^2)}{p_\perp^4} \longrightarrow \frac{\alpha_\mathrm{s}^2(p_{\perp 0}^2 + p_\perp^2)}{(p_{\perp 0}^2 + p_\perp^2)^2} \; . \tag{5}$$

Technically the dampening is implemented as an extra factor multiplying the standard QCD $2 \to 2$ cross sections, but could equally well have been associated with a dampening of the PDFs; it is only the product of these that enters in measurable quantities.

Empirically, a $p_{\perp 0}$ scale of $2 - 3$ GeV is required to describe data. This scale is larger than expected from the proton size alone, and is also in a regime where normally one would expect perturbation theory to be valid. The $p_{\perp 0}$ scale appears to increase slowly with energy, which is consistent with the growth of the number of gluons at smaller x values, leading to a closer-packing of partons and thereby a reduced screening distance d. A similar parametrization is chosen as for the rise of total cross section

$$p_{\perp 0}(E_{\mathrm{CM}}) = p_{\perp 0}^{\mathrm{ref}} \times \left(\frac{E_{\mathrm{CM}}}{E_{\mathrm{CM}}^{\mathrm{ref}}} \right)^{E_{\mathrm{CM}}^{\mathrm{pow}}}, \tag{6}$$

with $E_{\mathrm{CM}}^{\mathrm{pow}}$ and $p_{\perp 0}^{\mathrm{ref}}$ being tunable parameters and $E_{\mathrm{CM}}^{\mathrm{ref}}$ a reference energy scale.

With the protons being extended objects, the amount of overlap between two incoming ones strongly depends on the impact parameter b. A small b will allow for many parton-parton collisions, i.e. a high level of MPI activity, and a close-to-unity probability for the incoming protons to interact. A large b, on the other hand, gives less average activity and a higher likelihood that the protons pass by each other unaffected. Diffractive events predominantly occur in peripheral collisions, a concept well-known already from the optical point of view. In our approach it comes out naturally since we only allow one interaction to occur, namely the hard process of interest; if there is a second one this will fill the rapidity gap and kill the diffractive nature.

The shape of the proton and the resulting overlap – the convolution of the two incoming proton distributions – is not known in any detail. The proton electric charge distribution may give some hints, but measures quarks only and not gluons, and is in the static limit. Instead a few different simple parametrizations can be chosen:

- a simple Gaussian, offering no free parameters,
- a double Gaussian, i.e. a sum of two Gaussians with different radii and proton momentum fractions, and
- an overlap of the form $\exp(-b^p)$ (which does not correspond to a simple shape for the individual proton), with p a free parameter.

(A further option is a Gaussian with an x-dependent width, but this has not been implemented in a diffractive context.) All are normalized to unit momentum sum for the incoming partons, and an overall radius normalization factor is fixed by the total cross section.

The more uneven the matter distribution, the broader will the charged multiplicity distribution be. Notably the higher the overlap for central collisions, the higher the tail to very large multiplicities. Also other measures, like forward-backward correlations, probe the distribution. Unfortunately it is always indirectly, and closely correlated with other model details. As an example we can mention that the earliest tunes worked with a much lower $p_{\perp 0}$ than today and with double Gaussians rather far away from the single-Gaussian behaviour. This changed when more modern PDFs started to assume a steeper rise of the gluon PDF at small x, and when the PYTHIA parton showers were extended to apply to all MPIs rather than only the hardest one, and for some other improvements over the years. Currently best fits are not very far away from a simple Gaussian, e.g. with an overlap like $\exp(-b^{1.85})$, but still on the side of more peaked than a Gaussian.

An event that contains a high-$p_\perp$ interaction is likely to be more central than one that does not, since the former has more MPIs and therefore more chances that the hardest of these reaches a high $p_\perp$. This bias effect is included in the choice of a b for an event where the hardest interaction has been given, and is used in the subsequent generation of MPIs. For the current study of hard

diffraction this means that the hard process is initially picked biased towards smaller b values, but afterwards the central b region is strongly suppressed because the likelihood of several MPIs is so big there.

Starting from a hard interaction scale, and a selected b, the probability for an MPI at a lower scale has the characteristic form

$$\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}p_\perp^2} = O(b) \, \frac{1}{\sigma_{\mathrm{ref}}} \, \frac{\mathrm{d}\sigma_{\mathrm{QCD}}}{\mathrm{d}p_\perp^2} \ . \tag{7}$$

Here $O(b)$ is the overlap enhancement/depletion factor, $\mathrm{d}\sigma_{\mathrm{QCD}}$ the differential cross section for all $2 \to 2$ QCD processes, and σ_{ref} the total cross section for the event classes affected by the QCD processes. Historically σ_{ref} has been equated with the nondiffractive cross section in PYTHIA, on the assumption that diffraction only corresponds to a negligible fraction of $\mathrm{d}\sigma_{\mathrm{QCD}}$. Within the current framework a reformulation to use the full inelastic cross section would make sense, but would require further work and retuning, and is therefore left aside for now.

Given eq. (7) as a starting point, MPIs can be generated in a falling $p_\perp$ sequence, using a Sudakov-style formalism akin to what is used in parton showers. Actually, in the complete generation the MPI, ISR and FSR activity is interleaved into one common $p_\perp$-ordered chain of interactions and branchings, with one common "Sudakov form factor", down to the respective cutoff scales.

In the current case, the MPI formalism is used twice. Firstly, to determine whether an event is diffractive, and if not to generate the complete nondiffractive event. Secondly, for diffractive events, to determine the amount of MPI activity within the $\mathbb{P}$p system. Here eq. (7) can be reused, but with new meaning for the components of the equation.

- The $\mathrm{d}\sigma_{\mathrm{QCD}}/\mathrm{d}p_\perp$ is now evaluated using the $\mathbb{P}$ PDF on one side, but with the same damping as in eq. (5), where E_{CM} in eq. (6) is now the $\mathbb{P}$p invariant mass. If the $\mathbb{P}$ is supposed to have a smaller size than the proton then this could be an argument for a higher $p_{\perp 0}$ in this situation, but we have not here pursued this.

- The σ_{ref} now represents the $\mathbb{P}$p total cross section, an unknown quantity that relates to the normalizations of the $\mathbb{P}$ flux and $\mathbb{P}$ PDF. By default is is chosen to have a fixed value of 10 mb, higher than is normally quoted in literature. This way, with other quantities at their default settings, the charged multiplicity of a $\mathbb{P}$p collision agrees reasonably well with that of a nondiffractive pp one at the same invariant mass. This may not be the best of arguments, but is a reasonable first choice that is experimentally testable, at least in principle.

- The $O(b)$ factor may be changed, see next.

The impact parameter $b_{\mathbb{P}\mathrm{p}}$ of the $\mathbb{P}$p subcollision does not have to agree with the b_{pp} of the whole pp collision. It introduces the transverse matter profile of the Pomeron, even less known than that of the proton. Generally a Pomeron is supposed to be a smaller object in a localized part of the proton, but one should keep an open mind. For lack of better, three possibilities have been implemented, which can be compared to gauge the impact of this uncertainty.

- $b_{\mathbb{P}\mathrm{p}} = b_{\mathrm{pp}}$. This implicitly assumes that a Pomeron is as big as a proton and centered in the same place. Since small b_{pp} values already have been suppressed, by the MPI selection step, it implies that few events will have high enhancement factors.

- $b_{\mathbb{P}\mathrm{p}} = \sqrt{b_{\mathrm{pp}}}$ (where normalization is such that $\langle b \rangle = 1$ for minimum-bias events). This can crudely be motivated as follows. In the limit that the $\mathbb{P}$ is very tiny, such that the proton matter profile varies slowly over the width of the $\mathbb{P}$, then what matters is where the Pomeron strikes the other proton. Thus the variation of $O(b)$ with b is that of one proton, not two, and so the square root of the normal variation, loosely speaking. Technically this is messy to implement, but the current simple recipe provides the main effect of reducing the variation, bringing all b values closer to the average.

- Pick a completely new $b_{\mathbb{P}p}$, as was done with b_{pp} in the first place. This allows a broad spread from central to peripheral values, and thereby also a larger and more varying MPI activity inside the diffractive system than the other two options, and thereby offers a useful contrast.

3 Validation

In this section we summarise some of the tests and sanity checks we have performed on the model implementation. This provide insight into how the model operates and with what general results, but also highlights the uncertain nature of many of the components of the model.

In the model we have two options for when an event is classified as diffractive: either right after the event has passed the PDF selection criterion, eq. (4), or after passing the further MPI criterion. Results using only the former will from now on be denoted "PDF selected", and with the latter in addition "MPI selected". Our full model for hard diffraction corresponds to the latter, but the intermediate level is helpful in separating the effects of these two rather different physics components.

Notably, many distributions tend to be mainly determined by one of the two criteria. Those mainly sensitive to the PDF selection include the $x_{\mathbb{P}}$ and thereby the mass of the diffractive system, and the squared momentum transfer t of the process and thereby the scattering angle θ of the outgoing proton. In particular we will explore the dependence on Pomeron fluxes and PDFs. Aspects that depend on the details of the MPI model include several particle distributions, such as multiplicities, and that will also be highlighted.

The key number where both components are comparably important is the overall diffractive rate, where each of them gives an order-of-magnitude suppression, resulting in a $\sim 1\%$ fraction of hard events being of a diffractive nature. This number thereby receives a considerable overall uncertainty.

3.1 The Pomeron flux and PDF

We begin by studying the effects of variations of the $\mathbb{P}$ parametrizations. In Figs. 1a and 1b the seven different Pomeron fluxes are compared. As can be seen there is a considerable spread. Even in the region of medium $x_{\mathbb{P}}$ values, $x_{\mathbb{P}} \sim 0.1$, this corresponds to more than a factor of two between the extremes. The dramatic differences at large $x_{\mathbb{P}}$ are not readily visible, since a large-$x_{\mathbb{P}}$ event usually corresponds to a small rapidity gap and therefore is difficult to discern from non-diffractive events. The limit of small $x_{\mathbb{P}}$ generally is more interesting, tying in with the intercept of the Pomeron trajectory, but plays a lesser role for the current study of hard diffraction.

Turning to the Pomeron PDFs, a detailed comparison would entail the separate quark and gluon distributions at varying Q^2 scales. To simplify we show the QCD-charge-weighted sum

$$F_{\mathbb{P}}(x, Q^2) = \frac{4}{9} \sum_{i=q,\bar{q}} x f_{i/\mathbb{P}}(x, Q^2) + x g_{\mathbb{P}}(x, Q^2) \tag{8}$$

at a single value $Q^2 = 100 \text{ GeV}^2$, Figs. 1c and 1d. We notice that they all tend to be significantly harder than the corresponding proton PDF, here represented by the NNPDF 2.3 QCD+QED LO one. (The PomFix option is just a toy one, shown for completeness, but not used in the following.) For the gluon on its own, the $\mathbb{P}$ is significantly harder than the p, consistent with the idealized picture of a $\mathbb{P}$ as a glueball state with two "valence gluons", Figs. 2a and 2b. Surprisingly, also the quark PDFs of the $\mathbb{P}$ (Figs. 2c and 2d) are harder than proton ones, suggesting the presence of "valence quarks" in the $\mathbb{P}$, although an order of magnitude below the gluons. Another observation
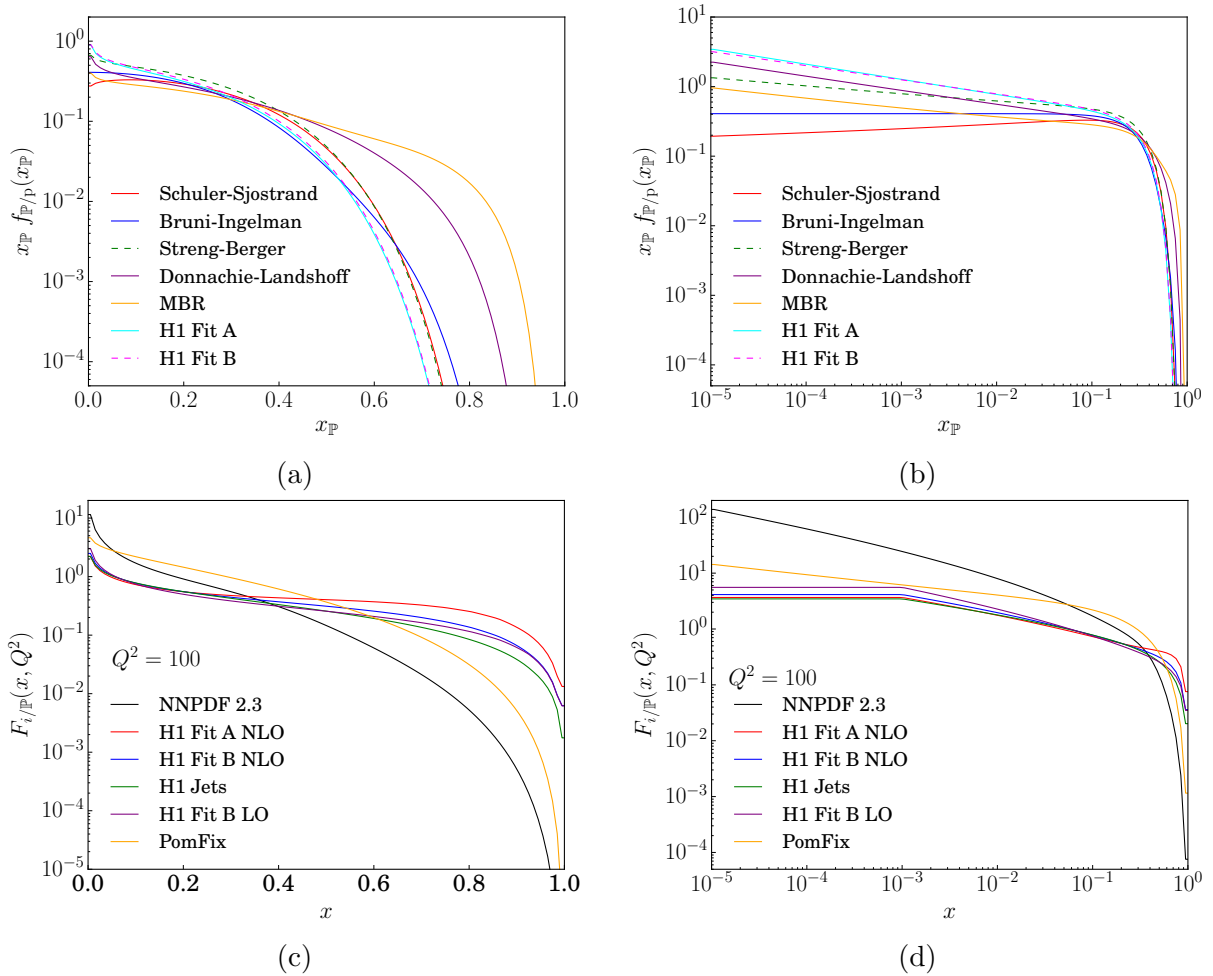
Figure 1: The seven different Pomeron fluxes included in PYTHIA on linear (a) and logarithmic scale (b). Note that the MBR flux has not been renormalized (see [28]). The QCD charge-weighted sum, eq. (8), of the five different Pomeron PDFs compared to the NNPDF 2.3 proton PDF on linear (c) and logarithmic scale (d).

is that the $\mathbb{P}$ PDF sets we compare are all based on H1 analyses, with largely the same data and with correlated assumptions for the definition of diffractive events. This is especially notable in the quark distributions, which are close to identical. Also the close affinity of gluons at lower x values should not be overstressed. Finally, note that the H1 parametrizations only apply down to $x = 10^{-3}$, and are frozen below that. This is likely to underestimate the low-x rise of PDFs, which as well could have been of the same shape as in the proton.

In the end, what matters is the convolution of the $\mathbb{P}$ flux with its PDFs, and that is shown in Fig. 3. There would be too many combinations possible to show individually, so we only indicate the range of possibilities and a few specific combinations. This may be on the extreme side, since some fluxes and PDFs come as fixed pairs, not really intended to be mixed freely. The key feature to note is that in this convolution the Pomeron part is now falling steeper at large x than the proton as a whole. This has the immediate consequence that diffractive hard subcollisions are not necessarily going to be produced more in the forwards direction than the bulk of corresponding nondiffractive ones, but on the contrary may be more central. The difference is not all that dramatic, however. It is also partly compensated by a somewhat slower increase of the $\mathbb{P}$ towards lower x values, a feature that derives from the artificial freezing of the $\mathbb{P}$ PDF below $x = 10^{-3}$.
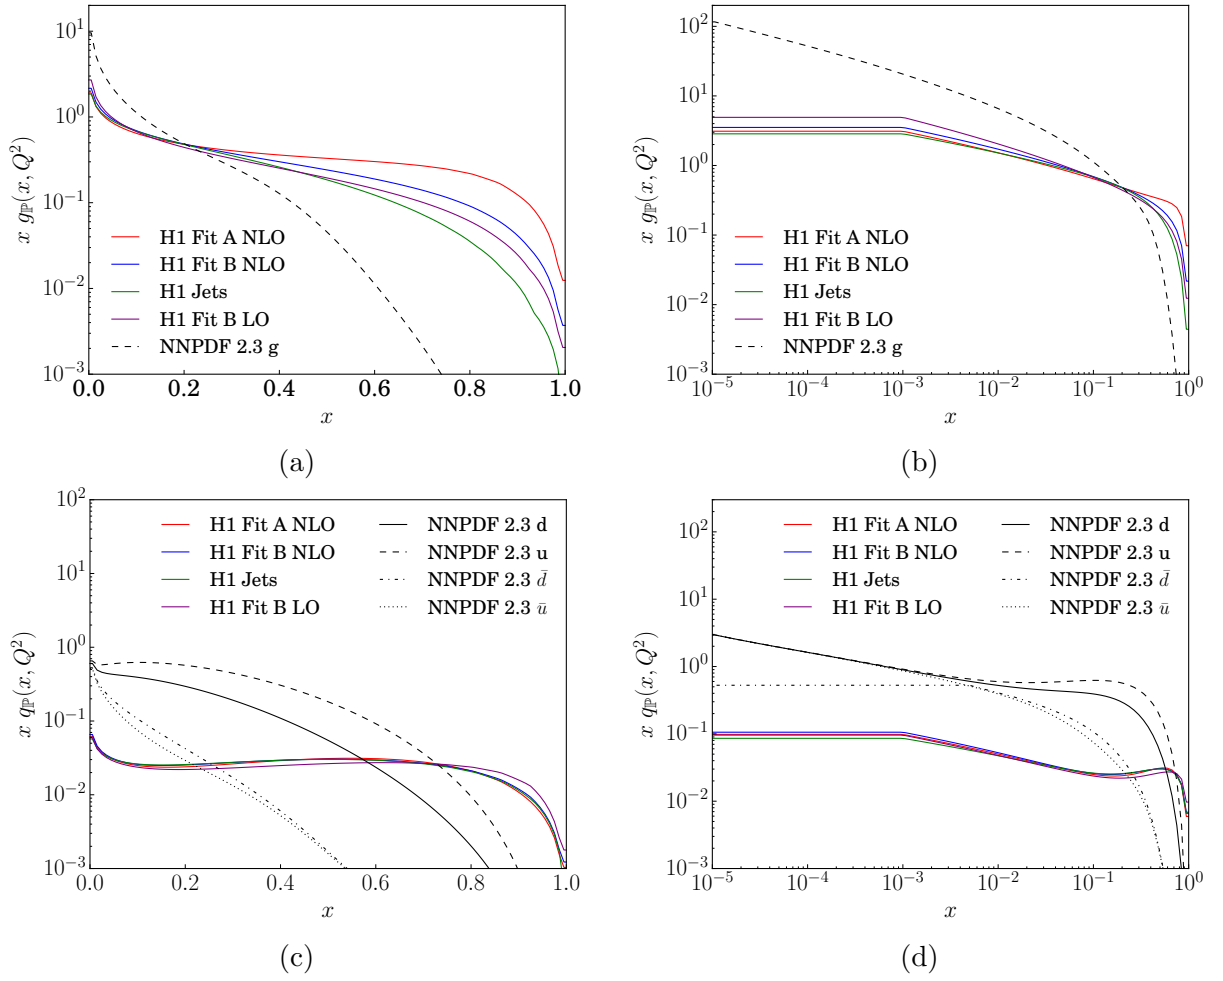
Figure 2: The $\mathbb{P}$ gluon distribution on linear (a) and logarithmic (b) scale. The $\mathbb{P}$ quark and antiquark distributions on linear (c) and logarithmic (d) scale. Both compared to the NNPDF 2.3 proton PDF distributions. Note that for the $\mathbb{P}$ we have d = u = s = $\bar{\text{d}}$ = $\bar{\text{u}}$ = $\bar{\text{s}}$(= c = $\bar{\text{c}}$), where the c, $\bar{\text{c}}$ are only included in H1Jets.
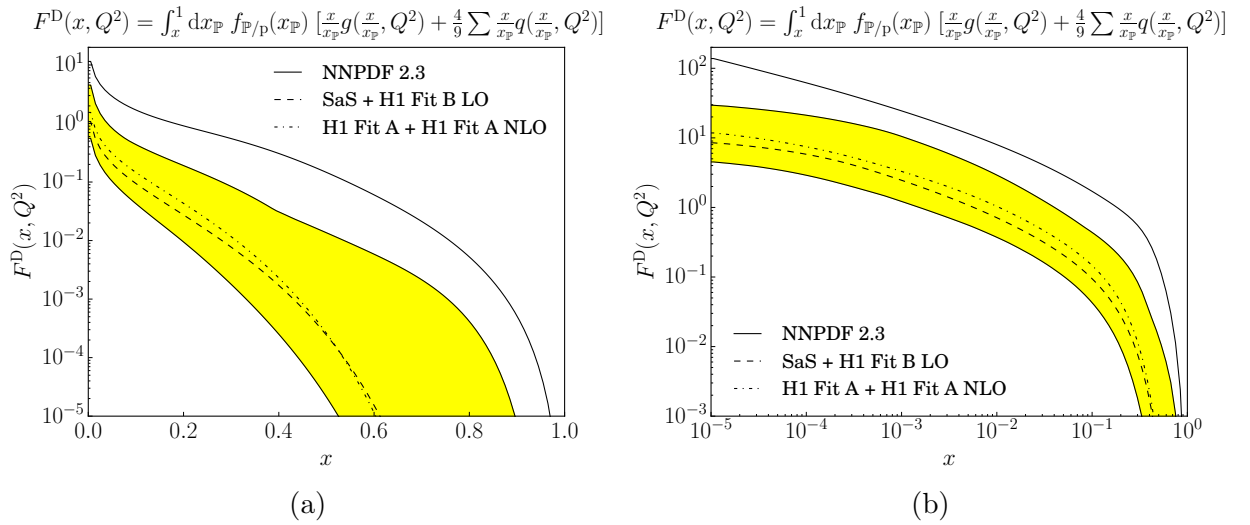
Figure 3: The convolution of Pomeron fluxes and PDFs for a few cases, with the range between the extremes marked in yellow; (a) linear and (b) logarithmic x scale.

Diffractive fractions pp collisions at $\sqrt{s} = 8$ TeV		
$\mathbb{P}$ PDF $\mathbb{P}$ flux	PDF selection	MPI selection
H1 Fit B LO SaS	(14.33 ± 0.11) %	(0.98 ± 0.03) %
H1 Fit B LO MBR	(14.79 ± 0.11) %	(0.96 ± 0.03) %
H1 Jets SaS	(13.70 ± 0.11) %	(0.92 ± 0.03) %
H1 Fit A NLO H1 Fit A	(20.55 ± 0.13) %	(1.35 ± 0.04) %
H1 Fit B LO H1 Fit A	(18.49 ± 0.12) %	(1.32 ± 0.04) %

Table 1: Diffractive fractions for the $2 \to 2$ QCD processes with $p_\perp > 20$ GeV obtained with PYTHIA 8. The samples have been produced without any phase-space cuts.

The lack of major shape differences between the $\mathbb{P}$ part and the rest will be visible in the more detailed studies later on.

Because of the close similarity of the different (but related) $\mathbb{P}$ PDFs at low-to-medium x, the bulk of the differences come from the $\mathbb{P}$ fluxes. We have chosen to exemplify this for $2 \to 2$ QCD processes with $p_\perp > 20$ GeV in $\sqrt{s} = 8$ TeV pp collisions, with the diffractive fractions for a few combinations shown in Table 1.

Note that changing the Pomeron parametrizations changes the fraction of events passing the PDF selection, but that the suppression factor introduced by the dynamical gap survival is about ~ 0.07 for all combinations in Table 1. This reflects the fact that neither the MPI model nor the proton PDF are influenced by the Pomeron parametrization, hence the probability for obtaining no additional MPIs in the pp system should not change. (This does not have to hold in general, but here we compare very similar distributions of x and $p_\perp$ values of the hard interaction, and then also the MPI effects are closely the same.) Differential distributions of the diffractive events are affected, since the kinematics of the $\mathbb{P}$p system is set up using the Pomeron flux parametrizations. A subset of these distributions is shown in Fig. 4, for the some of the same combinations as in Table 1. As expected, $\mathbb{P}$ PDF variations do not have a large impact on the shapes (cf. Fig. 4), while the $\mathbb{P}$ flux gives rise to large effects in $x_{\mathbb{P}}$, hence on the broadening of the mass spectrum and on the tails of the t and θ distributions. In view of these observations, we do not expect to be able to discrimate between the available Pomeron PDFs when comparing to data. Thus we will leave out this variation from now on, and focus on variations in the Pomeron flux.

The diffractive event fraction is not process-independent. One reason is that processes may be dominated by different initial states, another that different x and Q^2 scales are probed. In Table 2 we show the fraction of events passing either selection for various hard processes available in PYTHIA 8 using the SaS flux and the H1 Fit B LO PDF. Firstly we note that a smaller fraction of events pass the PDF selection than in Table 1, owing to the larger x needed to produce these particles, cf. Fig. 1. This is why top, being the heaviest, has the smallest diffractive fraction. In addition there is a notable difference between the gluon-dominated Higgs production and the quark-induced production of $W^{\pm}/\gamma^*/Z^0$, owing to the hard gluon PDF in the $\mathbb{P}$. If top production is considered separately for $q\bar{q} \to t\bar{t}$ and $gg \to t\bar{t}$, the PDF survival rate is (9.74 ± 0.09)% and (10.55 ± 0.10)%, respectively, displaying the difference between the two production channels.
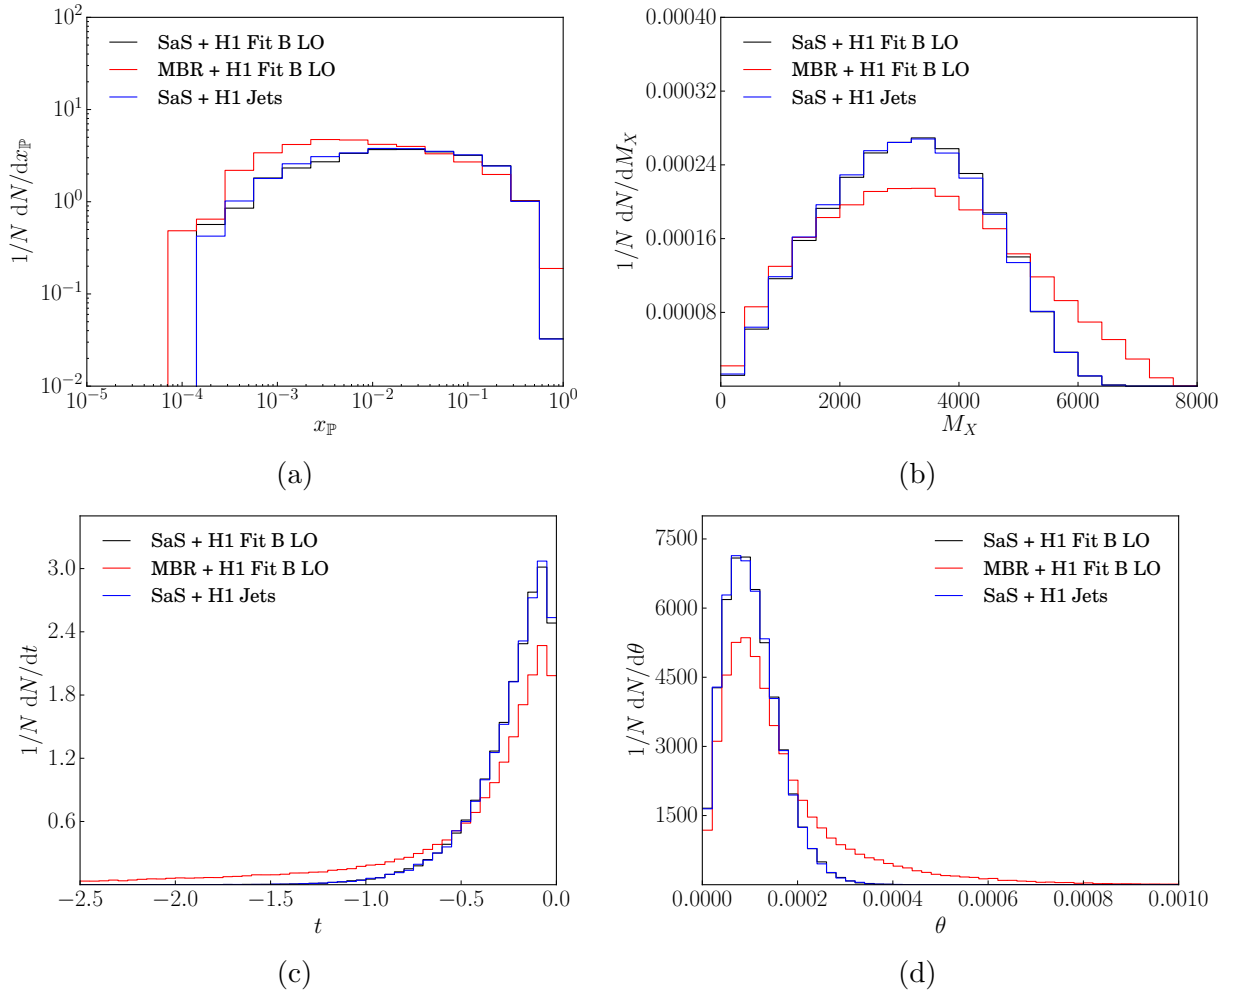
Figure 4: Some kinematics distributions obtained with variations of the Pomeron parametrizations: (a) $x_{\mathbb{P}}$, (b) M_X, (c) t and (d) θ.

Diffractive fractions pp collisions at $\sqrt{s} = 8$ TeV		
	PDF selection	MPI selection
$q\bar{q} \to W^{\pm}$	(11.16 ± 0.10) %	(0.70 ± 0.03) %
$q\bar{q} \to \gamma^*/Z^0$	(10.69 ± 0.10) %	(0.76 ± 0.03) %
Single top and top pair production	$(\ 8.51 \pm 0.09)$ %	(0.62 ± 0.02) %
SM Higgs production	(12.37 ± 0.10) %	(0.86 ± 0.03) %

Table 2: Diffractive fractions obtained with PYTHIA without any phasespace cuts at $\sqrt{s} = 8$ TeV for various hard processes. PYTHIA is run with the SaS flux and the H1 Fit B LO PDF.
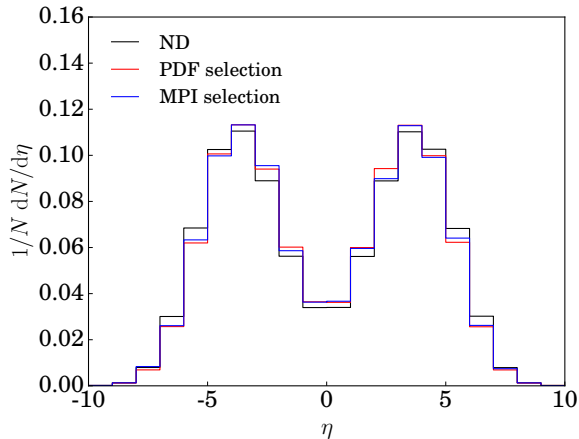
12

Figure 5: The rapidity of the W-boson produced in $q\overline{q} \to W^{\pm}$ at $\sqrt{s} = 8$ TeV.

In Fig. 5 we show the rapidity of the W-boson produced in the process $q\overline{q} \to W^{\pm}$ at an 8 TeV pp collision, comparing three samples; nondiffractive, PDF selected and MPI selected. It is observed that the diffractive W's are slighly more central than the nondiffractive in the CM frame, as expected from Fig. 3. The differences are small, however, being on the order of (5-10)%, and might reduce when phase-space cuts are introduced. We will study this process further in Sec. 4.1.

3.2   The dynamical gap survival and MPI models

In the above section we studied how the parametrization of the Pomeron flux and PDF affected the diffractive fractions and distributions, and notably by the choice of $\mathbb{P}$ flux. By contrast, we saw that the survival fraction in the MPI selection step was not significantly affected by these choices. A dependence does enter both via the x and the $p_{\perp}$ distributions of a process: larger x scales leaves less energy for MPIs and thereby gives a higher MPI survival probability, whereas larger $p_{\perp}$ values gives a longer MPI evolution range and thereby a lower MPI survival. Such effects are not too prominent, however, and tend to be overshadowed by the sensitivity to the parameters of the MPI model. These enter twice. Firstly, for the MPI selection, since the dynamical gap survival is tied to the number of MPIs in the pp system. Secondly, for the properties of the diffractive system, where the number of MPIs affects e.g. charged multiplicities.

The probability for obtaining MPIs is given by eq. (7), and hence depends on both the overlap function and the regulator $p_{\perp 0}^{\text{ref}}$. The related parameters are primarily tuned to minimum bias and underlying event data, e.g. charged particle pseudorapidity $dn/d\eta$, multiplicity $P(n)$ and transverse momenta $dn/dp_{\perp}$ and $\langle p_{\perp} \rangle (n_{\text{ch}})$ spectra of charged particles. This means that a change of MPI parameters for the diffractive studies would spoil agreement with nondiffractive data. Nevertheless, it is interesting to study how the survival rate changes with these parameters for the pp collision itself.

The MPI modelling of the $\mathbb{P}$p collision can be decoupled from that of the pp one. Then the MPI survival rate would not be affected by changes, but only the particle distributions in the diffractive system. One inevitable free parameter is the effective $\mathbb{P}$p total cross section. It is currently set always to be 10 mb, but could be made to depend on the mass of the diffractive system. Also the relative normalization of $\mathbb{P}$ flux and PDFs can influence the event activity. We will study the normalization dependence in the last part of this section.

To begin with, consider the impact-parameter picture associated with hard collisions in our
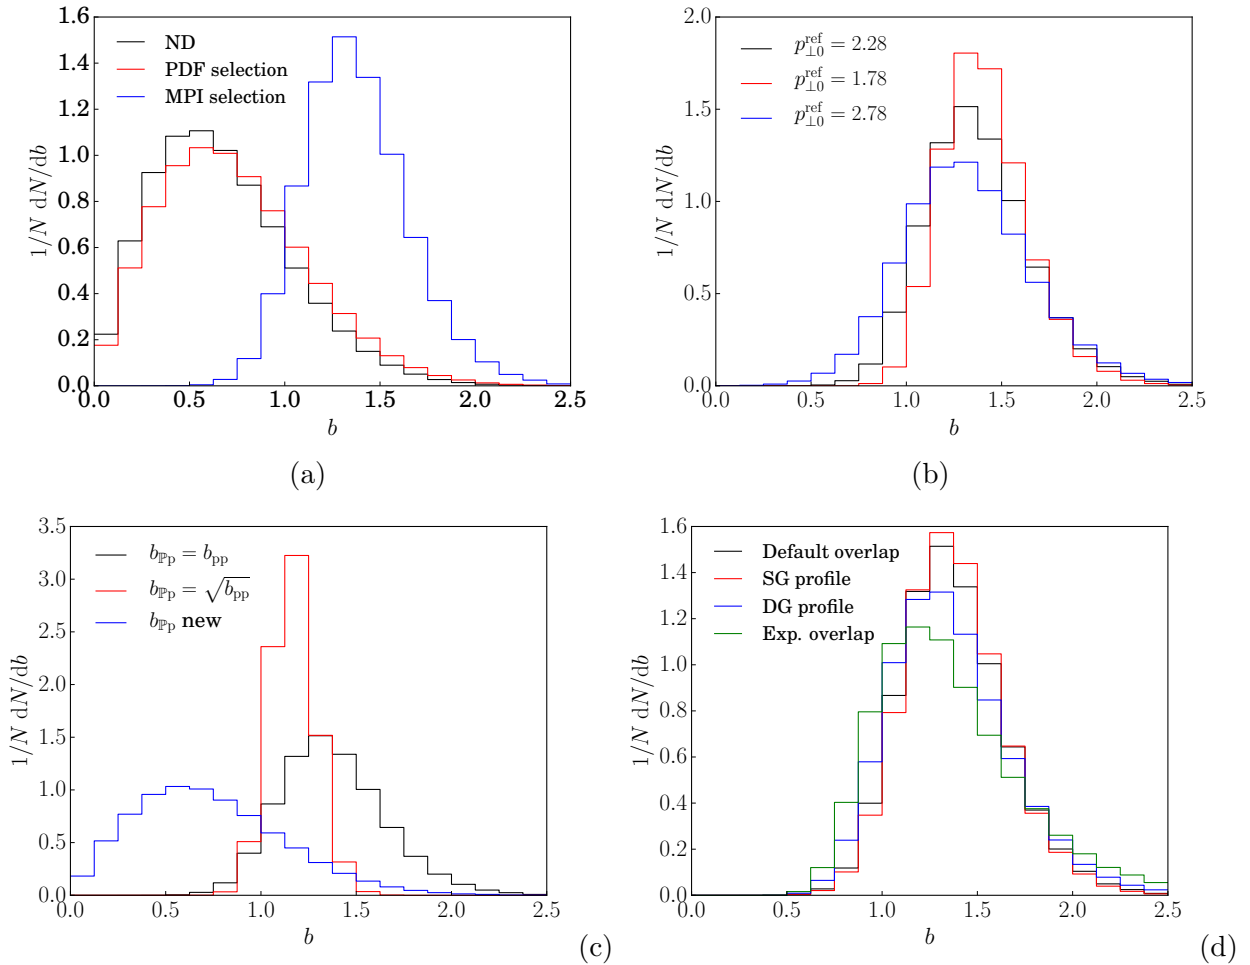
Figure 6: Impact-parameter distribution of $2 \to 2$ QCD processes with $p_\perp > 20$ GeV in $\sqrt{s} = 8$ TeV pp collisions. (a) The change during the selection steps. (b) The dependence on $p_{\perp 0}^{\mathrm{ref}}$. (c) The distribution in the $\mathbb{P}$p subcollision. (d) The dependence on impact-parameter profile.

model, Fig. 6a. The b scale is normalized such that $\langle b \rangle = 1$ for inclusive minimum-bias events. Events with a hard interaction tend to be more central than that, since central events have more MPIs in general and thereby a bigger likelihood that at least one of them is at large $p_\perp$. The PDF selection step does not have a significant impact, but the MPI one kills most low-b events and pushes $\langle b \rangle$ above unity. The reason is obvious: for central events the average number of MPIs is high, and thus the likelihood of only having the trigger hard process and no further MPIs is small, while more peripheral collisions give fewer MPIs and thereby a higher surviving fraction. Ultimately, when $\langle n_{\mathrm{MPI}}(b) \rangle \ll 1$, most protons pass by each other without colliding at all. Thus the interesting region for diffraction is where $\langle n_{\mathrm{MPI}}(b) \rangle \sim 1$.

The $p_{\perp 0}^{\mathrm{ref}}$ regulator is by default 2.28 GeV. Since an increase in this parameter gives less MPI in the pp system, we expect an increase in the diffractive fractions, and vice versa. Table 3 confirms that this is indeed the case: variations of ± 0.5 GeV around the default $p_{\perp 0}^{\mathrm{ref}}$ value gives about a factor of two in the MPI selection rate. This major $p_{\perp 0}^{\mathrm{ref}}$ dependence holds also for many other nondiffractive event properties, however; keeping everything else fixed even a variation of ± 0.1 GeV would be unacceptable. In Fig. 7 we show the charged multiplicity distribution, when we change the regulator $p_{\perp 0}^{\mathrm{ref}}$ for both diffractive and nondiffractive events, with minor/major effects for the former/latter. The stability in the diffractive case is because a change in the regulator also

Diffractive fractions pp collisions at $\sqrt{s} = 8$ TeV		
	PDF selection	MPI selection
$p_{\perp 0} = 1.78$	(14.50 ± 0.11) %	(0.39 ± 0.02) %
$p_{\perp 0} = 2.28$	(14.33 ± 0.11) %	(0.98 ± 0.03) %
$p_{\perp 0} = 2.78$	(14.19 ± 0.11) %	(2.00 ± 0.04) %

Table 3: Diffractive fractions for the $2 \to 2$ QCD processes with $p_\perp > 20$ GeV in $\sqrt{s} = 8$ TeV pp collisions. PYTHIA is run with the SaS flux and the H1 Fit B LO PDF.
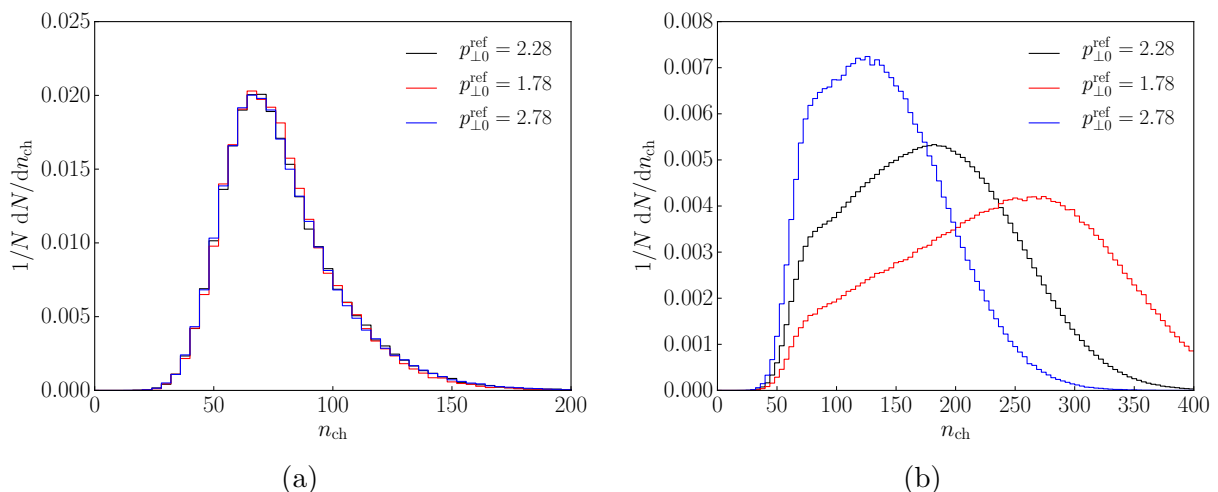

Figure 7: Charged multiplicity distributions in the (a) $\mathbb{P}$p subsystem for diffractive events, (b) pp system for nondiffractive events, in $2 \to 2$ QCD processes with $p_\perp > 20$ GeV as before.

affects the impact parameter picture. Specifically, in this case $b_{\mathbb{P}p} = b_{pp}$ has been assumed. A lower value of the regulator, giving rise to a larger number of MPIs in the pp system, pushes $\langle b_{pp} \rangle$ to larger values for those events that survive the diffractive MPI criterion, Fig. 6b. More precisely, the change is to b values where the average pp MPI activity is restored to its original level. With $b_{\mathbb{P}p} = b_{pp}$ the same then holds when MPI activity is generated in the diffractive system, such that the effects of a smaller regulator and a larger impact parameter almost completely cancel.

As we have already discussed, the modelling of the $\mathbb{P}$ size could also affect the MPI machinery for the $\mathbb{P}$p subcollision via the impact parameter $b_{\mathbb{P}p}$. The currently implemented three alternatives are compared in Fig. 6c. The maybe less realistic last option of picking a new $b_{\mathbb{P}p}$ value at random implies a significant fraction of events with small $b_{\mathbb{P}p}$ and thereby the possibility of many MPIs. The average $\langle n_{\mathrm{MPI}} \rangle$ for the three options is 1.66, 2.04 and 4.09, respectively, thus giving rise to 0.66, 1.04 and 3.09 additional MPIs besides the hardest process. This is reflected notably in the charged multiplicity distribution, Fig. 8a.

The MPI survival rate is highly dependent on the proton matter profile, Table 4 and Fig. 6d. Diffraction thrives when $\langle n_{\mathrm{MPI}}(b) \rangle \sim 1$, so this b region should be as broad as possible for a large diffractive rate. Conversely, a sharp proton edge implies less diffraction. The default overlap function $\exp(-b^{1.85})$ is close to a Gaussian, and the two have about the same MPI selection rate. The double Gaussian and the exponential overlap are examples of broader distributions, thus with more diffraction, whereas the option without any b dependence represents the other extreme (not shown in Fig. 6d), with less diffraction. Overall the variation is not so dramatic, however, if only experimentally acceptable variations are considered.
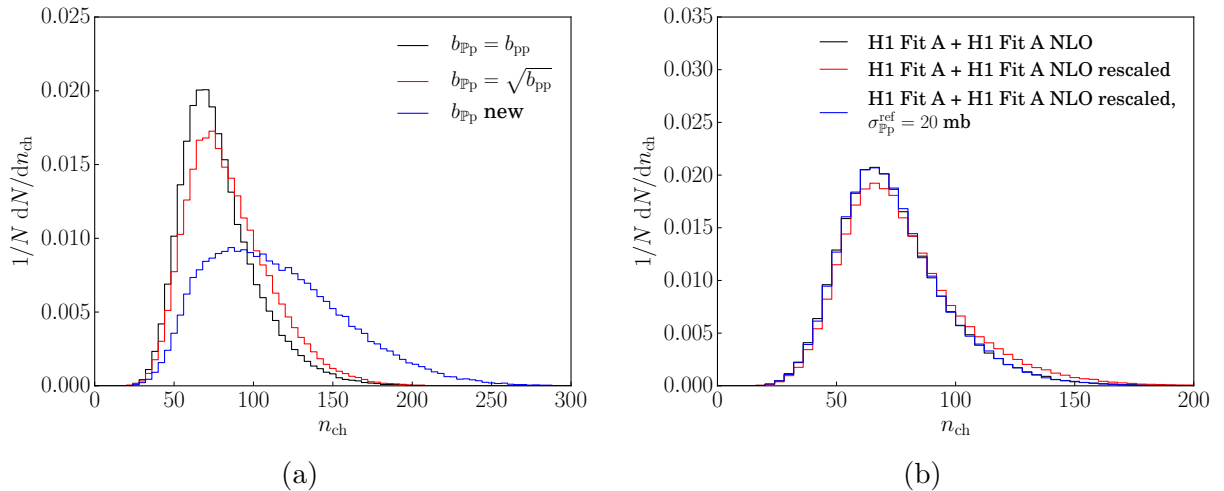
15

Figure 8: Charged multiplicity distribution distributions for the $\mathbb{P}$p diffractive subsystem, for events with $2 \to 2$ QCD processes with $p_\perp > 20$ GeV as before. (a) For three different $b_{\mathbb{P}p}$ impact-parameter profiles. (b) With or without rescaled $\mathbb{P}$ flux and PDFs, see text.

Diffractive fractions		
pp collisions at $\sqrt{s} = 8$ TeV		
	PDF selection	MPI selection
No impact parameter dependence	(14.36 ± 0.11) %	(0.38 ± 0.02) %
Single gaussian matter profile	(14.25 ± 0.11) %	(0.93 ± 0.03) %
Double gaussian matter profile	(14.24 ± 0.11) %	(1.04 ± 0.03) %
Default overlap	(14.33 ± 0.11) %	(0.98 ± 0.03) %
Exponential overlap	(14.50 ± 0.11) %	(1.28 ± 0.04) %

Table 4: Diffractive fractions for the $2 \to 2$ QCD processes with $p_\perp > 20$ GeV in $\sqrt{s} = 8$ TeV pp collisions. PYTHIA is run with the SaS flux and the H1 Fit B LO PDF.

Finally we turn to the relative normalization of the $\mathbb{P}$ PDF and flux. From eq. (3) we know that the PDF selection step depends on the convolution of the $\mathbb{P}$ flux and PDFs. Thus it has no net effect if the flux is scaled down by a factor of two and the PDFs are scaled up by the same amount, so as to bring the H1 PDFs to be approximately normalized to unit momentum sum. It does have consequences for the MPI selection step, however, since the average MPI rate comes up in the $\mathbb{P}$p system.

Compared with the (1.35 ± 0.04) % MPI selection rate in Table 1 for the H1 Fit A flux+PDF combination, such a rescaling changes the rate to (1.40 ± 0.04) %, ie. no effects are seen. The rescaling however, does change the multiplicity distribution, Fig. 8b, as a consequence of the increased $d\sigma_{\mathrm{MPI}}$ in eq. (7). This could be compensated by a corresponding increase of σ_{ref} from the default 10 mb to 20 mb, thereby restoring both the MPI selection rate and the multiplicity distribution, cf. the blue line in Fig. 8b.

3.3 Energy and scale dependence

Here we study the model dependence on the scales in the hard process and the energy of the collision.

In Fig. 9 the diffractive fractions are compared at different collision energies, $\sqrt{s}$, for $2 \to 2$

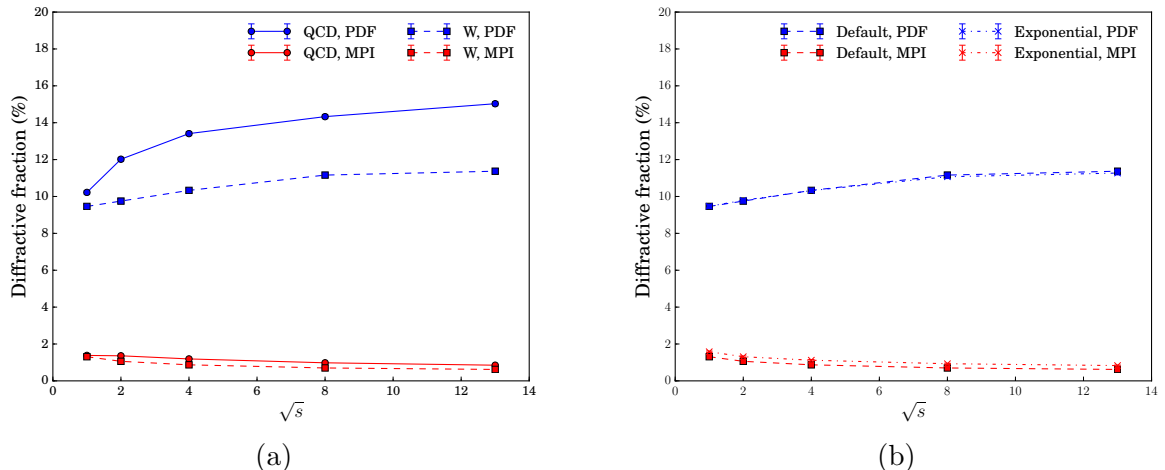Figure 9: (a) The diffractive fractions obtained in $2 \to 2$ QCD processes with $p_\perp > 20$ GeV (circles and solid lines) and $q\overline{q} \to W^\pm$ (squares and dashed lines) in pp collisions at different energies. (b) The diffractive fractions obtained in $q\overline{q} \to W^\pm$ with the default overlap function (squares and dashed lines) and the exponential overlap function (crosses and dashed-dotted lines). PYTHIA is run with the SaS flux and the H1 Fit B LO PDF.

Diffractive fractions pp collisions at $\sqrt{s} = 8$ TeV		
	PDF selection	MPI selection
$M_W = 50$ GeV	(11.52 ± 0.10) %	(0.72 ± 0.03) %
$M_W = 80.385$ GeV	(10.69 ± 0.10) %	(0.70 ± 0.03) %
$M_W = 150$ GeV	(10.46 ± 0.10) %	(0.72 ± 0.03) %
$M_W = 500$ GeV	$(\ 9.47 \pm 0.09)$ %	(0.65 ± 0.03) %

Table 5: Diffractive fractions for the process $q\overline{q} \to Z^0$ in $\sqrt{s} = 8$ TeV pp collisions. PYTHIA is run with the SaS flux and the H1 Fit B LO PDF.

QCD processes with $p_\perp > 20$ GeV, and for $W^\pm$ production. In the PDF selection step the diffractive rate increases with energy. The difference between the two processes indicates that this rise can depend on the incoming flavours and the relevant ranges of x values. Depending on the $\mathbb{P}$ flux and PDF, such as a freezing of the latter at small x, the fraction might even decrease with energy.

A larger collision energies also implies a higher average number of MPIs, in addition to the hardest collision, thus implying a reduced fraction of events passing the MPI criterion, see Fig. 9. There is a compensatory effect of diffraction shifting to larger impact parameters, as already discussed for the $p_{\perp 0}^{\mathrm{ref}}$ variations. For the close-to-Gaussian default overlap the relative size of the $\langle n_{\mathrm{MPI}} \rangle \approx 1$ region decreases with energy, however, resulting in the trend shown. By comparison an exponential overlap decreases slower than the close-to-Gaussian, hence resulting in less suppression with increasing energy.

Finally, Table 5 shows the number of events passing the PDF and MPI selections when the mass of the produced particle is changed. In the PDF selection step heavier particles are less likely to be produced diffractively, as they require larger x-values, where the probability for diffraction is lower (cf. Fig. 3). The same trend was observed in Table 2, but was there mixed up by the use of different production channels. After the MPI selection step the mass dependence is not as
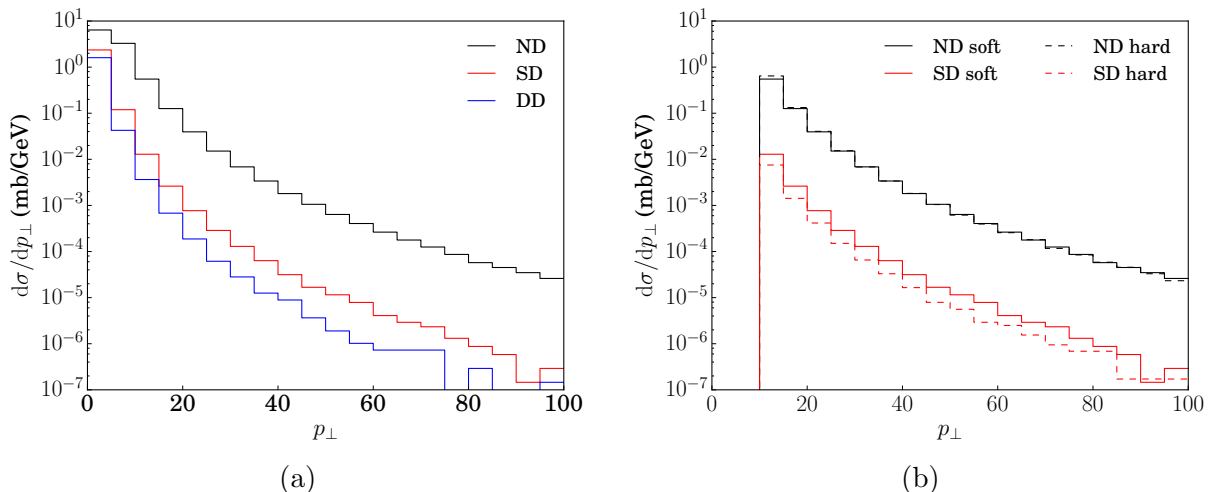
Figure 10: The $p_\perp$ of the hardest process obtained with (a) the soft (or inclusive) diffraction framework, and (b) both the soft and hard diffraction frameworks for events with $p_\perp > 10$ GeV.

Cross sections		
pp collisions at $\sqrt{s} = 8$ TeV		
	Soft diffraction	Hard diffraction
ND sample, $p_\perp > 10$ GeV (mb)	3.730	4.239
ND sample, $p_\perp > 20$ GeV (mb)	0.348	0.353
SD sample, $p_\perp > 10$ GeV (mb)	0.084	0.048
SD sample, $p_\perp > 20$ GeV (mb)	0.0066	0.0035

Table 6: Cross sections obtained with the two diffractive frameworks. Extracted from Fig. 10 by integration.

clearly visible. A partial compensation can indeed occur, since a higher subcollision mass implies more momentum taken out of the incoming protons and thereby less left for subsequent collisions.

3.4 Comparison with soft diffraction

The new model for hard diffraction complements the existing one for soft (or rather inclusive) diffraction. The latter already has a hard component arising from the MPI model, which is used to pick the hardest process and all subsequent scatterings in the $\mathbb{P}p$ system, except for low-mass diffractive systems where no perturbative framework can be applied. The soft diffractive model only allows for $2 \to 2$ QCD processes, unlike the new hard one, but for QCD processes a comparison between the two is meaningful. To this end, the $p_\perp$ of the hardest process in an event will be used. This is not a physically measurable observable, unlike e.g. the closely related $p_\perp$ of the hardest jet in an event, but for the relative comparison of hard and soft diffraction it is cleaner.

The MPI framework predominantly gives low-$p_\perp$ interactions, be it for diffractive or nondiffractive events. Thus only a small fraction of the events will have $p_\perp$ values at ~ 10 GeV or more, see Fig. 10a. Note that the $p_\perp$ spectrum falls faster for diffractive than nondiffractive events, mainly as a consequence of the former having a $\mathbb{P}p$ invariant mass spectrum peaked towards lower values.

In Fig. 10b the $p_\perp$ of the hardest process for the two samples is compared. One is obtained by generating inclusive (soft) events and keeping only those with large enough $p_\perp$, the other by

generating only hard events above 10 GeV. Nondiffractive events are shown as a sanity check, as for them the two approaches should give the same results. A closer look at integrated cross sections, Table 6, shows a small discrepancy for the $p_\perp > 10$ GeV case, while the $p_\perp > 20$ GeV agree much better. This discrepancy is caused by not having a "Sudakov factor" in the hard model. That is, in the soft model the rate at lower $p_\perp$ scales is reduced by the requirement of not having an interaction at a higher $p_\perp$ scale, whereas no such reduction is implemented in the hard framework, which only uses pure matrix elements.

The single diffractive events show differences in the normalisation, while the shape of the $p_\perp$ distributions agree between the two frameworks. The normalisation differences arise from the two different ways of handling the survival rate. The soft diffractive framework assumes an effective flux of $\mathbb{P}$'s inside the proton, rescaled to get the desired total diffractive cross section, and thereby implicitly includes an average rapidity gap survival factor. The hard diffractive framework has a higher initial $\mathbb{P}$ flux but then explicitly implements a dynamical event-by-event survival factor. As it works out, single diffractive high-$p_\perp$ events are somewhat more suppressed in the latter case. This is indeed what we would expect: there should be more MPIs in high-$p_\perp$ (and high-mass) events than in low-$p_\perp$ ones, and thus more MPI survival suppression. Put another way, the soft implementation overestimates the suppression at low $p_\perp$ and underestimates it at high $p_\perp$. (Assuming our new model is the right way to view the matter.)

In the future it would be desirable to include such dynamical effects also in the soft framework, so that the two descriptions can be made to agree in the high-$p_\perp$ region. This is not a trivial task, however.

4 Comparisons with data

In this section we compare the new model for hard diffraction with some available data. While many results have been presented for soft diffractive processes, less is available on hard ones.

At the Tevatron, both the CDF and D0 collaborations studied hard diffractive events. We have chosen here to compare with two analyses, one in which only the diffractive fractions are measured, the other in which also the distributions of the hard collisions are reported.

At the LHC, diffraction has been studied both by ATLAS [37, 38, 39] and CMS [40, 41, 42]. One key observation there is that the PYTHIA default $\mathbb{P}$ flux shape does not describe the rapidity gap distribution so well, suggesting that a new parametrization may be needed. In other respects the model seems to do a reasonable job. For hard diffraction we will compare to the latest ATLAS study, [39], and a similar CMS study, [41].

Unfortunately, neither of the studies at hand are implemented as Rivet [43] analyses, so we have tried to apply the relevant experimental cuts as best as we can. This makes comparisons with data less than reliable, and results should therefore be taken as a first indication only. At least for LHC the intention is that the new PYTHIA options can be directly tested by the experimental community, to allow more precise comparisons in the future.

4.1 Diffractive W/Z production at the Tevatron

CDF has measured the fraction of events with a diffractively produced W/Z boson at $\sqrt{s} = 1.96$ TeV [44]. The surviving antiproton was measured in a Roman Pot forward spectrometer, and the boson decay products in the central detector. The observed fraction of events with forward antiprotons was doubled, to compensate for there being no Roman Pots on the proton side. Only the e and μ leptonic decays of the bosons were taken into account. The cuts used in the analysis are listed in Table 7, along with the number of events that survive after each step. To this end, the internal W- and Z-finder projections available in Rivet [43] have been used as a starting point;

CDF cuts	W sample			Z sample		
	ND	SD	$\frac{\text{SD}\times 2}{\text{ND}}$ (%)	ND	SD	$\frac{\text{SD}\times 2}{\text{ND}}$ (%)
Lepton $E_T^e(p_T^\mu) > 25$ GeV	670602	2827	0.84	667851	2466	0.74
Missing $E_T > 25$ GeV	595236	2490	0.84	-	-	-
One electron in $|\eta| < 2.8$	-	-	-	642250	2366	0.74
One lepton in $|\eta| < 1.1$	331316	1374	0.83	366566	1397	0.76
$M_T^W = [40, 120]$ GeV	327671	1361	0.83	-	-	-
$M_Z = [66,116]$ GeV	-	-	-	36814	1397	0.76
$|t| < 1$ GeV^2	-	1348	0.82	-	1383	0.75
$x_\mathbb{P} = [0.03,0.1]$	-	366	0.23	-	346	0.19

Table 7: Cuts used in [44]. Number of events listed in each of the samples are based on Monte Carlo truth obtained when generating 10^6 inclusive events. A blank means that a specific cut was not relevant.

$\mathbb{P}$ PDF $\mathbb{P}$ flux	$(\text{p}\overline{\text{p}} \to \overline{\text{p}}' + \text{W}) \times 2$	$(\text{p}\overline{\text{p}} \to \overline{\text{p}}' + \text{Z}) \times 2$
CDF	(1.0 ± 0.11) %	(0.88 ± 0.22) %
H1 Fit B LO SaS	(0.19 ± 0.03) %	(0.24 ± 0.04) %
H1 Fit B LO MBR	(0.29 ± 0.04) %	(0.20 ± 0.03) %
H1 Jets SaS	(0.29 ± 0.04) %	(0.24 ± 0.04) %
H1 Fit A NLO H1 Fit A	(0.46 ± 0.05) %	(0.35 ± 0.04) %
H1 Fit B LO H1 Fit A	(0.44 ± 0.05) %	(0.29 ± 0.04) %

Table 8: Diffractive fractions for the W $\to l\nu$ and Z $\to l^+l^-$, $l = e, \mu$ in $\sqrt{s} = 1.96$ TeV p$\overline{\text{p}}$ collisions.

these have previously been validated for other CDF analyses. In addition the diffractive properties are derived from the measured antiproton as

$$t = -p_\perp^2 \tag{9}$$

$$x_\mathbb{P}^{\text{RPS}} = 1 - \frac{2|p_z|}{\sqrt{s}} \tag{10}$$

which has been compared to Monte Carlo truth, giving good agreement.

The results in Table 7 are obtained with PYTHIA 8 using the SaS flux and the H1 Fit B LO PDF, starting out from an inclusive MPI-selected sample. We note that a large fraction of the diffractive events do not pass the experimental $x_\mathbb{P}$ cut. Therefore, although we begin with a "Monte Carlo truth" fraction of $\sim 1\%$ diffractive W/Z, this is reduced to $\sim 0.2\%$ by the $x_\mathbb{P}$ cut. Results look better for other choices of $\mathbb{P}$ flux, see Table 8, but even at best still with a factor two discrepancy. Note that it is the fluxes that rise fastest in the low-$x_\mathbb{P}$ region that gives fractions closer to data.

The diffractive fraction can also be increased by changing the free parameters of the MPI framework, with the caveat that nondiffractive events will then no longer describe data as well.

Parameter	$(p\overline{p} \to \overline{p}' + W) \times 2$	$(p\overline{p} \to \overline{p}' + Z) \times 2$
CDF	(1.0 ± 0.11) %	(0.88 ± 0.22) %
$p_{\perp 0}^{\text{ref}} = 2.78$ GeV	(0.59 ± 0.06) %	(0.49 ± 0.05) %
Exponential overlap	(0.25 ± 0.04) %	(0.24 ± 0.04) %

Table 9: Diffractive fractions for the $W \to l\nu$ and $Z \to l^+l^-$, $l = e, \mu$ in $\sqrt{s} = 1.96$ TeV $p\overline{p}$ collisions.

CDF cuts	
Jet $E_T^{1,2}$	> 7 GeV
Jet E_T^3	> 5 GeV
Jet $|\eta^{1,2,3}|$	< 4.2
ΔR	0.7
$|t|$	< 1 GeV^2
$x_{\mathbb{P}}^{\text{RPS}}$	$[0.035, 0.095]$

Table 10: Cuts used in [21].

Table 9 shows the diffractive fractions obtained when varying some of the MPI parameters. This variation is still not sufficient when combined with the default flux and PDF in PYTHIA 8. If combined with some of the fluxes in Table 8 it would be possible to obtain fractions close to the experimentally observed values, however.

4.2 Diffractive dijets at the Tevatron

Another interesting measurement performed at CDF was the process $\overline{p}p \to \overline{p} + X_p$, $X_p \to X + J + J$, ie. SD dijet production with a leading antiproton. CDF measured this at three different energies, $\sqrt{s} = 630$, 1800 and 1960 GeV [45, 21, 46]. Here not only the diffractive fractions were measured, but a number of differential distributions. Large discrepancies were found between the diffractive structure functions determined from CDF data and those extracted by the H1 Collaboration from diffractive deep inelastic scattering data at HERA. The discrepancies are both in normalisation and shape and were interpreted as a breakdown of factorization.

Our comparison focuses on the 1800 GeV data ([21]), since this also includes a measurement of the diffractive structure function. The cuts used in the analysis are listed in Table 10. The jets are identified with the CDF cone algorithm as implemented in Rivet [43], with a cone radius of 0.7. Jet energy scale corrections for underying-event activity are done separately for diffractive and nondiffractive events, as outlined in the CDF article, but only has a minor impact on relative rates. The momentum transfer of the antiproton is evaluated using eq. (9) and the momentum loss of the antiproton using eq. (10).

We begin by evaluating the suppression factor introduced by the MPI framework. This is evaluated by running two samples of 10^6 events, one with and one without the MPI criterion, both using the cuts of Table 10 and the SaS flux and the H1 Fit B LO PDF. We obtain a suppression factor of 0.11, to be compared with the quoted discrepancies from CDF of $0.06\pm0.02\,(0.05\pm0.02)$ when using the H1 Fit 2 (Fit 3), respectively [21]. A similar suppression factor as for SaS is obtained when using the H1 Fit B flux, based on the same parametrization as the H1 Fit 2 and 3 fluxes, although with different values for the free parameters of the model. Using this flux, however, allows for approximately two times more events passing the experimental cuts. This is due to the fact that the H1 Fit B flux is less restrictive in the low-$x_{\mathbb{P}}$ region, where the experiment is performed. Hence we expect better agreement with data when using the H1 Fit B flux, as

compared to SaS.

Results on kinematical distributions using both the SaS and the H1 Fit B flux are shown in Fig. 11. The SD E_T^* distribution has a steeper falloff than the ND distribution, indicating a lower center-of-mass energy in the collision. Likewise the η^* distribution is shifted towards positive η, the proton direction, indicating a boost of the center-of-mass system. The final kinematical distribution here is the difference in azimuthal angle between the two leading jets. This observable was not shown in the 1800 GeV analysis but in the 1960 GeV one. The SD events there show a tendency to be more back-to-back than the ND ones. This can also be attributed to the lower energy in the $\mathbb{P}$p collision than in the full p$\overline{p}$ system, leaving less space for initial-state radiation.

The momentum fraction of the antiproton carried by the subcollision parton can be evaluated from the jets using

$$x = \frac{1}{\sqrt{s}} \sum_{i=1}^{3} E_T^i e^{-\eta^i}, \tag{11}$$

where the sum is over the two leading jets, plus a third if it has $E_T > 5$ GeV. The result is shown in Fig. 12, for the two $\mathbb{P}$ fluxes used in Fig. 11. As expected the SaS flux, Fig. 12a, suppress the diffractive events too much, as the suppression factor is too large compared to experimental value from CDF. With this flux, the PDF selected samples lie above the CDF data, but then drop by an order of magnitude by the MPI selection, to lie well below the data, by a factor of five. There is also some discrepancy in shape. Changing to the H1 Fit B flux, Fig. 12b, the PDF selected sample lies above the data as expected, with the MPI selected sample a bit below, although only by a factor of three. The suppression is still too large, and shapes still disagree, but not as markedly as in Fig. 12a.

There are some aspects of the CDF article that we don't understand, however. The key Fig. 4 of [21] is intended to show the H1 predictions for the diffractive structure function along with the experimentally measured one. The information provided on how the former prediction is obtained is inconsistent with the curve shown, however, in normalization and shape. In the end we therefore put more faith in the suppression factor between CDF and HERA, already presented above, than in absolute numbers. Assuming we could have reproduced the CDF curve intended to represent the predictions of the H1 PDFs, that then is suppressed by an average multiplicative factor of $0.05 - 0.06$ in data but 0.11 in our model, we should have been a factor of ~ 2 above data, which is inconsistent with the outcome in Fig. 12.

4.3 CMS diffractive contribution to dijet production

CMS has studied the diffractive contribution to dijet events at $\sqrt{s} = 7$ TeV pp collisions [41], The cross section is presented as a function of $\widetilde{\xi}$, an approximation to the fractional momentum loss of the scattered proton correspinding to the $x_{\mathbb{P}}$ variable. Dijets were selected with $p_\perp > 20$ GeV in the $|\eta| < 4.4$ range using the anti-$k_\perp$ algorithm with a cone size of $R = 0.5$ [47]. $\widetilde{\xi}$ was reconstructed using particles in the region $|\eta| < 2.4$ with $p_\perp > 0.2$ GeV for charged particles as well as particles in the range $3.0 < |\eta| < 4.9$ with $E > 4$ GeV. To enhance the diffractive contribution additional requirements was introduced, such that the minimum rapidity gap was of 1.9 units (no particles was allowed in the region $|\eta| > 3$). Finally a cut on $\widetilde{\xi} < 0.01$ was introduced.

With these cuts, rapidity gap survival probabilities are in the range 0.08 ± 0.04 (NLO) to 0.12 ± 0.05 (LO), where the NLO gap survival probability was found using POMPYT and POWHEG[48]+PYTHIA 8 and the LO gap survival probability was found using POMPYT and POMWIG.

Implementing the same cuts in PYTHIA 8, using the SaS flux and the H1 Fit B LO PDF gives a rapidity gap survival probability of 0.06, compatible with the CMS results. Changing from the
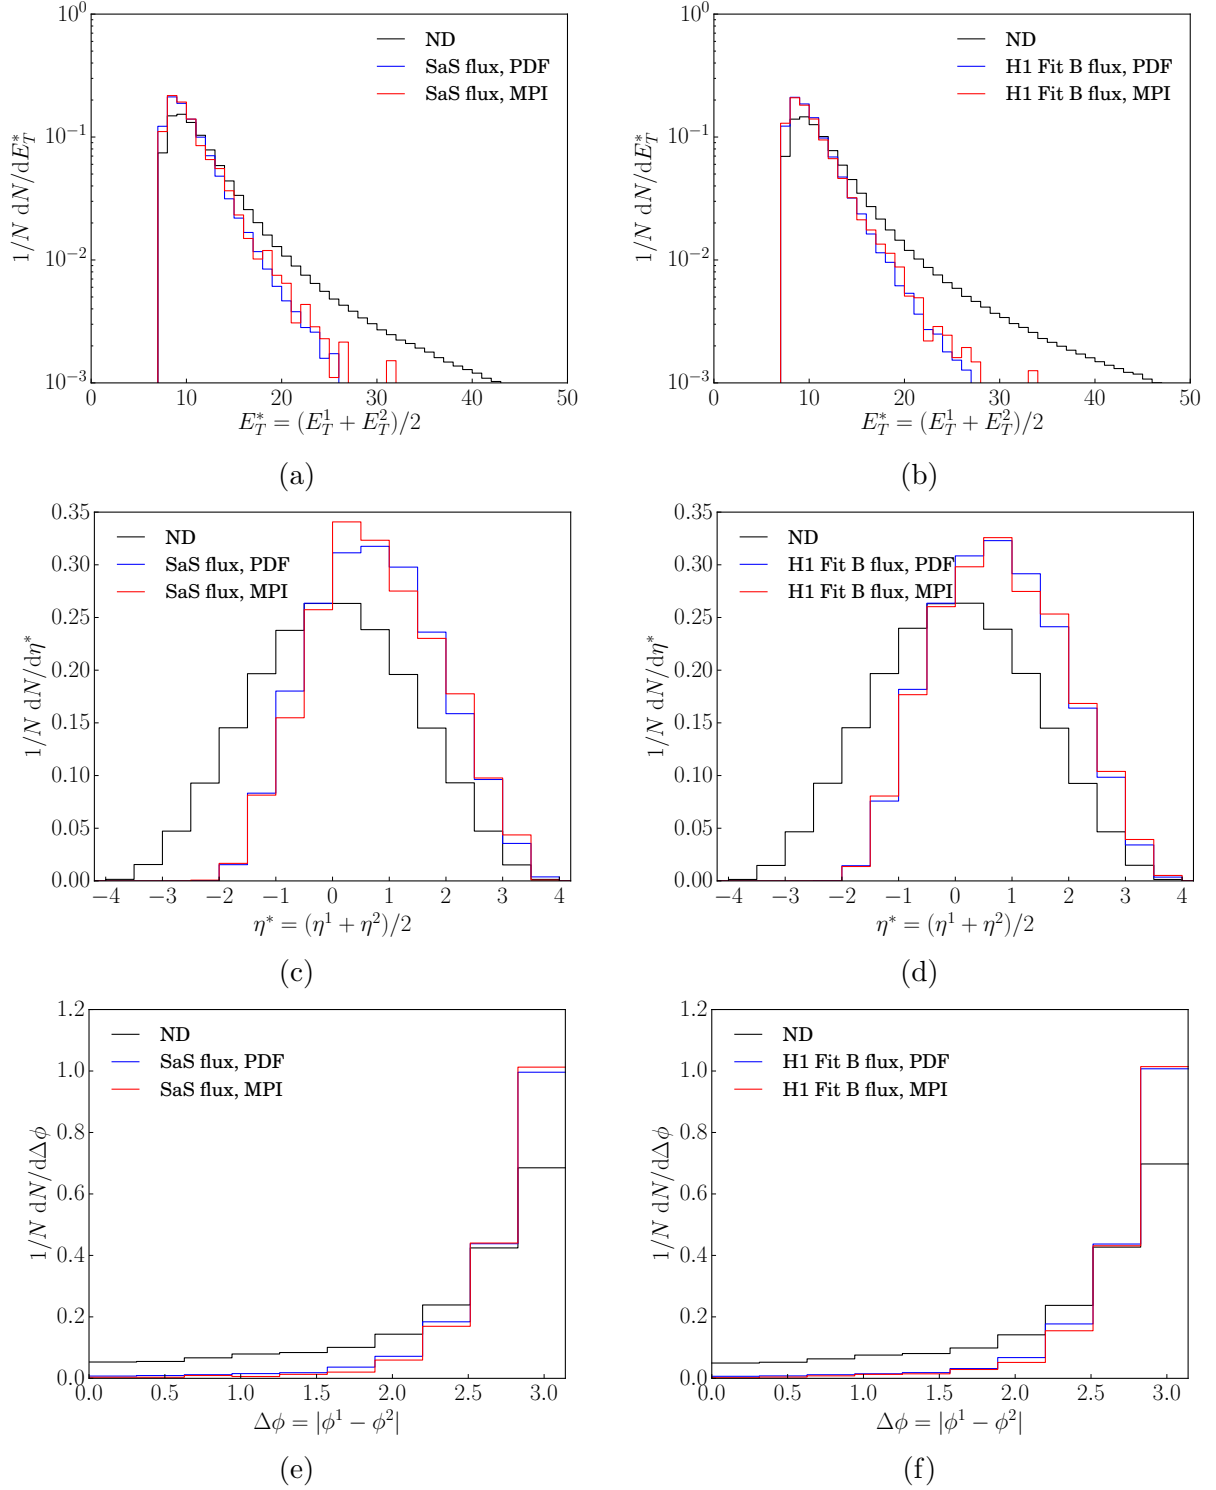
Figure 11: The mean E_T of the leading jets in both SD and ND events using (a) the SaS and (b) the H1 Fit B flux. The mean η of the leading jets in both SD and ND events using (c) the SaS and (d) the H1 Fit B flux. The mean difference in ϕ between the leading jets in both SD and ND events using (e) the SaS and (f) the H1 Fit B flux.
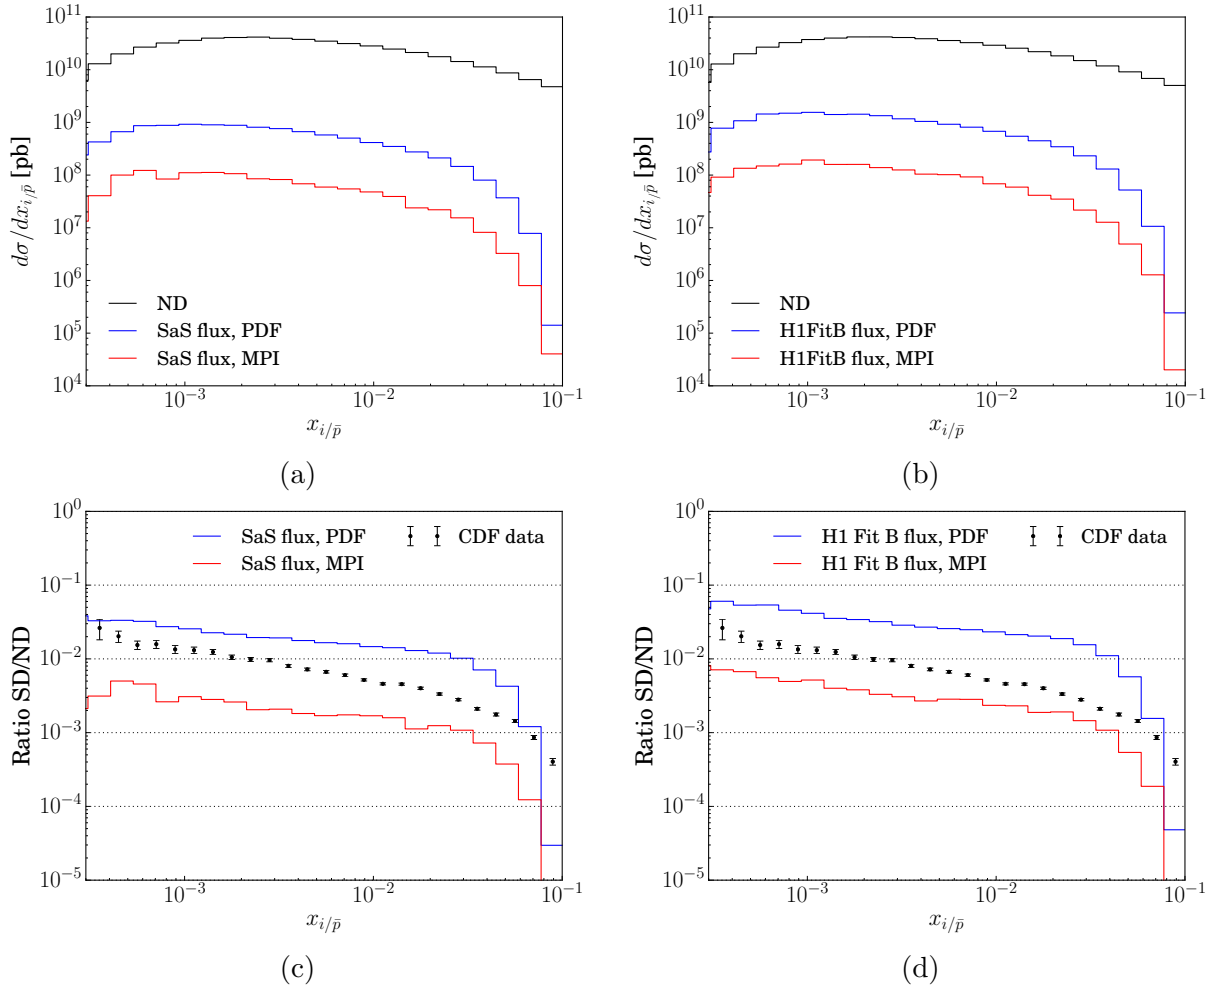
Figure 12: The antiproton momentum fraction carried by the parton entering the hard collision, for PYTHIA 8 compared with CDF data. PYTHIA is run with the H1 Fit B LO PDF and (a) the SaS or (b) H1 Fit B flux. (c) and (d) shows the ratio SD to ND using (a) and (b).

SaS flux to the H1 Fit B flux gives the same suppression factor, but allows for more events to pass the experimental cuts. We thus see the same trend as in the CDF analysis, where the SaS flux is too restrictive at low $x_{\mathbb{P}}$.

4.4 ATLAS dijets with large rapidity gaps

Recently, the ATLAS collaboration published a study of dijets with large rapidity gaps in $\sqrt{s} = 7$ TeV pp collisions [39]. Dijets were selected with $p_\perp > 20$ GeV in the $|\eta| < 4.4$ range, and the cross section was measured in terms of $\Delta\eta^F$, the size of the observed rapidity gap, as well as in $\widetilde{\xi} = \sum p_\perp^i e^{\pm\eta^i}/\sqrt{s}$, the estimate of the fractional momentum loss deduced from charged and neutral particles in the ATLAS detector (the sign on η depends on where in the detector the largest gap is located). Cuts used in the analysis are listed in Table 11.

Experimental results were compared with the PYTHIA 8 soft diffractive framework, which predicts both the ND, SD and DD contributions to the dijet production. Three different flux models were compared: SaS, Donnachie-Landshoff and MBR. All three predict cross sections in the range of the data, without any need for additional gap survival probability factors. The POMWIG generator [13], on the other hand, needed an additional suppression of $S^2 = 0.16 \pm$

24

Jet cuts	
Jet $E_T^{1,2}$	> 20 GeV
Jet $|\eta^{1,2}|$	< 4.4
Anti-$k_\perp$ ΔR	0.6
Neutral particles	
$|p|$	> 200 MeV
$|\eta|$	< 4.8
Charged particles	
$|p|$	> 500 MeV or
$p_\perp$	> 200 MeV
$|\eta|$	< 4.8

Table 11: Cuts used in [39].

$0.04(\text{stat}) \pm 0.08(\text{sys})$ in order to describe data.

In this section we use the new model for hard diffraction to study the same cross sections. The new model currently only includes the SD contribution, hence we will not be able to describe all aspects of data, especially in the high-$\Delta\eta^F$ and low-$\widetilde{\xi}$-regions, where the SD and DD contributions are comparable in size, at least according to the soft diffraction model available in PYTHIA 8. We could also expect the normalisation of the SD events obtained with the hard diffraction framework to be lower than in the soft one and thus in data, because of the difference in normalisation between the two frameworks (cf. Sec. 3.4). The ND contribution should not differ from the ATLAS analysis, however, since no changes have been implemented in this framework.

The ND distribution was normalized to data, where the normalization factor was found using the first bin of the $\Delta\eta^F$ distribution. This approach has also been used in our analysis, although when generating an inclusive sample (e.g. the purple distribution in Figs. 13b and 13d) this normalization is applied to the full sample, unlike in the ATLAS paper. In this sample, no classification of events occurs, hence the normalization cannot be performed only on the ND sample. In the exclusive samples, the distinction between ND and SD is performed, and we can apply the normalization to only the ND sample (cf. the black distribution in Figs. 13b and 13d).

In Fig. 13 we show the results obtained with the model for hard diffraction. Three samples are compared: ND, PDF-selected SD and MPI-selected SD. Note that the MPI-selected sample lies about a factor of 10 below the PDF-selected one, as usual, and that the suppression due to the MPI-framework is constant over both intervals. The new model undershoots the data in the regions where the DD contribution is non-negligible ($\Delta\eta^F > 1$ and $\log_{10}\widetilde{\xi} < -0.5$). When this contribution is included in the framework, a better agreement with data should be possible, and overall the picture should be consistent with the soft diffractive framework.

5 Summary and outlook

In this article we have studied hard diffraction by combining two concepts, the Ingelman–Schlein picture of a Pomeron and the PYTHIA model for multiparton interactions. The Pomeron fluxes and PDFs are mainly extracted from HERA data, while the MPI picture (and several other relevant physics components) makes use of a broader spectrum of Tevatron and LHC data. This combination allows us, in principle, to predict all physical quantities of hard diffractive events, from rapidity gap sizes to charged multiplicity distributions, but most importantly the fraction of diffractive events for any hard process.
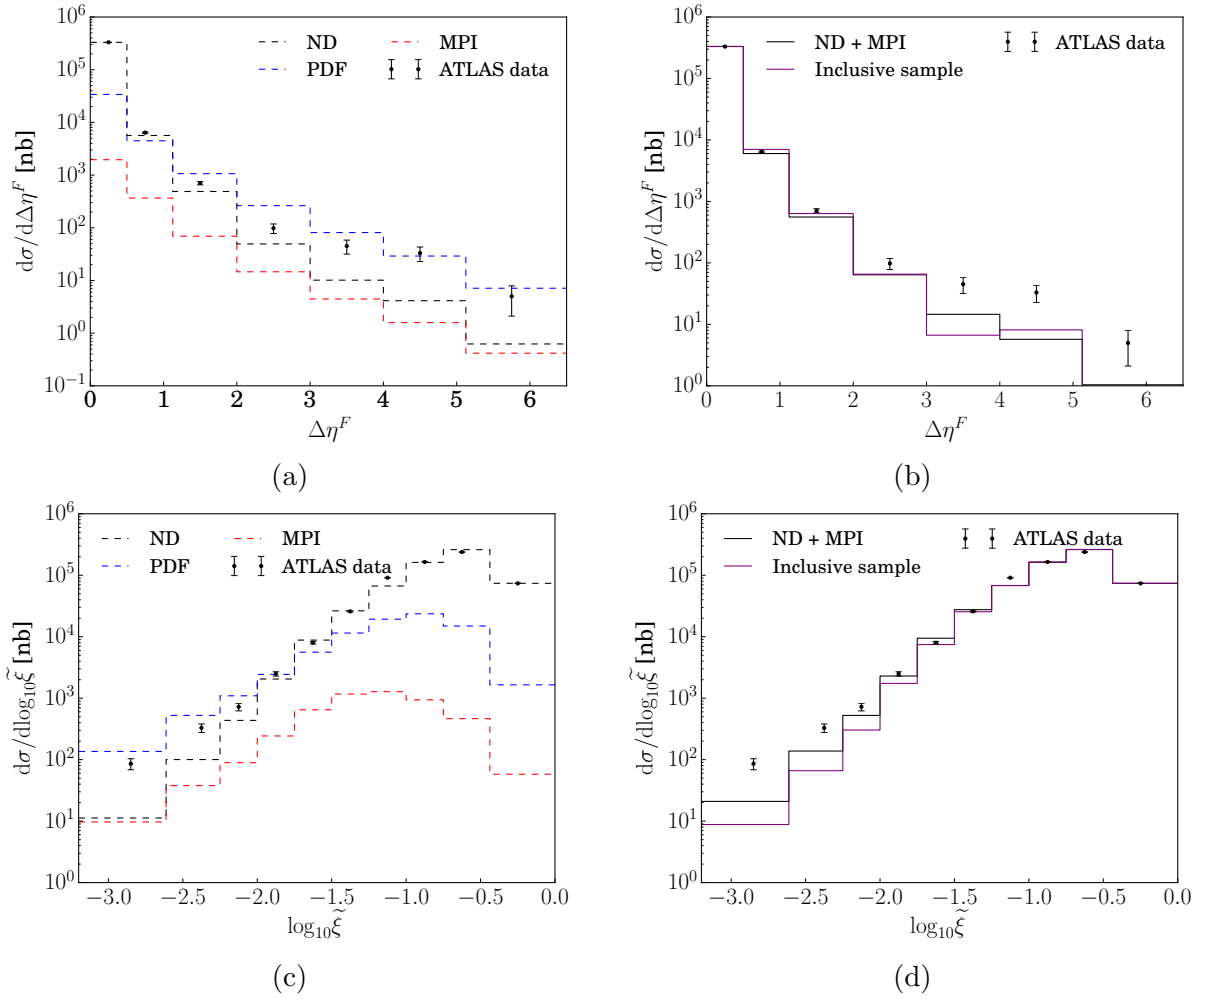
Figure 13: The dijet cross sections as a function of the size of the rapidity gap (a), (b) and the fractional momentum loss of the proton (c), (d). Compared to the hard diffraction model of PYTHIA 8 using the SaS flux and H1 Fit B LO PDF. In (b) and (d) the ND + MPI sample is a sum of the black and red dotted lines from (a) and (c), whereas the inclusive sample are generated directly with PYTHIA 8. Only statistical errors are included in the ATLAS errorbars.

Reality is not quite as simple, however. In this article we have studied the different assumptions that go into a detailed framework, and explored the inherent uncertainties. One part concerns the assumed Pomeron flux and PDFs, where particularly the latter is dominated by one source only, namely the H1 analyses, making it difficult to assess to full range of uncertainty. Another part concerns the MPI framework, which enters twice. When used the first time, to determine the diffractive MPI survival, it involves parameters already tuned to nondiffractive data, so narrowly constrained in principle. There could still be leeway, e.g. if we were to use other parton showers that give less/more activity at small $p_\perp$ scales, the average number of MPIs would have to rise/drop to compensate. Thus our studies focus on the sensitivity of some key parameters of the framework. When the MPIs are used the second time, inside the diffractive subsystem itself, the level of uncertainty is considerably higher. A key example is the impact-parameter picture of the $\mathbb{P}$p subcollision, notably how impact parameters are related between the pp and $\mathbb{P}$p steps of an event.

Our studies puts the finger on our still limited understanding of diffraction, also when restricted to the Pomeron framework, which is only one model class for diffraction. Further, we provide

computer code that can be used to compare with data for hard diffractive processes at the LHC. It thus can be used as a "straw man" model, where differences between predictions and data can help pave the way for a deeper understanding and more accurate models. Specifically, with a generator it is possible both to emulate the experimental diffractive trigger and to compare the resulting event properties, both of which are considerably more complicated for analytical models.

Comparisons with data have shown qualitative agreements in many respects, but maybe less so than one could have hoped for. For the Tevatron we face the problem of trying to understand 15 years old analyses, with uncertain results. The main message probably is that the overall Tevatron suppression factor of $\sim 10 - 20$, relative to HERA-based extrapolations, agrees well with what our model gives from the MPI selection step. For the future it will therefore be more interesting to compare with LHC studies, in particular those available in Rivet.

It is well known that the existing PYTHIA model for soft diffraction is not fully describing the existing LHC data; at places the difference can be up to a factor of two. Similarly we have seen less-than-perfect agreement for the hard diffractive processes studied in this article. There is therefore room for improvements in both areas, and also for work to bring the two approaches in closer contact. As one simple example, the soft model currently does not involve a MPI survival step, and therefore the Pomeron flux does not have to be normalized in the same way in the two cases. The intention is to study such issues closer, and to provide an improved description of diffractive cross sections, both integrated and differential ones.

Acknowledgements

References

[1] M. L. Good and W. D. Walker, Phys. Rev. **120** (1960) 1857.

[2] M. L. Perl, "High Energy Hadronic Physics" (Wiley-Interscience, 1974, ISBN 0-471-68049-4).

[3] J. R. Forshaw and D. A. Ross, "Quantum Chromodynamics and the Pomeron," (Cambridge University Press, 1997, ISBN 0 521 56880 3).

[4] H. Jung, hep-ph/9809374.

[5] V. Barone and E. Predazzi, "High-Energy Particle Diffraction" (Springer, 2002, ISBN 3-540-42107-6).

[6] S. Donnachie, H. G. Dosch, O. Nachtmann and P. Landshoff, "Pomeron Physics and QCD" (Cambridge University Press, 2002, ISBN 0 521 78039 X).

[7] H. Jung, R. B. Peschanski and C. Royon, Acta Phys. Polon. B **33** (2002) 3645 [hep-ph/0209143].

[8] M. G. Albrow, T. D. Coughlin and J. R. Forshaw, Prog. Part. Nucl. Phys. **65** (2010) 149 [arXiv:1006.1289 [hep-ph]].

[9] G. Ingelman and P. E. Schlein, Phys. Lett. B **152** (1985) 256.

[10] R. Bonino *et al.* [UA8 Collaboration], Phys. Lett. B **211** (1988) 239.

[11] P. Bruni and G. Ingelman, Conf. Proc. C **930722** (1993) 595.

[12] T. Sjöstrand, Comput. Phys. Commun. **82** (1994) 74.

[13] B. E. Cox and J. R. Forshaw, Comput. Phys. Commun. **144** (2002) 104 [hep-ph/0010303].

[14] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour and B. R. Webber, JHEP **0101** (2001) 010 [hep-ph/0011363].

[15] T. Sjöstrand and M. van Zijl, Phys. Rev. D **36** (1987) 2019.

[16] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, Phys. Rept. **97** (1983) 31.

[17] W. Buchmuller and A. Hebecker, Phys. Lett. B **355** (1995) 573 doi:10.1016/0370-2693(95)00721-V [hep-ph/9504374].

[18] A. Edin, G. Ingelman and J. Rathsman, Phys. Lett. B **366** (1996) 371 doi:10.1016/0370-2693(95)01391-1 [hep-ph/9508386].

[19] Y. L. Dokshitzer, V. A. Khoze and T. Sjöstrand, Phys. Lett. B **274** (1992) 116.

[20] J. D. Bjorken, Phys. Rev. D **47** (1993) 101.

[21] T. Affolder *et al.* [CDF Collaboration], Phys. Rev. Lett. **84** (2000) 5043.

[22] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna and S. Prestel *et al.*, Comput. Phys. Commun. **191** (2015) 159 [arXiv:1410.3012 [hep-ph]].

[23] F. D. Aaron *et al.* [H1 and ZEUS Collaborations], Eur. Phys. J. C **72** (2012) 2175 [arXiv:1207.4864 [hep-ex]].

[24] G. A. Schuler and T. Sjöstrand, Phys. Rev. D **49** (1994) 2257.

[25] P. Bruni and G. Ingelman, Phys. Lett. B **311** (1993) 317.

[26] H. Jung, Comput. Phys. Commun. **86** (1995) 147.

[27] A. Donnachie and P. V. Landshoff, Nucl. Phys. B **244** (1984) 322.

[28] R. Ciesielski and K. Goulianos, PoS ICHEP **2012** (2013) 301 [arXiv:1205.1446 [hep-ph]].

[29] A. Aktas *et al.* [H1 Collaboration], Eur. Phys. J. C **48** (2006) 715 [hep-ex/0606004].

[30] A. Aktas *et al.* [H1 Collaboration], Eur. Phys. J. C **48** (2006) 749 [hep-ex/0606003].

[31] A. Aktas *et al.* [H1 Collaboration], JHEP **0710** (2007) 042 [arXiv:0708.3217 [hep-ex]].

[32] A. Donnachie and P. V. Landshoff, Drell-Yan Pairs," Nucl. Phys. B **303** (1988) 634. doi:10.1016/0550-3213(88)90423-3

[33] A. D. Martin, M. G. Ryskin and G. Watt, Phys. Lett. B **644** (2007) 131 doi:10.1016/j.physletb.2006.11.032 [hep-ph/0609273].

[34] M. R. Whalley, D. Bourilkov and R. C. Group, hep-ph/0508110.

[35] A. Buckley, J. Ferrando, S. Lloyd, K. Nordstöm, B. Page, M. Rüfenacht, M. Schönherr and G. Watt, Eur. Phys. J. C **75** (2015) 3, 132 [arXiv:1412.7420 [hep-ph]].

[36] R. D. Ball *et al.*, Nucl. Phys. B **867** (2013) 244 [arXiv:1207.1303 [hep-ph]].

[37] G. Aad *et al.* [ATLAS Collaboration], Nature Commun. **2** (2011) 463 doi:10.1038/ncomms1472 [arXiv:1104.0326 [hep-ex]].

[38] G. Aad *et al.* [ATLAS Collaboration], Eur. Phys. J. C **72** (2012) 1926 doi:10.1140/epjc/s10052-012-1926-0 [arXiv:1201.2808 [hep-ex]].

[39] G. Aad *et al.* [ATLAS Collaboration], arXiv:1511.00502 [hep-ex].

[40] S. Chatrchyan *et al.* [CMS Collaboration], Eur. Phys. J. C **72** (2012) 1839 doi:10.1140/epjc/s10052-011-1839-3 [arXiv:1110.0181 [hep-ex]].

[41] S. Chatrchyan *et al.* [CMS Collaboration], Phys. Rev. D **87** (2013) 1, 012006 doi:10.1103/PhysRevD.87.012006 [arXiv:1209.1805 [hep-ex]].

[42] V. Khachatryan *et al.* [CMS Collaboration], Phys. Rev. D **92** (2015) 1, 012003 doi:10.1103/PhysRevD.92.012003 [arXiv:1503.08689 [hep-ex]].

[43] A. Buckley, J. Butterworth, L. Lönnblad, D. Grellscheid, H. Hoeth, J. Monk, H. Schulz and F. Siegert, Comput. Phys. Commun. **184** (2013) 2803 [arXiv:1003.0694 [hep-ph]].

[44] T. Aaltonen *et al.* [CDF Collaboration], Phys. Rev. D **82** (2010) 112004 [arXiv:1007.5048 [hep-ex]].

[45] D. Acosta *et al.* [CDF Collaboration], Phys. Rev. Lett. **88** (2002) 151802 doi:10.1103/PhysRevLett.88.151802 [hep-ex/0109025].

[46] T. Aaltonen *et al.* [CDF Collaboration], Phys. Rev. D **86** (2012) 032009 doi:10.1103/PhysRevD.86.032009 [arXiv:1206.3955 [hep-ex]].

[47] M. Cacciari, G. P. Salam and G. Soyez, JHEP **0804** (2008) 063 doi:10.1088/1126-6708/2008/04/063 [arXiv:0802.1189 [hep-ph]].

[48] S. Alioli, K. Hamilton, P. Nason, C. Oleari and E. Re, JHEP **1104** (2011) 081 doi:10.1007/JHEP04(2011)081 [arXiv:1012.3380 [hep-ph]].