

SOFT QCD IN ATLAS: MEASUREMENTS AND MODELLING OF MULTI-PARTON INTERACTIONS*

ANDY BUCKLEY

PPE Group, School of Physics and Astronomy, University of Edinburgh
Mayfield Road, Edinburgh EH9 3JZ, United Kingdom

(Received November 7, 2011)

Soft QCD contributes to all observables at the LHC, due to the presences of underlying event (UE) and pile-up in all events. Both these processes are dominated by multi-parton interactions (MPI), *i.e.* the result of proton collisions containing more than one partonic interaction due to collective and beam remnant effects. While there is undoubtedly interesting physics involved in MPI, the primary interest of LHC experiments is to characterise and model the behaviour of UE and pile-up sufficiently well that their influence may be cleanly subtracted in the process of searching for new physics signatures at 7 TeV and beyond. I summarise the soft QCD measurements made by ATLAS using the 2010 and early 2011 datasets, and the use of this data to improve Monte Carlo generator models of MPI for use in forthcoming simulation campaigns.

DOI:10.5506/APhysPolB.42.2669

PACS numbers: 12.38.Qk, 13.75.Cs, 13.85.Lg, 13.85.Hd

1. Introduction

A consequence of doing physics at a hadron collider is that one has to understand the incoming hadrons rather well. A multitude of LHC “new physics” processes are illustrated by diagrams in which a pair of partons neatly extract themselves from their parent protons without consequence, the only inconvenience being the loss of longitudinal momentum information thanks to the probabilistic nature of parton distributions; however, real life is not so clean, in particular because the rest of the proton constituents and the beam remnants left after parton extraction cannot be so easily ignored. Including such effects leads to a model in which multiple partonic interactions may occur in each event, and where the influence of the colour charge flows associated with those multiple scatterings and beam remnants can also be substantial.

* Presented at the LI Cracow School of Theoretical Physics “Soft Side of the LHC”, Zakopane, Poland, June 11–19, 2011.

Of course, the LHC was not built to run at 7 TeV (and eventually 14 TeV) to provide greater insight into the soft structure of protons — although that will be a welcome consequence of its existence. The main reason for LHC collaborations to be interested in multi-parton interactions is that they form troublesome backgrounds in the core LHC task of searching for signatures of new physics. These backgrounds occur in two forms: first, the additional parton interactions in hard processes adds activity to that event. This additional activity is, as a consequence of relative cross-sections, predominantly QCD based and can change the energies and distributions of QCD jets, add new jets, and fake electron signatures: this is the “underlying event” (UE).

The second way in which MPI can affect hard-process physics searches is via pile-up: the overlay of multiple pp interactions in a single bunch-crossing. Again as a consequence of relative cross-sections, pile-up events are overwhelmingly dominated by soft QCD scattering (minimum bias), and are typically modelled as pure MPI scatterings with no “hard” process. As LHC luminosities increase, the Poisson mean number of pile-up interactions per bunch crossing also increases, from $\mu \sim 0.1$ in the earliest LHC runs, to $\mu \sim 10$ in early 2011, $\mu \sim 30$ at the end of the 2011 run, and eventually μ values of several hundred in the LHC luminosity upgrade scenario. A typical p_T density contribution of 1 GeV per unit in η - ϕ for each pile-up event can significantly change event characteristics at high μ . Unlike UE, pile-up can be directly reduced by use of track-to-vertex matching; however, this is far from a panacea, as the reduction of tracking efficiency and increase in overlapping primary vertices at high occupancy, and the need to account for uncharged pile-up particles and to use calorimeter elements out of the tracking acceptance means that pile-up must be understood well enough that subtraction may be attempted.

So both pile-up and UE at the LHC require a level of understanding of multiple interactions in proton–proton collisions that was not established at the Tevatron. In the remainder of this contribution I summarise the current state of MPI modelling in Monte Carlo event generators and the degrees of freedom in these models, present the ATLAS QCD measurements which are most useful in constraining their free parameters, and show the resulting MC tunings which will be used in the next year of ATLAS (and other) physics studies.

2. Monte Carlo modelling of soft QCD

MC event generators are a crucial tool for experimental particle physics: despite occasional protestations to the contrary, it is sufficiently difficult to cleanly disentangle contributing processes in an experimental analysis that some degree of reliance on simulation — of Standard Model processes at least — is virtually inevitable. Even “data-driven” background estimations often use the data just to fix the normalisation of simulated distributions.

2.1. General-purpose MC event generators

Event generators come in several forms, from parton-level codes which calculate only total or perhaps differential cross-sections, to “general-purpose” codes in which the partons are hadronised, hadrons are decayed, and MPI effects are simulated. The real power of general-purpose generators from an experimentalist’s point of view is that they do not just produce asymptotic distributions, but that they typically use sampling to produce physical-looking simulated events. Fully exclusive simulation of this kind is key to the design of detectors and analyses, and the *unfolding* of detector effects on measured observables [1].

In addition to MPI simulation, general-purpose generators enhance the matrix element and phase space sampling of parton-level generators by using (matrix element matched) perturbative parton showers to stochastically approximate full QCD resummation, and by use of non-perturbative hadronisation models and hadron decays to produce realistic particle kinematics and identified multiplicities. As MPI and hadronisation are modelled phenomenologically rather than from first principles, the bulk of event generator free parameters are concentrated in these modelling areas. The process of parameter optimisation by comparison to experimental data is known as “tuning”.

2.2. MPI modelling in general-purpose generators

We already mentioned in the introduction an unrealistic (but popular!) simplified model of hadron collisions in which only a pair of asymptotically free partons from the two incoming protons interact in hard scattering. In this model, the only theoretical concession to proton structure is that the incoming parton energies are given by parton density functions (PDFs).

A slightly more realistic model is to treat protons as bags of non-interacting partons, of which more than one pair can interact in a given collision. In fact, structure function data from the HERA collider [2] was the historical catalyst for adoption of such a model, as the strong rise of the F_2 structure function at low x would without application of unitarization procedures lead to the the partonic jet cross-section could exceed the measured total interaction cross-section for sufficiently low minimal jet $p_{\perp}$ s, as shown in figure 1. The steep rise in partonic jet cross-section is driven by interactions at low values of momentum exchange (below a few GeV), where both the strong coupling and the PDFs diverge. This apparent contradiction — how can the partonic cross-section exceed top-down analyticity constraints like the Froissart–Martin bound? — is due to the neglected bulk interactions of the incoming hadrons. Hence these proton bulk effects must also be included in any model of hadron collisions which wishes to be infra-red complete.




Fig. 1. Comparison of cross-section predictions, with the dashed and dot-dashed (black) lines indicating the Donnachie–Landshoff 1992 and 2004 total pp cross-section parameterisations, constrained by analyticity arguments, and the steeper solid (blue) line showing the partonic dijet cross-section above 2 GeV. The energy at which the partonic cross-section exceeds the total cross-section reduces with the partonic $p_{\perp}$ cut.

Such a model is the Sjöstrand–van Zijl model implemented in the PYTHIA generator in 1987 [3]. This model defines the standard template of MPI modelling adopted by general-purpose MC generators since that time: the excess of partonic cross-section is interpreted as the mean number of partonic interactions in a hadron collision at that energy, *i.e.* $\bar{n} = \hat{\sigma}_{\text{jet}}/\sigma_{\text{inel}}$. An eikonal formalism is then applied to generate Poisson-distributed numbers of multi-parton interactions (MPI) from this $\bar{n}$, with use of a pp transverse impact parameter and nucleon form factor. The effect of this latter feature is to reproduce the experimental jet pedestal effect, where the level of MPI activity approximately plateaus as a function of the scale of the hardest scattering in the event: in the eikonal model, this corresponds to an increasing overlap of proton form factors until the sampled impact parameter $b = 0$ with high probability and all collisions are fully overlapping (“central”, in the terminology of double parton scattering and heavy ion physics).

Since the divergence of the partonic cross-section is driven by IR divergences in the matrix elements and PDFs, and is regularised by the Poisson distribution of number of interactions, an MPI model requires a mechanism to suppress the partonic divergence. In the first PYTHIA model, and in the JIMMY MPI model developed as a similar extension to the HERWIG event generator [4], this is achieved with a simple cutoff on the scattering p_T , denoted as $p_{\perp}^{\text{min}}$. In later models such as the current versions of PYTHIA6

and Pythia 8, a smoother regularisation of the divergence is used, with the matrix element $1/p_T^2$ divergence replaced with

$$\frac{1}{p_T^2} \rightarrow \frac{p_T^2}{(p_T^2 + (p_{\perp}^{\min})^2)^2}. \quad (1)$$

This *ad hoc* form is not theoretically motivated and represents an IR continuation of perturbative QCD scattering into the regime where the strong coupling diverges. Not all MC generators use this form: JIMMY retains the sharp cut-off regularisation, while Herwig++ attempts a more theory-driven continuation: the optical theorem is used to relate soft inelastic scatterings to the slope of the elastic scattering cross-section, with experimental input via CDF data and the Donnachie–Landshoff (DL) total pp cross-section parameterisations [5, 6, 7]. In all cases, a higher value of $p_{\perp}^{\min}$ introduces more screening of the soft divergence and hence results in less MPI activity than a lower value.

Another major feature of the PYTHIA MPI model is that $p_{\perp}^{\min}$ evolves as a function of the centre-of-mass collider energy $\sqrt{s}$. The form of this evolution is again not robustly predicted by QCD theory, but a Regge-inspired ansatz has long been adopted, in which $p_{\perp}^{\min}$ evolves with a power law in s , similarly to the Pomeron term in the DL total cross-section parameterisation. The specific form used in the PYTHIA generators is

$$p_{\perp}^{\min}(\sqrt{s}) = p_{\perp}^{\min}(1800 \text{ GeV}) \cdot \left(\frac{\sqrt{s}}{1800 \text{ GeV}} \right)^{e/2}. \quad (2)$$

Here the tuning parameters are $p_{\perp}^{\min}(1800 \text{ GeV})$, and the exponent, e , whose value in the Donnachie–Landshoff fit would be 0.16, but which favours a higher value of $e \sim 0.25$ in pre-LHC MPI tunes [8]. A higher value of e means that $p_{\perp}^{\min}$ will be higher at LHC energies, and hence LHC MPI activity will be reduced. The reference energy is set to 1800 GeV simply because the first fits were derived as deviations from the Tevatron Run I data at that energy: a different energy could be used, *e.g.* at 7 TeV for LHC-driven tunes, but would solely result in an algebraic transformation of the parameters which would make comparison with old tunes more difficult¹. Again, not all generators follow the PYTHIA example: particularly in the HERWIG generator family, the JIMMY MPI model makes no prediction for energy extrapolation and the more evolved Herwig++ model explicitly attempts to fit multi-energy data with a single value of $p_{\perp}^{\min}$. The latest tunes of Herwig++, however, have decided that the most minimal form of this model is insufficient to describe all data and have hence introduced an energy evolution parameterisation of their own.

¹ Exactly this has occurred with the PYTHIA6 “Perugia2011” tune set.

The remaining common features of Monte Carlo MPI models which are of relevance to tuning to LHC and other data are the proton form factor and the oft-mentioned “colour reconnection” mechanism. The first of these is crucial to description of the transition between “minimum bias” physics (*i.e.* the bulk of hadron collider events in which low multiplicity, low $p_\perp$ inelastic scattering dominates) and underlying event physics (in which the MPI interactions are in addition to a hard partonic scattering). The PYTHIA family, always keen to provide a wide range of phenomenological handles, offers a range of form factor parameterisation options, from single- and double-Gaussian parameterisations of the form factor itself (with tweakable relative widths and populations in the double-Gaussian case) to a general $O(b) \propto \exp(-b^n)$ form for the *overlap* function $O(b)$. The HERWIG family, by comparison, fixes the form factor shape to the Fourier transform of the proton electromagnetic form factor,

$$G(\vec{b}) = \int \frac{d^2\vec{k}}{(2\pi)^2} \frac{e^{i\vec{k}\cdot\vec{b}}}{(1 + k^2/\mu^2)^2}, \quad (3)$$

where μ^2 is an inverse radius scale-factor introduced to account for the possible difference in distribution of electric and colour charge: as in PYTHIA this width parameter is considered free. In recent versions of Herwig++ and Pythia 8, a refinement of this scheme has been introduced in which the density of the matter distribution is dependent on the momentum fraction x of the hardest scattering: this “hot spot” model is both supported by data [9, 10] and theoretically motivated [11, 12].

The consistency requirements of explicit event generator implementation force the introduction of additional complexities, since the colour charges of the resulting beam remnants must be resolved into colour singlet final-state hadrons. The initially simple form of this connection between MPI scatterings has been refined in recent years by the work of Skands, Sjöstrand, and Corke in implementing colour string reconnection, MPI rescattering, and x -dependent proton size models in PYTHIA6 and latterly Pythia 8. Colour reconnection is the final aspect of MPI modelling that we will discuss here. The motivation for this is that with many colour strings/dipoles being created by the multiple scattering, some form of annealing may take place on the timescale of hadronisation to form more energy/action-efficient topological configurations. This model was originally introduced as part of PYTHIA hadronisation, with an addition refinement to re-suppress the effect of such annealing for high- $p_\perp$ colour strings (motivated by the idea that such strings will have less time to participate in annealing) [13], and has recently been introduced into Herwig++, although differently formulated for cluster rather than string hadronisation [14]. Pythia 8 has additionally introduced a related form of topological reconfiguration called “rescattering”, whereby MPI

interactions may interact at a diagrammatic level: we will not consider this further in this note. Colour reconnection introduces one or more parameters related to the strength of the reconnection probability in the annealing process.

This concludes the summary of MC MPI modelling most prevalent in the general-purpose MC generators in use at the LHC. Alternative, although in most cases strikingly similar, models have also been developed, notably with inspiration from Regge models such as in PHOJET [15], nuclear collective excitations as in EPOS [16], the use of CCFM parton shower evolution as in CASCADE [17], and dipoles as in the DIPSY [18] model. However, the eikonal partonic scattering model pioneered in PYTHIA remains the mainstay of MPI simulation in general purpose event generators such as JIMMY [4], Herwig++ [19, 20], PYTHIA 6 and 8 [13, 21], and Sherpa [22]. As a result, this is the model with most current influence on LHC signal and background MPI simulation, and is the one on which the phenomenological aspects of the remainder of this note will concentrate.

3. ATLAS measurements of soft QCD observables

The crucial inputs to improvements in the quality of soft QCD modelling at the LHC are of course measurements of observables at the LHC which are particularly sensitive to MPI model features. There is no observable which is a purely MPI phenomenon — quantum mechanics tells us that we must consider all compatible processes as contributing towards any observable, and in the case of soft QCD observables competing effects such as initial state radiation (ISR), diffractive process matrix elements, hadronisation, *etc.* are all potential contributors to nominally “MPI” observables. The issue of decoupling these modelling aspects in the process of model tuning will be addressed in Section 4.

In the ATLAS experiment, the analyses of most importance to constraining MPI models are as follows:

- Diffractive and inelastic cross-section measurements;
- Minimum bias measurements at 900 GeV and 7 TeV;
- Underlying event measurements with leading cluster and leading track at 900 GeV and 7 TeV.

We will now briefly summarise each of these analyses, as well as in-progress analyses expected to contribute to future soft QCD phenomenology studies.

3.1. Diffractive and inelastic cross-section measurements

As mentioned in Section 2.2, MPI models make extensive use of parameterisations of total pp cross-section, particularly one of the Donnachie–

Landshoff parameterisations. All MPI models choose the mean number of partonic scatters based on the ratio of jet cross-section to σ_{inel} , the PYTHIA MPI models evolve their regularisation scale $p_{\perp}^{\text{min}}$ with an ansatz inspired by the DL Pomeron slope, and Herwig++'s MPI model makes an explicit analytic connection to the elastic slope determined from the DL parameterisation via the eikonal formalism. It is hence important to constrain the (components of the) pp total cross-section at 7 TeV based on experimental data.

The ATLAS measurement of the inelastic pp cross-section at 7 TeV is based on use of the forward minimum bias trigger scintillator (MBTS) detectors — a pair of 16-element scintillators located at the calorimeter end-caps on both sides of the interaction point at $z = \pm 3.56$ m, covering the range $2.09 < |\eta| < 3.84$ — and a luminosity measurement to a precision of 3.4% with the LUCID Čerenkov detector at $z = \pm 17$ m. The experimental definition of an inelastic event is that at least two of the 32 MBTS segments have a charge above the noise threshold, *i.e.* that there is measurable proton dissociation on at least one side of the detector. The measured inelastic cross-section at 7 TeV was measured within acceptance to be 60.33 ± 2.10 mb, slightly in tension with previous fits as shown in figure 2, with an extrapolation to the elastic limit indicating better agreement but with a much-increased systematic error due to the extrapolation. Additionally, a subset




Fig. 2. Inelastic pp cross-section as a function of centre-of-mass energy $\sqrt{s}$, with ATLAS' measurement indicated with the solid (red) dot at 7 TeV, by comparison with the parameterisation predictions shown with thick (red) lines. The triangular (blue) point and associated vertical (blue) bar is the ATLAS measurement extrapolated to full elastic acceptance, for comparison with the long-dashed and dot-dashed thin (blue) lines and shaded areas.

of the events were identified as “single-sided” when at least two MBTS segments were activated on one side, and none on the other side: single-sided events are expected to be dominated by single diffractive pp scattering. The measured fraction of single-sided to inclusive inelastic events R_{ss} was measured as just over 10%, which is shown by comparison to various MC models as a function of their diffractive cross-section fraction f_D in figure 3. A slight reduction in model diffractive cross-sections is favoured.




Fig. 3. Comparison of the measured same-side-to-inclusive inelastic event rate ratio, measured as $R_{ss} = [10.02^{+0.01}_{-0.04}]%$ by ATLAS, to various model predictions as a function of the diffractive fraction of the inelastic cross-section f_D in the models. The markers indicate the default values of f_D for each model: the indication is that at 7 TeV the fraction of diffractive contribution to the inelastic cross-section should be reduced from $\sim 32\%$ to $\sim 27\%$ in most models, except PHOJET in which an increase of f_D from $\sim 20\%$ is required.

3.2. Minimum bias measurements

“Minimum bias” is a much misapplied term in soft QCD physics: depending on whether one is speaking from an experimental or theoretical perspective, it respectively indicates a class of observables constructed on events selected using minimally strict conditions (either trigger or offline), or it is used to classify an event type in which there is generally no very hard interaction and where soft multiple scattering is the dominant physical process. These two concepts are closely related: if one minimally biases experimental selection criteria, then the majority of events will be dominated by such an interaction mode — but the distinction is still useful to draw, not least because an experimental minimum bias selection will also select

all kinds of “hard” processes, and because the phrase is also sometimes used to indicate only “non-single-diffractive” (NSD) processes: another case of misleading leakage from the calculational division of process types into the classification of real-world collider events.

All the measurements discussed here are measured using the ATLAS “minimum bias” trigger stream, but the specific collection of observables usually regarded as “minimum bias” are simple observables such as the η and $p_{\perp}$ distributions of tracks and calorimeter clusters, and the distribution of the number of (charged) particles in an event or the correlation of other mean properties with such an event-level property. Observables which further explicitly “bias” the event selection, such as the “underlying event” observables discussed in the next section, are considered distinct.

ATLAS measurements of these minimum bias observables have been made at 900 GeV and 2.36 and 7 TeV, using the low pile-up 2010 dataset to obtain clean measurements. Again, the MBTS trigger scintillators were key to the analysis: at least one MBTS hit was required on each side of the detector, in addition to a number n_{trk} charged particles above a track $p_{\perp}$ cut within the tracker acceptance of $|\eta| < 2.5$. Various values of the track $p_{\perp}$ cut and n_{trk} were used to change the phase space within which the observables mentioned above are computed: the more particles in the tracker acceptance, and the higher their $p_{\perp}$ cut, the more the events are expected to be dominated by perturbative QCD. Jet structure is expected to start emerging with the restriction of phase space, but these observables do not highlight that transition. The use of the two-sided MBTS requirement is a purely experimental version of the model-dependent NSD definition used at previous colliders, and is expected to suppress (but not eliminate) single-diffractive and elastic scattering events.

The $(p_{\perp}^{\text{trk}}, n_{\text{trk}})$ phase spaces used in this measurement are as follows: (100 MeV, 1), (100 MeV, 2), (100 MeV, 20), (500 MeV, 1), (500 MeV, 6), (2500 MeV, 1). Examples are shown in figure 4.

3.3. Underlying event measurements

The “underlying event” (UE) is the name that we give to all elements of a hadron collision which cannot be directly identified with the hard scattering process. This is a rather naïve view, and it is completely correct to say that “there is no underlying event; there is only event”² — however, it reflects the perspective that must be taken to make progress towards new physics discoveries: that there are hard interactions of interest whose clear experimental signatures are complicated and diluted by extra contributions related to MPI and ISR. In terms of soft QCD measurements however, UE

² © Rick Field, MPI@LHC 2008.

is almost always taken to mean observables which have been constructed to focus on non-hard aspects of event structure, and, in particular, to study the evolution of such aspects as a function of the hard scale. It should be said right away that the UE is not necessarily “soft” — fluctuations in MPI and ISR may produce new semi-hard jets, particularly if the hardest scattering in the event is very hard, *e.g.* a TeV-scale QCD or EW interaction.

ATLAS has published two UE measurements based on the 2010 dataset (again to avoid the pile-up contamination of the 2011 runs). Both use the topological decomposition illustrated in figure 5, which was first established by the pioneering UE measurements of the CDF experiment. In this construction, events are azimuthally oriented along an axis which represents the flow of energy in the hardest scattering in the event, so that aspects of the UE may be seen most clearly (with minimal contamination from the hard process) in the transverse directions. This leading axis could be determined using *e.g.* a tensorial diagonalisation of some kind, but is more usually taken to simply be the direction of the leading jet or reconstructed boson. Most UE observables are constructed to show the dependence of the $p_{\perp}$ and multiplicity observed in each region as a function of the $p_{\perp}$ of the hard process. If the $p_{\perp}$ of the hard process may be safely used down to the lowest scales, UE observables hence show the transition of MPI from “minimum bias” physics into the hard scattering regime.




Fig. 5. Topological decomposition of dijet/Drell–Yan events for underlying event measurements.

The ATLAS measurements use two different detector elements to make their measurements of UE quantities at 900 GeV and 7 TeV: the first follows the lead of the CDF measurements by using tracking information in

the azimuthal regions, whereas the second is the first measurement of UE properties using calorimeter clusters. In both cases, to avoid systematics issues with calorimeter jets in the ATLAS commissioning phase the leading object in the event is taken to be a “single particle” — a charged track or calorimeter cluster respectively — rather than a jet. This limits the range of




Fig. 6. Leading-track-based underlying event measurements at 900 GeV (left column) and 7 TeV (right column). The first row is the dependence of the mean number of particles with $p_{\perp} > 500$ MeV in the transverse regions as a function of the $p_{\perp}$ of the leading track; the second row shows the same evolution for the mean $p_{\perp}$ of those particles; and the final row shows the scale evolution of mean $p_{\perp}$ for the more inclusive track $p_{\perp}$ requirement of 100 MeV.

validity of the measurement, since at some scale the leading particle will not

necessarily be in the leading jet, and so both measurements are made using the minimum bias trigger stream, with a scale reach only up to ~ 20 GeV. Future ATLAS UE measurements will extend this programme to use leading track jets, calorimeter jets, and Z and W events.

Similarly to the minimum bias analysis, the leading track analysis uses two different track $p_{\perp}$ cuts, 100 and 500 MeV. The leading cluster analysis uses all clusters. Examples of observables from these analyses are shown in figure 6 and 7. The dominant features are the “ramp and plateau” structure in the UE plots against leading object scale: this is the “pedestal” structure driven by the increase and saturation of hadron form factor overlap as the hard event scale increases. Several connections between the plots are worth nothing:

- the plateau heights represent roughly twice the charged particle number and $p_{\perp}$ density as seen in the minimum bias analysis with the same cuts;
- the level of UE plateau activity increases by a factor of two between 900 GeV and 7 TeV;
- a factor of 1.5 increase in $\sum p_{\perp}$ is seen in reducing the track $p_{\perp}$ cut from 500 to 100 MeV;
- and the cluster-based analysis (which also measures neutral particles) has a higher $p_{\perp}$ density plateau value than the most inclusive track-based observable, by a factor of roughly 30%.

3.4. Other soft QCD analyses

Finally, we summarise further analysis efforts which are either in progress or which simply cannot be given due weight in this contribution.

Several other ATLAS analyses are in progress with direct relevance to soft QCD and MPI. The most obvious are extensions of the leading track/cluster underlying event analyses to use composite leading objects such as track and calorimeter jets: both are ongoing, as is a UE measurement in $Z \rightarrow e^+e^-/\mu^+\mu^-$ events.




Fig. 7. Underlying event measurements at 900 GeV and 7 TeV. The first row shows the correlation of mean track $p_{\perp}$ and region multiplicity in the transverse region of the leading track UE analysis. The second row shows the emergence of azimuthal jet structure in the leading cluster analysis as the leading cluster scale increases; and the final row shows the sum of transverse region cluster $p_{\perp}$ density as a function of leading cluster $p_{\perp}$.

There are also a series of analyses probing the flavour and correlation structure of minimum bias events: specifically the measurement of production rates and $p_{\perp}$ s of Λ and K_s hadrons, and correlations between both pairs of charged particles and between matched forward/backward pseudorapidity intervals.

A final set of analyses is focused on looking for specific soft QCD processes: jet gap analyses provide a direct study of diffractive topologies at 7 TeV, while a set of explicit searches for hard double parton scattering (DPS) in various event types will investigate the validity and universality of the “ σ_{eff} ” model used so far to formalise DPS calculations.

While not an explicit physics analysis, commissioning of pile-up modelling which provides a good description of LHC data is also a major driver of MPI model exploration and tune iteration.

4. Testing and tuning of MPI models

As already alluded to in Section 2.2, the various families of general-purpose event generators have distinct philosophies of which modelling aspects may be tuned, and which should be robustly predicted from theoretical inputs.

A particularly controversial area is that of tuning the parton showers in the PYTHIA generator family: parton showers are an iterative process whereby partons recursively radiate by means of *splitting functions* which are collinear expansions of full LO QCD matrix elements. As the evolution scales and splitting functions are definitively perturbative, generator authors disagree about whether fudge factors can be justified on the scales and couplings in the various types of parton shower — although the rapid rise of shower-matched higher-order matrix element generators such as POWHEG, AlpGen, and MC@NLO does appear to be driving the field in a more theoretically constrained direction.

Much less controversial is tuning of processes which involve the divergent strong coupling, such as hadronisation and MPI. As there is no IR-complete model of QCD, any process which has to approach or transcend Λ_{QCD} must be phenomenologically constructed and hence possess free parameters to be tuned to experimental data. The only disagreement is in the number of parameters available: again, the PYTHIA family favours configurability, while the HERWIG and Sherpa generators attempt to be more minimal. Each has their place in collider physics research, and much of the ubiquity of PYTHIA6 in the LHC experiments is due to the ability to make its output look very much like data — the price for this is reduced predictivity.

Typically hadronisation models introduce between 10 and 30 parameters: string hadronisation models tend to require more parameters for flavour structure, as the string breaking is not predictive about this, while cluster

hadronisation requires more parameters to fix the kinematics of cluster splitting. MPI models add an extra 5 or more parameters, the number depending strongly on the degree of available refinement in parameterising the proton form factor and the colour reconnection mechanism. Tuning 30+ parameters all at once is not a computationally feasible approach, even with modern semi-automated tools and most certainly not if the tuning is done by hand. Hence some factorisation of parameters is required: hadronisation, for example, if assumed to be universal between lepton and hadron colliders, can be cleanly tuned to LEP, JADE, and SLD data without any need for a functioning MPI configuration. A more specific example is that of the b and c quark fragmentation functions, which usually receive special treatment and can be tuned in isolation (if necessary) once a base tune of the light quark fragmentation has been established. The observables most sensitive to MPI (and shower) configuration may then be tuned using the final state setup constructed from the e^+e^- data.

This approach is the one taken by ATLAS' tuning group, with the constraint to data being made with the Rivet [23] analysis toolkit and Professor [8] tune optimisation program. This toolchain has been key to systematising the process of event generator tuning for the LHC, as Rivet provides standardised and validated Monte Carlo versions of all the relevant experimental analyses, and Professor is a numerically efficient and scalable system for numerical optimisation of the event generator parameters with respect to the reference data. The specific method used by Professor is as follows:

- Randomly sample points in the (possibly factorised) parameter space;
- At each point run a full set of high-statistics MC runs, for every collider configuration and process type that should contribute to the tuning. This part may take several days for each run — hence serial optimisation with a standard gradient descent minimiser is impractical — and hence the scalability of the Professor system relies on the ability to batch-parallelise this generation step;
- For each bin of each distribution taken independently, use the pseudoinverse method via a singular value decomposition to algebraically determine optimal coefficients for an arbitrary-order polynomial parameterisation of the bin value as a function of the MC parameters. Special treatments are also made for the statistical and theoretical errors;
- The many bin parameterisations are aggregated and compared to the reference data to compute a goodness of fit measure. This is then numerically optimised, since evaluating the generator observable predictions is now exceedingly fast: the result is an optimal generator tune.

The same speed of evaluation of the parameterised MC generator means that Professor can also:

- Provide an interactive GUI explorer for generator configurations;
- Use multiple equivalent parameterisations to obtain an error estimate on the accuracy of the parameterisation;
- Compute objective error tunes (or “eigentunes”) similar to the error sets produced in PDF fitting using a Hessian formalism.

4.1. ATLAS tunes of PYTHIA6, Pythia 8, and HERWIG/JIMMY

ATLAS has made several iterations of MC tunes, particularly for the PYTHIA6 generator which until recently has been responsible for the vast bulk of ATLAS simulation production. ATLAS’ involvement in generator tuning began with the PYTHIA6 and HERWIG/JIMMY “MC08” tunes in 2008, and began using the automated Rivet and Professor tools with elements of the MC09 pre-LHC tunes. The advent of early LHC data provided the first constraints on MPI models at 7 TeV (as well as an extra low-energy pp point at 900 GeV), and the ATLAS minimum bias observables of Section 3.2, as well as the UE data at 900 GeV, were used to drive the PYTHIA6 AMBT1 [24], and the HERWIG/JIMMY AUET1 [25] tunes in 2010.

The most recent tuning series from ATLAS has for the first time incorporated Pythia 8 into the tuning, as part of the general migration of MC simulation to use the C++ era replacements for the venerable Fortran generators. The second round of ATLAS tuning to its own data tunes both the initial state shower and the MPI model in PYTHIA6, in an attempt to describe both hard and soft QCD modelling with a single configuration. The tunings of Pythia 8 and HERWIG/JIMMY were restricted to the MPI modelling only, in the case of Pythia 8 because its description of jet structure was already very good, unlike PYTHIA6 AMBT1. In all cases, the tunings were performed for a range of leading order and MC-adapted LO (or “mLO”) PDFs, with NLO PDFs and tuning of hybrid generators such as AlpGen and POWHEG being reserved for a later study.

The shower tuning stage for PYTHIA6 used CDF and ATLAS jet shape and track-jet fragmentation data [26,27,28], D0 and ATLAS dijet azimuthal decorrelation data [29,30], and CDF $Z p_{\perp}$ data [31]. The latter study turns out to unfortunately bias the shower tune to the detriment of the ATLAS data $Z p_{\perp}$, which will be addressed in the *next* round of ATLAS shower tuning. The hadronisation, including specific b -fragmentation behaviour, had previously been tuned to LEP data, and the fitting weights were chosen to bias the fit toward a good description of the ATLAS observables. The results

were largely successful, with particular improvement of jet structure observable description although at the cost of multiple Λ_{QCD} values in the code — this latter point has much relevance to the interaction of showers with higher-order matrix element generators and is under active investigation. At the close of the AUET2 tune construction (which also includes and MPI tune), it was observed that the Perugia2010 tune series which had inspired the shower treatment resulted in some perverse behaviours of PYTHIA6-derived non-perturbative corrections in QCD jet studies: as a result of this, a second tuning round for the ATLAS MC11 production was started, which used the more conventional shower configuration of AMBT1 but tuned the three parameters required to optimise jet structure and near- π dijet decorrelation description. This tune series was also successful at describing jet data and, denoted as AMBT2B/AUET2B, was extended to include the LO*, LO**, CT09MC2, CTEQ6L1, and MSTW2008LO PDFs [32], all of which were then used as a base for MPI tuning.

For the MPI tuning of all three generators, the ATLAS data described in Sections 3.2 and 3.3 were used, in addition to the full range of minimum bias and underlying event measurements from CDF: minimum bias [33, 34], leading track UE at 1800 GeV [35], leading jet UE in jet events at 630, 1800, and 1960 GeV [36, 37], and the UE in $Z \rightarrow e^+e^-$ Drell–Yan events at 1960 GeV [37]. The weights in the fit were again chosen to bias the fit toward a good description of the ATLAS observables at the potential expense of the Tevatron ones, and for the 7 TeV data more than the 900 GeV data: this is designed to optimise the description of the observables that ATLAS most needs to simulate for the next two years. As the JIMMY MPI model specifically has no treatment of purely soft scattering, *i.e.* it is only formulated in the presence of a hard scattering and makes no attempt to IR-complete the MPI scatterings below $p_{\perp}^{\text{min}}$, it was only tuned to UE data. The Pythia-like energy evolution ansatz from Eq. (2) was manually applied to JIMMY’s $p_{\perp}^{\text{min}}$ value by means of sampled meta-parameters.

The results of the MPI fitting for all three generators are very strongly dependent on the PDF being used, since the MPI secondary scatterings in the models are mostly driven by the nature of the low- x gluon PDF, for $Q^2 \sim 10 \text{ GeV}^2$ and $x \sim 10^{-4}$. The modified leading order PDFs, intended specifically to make LO event generator simulation of hard processes more “NLO-like”, typically have larger low- x gluon fractions than the standard LO and NLO PDFs, and hence the $p_{\perp}^{\text{min}}$ screening parameter naturally increases to produce the same level of activity as for a less MPI-active PDF. Plots of the minimum bias tunes for both Pythia-family generators are shown in figures 8, 9, and underlying event tunes for all three generators in figures 10, 11, and 12.




Fig. 8. ATLAS minimum bias tunes of PYTHIA6 (left column) for LO PDFs and Pythia 8 (right column) for an LO and an mLO PDF, compared to ATLAS minimum bias observables (and the 4C author tune of Pythia 8). The data description is generally good, with most difficulty being observed in the description of the $p_{\perp}$ spectrum. This is particularly the case for the LO** PDF tune of Pythia 8 — a similar effect is seen for PYTHIA6 in the next figure.




Fig. 9. ATLAS tunes of PYTHIA6 to mLO PDFs, compared to ATLAS minimum bias $p_{\perp}$ spectrum data, for two different track $p_{\perp}$ cuts. In both cases large excesses above data are seen from 5–30 GeV, up to a factor of 1.8 excess for the lower track $p_{\perp}$ cut. This feature was found to be a regular result of using mLO PDFs, and could not be significantly altered by any use of MPI model parameters.

The major results from the latest ATLAS MPI tuning are as follows:

- A fully consistent description of MB and UE observables could not be obtained with either PYTHIA6 or Pythia 8. The ATLAS UE observables favour slightly more “active” MPI parameter configurations than the MB ones do. This appears to require a modelling extension, and it is possible that the addition of an x -dependent proton matter distribution to the latest Pythia 8 versions may help to describe this data. The result was that separate MB and UE tunes were constructed for each generator, hence the distinction between PYTHIA6’s AUET2B and AMBT2B tunes.
- For all three generators, distinct groupings of MPI cutoff values were discovered, with all mLO PDFs preferring higher values of $p_{\perp}^{\min}$ than the groups of LO PDFs. This is as expected, but is the first explicit and quantitative demonstration of this model behaviour.
- Surprising features were observed for some minimum bias observables in both PYTHIA6 and Pythia 8 when using mLO PDFs. In particular, an excess of MC above data by a factor of 1.8 was seen in the MB track $p_{\perp}$ spectrum for both PYTHIA6 and Pythia 8 when using the LO*, LO**, and CT09MC2 PDFs. This effect could not be changed by use of any MPI model parameters, and attempts to identify which PDF features were responsible were unable to pinpoint any single PDF aspect which was reliably correlated to the anomalous behaviours. The decision was made to not use mLO PDFs for minimum bias simulation




Fig. 10. ATLAS tunes of PYTHIA6 compared to ATLAS underlying event data, for LO (left column) and mLO (right column) PDF tunes. Both types of PDF describe UE data well, and so there is no significant MPI-oriented problem to using mLO PDFs for simulation of the UE in hard-scale event simulation, unlike the case for minimum bias generation seen in the previous figure. There does appear to be a slight effect of mLO PDFs in that all tunes using them slightly undershoot the “turn-over” region of the N_{ch} and $\Sigma p_{\perp}$ profiles.




Fig. 11. ATLAS systematic MPI error tunes of PYTHIA6 compared to ATLAS UE data at 7 TeV. These tunes have been constructed similarly to PDF Hessian error sets, by requiring fixed deviations in goodness of fit from the optimised tunes along diagonalised principle directions in the parameter space, and hence provide a set of tunes which quantitatively represent the uncertainties in PYTHIA6 MPI modelling and ATLAS UE data.

— while some artefacts were also observed in UE observables, they were much less extreme than in MB, to the extent that they should be a very minor issue compared to the benefits to the hard process simulation of using an mLO PDF.

- For all generators, leading order PDFs were capable of the best description of MPI observables. The CTEQ6L1 PDF was seen to be particularly good at describing MPI data with both Pythia generators. This is consistent with theory, as there is no motivation for mLO PDFs to improve the description of the dijet matrix elements used for MPI scattering simulation, but is again the first observation that not all PDF effects can be “tuned away” by MPI models.
- No model is currently capable of describing either MB or UE data for tracks with $p_{\perp}$ below 500 MeV.

The AUET2 tunes of HERWIG/JIMMY are the last such tunes which will be constructed by ATLAS: they provide generally good descriptions of the UE for signal processes, and cover sufficient PDFs for use in PDF systematics studies or for combination with NLO generators such as MC@NLO. There is, however, not enough flexibility in the JIMMY model to describe the detailed structure of UE observables: essentially the 3 available parameters ($p_{\perp}^{\text{min}}$, $\sqrt{s}$ evolution exponent, and matter radius) all move the UE plateaus up




Fig. 12. ATLAS underlying event tunes of Pythia 8 and HERWIG/JIMMY, compared to the same ATLAS UE data at 7 TeV as shown in earlier figures for PYTHIA6. The quality of description is again good for Pythia 8, but for HERWIG/JIMMY while the overall level of plateau activity is seen to be correct, its modelling both undershoots the turn-over region (bottom left) and does not quite capture the necessary balance between charged particle multiplicity and $p_{\perp}$ (bottom right). These limitations could not be addressed by tuning of the JIMMY MPI model, and no further JIMMY tuning will be pursued by ATLAS.

and down but without any more nuanced changes in shape. The emphasis for MPI models other than JIMMY will now move to Herwig++ and Sherpa, and HERWIG dependence is itself being phased out of ATLAS production in favour of the newer and more capable C++ generators. It is notable, however, that despite the simplicity of the JIMMY model — no colour reconnection, a very minimal matter distribution parameterisation, *etc.* — it does describe most UE data very well! In particular, the indication appears to be that the detail of how $p_{\perp}^{\min}$ is used to regularise MPI scattering cross-sections is not particularly crucial, at least above a track $p_{\perp}$ of ~ 500 MeV.

We should also mention the successfulness of the DL Pomeron-inspired energy evolution ansatz for minimum bias data. This has been studied by Schulz and Skands [38], by using the same tuning machinery as used by ATLAS to fit values of $p_{\perp}^{\min}$ for PYTHIA6 against all available minimum bias data at a range of energies. There is no assumption of an energy evolution form in this study, yet the usual ansatz does give a strikingly high quality fit to this evolution form, as shown in figure 13. However, it is not clear whether underlying event observables are also compatible with this form, as they certainly do not appear to be compatible with minimum bias description at the *same* energy with current models.




Fig. 13. The evolution of tuned values of $p_{\perp}^{\min}$ in the PYTHIA6 MPI model as a function of collider centre-of-mass energy $\sqrt{s}$, for minimum bias data with a fiducial charged multiplicity cut of $N_{\text{ch}} \geq 1$. A power law form of evolution, as used by the Pomeron-inspired ansatz in the Pythia family of event generators, gives a remarkable description of empirical evolution of the MPI cutoff.

Finally, we note that the ATLAS measurements of inelastic and diffractive cross-sections, discussed in Section 3.1, have also had an impact on MPI modelling, although not directly in ATLAS studies. The “4C” author tune of Pythia 8 specifically adds a damping of diffractive cross-section evolution into its improved inclusive diffraction model, to better describe this and other data, and this tune is in use in ATLAS production for the simulation of pile-up minimum bias events.

5. Conclusions

We have described the modelling of multi-parton interactions in Monte Carlo generators, and the measurements and generator tuning activities in the ATLAS experiment which have been used to increase our understanding and modelling ability for multiple scattering effects at the LHC. Tunes have been constructed using a wide range of PDFs for three event generators, PYTHIA6, Pythia 8, and HERWIG/JIMMY, in addition to ATLAS' use of author-supplied tunes of the Herwig++ and Sherpa generators. Interesting dependences on PDF details have been observed, in particular several strong and anomalous effects of MC-adapted (mLO) PDFs on MPI observables. No model currently describes MPI-dominated data well below a particle $p_{\perp}$ cut of 500 MeV.

The increased constraint on the energy evolution of inclusive MPI will also assist the LHC program when the centre-of-mass energy is increased to 8, 9, and/or 14 TeV in the coming years. This improved description of ATLAS data has not been entirely without cost: particularly where ATLAS underlying event data have been used, there are signs of tension with CDF data and between MB and UE observables, forcing a split into MB and UE tune families. This matter, and that of improved hadronisation and jet structure description, will be taken up in the next set of ATLAS tunes — including data with greater hard-scattering scales in UE and including flavour constraints via identified particle data — and model developments.

A.B. wishes to thank the organisers of the Kraków Theoretical Physics Summer School for the invitation and hospitality in Zakopane, other members of the ATLAS soft QCD and MC tuning subgroups for their many contributions to the work described here. Many thanks also to the Edinburgh University PPE group for additional support, and the Durham University IPPP and Scottish Universities Physics Alliance for an Associateship grant and Advanced Fellowship respectively.

REFERENCES

- [1] B. Wynne, Imagirol: an Implementation of Bayesian Iterative Unfolding, Technical Report ATL-COM-PHYS-2011-674, CERN, Geneva, June 2011.
- [2] M. Derrick *et al.* [ZEUS Collaboration], *Phys. Lett.* **B316**, 412 (1993).
- [3] T. Sjostrand, M. van Zijl, *Phys. Rev.* **D36**, 2019 (1987).
- [4] J.M. Butterworth, J.R. Forshaw, M.H. Seymour, *Z. Phys.* **C72**, 637 (1996).
- [5] A. Donnachie, P. Landshoff, *Phys. Lett.* **B296**, 227 (1992).
- [6] A. Donnachie, P. Landshoff, *Phys. Lett.* **B437**, 408 (1998).

- [7] A. Donnachie, P. Landshoff, *Phys. Lett.* **B518**, 63 (2001).
- [8] A. Buckley *et al.*, *Eur. Phys. J.* **C65**, 331 (2010).
- [9] M. Bahr, S. Gieseke, M.H. Seymour, *J. High Energy Phys.* **0807**, 076 (2008).
- [10] M. Bahr, J.M. Butterworth, S. Gieseke, M.H. Seymour, [arXiv:0905.4671v1 \[hep-ph\]](#).
- [11] A. Corsetti, R. Godbole, G. Pancheri, *Phys. Lett.* **B435**, 441 (1998).
- [12] R. Godbole, A. Grau, G. Pancheri, Y. Srivastava, *Phys. Rev.* **D72**, 076001 (2005).
- [13] T. Sjöstrand, S. Mrenna, P. Skands, *J. High Energy Phys.* **05**, 026 (2006).
- [14] S. Gieseke, C. Rohr, A. Siodmok, [arXiv:1110.2675v1 \[hep-ph\]](#).
- [15] R. Engel, *Z. Phys.* **C60**, 203 (1995).
- [16] K. Werner, *Nucl. Phys. Proc. Suppl.* **175-176**, 81 (2008).
- [17] H. Jung, G. Salam, *Eur. Phys. J.* **C19**, 351 (2001).
- [18] C. Flensburg, G. Gustafson, L. Lonnblad, *J. High Energy Phys.* **1108**, 103 (2011).
- [19] M. Bahr *et al.*, *Eur. Phys. J.* **C58**, 639 (2008).
- [20] S. Gieseke *et al.*, [arXiv:1102.1672v1 \[hep-ph\]](#).
- [21] T. Sjostrand, S. Mrenna, P.Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [22] T. Gleisberg *et al.*, *J. High Energy Phys.* **0902**, 007 (2009).
- [23] A. Buckley *et al.*, [arXiv:1003.0694v6 \[hep-ph\]](#).
- [24] ATLAS Collaboration, ATLAS Note (2010) ATLAS-CONF-2010-031.
- [25] ATLAS Collaboration, First Tuning of HERWIG/JIMMY to ATLAS data, Technical Report ATL-PHYS-PUB-2010-014, CERN, Geneva, Oct. 2010.
- [26] D. Acosta *et al.* [CDF Collaboration], *Phys. Rev.* **D71**, 112002 (2005).
- [27] G. Aad *et al.* [ATLAS Collaboration], *Phys.Rev.* **D83**, 052003 (2011).
- [28] G. Aad *et al.* [ATLAS Collaboration], [arXiv:1109.5816v1 \[hep-ex\]](#).
- [29] V.M. Abazov [D0 Collaboration], *Phys. Rev. Lett.* **94**, 221801 (2005).
- [30] G. Aad *et al.* [ATLAS Collaboration], *Phys. Rev. Lett.* **106**, 172002 (2011).
- [31] A.A. Affolder *et al.* [CDF Collaboration], *Phys. Rev. Lett.* **84**, 845 (2000).
- [32] A. Sherstnev, R.S. Thorne, *Eur. Phys. J.* **C55**, 553 (2008).
- [33] D.E. Acosta *et al.* [CDF Collaboration], *Phys. Rev.* **D65**, 072005 (2002).
- [34] T. Aaltonen *et al.* [CDF Collaboration], *Phys. Rev.* **D79**, 112005 (2009) [[arXiv:0904.1098 \[hep-ex\]](#)].
- [35] A.A. Affolder *et al.* [CDF Collaboration], *Phys. Rev.* **D65**, 092002 (2002).
- [36] D.E. Acosta *et al.* [CDF Collaboration], *Phys. Rev.* **D70**, 072002 (2004).
- [37] T. Aaltonen *et al.* [CDF Collaboration], *Phys. Rev.* **D82**, 034001 (2010).
- [38] H. Schulz, P. Skands, *Eur. Phys. J.* **C71**, 1644 (2011).