

41. MONTE CARLO EVENT GENERATORS

Revised September 2013 by P. Nason (INFN, Milan) and P.Z. Skands (Monash University).

General-purpose Monte Carlo (GPMC) generators like HERWIG [1], HERWIG++ [2], PYTHIA 6 [3], PYTHIA 8 [4], and SHERPA [5], provide fully exclusive simulations of high-energy collisions. They play an essential role in QCD modeling (in particular for aspects beyond fixed-order perturbative QCD), in data analysis, where they are used together with detector simulation to provide a realistic estimate of the detector response to collision events, and in the planning of new experiments, where they are used to estimate signals and backgrounds in high-energy processes. They are built from several components, that describe the physics starting from very short distance scales, up to the typical scale of hadron formation and decay. Since QCD is weakly interacting at short distances (below a femtometer), the components of the GPMC dealing with short-distance physics are based upon perturbation theory. At larger distances, all soft hadronic phenomena, like hadronization and the formation of the underlying event, cannot be computed from first principles, and one must rely upon QCD-inspired models.

The purpose of this review is to illustrate the main components of these generators. It is divided into four sections. The first one deals with short-distance, perturbative phenomena. The basic concepts leading to the simulations of the dominant QCD processes are illustrated here. In the second section, hadronization phenomena are treated. The two most popular hadronization models for the formation of primary hadrons, the string and cluster models, are illustrated. The basics of the implementation of primary-hadron decays into stable ones is also illustrated here. In the third section, models for soft hadron physics are discussed. These include models for the underlying event, and for minimum-bias interactions. Issues of Bose-Einstein and color-reconnection effects are also discussed here. The fourth section briefly introduces the problem of MC tuning.

We use natural units throughout, such that $c = 1$ and $\hbar = 1$, with energy, momenta and masses measured in GeV, and time and distances measured in GeV^{-1} .

41.1. Short-distance physics in GPMC generators

The short-distance components of a GPMC generator deal with the computation of the primary process at hand, with decays of short-lived particles, and with the generation of QCD and QED radiation, on time scales below $1/\Lambda$, with Λ denoting a typical hadronic scale of a few hundred MeV, corresponding roughly to an inverse femtometer. In e^+e^- annihilation, for example, the short-distance physics describes the evolution of the system from the instant when the e^+e^- pair annihilates up to a time when the size of the produced system is just below a femtometer.

In the present discussion we take the momentum scale of the primary process to be $Q \gg \Lambda$, so that the corresponding time and distance scale $1/Q$ is small. Soft- and collinear-safe inclusive observables, such as total decay widths or inclusive cross sections, can then be reliably computed in QCD perturbation theory (pQCD), with the perturbative expansion truncated at any fixed order n , and the remainder suppressed by $\alpha_S(Q)^{n+1}$.

Less inclusive observables, however, can receive large enhancements that destroy the convergence of the fixed-order expansion. This is due to the presence of collinear and infrared singularities in QCD. Thus, for example, a correction in which a parton from the primary interaction splits collinearly into two partons of comparable energy, is of order $\alpha_S(Q) \ln(Q/\Lambda)$, where the logarithm arises from an integral over a singularity regulated by the hadronic scale Λ . Since $\alpha_S(Q) \propto 1/\ln(Q/\Lambda)$, the corresponding cross section receives a correction of order unity. Two subsequent collinear splittings yield $\alpha_S^2(Q) \ln^2(Q/\Lambda)$, and so on. Thus, corrections of order unity arise at all orders in perturbation theory. The dominant region of phase space is the one where radiation is strongly ordered in a measure of hardness. This means that, from a typical final-state configuration, by clustering together final-state parton pairs with the smallest hardness recursively, we can reconstruct a branching tree, that may be viewed as the splitting history of the event. This history necessarily has some dependence on how we define hardness. For example, we can define it as the energy of the incoming parton times the splitting angle, or as its virtuality, or as the transverse momentum of the splitting partons with respect to the incoming one. These definitions, however, are all equivalent in the collinear region. In fact, in the small-angle limit, the virtuality of a parton of energy E , splitting into two on-shell partons, is given by

$$p^2 = E^2 z(1-z)(1-\cos\theta) \approx \frac{z(1-z)}{2} E^2 \theta^2, \quad (41.1)$$

where z and $1-z$ are the energy fractions carried by the produced partons, and θ is their relative angle. The transverse momentum of the final partons relative to the direction of the incoming one is given by

$$p_T^2 \approx z^2(1-z)^2 E^2 \theta^2. \quad (41.2)$$

Thus, significant differences between these measures only arise in regions with very small z or $1-z$ values. In QCD, because of soft divergences, these regions are in fact important, and the choice of the appropriate ordering variable is very relevant (see Sec. 41.3).

The so called KLN theorem [6,7] guarantees that large logarithmically divergent corrections, arising from final-state collinear splitting and from soft emissions, cancel against the virtual corrections in the total cross section, order by order in perturbation theory. Furthermore, the factorization theorem guarantees that initial-state collinear singularities can be factorized into the parton density functions (PDFs). Therefore, the cross section for the basic process remains accurate up to corrections of higher orders in $\alpha_S(Q)$, provided it is interpreted as an inclusive cross section, rather than as a bare partonic cross section. Thus, for example, the leading order (LO) cross section for $e^+e^- \rightarrow q\bar{q}$ is a good LO estimate of the e^+e^- cross section for the production of a pair of quarks accompanied by an arbitrary number of collinear and soft gluons, but is not a good estimate of the cross section for the production of a $q\bar{q}$ pair with no extra radiation.

Shower algorithms are used to compute the cross section for generic hard processes including all leading-logarithmic (LL) corrections. These algorithms begin with the generation of the kinematics of the basic process, performed with a probability proportional to its LO partonic cross section, which is interpreted physically as the

inclusive cross section for the basic process, followed by an arbitrary sequence of shower splittings. A probability is then assigned to each splitting sequence. Thus, the initial LO cross section is partitioned into the cross sections for a multitude of final states of arbitrary multiplicity. The sum of all these partial cross sections equals that of the primary process. This property of the GPMCs reflects the KLN cancellation mentioned earlier, and it is often called “unitarity of the shower process”, a name that reminds us that the KLN cancellation itself is a consequence of unitarity. The fact that a quantum mechanical process can be described in terms of composition of probabilities, rather than amplitudes, follows from the LL approximation. In fact, in the dominant, strongly ordered region, subsequent splittings are separated by increasingly large times and distances, and this suppresses interference effects.

We now illustrate the basic parton-shower algorithm, as first introduced in Ref. 8. The purpose of this illustration is to give a schematic representation of how shower algorithms work, to introduce some concepts that will be referred to in the following, and to show the relationship between shower algorithms and Feynman-diagram results. For simplicity, we consider the example of e^+e^- annihilation into $q\bar{q}$ pairs. With each dominant (i.e. strongly ordered) final-state configuration one can associate an ordered tree diagram, by recursively clustering together final-state parton pairs with the smallest hardness, and ending up with the hard production vertex (i.e. the $\gamma^* \rightarrow q\bar{q}$). The momenta of all intermediate lines of the tree diagram are then uniquely determined from the final-state momenta. Hardnesses in the graph are also strongly ordered. One assigns to each splitting vertex the hardness t , the energy fractions z and $1 - z$ of the two generated partons, and the azimuth ϕ of the splitting process with respect to the momentum of the incoming parton. For definiteness, we assume that z and ϕ are defined in the center-of-mass (CM) frame of the e^+e^- collision, although other definitions are possible that differ only beyond the LL approximation. The differential cross section for a given final state is given by the product of the differential cross section for the initial $e^+e^- \rightarrow q\bar{q}$ process, multiplied by a factor

$$\Delta_i(t, t') \frac{\alpha_S(t)}{2\pi} P_{i,jk}(z) \frac{dt}{t} dz \frac{d\phi}{2\pi} \quad (41.3)$$

for each intermediate line ending in a splitting vertex. We have denoted with t' the maximal hardness that is allowed for the line, with t its hardness, and z and ϕ refer to the splitting process. $\Delta(t, t')$ is the so-called Sudakov form factor

$$\Delta_i(t, t') = \exp \left[- \int_t^{t'} \frac{dq^2}{q^2} \frac{\alpha_S(q^2)}{2\pi} \sum_{jk} P_{i,jk}(z) dz \frac{d\phi}{2\pi} \right]. \quad (41.4)$$

The suffixes i and jk represent the parton species of the incoming and final partons, respectively, and $P_{i,jk}(z)$ are the Altarelli-Parisi [9] splitting kernels. Final-state lines that do not undergo any further splitting are associated with a factor

$$\Delta_i(t_0, t'), \quad (41.5)$$

where t_0 is an infrared cutoff defined by the shower hadronization scale (at which the charges are screened by hadronization) or, for an unstable particle, its width (a source cannot emit radiation with a period exceeding its lifetime).

Notice that the definition of the Sudakov form factor is such that

$$\Delta_i(t_2, t_1) + \int_{t_2}^{t_1} \frac{dt}{t} dz \frac{d\phi}{2\pi} \sum_{jk} \Delta_i(t, t_1) \frac{\alpha_S(t)}{2\pi} P_{i,jk}(z) = 1. \quad (41.6)$$

This implies that the cross section for developing the shower up to a given stage does not depend on what happens next, since subsequent factors for further splitting or not splitting add up to one.

The shower cross section can then be formulated in a probabilistic way. The Sudakov form factor $\Delta_i(t_2, t_1)$ is interpreted as the probability for a splitting *not* to occur, for a parton of type i , starting from a branching vertex at the scale t_1 , down to a scale t_2 . Notice that $0 < \Delta_i(t_2, t_1) \leq 1$, where the upper extreme is reached for $t_2 = t_1$, and the lower extreme is approached for $t_2 = t_0$. From Eq. (41.4), it seems that the Sudakov form factor should vanish if $t_2 = 0$. However, because of the presence of the running coupling in the integrand, t_2 cannot be taken smaller than some cutoff scale of the order of Λ , so that at its lower extreme the Sudakov form factor is small, but not zero. Event generation then proceeds as follows. One gets a uniform random number $0 \leq r \leq 1$, and seeks a solution of the equation $r = \Delta_i(t_2, t_1)$ as a function of t_2 . If r is too small and no solution exists, no splitting is generated, and the line is interpreted as a final parton. If a solution t_2 exists, a branching is generated at the scale t_2 . Its z value and the final parton species jk are generated with a probability proportional to $P_{i,jk}(z)$. The azimuth is generated uniformly. This procedure is started with each of the primary process partons, and is applied recursively to all generated partons. It may generate an arbitrary number of partons, and it stops when no final-state partons undergo further splitting.

The four-momenta of the final-state partons are reconstructed from the momenta of the initiating ones, and from the whole sequence of splitting variables, subject to overall momentum conservation. Different algorithms employ different strategies to treat recoil effects due to momentum conservation, which may be applied either locally for each parton or dipole splitting, or globally for the entire set of partons (a procedure called *momentum reshuffling*.) This has a subleading effect with respect to the collinear approximation.

We emphasize that the shower cross section described above can be derived from perturbative QCD by keeping only the collinear-dominant real and virtual contributions to the cross section. In particular, up to terms that vanish after azimuthal averaging, the product of the cross section for the basic process, times the factors

$$\frac{\alpha_S}{2\pi} \frac{dt}{t} dz \frac{d\phi}{2\pi} P_{i,jk}(z) \quad (41.7)$$

at each branching vertex, gives the leading collinear contribution to the tree-level cross section for the same process. The dominant virtual corrections in the same approximation are provided by the running coupling at each vertex and by the Sudakov form factors in the intermediate lines.

41.1.1. Angular correlations :

In gluon splitting processes ($g \rightarrow q\bar{q}$, $g \rightarrow gg$) in the collinear approximation, the distribution of the split pair is not uniform in azimuth, and the Altarelli-Parisi splitting functions are recovered only after azimuthal averaging. This dependence is due to the interference of positive and negative helicity states for the gluon that undergoes splitting. Spin correlations propagate through the splitting process, and determine acausal correlations of the EPR kind [10]. A method to partially account for these effects was introduced in Ref. 11, in which the azimuthal correlation between two successive splittings is computed by averaging over polarizations. This can then be applied at each branching step. Acausal correlations are argued to be small, and are discarded with this method, that is still used in the PYTHIA code [3]. A method that fully includes spin correlation effects was later proposed by Collins [12], and has been implemented in the fortran HERWIG code [13].

41.1.2. Initial-state radiation :

Initial-state radiation (ISR) arises because incoming charged particles can radiate before entering the hard-scattering process. In doing so, they acquire a non-vanishing transverse momentum, and their virtuality becomes negative (spacelike). The dominant logarithmic region is the collinear one, where virtualities become larger and larger in absolute value with each emission, up to a limit given by the hardness of the basic process itself. A shower that starts by considering the highest virtualities first would thus have to work backward in time for ISR. A corresponding backwards-evolution algorithm was formulated by Sjöstrand [14], and was basically adopted in all shower models.

The key point in backwards evolution is that the evolution probability depends on the amount of partons that could have given rise to the one being evolved. This is reflected by introducing the ratio of the PDF after the branching to the PDF before the branching in the definition of the backward-evolution Sudakov form factor,

$$\Delta_i^{\text{ISR}}(t, t') = \exp \left[- \int_{t'}^t \frac{dt''}{t''} \frac{\alpha_S(t'')}{2\pi} \int_x^1 \frac{dz}{z} \sum_{jk} P_{j,ik}(z) \frac{f_j(t'', x/z)}{f_i(t'', x)} \right]. \quad (41.8)$$

Notice that there are two uses of the PDFs: they are used to compute the cross section for the basic hard process, and they control ISR via backward evolution. Since the evolution is generated with leading-logarithmic accuracy, it is acceptable to use two different PDF sets for these two tasks, provided they agree at the LO level.

In the context of GPMC evolution, each ISR emission generates a finite amount of transverse momentum. Details on how the recoils generated by these transverse “kicks” are distributed among other partons in the event, in particular the ones involved in the hard process, constitute one of the main areas of difference between existing algorithms, see Ref. 15. An additional $\mathcal{O}(1 \text{ GeV})$ of “primordial k_T ” is typically added, to represent the sum of unresolved and/or non-perturbative motion below the shower cutoff scale.

41.1.3. *Soft emissions and QCD coherence* :

In massless field theories like QCD, there are two sources of large logarithms of infrared origin. One has to do with collinear singularities, which arise when two final-state particles become collinear, or when a final-state particle becomes collinear to an initial-state one. The other has to do with the emission of soft gluons at arbitrary angles. Because of that, it turns out that in QCD perturbation theory two powers of large logarithms can arise for each power of α_S . The expansion in leading soft and collinear logarithms is often referred to as the double-logarithmic expansion.

Within the conventional parton-shower formalism, based on collinear factorization, it was shown in a sequel of publications (see Ref. 16 and references therein) that the double-logarithmic region can be correctly described by using the angle of the emissions as the ordering variable, rather than the virtuality, and that the argument of α_S at the splitting vertex should be the relative parton transverse momentum after the splitting. Physically, the ordering in angle approximates the coherent interference arising from large-angle soft emission from a bunch of collinear partons. Without this effect, the particle multiplicity would grow too rapidly with energy, in conflict with e^+e^- data. For this reason, angular ordering is used as the evolution variable in both the HERWIG [16] and HERWIG++ [17] programs, and an angular veto is imposed on the virtuality-ordered evolution in PYTHIA 6 [18].

A radical alternative formulation of QCD cascades first proposed in Ref. 19 focuses upon soft emission, rather than collinear emission, as the basic splitting mechanism. It then becomes natural to consider a branching process where it is a parton pair (i.e. a dipole) rather than a single parton, that emits a soft parton. Adding a suitable correction for non-soft, collinear partons, one can achieve in this framework the correct logarithmic structure for both soft and collinear emissions in the limit of large number of colors N_c , without any explicit angular-ordering requirement. The ARIADNE [20] and VINCIA [21] programs are based on this approach. In SHERPA, the default shower [22] is also of a dipole type [23], while the $p_\perp$ -ordered showers in PYTHIA 6 and 8 represent a hybrid, combining collinear splitting kernels with dipole kinematics [24].

41.1.4. *Massive quarks* :

Quark masses act as cut-off on collinear singularities. If the mass of a quark is below, or of the order of Λ , its effect in the shower is small. For larger quark masses, like in c , b , or t production, it is the mass, rather than the typical hadronic scale, that cuts off collinear radiation. For a quark with energy E and mass m_Q , the divergent behavior $d\theta/\theta$ of the collinear splitting process is regulated for $\theta \leq \theta_0 = m_Q/E$. We thus expect less collinear activity for heavy quarks than for light ones, which in turn is the reason why heavy quarks carry a larger fraction of the momentum acquired in the hard production process.

This feature can be implemented with different levels of sophistication. Using the fact that soft emission exhibits a zero at zero emission angle, older parton shower algorithms simply limited the shower emission to be not smaller than the angle θ_0 . More modern approaches are used in both PYTHIA, where mass effects are included using a kind of matrix-element correction method [25], and in HERWIG++ and SHERPA, where a generalization of the Altarelli-Parisi splitting kernel is used for massive quarks [26].

41.1.5. Color information :

Shower MC generators track large- N_c color information during the development of the shower. In the large- N_c limit, quarks or antiquarks are represented by a color line, i.e. a line with an arrow indicating the direction of color flow. Gluons are represented by a pair of color lines with opposite arrows. The rules for color propagation are:



During the shower development, partons are connected by color lines. We can have a quark directly connected by a color line to an antiquark, or via an arbitrary number of intermediate gluons, as shown in Fig. 41.1.




Figure 41.1: Color development of a shower in e^+e^- annihilation. Systems of color-connected partons are indicated by the dashed lines.

It is also possible for a set of gluons to be connected cyclically in color, as e.g. in the decay $\Upsilon \rightarrow ggg$.

The color information is used in angular-ordered showers, where the angle of color-connected partons determines the initial angle for the shower development, and in dipole showers, where dipoles are always color-connected partons. It is also used in hadronization models, where the initial strings or clusters used for hadronization are formed by systems of color-connected partons.

41.1.6. Electromagnetic corrections :

The physics of photon emission from light charged particles can also be treated with a shower MC algorithm. High-energy electrons and quarks, for example, are accompanied by bremsstrahlung photons. Also here, similarly to the QCD case, electromagnetic corrections are of order $\alpha_{\text{em}} \ln(Q/m)$, where m is the mass of the radiating particle, or even of order $\alpha_{\text{em}} \ln(Q/m) \ln(E_\gamma/E)$ in the region where soft photon emission is important, so that, especially for the case of electrons, their inclusion in the simulation process is mandatory. This is done in most of the GPMC's (for a recent comparative study see [27]) . The specialized generator PHOTOS [28] is sometimes used as an afterburner for an improved treatment of QED radiation in non-hadronic resonance decays.

In case of photons emitted by leptons the shower can be continued down to virtualities arbitrarily close to the lepton mass shell (unlike the case in QCD). In practice, photon radiation must be cut off below a certain energy, in order for the shower algorithm to terminate. Therefore, there is always a minimum energy for emitted photons that

depends upon the implementations [27] (and so does the MC truth for a charged lepton). In the case of electrons, this energy is typically of the order of its mass. Electromagnetic radiation below this scale is not enhanced by collinear singularities, and is thus bound to be soft, so that the electron momentum is not affected by it.

For photons emitted from quarks, we have instead the obvious limitation that the photon wavelength cannot exceed the typical hadronic size. Longer-wavelength photons are in fact emitted by hadrons, rather than quarks. This last effect is in practice never modeled by existing shower MC implementations. Thus, electromagnetic radiation from quarks is cut off at a typical hadronic scale. Finally, hadron (and τ) decays involving charged particles can produce additional soft bremsstrahlung. This is implemented in a general way in HERWIG++ [29] and SHERPA [30].

41.1.7. *Beyond-the-Standard-Model Physics :*

The inclusion of processes for physics beyond the Standard Model (BSM) in event generators is to some extent only a matter of implementing the relevant hard processes and (chains of) decays, with the level of difficulty depending on the complexity of the model and the degree of automation [31,32]. Notable exceptions are long-lived colored particles [33], particles in exotic color representations, and particles showering under new gauge symmetries, with a growing set of implementations documented in the individual GPMC manuals. Further complications that may be relevant are finite-width effects (discussed in Sec. 41.1.8) and the assumed threshold behavior.

In addition to code-specific implementations [15], there are a few commonly adopted standards that are useful for transferring information and events between codes. Currently, the most important of these is the Les Houches Event File (LHEF) standard [34], normally used to transfer parton-level events from a hard-process generator to a shower generator. Another important standard is the Supersymmetry Les Houches Accord (SLHA) format [35], originally used to transfer information on supersymmetric particle spectra and couplings, but by now extended to apply also to more general BSM frameworks and incorporated within the LHEF standard [36].

41.1.8. *Decay Chains and Particle Widths :*

In most BSM processes and some SM ones, an important aspect of the event simulation is how decays of short-lived particles, such as top quarks, EW and Higgs bosons, and new BSM resonances, are handled. We here briefly summarize the spectrum of possibilities, but emphasize that there is no universal standard. Users are advised to check whether the treatment of a given code is adequate for the physics study at hand.

The appearance of an unstable resonance as a physical particle at some intermediate stage of the event generation implies that its production and decay processes are treated as being factorized. This is valid up to corrections of order Γ/m_0 , with Γ the width and m_0 the pole mass. States whose widths are a substantial fraction of their mass should not be treated as “physical particles,” but rather as intrinsically off-shell internal propagator lines.

For states treated as physical particles, two aspects are relevant: the mass distribution of the decaying particle itself and the distributions of its decay products. For the

former, matrix-element generators often use a simple δ function at m_0 . The next level up, typically used in GPMCs, is to use a Breit-Wigner distribution (relativistic or non-relativistic), which formally resums higher-order virtual corrections to the mass distribution. Note, however, that this still only generates an improved picture for *moderate* fluctuations away from m_0 . Similarly to above, particles that are significantly off-shell (in units of Γ) should not be treated as resonant, but rather as internal off-shell propagator lines. In most GPMCs, further refinements are included, for instance by letting Γ be a function of m (“running widths”) and by limiting the magnitude of the allowed fluctuations away from m_0 .

For the distributions of the decay products, the simplest treatment is again to assign them their respective m_0 values, with a uniform phase-space distribution. A more sophisticated treatment distributes the decay products according to the differential decay matrix elements, capturing at least the internal dynamics and helicity structure of the decay process, including EPR-like correlations. Further refinements include polarizations of the external states [37] and assigning the decay products their own Breit-Wigner distributions, the latter of which opens the possibility to include also intrinsically off-shell decay channels, like $H \rightarrow WW^*$.

During subsequent showering of the decay products, most parton-shower models will preserve their total invariant mass, so as not to skew the original resonance shape.

When computing partial widths and/or modifying decay tables, one should be aware of the danger of double-counting intermediate on-shell particles, see Sec. 41.2.3.

41.1.9. *Matching with Matrix Elements* :

Shower algorithms are based upon a combination of the collinear (small-angle) and soft (small-energy) approximations and are thus normally inaccurate for hard, wide-angle emissions (i.e., additional well-resolved jets). They also contain only the leading singular pieces of next-to-leading order (NLO) and higher corrections to the basic process.

Traditional GPMCs, like HERWIG and PYTHIA, have included for a long time the so called Matrix Element Corrections (MEC), first formulated in Ref. 38 with later developments summarized in Ref. 15. They are typically available for $2 \rightarrow 1$ or $1 \rightarrow 2$ processes, like DIS, vector boson and Higgs production and decays, and top decays. The MEC corrects the emission of the hardest jet at large angles, so that it becomes exact at LO. A generalization of the method to multiple emissions was formulated recently [39].

Aside from MECs implemented directly in the GPMCs, the improvements on the parton-shower description of hard collisions have been made in two main directions: the so called Matrix Elements and Parton Shower matching (ME+PS from now on), and the matching of NLO calculations and Parton Showers (NLO+PS). We now discuss each of these, and then briefly summarise techniques becoming available for combining them.

The ME+PS method allows one to use tree-level matrix elements for hard, large-angle emissions. It was first formulated in the so-called CKKW paper [40], and several variants have appeared, including the CKKW-L, MLM, and pseudoshower methods, see Refs. 41, 15 for summaries. Truncated showers are required [42] to maintain color coherence when interfacing to angular-ordered parton showers, and care must be taken to use consistent α_S choices for the real (ME-driven) and virtual (PS-driven) corrections [43].

In the ME+PS method one typically starts by generating LO matrix elements for the production of the basic process plus a certain number $\leq n$ of other partons. A minimum separation is imposed on the produced partons, requiring, for example, that the relative transverse momentum in any pair of partons is above a given cut Q_{cut} . One then reweights these amplitudes in such a way that, in the strongly ordered region, the virtual effects that are included in the shower algorithm (i.e. running couplings and Sudakov form factors) are also accounted for. At this stage, before parton showers are added, the generated configurations are tree-level accurate at large angle, and at small angle they match the results of the shower algorithm, except that there are no emissions below the scale Q_{cut} , and no final states with more than n partons. These kinematic configurations are thus fed into a GPMC, that must generate all splittings with relative transverse momentum below the scale Q_{cut} , for initial events with less than n partons, or below the scale of the smallest pair transverse momentum, for events with n partons. The matching parameter Q_{cut} must be chosen to be large enough for fixed-order perturbation theory to hold, but small enough so that the shower is accurate for emissions below it. Notice that the accuracy achieved with MEC is equivalent to that of ME+PS with $n = 1$, where MEC has the advantage of not having a matching parameter Q_{cut} .

The popularity of the ME+PS method is due to the fact that processes with many jets appear often as backgrounds to new-physics searches. These jets are typically required to be well separated, and to have large transverse momenta. These kinematical configurations are exactly those for which pure shower algorithms are unreliable, hence it is mandatory to describe them using at least LO matrix elements.

Several ME+PS implementations use existing LO generators, like **ALPGEN** [52], **MADGRAPH** [53], and others summarized in Ref. 41, for the calculation of the matrix elements, and feed the partonic events to a GPMC like **PYTHIA** or **HERWIG** using the Les Houches Interface for User Processes (LHI/LHEF) [54,34]. **SHERPA** and **HERWIG++** also include their own matrix-element generators.

The NLO+PS methods promote the accuracy of the generation of the basic process from LO to NLO in QCD. They must thus include the radiation of one extra parton with tree-level accuracy, since this radiation constitutes a NLO correction to the basic process. They must also include NLO virtual corrections. They can be viewed as an extension of the MEC methods with the inclusion of NLO virtual corrections. They are however more general, since they are applicable to processes of arbitrary complexity. Two of these methods are now widely used: **MC@NLO** [44] and **POWHEG** [42,45], with several alternative methods now also being pursued, see Ref. 15 and references therein.

NLO+PS generators produce NLO accurate distributions for inclusive quantities, and generate the hardest jet with tree-level accuracy. It should be recalled, though, that in $2 \rightarrow 1$ processes like Z/W production, GPMCs including MEC and weighted by a constant K factor may perform nearly as well, and, if suitably tuned, may even yield a better description of data. In this context, note also that the optimal tuning of an NLO+PS generator may well be different from that of the pure PS.

Several NLO+PS processes are implemented in the **MC@NLO** program [44], together with the new **AMC@NLO** development [55], and in the **POWHEG BOX** framework [45]. **HERWIG++** also includes its own **POWHEG** implementation, suitably adapted with the inclusion of

vetoed and truncated showers, for several processes. SHERPA instead implements a variant of the MC@NLO method.

For applications that require an accurate description of more than one hard, large-angle jet associated with the primary process, ME+PS schemes are still superior to NLO+PS ones. Ideally, one would like to improve NLO generators in such a way that also the production of associated jets achieves NLO accuracy. The FFX [47], UNLOPS [48], MiNLO [49] and MEPS@NLO [50] methods address this problem. In turn, its solution is a prerequisite for the construction of NNLO+PS generators, that in fact have already appeared for the $gg \rightarrow H$ and Drell-Yan processes (see ref. [51] and references therein).

41.2. Hadronization Models

In the context of GPMCs, *hadronization* denotes the process by which a set of colored partons (*after* showering) is transformed into a set of color-singlet *primary* hadrons, which may then subsequently decay further (to *secondary* hadrons). This non-perturbative transition takes place at the *hadronization scale* Q_{had} , which by construction is identical to the infrared cutoff of the parton shower. In the absence of a first-principles solution to the relevant dynamics, GPMCs use QCD-inspired phenomenological models to describe this transition.

A key difference between MC hadronization models and the fragmentation-function (FF) formalism used to describe inclusive hadron spectra in perturbative QCD (see Chap. 9 and Chap. 20 of PDG book) is that the former is always defined at the hadronization scale, while the latter can be defined at an arbitrary perturbative scale Q . They can therefore only be compared directly if the perturbative evolution between Q and Q_{had} is taken into account. FFs are calculable in pQCD, given a non-perturbative initial condition obtained by fits to hadron spectra. In the MC context, one can prove that the correct QCD evolution of the FFs arises from the shower formalism, with the hadronization model providing an explicit parametrization of the non-perturbative component. However, the MC modeling of shower and hadronization includes much more information on the final state since it is fully exclusive (i.e., it addresses all particles in the final state explicitly), while FFs only describe inclusive spectra. This exclusivity also enables MC models to make use of the color-flow information coming from the perturbative shower evolution (see Sec. 41.1.5) to determine between which partons the confining potentials should arise. This is the non-perturbative analog of QCD coherence [56].

Given an exact hadronization model, its dependence on the hadronization scale Q_{had} should in principle be compensated by the corresponding scale dependence of the shower algorithm, which stops generating branchings at the scale Q_{had} . However, due to their complicated and fully exclusive nature, it is generally not possible to enforce this compensation automatically in MC models. One must therefore be aware that the model must be “retuned” by hand if changes are made to the perturbative evolution, in particular if the infrared cutoff is modified. Tuning is discussed briefly in Sec. 41.4.

An important result in “quenched” lattice QCD (see Chap. 18 of PDG book) is that the potential of the color-dipole field between a charge and an anticharge appears to grow

linearly with the separation of the charges, at distances greater than about a femtometer. This is known as “linear confinement”, and it forms the starting point for the *string model of hadronization*, discussed below in Sec. 41.2.1. Alternatively, a property of perturbative QCD called “preconfinement” is the basis of the *cluster model of hadronization*, discussed in Sec. 41.2.2.

Finally, it should be emphasized that the so-called “parton level” that can be obtained by switching off hadronization in a GPMC, is not a universal concept, since each model defines the hadronization scale differently (e.g. by a cutoff in $p_\perp$, invariant mass, etc., with different tunes using different values for the cutoff). Comparisons to distributions at this level may therefore be used to provide an idea of the overall impact of hadronization corrections within a given model, but should be avoided in the context of physical observables.

41.2.1. The String Model :

Starting from early concepts [57], several hadronization models based on strings have been proposed [15]. Of these, the most widely used today is the so-called Lund model [58,59], implemented in PYTHIA [3,4]. We concentrate on that particular model here, though many of the overall concepts would be shared by any string-inspired method.

Consider a color-connected quark-antiquark pair with no intermediate gluons emerging from the parton shower (like the $\bar{q}q$ pair in the center of Fig. 41.1), e.g. a red q and an antired $\bar{q}$. As the charges move apart, linear confinement implies that a potential $V(r) = \kappa r$ is reached for large distances r . (At short distances, there is a Coulomb term $\propto 1/r$ as well, but this is neglected in the Lund string.) This potential describes a string with tension $\kappa \sim 1 \text{ GeV/fm} \sim 0.2 \text{ GeV}^2$. The physical picture is that of a color flux tube being stretched between the q and the $\bar{q}$.




Figure 41.2: Illustration of string breaking by quark pair-creation in the string field.

As the string grows, the non-perturbative creation of quark-antiquark pairs can break the string, via the process $(q\bar{q}) \rightarrow (q\bar{q}') + (q'\bar{q})$, illustrated in Fig. 41.2. More complicated color-connected quark-antiquark configurations involving intermediate gluons (like the $\bar{q}gggq$ and $\bar{q}gq$ systems on the left and right part of Fig. 41.1) are treated by representing gluons as transverse “kinks.” Thus soft gluons effectively build up a transverse structure in the originally one-dimensional object, with infinitely soft ones smoothly absorbed into the string. For strings with finite-energy kinks, the space-time evolution is slightly more involved [59], but the main point is that there are no separate free parameters for gluon

jets. Differences with respect to quark fragmentation arise simply because quarks are only connected to a single string piece, while gluons have one on either side, increasing their relative energy loss (per unit invariant time) by a factor of 2, similar to the ratio of color Casimirs $C_A/C_F = 2.25$.

Since the string breaks are causally disconnected (as can be realized from space-time diagrams [59]) , they do not have to be considered in any specific time-ordered sequence. In the Lund model, the string breaks are generated starting with the leading (“outermost”) hadrons, containing the endpoint quarks, and iterating inwards towards the center of the string, alternating randomly between the left and right sides. One can thereby split off a single on-shell hadron in each step, making it straightforward to ensure that only states consistent with known hadron states are produced.

For each breakup vertex, quantum mechanical tunneling is assumed to control the masses and $p_\perp$ kicks that can be produced, leading to a Gaussian suppression

$$\text{Prob}(m_q^2, p_{\perp q}^2) \propto \exp\left(\frac{-\pi m_q^2}{\kappa}\right) \exp\left(\frac{-\pi p_{\perp q}^2}{\kappa}\right) , \quad (41.10)$$

where m_q is the mass of the produced quark flavor and $p_\perp$ is the non-perturbative transverse momentum imparted to it by the breakup process (the antiquark has the same mass and opposite $p_\perp$), with a universal average value of $\langle p_{\perp q}^2 \rangle = \kappa/\pi \sim (250 \text{ MeV})^2$. The charm and bottom masses are sufficiently heavy that they are not produced at all in the soft fragmentation. The transverse direction is defined with respect to the string axis, so the $p_\perp$ in a frame where the string is moving will be modified by a Lorentz boost. Note that the effective amount of “non-perturbative” $p_\perp$, in a Monte Carlo model with a fixed shower cutoff Q_{had} , may be larger than the purely non-perturbative κ/π above, to account for effects of additional unresolved soft-gluon radiation below Q_{had} . In principle, the magnitude of this additional component should scale with the cutoff, but in practice it is up to the user to enforce this by retuning the relevant parameter when changing the hadronization scale.

Since quark masses are difficult to define for light quarks, the value of the strangeness suppression is determined from experimental observables, such as the K/π and K^*/ρ ratios. Note that the parton-shower evolution generates a small amount of strangeness as well, through perturbative $g \rightarrow s\bar{s}$ splittings.

Baryon production can also be incorporated, by allowing string breaks to produce pairs of *diquarks*, loosely bound states of two quarks in an overall $\bar{3}$ representation. Again, since diquark masses are difficult to define, the relative rate of diquark to quark production is extracted, e.g. from the p/π ratio. Since the perturbative shower splittings do not produce diquarks, the optimal value for this parameter is mildly correlated with the amount of $g \rightarrow q\bar{q}$ splittings produced by the shower. More advanced scenarios for baryon production have also been proposed, see Ref. 59. Within the PYTHIA framework, a fragmentation model including baryon string junctions [60] is also available.

The next step of the algorithm is the assignment of the produced quarks within hadron multiplets. Using a nonrelativistic classification of spin states, the fragmenting q may

combine with the $\bar{q}'$ from a newly created breakup to produce a meson — or baryon, if diquarks are involved — of a given spin S and angular momentum L . The lowest-lying pseudoscalar and vector meson multiplets, and spin-1/2 and -3/2 baryons, are assumed to dominate in a string framework¹, but individual rates are not predicted by the model. This is therefore the sector that contains the largest amount of free parameters.

From spin counting, the ratio V/P of vectors to pseudoscalars is expected to be 3, but in practice this is only approximately true for B mesons. For lighter flavors, the difference in phase space caused by the V – P mass splittings implies a suppression of vector production. When extracting the corresponding parameters from data, it is advisable to begin with the heaviest states, since so-called feed-down from the decays of higher-lying hadron states complicates the extraction for lighter particles, see Sec. 41.2.3. For baryons, separate parameters control the relative rates of spin-1 diquarks vs. spin-0 ones and, likewise, have to be extracted from data.

With $p_{\perp}^2$ and m^2 now fixed, the final step is to select the fraction, z , of the fragmenting endpoint quark’s longitudinal momentum that is carried by the created hadron, an aspect for which the string model is highly predictive. The requirement that the fragmentation be independent of the sequence in which breakups are considered (causality) imposes a “left-right symmetry” on the possible form of the fragmentation function, $f(z)$, with the solution

$$f(z) \propto \frac{1}{z} (1-z)^a \exp\left(-\frac{b(m_h^2 + p_{\perp h}^2)}{z}\right), \quad (41.11)$$

which is known as the Lund symmetric fragmentation function (normalized to unit integral). The dimensionless parameter a dampens the hard tail of the fragmentation function, towards $z \rightarrow 1$, and may in principle be flavor-dependent, while b , with dimension GeV^{-2} , is a universal constant related to the string tension [59] which determines the behavior in the soft limit, $z \rightarrow 0$. Note that the explicit mass dependence in $f(z)$ implies a harder fragmentation function for heavier hadrons (in the rest frame of the string).

As a by-product, the probability distribution in invariant time τ of $q'\bar{q}$ breakup vertices, or equivalently $\Gamma = (\kappa\tau)^2$, is also obtained, with $dP/d\Gamma \propto \Gamma^a \exp(-b\Gamma)$ implying an area law for the color flux, and the average breakup time lying along a hyperbola of constant invariant time $\tau_0 \sim 10^{-23}\text{s}$ [59].

For massive endpoints (e.g. c and b quarks, or hypothetical hadronizing new-physics particles), which do not move along straight lightcone sections, the exponential suppression with string area leads to modifications of the form $f(z) \rightarrow f(z)/z^{bm_Q^2}$, with m_Q the mass of the heavy quark [61]. Although different forms can also be used to

¹ The PYTHIA implementation includes the lightest pseudoscalar and vector mesons, with the four $L = 1$ multiplets (scalar, tensor, and 2 pseudovectors) available but disabled by default, largely because several states are poorly known and thus may result in a worse overall description when included. For baryons, the lightest spin-1/2 and -3/2 multiplets are included.

describe inclusive heavy-meson spectra (see Sec 20.9 of PDG book), such choices are not consistent with causality in the string framework and hence are theoretically disfavored in this context, one well-known example being the Peterson formula [62],

$$f(z) \propto \frac{1}{z} \left(1 - \frac{1}{z} - \frac{\epsilon_Q}{1-z} \right)^{-2}, \quad (41.12)$$

with ϵ_Q a free parameter expected to scale $\propto 1/m_Q^2$.

41.2.2. The Cluster Model :

The cluster hadronization model is based on *preconfinement*, i.e., on the observation [63,64] that the color structure of a perturbative QCD shower evolution at any scale Q_0 is such that color-singlet subsystems of partons (labeled “clusters”) occur with a universal invariant mass distribution that only depends on Q_0 and on Λ_{QCD} , not on the starting scale Q , for $Q \gg Q_0 \gg \Lambda_{\text{QCD}}$. Further, this mass distribution is power-suppressed at large masses.

Following early models based on this universality [8,65], the cluster model developed by Webber [66] has for many years been a hallmark of the **HERWIG** and **HERWIG++** generators, with an alternative implementation [67] now available in the **SHERPA** generator. The key idea, in addition to preconfinement, is to force “by hand” all gluons to split into quark-antiquark pairs at the end of the parton shower. Compared with the string description, this effectively amounts to viewing gluons as “seeds” for string breaks, rather than as kinks in a continuous object. After the splittings, a new set of low-mass color-singlet clusters is obtained, formed only by quark-antiquark pairs. These can be decayed to on-shell hadrons in a simple manner.

The algorithm starts by generating the forced $g \rightarrow q\bar{q}$ breakups, and by assigning flavors and momenta to the produced quark pairs. For a typical shower cutoff corresponding to a gluon virtuality of $Q_{\text{had}} \sim 1 \text{ GeV}$, the $p_{\perp}$ generated by the splittings can be neglected. The constituent light-quark masses, $m_{u,d} \sim 300 \text{ MeV}$ and $m_s \sim 450 \text{ MeV}$, imply a suppression (typically even an absence) of strangeness production. In principle, the model also allows for diquarks to be produced at this stage, but due to the larger constituent masses this would only become relevant for shower cutoffs larger than 1 GeV .

If a cluster formed in this way has an invariant mass above some cutoff value, typically 3–4 GeV, it is forced to undergo sequential $1 \rightarrow 2$ cluster breakups, along an axis defined by the constituent partons of the original cluster, until all sub-cluster masses fall below the cutoff value. Due to the preservation of the original axis in these breakups, this treatment has some resemblance to the string-like picture, though the non-perturbative $p_{\perp}$ kicks generated in this way are generally larger, up to half the allowed cluster mass.

Next, on the low-mass side of the spectrum, some clusters are allowed to decay directly to a single hadron, with nearby clusters absorbing any excess momentum. This improves the description of the high- z part of the fragmentation spectrum — where the hadron carries almost all the momentum of its parent jet — at the cost of introducing one additional parameter, controlling the probability for single-hadron cluster decay.

Having obtained a final distribution of small-mass clusters, now with a strict cutoff at 3–4 GeV and with the component destined to decay to single hadrons already removed,

the remaining clusters are interpreted as a smoothed-out spectrum of excited mesons, each of which decays isotropically to two hadrons, with relative probabilities proportional to the available phase space for each possible two-hadron combination that is consistent with the cluster’s internal flavors, including spin degeneracy. It is important that all the light members (containing only uds) of each hadron multiplet be included, as the absence of members can lead to unphysical isospin or $SU(3)$ flavor violation. Typically, the lightest pseudoscalar, vector, scalar, even and odd charge conjugation pseudovector, and tensor multiplets of light mesons are included. In addition, some excited vector multiplets of light mesons may be available. For baryons, usually only the lightest flavor-octet, -decuplet and -singlet baryons are present, although both the **HERWIG++** and **SHERPA** implementations now include some heavier baryon multiplets as well.

Differently from the string model, the mechanism of phase-space suppression employed here leads to a natural enhancement of the lighter pseudoscalars, and no parameters beyond the spectrum of hadron masses need to be introduced at this point. The phase space also limits the transverse momenta of the produced hadrons relative to the jet axis.

Note that, since the masses and decays of excited heavy-flavor hadrons in particular are not well known, there is some freedom in the model to adjust these, which in turn will affect their relative phase-space populations.

41.2.3. *Hadron and τ Decays :*

Of the so-called primary hadrons, originating directly from string breaks and/or cluster decays (see above), many are unstable and so decay further, until a set of particles is obtained that can be considered stable on time scales relevant to the given measurement². The decay modeling can therefore have a significant impact on final particle yields and spectra, especially for the lowest-lying hadronic states, which receive the largest relative contributions from decays (feed-down). Note that the interplay between primary production and feed-down implies that the hadronization parameters should be retuned if significant changes to the decay treatment are made.

Particle summary tables, such as those given elsewhere in this *Review*, represent a condensed summary of the available experimental measurements and hence may be incomplete and/or exhibit inconsistencies within the experimental precision. In an MC decay package, on the other hand, all information must be quantified and consistent, with all branching ratios summing to unity. When adapting particle summary information for use in a decay package, a number of choices must therefore be made. The amount of ambiguity increases as more excited hadron multiplets are added to the simulation, about which less and less is known from experiment, with each GPMC making its own choices.

A related choice is how to distribute the decay products differentially in phase space, in particular which matrix elements to use. Historically, MC generators contained matrix elements only for selected (generator-specific) classes of hadron and τ decays, coupled with a Breit-Wigner smearing of the masses, truncated at the edges of the physical decay phase space (the treatment of decay thresholds can be important for certain modes [15]) .

² E.g., a typical hadron-collider definition of a “stable particle” is $c\tau \geq 10$ mm, which includes the weakly-decaying strange hadrons (K , Λ , $\Sigma^\pm$, $\bar{\Sigma}^\pm$, Ξ , Ω).

A more sophisticated treatment can then be obtained by reweighting the generated events using the obtained particle four-momenta and/or by using specialized external packages such as `EVTGEN` [68] for hadron decays and `TAUOLA` [69] for τ decays.

More recently, `HERWIG++` and `SHERPA` include helicity-dependence in τ decays [70,71], with a more limited treatment available in `PYTHIA 8` [4]. The `HERWIG++` and `SHERPA` generators have also included significantly improved internal simulations of hadronic decays, which include spin correlations between those decays for which matrix elements are used. Photon-bremsstrahlung effects are discussed in Sec. 41.1.6.

`HERWIG++` and `PYTHIA` include the probability for B mesons to oscillate into $\bar{B}$ ones before decay. `SHERPA` and `EVTGEN` also include CP-violating effects and, for common decay modes of the neutral meson and its antiparticle, the interference between the direct decay and oscillation followed by decay.

We end on a note of warning on double counting. This may occur if a particle can decay via an intermediate on-shell resonance. An example is $a_1 \rightarrow \pi\pi\pi$ which may proceed via $a_1 \rightarrow \rho\pi$, $\rho \rightarrow \pi\pi$. If these decay channels of the a_1 are both included, each with their full partial width, a double counting of the on-shell $a_1 \rightarrow \rho\pi$ contribution would result. Such cases are normally dealt with consistently in the default MC generator packages, so this warning is mostly for users that wish to edit decay tables on their own.

41.3. Models for Soft Hadron-Hadron Physics

41.3.1. Minimum-Bias and Diffraction :

The term “minimum bias” (MB) originates from the experimental requirement of a minimal number of tracks (or hits) in a given instrumented region. In order to make MC predictions for such observables, all possible contributions to the relevant phase-space region must be accounted for. There are essentially four types of physics processes, which together make up the total hadron-hadron (hh) cross section: 1) elastic scattering³: $hh \rightarrow hh$, 2) single diffractive dissociation: $hh \rightarrow h + \text{gap} + X$, with X denoting anything that is not the original beam particle, and “gap” denoting a rapidity region devoid of observed activity; 3) double diffractive dissociation: $hh \rightarrow X + \text{gap} + X$, and 4) inelastic non-diffractive scattering: everything else. A fifth class may also be defined, called central diffraction ($hh \rightarrow h + \text{gap} + X + \text{gap} + h$). Some differences exist between theoretical and experimental terminology [72]. In the experimental setting, diffraction is defined by an observable gap, of some minimal size in rapidity. In the MC context, each diffractive physics process typically produces a whole spectrum of gaps, with small ones suppressed but not excluded.

The inelastic non-diffractive part of the cross section is typically modeled either by smoothly regulating and extending the perturbative QCD scattering cross sections all the way to zero $p_\perp$ [73] (`PYTHIA 6`, `PYTHIA 8`, and `SHERPA`), or by regulating the QCD cross sections with a sharp cutoff [74] (`HERWIG+JIMMY`) and adding a separate class of intrinsically soft scatterings below that scale [75] (`HERWIG++`). See also Sec. 41.3.2.

³ The QED elastic-scattering cross section diverges and is normally a non-default option in MC models.

In all cases, the three most important ingredients are: 1) the IR regularization of the perturbative scattering cross sections, including their PDF dependence, 2) the assumed matter distribution of the colliding hadrons, possibly including multi-parton correlations [60] and/or x dependence [76], and 3) additional soft-QCD effects such as color reconnections and/or other collective effects, discussed in Sec. 41.3.3.

Currently, there are essentially three methods for simulating diffraction in the main MC models: 1) in **PYTHIA 6**, one picks a diffractive mass according to parametrized cross sections $\propto dM^2/M^2$ [77]. This mass is represented as a string, which is fragmented as described in Sec. 41.2.1, though differences in the effective scale of the hadronization may necessitate a (re)tuning of the fragmentation parameters for diffraction; 2) in **PYTHIA 8**, the high-mass tail beyond $M \sim 10$ GeV is augmented by a partonic description in terms of pomeron PDFs [78], allowing diffractive jet production including showers and underlying event [79]; 3) the **PHOJET** and **DPMJET** programs also include central diffraction and rely directly on a formulation in terms of pomerons (color-singlet multi-gluon states) [80–82]. Cut pomerons correspond to exchanges of soft gluons while uncut ones give elastic and diffractive topologies as well as virtual corrections that help preserve unitarity. So-called “hard pomerons” provide a transition to the perturbative regime. Fragmentation is still handled using the Lund string model, so there is some overlap with the above models at the hadronization stage. In addition, a pomeron-based package exists for **HERWIG** [83], and an effort is underway to construct an MC implementation of the “KMR” model [84] within the **SHERPA** generator. Color reconnections (Sec. 41.3.3) may also play a role in creating rapidity gaps and the underlying event (Sec. 41.3.2) in destroying them.

41.3.2. Underlying Event and Jet Pedestals :

In the GPMC context, “underlying event” (UE) denotes any additional activity *beyond* the basic process and its associated ISR and FSR activity. The dominant contribution is believed to come from additional color exchanges between the beam particles, which in models are represented either as multiple parton-parton interactions (MPI) or as so-called cut pomerons (Sec. 41.3.1). The experimentally observed fact that the UE is more active than average (MB) events at the same CM energy is called the “jet pedestal” effect.

The most clearly identifiable consequence of MPI is arguably the possibility of observing several hard parton-parton interactions in one and the same hadron-hadron event. In the most likely case that they are all QCD $2 \rightarrow 2$ interactions, this produces two or more back-to-back jet pairs, with each pair having a small value of $\text{sum}(\vec{p}_\perp)$. The fraction of MPI that give rise to additional reconstructible jets is, however, small. Soft interactions, that exchange color and a small amount of momentum without giving rise to observable jets, are much more plentiful, and can give significant corrections to the color flow and total scattered energy of the event. This affects the final-state activity in a more global way, increasing hadron-multiplicity and summed E_T distributions, and contributing to the break-up of the beam remnants in the forward direction.

The first detailed Monte Carlo model for perturbative MPI was proposed in Ref. 73, and with some variation this still forms the basis for most modern implementations. Some useful additional references can be found in Ref. 15. The first crucial observation is that the t -channel propagators appearing in perturbative QCD $2 \rightarrow 2$ scattering almost

go on shell at low $p_\perp$, causing the differential cross sections to behave roughly as

$$d\sigma_{2\rightarrow 2} \propto \frac{dt}{t^2} \sim \frac{dp_\perp^2}{p_\perp^4} . \quad (41.13)$$

This cross section represents the inclusive scattering of partons against partons in perturbative QCD, summed over all partons. Thus, if a single hadron-hadron scattering contains *two* parton-parton interactions, that event will contribute twice to the parton-parton cross section $\sigma_{2\rightarrow 2}$ but only once to the hadron-hadron one σ_{tot} , and so on. In the limit that all the parton-parton interactions are independent and equivalent, one has

$$\sigma_{2\rightarrow 2} = \langle n \rangle \sigma_{\text{tot}} , \quad (41.14)$$

with $\langle n \rangle$ the average number of parton-parton interactions, typically defined with some minimal $p_\perp > p_{\perp\text{min}}$ to render the parton-parton cross section finite. The probability for having n parton-parton scatterings in a single hadron-hadron collision then follows a Poisson distribution,

$$\mathcal{P}_n = \langle n \rangle^n \frac{\exp(-\langle n \rangle)}{n!} . \quad (41.15)$$

This simple argument expresses unitarity; instead of the total hadron-hadron interaction cross section diverging as the parton-parton $p_\perp \rightarrow 0$ (which would violate unitarity), we have restated the problem so that it is now the *number of parton-parton interactions per hadron-hadron collision* that diverges, with the total hadron-hadron cross section remaining finite. At LHC energies, the parton-parton scattering cross sections computed using the LO QCD cross section folded with modern PDFs become larger than the total pp one for $p_{\perp\text{min}}$ values of order 4–5 GeV (see e.g. [85,86]) . One therefore expects the average number of perturbative MPI to exceed unity at around that scale.

Two important ingredients remain to fully regulate the remaining divergence. Firstly, the interactions cannot use up more momentum than is available in the parent hadron. This suppresses the large- n tail of the estimate above. In PYTHIA-based models, the MPI are ordered in $p_\perp$, and the parton densities for each successive interaction are explicitly constructed so that the sum of x fractions can never be greater than unity. In the HERWIG models, the Poisson estimate of $\langle n \rangle$ above is used as an initial guess, but the generation of actual MPI is stopped once the energy-momentum conservation limit is reached. Both of these approaches generate momentum (conservation) correlations among the MPI.

The second ingredient invoked to suppress the number of interactions, at low $p_\perp$ and x , is color screening; if the wavelength $\sim 1/p_\perp$ of an exchanged colored parton becomes larger than a typical color-anticolor separation distance, it will only see an *average* color charge that vanishes in the limit $p_\perp \rightarrow 0$, hence leading to suppressed interactions. This provides an infrared cutoff for MPI similar to that provided by the hadronization scale for parton showers. A first estimate of the color-screening cutoff would be the proton size, $p_{\perp\text{min}} \approx \hbar/r_p \approx 0.3 \text{ GeV} \approx \Lambda_{\text{QCD}}$, but empirically this appears to be far too low. In current models, one replaces the proton radius r_p in the above formula by a “typical color screening distance,” i.e., an average size of a region within which the net compensation

of a given color charge occurs. This number is not known from first principles [84] and is perceived of simply as an effective cutoff parameter. The simplest choice is to introduce a step function $\Theta(p_\perp - p_{\perp\min})$. Alternatively, one may note that the jet cross section is divergent like $\alpha_S^2(p_\perp^2)/p_\perp^4$, cf. Eq. (41.13), and that therefore a factor

$$\frac{\alpha_S^2(p_{\perp 0}^2 + p_\perp^2)}{\alpha_S^2(p_\perp^2)} \frac{p_\perp^4}{(p_{\perp 0}^2 + p_\perp^2)^2} \quad (41.16)$$

would smoothly regulate the divergences, now with $p_{\perp 0}$ as the free parameter. Regardless of whether it is imposed as a smooth (PYTHIA and SHERPA) or steep (HERWIG++) function, this is effectively the main “tuning” parameter in such models.

Note that the numerical value obtained for the cross section depends upon the PDF set used, and therefore the optimal value to use for the cutoff will also depend on this choice. Note also that the cutoff does not have to be energy-independent. Higher energies imply that parton densities can be probed at smaller x values, where the number of partons rapidly increases. Partons then become closer packed and the color screening distance d decreases. The uncertainty on the energy and/or x scaling of the cutoff is a major concern when extrapolating between different collider energies [87].

We now turn to the origin of the observational fact that hard jets appear to sit on top of a higher “pedestal” of underlying activity than events with no hard jets. This is interpreted as a consequence of impact-parameter-dependence: in peripheral collisions, only a small fraction of events contain any high- $p_\perp$ activity, whereas central collisions are more likely to contain at least one hard scattering; a high- $p_\perp$ triggered sample will therefore be biased towards small impact parameters, b . The ability of a model to describe the shape of the pedestal (e.g. to describe both MB and UE distributions simultaneously) therefore depends upon its modeling of the b -dependence, and correspondingly the impact-parameter shape constitutes another main tuning parameter.

For each impact parameter b , the number of interactions $\tilde{n}(b)$ can still be assumed to be distributed according to Eq. (41.15), again modulo momentum conservation, but now with the mean value of the Poisson distribution depending on impact parameter, $\langle \tilde{n}(b) \rangle$. This causes the final n -distribution (integrated over b) to be wider than a Poissonian.

Finally, there are two perturbative modeling aspects which go beyond the introduction of MPI themselves: 1) parton showers off the MPI, and 2) perturbative parton-rescattering effects. Without showers, MPI models would generate very sharp peaks for back-to-back MPI jets, caused by unshowered partons passed directly to the hadronization model. However, with the exception of the oldest PYTHIA6 model, all GPMC models do include such showers [15], and hence should exhibit more realistic (i.e., broader and more decorrelated) MPI jets. On the initial-state side, the main questions are whether and how correlated multi-parton densities are taken into account and, as discussed previously, how the showers are regulated at low $p_\perp$ and/or low x . Although none of the MC models currently impose a rigorous correlated multi-parton evolution, all of them include some elementary aspects. The most significant for parton-level results is arguably momentum conservation, which is enforced explicitly in all the models. The so-called “interleaved” models [24] attempt to go a step further, generating an explicitly correlated multi-parton

evolution in which flavor sum rules are imposed to conserve, e.g. the total numbers of valence and sea quarks [60].

Perturbative rescattering in the final state can occur if partons are allowed to undergo several distinct interactions, with showering activity possibly taking place in-between. This has so far not been studied extensively, but a first exploratory model is available [88]. In the initial state, parton rescattering/recombination effects have so far not been included in any of the GPMC models.

41.3.3. *Bose-Einstein and Color-Reconnection Effects :*

In the context of e^+e^- collisions, Bose-Einstein (BE) correlations have mostly been discussed as a source of uncertainty on high-precision W mass determinations at LEP [89]. In hadron-hadron (and nucleus-nucleus) collisions, however, BE correlations are used extensively to study the space-time structure of hadronizing matter (“femtoscopy”).

In MC models of hadronization, each string break or particle/cluster decay is normally factorized from all other ones. This reduces the number of variables that must be considered simultaneously, but also makes the introduction of correlations among particles from different breaks/decays intrinsically difficult to address. In the context of GPMCs, a few semi-classical models are available within the PYTHIA 6 and 8 generators [90], in which the BE effect is mimicked by an attractive interaction between pairs of identical particles in the final state, with no higher correlations included. This “force” acts after the decays of very short-lived particles, like ρ , but before decays of longer-lived ones, like π^0 . The main differences between the variants of this model is the assumed shape of the correlation function and how overall momentum conservation is handled.

As discussed in Sec. 41.2, leading-color (“planar”) color flows are used to set up the hadronizing systems (clusters or strings) at the hadronization stage. If the systems do not overlap significantly in space and time, subleading-color ambiguities and/or non-perturbative reconnections are expected to be small. However, if the density of displaced color charges is sufficiently high that several systems can overlap significantly, full-color and/or reconnection effects should become progressively larger.

In the specific context of MPI, a crucial question is how color is neutralized *between* different MPI systems, including the remnants. The large rapidity differences involved imply large invariant masses (though normally low $p_\perp$), and hence large amounts of (soft) particle production. Indeed, in the context of soft-inclusive physics, it is these “inter-system” strings/clusters that furnish the dominant particle-production mechanism, and hence their modeling is an essential part of the soft-physics description, affecting topics such as MB/UE multiplicity and $p_\perp$ distributions, rapidity gaps, and precision mass measurements. More comprehensive reviews of color-reconnection effects can be found in Refs. 15,91.

41.4. Parameters and Tuning

The accuracy that can be achieved by a GPMC model depends on the sophistication of theory models it incorporates, on the inclusiveness of the observable(s) under study, and on the available constraints on its free parameters. Using existing data (or more precise theory calculations) to constrain the model parameters is referred to as generator tuning.

Typically, the most inclusive event properties are determined by only a few, very important parameters, such as the value of α_S , for perturbative corrections, and the shape of the fragmentation functions, for non-perturbative ones. More parameters may then be introduced to describe successively more exclusive aspects, but these should have progressively less impact on the overall modelling. One may therefore take a factorized approach, first constraining the perturbative parameters and thereafter the non-perturbative ones, in order of decreasing significance to the overall modeling. Furthermore, by identifying which measurements are most sensitive to each parameter, this ordering can be reflected in the way that data is selected and applied to constrain the models. Thus, the most inclusive measurements should be used first, to constrain the most inclusive parameters, and so on for progressively more exclusive aspects.

At LO $\times$ LL, perturbation theory is doing well if it agrees with an IR safe measurement within 10%. It would therefore not make much sense to tune a GPMC beyond roughly 5% (it might even be dangerous, due to overfitting). The advent of NLO Monte Carlos may reduce this number slightly, but only for quantities for which one expects NLO precision. For LO Monte Carlos, distributions should be normalized to unity, since the NLO normalization is not tunable. For quantities governed by non-perturbative physics, uncertainties are larger. For some quantities, e.g. ones for which the underlying modeling is known to be poor, an order-of-magnitude agreement or worse may have to be accepted.

In the context of LO $\times$ LL GPMC tuning, subleading aspects of coupling-constant and PDF choices are relevant. In particular, one should be aware that the choice of QCD Λ parameter $\Lambda_{\text{MC}} = 1.569\Lambda_{\overline{\text{MS}}}$ (for 5 active flavors) improves the predictions of coherent shower algorithms at the NLL level [92], and hence this scheme is typically considered the baseline for shower tuning. The question of LO vs. NLO PDFs is more involved [15], but it should be emphasized that the low- x gluon in particular is important for determining the level of the underlying event in MPI models (Sec. 41.3.2), and hence the MB/UE tuning (and energy scaling [87]) is linked to the choice of PDF in such models. Further issues and an example of a specific recipe that could be followed in a realistic set-up can be found in Ref. 93. A useful online resource can be found at the mcplots.cern.ch web site [94], based on the RIVET tool [95].

Recent years have seen the emergence of automated tools to reduce the amount of both computer and manpower required for tuning [96]. Automating the human expert input is more difficult. In the tools currently on the market, this is addressed by a combination of input solicited from the GPMC authors (e.g., which parameters and ranges to consider, which observables constitute a complete set, etc) and a set of weights determining the relative priority given to each bin in each distribution. Studies of sensitivities and correlations also play an important role. Overall, the quality of the resulting tunes is by now competitive. The field is still burgeoning, with future sophistications to be expected.

References:

1. G. Corcella *et al.*, JHEP **0101**, 010 (2001), [hep-ph/0011363](#).
2. M.Bähr *et al.*, Eur. Phys. J. **C58**, 639 (2008), [arXiv:0803.0883](#).
3. T. Sjöstrand, S. Mrenna, and P. Z. Skands, JHEP **05**, 026 (2006), [hep-ph/0603175](#).
4. T. Sjöstrand *et al.*, Comp. Phys. Comm. **191**, 159 (2015), [arXiv:1410.3012](#).
5. T. Gleisberg *et al.*, JHEP **0402**, 056 (2004), [hep-ph/0311263](#).
6. T. Kinoshita, J. Math. Phys. **3**, 650 (1962).
7. T. Lee and M. Nauenberg, Phys. Rev. **133**, 1549 (1964).
8. G.C. Fox and S. Wolfram, Nucl. Phys. **B168**, 285 (1980).
9. G. Altarelli and G. Parisi, Nucl. Phys. **B126**, 298 (1977).
10. A. Einstein, B. Podolsky, and N. Rosen, Phys. Rev. **47**, 777 (1935).
11. B.R. Webber, Phys. Lett. **B193**, 91 (1987).
12. J.C. Collins, Nucl. Phys. **B304**, 794 (1988).
13. I.G. Knowles, Comp. Phys. Comm. **58**, 271 (1990).
14. T. Sjöstrand, Phys. Lett. **B157**, 321 (1985).
15. A. Buckley *et al.*, Phys. Reports **504**, 145 (2011), [arXiv:1101.2599](#).
16. G. Marchesini and B.R. Webber, Nucl. Phys. **B310**, 461 (1988).
17. S. Gieseke, P. Stephens, and B. Webber, JHEP **0312**, 045 (2003), [hep-ph/0310083](#).
18. M. Bengtsson and T. Sjöstrand, Nucl. Phys. **B289**, 810 (1987).
19. G. Gustafson and U. Pettersson, Nucl. Phys. **B306**, 746 (1988).
20. L. Lönnblad, Comp. Phys. Comm. **71**, 15 (1992).
21. W.T. Giele, D.A. Kosower, and P.Z. Skands, Phys. Rev. **D78**, 014026 (2008), [arXiv:0707.3652](#).
22. S. Schumann and F. Krauss, JHEP **0803**, 038 (2008), [arXiv:0709.1027](#).
23. Z. Nagy and D.E. Soper, JHEP **0510**, 024 (2005), [hep-ph/0503053](#).
24. T. Sjöstrand and P.Z. Skands, Eur. Phys. J. **C39**, 129 (2005), [hep-ph/0408302](#).
25. E. Norrbin and T. Sjöstrand, Nucl. Phys. **B603**, 297 (2001), [hep-ph/0010012](#).
26. S. Catani *et al.*, Nucl. Phys. **B627**, 189 (2002), [hep-ph/0201036](#).
27. J. Cembranos *et al.*, (2013), [arXiv:1305.2124](#).
28. N. Davidson, T. Przedzinski, and Z. Was, (2010), [arXiv:1011.0937](#).
29. K. Hamilton and P. Richardson, JHEP **0607**, 010 (2006), [hep-ph/0603034](#).
30. M. Schönherr and F. Krauss, JHEP **0812**, 018 (2008), [arXiv:0810.5071](#).
31. A. Semenov, Comp. Phys. Comm. **180**, 431 (2009), [arXiv:0805.0555](#).
32. N.D. Christensen and C. Duhr, Comp. Phys. Comm. **180**, 1614 (2009), [arXiv:0806.4194](#).
33. M. Fairbairn *et al.*, Phys. Reports **438**, 1 (2007), [hep-ph/0611040](#).
34. J. Alwall *et al.*, Comp. Phys. Comm. **176**, 300 (2007), [hep-ph/0609017](#).
35. P.Z. Skands *et al.*, JHEP **0407**, 036 (2004), [hep-ph/0311123](#).
36. J. Alwall *et al.*, (2007), [arXiv:0712.3311](#).
37. P. Richardson, JHEP **0111**, 029 (2001), [hep-ph/0110108](#).
38. M. Bengtsson and T. Sjöstrand, Phys. Lett. **B185**, 435 (1987).
39. W. T. Giele, D. A. Kosower and P. Z. Skands, Phys. Rev. **D84**, 054003 (2011), [arXiv:1102.2126](#).
40. S. Catani *et al.*, JHEP **11**, 063 (2001), [hep-ph/0109231](#).

41. J. Alwall *et al.*, Eur. Phys. J. **C53**, 473 (2008), [arXiv:0706.2569](#).
42. P. Nason, JHEP **11**, 040 (2004), [hep-ph/0409146](#).
43. B. Cooper *et al.*, Eur. Phys. J. **C72**, 2078 (2012), [arXiv:1109.5295](#).
44. S. Frixione and B.R. Webber, JHEP **06**, 029 (2002), [hep-ph/0204244](#).
45. S. Alioli *et al.*, JHEP **1006**, 043 (2010), [arXiv:1002.2581](#).
46. S. Alioli, K. Hamilton, and E. Re, JHEP **09**, 104 (2011), [arXiv:1108.0909](#).
47. R. Frederix and S. Frixione, JHEP **12**, 061 (2012), [arXiv:1209.6215](#).
48. L. Lönnblad and S. Prestel, JHEP **03**, 166 (2013), [arXiv:1211.7278](#).
49. K. Hamilton, P. Nason and G. Zanderighi, JHEP **10**, 155 (2012), [arXiv:1206.3572](#).
50. S. Höche *et al.*, JHEP **04**, 027 (2013), [arXiv:1207.5030](#).
51. K. Hamilton, P. Nason and G. Zanderighi, JHEP **05**, 140 (2015), [arXiv:1501.04637](#).
52. M.L. Mangano *et al.*, JHEP **0307**, 001 (2003), [hep-ph/0206293](#).
53. J. Alwall *et al.*, JHEP **1106**, 128 (2011), [arXiv:1106.0522](#).
54. E. Boos *et al.*, (2007), [hep-ph/0109068](#).
55. J. Alwall *et al.*, JHEP **07**, 079 (2014), [arXiv:1405.0301](#).
56. G. Gustafson, Phys. Lett. **B175**, 453 (1986).
57. X. Artru and G. Mennessier, Nucl. Phys. **B70**, 93 (1974).
58. B. Andersson *et al.*, Phys. Reports **97**, 31 (1983).
59. B. Andersson, Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. **7** (1997).
60. T. Sjöstrand and P.Z. Skands, JHEP **0403**, 053 (2004), [hep-ph/0402078](#).
61. M. Bowler, Z. Phys. **C11**, 169 (1981).
62. C. Peterson *et al.*, Phys. Rev. **D27**, 105 (1983).
63. D. Amati and G. Veneziano, Phys. Lett. **B83**, 87 (1979).
64. A. Bassetto, M. Ciafaloni, and G. Marchesini, Phys. Lett. **B83**, 207 (1979).
65. R.D. Field and S. Wolfram, Nucl. Phys. **B213**, 65 (1983).
66. B.R. Webber, Nucl. Phys. **B238**, 492 (1984).
67. J.-C. Winter, F. Krauss, and G. Soff, Eur. Phys. J. **C36**, 381 (2004), [hep-ph/0311085](#).
68. D. Lange, Nucl. Instrum. Methods **A462**, 152 (2001).
69. S. Jadach *et al.*, Comp. Phys. Comm. **76**, 361 (1993).
70. D. Grellscheid and P. Richardson, (2007), [arXiv:0710.1951](#).
71. T. Gleisberg *et al.*, JHEP **0902**, 007 (2009), [arXiv:0811.4622](#).
72. V. Khoze *et al.*, Eur. Phys. J. **C69**, 85 (2010), [arXiv:1005.4839](#).
73. T. Sjöstrand and M. van Zijl, Phys. Rev. **D36**, 2019 (1987).
74. J.M. Butterworth, J.R. Forshaw, and M.H. Seymour, Z. Phys. **C72**, 637 (1996), [hep-ph/9601371](#).
75. M. Bähr *et al.*, (2009), [arXiv:0905.4671](#).
76. R. Corke and T. Sjöstrand, JHEP **1105**, 009 (2011), [1101.5953](#).
77. G.A. Schuler and T. Sjöstrand, Phys. Rev. **D49**, 2257 (1994).
78. G. Ingelman and P. Schlein, Phys. Lett. **B152**, 256 (1985).
79. S. Navin, (2010), [arXiv:1005.3894](#).
80. P. Aurenche *et al.*, Comp. Phys. Comm. **83**, 107 (1994), [hep-ph/9402351](#).
81. F.W. Bopp, R. Engel, and J. Ranft, (1998), [hep-ph/9803437](#).
82. S. Roesler, R. Engel, and J. Ranft, p. 1033 (2000), [hep-ph/0012252](#).

83. B.E. Cox and J.R. Forshaw, *Comp. Phys. Comm.* **144**, 104 (2002), [hep-ph/0010303](#).
84. M. Ryskin, A. Martin, and V. Khoze, *Eur. Phys. J.* **C71**, 1617 (2011), [arXiv:1102.2844](#).
85. P. Skands, S. Carrazza and J. Rojo, *Eur. Phys. J.* **C74**, 3024 (2014), [arXiv:1404.5630](#).
86. M. Bähr, J.M. Butterworth, and M.H. Seymour, *JHEP* **01**, 065 (2009), [arXiv:0806.2949](#).
87. H. Schulz and P.Z. Skands, *Eur. Phys. J.* **C71**, 1644 (2011), [arXiv:1103.3649](#).
88. R. Corke and T. Sjöstrand, *JHEP* **01**, 035 (2009), [arXiv:0911.1901](#).
89. LEP Electroweak Working Group, (2005), [hep-ex/0511027](#).
90. L. Lönnblad and T. Sjöstrand, *Eur. Phys. J.* **C2**, 165 (1998), [hep-ph/9711460](#).
91. J. R. Christiansen and P. Z. Skands, *JHEP* **08**, 003 (2015), [arXiv:1505.01681](#).
92. S. Catani, B. R. Webber, and G. Marchesini, *Nucl. Phys.* **B349**, 635 (1991).
93. P.Z. Skands, (2011), [arXiv:1104.2863](#).
94. A. Karneyeu *et al.*, *Eur. Phys. J.* **C74**, 2714 (2013), [arXiv:1306.3436](#).
95. A. Buckley *et al.*, *Comp. Phys. Comm.* **184**, 2803 (2010), [arXiv:1003.0694](#).
96. A. Buckley *et al.*, *Eur. Phys. J.* **C65**, 331 (2010), [arXiv:0907.2973](#).