

43. Monte Carlo Event Generators

Revised Oct 2023 by P. Nason (INFN, Milano-Bicocca; Milano-Bicocca U.) and P. Skands (Monash U.).

General-purpose Monte Carlo (GPMC) generators like HERWIG [1–3], PYTHIA [4, 5], and SHERPA [6], provide detailed simulations of high-energy collisions. They play an essential role in QCD modeling (in particular for aspects beyond fixed-order perturbative QCD) and in data analysis and the planning of new experiments, where they are used together with detector simulation 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 in hadron collisions, cannot be computed from first principles at present, 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, the nonperturbative transition from partons to hadrons (“hadronization”) is treated. The two most popular hadronization models, the string and cluster models, are illustrated. The basics of the implementation of decay chains of unstable “primary” hadrons into stable “secondaries” is also illustrated here. In the third section, models for soft hadron physics are discussed. These include models for the underlying event and for low- p_{\perp} (“minimum-bias”) interactions. Aspects of collective effects, such as Bose-Einstein and color-reconnection effects, are also discussed here. The fourth section briefly introduces the challenges of MC uncertainty estimates and tuning.

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

43.1 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. QCD radiation is computable in perturbation theory as long as the time scales involved are well below $1/\Lambda$, where Λ is a typical hadronic scale of few hundred MeV. Because of the presence of logarithmic enhancements due to both collinear and soft emissions, this description involves an indefinite number of final-state particles that are emitted at time scales below $1/\Lambda$. In e^+e^- annihilation into hadrons, for example, the time scale of the primary process is of the order of the inverse of the annihilation energy Q . Collinear and soft emissions take place at all time scales between $1/Q$ and $1/\Lambda$. Technically, the computation of the dominant collinear and soft radiation is carried out by the so called shower algorithms. Historically, “Parton Shower” algorithms were first developed for resummation of collinear singularities. We will briefly describe this approach in this section. We stress, however, that many modern generators adopt approaches that focus initially upon soft singularities, leading to “Dipole Showers”, discussed in Sec. 43.1.3.

Collinear singularities arise when the angle between two emitted light partons becomes small. For example, in a process in which a quark and a gluon are emitted, if the angle θ among them is very small (and is smaller than the angles among all other pairs of light partons in the process) the

squared amplitude factorizes as follows

$$|M_{qg}|^2 d\Phi_{qg} \approx |M_q|^2 d\Phi_q \frac{\alpha_s}{2\pi} P_{q,qg}(z) dz \frac{d\phi d\theta^2}{2\pi \theta^2} \quad (43.1)$$

where M_{qg} , $d\Phi_{qg}$ are the amplitude and phase space describing the emission of both the gluon and the quark; M_q , $d\Phi_q$ are the amplitude and phase space when only the quark is emitted; $z = E_q/(E_q + E_g)$ is the fraction of energy carried by the quark; ϕ is the azimuth of the splitting plane, and $P_{q,qg}(z) = C_F(1+z^2)/(1-z)$ is the Altarelli-Parisi splitting kernel for gluon emission from a quark line, with color factor $C_F = 4/3$. The factorized form of Eq. (43.1) is due to the fact that for small angle the process is dominated by a single amplitude in which the splitting quark is almost on shell and hence propagates for long distances. We define the energy scale corresponding to the inverse of this distance as the *hardness* of the splitting process, so that larger hardness corresponds to shorter distance. We can define the hardness t as the product $E^2\theta^2$, or as the virtuality of the splitting parton p^2 , or as a measure of the relative transverse momentum in the splitting such as the k_t of an emitted parton relative to its parent, defined by

$$p^2 = 2E^2z(1-z)(1-\cos\theta) \approx z(1-z)E^2\theta^2, \quad k_t^2 = z^2(1-z)^2E^2\theta^2. \quad (43.2)$$

If the region of small values of z and $1-z$ was not important, these definitions would be equivalent. In QCD we also have soft divergences, arising when soft gluons are emitted. In Eq. (43.1) they appear as $z \rightarrow 1$, because of the $1/(1-z)$ singularity of $P_{q,qg}(z)$. Thus, we expect that the choice of the appropriate ordering variable will be relevant when dealing with soft divergences (see Sec. 43.3). The $d\theta^2/\theta^2$ factor in Eq. (43.1) can be equivalently written in terms of the hardness dt/t . After integration, it gives rise to a logarithmic factor $\log(Q^2/\Lambda^2)$. We can have many subsequent splittings, that we can describe by applying Eq. (43.1) recursively, as long as the splittings are strongly ordered in decreasing 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. We stress that all hardness values between the hardness of the primary process and the cutoff scale Λ are equally involved here. The collinear approximation is applied recursively to splitting processes that have much smaller hardness with respect to all previous ones.

By integrating over the phase space (and neglecting for the moment soft divergences), a process with n collinear splittings will be of order $(\alpha_S(Q^2) \log(Q^2/\Lambda^2))^n$ with respect to the primary process. Since $\alpha_S(Q^2) \propto 1/\log(Q^2/\Lambda^2)$ [7], these corrections are not small. The so-called KLN theorem [8,9] guarantees that large logarithmic enhancements arising from final-state collinear splitting cancel against the virtual corrections in inclusive cross sections, 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) [7]. 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. For example, the leading order (LO) cross section for $e^+e^- \rightarrow q\bar{q}$ is a good 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 a veto on extra radiation. In summary, perturbation theory at fixed order can yield increasingly accurate predictions for inclusive observables, but cannot be used to describe the indefinite sequence of collinear and soft radiations that accompany the hard partons.

Parton-Shower algorithms are used to compute the cross section for generic hard processes including all dominant collinear radiation. These algorithms begin with the generation of the

kinematics of the basic process, performed with a probability proportional to its LO partonic cross section. This is interpreted physically as the inclusive cross section for the basic process, followed by an arbitrary sequence of shower splittings. The algorithm then assigns a probability to each splitting sequence, so that the initial LO cross section is partitioned into the cross sections for a multitude of final states of arbitrary multiplicity, with their sum being equal to the cross section 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 collinear approximation. In fact, because of strong ordering, a radiated parton cannot be collinear to more than one parton in the amplitude, and this suppresses interference effects.

We now illustrate the basic parton-shower algorithm, as first introduced in Ref. [10]. (For more pedagogical introductions see Ref. [11] and references therein.) For simplicity, we consider the example of e^+e^- annihilation into $q\bar{q}$ pairs, where we only have to deal with final state radiation (FSR). We consider all final states that can be built by dressing the q and \bar{q} partons with an indefinite number of splitting processes. By recursively clustering together final state parton pairs with the smallest relative hardness, from each final state configuration we can construct two trees rooted at the q and \bar{q} partons. The momenta of all intermediate lines of the tree diagrams are then uniquely determined from the final-state momenta. Hardnesses in the trees are 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. 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

$$\Delta_i(t_m, t_n) \frac{\alpha_s(t)}{2\pi} P_{i,jk}(z) \frac{dt_m}{t_m} dz \frac{d\phi}{2\pi} \quad (43.3)$$

for each intermediate line arising from the n^{th} and ending in the m^{th} splitting vertex. $\Delta(t_m, t_n)$ is the so-called Sudakov form factor

$$\Delta_i(t_m, t_n) = \exp \left[- \int_{t_m}^{t_n} \frac{dq^2}{q^2} \frac{\alpha_s(q^2)}{2\pi} \int dz \frac{d\phi}{2\pi} \sum_{jk} P_{i,jk}(z) \right]. \quad (43.4)$$

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 [12] splitting kernels. Notice that the endpoints on the z integration depend upon the definition of hardness. For example, in case of virtuality or transverse momentum ordering, the z integration is automatically cut-off near the extremes, see eq. (1.2). When this is not the case (as, for example, for angular ordering) an explicit cut-off on z must be introduced, corresponding to the requirement that an emission must have some minimum energy to be distinguishable from no emission. For lines originating at the primary vertex, the scale t_n is replaced by the typical scale of the primary process and for lines ending without any further splitting the scale t_m is replaced by t_0 , 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).

Eq. (43.3) can be obtained by iterating formula Eq. (43.1) recursively, with two important corrections: a) the strong coupling is evaluated at a scale corresponding to the hardness of the splitting process; b) the presence of the Sudakov form factor. Both these modifications arise from the inclusion of all collinear-dominant virtual corrections.

Notice that the Sudakov form factor for a small hardness interval $\Delta_i(t, t + \delta t)$ is equal to one minus the integrated emission probability of Eq. (43.3), i.e. it can be interpreted as the probability of no emission in the interval $t, t + \delta t$. From this, it immediately follows that $\Delta_i(t_m, t_n)$ can be interpreted as the no-emission probability in the full t_m, t_n interval. This interpretation allows one to formulate the shower process as a probabilistic algorithm. We first notice that $0 < \Delta_i(t_m, t_n) \leq 1$, where the upper extreme is reached for $t_m = t_n$, and the lower extreme is approached for $t_m = t_0$. Starting from each of the partons in the primary process (e.g., $e^+e^- \rightarrow q\bar{q}$), event generation then proceeds recursively as follows. Given a parton exiting a vertex with hardness t_n , (taken to be of order the annihilation scale Q^2 for the first branching) one seeks a solution of the equation $r = \Delta_i(t_m, t_n)$, with $r \in [0, 1]$ a uniform random number, and solves it for the hardness of the next branching t_m . If $t_m \leq t_0$, no splitting is generated and the line is interpreted as a final parton. If $t_m > t_0$, a branching is generated at the scale t_m . Its z value and the final parton species jk are generated with a probability proportional to $P_{i,jk}(z)$. If angular correlations are neglected (see Sec. 43.1.1) the azimuth can be 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 initial 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 splitting, or globally for the entire set of partons (a procedure called *momentum reshuffling*).

We emphasize that the shower cross sections described above can be derived from perturbative QCD by keeping only the collinear-dominant real and virtual contributions to the cross section. As such it is unproductive for large-angle radiation. It is thus unsafe to rely upon Parton Shower Monte Carlo alone to compute backgrounds to new physics signals that are characterized by several widely separated jets.

A Shower Monte Carlo builds its final state as if it developed from an iterative process, often with each intermediate stage made available to the user. It should be remarked that the meaning of these intermediate stages is only relevant within the approximation adopted by the generator, and could also differ in different implementations.

43.1.1 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 [13]. A method to partially account for these effects was introduced in Ref. [14], 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 PYTHIA [4]. A method that fully includes spin correlation effects was later proposed [15], and has been implemented in HERWIG [3, 16]. For more recent developments, see refs. [17] and references therein.

43.1.2 Initial-state radiation

Initial-state radiation (ISR) arises because incoming particles may undergo collinear radiation before entering the hard-scattering process. In doing so, they acquire a non-vanishing transverse momentum, and their virtuality becomes negative (spacelike). In order to efficiently sample the hard collision, it turns out to be convenient to develop the ISR shower starting with the highest

hardness (i.e. with the hard process) and ending with the smallest (i.e. with the incoming parton in the hadron). Unlike the case of FSR, however, hardness ordering is opposite to time ordering in the ISR case. A corresponding backwards-evolution algorithm was formulated by Sjöstrand [18], and was basically adopted in all shower models. It can be illustrated by considering a primary interaction initiated by a quark where no collinear emissions of hardness $\geq t$ have taken place, and the same process where the quark also emits a collinear gluon of hardness t . The respective cross sections are proportional to

$$|M_q(x)|^2 dx f_q(x, t), \quad \text{and} \quad |M_q(x)|^2 dx \frac{\alpha_s(t)}{2\pi} f_q(x/z, t) P_{q,qq}(z) dz \frac{d\phi dt}{2\pi t}. \quad (43.5)$$

Here f_q is the quark PDF in the incoming hadron, x is the fraction of momentum of the incoming quark that enters the basic process, while x/z is the fraction of momentum of the incoming quark *before* it emits the collinear gluon. The elementary *emission probability* is the ratio of the second over the first expression in Eq. (43.5). In analogy with the final state radiation case, this ratio will appear in the exponent of the Sudakov form factor, that (after the inclusion of all splitting subprocesses) is given by

$$\Delta_i^{\text{ISR}}(t, t') = \exp \left[- \int_{t'}^t \frac{dt''}{t''} \frac{\alpha_s(t'')}{2\pi} \int_x^1 \frac{dz}{z} \sum_{jk} P_{j,ik}(z) \frac{f_j(t'', x/z)}{f_i(t'', x)} \right]. \quad (43.6)$$

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. [11]. 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.

43.1.3 Soft emissions and QCD coherence

Soft singularities arise in QCD due to the real or virtual emission of soft gluons. For example, the cross section for the emission of a soft gluon in e^+e^- annihilation into hadrons is given by

$$d\sigma_{q\bar{q}g} \approx d\sigma_{q\bar{q}} C_F (4\pi\alpha_s) \left[\frac{2 p_q \cdot p_{\bar{q}}}{p_q \cdot l p_{\bar{q}} \cdot l} \right] \frac{d^3l}{2l^0 (2\pi)^3} = d\sigma_{q\bar{q}} \frac{\alpha_s}{2\pi} C_F \frac{dl^0}{l^0} \frac{d\phi}{2\pi} \frac{d\cos\theta}{1 - \cos^2\theta}, \quad (43.7)$$

where p_q , $p_{\bar{q}}$ and l are the quark, antiquark and gluon momentum, and θ and ϕ are the polar and azimuthal angle of the gluon momentum with respect to the quark direction. Since the gluon is soft, we may assume that p_q and $p_{\bar{q}}$ are unaffected by the gluon emission. The soft singularity is manifest in the dl^0/l^0 factor. Notice that also collinear singularities are present at the same time when $\theta \rightarrow 0$ and $\theta \rightarrow \pi$, corresponding to the gluon becoming collinear to either the quark or the antiquark. It is easy to check that in the collinear limits Eq. (43.7) becomes equivalent to Eq. (43.1) with $P_{q,qq}(z) = C_F 2/(1-z)$, i.e. the limiting form of $P_{q,qq}(z)$ when z approaches 1. Thus, soft singularities coexist with collinear ones, so that two potentially large logarithms can arise simultaneously due to gluon emission.

Unlike the case of collinear emission, soft emission is not tied to a single emitting particle. The amplitude for the emission of a soft gluon from an external (incoming or outgoing) line with

momentum p is proportional to $p \cdot \epsilon/p \cdot l$. When squaring the amplitude, products like the one appearing in the square bracket of Eq. (43.1) arise for all pairs of external particles, with the product of a single emission amplitude with itself appearing only if $p^2 > 0$, i.e. for massive colored particles. Thus, interference plays here a crucial role. This is unlike the case of collinear singularities, where because of strong ordering a radiated parton cannot be collinear to more than one other parton.

It was shown in a set of publications (see Ref. [19]) that, within the conventional parton-shower formalism based on collinear factorization, the region of collinear and soft emissions can be correctly described by using the angle of the emissions as the ordering variable, rather than the virtuality, and by setting the argument of α_s at the splitting vertex equal to 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 default evolution variable in all versions of HERWIG (see Ref. [20]). To partially account for soft interference effects, an angular veto is imposed on the virtuality-ordered evolution in PYTHIA 6 [21]. In modern designs this type of fix is avoided, the soft problem being addressed from the start.

A radical alternative formulation of QCD cascades first proposed in Ref. [22] 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 simultaneously achieve the correct logarithmic structure for both the collinear and soft emissions in the so called leading color approximation, i.e. when terms suppressed by a power of the number of colors are neglected. The ARIADNE [23] and VINCIA [24] programs are based on this approach. Dipole-type showers [25] are also used by default in SHERPA [26] and exist as an option in HERWIG [27]. An alternative dipole-based model is available in PYTHIA and SHERPA via the DIRE [28] plugin. The p_\perp -ordered showers in PYTHIA 6 and 8 represent a hybrid, combining collinear splitting kernels with dipole kinematics [29].

43.1.4 Resummation

Shower Monte Carlo generators perform resummation of all-order collinear- and/or soft-enhanced perturbative contributions, and it is thus natural to compare them to QCD resummation calculations [7]. The latter start from the definition of specific infrared-safe observables, that develop towers of large logarithms in certain regions of phase space, typically organized as

$$A(\alpha_s) \exp[Lg_0(\alpha_s L) + g_1(\alpha_s L) + \alpha_s g_2(\alpha_s L) + \dots] \quad (43.8)$$

where L is the large logarithm, α_s is the strong coupling constant evaluated at some hard scale. The functions A has a power expansion in α_s starting with a constant, and the functions g_i have a power series expansion in their argument starting with the linear term. We talk about N^n LL accuracy if all g_i functions for $i \leq n$ are fully included in the resummation formula, and the A function is given up to the order $n - 1$ in α_s . In particular we talk about LL accuracy if the function g_0 is included, NLL accuracy if also g_1 is included. In both cases only the constant term in A is needed, while for N^2 LL (or NNLL) g_0 , g_1 and g_2 are included and A is given up to order α_s .

In general, a dedicated resummation calculation must be performed for each new observable. The predictions of shower MCs, on the other hand, are cast in terms of complete sets of final-state momenta, on which one can evaluate any observable; i.e., the shower algorithm itself is normally independent of the specific observable(s) under study. Because of this, it is not easy to qualifying the accuracy of a shower MC using the same criteria adopted in resummation calculation. In spite

of this fact, shower MCs perform generally quite well in the description of observables that require resummation. This is related to their inclusion of several universal but formally subleading aspects. There is no guarantee, however, that the shower MC should perform to a certain level of accuracy for all distributions that require resummation. Thus, a more systematic treatment of subleading effects is quite desirable.

Several studies have appeared in the literature, aiming at either improving the current shower algorithms or formulate totally new ones, in such a way that the theoretical accuracy can be discussed in more precise terms [17, 30–35]. In particular, in refs [17, 30, 31] (the PanScales collaboration), it was shown that criteria for defining the accuracy of Shower Generators can indeed be given on a quite general ground. They identified two such criteria: the first one regards the approximations to multi-parton real matrix elements generated by the shower algorithm, that are required to be accurate in the region of phase space that are relevant for N^n LL accuracy. The second one regards the implementation of the virtual corrections, and it requires that for a large class of shape variables the shower MC must agree with the resummed calculation at the N^n LL level. The PanScales collaboration has constructed shower algorithms that are (with some restrictions) NLL accurate according to these criteria for e^+e^- annihilation, hadronic collisions, Deep Inelastic Scattering and Vector Boson Fusion processes (see [36] and references therein). Other NLL accurate (according to the same criteria) shower algorithms have been proposed in refs. [32] and [35].


43.1.5 Massive quarks

Quark masses act as a 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 radiation is strongly suppressed 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 [37], and in HERWIG and SHERPA, where a generalization of the Altarelli-Parisi splitting kernel is used for massive quarks [38]. For newer approaches for QCD evolution in processes with massive partons see ref. [39] and references therein.

43.1.6 Color information

In event generators, quarks and antiquarks are represented by color lines, with arrows indicating the direction of color flow. In the limit of infinitely many colors (called the leading color approximation), each such line can be associated with a unique label; the probability for two quarks (or antiquarks) to have the same color (anticolor) vanishes. Moreover, in the same limit gluons can be represented by a pair of color lines with opposite arrows, as can be realized e.g. from the $SU(3)$ group relation $8 = 3 \otimes \bar{3} \ominus 1$. 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. 43.1. It is also possible for a set of gluons to be connected cyclically in color, as e.g. in the decay $\mathcal{T} \rightarrow ggg$. The color information is used in angular-ordered showers, where




Figure 43.1: Color development of a shower in e^+e^- annihilation. Color-neutral clusters of partons are indicated by the dashed under-brackets.

the angle of color-connected partons (i.e. partons connected by the same color line) determines the initial angle for the shower development, and in dipole showers, where dipoles are formed by color-connected partons. It is also used in hadronization models, where the initial strings or clusters used for hadronization are formed by color-neutral clusters of partons.

43.1.7 Electromagnetic (and weak) corrections

The physics of photon emission from charged particles, and of photon splittings to pairs of charged fermions, 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 $\alpha_{\text{em}} \ln(Q/m) \ln(E_\gamma/E)$ for soft γ emissions, so that, especially for the case of electrons, their inclusion in the simulation process is mandatory. This is done in most GPMCs (for a comparative study see [40]), which typically agree in collinear (DGLAP) limits but differ in whether and how soft (multipole) QED interference effects are handled. The specialized generator PHOTOS [41] is sometimes used as an afterburner for an improved treatment of QED radiation in non-hadronic resonance decays.

For photon emissions off leptons, the shower can be continued down to virtualities arbitrarily close to the lepton mass shell (unlike the case in QCD). In practice, an infrared cutoff is still required for the shower algorithm to terminate. Therefore, there is always an energy cut-off for emitted photons that depends upon the implementations [40]. 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, there is 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. In PYTHIA, this is done via interfacing with PHOTOS, while HERWIG and SHERPA contain their own dedicated implementations [42, 43] based upon the YFS formalism [44], which is also adopted in the KKMC-hh Monte Carlo [45].

At energies significantly above the electroweak (EW) scale, showers involving emissions (and splittings) of weak gauge and Higgs bosons can also be relevant, with recent implementations based on collinear limits ranging from weak-boson emissions from fermions in PYTHIA 8 [46], to triple-boson couplings in HERWIG 7 [47], to the full set of EW branching processes in VINCIA [48].

43.1.8 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 [49, 50]. Notable exceptions are long-lived colored particles [51], 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. 43.1.9) and the assumed threshold behavior.

In addition to code-specific implementations [11], 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 [52], 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 [53], 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 [54]. Unfortunately, at present no standard method for handling multiple shower starting scales has been defined in the LHEF, so that further code, specific to the given GPMC, needs to be added to the interface in this case.

43.1.9 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 an 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 instead be treated 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, 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 . We finally point out that NLO+PS generators have appeared that can deal with resonances including off-shell effects, non-resonance contributions and interference of radiation generated in resonance decay and production, see [55] and references therein. A new “interleaved” shower treatment also allows for interference between production and decay in the VINCIA model [48].

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 [56] 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 $H \rightarrow WW^*$.

GPMC manuals often give instructions on how to include new decay modes, at varying levels of

sophistications ranging from simple uniform phase-space sampling (which the user can reweight a posteriori) and step-function thresholds, to fully matrix-element weighted decay implementations including potential off-shell / threshold effects.

During subsequent showering of the decay products, most parton-shower models will preserve the total invariant mass of the decayed resonance, so as not to skew the original resonance shape. In the context of passing externally generated LHEF files [52] to a GPMC for showering, note that this is only possible if the intermediate resonances are present (with status code 2) in the LHEF event record [57].

43.1.10 Matching and Merging with Fixed-Order 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. [58] with later developments summarized in Ref. [11]. 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 in [59].

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 summarize 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 [60], and several variants have appeared, including the CKKW-L, MLM, pseudoshower and MESS methods, see Refs. [11, 61, 62] and references therein. So called “Truncated Showers” are required [63] 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 [64].

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 often appear as backgrounds to new-physics searches. These jets are typically required to be well sepa-

rated, and to have large transverse momenta. These kinematic 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 [65], MADGRAPH [66], and others summarized in Ref. [61], 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) [52, 57]. 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 [67] and POWHEG [63, 68], with several alternative methods now also being pursued, see [69] and references therein, and the recent method presented in [70].

NLO+PS generators produce NLO accurate distributions for inclusive quantities, and generate the hardest jet with tree-level accuracy. 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 [67], together with the AMC@NLO development [71], and in the POWHEG-BOX framework [68]. HERWIG supports its own variants of POWHEG and MC@NLO 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 FXFX [72], UNLOPS [73], MiNLO [74] and MEPS@NLO [75] methods address this problem. The solution of this problem is also a prerequisite for the construction of NNLO+PS generators, i.e. generators that, besides being NLO accurate for the production of an associated jet, are also NNLO accurate for fully inclusive observables. Three different approaches have appeared in the literature for dealing with this problem: the UN²NLOPS, based upon UNLOPS method [76]; methods extending the MiNLO approach by reweighting [77] and by the so called MiNNLO_{PS} technique [78]; and the GENEVA method [79]. Several processes have been implemented with the MiNLO related and GENEVA methods, mostly for the production of color singlet systems, with the exception of the recent MiNNLO_{PS} implementation of $t\bar{t}$ production (see refs. [69, 80, 81] for an extensive list of references).

43.2 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 “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 equal 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.

An important result in “quenched” lattice QCD (see Chap. 17 of PDG book) is that the potential energy between two partons with opposite color charges grows linearly with their separation, 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. 43.2.1. Alternatively, a property of perturbative QCD called “preconfinement” is the basis of the *cluster model of*

hadronization, discussed in Sec. 43.2.2.

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. 19 of PDG book) is that FFs can be defined at an arbitrary perturbative scale Q while MC hadronization models are intrinsically defined at the scale Q_{had} . Direct comparisons are therefore only meaningful 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 parameterization 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. 43.1.6) to determine between which partons confining potentials should arise. E.g., in the string picture, the nonperturbative limit of a QCD dipole is a string piece [82].

Given an ideal, exact hadronization model, its dependence on the 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 nonperturbative model parameters must be “retuned” by hand if the infrared cutoff is modified. Any other changes to the perturbative part of the calculation, such as matching to further (fixed-order or resummed) coefficients, may also necessitate a retuning. Tuning is discussed briefly in Sec. 43.4.

The impact of hadronization corrections in a GPMC can be estimated by switching off hadronization. It should be emphasized, however, that the so-called “parton level” obtained this way is not a fully universal concept. This is because Q_{had} is typically defined in a model-dependent way, via a cutoff in the respective shower model’s evolution variable, sometimes with supplementary cutoffs, e.g., in invariant mass, p_T , or the like and different tunes possibly using different numerical values for the cutoff(s). Variations of the main shower cutoff scale (e.g., by a factor of 2) may give an indication of the level of model dependence, which may be mild for sufficiently inclusive observables. Nevertheless our recommendation is to avoid correcting experimental measurements to “parton level” as the intrinsic model dependence will always imply a loss of accuracy.

43.2.1 The String Model

Starting from early concepts [83], several hadronization models based on strings have been proposed [11]. Of these, the most widely used today is the so-called Lund model [84, 85], implemented in PYTHIA [4, 5]. We concentrate on that particular model here, though many of the overall concepts would be shared by any string-inspired method.

Consider a quark and an antiquark that have a large relative momentum and which are in an overall color-singlet state, such as the $\bar{q}q$ pair produced at the end of the shower in the center of Fig. 43.1. 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} . As the string grows, the nonperturbative creation of quark-antiquark pairs can break the string, via the process illustrated in Fig. 43.2. If either of the resulting string pieces after the break still has a large invariant mass, new breaks will then continue to occur until every piece has a mass of the order of that of a typical hadron. In nature, quantum mechanics




Figure 43.2: Illustration of string breaking by quark pair-creation in the string field.

ensures that these masses are (Breit-Wigner-distributed around) those of the physical hadron states in QCD, while in an algorithmic implementation, this is less trivial to impose. Within the Lund model, the constraint that each final string piece must have a mass consistent with that of physical hadron is addressed by first noting that each breakup vertex is spacelike separated from all others (the string cannot break again in the forward lightcone of a point where it has already broken); since the breakups are therefore causally disconnected, they do not have to be considered in any particular order. The Lund model exploits this to perform the fragmentation in an outwards-in manner [85], splitting off one physical hadron at a time from either the left or right endpoint (chosen randomly in each step).

The first hadron to be generated (starting either from the left or right) is thus the “outermost” one (sometimes called the “first-rank” hadron), formed by combining the original hadronizing endpoint quark (or antiquark) q_0 with an antiquark (or quark) \bar{q}_1 produced by the first breakup. The new leftover quark (or antiquark) q_1 becomes the string endpoint for the next iteration, in a Markov chain which continues, alternating randomly between the left and right ends of the string, until finally a small last bit of string is decayed directly to two hadrons, with no energy left over.

Not only does this allow for the hadron that is split off in each step to be assigned a physical mass, the fact that the fragmentation spectrum should be independent of whether one performs the fragmentation from left to right or vice versa places a strong constraint on the form of the nonperturbative *string fragmentation function*¹, $f(z)$, which governs the probability for the hadron produced in a given step to take a fraction $z \in [0, 1]$ of the remaining energy. Thus, in the Lund model, causality dictates that the fragmentation function should be of the form,

$$f(z) \propto \frac{1}{z} (1-z)^a \exp\left(-\frac{b(m_h^2 + p_{T,h}^2)}{z}\right). \quad (43.9)$$

This 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 [85] which determines the behavior in the soft limit, $z \rightarrow 0$. Note that the dependence on the hadron mass, m_h , in $f(z)$ implies that heavier hadrons have higher $\langle z \rangle$. We return to the transverse momentum $p_{T,h}$ below.

The model is Lorentz invariant, so considerations involving boosted string systems are straightforward, involving the usual Lorentz effects.

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 $\tau_0 \sim 10^{-23}\text{s}$ [85].

¹Note that the string fragmentation function differs from the parton-to-hadron fragmentation functions discussed earlier. In principle, the two kinds should be related (with the latter evolved to a low scale of the order of the hadronization scale), but the relationship is not trivial and goes beyond the scope of this review.

For massive endpoints (e.g. c and b quarks), 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 [86]. Although different forms, such as the Peterson formula [87], can also be used to describe inclusive heavy-meson spectra (see Sec 19.8 of PDG book), such choices are not strictly consistent with causality in the string framework.

43.2.1.1 Strings with Gluons

In the string model, energetic gluons lead to transverse “kinks” on strings, illustrated in Fig. 43.3. The order of these kinks follows from the color ordering produced by the parton shower, cf. the $\bar{q}ggq$ and $\bar{q}gq$ systems on the left and right part of Fig. 43.1. (Modifications to this order, by possible color reconnection/rearrangement effects, are discussed in Sec. 43.3.2.) Thus gluons




Figure 43.3: Schematic illustration of an $e^+e^- \rightarrow qq\bar{q}$ configuration emerging from the parton shower (PS). Snapshots of string positions are shown at two different times (full and shaded lines respectively). The gluon forms a transverse kink which grows in the y direction until all the gluon’s kinetic energy has been used up.

effectively build up a transverse structure in the originally one-dimensional object, with infinitely soft ones smoothly absorbed into the string. Note: cyclic topologies made entirely of gluons (closed strings) are also possible, e.g. in decays such as $H \rightarrow gg$ or $\Upsilon \rightarrow ggg$. The space-time evolution is more involved when kinks are taken into account [85], but no additional free parameters need to be introduced. The main difference between quark and gluon hadronization stems from the fact that gluons are connected to two string pieces (one on either side), while quarks are only connected to a single string piece. Hence, the relative rate of energy loss per unit invariant time — and consequently also the rate of hadron production — is larger by a factor of 2 for gluons (analogously to the ratio of gluon to quark color charges $C_A/C_F = 2.25$).

43.2.1.2 Transverse Momentum and Flavors

For each breakup vertex, quantum mechanical tunneling is assumed to control the masses and p_T kicks (transverse to the string axis, in a frame in which the string itself has no transverse motion) that can be produced, leading to a Gaussian suppression

$$\text{Prob}(m_q^2, p_{T,q}^2) \propto \exp\left(\frac{-\pi m_q^2}{\kappa}\right) \exp\left(\frac{-\pi p_{T,q}^2}{\kappa}\right), \quad (43.10)$$

where m_q is the mass of the produced quark flavor and p_T is the nonperturbative transverse momentum imparted to it by the breakup process, with a universal average value of $\langle p_{T,q}^2 \rangle = \kappa/\pi \sim (250 \text{ MeV})^2$. The antiquark has the same mass and opposite p_T .

In an MC model with a fixed shower cutoff t_0 , the effective amount of p_T in string breaks may be larger than the purely nonperturbative κ/π above, to account for effects of additional (unresolved) radiation below t_0 .

From the mass term in Eq. (43.10), one concludes that charm and bottom quarks are too heavy to be produced in string breaks, while strange quarks will be suppressed relative to up and down

ones. Lacking unambiguous and precise mass definitions for light quarks, however, the effective amount of strangeness suppression is normally extracted from experimental data, using observables such as K/π , K^*/ρ , and ϕ/K^* ratios.

Baryon production can also be incorporated, at various levels of sophistication. The simplest option is to allow string breaks to produce a diquark-antidiquark pair, where a diquark is a loosely bound states of two quarks in an overall $\bar{3}$ representation. Again, the relative rate of such pairs is extracted from data, e.g. using p/π or Λ/K ratios. 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 in the shower. More sophisticated options, including the so-called “popcorn” mechanism, are discussed in Ref. [85]. Finally, the PYTHIA framework also allows for baryon string junctions [88], in which three string pieces (which can be thought of as red, green, and blue, respectively) join together. These are used, e.g., to model the fragmentation of baryon beam remnants, and can also be created (in pairs of junctions and antijunctions) in some color-reconnection scenarios [89], making the effective baryon-to-meson ratios in such models dependent on the amount and type of color reconnections that occur in each event

The next step of the algorithm is the assignment of the produced quarks within hadron multiplets. Using a nonrelativistic classification of spin states, the hadronizing q may combine with the \bar{q}' from a newly created breakup to produce a hadron 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 framework², but individual rates are not predicted by the model, and hence the ratios of spin-1 to spin-0 mesons and diquarks are free parameters (modulo a factor 3 from spin counting) which must be constrained by data. In that context, it is often 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. 43.2.3.

43.2.2 The Cluster Model

The cluster hadronization model is based on *preconfinement*, i.e., on the observation [90,91] 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 which is power suppressed at large masses. For any starting scale $Q \gg Q_0 \gg \Lambda_{\text{QCD}}$, only the number of such clusters depends on Q , while the shape of their mass distribution only depends on Q_0 and on Λ_{QCD} .

Following early models based on this universality [10,92], the cluster model developed by Weber [93] has for many years been a hallmark of the HERWIG generators, with an alternative implementation [94] 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, with the relative yields of different hadron species mainly governed by their masses and the size of the phase space.

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_{T} 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

²In PYTHIA, the four $L = 1$ meson multiplets (scalar, tensor, and 2 pseudovectors) are also available but are disabled by default, since many states (and their decays) are poorly known.

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 nonperturbative 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 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.

43.2.3 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 measurement. (A typical hadron-collider definition of a “stable particle” $c\tau \geq 10$ mm includes weakly-decaying strange hadrons K , Λ , Σ^{\pm} , $\bar{\Sigma}^{\pm}$, Ξ , Ω .) 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). This interplay also implies that hadronization parameters may need to 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 [11]). 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 [95] for hadron decays and TAUOLA [96] for τ decays.

More recently, HERWIG and SHERPA include helicity-dependence in τ decays [6, 97], with a more limited treatment available in PYTHIA 8 [5]. 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. 43.1.7.

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.

43.3 Models for Soft Hadron-Hadron Physics

43.3.1 Underlying Event

In the GPMC context, “underlying event” (UE) denotes any additional activity *beyond* the basic process and its associated ISR and FSR activity. The UE is thus only defined in the context of events selected with a “hard” (i.e., high- p_\perp) trigger which defines the basic process at hand. The dominant contribution to the UE is believed to come from additional color exchanges between the colliding hadronic states. These multiple exchanges can be modeled either as additional perturbative (mainly t -channel gluon) exchanges, called multiple parton-parton interactions (MPI), or nonperturbatively using so-called cut pomerons (roughly equivalent to exchange of gluons with $p_\perp \rightarrow 0$). The experimental observation that events with a hard trigger are accompanied by a higher-than-average level of associated activity 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. Typically, these are QCD $2 \rightarrow 2$ interactions, which produce additional 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. [98], and with some variation this still forms the basis for most modern implementations. Some useful additional references can be found in Ref. [11]. 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_T , causing the differential cross sections to behave roughly as

$$d\sigma_{2 \rightarrow 2} \propto \frac{dt}{t^2} \sim \frac{dp_T^2}{p_T^4}. \quad (43.11)$$

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

$$\sigma_{2 \rightarrow 2} = \langle n \rangle \sigma_{\text{tot}}, \quad (43.12)$$

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

$$\mathcal{P}_n = \langle n \rangle^n \frac{\exp(-\langle n \rangle)}{n!}. \quad (43.13)$$

This simple argument expresses unitarity; instead of the total hadron-hadron interaction cross section diverging as the parton-parton $p_T \rightarrow 0$ (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. [99, 100]). One therefore expects the average number of perturbative MPI to exceed unity at around that scale.

Two 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_T , 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_T and x , is color screening; if the wavelength $\sim 1/p_T$ 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 $p_T \rightarrow 0$. 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 [101] and is perceived of simply as an effective cutoff parameter. The simplest choice is to introduce a step function $\Theta(p_T - p_{\perp \text{min}})$. Alternatively, one may note that the jet cross section is divergent like $\alpha_s^2(p_T^2)/p_T^4$, cf. Eq. (43.11), and that therefore a factor

$$\frac{\alpha_s^2(p_{T,0}^2 + p_T^2)}{\alpha_s^2(p_T^2)} \frac{p_T^4}{(p_{T,0}^2 + p_T^2)^2} \quad (43.14)$$

would smoothly regulate the divergences, now with $p_{T,0}$ as the free parameter. Regardless of whether it is imposed as a smooth (PYTHIA and SHERPA) or step (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 [102].

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_T activity, whereas central collisions are more likely to contain at least one hard scattering; a high- p_T triggered sample will therefore be biased towards small impact parameters, b . The ability of a model to describe both the UE in events with a hard trigger as well as the activity in inclusive “minimum-bias” (MB) samples (see below) 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. (43.13), 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 [11], 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_T and/or low x . Although none of the general purpose 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 [29] 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 [88]. Fully correlated double-parton evolution is so far only available in the more specialised code dShower [103]; see [104] for comparisons between this and Pythia’s interleaved model.

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 [105]. In the initial state, parton rescattering/recombination effects have so far not been included in any of the GPMC models.

43.3.2 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 [106]. In hadron-hadron (and nucleus-nucleus) collisions, however, BE correlations are used extensively to study the space-time structure of hadronizing matter (“femtoscropy”).

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 in each step, but also makes it intrinsically difficult to introduce correlations among particles from different breaks/decays. In GPMCs, a few semi-classical models are available within the PYTHIA 6 and 8 generators [107], 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. Variants of this model differ mainly by the assumed shape of the correlation function and how overall momentum conservation is handled.

As discussed in Sec. 43.2, 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 nonperturbative 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_T), 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_T distributions, rapidity gaps, and precision mass measurements. Reviews of color-reconnection effects can be found in [11, 89].

43.3.3 Minimum-Bias Events 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 ($hh \rightarrow h + \text{gap} + X + \text{gap} + h$). Note that different terminologies exist [108]: in experimental settings, diffraction is typically defined by an observable gap, of some minimal size in rapidity, while in the MC context, each diffractive physics process 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 [98] (PYTHIA and SHERPA), or by regulating the QCD cross sections with a sharp cutoff [109] and adding a separate class of nonperturbative scatterings below that scale [110] (HERWIG). See also Sec. 43.3.1. In all cases, the 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 [88] and/or x dependence [111], and 3) additional soft-QCD effects such as color reconnections, discussed in Sec. 43.3.2.

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 parameterized cross sections $\propto dM^2/M^2$ [112]. This mass is represented as a string, which is hadronized as described in Sec. 43.2.1, though differences in the effective scale of the hadronization may necessitate a (re)tuning of the hadronization 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 [113], allowing diffractive jet production in-

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

cluding showers and underlying event [114]; 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) [115–117]. 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. Hadronization 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 [118], and an effort is underway to construct an MC implementation of the “KMR” model [101] within the SHERPA generator. Color reconnections (Sec. 43.3.2) may also play a role in creating rapidity gaps and the underlying event (Sec. 43.3.1) in filling them.

43.4 Uncertainties and Tuning

The accuracy that can be achieved by a GPMC model depends on the sophistication of the theory models it incorporates, on the available constraints on its free parameters, and on the nature of the observable(s) under study. Using existing data to constrain the model parameters is referred to as generator tuning. Although tuned models do tend to yield improved results also for observables that they have not been tuned to, the question of evaluating the remaining uncertainties reliably is still far from solved. It is worth noting, however, that all of the GPMCs now provide options for automatic evaluation of perturbative shower uncertainties (e.g., via renormalization-scale variations), in the form of vectors of alternative event weights [119–121]. One must be aware that these variations are not necessarily exhaustive and significant weight fluctuations can be a problem for processes with large shower phase spaces. Nonperturbative uncertainties are normally still evaluated by varying salient model parameters by hand, though this looks set to change with a new method based on reweighting [122]. A general method called eigentunes [123] is also available, based on global fits to data.

Typically, the overall event properties are determined by only a few, very important parameters, such as the effective value of α_s , for perturbative corrections, and the shape of the fragmentation functions, for nonperturbative ones. More parameters may then be introduced to describe successively more detailed aspects (e.g., the rates and decays of individual hadron species), but these should have progressively less impact on the overall modeling. One may therefore take a factorized approach, first constraining the perturbative parameters and thereafter the nonperturbative 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, measurements sensitive to global event properties would typically be applied first, to constrain the most inclusive parameters, and so on for progressively more exclusive aspects. For high- p_T processes, one should be aware that if the perturbative modelling is not adequate, there is a risk that one may end up effectively mistuning nonperturbative parameters to compensate for perturbative deficiencies. This is obviously wrong and care must be taken to avoid it. E.g., it would not make sense to tune to get the right total number of hadrons unless one also gets the right total number of jets. In this context, it may be more desirable to use ratios like the number of hadrons *per jet*.

Note also that the unitarity of shower and hadronization models implies that the Born-level cross-section normalization is not tunable, hence in tuning contexts one tends to focus on the shapes of distributions rather than their normalizations.

At $\text{LO} \times (\text{N})\text{LL}$, perturbation theory is doing well if it agrees with an IR safe shape measurement within $\sim 10\%$. It would therefore not make much sense to tune a GPMC to agree with individual data points beyond this level. To reflect this, and to protect against overfitting, a 5% sanity limit has sometimes been enforced in tuning studies [99, 124]. The advent of NLO Monte Carlos may

reduce this number slightly, but only for quantities for which one expects NLO precision. For quantities governed by nonperturbative 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 for a class of relevant observables [125]. The question of LO vs. NLO PDFs is more involved [11], but it should be emphasized that the gluon PDF at (very) low x is important for determining the level of the underlying event in MPI models (Sec. 43.3.1), and hence the MB/UE tuning (and energy scaling [102]) 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. [99]. A useful online resource can be found at the mcplots.cern.ch website [126], based on the RIVET tool [127].

Recent years have seen the emergence of automated tools to reduce the amount of both computer and manpower required for tuning [123, 128, 129]. 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. The final result is therefore still subjective but at least reproducible. Backed by careful demonstrations of sensitivities, correlations, and uncertainties, competitive tunes can be produced. The field is still burgeoning, with future sophistications to be expected.

References

- [1] G. Corcella *et al.*, *JHEP* **01**, 010 (2001), [[hep-ph/0011363](#)].
- [2] M. Bahr *et al.*, *Eur. Phys. J. C* **58**, 639 (2008), [[arXiv:0803.0883](#)].
- [3] J. Bellm *et al.*, *Eur. Phys. J. C* **76**, 4, 196 (2016), [[arXiv:1512.01178](#)].
- [4] T. Sjostrand, S. Mrenna and P. Z. Skands, *JHEP* **05**, 026 (2006), [[hep-ph/0603175](#)].
- [5] C. Bierlich *et al.*, *SciPost Phys. Codebases* **8** (2022), [[arXiv:2203.11601](#)].
- [6] T. Gleisberg *et al.*, *JHEP* **02**, 007 (2009), [[arXiv:0811.4622](#)].
- [7] QCD summary, PDG.
- [8] T. Kinoshita, *J. Math. Phys.* **3**, 650 (1962).
- [9] T. Lee and M. Nauenberg, *Phys. Rev.* **133**, B1549 (1964).
- [10] G. C. Fox and S. Wolfram, *Nucl. Phys. B* **168**, 285 (1980).
- [11] A. Buckley *et al.*, *Phys. Rept.* **504**, 145 (2011), [[arXiv:1101.2599](#)].
- [12] G. Altarelli and G. Parisi, *Nucl. Phys. B* **126**, 298 (1977).
- [13] A. Einstein, B. Podolsky and N. Rosen, *Phys. Rev.* **47**, 777 (1935).
- [14] B. Webber, *Phys. Lett. B* **193**, 91 (1987).
- [15] J. C. Collins, *Nucl. Phys. B* **304**, 794 (1988).
- [16] I. Knowles, *Comput. Phys. Commun.* **58**, 271 (1990).
- [17] A. Karlberg *et al.*, *Eur. Phys. J. C* **81**, 8, 681 (2021), [[arXiv:2103.16526](#)].
- [18] T. Sjostrand, *Phys. Lett. B* **157**, 321 (1985).
- [19] G. Marchesini and B. Webber, *Nucl. Phys. B* **310**, 461 (1988).
- [20] S. Gieseke, P. Stephens and B. Webber, *JHEP* **12**, 045 (2003), [[hep-ph/0310083](#)].

- [21] M. Bengtsson and T. Sjostrand, *Nucl. Phys. B* **289**, 810 (1987).
- [22] G. Gustafson and U. Pettersson, *Nucl. Phys. B* **306**, 746 (1988).
- [23] L. Lonnblad, *Comput. Phys. Commun.* **71**, 15 (1992).
- [24] W. T. Giele, D. A. Kosower and P. Z. Skands, *Phys. Rev. D* **78**, 014026 (2008), [[arXiv:0707.3652](#)].
- [25] Z. Nagy and D. E. Soper, *JHEP* **10**, 024 (2005), [[hep-ph/0503053](#)].
- [26] S. Schumann and F. Krauss, *JHEP* **03**, 038 (2008), [[arXiv:0709.1027](#)].
- [27] S. Platzer and S. Gieseke, *Eur. Phys. J. C* **72**, 2187 (2012), [[arXiv:1109.6256](#)].
- [28] S. Höche and S. Prestel, *Eur. Phys. J. C* **75**, 9, 461 (2015), [[arXiv:1506.05057](#)].
- [29] T. Sjostrand and P. Z. Skands, *Eur. Phys. J. C* **39**, 129 (2005), [[hep-ph/0408302](#)].
- [30] M. Dasgupta *et al.*, *JHEP* **09**, 033 (2018), [Erratum: *JHEP* **03**, 083 (2020)], [[arXiv:1805.09327](#)].
- [31] M. Dasgupta *et al.*, *Phys. Rev. Lett.* **125**, 5, 052002 (2020), [[arXiv:2002.11114](#)].
- [32] J. R. Forshaw, J. Holguin and S. Plätzer, *JHEP* **09**, 014 (2020), [[arXiv:2003.06400](#)].
- [33] J. R. Forshaw, J. Holguin and S. Plätzer, *JHEP* **08**, 145 (2019), [[arXiv:1905.08686](#)].
- [34] Z. Nagy and D. E. Soper, *Phys. Rev. D* **99**, 5, 054009 (2019), [[arXiv:1902.02105](#)].
- [35] F. Herren *et al.* (2022), [[arXiv:2208.06057](#)].
- [36] M. van Beekveld and S. Ferrario Ravasio (2023), [[arXiv:2305.08645](#)].
- [37] E. Norrbin and T. Sjostrand, *Nucl. Phys. B* **603**, 297 (2001), [[hep-ph/0010012](#)].
- [38] S. Catani *et al.*, *Nucl. Phys. B* **627**, 189 (2002), [[hep-ph/0201036](#)].
- [39] B. Assi and S. Höche (2023), [[arXiv:2307.00728](#)].
- [40] J. Cembranos *et al.*, *JHEP* **09**, 077 (2013), [[arXiv:1305.2124](#)].
- [41] N. Davidson, T. Przedzinski and Z. Was, *Comput. Phys. Commun.* **199**, 86 (2016), [[arXiv:1011.0937](#)].
- [42] K. Hamilton and P. Richardson, *JHEP* **07**, 010 (2006), [[hep-ph/0603034](#)].
- [43] M. Schonherr and F. Krauss, *JHEP* **12**, 018 (2008), [[arXiv:0810.5071](#)].
- [44] D. R. Yennie, S. C. Frautschi and H. Suura, *Annals Phys.* **13**, 379 (1961).
- [45] S. Jadach *et al.*, *Phys. Rev. D* **94**, 7, 074006 (2016), [[arXiv:1608.01260](#)].
- [46] J. R. Christiansen and T. Sjöstrand, *JHEP* **04**, 115 (2014), [[arXiv:1401.5238](#)].
- [47] M. R. Masouminia and P. Richardson (2021), [[arXiv:2108.10817](#)].
- [48] H. Brooks, P. Skands and R. Verheyen (2021), [[arXiv:2108.10786](#)].
- [49] A. Semenov, *Comput. Phys. Commun.* **180**, 431 (2009), [[arXiv:0805.0555](#)].
- [50] N. D. Christensen and C. Duhr, *Comput. Phys. Commun.* **180**, 1614 (2009), [[arXiv:0806.4194](#)].
- [51] M. Fairbairn *et al.*, *Phys. Rept.* **438**, 1 (2007), [[hep-ph/0611040](#)].
- [52] J. Alwall *et al.*, *Comput. Phys. Commun.* **176**, 300 (2007), [[hep-ph/0609017](#)].
- [53] P. Z. Skands *et al.*, *JHEP* **07**, 036 (2004), [[hep-ph/0311123](#)].
- [54] J. Alwall *et al.* (2007), [[arXiv:0712.3311](#)].
- [55] T. Ježo *et al.*, *Eur. Phys. J. C* **76**, 12, 691 (2016), [[arXiv:1607.04538](#)].
- [56] P. Richardson, *JHEP* **11**, 029 (2001), [[hep-ph/0110108](#)].

- [57] E. Boos *et al.*, in “2nd Les Houches Workshop on Physics at TeV Colliders,” (2001), [[hep-ph/0109068](#)].
- [58] M. Bengtsson and T. Sjostrand, *Phys. Lett. B* **185**, 435 (1987).
- [59] W. Giele, D. Kosower and P. Skands, *Phys. Rev. D* **84**, 054003 (2011), [[arXiv:1102.2126](#)].
- [60] S. Catani *et al.*, *JHEP* **11**, 063 (2001), [[hep-ph/0109231](#)].
- [61] J. Alwall *et al.*, *Eur. Phys. J. C* **53**, 473 (2008), [[arXiv:0706.2569](#)].
- [62] H. Brooks and C. T. Preuss, *Comput. Phys. Commun.* **264**, 107985 (2021), [[arXiv:2008.09468](#)].
- [63] P. Nason, *JHEP* **11**, 040 (2004), [[hep-ph/0409146](#)].
- [64] B. Cooper *et al.*, *Eur. Phys. J. C* **72**, 2078 (2012), [[arXiv:1109.5295](#)].
- [65] M. L. Mangano *et al.*, *JHEP* **07**, 001 (2003), [[hep-ph/0206293](#)].
- [66] J. Alwall *et al.*, *JHEP* **06**, 128 (2011), [[arXiv:1106.0522](#)].
- [67] S. Frixione and B. R. Webber, *JHEP* **06**, 029 (2002), [[hep-ph/0204244](#)].
- [68] S. Alioli *et al.*, *JHEP* **06**, 043 (2010), [[arXiv:1002.2581](#)].
- [69] J. M. Campbell *et al.*, in “Snowmass 2021,” (2022), [[arXiv:2203.11110](#)].
- [70] P. Nason and G. P. Salam, *JHEP* **01**, 067 (2022), [[arXiv:2111.03553](#)].
- [71] J. Alwall *et al.*, *JHEP* **07**, 079 (2014), [[arXiv:1405.0301](#)].
- [72] R. Frederix and S. Frixione, *JHEP* **12**, 061 (2012), [[arXiv:1209.6215](#)].
- [73] L. Lönnblad and S. Prestel, *JHEP* **03**, 166 (2013), [[arXiv:1211.7278](#)].
- [74] K. Hamilton, P. Nason and G. Zanderighi, *JHEP* **10**, 155 (2012), [[arXiv:1206.3572](#)].
- [75] S. Hoeche *et al.*, *JHEP* **04**, 027 (2013), [[arXiv:1207.5030](#)].
- [76] S. Höche, Y. Li and S. Prestel, *Phys. Rev. D* **91**, 7, 074015 (2015), [[arXiv:1405.3607](#)].
- [77] K. Hamilton *et al.*, *JHEP* **05**, 082 (2013), [[arXiv:1212.4504](#)].
- [78] P. F. Monni *et al.*, *JHEP* **05**, 143 (2020), [[arXiv:1908.06987](#)].
- [79] S. Alioli *et al.*, *JHEP* **06**, 089 (2014), [[arXiv:1311.0286](#)].
- [80] L. Buonocore *et al.* (2022), [[arXiv:2203.07240](#)].
- [81] S. Alioli *et al.*, *JHEP* **05**, 128 (2023), [[arXiv:2301.11875](#)].
- [82] G. Gustafson 193–198 (1986).
- [83] X. Artru and G. Mennessier, *Nucl. Phys. B* **70**, 93 (1974).
- [84] B. Andersson *et al.*, *Phys. Rept.* **97**, 31 (1983).
- [85] B. Andersson, *The Lund model*, volume 7, Cambridge University Press (2005), ISBN 978-0-521-01734-3, 978-0-521-42094-5, 978-0-511-88149-7.
- [86] M. Bowler, *Z. Phys. C* **11**, 169 (1981).
- [87] C. Peterson *et al.*, *Phys. Rev. D* **27**, 105 (1983).
- [88] T. Sjostrand and P. Z. Skands, *JHEP* **03**, 053 (2004), [[hep-ph/0402078](#)].
- [89] J. R. Christiansen and P. Z. Skands, *JHEP* **08**, 003 (2015), [[arXiv:1505.01681](#)].
- [90] D. Amati and G. Veneziano, *Phys. Lett. B* **83**, 87 (1979).
- [91] A. Bassetto, M. Ciafaloni and G. Marchesini, *Phys. Lett. B* **83**, 207 (1979).
- [92] R. D. Field and S. Wolfram, *Nucl. Phys. B* **213**, 65 (1983).
- [93] B. Webber, *Nucl. Phys. B* **238**, 492 (1984).

- [94] J.-C. Winter, F. Krauss and G. Soff, *Eur. Phys. J. C* **36**, 381 (2004), [[hep-ph/0311085](#)].
- [95] D. Lange, *Nucl. Instrum. Meth. A* **462**, 152 (2001).
- [96] S. Jadach *et al.*, *Comput. Phys. Commun.* **76**, 361 (1993).
- [97] D. Grellscheid and P. Richardson (2007), [[arXiv:0710.1951](#)].
- [98] T. Sjostrand and M. van Zijl, *Phys. Rev. D* **36**, 2019 (1987).
- [99] P. Skands, S. Carrazza and J. Rojo, *Eur. Phys. J. C* **74**, 8, 3024 (2014), [[arXiv:1404.5630](#)].
- [100] M. Bahr, J. M. Butterworth and M. H. Seymour, *JHEP* **01**, 065 (2009), [[arXiv:0806.2949](#)].
- [101] M. Ryskin, A. Martin and V. Khoze, *Eur. Phys. J. C* **71**, 1617 (2011), [[arXiv:1102.2844](#)].
- [102] H. Schulz and P. Skands, *Eur. Phys. J. C* **71**, 1644 (2011), [[arXiv:1103.3649](#)].
- [103] B. Cabouat and J. R. Gaunt, *JHEP* **10**, 012 (2020), [[arXiv:2008.01442](#)].
- [104] O. Fedkevych and J. R. Gaunt, *JHEP* **02**, 090 (2023), [[arXiv:2208.08197](#)].
- [105] R. Casalbuoni *et al.*, *JHEP* **01**, 035 (2009), [[arXiv:0812.1982](#)].
- [106] (2005), [[hep-ex/0511027](#)].
- [107] L. Lonnblad and T. Sjostrand, *Eur. Phys. J. C* **2**, 165 (1998), [[hep-ph/9711460](#)].
- [108] V. Khoze *et al.*, *Eur. Phys. J. C* **69**, 85 (2010), [[arXiv:1005.4839](#)].
- [109] J. Butterworth, J. R. Forshaw and M. Seymour, *Z. Phys. C* **72**, 637 (1996), [[hep-ph/9601371](#)].
- [110] M. Bahr *et al.*, in “1st International Workshop on Multiple Partonic Interactions at the LHC,” 239–248 (2009), [[arXiv:0905.4671](#)].
- [111] R. Corke and T. Sjostrand, *JHEP* **05**, 009 (2011), [[arXiv:1101.5953](#)].
- [112] G. A. Schuler and T. Sjostrand, *Phys. Rev. D* **49**, 2257 (1994).
- [113] G. Ingelman and P. Schlein, *Phys. Lett. B* **152**, 256 (1985).
- [114] S. Navin (2010), [[arXiv:1005.3894](#)].
- [115] P. Aurenche *et al.*, *Comput. Phys. Commun.* **83**, 107 (1994), [[hep-ph/9402351](#)].
- [116] F. W. Bopp, R. Engel and J. Ranft, in “LAFEX International School on High-Energy Physics (LISHEP 98) Session A: Particle Physics for High School Teachers - Session B: Advanced School in HEP - Session C: Workshop on Diffractive Physics,” 729–741 (1998), [[hep-ph/9803437](#)].
- [117] S. Roesler, R. Engel and J. Ranft, in “International Conference on Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications (MC 2000),” 1033–1038 (2000), [[hep-ph/0012252](#)].
- [118] B. E. Cox and J. R. Forshaw, *Comput. Phys. Commun.* **144**, 104 (2002), [[hep-ph/0010303](#)].
- [119] J. Bellm *et al.*, *Phys. Rev. D* **94**, 3, 034028 (2016), [[arXiv:1605.08256](#)].
- [120] S. Mrenna and P. Skands, *Phys. Rev. D* **94**, 7, 074005 (2016), [[arXiv:1605.08352](#)].
- [121] E. Bothmann, M. Schönherr and S. Schumann, *Eur. Phys. J. C* **76**, 11, 590 (2016), [[arXiv:1606.08753](#)].
- [122] C. Bierlich *et al.* (2023), [[arXiv:2308.13459](#)].
- [123] A. Buckley *et al.*, *Eur. Phys. J. C* **65**, 331 (2010), [[arXiv:0907.2973](#)].
- [124] S. Amoroso *et al.*, *JCAP* **05**, 007 (2019), [[arXiv:1812.07424](#)].
- [125] S. Catani, B. Webber and G. Marchesini, *Nucl. Phys. B* **349**, 635 (1991).
- [126] A. Karneyeu *et al.*, *Eur. Phys. J. C* **74**, 2714 (2014), [[arXiv:1306.3436](#)].

- [127] A. Buckley *et al.*, *Comput. Phys. Commun.* **184**, 2803 (2013), [arXiv:1003.0694].
- [128] J. Bellm and L. Gellersen, *Eur. Phys. J. C* **80**, 1, 54 (2020), [arXiv:1908.10811].
- [129] M. Krishnamoorthy *et al.*, *EPJ Web Conf.* **251**, 03060 (2021), [arXiv:2103.05748].