The "Perugia" Tunes

P. Skands
Theoretical Physics, Fermilab, MS106, Box 500, Batavia IL-60510, USA

Abstract
We present 7 new tunes of the $p_\perp$-ordered shower and underlying-event model in PYTHIA 6.4. These "Perugia" tunes update and supersede the older "S0" family. The new tunes include the updated LEP fragmentation and flavour parameters reported on by H. Hoeth at this workshop [1]. The hadron-collider specific parameters were then retuned (manually) using Tevatron min-bias data from 630, 1800, and 1960 GeV, Tevatron Drell-Yan data at 1800 and 1960 GeV, as well as SPS min-bias data at 200, 540, and 900 GeV. In addition to the central parameter set, related tunes exploring systematically soft, hard, parton density, and color structure variations are included. Based on these variations, a best-guess prediction of the charged track multiplicity in inelastic, nondiffractive minimum-bias events at the LHC is made.

1 Introduction

Perturbative calculations of collider observables rely on two important prerequisites: factorisation and infrared safety. These are the tools that permit us to relate the calculations to detector-level measured quantities, up to corrections of known dimensionality, which can then be suppressed (or enhanced!) by appropriate choices of the dimensionful scales appearing in the poblem. However, this approach does limit us to consider only a predefined class of observables, at a limited precision set by the aforementioned scales. In the context of the underlying event, say, we are faced with the fact that we do not (yet) have factorisation theorems for this component, while at the same time acknowledging that not all collider measurements can be made insensitive to it at a level comparable to the achievable experimental precision. And when considering observables such as track multiplicities, hadronisation corrections, or even short-distance resonance masses if the precision required is very high, we are confronted with quantities which may be experimentally well measured but which are explicitly sensitive to infrared physics.

Let us begin with factorisation. When applicable, factorisation allows us to subdivide the calculation of an observable (regardless of whether it is infrared safe or not) into a perturbatively calculable short-distance part and a universal long-distance part, the latter of which may be modeled and constrained by fits to data. However, in the context of hadron collisions the conceptual separation into "hard-scattering" and "underlying-event" components is not necessarily equivalent to a clean separation in terms of "hardness" (or perhaps more properly formation time), since what is labeled the "underlying event" may contain short-distance physics of its own. Indeed, from ISR energies [2] through the SPS [3, 4] to the Tevatron [5–9], and even in photoproduction at HERA [10], we see evidence of (perturbative) "minijets" in the underlying event, beyond what bremsstrahlung alone appears to be able to account for. It would therefore seem apparent that a

universal modeling of the underlying event must include at least some degree of correlation between the hard-scattering and underlying-event components. It is in this spirit that the concept of "interleaved evolution" [11] was developed as the cornerstone of the $p_\perp$-ordered models [11, 12] in both PYTHIA 6 [13] and, more recently, PYTHIA 8 [14].

The second tool, infrared safety, provides us with a class of observables which are insensitive to the details of the long-distance physics. This works up to corrections of order the long-distance scale divided by the short-distance scale, $Q^2_{\rm IR}/Q^2_{\rm UV}$, where $Q_{\rm UV}$ denotes a generic hard scale in the problem and $Q_{\rm IR} \sim \Lambda_{\rm QCD} \sim \mathcal{O}(1\ {\rm GeV})$. Since $Q_{\rm IR}/Q_{\rm UV} \to 0$ for large $Q_{\rm UV}$, such observables "decouple" from the infrared physics as long as all relevant scales are $\gg Q_{\rm IR}$. Only if we require a precision that begins to approach $Q_{\rm IR}$ should we begin to worry about nonperturbative effects for such observables. Infrared sensitive quantities, on the other hand, contain logarithms $\log^n(Q^2_{\rm UV}/Q^2_{\rm IR})$ which grow increasingly large as $Q_{\rm IR}/Q_{\rm UV} \to 0$. As an example, consider particle or track multiplicities; in the absence of nontrivial infrared effects, the number of partons that would be mapped to hadrons in a naïve local-parton-hadron-duality [15] picture depends logarithmically on the infrared cutoff.

Min-bias/UE physics can therefore be perceived of as offering an ideal lab for studying nonfactorized and nonperturbative phenomena with the highest possible statistics, giving crucial tests of our ability to model and understand these ubiquitous components. As a beneficial side effect, the improved models and tunes that result from this effort are important ingredients in the modeling of high-$p_\perp$ physics, in the context of which the underlying event and nonperturbative effects furnish a nontrivial "haze" into which the high-$p_\perp$ physics is embedded.

As part of the effort to spur more interplay between theorists and experimentalists in this field, we here report on a new set of tunes of the $p_\perp$-ordered PYTHIA framework, which update and supersede the older "S0" family of tunes. The new tunes have been made available via the routine PYTUNE starting from PYTHIA version 6.4.20.

We have here focused in particular on the energy scaling from lower energies towards the LHC and on attempting to provide at least some form of systematic uncertainty estimates, in the form of a small number of alternate parameter sets that represent systematic variations in some of the main tune parameters

We also present a few distributions that carry interesting and complementary information about the underlying physics, updating and complementing those contained in [16]. For brevity, this text only includes a representative selection, with more results available on the web [17].

The main point is that, while each plot represents a complicated cocktail of physics effects, such that any sufficiently general model presumably could be tuned to give an acceptable description observable by observable, it is very difficult to simultaneously describe the entire set. The real game is therefore not to study one distribution in detail, but to study the degree of simultaneous agreement or disagreement over many, mutually complementary, distributions.

We have tuned the Monte Carlo in four consecutive steps:

1. Final-State Radiation (FSR) and Hadronisation (HAD): using LEP data, tuned by Professor [1, 18].
2. Initial-State Radiation (ISR) and Primordial k_T: using the Drell-Yan $p_\perp$ spectrum at 1800 and 1960 GeV, as measured by CDF [19] and DØ [20], respectively. We treat the data

as fully corrected for photon bremsstrahlung effects in this case, i.e., we compare the measured points to the Monte Carlo distribution of the original Z boson. We believe this to be reasonably close to the definition used for the data points in both the CDF and DØ studies.

3. Underlying Event (UE) and Beam Remnants (BR): using $N_{\rm ch}$ [21], $dN_{\rm ch}/dp_\perp$ [22], and $\langle p_\perp \rangle\,(N_{\rm ch})$ [23] in min-bias events at 1800 and 1960 GeV, as measured by CDF. Note that the $N_{\rm ch}$ spectrum extending down to zero $p_\perp$ measured by the E735 Collaboration at 1800 GeV [24] was left out of the tuning, since we were not able to consolidate this measurement with the rest of the data. We do not know whether this is due to intrinsic limitations in the modeling or to a misinterpretation on our part of the measured result.
4. Energy Scaling: using $N_{\rm ch}$ in min-bias events at 200, 540, and 900 GeV, as measured by UA5 [25, 26], and at 630 and 1800 GeV, as measured by CDF [21]. Note that we include neither elastic nor diffractive Monte Carlo events in any of our comparisons, which could affect the validity of the modeling for the first few bins in multiplicity. We therefore assigned less importance to these bins when doing the tunes. The last two steps were iterated a few times.

Note that the clean separation between the first and second points assumes jet universality, i.e., that a Z^0, for instance, fragments in the same way at a hadron collider as it did at LEP. This is not an unreasonable first assumption, but it is still important to check it explicitly, e.g., by measuring strange to unstrange particle production ratios, vector to pseudoscalar meson ratios, and/or baryon to meson ratios *in situ* at hadron colliders.

Note also that we do not include any explicit "underlying-event" observables here. Instead, we rely on the large-multiplicity tail of minimum-bias events to mimic the underlying event. A similar procedure was followed for the older "S0" tune [27, 28], which turned out to give a very good simultaneous description of both minimum-bias and underlying-event physics at the Tevatron, despite only having been tuned on minimum-bias data there[1]. Conversely, Rick Field's "Tune A" [29–32] was originally only tuned on underlying-event data, but turned out to give a very good simultaneous description of minimum-bias physics. We perceive of this as good, if circumstantial, evidence of the universal properties of the PYTHIA modeling.

Additional important quantities to consider for further validation (and eventually tuning, e.g., in the Professor framework), would be observables involving explicit jet reconstruction and explicit underlying-event observables in leading-jet, dijet, jet + photon, and Drell-Yan events. Some of these have already been included in the Professor framework, see [1, 18]. See also the underlying-event sections in the HERA-and-the-LHC [33], Tevatron-for-LHC [32], and Les Houches write-ups [34].

2 Main Features of the Perugia Tunes

In comparison with tunes of the old (PYTHIA 6.2) framework [35], such as Tune A [29–32], all tunes of the new framework share a few common features. Let us first describe those, with plots to illustrate each point, and then turn to the properties of the individual tunes.

[1] Note: when extrapolating to other energies, the alternative scaling represented by "S0A" appears to be preferred over the default scaling used in "S0".

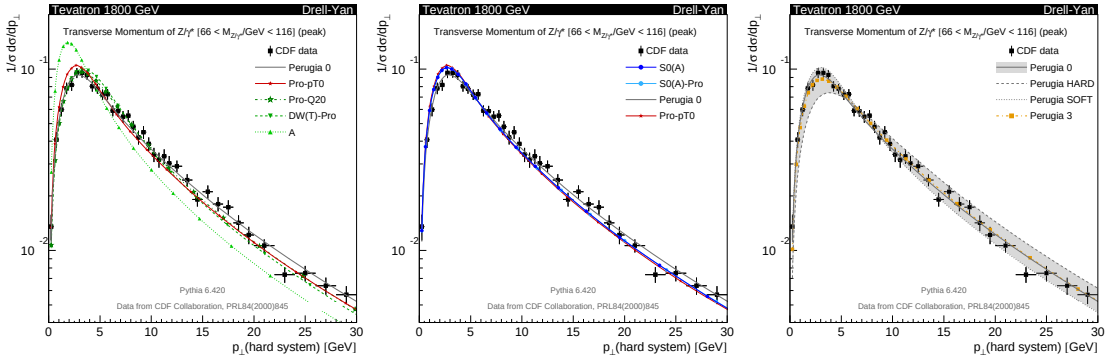

Fig. 1: Comparisons to the CDF Run I measurement of the $p_\perp$ of Drell-Yan pairs [19]. *Left:* a representative selection of models. *Center:* different tunes of the new framework. *Right:* the range spanned by the main Perugia variations. Comparisons to the DØ Run II measurement [20] and results with more tunes can be found at [17]. Note that the Monte Carlo curves shown are for the $p_\perp$ of the original boson rather than of the lepton pair after (QED) showering.

First of all, the new $p_\perp$-ordered showers [11] employ a dipole-style recoil model, which appears to make it very easy to obtain a good agreement with, e.g., the Drell-Yan $p_\perp$ spectrum. In the old model with default settings, the Drell-Yan spectrum is only well described if FSR off ISR jets is switched off. When switching this back on, which is of course necessary to obtain the desired perturbative broadening of the ISR jets, the old shower kinematics work in such a way that each FSR emission off a final-state parton from ISR effectively removes $p_\perp$ from the Z boson, shifting the spectrum towards lower values. This causes any tune of the old PYTHIA framework with default ISR settings — such as Tune A or the ATLAS DC2/"Rome" tune — to predict a too narrow spectrum for the Drell-Yan $p_\perp$ distribution, as illustrated in fig. 1.

To re-establish agreement with the measured spectrum without changing the recoil kinematics, the total amount of ISR in the old model had to be increased. This was done by choosing extremely low values of the renormalisation scale (and hence large α_s values) for ISR (tunes DW-Pro and Pro-Q20 in fig. 1). While this nominally works, the whole business does smell faintly of fixing one problem by introducing another and hence the default in PYTHIA has remained the unmodified Tune A, at the price of retaining the poor agreement with the Drell-Yan spectrum.

In the new $p_\perp$-ordered showers [11], however, FSR off ISR is treated within individual QCD dipoles and does not affect the Drell-Yan $p_\perp$. This appears to make the spectrum come out generically much closer to the data. The only change from the standard $\alpha_s(p_\perp)$ choice used in the S0 family of tunes was thus switching to the so-called CMW choice [36] for $\Lambda_{\rm QCD}$ for ISR in the Perugia tunes, rather than the $\overline{\rm MS}$ value used previously, similarly to what is done in HERWIG [37, 38]. The effect of this relatively small change can be seen by comparing S0(A), which uses the $\overline{\rm MS}$ value, to Perugia 0 in the middle plot on fig. 1. The extremal curves on the right plot are obtained by using $\alpha_s^{\rm CMW}(\frac{1}{2}p_\perp)$ (HARD) and $\alpha_s^{\overline{\rm MS}}(\sqrt{2}p_\perp)$ (SOFT).

Secondly, as mentioned above, we here include data from different colliders at different energies, in an attempt to fix the energy scaling better. Like Rick Field, we find that the default

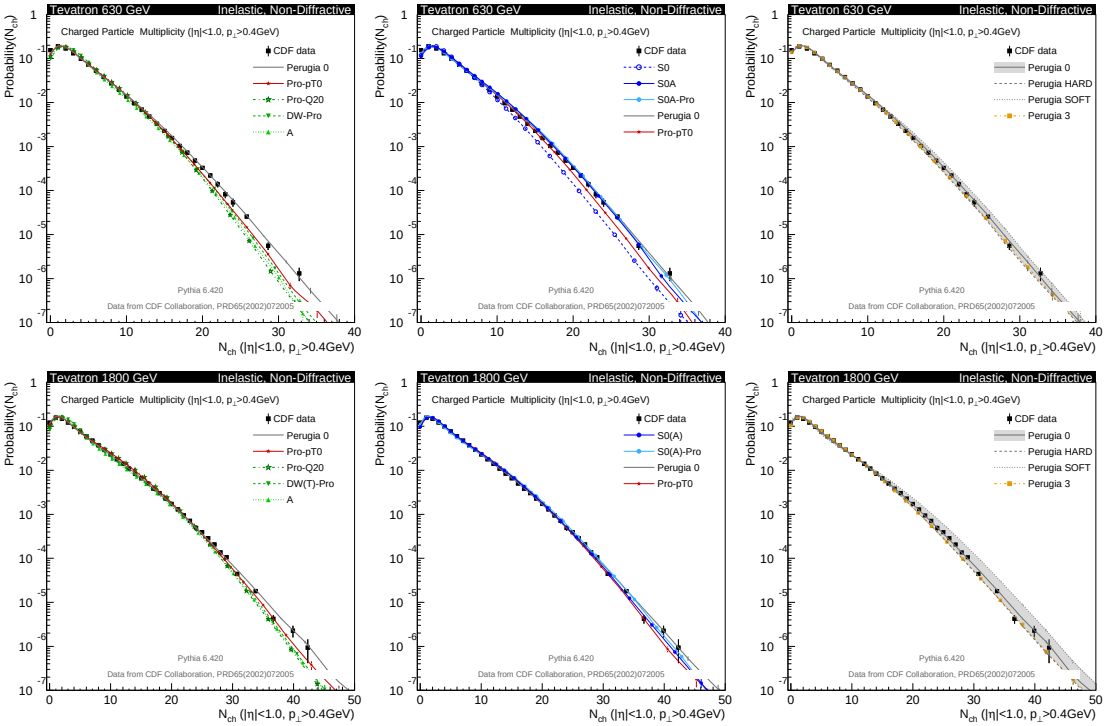

Fig. 2: Comparisons to the CDF measurements of the charged track multiplicity in minimum-bias $p\bar{p}$ collisions at 630 GeV (top row) and at 1800 GeV (bottom row). *Left:* a representative selection of models. *Center:* different tunes of the new framework. *Right:* the range spanned by the main Perugia variations. Results with more tunes can be found at [17].

energy scaling behaviour in PYTHIA results in the overall activity growing too fast with collider energy. This can be mitigated by increasing the dependence of the MPI infrared cutoff on collider energy. For Tune A, Rick Field increased the power of this dependence from $\propto E_{\mathrm{cm}}^{0.16}$ (the default, see [13]) to $\propto E_{\mathrm{cm}}^{0.25}$. The Perugia tunes incorporate a large range of values, between 0.22 and 0.32, with Perugia 0 using 0.26, i.e., very close to the Tune A value. Note that the default was originally motivated by the scaling of the total cross section, which grows like $\propto (E_{\mathrm{cm}}^2)^{0.08}$. It therefore seems that at least in the current models, the colour screening / infrared cutoff of the individual multi-parton interactions needs to scale significantly faster than the total cross section. A discussion of whether this tendency could be given a meaningful physical interpretation (e.g., in terms of low-x, saturation, or unitarisation effects) is beyond the scope of this contribution.

As evident from fig. 2, the Perugia tunes all describe the Tevatron N_{ch} distributions at 630 (top) and 1800 (bottom) GeV within an acceptable margin. Note that the charged track definition is here $p_\perp > 0.4$ GeV, $|\eta| < 1.0$, and particles with $c\tau \geq 10$mm treated as stable. To highlight the difference in the scaling, the middle plot shows both Tune S0 and Tune S0A at 630 GeV. These are identical at 1800 GeV and only differ by the energy scaling, with S0 using the default

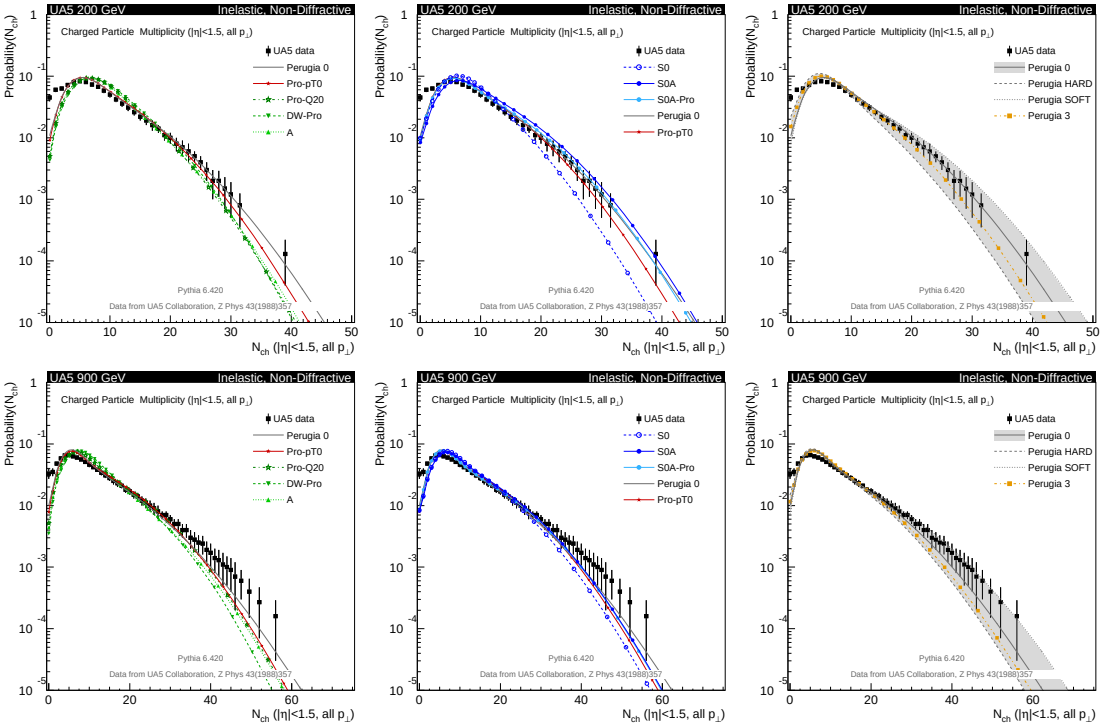

Fig. 3: Comparisons to the UA5 measurements of the charged track multiplicity in minimum-bias $p\bar{p}$ collisions at 200 GeV (top row) and at 900 GeV (bottom row). *Left:* a representative selection of models. *Center:* different tunes of the new framework. *Right:* the range spanned by the main Perugia variations. More results can be found at [17].

scaling mentioned above and S0A using the Tune A value. It is mainly the comparative failure of S0 with the default scaling to describe the 630 GeV data on the top middle plot in fig. 2 that drives the choice of a slower-than-default pace of the energy scaling of the activity (equivalent to a higher scaling power of the infrared cutoff, as discussed above).

A similar comparison to UA5 data at two different energies, but now in a slightly larger η region and including all $p_\perp$ is shown in fig. 3. Since the data here includes all $p_\perp$, the theoretical models have been allowed to deviate slightly more from the data than for the Tevatron and the first few bins were ignored, to partly reflect uncertainties associated with the production of very soft particles.

The good news, from the point of view of LHC physics, is that even the most extreme Perugia variants need to have a more slowly growing activity than the default. Thus, their extrapolations to the LHC produce *less* underlying event than those of their predecessors that used the default scaling, such as S0, DWT, or ATLAS-DC2/Rome.

Thirdly, while the charged particle $p_\perp$ spectrum (see [17, dN/dpT]) and $N_{\rm ch}$ distribution in Tune A was in almost perfect agreement with Tevatron min-bias data, the high-multiplicity behaviour of the $\langle N_{\rm ch}\rangle\,(p_\perp)$ distribution was slightly too high [23]. This slight discrepancy

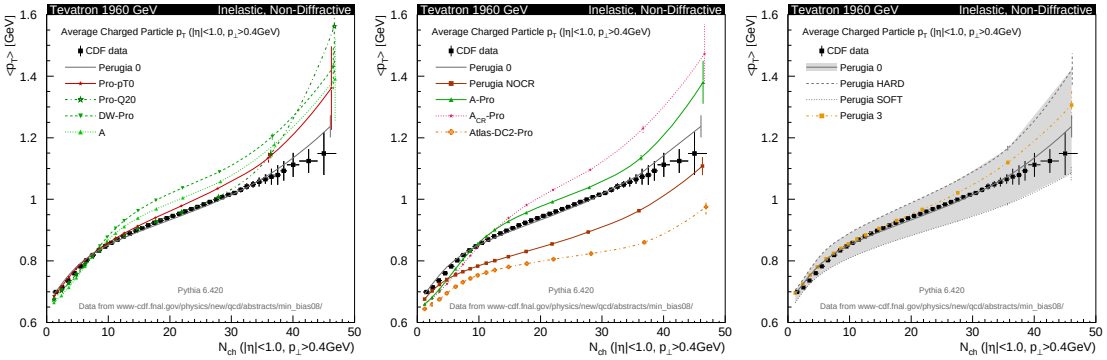

Fig. 4: Comparisons to the CDF Run II measurement of the average track $p_\perp$ as a function of track multiplicity in min-bias $p\bar{p}$ collisions. *Left:* a representative selection of models. *Center:* the impact of varying models of color (re-)connections on this distribution. *Right:* the range spanned by the main Perugia variations. The SOFT and HARD variations were here allowed to deviate by significantly more than the statistical precision due to the high sensitivity of the distribution and the large theoretical uncertainties. Results with more tunes can be found at [17].

carried over to the S0 family of tunes of the new framework, since these were tuned to Tune A, in the absence of published data. Fortunately, CDF data has now been made publicly available [23], and hence it was possible to take the actual data into consideration for the Perugia tunes, resulting in somewhat softer particle spectra in high-multiplicity events, cf. fig. 4. Note that this distribution is highly sensitive to the colour structure of the events, as emphasized in [27, 28, 35, 39].

Finally, the old framework did not include showering off the MPI in- and out-states[2]. The new framework does include such showers, which furnishes an additional fluctuating physics component. Relatively speaking, the new framework therefore needs *less* fluctuations from other sources in order to describe the same data. This is reflected in the tunes of the new framework generally having a less lumpy proton (smoother proton transverse density distributions) and fewer total numbers of MPI than the old one. We included illustrations of this in a special "theory" section of the web plots, cf. [17, Theory Plots] and [16, Fig. 4].

The showers off the MPI also lead to a greater degree of decorrelation and $p_\perp$ imbalance between the minijets produced by the underlying event, in contrast to the old framework where these remained almost exactly balanced and back-to-back. This should show up in minijet $\Delta\phi_{jj}$ and/or ΔR_{jj} distributions sensitive to the underlying event, such as in Z/W+jets with low $p_\perp$ cuts on the additional jets.

Further, since showers tend to produce shorter-range correlations than MPI, the new tunes also exhibit smaller long-range correlations than the old models. I.e., if there is a large fluctuation in one end of the detector, it is *less* likely in the new models that there is a large fluctuation in the same direction in the other end of the detector. The impact of this, if any, on the overall modeling

[2] It did, of course, include showers off the primary interaction. S. Mrenna has since implemented FSR off the MPI as an additional option in that framework, but tunes using that option have not yet been made.

and correction procedures derived from it, has not yet been studied. At the very least it furnishes a systematic difference between the models. For brevity, we do not include the plots here but refer to the web [17, FB Correlation] and to the original PYTHIA MPI paper for a definition and comparable plots [35].

3 Tune-by-Tune Descriptions

The starting point for all the Perugia tunes, apart from Perugia NOCR, was S0(A)-Pro, i.e., the original tunes S0 and S0A, revamped to include the Professor tuning of flavour and fragmentation parameters to LEP data [1]. The starting point for Perugia NOCR was NOCR-Pro. From these starting points, the main hadron collider parameters were retuned to better describe the above mentioned data sets. An overview of the tuned parameters and their values is given in table 1.

Perugia 0 (320): Uses $\Lambda_{\rm CMW}$ instead of $\Lambda_{\overline{\rm MS}}$, which results in near-perfect agreement with the Drell-Yan $p_\perp$ spectrum, both in the tail and in the peak, cf. fig. 1, middle plot. Also has slightly less colour reconnections, especially among high-$p_\perp$ string pieces, which improves the agreement both with the $\langle p_\perp \rangle\,(N_{\rm ch})$ distribution and with the high-$p_\perp$ tail of charged particle $p_\perp$ spectra, cf [17, dN/dpT (tail)]). Compared to S0A-Pro, this tune also has slightly more beam-remnant breakup (more baryon number transport), mostly in order to explore this possibility than due to any necessity of tuning. Without further changes, these modifications would lead to a greatly increased average multiplicity as well as larger multiplicity fluctuations. To keep the total multiplicity unchanged, cf. the solid grey curves labeled "Perugia 0" on the plots in the top row of fig. 2, the changes above were accompanied by an increase in the MPI infrared cutoff, which decreases the overall MPI-associated activity, and by a slightly smoother proton mass profile, which decreases the fluctuations. Finally, the energy scaling is closer to that of S0A than to the default one used for S0, cf. the middle panes in figs. 2 and 3.

Perugia HARD (321): Variant of Perugia 0 which has a higher amount of activity from perturbative physics and counter-balances that partly by having less particle production from nonperturbative sources. Thus, the $\Lambda_{\rm CMW}$ value is used for ISR, together with a renormalisation scale for ISR of $\mu_R = \frac{1}{2}p_\perp$, yielding a comparatively hard Drell-Yan $p_\perp$ spectrum, cf. the dashed curve labeled "HARD" in the right pane of fig. 1. It also has a slightly larger phase space for both ISR and FSR, uses higher-than-nominal values for FSR, and has a slightly harder hadronisation. To partly counter-balance these choices, it has less "primordial k_T", a higher infrared cutoff for the MPI, and more active color reconnections, yielding a comparatively high curve for $\langle p_\perp \rangle\,(N_{\rm ch})$, cf. fig. 4.

Perugia SOFT (322): Variant of Perugia 0 which has a lower amount of activity from perturbative physics and makes up for it partly by adding more particle production from nonperturbative sources. Thus, the $\Lambda_{\overline{\rm MS}}$ value is used for ISR, together with a renormalisation scale of $\mu_R = \sqrt{2}p_\perp$, yielding a comparatively soft Drell-Yan $p_\perp$ spectrum, cf. the dotted curve labeled "SOFT" in the right pane of fig. 1. It also has a slightly smaller phase space for both ISR and FSR, uses lower-than-nominal values for FSR, and has a slightly softer hadronisation. To partly

Parameter	Type	S0A-Pro	P-0	P-HARD	P-SOFT	P-3	P-NOCR	P-X	P-6
MSTP(51)	PDF	7	7	7	7	7	7	20650	10042
MSTP(52)	PDF	1	1	1	1	1	1	2	2
MSTP(64)	ISR	2	3	3	2	3	3	3	3
PARP(64)	ISR	1.0	1.0	0.25	2.0	1.0	1.0	2.0	1.0
MSTP(67)	ISR	2	2	2	2	2	2	2	2
PARP(67)	ISR	4.0	1.0	4.0	0.5	1.0	1.0	1.0	1.0
MSTP(70)	ISR	2	2	0	1	0	2	2	2
PARP(62)	ISR	-	-	1.25	-	1.25	-	-	-
PARP(81)	ISR	-	-	-	1.5	-	-	-	-
MSTP(72)	ISR	0	1	1	0	2	1	1	1
PARP(71)	FSR	4.0	2.0	4.0	1.0	2.0	2.0	2.0	2.0
PARJ(81)	FSR	0.257	0.257	0.3	0.2	0.257	0.257	0.257	0.257
PARJ(82)	FSR	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8
MSTP(81)	UE	21	21	21	21	21	21	21	21
PARP(82)	UE	1.85	2.0	2.3	1.9	2.2	1.95	2.2	1.95
PARP(89)	UE	1800	1800	1800	1800	1800	1800	1800	1800
PARP(90)	UE	0.25	0.26	0.30	0.24	0.32	0.24	0.23	0.22
MSTP(82)	UE	5	5	5	5	5	5	5	5
PARP(83)	UE	1.6	1.7	1.7	1.5	1.7	1.8	1.7	1.7
MSTP(88)	BR	0	0	0	0	0	0	0	0
PARP(79)	BR	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0
PARP(80)	BR	0.01	0.05	0.01	0.05	0.03	0.01	0.05	0.05
MSTP(91)	BR	1	1	1	1	1	1	1	1
PARP(91)	BR	2.0	2.0	1.0	2.0	1.5	2.0	2.0	2.0
PARP(93)	BR	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
MSTP(95)	CR	6	6	6	6	6	6	6	6
PARP(78)	CR	0.2	0.33	0.37	0.15	0.35	0.0	0.33	0.33
PARP(77)	CR	0.0	0.9	0.4	0.5	0.6	0.0	0.9	0.9
MSTJ(11)	HAD	5	5	5	5	5	5	5	5
PARJ(21)	HAD	0.313	0.313	0.34	0.28	0.313	0.313	0.313	0.313
PARJ(41)	HAD	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49
PARJ(42)	HAD	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
PARJ(46)	HAD	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
PARJ(47)	HAD	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

Table 1: Parameters of the Perugia tunes, omitting the LEP flavour parameters tuned by Professor [1] (common to all the “Pro” and “Perugia” tunes). The starting point, S0A-Pro, is shown for reference. (BR stands for Beam Remnants and CR stands for Colour Reconnections.)

counter-balance these choices, it has a more sharply peaked proton mass distribution, a more active beam remnant fragmentation (lots of baryon transport), a slightly lower infrared cutoff for the MPI, and slightly less active color reconnections, yielding a comparatively low curve for $\langle p_\perp \rangle\,(N_{\rm ch})$, cf. fig. 4.

Perugia 3 (323): Variant of Perugia 0 which has a different balance between MPI and ISR and a different energy scaling. Instead of a smooth dampening of ISR all the way to zero $p_\perp$, this tune uses a sharp cutoff at 1.25 GeV, which produces a slightly harder ISR spectrum. The additional ISR activity is counter-balanced by a higher infrared MPI cutoff. Since the ISR cutoff is independent of the collider CM energy in this tune, the multiplicity would nominally evolve very rapidly with energy. To offset this, the MPI cutoff itself must scale very quickly, hence this tune has a very large value of the scaling power of that cutoff. This leads to an interesting systematic difference in the scaling behaviour, with ISR becoming an increasingly more important source of particle production as the energy increases in this tune, relative to Perugia 0.

Perugia NOCR (324): An update of NOCR-Pro that attempts to fit the data sets as well as possible, without invoking any explicit colour reconnections. Can reach an acceptable agreement with most distributions, except for the $\langle p_\perp \rangle\,(N_{\rm ch})$ one, cf. fig. 4.

Perugia X (325): A Variant of Perugia 0 which uses the MRST LO* PDF set [40]. Due to the increased gluon densities, a slightly lower ISR renormalisation scale and a higher MPI cutoff than for Perugia 0 is used. Note that, since we are not yet sure the implications of using LO* for the MPI interactions have been fully understood, this tune should be considered experimental for the time being. See [17, Perugia PDFs] for distributions.

Perugia 6 (326): A Variant of Perugia 0 which uses the CTEQ6L1 PDF set [41]. Identical to Perugia 0 in all other respects, except for a slightly lower MPI infrared cutoff at the Tevatron and a lower scaling power of the MPI infrared cutoff. See [17, Perugia PDFs] for distributions.

4 Extrapolation to the LHC

Part of the motivation for updating the S0 family of tunes was specifically to improve the constraints on the energy scaling to come up with tunes that extrapolate more reliably to the LHC. This is not to say that the uncertainty is still not large, but as mentioned above, it does seem that, e.g., the default PYTHIA scaling has by now been convincingly ruled out, and so this is naturally reflected in the updated parameters.

Fig. 5 contains predictions for the Drell-Yan $p_\perp$ distribution (using the CDF cuts), the charged track multiplicity distribution in minimum-bias collisions, and the average track $p_\perp$ as a function of multiplicity at 14 TeV, for the central, hard, soft, and "3" variations of the Perugia tunes. We hope this helps to give a feeling for the kind of ranges spanned by the Perugia tunes (the PDF variations give almost identical results to Perugia 0 for these distributions). A full set of plots illustrating the extrapolations to the LHC for both the central region $|\eta| < 2.5$ as well as the region $1.8 < \eta < 4.9$ covered by LHCb can be found on the web [17].

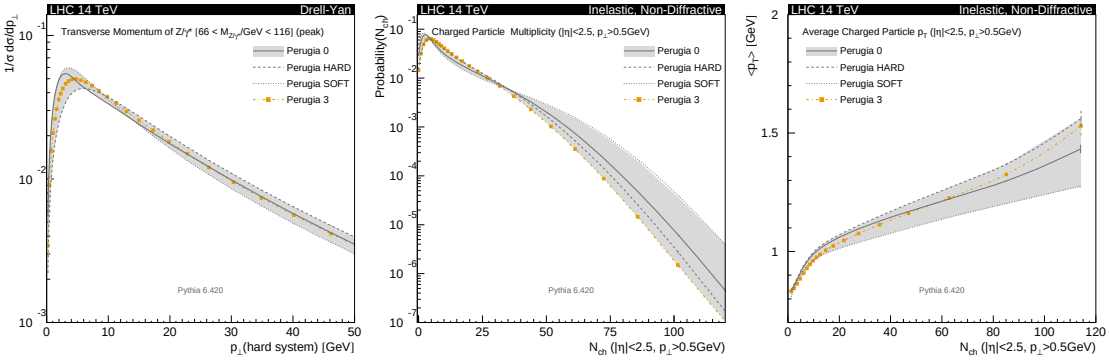

Fig. 5: Perugia "predictions" for the $p_\perp$ of Drell-Yan pairs (left), the charged track multiplicity in min-bias (center), and the average track $p_\perp$ in min-bias (right) at the LHC. See [17] for additional plots.

However, in addition to these plots, we thought it would be interesting to make at least one set of numerical predictions for an infrared sensitive quantity that could be tested with the very earliest LHC data. We therefore used the Perugia tunes and their variations to get an estimate for the mean multiplicity of charged tracks in (inelastic, nondiffractive) minimum-bias pp collisions at 10 and 14 TeV. The Perugia variations indicate an uncertainty of order 15% or less on the central values, which is probably an underestimate, due to the limited nature of the models. Nonetheless, having spent a significant amount of effort in making these estimates, given in tab. 2, we intend to stick by them until proved wrong. The acknowledgments therefore contain a recognition of a bet to that effect.

5 Conclusions

We have presented a set of updated parameter sets (tunes) for the interleaved $p_\perp$-ordered shower and underlying-event model in PYTHIA 6.4. These parameter sets include the revisions to the fragmentation and flavour parameters obtained by the Professor group and reported on elsewhere in these proceedings [1]. The new sets further include more Tevatron data and more data from different collider CM energies in an attempt to simultaneously improve the overall description at the Tevatron data while also improving the reliability of the extrapolations to the LHC. We have also attempted to deliver a first set of "tunes with uncertainty bands", by including alternative tunes with systematically different parameter choices. The new tunes are available from Pythia version 6.4.20, via the routine PYTUNE.

We note that these tunes still only included Drell-Yan and minimum-bias data directly; leading-jet, photon+jet, and underlying-event data was not considered explicitly. This is not expected to be a major problem due to the good universality properties that the PYTHIA modeling has so far exhibited, but it does mean that the performance of the tunes on such data sets should be tested, which will hopefully happen in the near future.

We hope these tunes will be useful to the RHIC, Tevatron, and LHC communities.

Predictions for Mean Densities of Charged Tracks

	$\frac{\langle N_{\rm ch}\rangle\,\vert_{N_{\rm ch}\geq 0}}{\Delta\eta\Delta\phi}$	$\frac{\langle N_{\rm ch}\rangle\,\vert_{N_{\rm ch}\geq 1}}{\Delta\eta\Delta\phi}$	$\frac{\langle N_{\rm ch}\rangle\,\vert_{N_{\rm ch}\geq 2}}{\Delta\eta\Delta\phi}$	$\frac{\langle N_{\rm ch}\rangle\,\vert_{N_{\rm ch}\geq 3}}{\Delta\eta\Delta\phi}$	$\frac{\langle N_{\rm ch}\rangle\,\vert_{N_{\rm ch}\geq 4}}{\Delta\eta\Delta\phi}$
LHC 10 TeV	0.40 ± 0.05	0.41 ± 0.05	0.43 ± 0.05	0.46 ± 0.06	0.50 ± 0.06
LHC 14 TeV	0.44 ± 0.05	0.45 ± 0.06	0.47 ± 0.06	0.51 ± 0.06	0.54 ± 0.07

Table 2: Best-guess predictions for the mean density of charged tracks for min-bias pp collisions at two LHC energies. These numbers should be compared to data corrected to 100% track finding efficiency for tracks with $|\eta| < 2.5$ and $p_\perp > 0.5$ GeV and 0% efficiency outside that region. The definition of a stable particle was set at $c\tau \geq 10$mm (e.g., the two tracks from a $\Lambda^0 \to p^+\pi^-$ decay were not counted). The $\pm$ values represent the estimated uncertainty, based on the Perugia tunes. Since the lowest multiplicity bins may receive large corrections from elastic/diffractive events, it is possible that it will be easier to compare the (inelastic nondiffractive) theory to the first data with one or more of the lowest multiplicity bins excluded, hence we have here recomputed the means with up to the first 4 bins excluded. (These predictions were first shown at the 2009 Aspen Winter Conference.)

Acknowledgments

We are grateful to the organizers of this very enjoyable workshop, which brought people from different communities together, and helped us take some steps towards finding a common language. In combination with the writeup of this article, I agreed to owing Lisa Randall a bottle of champagne if the first published measurement of any number in table 2 is outside the range given. Conversely, she agrees to owing me a bottle if the corresponding number happens to be right.

References

[1] H. Hoeth, these proceedings.

[2] AFS, T. Akesson *et al.*, Z. Phys. **C34**, 163 (1987).

[3] UA1, C.-E. Wulz, in proceedings of the 22nd Rencontres de Moriond, Les Arcs, France, 15-21 March 1987.

[4] UA2, J. Alitti *et al.*, Phys. Lett. **B268**, 145 (1991).

[5] CDF, F. Abe *et al.*, Phys. Rev. **D47**, 4857 (1993).

[6] CDF, F. Abe *et al.*, Phys. Rev. Lett. **79**, 584 (1997).

[7] CDF, F. Abe *et al.*, Phys. Rev. **D56**, 3811 (1997).

[8] DØ, V. M. Abazov *et al.*, Phys. Rev. **D67**, 052001 (2003), hep-ex/0207046.

[9] D. Bandurin *et al.*, A study of γ + 3 jet events at DØ, in progress.

[10] ZEUS, C. Gwenlan, Acta Phys. Polon. **B33**, 3123 (2002).

[11] T. Sjöstrand and P. Z. Skands, Eur. Phys. J. **C39**, 129 (2005), hep-ph/0408302.

[12] T. Sjostrand and P. Z. Skands, JHEP **03**, 053 (2004), hep-ph/0402078.

[13] T. Sjöstrand, S. Mrenna, and P. Skands, JHEP **05**, 026 (2006), hep-ph/0603175.

[14] T. Sjöstrand, S. Mrenna, and P. Skands, Comput. Phys. Commun. **178**, 852 (2008), 0710.3820.

[15] Y. I. Azimov, Y. L. Dokshitzer, V. A. Khoze, and S. I. Troyan, Z. Phys. **C27**, 65 (1985).

[16] P. Z. Skands, in C. Buttar et al., arXiv:0803.0678 [hep-ph].

[17] P. Skands, Peter's pythia plots, see `http://home.fnal.gov/~skands/leshouches-plots/`.

[18] A. Buckley, H. Hoeth, H. Schulz, and J. E. von Seggern, (2009), arXiv:0902.4403 [hep-ph].

[19] CDF, A. A. Affolder *et al.*, Phys. Rev. Lett. **84**, 845 (2000), hep-ex/0001021.

[20] D0, V. M. Abazov *et al.*, Phys. Rev. Lett. **100**, 102002 (2008), 0712.0803.

[21] CDF, D. E. Acosta *et al.*, Phys. Rev. **D65**, 072005 (2002).

[22] CDF, F. Abe *et al.*, Phys. Rev. Lett. **61**, 1819 (1988).

[23] N. Moggi, Inclusive pp differential cross-sections, see `http://www-cdf.fnal.gov/physics/new/qcd/abstracts/min_bias08/publicpage.html`, 2008.

[24] T. Alexopoulos *et al.*, Phys. Lett. **B435**, 453 (1998).

[25] UA5, G. J. Alner *et al.*, Phys. Rept. **154**, 247 (1987).

[26] UA5, R. E. Ansorge *et al.*, Z. Phys. **C43**, 357 (1989).

[27] M. Sandhoff and P. Skands, presented at Les Houches Workshop on Physics at TeV Colliders, Les Houches, France, 2-20 May 2005, in hep-ph/0604120.

[28] P. Skands and D. Wicke, Eur. Phys. J. **C52**, 133 (2007), hep-ph/0703081.

[29] R. D. Field, hep-ph/0201192 CDF Note 6403; further recent talks available from webpage `http://www.phys.ufl.edu/~rfield/cdf/`.

[30] R. Field and R. C. Group, (2005), hep-ph/0510198.

[31] CDF, R. Field, AIP Conf. Proc. **828**, 163 (2006).

[32] TeV4LHC QCD Working Group, M. G. Albrow *et al.*, (2006), hep-ph/0610012.

[33] S. Alekhin *et al.*, (2005), hep-ph/0601012.

[34] C. Buttar *et al.*, (2008), 0803.0678.

[35] T. Sjöstrand and M. van Zijl, Phys. Rev. **D36**, 2019 (1987).

[36] S. Catani, B. R. Webber, and G. Marchesini, Nucl. Phys. **B349**, 635 (1991).

[37] G. Corcella *et al.*, JHEP **01**, 010 (2001), hep-ph/0011363.

[38] M. Bähr *et al.*, Eur. Phys. J. **C58**, 639 (2008), 0803.0883.

[39] D. Wicke and P. Z. Skands, (2008), 0807.3248.

[40] A. Sherstnev and R. S. Thorne, Eur. Phys. J. **C55**, 553 (2008), 0711.2473.

[41] J. Pumplin *et al.*, JHEP **07**, 012 (2002), hep-ph/0201195.