

Event Generator Physics

Peter Skands

Theoretical Physics, CERN

Geneva, Switzerland

Abstract

These lectures were given at the Ecole Joliot Curie in the autumn of 2013 and are directed at a level suitable for graduate students in High Energy Physics. They are intended to give an introduction to the theory and phenomenology of general-purpose Monte Carlo event generators like HERWIG, PYTHIA, or SHERPA, focusing on collider physics applications. The principles underlying parton-shower calculations are examined, followed by an introduction to models of hadronization and the underlying event. We end with some brief comments on generator tuning. The aim is to bring the reader to a level where informed decisions can be made concerning different approaches and their uncertainties.

Contents

1	Introduction	2
1.1	The Monte Carlo Method	3
1.2	A Brief History of Event Generators	4
2	Perturbative QCD and Parton Showers	5
2.1	Scaling and the Strong Coupling	5
2.2	Factorization and PDFs	7
2.3	The Hard Scattering: Fixed-Order QCD	8
2.4	Infrared Structure of QCD	9
2.5	Theoretical Basis of Parton Showers	12
2.6	Perturbation Theory with Markov Chains	13
2.7	Matrix-Element Matching at LO and NLO	16
3	Hadronization and Soft Hadron-Hadron Physics	20
3.1	String Model	21
3.2	Soft Hadron-Hadron Processes	26
3.3	Multiple Parton Interactions	26
3.4	Tuning	28

1 Introduction

When probed at very short wavelengths, QCD is essentially a theory of free ‘partons’ — quarks and gluons — which only scatter off one another through relatively small quantum corrections, that can be systematically calculated. At longer wavelengths, of order the size of the proton $\sim 1\text{fm} = 10^{-15}\text{m}$, however, we see strongly bound towers of hadron resonances emerge, with string-like potentials building up if we try to separate their partonic constituents. Due to our inability to perform analytic calculations in strongly coupled field theories, QCD is therefore still only partially solved. Nonetheless, all its features, across all distance scales, are believed to be encoded in a single one-line formula of alluring simplicity; the Lagrangian of QCD.

The consequence for collider physics is that some parts of QCD can be calculated in terms of the fundamental parameters of the Lagrangian, whereas others must be expressed through models or functions whose effective parameters are not a priori calculable but which can be constrained by fits to data.

However, even in the absence of a perturbative expansion, there are still several strong theorems which hold, and which can be used to give relations between seemingly different processes. (This is, e.g., the reason it makes sense to constrain parton distribution functions in ep collisions and then re-use the same ones for pp collisions.) Thus, in the sections dealing with phenomenological models we shall emphasize that the loss of a factorized perturbative expansion is not equivalent to a total loss of predictivity.

The field of QCD phenomenology is currently experiencing something of a revolution. On the perturbative side, new methods to compute scattering amplitudes with very high particle multiplicities are being developed, together with advanced techniques for combining such amplitudes with all-orders resummation frameworks (like parton showers). On the non-perturbative side, the wealth of data on soft-physics processes from the LHC is forcing us to reconsider the reliability of the standard fragmentation models, and relations between heavy-ion collisions and pp ones are providing new insights into the collective behavior of hadronic matter. The study of cosmic rays impinging on the Earth’s atmosphere challenges our ability to extrapolate fragmentation models to ultra-high energies. And finally, dark-matter annihilation processes in space may produce hadrons, whose spectra are sensitive to the modeling of fragmentation.

In the following, we shall focus on mainstream QCD phenomenology, in the context of collider physics. This includes factorization, parton showers and matching, hadronization, the so-called underlying event, and event-generator tuning. While not covering everything, hopefully these topics can also serve at least as stepping stones to more specialized issues that have been left out.

Several of the discussions below rely on material from the PDG review on Monte Carlo Event Generators [1] and on the more comprehensive review by the *MCnet* collaboration [2]. The latter also contains brief descriptions of the physics implementations of each of the main general-purpose event generators, together with a guide on how to use them, and a collection of comparisons to important experimental distributions. Additional material can also be found in longer writeups from lectures given at the 2010 ESHEP [3] and 2012 TASI schools [4].

Relative uncertainty with n points	1-Dim	d -Dim	$n_{\text{eval}}/\text{point}$
Trapezoidal Rule	$1/n^2$	$1/n^{2/d}$	2^d
Simpson's Rule	$1/n^4$	$1/n^{4/d}$	3^d
Monte Carlo	$1/\sqrt{n}$	$1/\sqrt{n}$	1

Table 1: Relative uncertainty after n evaluations, in 1 and d dimensions, for two traditional numerical integration methods and stochastic Monte Carlo. The last column shows the number of function evaluations that are required per point.

1.1 The Monte Carlo Method

A ubiquitous problem in fundamental physics is the following: given a source located some distance from a detector, predict the number of counts that should be observed within the solid angle spanned by the detector (or within a bin of its phase-space acceptance), as a function of the properties of the source, the intervening medium, and the efficiency of the detector. I.e., the task is to compute integrals of the form

$$N_{\text{Count}}(\Delta\Omega) = \int_{\Delta\Omega} d\Omega \frac{d\sigma}{d\Omega}, \quad (1)$$

with $d\sigma$ a differential cross section (or rate) for the process of interest.

In particle physics, phase space has three dimensions per final-state particle (minus four for overall four-momentum-conservation). Thus, for problems with more than a few outgoing particles, the dimensionality of phase space increases rapidly. At LEP, for instance, the total multiplicity of neutral + charged hadrons (before weak decays) was typically ~ 30 particles, so about 86 dimensions.

The standard 1D numerical-integration methods give very slow convergence rates for higher-dimensional problems. For illustration, a table of convergence rates in 1 and d dimensions is given in tab. 1, comparing the Trapezoidal (2-point) rule and Simpson's (3-point) rule to random-number-based Monte Carlo. Going from 1 to d dimensions, the convergence rate of the n -point rules all degrade (while the number of function evaluations required for each "point" simultaneously increases). The MC convergence rate, on the other hand, remains the simple stochastic $1/\sqrt{n}$, and each point still only requires one function evaluation. These are some of the main reasons that MC is the preferred numerical integration technique for high-dimensional problems. In particle physics, virtually all numerical cross-section calculations are based on Monte Carlo techniques in one form or another, examples of simple generic algorithms being RAMBO [5] (flat phase-space scan) or VEGAS [6, 7] (automated importance-sampling).

Finally, the functional form of $d\sigma/d\Omega$ can become very complicated, and substantial simplifications can be obtained by splitting it into many, nested, pieces, each with fewer dimensions and simpler structures. This can formally be achieved by writing it on a factorized form and casting the result as a Markov Chain, admitting iterative solutions. Moreover, the random phase-space vectors this generates can be re-used in many ways, for instance to compute many different observables simultaneously and/or to hand "events" to propagation and detector-simulation codes.



“This risk, that convergence is only given with a certain probability, is inherent in Monte Carlo calculations and is the reason why this technique was named after the world’s most famous gambling casino. Indeed, the name is doubly appropriate because the style of gambling in the Monte Carlo casino, not to be confused with the noisy and tasteless gambling houses of Las Vegas and Reno, is serious and sophisticated.”

F. James, “Monte Carlo theory and practice”,
Rept. Prog. Phys. 43 (1980) 1145

Figure 1: *Left:* The casino in Monaco. *Right:* extract from [8] concerning the nature of Monte Carlo techniques.

The price of using random numbers is that we must generalize our notion of convergence. In calculus, we say that a sequence $\{A\}$ converges to B if an n exists for which the difference $|A_{i>n} - B| < \epsilon$ for any $\epsilon > 0$. In random-number-based techniques, we cannot completely rule out the possibility of very pathological sequences of “dice rolls” leading to large deviations from the true goal, hence we are restricted to say that $\{A\}$ converges to B if an n exists for which *the probability* for $|A_{i>n} - B| < \epsilon$, for any $\epsilon > 0$, is greater than P , for any $P \in [0, 1]$ [8]. This risk, that convergence is only given with a certain probability, is the reason why Monte Carlo techniques were named after the famous casino in Monaco, illustrated in fig. 1.

1.2 A Brief History of Event Generators

The task of a Monte Carlo event generator is to calculate everything that happens in a high-energy collision, from the hard short-distance physics to the long wavelengths of hadronization and hadron decays. Obviously, this requires some compromises to be made. General-purpose generators like HERWIG [9], PYTHIA [10], and SHERPA [11], start from low-order (LO or NLO) descriptions of the perturbative hard physics and then attempt to include the “most significant” corrections, such as higher-order matrix-element corrections and parton showers, resonance decays and finite-width effects, underlying event, beam remnants, hadronization, and hadron decays. These corrections must be taken into account to complete our understanding of QCD and connect the short-distance physics with macroscopic experiments.

Each of the generators had slightly different origins, which carries through to the emphasis placed on various physics aspects today:

- PYTHIA. Successor to JETSET (begun in 1978). Originated in hadronization studies. Main feature: the Lund string fragmentation model.
- HERWIG. Successor to EARWIG (begun in 1984). Originated in perturbative coherence studies. Main feature: angular-ordered parton showers.
- SHERPA. Begun in 2000. Originated in studies of the matching of hard-emission

matrix elements with parton showers. Main feature: CKKW matching.

There is also a large number of more specialized generators, mainly for hard processes within and beyond the SM, a few that offer alternative shower models, and ones specializing in soft-inclusive and/or heavy-ion physics.

An important aspect of contemporary generators is the ability to combine specialized ones with general-purpose ones, via interfaces. The most common interface between partonic hard-process and parton-shower generators is the Les Houches Event File (LHEF) standard, defined in [12, 13] and “spoken” by most modern generator tools. For interfaces to experimental analysis packages (like RIVET [14]) and detector simulations (like GEANT [15]), typically the HepMC standard is used [16].

2 Perturbative QCD and Parton Showers

In this section, we discuss the perturbative aspects of QCD and their modeling in Monte Carlo Event Generators. We focus in particular on initial- and final-state radiation (ISR and FSR), as represented by parton showers.

We start off, in section 2.1, by giving some general remarks about scaling in QCD and the strong coupling, before turning to the main issues in sections 2.2 – 2.6. We then round off the discussion with a brief summary of how matrix-element and parton-shower calculations can be combined to give a more accurate description of hard, wide-angle radiation, in section 2.7.

We emphasize that the focus on perturbative QCD (pQCD) in this section implies that we take all dimensionful scales to be much larger than $\mathcal{O}(1 \text{ GeV})$.

2.1 Scaling and the Strong Coupling

To first approximation, QCD is *scale invariant*. That is, if one “zooms in” on a QCD jet, one will find a repeated self-similar pattern of jets within jets within jets, reminiscent of fractals. In the context of QCD, this property was originally called Bjorken scaling. The running of the coupling, $\alpha_s(Q^2)$, breaks this invariance, but the limit of exact Bjorken scaling — QCD at fixed coupling — still provides useful intuition about what to expect.

In the scale-invariant limit, properties of interactions are determined only by *dimensionless* kinematic quantities, such as scattering angles (pseudorapidities) and ratios of energy scales. For applications of QCD to high-energy physics, an important consequence of Bjorken scaling is thus that the rate of bremsstrahlung jets, with a given transverse momentum, scales in direct proportion to the hardness of the processes that produce them. This agrees well with our intuition about accelerated charges; the harder you “kick” them, the harder the radiation they produce.

On top of that, the running coupling will introduce a dependence on the absolute scale, implying more radiation at low scales than at high ones. The running is logarithmic with energy and is governed by the so-called *beta function*,

$$Q^2 \frac{\partial \alpha_s}{\partial Q^2} = \frac{\partial \alpha_s}{\partial \ln Q^2} = \beta(\alpha_s) = -\alpha_s^2(\beta_0 + \beta_1 \alpha_s + \beta_2 \alpha_s^2 + \dots), \quad (2)$$

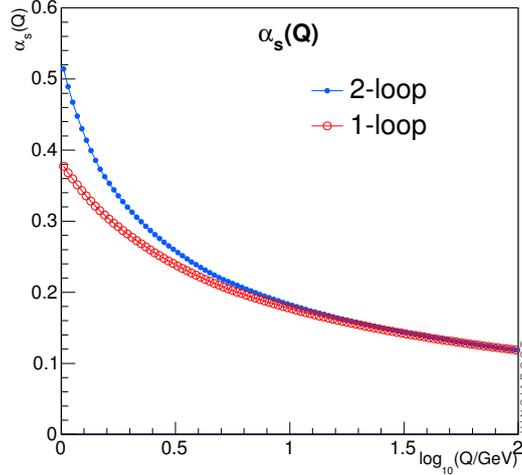


Figure 2: Illustration of the running of α_s at 1- (open circles) and 2-loop order (filled circles), starting from the same value of $\alpha_s(M_Z) = 0.12$.

with LO (1-loop) and NLO (2-loop) coefficients

$$\beta_0 = \frac{11C_A - 4T_R n_f}{12\pi} = \frac{33 - 2n_f}{12\pi}, \quad (3)$$

$$\beta_1 = \frac{17C_A^2 - 10T_R C_A n_f - 6T_R C_F n_f}{24\pi^2} = \frac{153 - 19n_f}{24\pi^2}, \quad (4)$$

with $C_A = N_C = 3$, $C_F = (N_C^2 - 1)/(2N_C) = 4/3$, and $T_R = 1/2$ for QCD in the standard normalization conventions (see, e.g., [17]). In the β_0 coefficient, the first term is due to gluon loops while the second is due to quark ones (proportional to the number of quark flavours running in the loops, n_f). Similarly, the first term of the β_1 coefficient arises from double gluon loops, while the second and third represent mixed quark-gluon ones. See, e.g., the PDG review on QCD [1] for more.

Numerically, the value of the strong coupling is usually specified by giving its value at the specific reference scale $Q^2 = M_Z^2$, from which we can obtain its value at any other scale by solving eq. (2),

$$\alpha_s(Q^2) = \alpha_s(M_Z^2) \frac{1}{1 + \beta_0 \alpha_s(M_Z^2) \ln \frac{Q^2}{M_Z^2} + \mathcal{O}(\alpha_s^2)}, \quad (5)$$

with higher-order relations available, e.g., in [18]. A comparison of one- and two-loop running, in both cases starting from $\alpha_s(M_Z) = 0.12$, is given in fig. 2. As is evident from the figure, the 2-loop running is somewhat faster than the 1-loop one.

Now consider what happens when we run the coupling towards smaller energies. Taken at face value, the numerical value of the coupling diverges rapidly at scales below ~ 1 GeV. To make this divergence explicit, one can rewrite eq. (5) in the following form (ignoring terms beyond β_0 for simplicity),

$$\alpha_s(Q^2) = \frac{1}{\beta_0 \ln \frac{Q^2}{\Lambda^2}}, \quad (6)$$

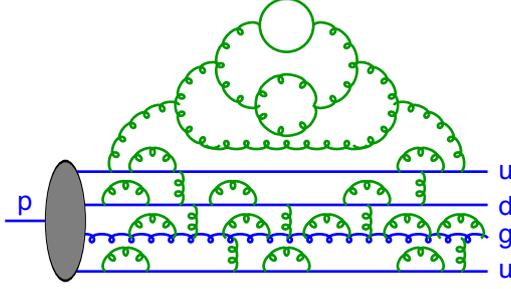


Figure 3: Illustration of partonic fluctuations inside a proton beam (from [19]).

where $\Lambda \sim 200$ MeV specifies the energy scale at which the perturbative coupling would nominally become infinite, called the Landau pole. Note, however, that this only parametrizes the purely *perturbative* result, which is not reliable at strong coupling, so eq. (6) should not be taken to imply that the physical behavior of full QCD should exhibit a divergence for $Q \rightarrow \Lambda$.

2.2 Factorization and PDFs

In high-energy scattering problems involving hadrons in the initial state, we immediately face the complication that hadrons are composite, with a time-dependent structure illustrated in fig. 3; there are partons within clouds of further partons, constantly being emitted and absorbed. Thus, before we can use perturbatively calculated partonic scattering matrix elements, we must first address the partonic structure of the colliding hadron(s).

The *factorization theorem* [20] expresses the independence of long-wavelength (soft) structure on the nature of the short-distance (hard) process. Originally formulated for DIS, factorization allows us to write the cross section for lepton-hadron scattering as a convolution of a non-perturbative but universal (i.e., process-independent) parton density function (PDF) and a perturbatively calculable partonic scattering cross section. Denoting the fraction of the hadron momentum carried by parton i by x_i ,

$$\vec{p}_i = x_i \vec{p}_h, \quad (7)$$

we may write the lepton-hadron cross section as (see, e.g., [17, 21]),

$$\sigma_{\ell h} = \sum_i \int_0^1 dx_i \int d\Phi_f f_{i/h}(x_i, \mu_F^2) \frac{d\hat{\sigma}_{\ell i \rightarrow f}(x_i, \Phi_f, \mu_F^2)}{dx_i d\Phi_f}, \quad (8)$$

with i an index running over all parton types in the incoming hadron and f enumerating all possible (partonic) final states, with Lorentz-invariant phase space, Φ_f .

The *parton density functions* (PDFs), $f_{i/h}$, parametrize the distribution of partons inside the target hadron, h . They are not a priori calculable and must be constrained by fits to data. For a discussion of PDFs in the context of MC generators, see, e.g., my TASI lectures [4]. A complementary useful discussion of the pros and cons of different choices of the factorization scale, μ_F , can be found in the TASI lectures by Tilman Plehn [22]. A more formal introduction to factorization and PDFs is given in the TASI lectures by George Sterman [23].

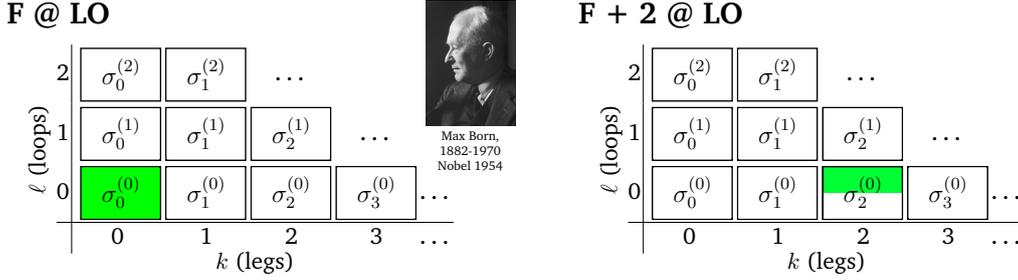


Figure 4: Coefficients of the perturbative series covered by LO calculations. *Left:* F production at lowest order. *Right:* $F + 2$ jets at LO, with the half-shaded box illustrating the restriction to the region of phase space with 2 resolved jets. (Photo from nobelprize.org).

The *partonic cross section*, $d\hat{\sigma}$, knows nothing of the target hadron apart from the fact that it contained the struck parton. It is calculable within perturbation theory, as we shall presently discuss.

2.3 The Hard Scattering: Fixed-Order QCD

Our main tool for solving QCD at high energy scales, $Q \gg \Lambda_{\text{QCD}}$, is perturbative quantum field theory, the starting point for which is Matrix Elements (MEs) which can be calculated systematically at fixed orders in the strong coupling α_s . At least at lowest order (LO), the procedure is standard textbook material [24] and it has also by now been highly automated, by the advent of tools like MADGRAPH [25, 26] and others [9, 27–34]. Here, we require only that the reader has a basic familiarity with the methods involved (based, e.g., on [17, 24]).

Importantly, only *infrared safe* observables are perturbatively calculable. IR safety implies two things (see, e.g., [35, Chapter 5]): the observable must not change value if we add any number of infinitely soft partons (invariance under *soft* radiation), and it must also not change value if we split any parton into an arbitrary number of collinearly moving ones, preserving the total energy and momentum (invariance under *collinear* splittings). These requirements are intimately related to the conditions for the sum over degenerate quantum states to produce finite total cross sections, as encapsulated by unitarity (probability conservation) and expressed by the Bloch-Nordsieck [36] and Kinoshita-Lee-Nauenberg [37, 38] theorems.

A convenient way of illustrating the terms of the perturbative series that a given calculation includes is shown in fig. 4, for an arbitrary final state: F . In the left-hand pane, the shaded box corresponds to the lowest-order “Born-level” matrix element squared. This coefficient is non-singular and hence can be integrated over all of phase space, which we illustrate by letting the shaded area fill all of the relevant box. A different kind of leading-order calculation is illustrated in the right-hand pane of fig. 4, where the shaded box corresponds to the LO matrix element squared for $F + 2$ jets. This coefficient diverges in the part of phase space where one or both of the jets are soft or collinear, and hence integrations can only cover the hard part of phase space, which we reflect by only shading the upper half of the relevant box.

Since we talk about *collinear* and *soft* divergences, cuts on *angles* and *energies* and/or cuts on combinations, like *transverse momenta*, can be used to cut away the

problematic regions of phase space. Recall, however, that pQCD is approximately scale invariant. This strongly implies that any regularization cut on a dimensionful quantity, like an energy or a transverse momentum, should in fact be formulated as a *ratio* of scales, rather than as an absolute number. For example, a jet with $p_{\perp} = 50 \text{ GeV}$ would be considered hard and well-separated if produced in association with an ordinary Z boson (with hard scale $Q_F = M_Z = 91.2 \text{ GeV}$), while the same jet would be considered soft if produced in association with a 900-GeV Z' boson.

The point is that, if $p_{\perp} \ll Q_F$, logarithmic cross-section enhancements of the type

$$\alpha_s^n \ln^{m \leq 2n} \left(\frac{Q_F^2}{p_{\perp}^2} \right) \quad (9)$$

will generate progressively larger corrections, order by order, which invalidate any attempt at a fixed-order truncation of the perturbative series. A good rule of thumb is that you must place your cuts on the matrix element such that $\sigma_{F+1}(Q_{\text{cut}}) \lesssim \alpha_s \sigma_F$, for fixed-order perturbation theory to be reliable (i.e., each successive term in your Taylor series should be suppressed by $\mathcal{O}(\alpha_s)$).

In the discussion of parton showers in section 2.5, we shall see how the region of applicability of perturbation theory can be extended to include situations with large scale hierarchies. First, however, we must understand better the origin and physics of the infrared singularities which are responsible for the logarithms in eq. (9).

2.4 Infrared Structure of QCD

The infrared structure of perturbative QCD amplitudes is governed by extremely simple physics: propagators, whose $(p^2 - m^2)$ denominators vanish when they go on shell. These terms are independent of the particular process they are embedded in, and hence the IR singularities they generate are universal. They can therefore be classified using a set of universal, process-independent functions.

Historically, the first such set of functions to be written down were the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) splitting kernels [39–41]. They can be derived in the limit that two (or more) partons become *collinear* to each other, so that their invariant mass vanishes. If the two partons could physically have originated from a single one, then there is at least one Feynman diagram which has the corresponding on-shell propagator, and, since this diagram will be singular, interference effects with any other possible (non-singular) quantum histories are irrelevant — only the singular diagram squared need be retained. Moreover, the presence of the on-shell propagator implies that the amplitude squared factorizes as follows,

$$|\mathcal{M}_{F+1}(\dots, p_i, p_j \dots)|^2 \xrightarrow{i||j} g_s^2 C_{ij} \frac{P_{I \rightarrow ij}(z)}{s_{ij}} |\mathcal{M}_F(\dots, p_i + p_j, \dots)|^2, \quad (10)$$

where $g_s^2 = 4\pi\alpha_s$ is the strong coupling, $P_{I \rightarrow ij}(z)$ is the splitting kernel¹ for $I \rightarrow ij$, and $z \in [0, 1]$ denotes the fraction of the momentum of the parent parton carried by parton i , with parton j taking a fraction $1 - z$ (so that $p_I = p_i + p_j$). The colour factor, C_{ij} , is N_C for $g \rightarrow gg$, C_F for $q \rightarrow qq$, and T_R for $g \rightarrow q\bar{q}$.

¹See [2, 17, 42] for their functional forms. They are not needed to follow this writeup.

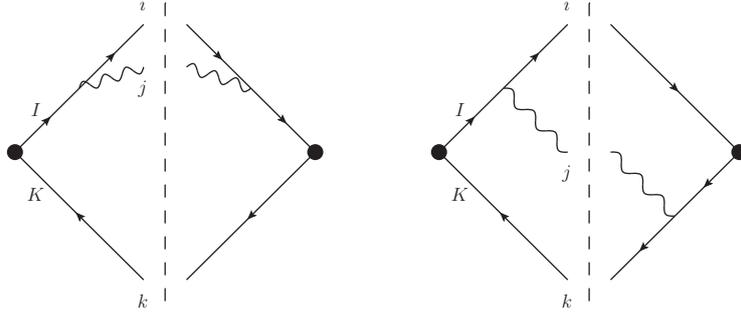


Figure 5: Squared diagrams illustrating collinear (*left*) and soft (*right*) singularities.

The basic origin of the collinear singularities is illustrated by the diagram squared in the left-hand pane of fig. 5. In the limit that partons i and j are collinear to each other, $s_{ij} \rightarrow 0$, the propagator of the parent parton, I , goes on shell; the singularity of the associated propagator factor is the origin of the $1/s_{ij}$ singularity in eq. (10).

The right-hand pane of fig. 5 illustrates the other possible IR singularity, which is not directly captured by the DGLAP functions: interference between a diagram with emission from parton I and one with emission from parton K . The resulting term has propagator singularities when both partons I and K go on shell, which can happen simultaneously if parton j is soft (i.e., has vanishing energy). This generates a universal singularity called the soft eikonal factor or the dipole factor. Thus, for each pair of colour-connected partons I and K in F , the squared amplitude for $F + 1$ gluon, $|\mathcal{M}_{F+1}|^2$, will include a factor

$$|\mathcal{M}_{F+1}(\dots, p_i, p_j, p_k, \dots)|^2 \xrightarrow{j_g \rightarrow 0} g_s^2 N_C \frac{2s_{ik}}{s_{ij}s_{jk}} |\mathcal{M}_F(\dots, p_I, p_K, \dots)|^2, \quad (11)$$

where i and k represent partons I and K after the branching (in the soft limit, $i = I$ and $k = K$). Mass effects are neglected in the above formula but can be taken systematically into account [43, 44]. The colour factor, N_C , is valid for the leading-colour contribution, regardless of whether the i and k partons are quarks or gluons; the difference in colour factors, $C_A \sim 2C_F$ is due mainly to gluons being colour-connected to two partons, while quarks are connected to only one. Subleading-colour effects go beyond the discussion here.

To give an explicit example of how the soft and collinear singularities arise in the context of a concrete physical process, consider the matrix elements squared for the process $Z^0 \rightarrow qg\bar{q}$, which is representative of a gluon being emitted from a colour-connected $q\bar{q}$ pair. For massless partons, this matrix element can be written [45]

$$\frac{|\mathcal{M}(Z^0 \rightarrow q_i g_j \bar{q}_k)|^2}{|\mathcal{M}(Z^0 \rightarrow q_I \bar{q}_K)|^2} = g_s^2 2C_F \left[\underbrace{\frac{2s_{ik}}{s_{ij}s_{jk}}}_{\text{eikonal}} + \frac{1}{s_{IK}} \underbrace{\left(\frac{s_{jk}}{s_{ij}} + \frac{s_{ij}}{s_{jk}} \right)}_{\text{collinear}} \right]. \quad (12)$$

We recognize the universal eikonal soft factor from eq. (11) in the first term. The two additional terms are less singular, and are required to obtain the correct collinear (DGLAP) limits as $s_{ij} \rightarrow 0$ or $s_{jk} \rightarrow 0$.

Note that we could equally well have considered the process $H^0 \rightarrow qg\bar{q}$. This merely corresponds to choosing the initial $q\bar{q}$ pair to be in an overall scalar state, rather than in a vector one, which should not affect the (universal) singularities². The result would then have come out as [44],

$$\frac{|\mathcal{M}(H^0 \rightarrow q_i g_j \bar{q}_k)|^2}{|\mathcal{M}(H^0 \rightarrow q_I \bar{q}_K)|^2} = g_s^2 2C_F \left[\underbrace{\frac{2s_{ik}}{s_{ij}s_{jk}}}_{\text{eikonal}} + \frac{1}{s_{IK}} \underbrace{\left(\frac{s_{jk}}{s_{ij}} + \frac{s_{ij}}{s_{jk}} \right)}_{\text{collinear}} + \underbrace{\frac{2}{s_{IK}}}_{\text{finite}} \right]. \quad (13)$$

As expected, the singular terms are all the same as in eq. (12), but there is an additional term $+2/s_{IK}$. This term is non-singular (“finite”) over all of phase space. Thus, we here see an explicit example that the singularities are process-independent while non-singular terms are process-dependent.

The logarithms discussed in the previous section, and expressed by eq. (9) in particular, arise if we try to integrate matrix elements like eqs. (12) or (13) over the entire phase space. If the singularities at $s_{ij} \rightarrow 0$, $s_{jk} \rightarrow 0$ are included in the integration, we would end up with a double pole. If we instead regulate the divergence by cutting off the integration at some minimal *perturbative cutoff scale* μ_{IR}^2 , we end up with a logarithm squared of that scale. This is a typical example of “large logarithms” being generated by the presence of scale hierarchies.

We now understand the fundamental origin of the IR singularities, why they are universal, and why amplitudes factorize in the soft and collinear limits — the singularities are simply generated by intermediate parton propagators going on shell, which is independent of the nature of the hard process, and hence can be factorized from it. Equivalently, when a parton goes on shell, it can propagate over very long distances. Its subsequent breakup can therefore not depend on any of the short-distance physics of the primary process (apart from what may be encoded in the parton itself via entanglement, such as helicity and/or colour correlations).

Before we continue, it is worth noting that eq. (11) is often rewritten in other forms to emphasize specific aspects of it. One such rewriting is thus to reformulate the invariants s_{ij} in terms of energies and angles,

$$s_{ij} = 2E_i E_j (1 - \cos \theta_{ij}) . \quad (14)$$

$$\implies \frac{ds_{ij}}{s_{ij}} \frac{ds_{jk}}{s_{jk}} \propto \frac{dE_j}{E_j} \frac{d\theta_{ij}}{\theta_{ij}} + \frac{dE_j}{E_j} \frac{d\theta_{jk}}{\theta_{jk}} . \quad (15)$$

This rewriting enables an intuitively appealing categorization of the singularities as related to vanishing energies and angles, explaining why they are called *soft* and *collinear*, respectively. Nonetheless, Lorentz non-invariant formulations come with similar caveats as do gauge non-invariant formulations of quantum field theory: while they can be practical to work with at intermediate stages of a calculation, one should be careful with any physical conclusions that rely explicitly on them.

In this writeup, we shall therefore restrict ourselves to a Lorentz-invariant formalism based directly on the invariants, as pioneered by the dipole formulation of QCD cascades by Gustafson [45] and formalized for QCD amplitudes by Kosower,

²At least not at the spin-averaged level. See [42] for a treatment of helicity-dependent structures.

in *antenna factorization* [46–48]. (Note: to avoid confusion with other dipole-like objects, we shall usually only use the word *antennae* in this writeup.) The functions in square brackets in eqs. (12) & (13) are then called *antenna functions*, and describe the IR structure of coherent $2 \rightarrow 3$ parton branching processes (here $q_I \bar{q}_K \rightarrow q_i g_j \bar{q}_k$). A convenient set of $2 \rightarrow 3$ (and $2 \rightarrow 4$) antenna functions that parametrize all the (soft and collinear) singularities of QCD amplitudes through NNLO is given by “GGG” in [48], and there is a corresponding simple factorization of phase space, based on $2 \rightarrow 3$ (and $2 \rightarrow 4$) on-shell kinematics (see section 2.6).

Another widely used set of functions and phase-space factorizations are the so-called *Catani-Seymour* (CS) dipoles [49,50]. The distinction between CS dipoles and antennae is basically that one antenna is made up of two CS dipole “ends”. At NLO, however, there is no fundamental incompatibility — the antennae we use here can always be partitioned into two CS dipole ends, if so desired.

2.5 Theoretical Basis of Parton Showers

Assume that we have computed the Born-level (LO) cross section for some process, F , and that this process contains a number of coloured partons. For each pair of (massless) colour-anticolour charges in F , we then know from the preceding section that the LO cross section for $F + 1$ partons, $d\sigma_{F+1}^{(0)}$, will include a factor

$$|\mathcal{M}_{F+1}^{(0)}|^2 = g_s^2 N_C \left[\frac{2s_{ik}}{s_{ij}s_{jk}} + \text{less singular terms} \right] |\mathcal{M}_F^{(0)}|^2. \quad (16)$$

As hinted at above, the real power of this result lies in the fact that it is *universal*. Thus, for *any* process F , we can apply eq. (16) in order to get an approximation for $d\sigma_{F+1}$. We may then, for instance, take our newly obtained expression for $F + 1$ as our arbitrary process and crank eq. (16) again, to obtain an approximation for $d\sigma_{F+2}$, and so forth. This provides a very simple recursive strategy for generating approximations to tree-level (LO) cross sections with arbitrary numbers of additional legs. (The quality of the approximations will be discussed below.)

In order to obtain a truly *all-orders* calculation, however, tree-level expressions are not enough, especially since our aim in this section is to address emissions of relatively soft and/or collinear bremsstrahlung partons, for which the tree-level expressions are badly divergent, as discussed above. The key to obtaining meaningful and finite cross sections that continue to be valid deep into the infrared regions (though still only for scales $\gg \Lambda_{\text{QCD}}$) is *unitarity*.

Unitarity (see section 2.3) implies that the singularities caused by integration over the tree-level matrix elements must be canceled, order by order, by equal but opposite-sign singularities in the virtual corrections at the same order. That is, from eq. (16), we know that the 1-loop correction to $d\sigma_F$ *must* contain a term,

$$2\text{Re}[\mathcal{M}_F^{(0)} \mathcal{M}_F^{(1)*}] \supset -g_s^2 N_C |\mathcal{M}_F^{(0)}|^2 \int \frac{ds_{ij} ds_{jk}}{16\pi^2 s_{ijk}} \left(\frac{2s_{ik}}{s_{ij}s_{jk}} + \text{less singular terms} \right), \quad (17)$$

that cancels the divergence coming from eq. (16) itself. Further, since this is universally true³, we may apply eq. (17) again to get an approximation to the corrections

³Up to process-dependent non-singular terms similar to those appearing in eqs. (12) & (13).

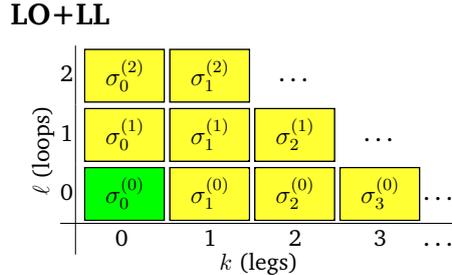


Figure 6: Coefficients of the perturbative series covered by LO+LL approximations. Green (darker) shading represents the full perturbative coefficient at the respective k and ℓ . Yellow (lighter) shading represents an LL approximation to it.

generated by eq. (16) at the next order and so on. By adding such terms explicitly, order by order, we may bootstrap our way around the entire perturbative series, using eq. (16) to move horizontally and eq. (17) to move back up along diagonals of constant $n = k + \ell$. With real-virtual cancellations now restored, the entire perturbative phase space can be covered also for $k \geq 1$, resulting in the picture shown in fig. 6, with the lighter (yellow) shading used to emphasize that only the leading singular parts (denoted LL, for leading-logarithmic) are included for $k + \ell \geq 1$.

Physically speaking, unitarity thus implies that there is a *detailed balance* between the real and virtual corrections: the negative contribution to the n -parton bin must be exactly minus the integral of the positive contribution to the $(n + 1)$ -parton one (modulo non-singular terms), and so on.

In the framework of parton-shower algorithms, this relationship is represented in terms of an *evolution* of the event structure, with precisely balanced gain/loss terms driven by the bremsstrahlung singularities, as represented e.g. by DGLAP (HERWIG [51, 52] and PYTHIA [53, 54]) or dipole/antenna functions (ARIADNE [45, 55], SHERPA [56, 57], and VINCIA [58, 59]). We consider how to recast perturbation theory in this way, as a Markov-Chain evolution process, in the next subsection.

The quality of the approximation depends on how many terms besides the leading one shown in eq. (16) are included in the game. Parton showers normally include at least the soft and collinear terms of eqs. (10) & (11), as well as several formally sub-leading improvements, such as clever renormalization-scale choices [59, 60], explicit momentum conservation, spin-correlation effects [61–63], higher-order coherence effects [64], finite-width effects [65], etc. Thus, the final result is typically much closer to reality than a corresponding “strictly LL” analytical calculation would be.

2.6 Perturbation Theory with Markov Chains

Consider again the Born-level cross section for an arbitrary hard process, F , differentially in an arbitrary (IR-safe) observable \mathcal{O} :

$$\left. \frac{d\sigma_F^{(0)}}{d\mathcal{O}} \right|_{\text{Born}} = \int d\Phi_F |\mathcal{M}_F^{(0)}|^2 \delta(\mathcal{O} - \mathcal{O}(\Phi_F)), \quad (18)$$

where the integration runs over all Born-level momentum configurations, Φ_F (for hadron collisions, include integrals over parton luminosities as well), and the δ func-

tion projects out the cross section for each value of \mathcal{O} . To make the connection to parton showers, we insert an operator, \mathcal{S} , that acts on the Born-level final state *before* the observable is evaluated, i.e.,

$$\left. \frac{d\sigma_F}{d\mathcal{O}} \right|_{\text{Shower}} = \int d\Phi_F |\mathcal{M}_F^{(0)}|^2 \mathcal{S}(\Phi_F, \mathcal{O}) . \quad (19)$$

Formally, this operator — the evolution operator — will be responsible for generating all (real and virtual) higher-order corrections to the Born-level expression. The measurement δ function appearing explicitly in eq. (18) is now implicit in \mathcal{S} .

Physically, the job of \mathcal{S} is to generate stochastic bremsstrahlung emissions according to the radiation functions discussed previously (DGLAP or dipole/antenna functions), subject to the modifications imposed by unitarity. In a shower context, we thus interpret the radiation functions as the basic *probability density* for a radiator (parton or dipole/antenna) to undergo a branching, per unit phase-space volume,

$$\frac{dP(\Phi)}{d\Phi} = g_s^2 \mathcal{C}_r A_r(\Phi) , \quad (20)$$

where \mathcal{C}_r and A_r stand for generic colour factors and radiation functions. (If there are several partons / dipoles / antennae, the total probability for branching of the event as a whole is obtained as a sum of such terms, see below.)

Unitarity implies that the phase-space density, eq. (20), must be modified to take detailed balance into account. This is natural in a picture in which each radiator undergoes an evolution, during which it may branch or not, as a function of some *evolution scale*, Q_E . As Q_E is lowered, the number of unbranched radiators decreases, while the number of branched ones increases, reflecting unitarity. The resummed branching probability, which takes into account the probability that a radiator remains unbranched between the scales Q_{E1} and Q_{E2} , is then

$$\frac{dP_{\text{res}}(\Phi)}{d\Phi} = g_s^2 \mathcal{C}_r A_r(\Phi) \Delta(\Phi, Q_{E1}^2, Q_{E2}^2) , \quad (21)$$

with the probability that there is *no evolution* (i.e., no emissions) between the scales Q_{E1} and Q_{E2} given by the so-called *Sudakov factor*, Δ , defined by

$$\Delta(\Phi_F, Q_{E1}^2, Q_{E2}^2) = \exp \left[- \sum_r \int_{Q_{E2}^2}^{Q_{E1}^2} \frac{d\Phi_{F+1}^r}{d\Phi_F} g_s^2 \mathcal{C}_r A_r(\Phi_{F+1}) \right] . \quad (22)$$

This exponential is completely analogous to the one arising, e.g., in the problem of nuclear decay, with the difference that our “decay constant”, $dP/d\Phi$, here depends explicitly on time (as measured by the evolution variable). The shower is thus naturally ordered such that emissions with large values of Q_E (corresponding to short time- and distance-scales) are considered first. Note that, in regions of phase space where A_r is small, Δ will be close to unity; the corrections from unitarity are small, and there is little or no evolution. Conversely, in the singular regions, where $Q_E \rightarrow 0$ and $A_r \rightarrow \infty$, we have $\Delta \rightarrow 0$, implying that *all* events evolve.

The factor $d\Phi_{F+1}^r/d\Phi_F$ in eq. (22) defines the chosen phase-space factorization. Our favorite is the so-called antenna factorization, whose principal virtue is that it

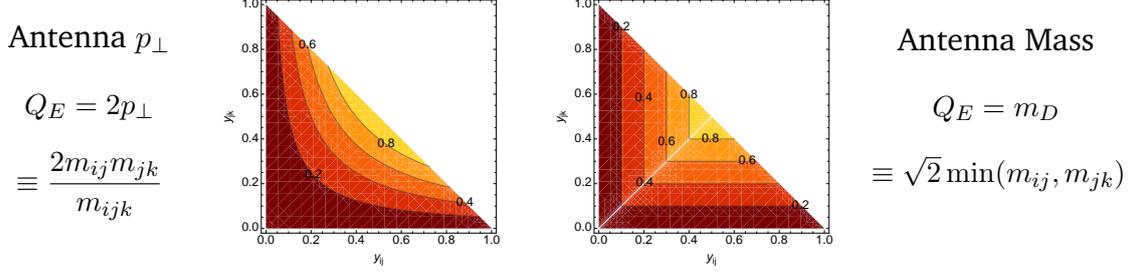


Figure 7: Example showing contours of two different evolution variables, over the triangular antenna phase space (the axes are the branching invariants, s_{ij} and s_{jk} , normalized by s_{ijk}).

is the simplest Lorentz invariant factorization which covers all of phase space while only involving on-shell momenta. For completeness, its form is

$$\frac{d\Phi_{F+1}^r}{d\Phi_F} = \frac{d\Phi_3^r}{d\Phi_2} = ds_{ij} ds_{jk} \frac{d\phi}{2\pi} \frac{1}{16\pi^2 s_r}, \quad (23)$$

i.e., it is a $2 \rightarrow 3$ factorization which involves just one parent colour-anticolour pair for each r , with invariant mass squared $s_r = (p_i + p_j + p_k)^2$. (Other choices, e.g. more global ones involving all partons in the event, or less global ones with a single parton playing the dominant role as emitter, are also possible.)

Depending on the shower algorithm, Q_E may be defined as a parton virtuality (virtuality-ordered showers [53]), as a transverse-momentum scale (p_\perp -ordered showers [45, 54]), or as a combination of energies times angles (angular ordering [64]). During the shower evolution, each model effectively “sweeps” over phase space in the order implied by these contours, illustrated in fig. 7. E.g., a p_\perp -ordered antenna shower (left) will treat a hard-collinear branching as occurring “earlier” than a soft one, while a mass-ordered antenna shower (right) will tend to do the opposite. This affects the tower of virtual corrections generated by the Sudakov factors for each shower model.

Skipping intermediate steps, the form of the all-orders pure-shower Markov chain, for the evolution of an event between two evolution scales $Q_{E1} > Q_{E2}$, is,

$$\begin{aligned} \mathcal{S}(\Phi_F, Q_{E1}^2, Q_{E2}^2, \mathcal{O}) &= \underbrace{\Delta(\Phi_F, Q_{E1}^2, Q_{E2}^2) \delta(\mathcal{O} - \mathcal{O}(\Phi_F))}_{F + 0 \text{ exclusive above } Q_{E2}} \\ &+ \underbrace{\sum_r \int_{Q_{E2}^2}^{Q_{E1}^2} \frac{d\Phi_{F+1}^r}{d\Phi_F} g_s^2 C_r A_r(\Phi_{F+1}) \Delta(\Phi_F, Q_{E1}^2, Q_{F+1}^2) \mathcal{S}(\Phi_{F+1}, Q_{F+1}^2, Q_{E2}^2, \mathcal{O})}_{F + 1 \text{ inclusive above } Q_{E2}}. \end{aligned} \quad (24)$$

The term on the first line of eq. (24) thus represents all events that *did not* evolve as the resolution scale was lowered from Q_{E1} to Q_{E2} , while the second line contains a sum and phase-space integral over those events that *did* evolve — including the insertion of $\mathcal{S}(\Phi_{F+1})$ representing the possible further evolution of the event and completing the iterative definition of the Markov chain.

Algorithmically, branchings are generated with this distribution, starting from a uniformly distributed random number $\mathcal{R} \in [0, 1]$, by solving the equation,

$$\mathcal{R} = \Delta(Q_{E1}^2, Q_E^2) \quad (25)$$

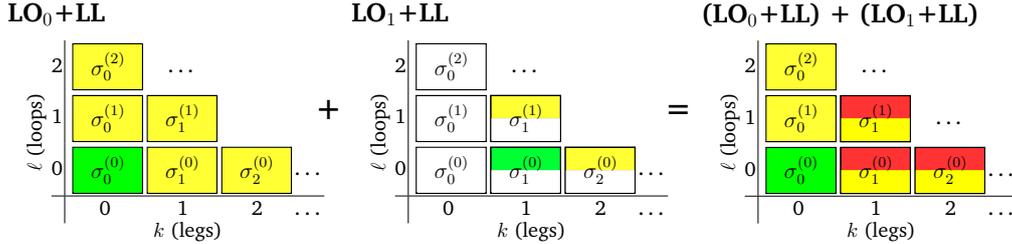


Figure 8: The double-counting problem caused by naively adding ME calculations with different numbers of legs. The notation LO_n denotes the LO matrix element for $F + n$ partons.

for Q_E . Two complementary phase-space variables are then generated (usually labeled z and ϕ), according to the shape of the radiation function over a phase-space contour of constant Q_E . From these, the model-independent set of trial phase-space variables (s_{ij}, s_{jk}, ϕ) are determined by inversion of the defining relations $Q_E(s_{ij}, s_{jk})$ and $z(s_{ij}, s_{jk})$, and the full kinematics (i.e., four-momenta) of the trial branching can then be constructed, see e.g. [58, 59, 66]. If the branching is accepted (MC overestimates used in the previous steps result in an accept probability < 1 here), the mother parton(s) are replaced by the daughter ones in the event record. Otherwise the pre-branching system is kept. In both cases, the evolution is restarted from the current value of Q_E .

2.7 Matrix-Element Matching at LO and NLO

The basic problem that leads to matrix-element/parton-shower matching can be illustrated in a very simple way. Assume again that we have performed an LO+LL calculation for some process, F , as in the left-hand pane of fig. 8, and that we also have an LO+LL calculation based on the LO matrix element for $F + 1$ (restricted to cover only the phase-space region with at least one hard resolved jet, to avoid the singularities), illustrated by the half-shaded boxes in the middle pane of fig. 8.

Desiring to combine these two samples, we might attempt simply to add them together. However, the LL terms for $F + 1$ would then be counted twice: once from the shower off F and once from the matrix element for $F + 1$, illustrated by the dark shaded (red) areas of the right-hand pane of fig. 8. This *double-counting* problem would grow worse if we attempted to add more matrix elements, with more legs. The cause is very simple. Each such calculation corresponds to an *inclusive* cross section, and hence naive addition would give

$$\sigma_{\text{tot}} = \sigma_{0;\text{incl}} + \sigma_{1;\text{incl}} = \sigma_{0;\text{excl}} + 2\sigma_{1;\text{incl}}. \quad (26)$$

Recall the definition of inclusive and exclusive cross sections: “ F inclusive” \leftrightarrow “ F plus anything”. “ F exclusive” \leftrightarrow “ F and only F ”. Thus, $\sigma_{F;\text{incl}} = \sum_{k=0}^{\infty} \sigma_{F+k;\text{excl}}$.

Instead, we must *match* the coefficients calculated by the two parts of the full calculation — showers and matrix elements — more systematically, for each order in perturbation theory, so that the nesting of inclusive and exclusive cross sections is respected without overcounting. Given a parton shower and a matrix-element generator, there are fundamentally three different ways in which we can consider matching the two [66], as will now be discussed.

ME CORRECTIONS

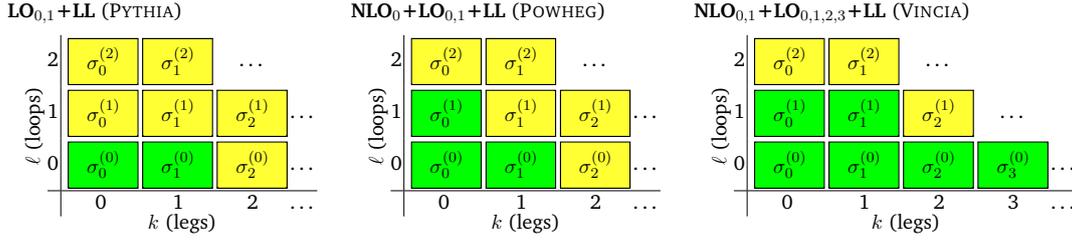


Figure 9: ME Corrections. PYTHIA (left), POWHEG (middle), and VINCIA (right). Only one event sample is produced by each of these methods, hence no sub-components are shown. The notation (N)LO_n denotes the (N)LO matrix element for $F + n$ partons.

2.7.1 ME Corrections

The oldest, and in my view most attractive, approach [53,67] consists of working out the shower approximation to a given fixed order, and correcting the shower splitting functions at that order by a multiplicative factor given by the ratio of the matrix element to the shower approximation, phase-space point by phase-space point. We may sketch this as

$$\text{Matched} = \underbrace{\text{shower}}_{\text{Approximate}} \times \underbrace{\frac{\text{Exact}}{\text{Approximate}}}_{\text{multiplicative correction}}. \quad (27)$$

When these correction factors are inserted back into the shower evolution, they guarantee that the shower evolution off $n - 1$ partons correctly reproduces the n -parton matrix elements, without the need to generate a separate n -parton sample. That is, the shower approximation is essentially used as a pre-weighted (stratified) all-orders phase-space generator, on which a more exact answer can subsequently be imprinted order by order in perturbation theory. Since the shower is already optimized for exactly the kind of singular structures that occur in QCD, very fast computational speeds can therefore be obtained with this method [68].

In the original approach [53,67], used by PYTHIA [10,69], this was only worked out for one additional emission beyond the basic resonance-decay [70] or hard-scattering process. In POWHEG [71,72], it was extended to include also virtual corrections to the Born-level matrix element. Finally, in VINCIA [42,59,66], it has been extended to include arbitrary numbers of emissions, though that method has so far only been applied to final-state showers.

An illustration of the perturbative coefficients included in each of these approaches is given in fig. 9, as usual with green (darker shaded) boxes representing exact coefficients and yellow (light shaded) boxes representing logarithmic approximations.

Finally, two more properties unique to this method deserve mention. Firstly, since the corrections modify the actual shower evolution kernels, the corrections are automatically *resummed* in the Sudakov exponential, which should improve the logarithmic precision once $k \geq 2$ is included, and secondly, since the shower is *unitary*, an initially unweighted sample of $(n - 1)$ -parton configurations remains unweighted, with no need for a separate event-unweighting or event-rejection step.

SUBTRACTION: MC@NLO

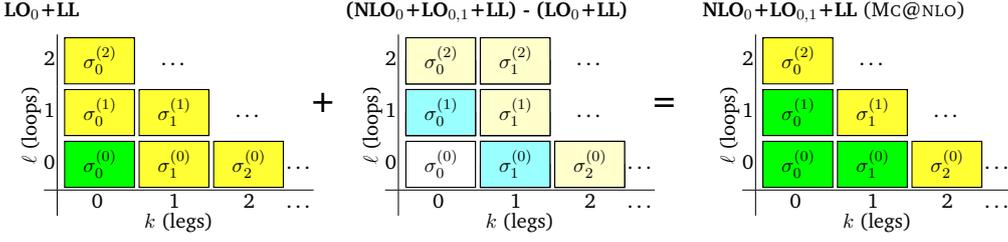


Figure 10: In the middle pane, cyan boxes denote non-singular correction terms, while the egg-colored ones denote showers off such corrections, which cannot lead to double-counting at the LL level. The notation (N)LO_{*n*} denotes the (N)LO matrix element for $F + n$ partons.

2.7.2 Subtraction

Another way of matching two calculations is by subtracting one from the other and correcting by the difference, schematically

$$\text{Matched} = \overbrace{\text{Approximate}}^{\text{shower expansion}} + \overbrace{(\text{Exact} - \text{Approximate})}^{\text{additive correction}}. \quad (28)$$

This looks very much like the structure of a subtraction-based NLO fixed-order calculation, with the shower approximation playing the role of subtraction terms. This forms the basis of the MC@NLO [73,74] strategy, illustrated in the top pane of fig. 10.

In this approach, the subtraction terms depend explicitly on the shower model, and negative-weight events will generally occur, for instance in phase-space points where the approximation is larger than the ME answer. Regarding negative weights, imagine a worst-case scenario in which 100 positive-weight events have been generated, along with 99 negative-weight ones (assuming each weight has the same absolute value). The statistical precision of the MC answer would be equivalent to one event, for 200 generated, i.e., a big loss in convergence rate. In practice, generators like MC@NLO “only” produce around 10% or less events with negative weights, so the convergence rate should not be severely affected for ordinary applications.

Nevertheless, the problem of negative weights motivated the development of the POWHEG approach (discussed above), which was constructed specifically to prevent negative-weight events from occurring and simultaneously to be more independent of which parton-shower algorithm it is used with.

2.7.3 Slicing

The most commonly encountered matching type is currently based on separating (slicing) phase space into two regions, one of which is supposed to be mainly described by hard matrix elements and the other, by the shower.

This type of approach was first used in HERWIG [75], to include matrix-element corrections for one emission beyond the basic hard process [76,77]. This is illustrated in fig. 11. The method has since been generalized by several independent groups to include LO matrix elements for arbitrary numbers of additional legs, the most well-known of these being the CKKW [78], CKKW-L [79,80], and MLM [81,

SLICING: HERWIG

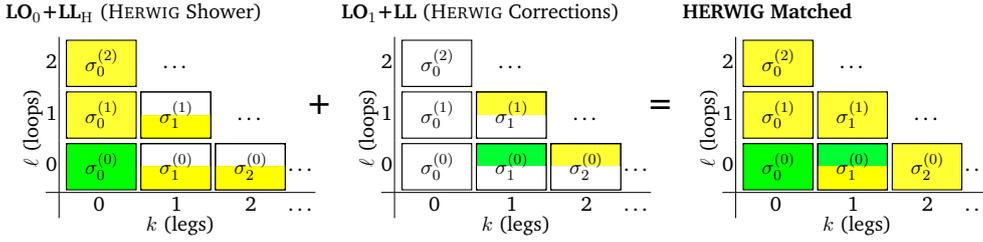


Figure 11: HERWIG’s original matching scheme [76, 77], in which the dead zone of the HERWIG shower (represented by LL_H) was used as an effective “matching scale”.

82] approaches. There is also a unitarity-improved approach to LO slicing, called UMEPS [83]. Finally, there are by now also several proposals for how to combine NLO samples with multiple parton multiplicities, such as MENLOPS [84], FFX [85], UNLOPS [86], and others [59, 87, 88]. Here, we focus on slicing for LO approaches.

The defining characteristic of slicing schemes is that the shower approximation is set to zero above some scale, either due to the presence of dead zones in the shower, as in HERWIG, or by explicitly vetoing any emissions above a certain *matching scale*. The empty part of phase space can then be filled by separate events generated according to higher-multiplicity tree-level matrix elements. The practical limit on the number of additional partons is around 3 or 4, due to computational complexity (which rises very rapidly in these approaches, see e.g., [42, 68]).

In order to match smoothly with the shower calculation, the higher-multiplicity matrix elements must be associated with Sudakov form factors (representing the virtual corrections that would have been generated if a shower had produced the same phase-space configuration), and their α_s factors must be chosen so that, at least at the matching scale, they become identical to the choices made on the shower side [89]. This can be done by constructing “fake parton-shower histories” for the higher-multiplicity matrix elements. By applying a sequential jet clustering algorithm, a tree-like branching structure can be created that has the same dominant structure as that of a parton shower. Given the fake shower tree, α_s factors can be chosen for each vertex with argument $\alpha_s(p_\perp)$ and Sudakov factors can be computed for each internal line in the tree. In the CKKW method, these Sudakov factors are estimated analytically, while the MLM and CKKW-L, and UMEPS methods compute them numerically, from the actual shower evolution.

Below the matching scale, the small difference between the matrix elements and the shower approximation can be dropped (since their leading singularities are identical and this region by definition includes no hard jets), yielding the pure shower answer in that region. Above the matching scale, the matched result is identical to the matrix element (ME), modulo higher-order (Sudakov and α_s) corrections. This type of strategy is illustrated in fig. 12.

As emphasized above, since this strategy is discontinuous across phase space, a main point here is to ensure that the behavior across the matching scale be as smooth as possible. Technical details of the implementation are important, and the dependence on the unphysical matching scale may be larger than expected unless the implementation matches the theoretical algorithm precisely [79, 80, 83, 90]. Fur-

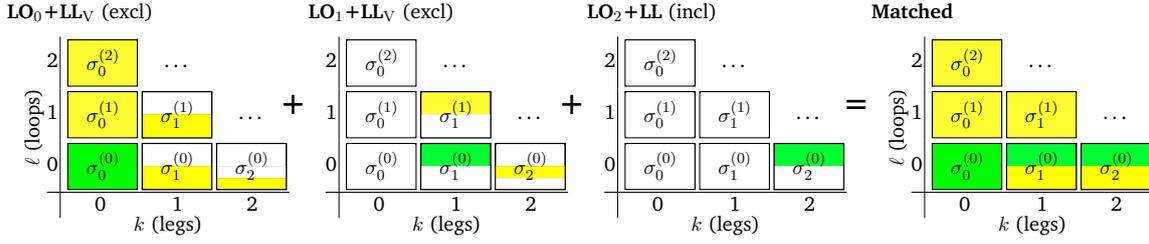


Figure 12: Slicing, including two additional tree-level matrix elements beyond the basic process, with an explicit veto on the parton showers, denoted LL_V .

thermore, since the Sudakov factors are generally computed using showers (MLM, L-CKKW, UMEPS), while the real corrections are computed using matrix elements, care must be taken not to (re-)introduce differences that could break the detailed real-virtual balance that ensures unitarity among the singular parts, see e.g., [87, 89].

It is advisable not to choose the matching scale too low. This is again essentially due to the approximate scale invariance of QCD imploring us to write the matching scale as a ratio, rather than as an absolute number. If one uses a very low matching scale, the higher-multiplicity matrix elements will already be quite singular, leading to very large LO cross sections before matching. After matching, these large cross sections are tamed by the Sudakov factors produced by the matching scheme, and hence the final cross sections may still look reasonable. But the higher-multiplicity matrix elements in general contain subleading singularity structures, beyond those accounted for by the shower, and hence unitarity has been compromised [83, 87]. We therefore recommend not to take the matching scale lower than about an order of magnitude below the characteristic scale of the hard process.

3 Hadronization and Soft Hadron-Hadron Physics

We here give a very brief overview of the main aspects of soft QCD that are relevant for hadron-hadron collisions, such as hadronization, minimum-bias and soft-inclusive physics, and the so-called underlying event. This will be kept at a pedestrian level and is largely based on the reviews in [1–3].

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

The problem can be stated as follows: given a set of partons resolved at a scale of $Q_{\text{had}} \sim 1$ GeV, we need a “mapping” from this set onto a set of on-shell colour-singlet (i.e., confined) hadronic states. MC models do this in three steps:

1. Map the partonic system onto a continuum of high-mass hadronic states (called “strings” or “clusters”).

2. Iteratively map strings/clusters onto discrete set of primary hadrons (via string breaks / cluster splittings / cluster decays).
3. Sequential decays into secondaries ($\rho \rightarrow \pi\pi$, $\Lambda \rightarrow n\pi$, $\pi^0 \rightarrow \gamma\gamma$, ...).

The physics governing this mapping is non-perturbative. However, we do have some knowledge of the properties that such a solution must have. For instance, Poincaré invariance, unitarity, and causality are all concepts that apply beyond perturbation theory. In addition, lattice QCD provides us a means of making explicit quantitative studies in a genuinely non-perturbative setting (albeit only of certain questions).

An important result in “quenched” lattice QCD⁴ is that the potential between a quark and an antiquark appears to grow linearly with the 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 section 3.1. Alternatively, a property of perturbative QCD called “preconfinement” [91] is the basis of the *cluster model of hadronization*, described in [1, 2].

In the generator landscape, PYTHIA uses string fragmentation, while HERWIG and SHERPA use cluster fragmentation. Note that the so-called *parton level* that can be obtained by switching off hadronization in an MC generator, is not a universal concept, since each model defines the hadronization scale differently, with different tunes using different values for it. Comparisons to distributions at this level (i.e., with hadronization switched off) may therefore be used to provide an idea of the overall impact of hadronization corrections within a given model, but should be avoided in the context of physical observables. Note also that the corresponding MC *fragmentation functions* are intrinsically defined at the hadronization scale. They can therefore not be compared directly to those that are used in fixed-order / analytical-resummation contexts, which are typically defined at a factorization scale of order the scale of the hard process.

Soft hadron-hadron processes, like diffraction, minimum-bias and the underlying event, are discussed in sections 3.2 & 3.3. We round off with a discussion of Monte Carlo tuning in section 3.4.

3.1 String Model

Starting from early concepts developed by Artru and Mennessier [92], several hadronization models based on strings were proposed in the late 70’ies and early 80’ies. Of these, the most widely used today is the so-called Lund model [93], implemented in the PYTHIA code. We concentrate on that particular model here, though many of the overall concepts would be shared by any string-inspired method.

Consider the production of a $q\bar{q}$ pair from vacuum, for instance in the process $e^+e^- \rightarrow \gamma^*/Z \rightarrow q\bar{q} \rightarrow \text{hadrons}$. As the quarks move apart, linear confinement implies that a potential

$$V(r) = \kappa r \tag{29}$$

is asymptotically reached for large distances, r . Such a potential describes a string with tension (energy per unit length) κ , which has been determined (from hadron

⁴Quenched QCD implies no “dynamical” quarks, i.e., no $g \rightarrow q\bar{q}$ splittings allowed.

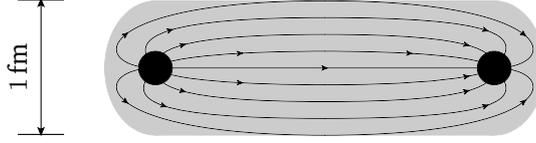


Figure 13: Illustration of the transition between a Coulomb potential at short distances to the string-like one of eq. (29) at large $q\bar{q}$ separations.

mass spectroscopy) to have the value $\kappa \sim 1 \text{ GeV/fm} \sim 0.2 \text{ GeV}^2$.

The string can be thought of as parameterizing the position of the axis of a cylindrically symmetric flux tube, illustrated in fig. 13. Such simple $q - \bar{q}$ strings form the starting point for the string model. More complicated multi-parton topologies are treated by representing gluons as transverse “kinks”, e.g., $q - g - \bar{q}$. The space-time evolution is slightly more involved [93], and modifications to the fragmentation model to handle stepping across gluon corners have to be included, but the main point is that there are no separate free parameters for gluon jets. Differences with respect to quark fragmentation arise simply because quarks are only connected to a single string piece, while gluons have one on either side, increasing the energy loss per unit (invariant) time from a gluon to the string by a factor of 2 relative to quarks, which can be compared to the ratio of colour Casimirs $C_A/C_F = 2.25$. Another appealing feature of the model is that low-energy gluons are absorbed smoothly into the string, without leading to modifications. This improves the stability of the model with respect to variations of the infrared behaviour of the parton shower.

As the partonic string endpoints move apart, their kinetic energy is gradually converted to potential energy, stored in the growing string spanned between them. In the “quenched” approximation, in which $g \rightarrow q\bar{q}$ splittings are not allowed, this process would continue until the endpoints have lost *all* their momentum, at which point they would reverse direction and be accelerated by the now shrinking string.

In the real world, quark-antiquark fluctuations inside the string field can make the transition to become real particles by absorbing energy from the string, thereby screening the original endpoint charges from each other and breaking the string into two separate colour-singlet pieces, $(q\bar{q}) \rightarrow (q\bar{q}') + (q'\bar{q})$, illustrated in fig. 14 a. This process then continues until only ordinary hadrons remain. (We will give more details on the individual string breaks below.)

Since the string breaks are causally disconnected (as can easily be realized from space-time diagrams [93]), they do not have to be considered in any specific time-ordered sequence. In the Lund model, the string breaks are generated starting with the leading (“outermost”) hadrons, containing the endpoint quarks, and iterating inwards towards the center of the string, alternating randomly between fragmentation off the left- and right-hand sides, respectively, fig. 14 b. One can thereby split off a single well-defined hadron in each step, with a mass that, for unstable hadrons, is selected according to a Breit-Wigner distribution.

The details of the individual string breaks are not known from first principles. The Lund model invokes the idea of quantum mechanical tunneling, which leads to

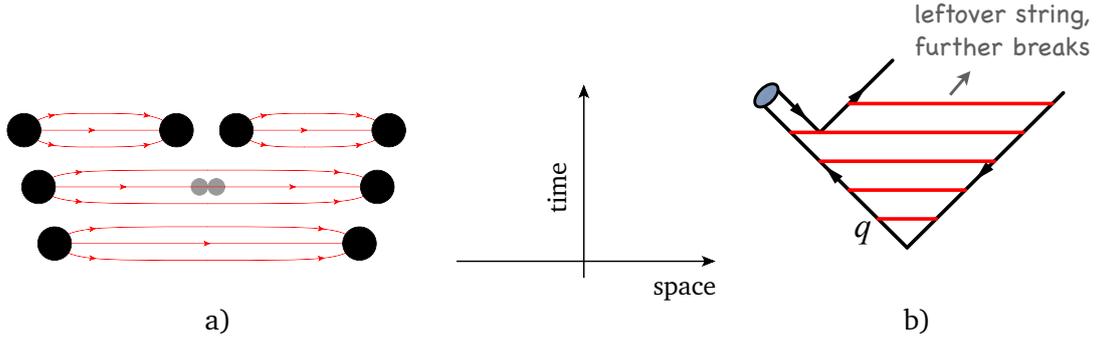


Figure 14: a) String breaking by quark pair creation in the string field. b) Fragmentation from the outside-in, splitting off a single on-shell hadron in each step.

a Gaussian suppression of the energies and masses imparted to the produced quarks,

$$\text{Prob}(m_q^2, p_{\perp q}^2) \propto \exp\left(\frac{-\pi m_q^2}{\kappa}\right) \exp\left(\frac{-\pi p_{\perp q}^2}{\kappa}\right), \quad (30)$$

where m_q is the mass of the produced quark and p_{\perp} is the transverse momentum imparted to it by the breakup process (with the \bar{q} having the opposite p_{\perp}).

Due to the factorization of the p_{\perp} and m dependence implied by eq. (30), the p_{\perp} spectrum of produced quarks in this model is independent of the quark flavour, with a universal average value of

$$\langle p_{\perp q}^2 \rangle = \sigma^2 = \kappa/\pi \sim (250 \text{ MeV})^2. \quad (31)$$

Bear in mind that “transverse” is here defined with respect to the string axis. Thus, the p_{\perp} in a frame where the string is moving is modified by a Lorentz boost factor. Also bear in mind that σ^2 is here a purely non-perturbative parameter. In a Monte Carlo model with a fixed shower cutoff Q_{had} , additional unresolved soft-gluon radiation below Q_{had} may give a larger effective amount of “non-perturbative” p_{\perp} . Note that, since hadrons receive p_{\perp} contributions from two breakups, one on either side, their average transverse momentum squared will be twice as large,

$$\langle p_{\perp h}^2 \rangle = 2\sigma^2. \quad (32)$$

The mass suppression implied by eq. (30) is less straightforward to interpret. Since light-quark masses are ambiguous, the value of the strangeness suppression must effectively be extracted from experimental measurements, e.g., of the K/π ratio, with a resulting suppression of roughly $2s/(u+d) \sim 0.2$. The expected suppression of c quarks is of order 10^{-11} . Heavy quarks can therefore safely be considered to be produced only in the perturbative stages and not by the soft fragmentation.

Baryon production can be incorporated in the same basic picture [94], by allowing string breaks to occur also by the production of pairs of so-called *diquarks*, loosely bound states of two quarks in an overall $\bar{3}$ representation (e.g., red + blue = antigreen). Again, the relative rate of diquark-to-quark production must be extracted from measurements, e.g., of the p/π ratio. More advanced scenarios for baryon production have also been proposed, e.g. the so-called *popcorn* model [95, 96], which

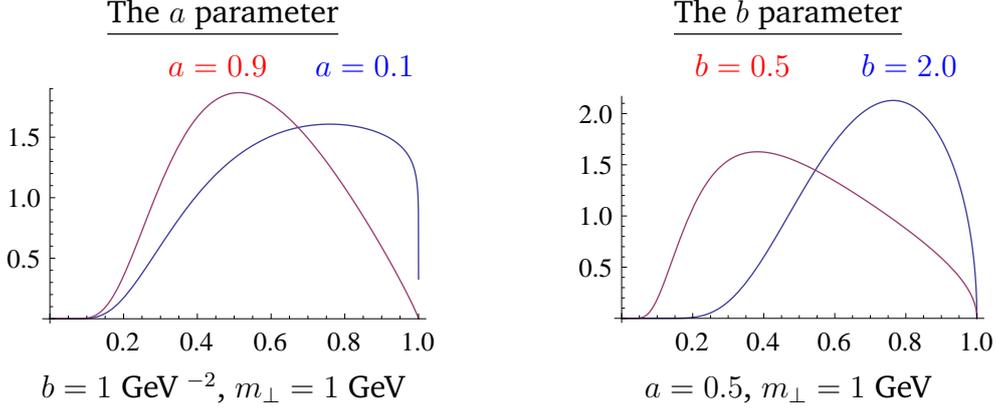


Figure 15: Illustration of the Lund symmetric fragmentation function (normalized to unity). *Left:* variation of the a parameter, from 0.1 (blue) to 0.9 (red), with fixed b . *Right:* variation of the b parameter, from 0.5 (red) to 2 (blue) GeV^{-2} , with fixed a .

complements the diquark picture and acts to decrease the correlations among neighboring baryon-antibaryon pairs by allowing mesons to be formed inbetween them, and a model based on *string junctions* [97] which can be applied to the description of beam remnants, and then acts to increase baryon stopping [98].

This brings us to the next step of the algorithm: assignment of the produced quarks within hadron multiplets. Using a nonrelativistic classification of spin states, the fragmenting q (\bar{q}) may combine with the \bar{q}' (q') from a newly created breakup to produce either a vector or a pseudoscalar meson, or, if diquarks are involved, either a spin-1/2 or spin-3/2 baryon. Unfortunately, the string model is entirely unproductive in this respect, and this is therefore the sector that contains the largest amount of free parameters. From spin counting alone, one would expect the ratio V/S of vectors to pseudoscalars to be 3, but this is modified by the V - S mass splittings, which implies a phase-space suppression of vector production, with corresponding suppression parameters to be extracted from data. The production of higher meson resonances is assumed to be low in a string framework⁵. For diquarks, separate parameters control the relative rates of spin-1 diquarks vs. spin-0 ones and, likewise, have to be extracted from data.

With p_{\perp}^2 and m^2 now fixed, the final step is to select the fraction, z , of the fragmenting endpoint quark’s longitudinal momentum that is carried by the created hadron. In this respect, the string picture is substantially more predictive than for the flavour selection. Firstly, the requirement that the fragmentation be independent of the sequence in which breakups are considered (causality) imposes a “left-right symmetry” on the fragmentation function, $f(z)$, with the solution

$$f(z) \propto \frac{1}{z}(1-z)^a \exp\left(-\frac{b(m_h^2 + p_{\perp h}^2)}{z}\right), \quad (33)$$

which is known as the *Lund symmetric fragmentation function* (normalized to unit integral). The a and b parameters, illustrated in fig. 15, are the only free parameters

⁵The four $L = 1$ multiplets are implemented in PYTHIA, but are disabled by default, largely because several states are poorly known and thus may result in a worse overall description when included.

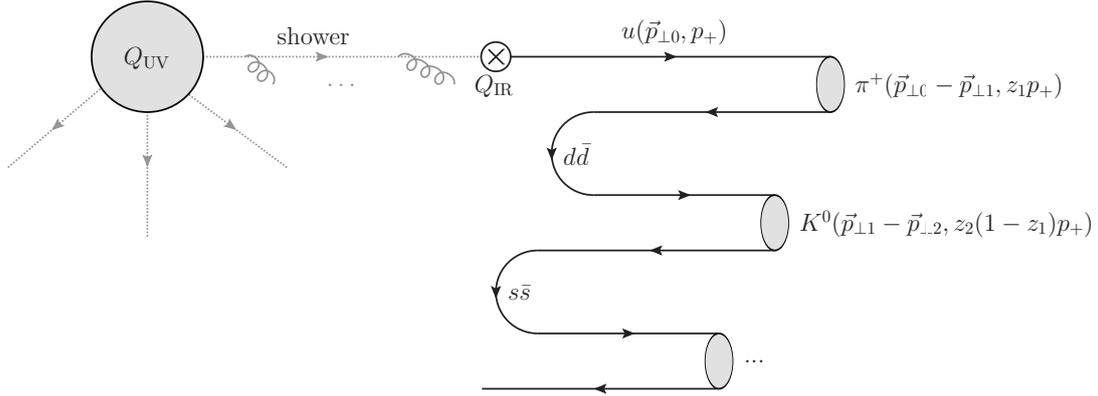


Figure 16: Illustration of the iterative selection of flavours and momenta in the Lund model.

of the fragmentation function (though a can be flavour-dependent). Note that the explicit mass dependence in $f(z)$ implies a harder fragmentation function for heavier hadrons (in the rest frame of the string).

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 [99], $f(z) \rightarrow f(z)/z^{b m_Q^2}$, with m_Q the heavy-quark mass. Strictly speaking, this is the only fragmentation function that is consistent with causality in the string model, though a few alternative forms are typically provided as well.

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 colour flux [100], and the average breakup time lying along a hyperbola of constant invariant time $\tau_0 \sim 10^{-23}\text{s}$ [93].

We may also ask, e.g., how many units of rapidity does the particle production from a string span? Measuring p_z along the string direction and defining rapidity by

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \quad (34)$$

the absolute highest rapidity that can be reached, by a pion traveling exactly along the string direction and taking all of the endpoint quark's energy, is $y_{\max} = \ln(2E_q/m_\pi)$. I.e., the rapidity region covered by a fragmenting string scales logarithmically with the energy, and since the density of hadrons produced per unit rapidity is roughly constant (modulo endpoint effects), the average number of hadrons produced by string fragmentation likewise scales logarithmically with energy.

The iterative selection of flavours, p_\perp , and z values is illustrated in fig. 16. A parton produced in a hard process at some high scale Q_{UV} emerges from the parton shower, at the hadronization scale Q_{IR} , with 3-momentum $\vec{p} = (\vec{p}_{\perp 0}, p_+)$, where the “+” on the third component denotes “light-cone” momentum, $p_\pm = E \pm p_z$. Next, an adjacent $d\bar{d}$ pair from the vacuum is created, with relative transverse momenta $\pm p_{\perp 1}$. The fragmenting quark combines with the \bar{d} from the breakup to form a π^+ , which carries off a fraction z_1 of the total lightcone momentum p_+ . The next hadron carries off a fraction z_2 of the remaining momentum, etc.

3.2 Soft Hadron-Hadron Processes

The total hadron-hadron (hh) cross section is around 100 mb at LHC energies, growing slowly with the CM energy, $\sigma_{\text{tot}}(s) \propto s^{0.08}$. There are essentially four types of physics processes, which together make up σ_{tot} :

1. Elastic scattering: $hh \rightarrow hh$;
2. Single diffractive dissociation: $hh \rightarrow h + \text{gap} + X$, with “gap” denoting an empty rapidity region, and X anything that is not the original beam particle;
3. Double diffractive dissociation: $hh \rightarrow X + \text{gap} + X$ (both hadrons “blow up”);
4. Inelastic non-diffractive scattering: everything else.

In principle, higher “multi-gap” diffractive components may be defined as well, the most important one being central diffraction: $hh \rightarrow h + \text{gap} + X + \text{gap} + h$.

Some important differences exist between theoretical and experimental terminology [101]. In the experimental setting, diffraction is defined by an observable rapidity gap, with $|\Delta y|_{\text{gap}} \gtrsim 3$ typically giving clean diffractive samples. In the MC context, however, each diffractive process type produces a whole spectrum of gaps, with small ones suppressed but not excluded. Likewise, events of non-diffractive origin may produce accidental rapidity gaps, now with large ones suppressed (exponentially) but not excluded, and in the transition region there could even be quantum mechanical interference between the two. Due to this unphysical model dependence of theoretical definitions of diffraction, we strongly advise to phrase measurements first and foremost in terms of physical observables, and only seek to connect with theory models as a second, separate, step.

Another potentially confusing term is “minimum bias” (MB). This originates from the experimental requirement of a minimal energy or number of hits in a given (experiment-dependent) instrumented region near the beam, used to determine whether there was any non-trivial activity in the event, or not. This represents the smallest possible “trigger bias” that the corresponding experiment is capable of. There is thus no universal definition of “min-bias”; each experiment has its own.

3.3 Multiple Parton Interactions

In this subsection, we focus on the physics of multiple parton interactions (MPI) as a theoretical basis for understanding both inelastic, non-diffractive processes (minimum-bias), as well as the so-called underlying event (a.k.a. the jet pedestal effect). Keep in mind, however, that especially at low multiplicities, and when gaps are present, the contributions from diffractive processes should not be ignored.

Due to the simple fact that hadrons are composite, multi-parton interactions (several distinct parton-parton interactions in one and the same hadron-hadron collision) will always be there — but how many, and how much additional energy and tracks do they deposit in a given measurement region? The first detailed Monte Carlo model for perturbative MPI was proposed by Sjöstrand in [102], and with some variation this still forms the basis for most modern implementations [2].

The first crucial observation is that the t -channel propagators appearing in perturbative QCD $2 \rightarrow 2$ scattering almost go on shell at low p_\perp , causing the differential cross sections to become very large, behaving roughly as

$$d\sigma_{2 \rightarrow 2} \propto \frac{dt}{t^2} \sim \frac{dp_\perp^2}{p_\perp^4}. \quad (35)$$

At LHC energies, this *parton-parton* cross section becomes larger than the total *hadron-hadron* cross section at p_\perp scales of order 4 – 5 GeV [103]. In the context of MPI models, this is interpreted straightforwardly to mean that *each* hadron-hadron collision contains *several* parton-parton collisions.

In the limit that all the partonic interactions are independent and equivalent, one would simply have a Poisson distribution in the number of MPI, with average

$$\langle n \rangle(p_{\perp \min}) = \frac{\sigma_{2 \rightarrow 2}(p_{\perp \min})}{\sigma_{\text{tot}}}, \quad (36)$$

with $p_{\perp \min}$ a lower cutoff scale which we shall return to below, and σ_{tot} a measure of the inelastic hadron-hadron cross section. This simple reinterpretation in fact expresses unitarity; instead of the total interaction cross section diverging as $p_{\perp \min} \rightarrow 0$ (which would violate unitarity), we have restated the problem so that it is now the *number of MPI per collision* that diverges, with the total cross section remaining finite.

Two important ingredients remain to fully regulate the remaining divergence. Firstly, the interactions cannot use up more momentum than is available in the parent hadron. This suppresses the large- n tail of the estimate above. In PYTHIA-based models, the MPI are ordered in p_\perp [54, 102], 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 [104, 105], instead the uncorrelated 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.

The second ingredient invoked to suppress the number of interactions, at low p_\perp and x , is colour screening; if the wavelength $\sim 1/p_\perp$ of an exchanged coloured parton becomes larger than a typical colour-anticolour separation distance, it will only see an *average* colour charge that vanishes in the limit $p_\perp \rightarrow 0$, hence leading to suppressed interactions. This provides an infrared cutoff for MPI similar to that provided by the hadronization scale for parton showers. A first estimate of the colour-screening cutoff would be the proton size, $p_{\perp \min} \approx \hbar/r_p \approx 0.3 \text{ GeV} \approx \Lambda_{\text{QCD}}$, 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 colour screening distance,” i.e., an average size of a region within which the net compensation of a given colour charge occurs. This number is not known from first principles, though it may be related to saturation [106]. In current MPI models, it is perceived of simply as an effective cutoff parameter, to be determined from data.

Note that the partonic cross sections depend upon the PDF set used, and therefore the optimal value to use for the cutoff will also depend on this choice [107]. 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 colour-screening

distance d decreases. The uncertainty on the scaling of the cutoff is a major concern when extrapolating between different collider energies [107–109].

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. That is, the so-called “underlying event” (UE) is much more active, with larger fluctuations, than the average min-bias event. This is interpreted as a consequence of impact-parameter-dependence: in peripheral collisions, only a small fraction of events contain any high- p_{\perp} activity, whereas central collisions are more likely to contain at least one hard scattering; a high- p_{\perp} triggered sample will therefore be biased towards small impact parameters, b , with a large number of MPI (and associated increased activity). The ability of a model to describe the shape of the pedestal (e.g. to describe both MB and UE distributions simultaneously) therefore depends upon its modeling of the b -dependence, and correspondingly the impact-parameter shape constitutes another main tuning parameter. A detailed discussion of impact-parameter dependent models goes beyond the scope of these lectures; see [98, 102].

For hard processes at the LHC at 13 TeV, the transverse energy, E_T , in the UE is expected to be about 3.3 GeV per unit $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ area [109], though with large event-to-event fluctuations of order ± 2 GeV [110]. Thus, for example, the E_T originating from the UE, in a cone with radius 0.4 can be estimated to be $E_{T\text{UE}}(R = 0.4) \sim 1.6 \pm 1$ GeV, while the E_T in cones with radii 0.7 and 1.0 would be $E_{T\text{UE}}(R = 0.7) \sim 5 \pm 3$ GeV and $E_{T\text{UE}}(R = 1.0) \sim 10 \pm 6$ GeV, respectively.

3.4 Tuning

A main virtue of general-purpose Monte Carlo event generators is their ability to provide a complete and fully differential picture of collider final states, down to the level of individual particles. As has been emphasized in these lectures, the achievable accuracy depends both on the inclusiveness of the chosen observable and on the sophistication of the simulation itself. An important driver for the latter is obviously the development of improved theoretical models, e.g., by including matching to higher-order matrix elements, more accurate resummations, or better non-perturbative models, as discussed in the previous sections; but it also depends crucially on the available constraints on the remaining free parameters of the model. Using existing data (or more precise calculations) to constrain these is referred to as generator tuning.

Keep in mind that generators attempt to deliver a *global* description of the data; a tune is no good if it fits one distribution perfectly, but not any others. It is therefore crucial to study the simultaneous degree of agreement or disagreement over many, mutually complementary, distributions. A useful online resource for making such comparisons can be found at the [MC PLOTS](#) web site [111] (which relies on computing power donated by volunteers, via the [LHC@home](#) project [112]). The analyses come from the comprehensive RIVET analysis toolkit [14], which can also be run stand-alone to make your own MC tests and comparisons.

Although MC models may appear to have a bewildering number of independently adjustable parameters, it is worth noting that most of these only control relatively small (exclusive) details of the event generation. The majority of the (inclusive)

physics is determined by only a few, very important ones, such as the value of the strong coupling, in the perturbative domain, and the form of the fragmentation function for massless partons, in the non-perturbative one.

Armed with a good understanding of the underlying model, an expert would therefore normally take a highly factorized approach to constraining the parameters, first constraining the perturbative ones (using IR safe observables and/or more precise theory calculations) and thereafter the non-perturbative ones, each ordered in a measure of their relative significance to the overall modeling. This allows one to concentrate on just a few parameters and a few carefully chosen distributions at a time, reducing the full parameter space to manageable-sized chunks. Still, each step will often involve more than one single parameter, and non-factorizable correlations may still necessitate additional iterations from the beginning before a fully satisfactory set of parameters is obtained.

Recent years have seen the emergence of automated tools that attempt to reduce the amount of both computer and manpower required for this task, for instance by making full generator runs only for a limited set of parameter points, and then interpolating between these to obtain approximations to what the true generator result would have been for any intermediate parameter point, as, e.g., in PROFESSOR [113]. Automating the human expert input is more difficult. Currently, this is addressed by a combination of input solicited from the generator authors (e.g., which parameters and ranges to consider, which observables constitute a complete set, etc) and the elaborate construction of non-trivial weighting functions that determine how much weight is assigned to each individual bin in each distribution. The field is still burgeoning, and future sophistications are to be expected. Nevertheless, at this point the overall quality of the tunes obtained with automated methods appear to at least be competitive with the manual ones.

However, though we have very good LHC tunes for essentially all the general-purpose generators by now, there are two important aspects which have so far been sorely neglected, and which it is becoming increasingly urgent to address. The first is that a central tune is not really worth much, unless you know what the uncertainty on it is. A few individual proposals for systematic tuning variations have been made [108, 114], but so far there is no general approach for establishing MC uncertainties by tune variations. The second issue is that virtually all generator tuning is done at the “pure” LL shower level, and not much is known about what happens to the tuning when matrix-element matching is subsequently included.

Finally, rather than performing one global tune to all the data, as is usually done, a more systematic check on the validity of the underlying physics model could be obtained by instead performing several independent optimizations of the model parameters for a range of different phase-space windows and/or collider environments. In regions in which consistent parameter sets are obtained (with reasonable $\Delta\chi^2$ values), the underlying model can be considered as interpolating well, i.e., it is universal. If not, a breakdown in the ability of the model to span different physical regimes has been identified, and can be addressed, with the nature of the deviations giving clues as to the nature of the breakdown. With the advent of automated tools, such systematic studies are now becoming feasible, with a first example given in [107].

We round off by giving a sketch of a reasonably complete tuning procedure, with-

out going into details about the parameters that control each of these sectors in individual Monte Carlo models:

1) Keep in mind that inabilities of models to describe data is a vital part of the feedback cycle between theory and experiment. Also keep in mind that perturbation theory at (N)LO+LL is doing *very well* if it gets within 10% of a given IR safe measurement. An agreement of 5% should be considered the absolute sanity limit, beyond which it does not make any sense to tune further. 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.

2) Final-State Radiation and Hadronization: mainly using LEP and other e^+e^- collider data. On the IR safe side, there are event shapes and jet observables. On the IR sensitive side, multiplicities and particle spectra. Pay attention to the high- z tail of the fragmentation, where a single hadron carries a large fraction of an entire jet’s momentum (most likely to give “fakes”). Depending on the focus of the tuning, attention should also be paid to identified-particle rates and ratios (perhaps with a nod to hadron-collider measurements), and to fragmentation in events containing heavy quarks and/or gluon jets. Usually, more weight is given to those particles that are most copiously produced. The scaling properties of IR safe vs. IR sensitive contributions can be tested by comparing data at several different e^+e^- collider energies.

3) Initial-State Radiation, and “Primordial⁶ k_T ”: the main constraining distribution is the dilepton p_\perp distribution in Drell-Yan events in hadron-hadron collisions. Ideally, one would like to use several different Q^2 values, and/or complementary processes, like $p_\perp(\text{dijet})$ or $p_\perp(t\bar{t})$. For any observables containing explicit jets, be aware that the UE can produce small horizontal shifts in jet p_\perp distributions, which may in turn result in larger-than-expected vertical changes if the distributions are falling sharply. Also note that the ISR evolution is sensitive to the choice of PDFs.

4) Initial-Final Connections: (radiation from colour lines connected to the initial state): jet broadening in hadron collisions. This is one of the most poorly controlled parts of most MC models, highly sensitive to the treatment of coherence. Keep in mind that it is *not* directly constrained by pure final-state observables, such as LEP fragmentation, nor by pure initial-state observables, such as the Drell-Yan p_\perp spectrum, which is why we list it as a separate item here. The modeling of this aspect can have important effects on specific observables, a recent example being the $t\bar{t}$ forward-backward asymmetry at the Tevatron [115].

5) Underlying Event: Good constraints on the overall level of the underlying event can be obtained by counting the summed transverse energy (more IR safe) and/or particle multiplicities and average transverse momenta (more IR sensitive) in regions *transverse* to a hard trigger jet (more IR safe) or particle (more IR sensitive), see e.g. [116]. Constraints on the *fluctuations* of the underlying event are also important, and can be obtained, e.g., by comparing to measurements of the RMS of such distributions [110]. Again, note that the UE is sensitive to the choice of PDFs [107].

6) Colour (Re-)Connections and other Final-State Interactions: By Final-State Interactions, we intend a broad spectrum of possible collective effects that

⁶Primordial k_T : an additional soft p_\perp component that is injected on top of the p_\perp generated by the initial-state shower itself, see [2, Section 7.1].

may be included to a greater or lesser extent in various models. These effects include: Bose-Einstein correlations (see, e.g., [117]), rescattering (see, e.g., [118]), colour reconnections / string interactions (see, e.g., [119–121]), hydrodynamics (see, e.g., [122]), etc. As a rule, these effects are soft and/or non-perturbative and hence should not modify hard IR safe observables appreciably. They can, however, have *drastic* effects on IR sensitive ones, such as particle multiplicities, momentum distributions, and correlations, wherefore useful constraints are typically furnished by measurements of spectra and correlations as functions of quantities believed to serve as indicators of the strength of these phenomena (such as event multiplicity), and/or by collective-flow-type measurements. Finally, if the model includes a universal description of underlying event and soft-inclusive QCD, as many MPI-based models do, then minimum-bias data can also be used as a control sample, though one must then be careful either to address diffractive contributions properly or to include only gap-suppressed data samples. A complete MB and UE model should also be able to describe the rise of the pedestal from MB to UE, e.g., in the transverse UE observables (see above).

7) Beam Remnants: Constraints on beam remnant fragmentation (see, e.g., [98]) are most easily obtained in the forward region, but, e.g., the amount of baryon transport from the remnant to a given rapidity region can also be used to probe how much the colour structure of the remnant was effectively disturbed, with more baryon transport indicating a larger amount of “beam baryon blowup”.

References

- [1] **Particle Data Group** Collaboration, J. Beringer *et al.*, “Review of Particle Physics (RPP),” *Phys.Rev.* **D86** (2012) 010001.
- [2] **MCnet** Collaboration, A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Höche, *et al.*, “General-purpose event generators for LHC physics,” *Phys.Rept.* **504** (2011) 145, [1101.2599](#).
- [3] P. Z. Skands, “QCD for Collider Physics,” [1104.2863](#). Lectures given at ESHEP 2010, Raseborg, Finland, arXiv:1202.1629.
- [4] P. Skands, “Introduction to QCD,” [1207.2389](#). Lectures given at TASI 2012, Boulder, Colorado.
- [5] R. Kleiss, W. Stirling, and S. Ellis, “A New Monte Carlo Treatment of Multiparticle Phase Space at High-Energies,” *Comput.Phys.Commun.* **40** (1986) 359.
- [6] G. Lepage, “A New Algorithm for Adaptive Multidimensional Integration,” *J.Comput.Phys.* **27** (1978) 192. Revised version.
- [7] G. P. Lepage, “VEGAS - An Adaptive Multi-dimensional Integration Program, ,” CLNS-80/447.
- [8] F. James, “Monte Carlo theory and practice,” *Rept. Prog. Phys.* **43** (1980) 1145.
- [9] M. Bähr, S. Gieseke, M. Gigg, D. Grellscheid, K. Hamilton, *et al.*, “Herwig++ Physics and Manual,” *Eur.Phys.J.* **C58** (2008) 639, [0803.0883](#).

- [10] T. Sjöstrand, S. Mrenna, and P. Skands, “A brief introduction to PYTHIA 8.1,” *Comput. Phys. Commun.* **178** (2008) 852, [0710.3820](#).
- [11] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter, “Event generation with SHERPA 1.1,” *JHEP* **02** (2009) 007, [0811.4622](#).
- [12] E. Boos *et al.*, “Generic user process interface for event generators, ,” [hep-ph/0109068](#).
- [13] J. Alwall *et al.*, “A standard format for Les Houches Event Files,” *Comput. Phys. Commun.* **176** (2007) 300, [hep-ph/0609017](#).
- [14] A. Buckley, J. Butterworth, L. Lönnblad, H. Hoeth, J. Monk, *et al.*, “Rivet user manual,” [1003.0694](#).
- [15] **GEANT4** Collaboration, S. Agostinelli *et al.*, “GEANT4: A Simulation toolkit,” *Nucl.Instrum.Meth.* **A506** (2003) 250.
- [16] M. Dobbs and J. B. Hansen, “The HepMC C++ Monte Carlo event record for High Energy Physics,” *Comput.Phys.Commun.* **134** (2001) 41–46.
- [17] G. Dissertori, I. Knowles, and M. Schmelling, *Quantum Chromodynamics — High energy experiments and theory*. Oxford Science Publications, 2003.
- [18] R. Ellis, W. Stirling, and B. Webber, *QCD and collider physics*. Camb. Monogr. Part. Phys. Nucl. Phys. Cosm., 1996.
- [19] T. Sjöstrand, “Monte Carlo Generators,” [hep-ph/0611247](#).
- [20] J. C. Collins and D. E. Soper, “Parton Distribution and Decay Functions,” *Nucl.Phys.* **B194** (1982) 445.
- [21] **CTEQ** Collaboration, R. Brock *et al.*, “Handbook of perturbative QCD: Version 1.1,” *Rev.Mod.Phys.* **67** (1995) 157–248.
- [22] T. Plehn, “LHC Phenomenology for Physics Hunters,” [0810.2281](#).
- [23] G. F. Sterman, “Partons, factorization and resummation, TASI 95,” [hep-ph/9606312](#).
- [24] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*. Addison-Wesley, 1995.
- [25] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, *et al.*, “MadGraph/MadEvent v4: The New Web Generation,” *JHEP* **0709** (2007) 028, [0706.2334](#).
- [26] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, “MadGraph 5 : Going Beyond,” *JHEP* **1106** (2011) 128, [1106.0522](#).
- [27] A. Kanaki and C. G. Papadopoulos, “HELAC: A Package to compute electroweak helicity amplitudes,” *Comput.Phys.Commun.* **132** (2000) 306, [hep-ph/0002082](#).
- [28] F. Krauss, R. Kuhn, and G. Soff, “AMEGIC++ 1.0: A Matrix element generator in C++,” *JHEP* **0202** (2002) 044, [hep-ph/0109036](#).
- [29] M. Moretti, T. Ohl, and J. Reuter, “O’Mega: An Optimizing matrix element generator,” [hep-ph/0102195](#).

- [30] **CompHEP** Collaboration, E. Boos *et al.*, “CompHEP 4.4: Automatic computations from Lagrangians to events,” *Nucl.Instrum.Meth.* **A534** (2004) 250, [hep-ph/0403113](#).
- [31] A. Pukhov, “Calcchep 2.3: MSSM, structure functions, event generation, 1, and generation of matrix elements for other packages,” [hep-ph/0412191](#).
- [32] W. Kilian, T. Ohl, and J. Reuter, “WHIZARD: Simulating Multi-Particle Processes at LHC and ILC,” *Eur.Phys.J.* **C71** (2011) 1742, [0708.4233](#).
- [33] A. Cafarella, C. G. Papadopoulos, and M. Worek, “Helac-Phegas: A Generator for all parton level processes,” *Comput.Phys.Commun.* **180** (2009) 1941, [0710.2427](#).
- [34] T. Gleisberg and S. Höche, “Comix, a new matrix element generator,” *JHEP* **0812** (2008) 039, [0808.3674](#).
- [35] G. P. Salam, “Elements of QCD for hadron colliders,” [1011.5131](#).
- [36] F. Bloch and A. Nordsieck, “Note on the Radiation Field of the electron,” *Phys.Rev.* **52** (1937) 54–59.
- [37] T. Kinoshita, “Mass singularities of Feynman amplitudes,” *J. Math. Phys.* **3** (1962) 650–677.
- [38] T. D. Lee and M. Nauenberg, “Degenerate Systems and Mass Singularities,” *Phys. Rev.* **133** (1964) B1549–B1562.
- [39] V. N. Gribov and L. N. Lipatov, “Deep inelastic e - p scattering in perturbation theory,” *Sov. J. Nucl. Phys.* **15** (1972) 438.
- [40] G. Altarelli and G. Parisi, “Asymptotic freedom in parton language,” *Nucl. Phys.* **B126** (1977) 298.
- [41] Y. L. Dokshitzer, “Calculation of the structure functions for deep inelastic scattering and e^+e^- annihilation by perturbation theory in quantum chromodynamics,” *Sov. Phys. JETP* **46** (1977) 641–653.
- [42] A. J. Larkoski, J. J. Lopez-Villarejo, and P. Skands, “Helicity-Dependent Showers and Matching with VINCIA,” *Phys.Rev.* **D87** (2013) 054033, [1301.0933](#).
- [43] A. Gehrmann-De Ridder and M. Ritzmann, “NLO Antenna Subtraction with Massive Fermions,” *JHEP* **0907** (2009) 041, [0904.3297](#).
- [44] A. Gehrmann-De Ridder, M. Ritzmann, and P. Skands, “Timelike Dipole-Antenna Showers with Massive Fermions,” *Phys.Rev.* **D85** (2012) 014013, [1108.6172](#).
- [45] G. Gustafson and U. Pettersson, “Dipole formulation of QCD cascades,” *Nucl. Phys.* **B306** (1988) 746.
- [46] D. A. Kosower, “Antenna factorization of gauge-theory amplitudes,” *Phys. Rev.* **D57** (1998) 5410–5416, [hep-ph/9710213](#).
- [47] D. A. Kosower, “Antenna factorization in strongly-ordered limits,” *Phys. Rev.* **D71** (2005) 045016, [hep-ph/0311272](#).
- [48] A. Gehrmann-De Ridder, T. Gehrmann, and E. W. N. Glover, “Antenna subtraction at NNLO,” *JHEP* **09** (2005) 056, [hep-ph/0505111](#).

- [49] S. Catani and M. H. Seymour, “A general algorithm for calculating jet cross sections in NLO QCD,” *Nucl. Phys.* **B485** (1997) 291–419, [hep-ph/9605323](#).
- [50] S. Catani and M. H. Seymour, “The Dipole Formalism for the Calculation of QCD Jet Cross Sections at Next-to-Leading Order,” *Phys. Lett.* **B378** (1996) 287, [hep-ph/9602277](#).
- [51] G. Marchesini and B. R. Webber, “Monte Carlo Simulation of General Hard Processes with Coherent QCD Radiation,” *Nucl. Phys.* **B310** (1988) 461.
- [52] S. Gieseke, P. Stephens, and B. Webber, “New formalism for QCD parton showers,” *JHEP* **12** (2003) 045, [hep-ph/0310083](#).
- [53] M. Bengtsson and T. Sjöstrand, “A comparative study of coherent and non-coherent parton shower evolution,” *Nucl. Phys.* **B289** (1987) 810.
- [54] T. Sjöstrand and P. Z. Skands, “Transverse-momentum-ordered showers and interleaved multiple interactions,” *Eur. Phys. J.* **C39** (2005) 129–154, [hep-ph/0408302](#).
- [55] L. Lönnblad, “Ariadne version 4: A program for simulation of QCD cascades implementing the colour dipole model,” *Comput. Phys. Commun.* **71** (1992) 15.
- [56] Z. Nagy and D. E. Soper, “Matching parton showers to NLO computations,” *JHEP* **10** (2005) 024, [hep-ph/0503053](#).
- [57] S. Schumann and F. Krauss, “A Parton shower algorithm based on Catani-Seymour dipole factorisation,” *JHEP* **0803** (2008) 038, [0709.1027](#).
- [58] W. T. Giele, D. A. Kosower, and P. Z. Skands, “A Simple shower and matching algorithm,” *Phys. Rev.* **D78** (2008) 014026, [0707.3652](#).
- [59] L. Hartgring, E. Laenen, and P. Skands, “Antenna Showers with One-Loop Matrix Elements,” [1303.4974](#).
- [60] S. Catani, B. R. Webber, and G. Marchesini, “QCD coherent branching and semiinclusive processes at large x ,” *Nucl. Phys.* **B349** (1991) 635–654.
- [61] J. C. Collins, “Spin correlations in Monte Carlo event generators,” *Nucl. Phys.* **B304** (1988) 794.
- [62] I. Knowles, “Spin Correlations in Parton - Parton Scattering,” *Nucl.Phys.* **B310** (1988) 571.
- [63] P. Richardson, “Spin correlations in Monte Carlo simulations,” *JHEP* **11** (2001) 029, [hep-ph/0110108](#).
- [64] G. Marchesini and B. Webber, “Simulation of QCD Jets Including Soft Gluon Interference,” *Nucl.Phys.* **B238** (1984) 1.
- [65] M. A. Gigg and P. Richardson, “Simulation of Finite Width Effects in Physics Beyond the Standard Model,” [0805.3037](#).
- [66] W. Giele, D. Kosower, and P. Skands, “Higher-Order Corrections to Timelike Jets,” *Phys.Rev.* **D84** (2011) 054003, [1102.2126](#).

- [67] M. Bengtsson and T. Sjöstrand, “Coherent Parton Showers Versus Matrix Elements: Implications of PETRA - PEP Data,” *Phys. Lett.* **B185** (1987) 435.
- [68] J. Lopez-Villarejo and P. Skands, “Efficient Matrix-Element Matching with Sector Showers,” *JHEP* **1111** (2011) 150, [1109.3608](#).
- [69] T. Sjöstrand, S. Mrenna, and P. Skands, “PYTHIA 6.4 physics and manual,” *JHEP* **05** (2006) 026, [hep-ph/0603175](#).
- [70] E. Norrbin and T. Sjöstrand, “QCD radiation off heavy particles,” *Nucl. Phys.* **B603** (2001) 297–342, [hep-ph/0010012](#).
- [71] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with parton shower simulations: the POWHEG method,” *JHEP* **11** (2007) 070, [0709.2092](#).
- [72] S. Alioli, P. Nason, C. Oleari, and E. Re, “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX,” *JHEP* **06** (2010) 043, [1002.2581](#).
- [73] S. Frixione and B. R. Webber, “Matching NLO QCD computations and parton shower simulations,” *JHEP* **06** (2002) 029, [hep-ph/0204244](#).
- [74] S. Frixione and B. R. Webber, “The MC@NLO 3.4 Event Generator,” [0812.0770](#).
- [75] G. Corcella *et al.*, “HERWIG 6: an event generator for hadron emission reactions with interfering gluons,” *JHEP* **01** (2001) 010, [hep-ph/0011363](#).
- [76] M. H. Seymour, “A simple prescription for first-order corrections to quark scattering and annihilation processes,” *Nucl. Phys.* **B436** (1995) 443, [hep-ph/9410244](#).
- [77] M. H. Seymour, “Matrix-element corrections to parton shower algorithms,” *Comp. Phys. Commun.* **90** (1995) 95, [hep-ph/9410414](#).
- [78] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, “QCD matrix elements + parton showers,” *JHEP* **11** (2001) 063, [hep-ph/0109231](#).
- [79] L. Lönnblad, “Correcting the colour-dipole cascade model with fixed order matrix elements,” *JHEP* **05** (2002) 046, [hep-ph/0112284](#).
- [80] N. Lavesson and L. Lönnblad, “ W + jets matrix elements and the dipole cascade,” *JHEP* **07** (2005) 054, [hep-ph/0503293](#).
- [81] M. L. Mangano, M. Moretti, F. Piccinini, and M. Treccani, “Matching matrix elements and shower evolution for top-pair production in hadronic collisions,” *JHEP* **01** (2007) 013, [hep-ph/0611129](#).
- [82] S. Mrenna and P. Richardson, “Matching matrix elements and parton showers with HERWIG and PYTHIA,” *JHEP* **05** (2004) 040, [hep-ph/0312274](#).
- [83] L. Lönnblad and S. Prestel, “Unitarising Matrix Element + Parton Shower merging,” *JHEP* **1302** (2013) 094, [1211.4827](#).
- [84] K. Hamilton and P. Nason, “Improving NLO-parton shower matched simulations with higher order matrix elements,” *JHEP* **06** (2010) 039, [1004.1764](#).
- [85] R. Frederix and S. Frixione, “Merging meets matching in MC@NLO,” *JHEP* **1212** (2012) 061, [1209.6215](#).

- [86] L. Lönnblad and S. Prestel, “Merging Multi-leg NLO Matrix Elements with Parton Showers,” *JHEP* **1303** (2013) 166, [1211.7278](#).
- [87] S. Plätzer, “Controlling inclusive cross sections in parton shower + matrix element merging,” *JHEP* **1308** (2013) 114, [1211.5467](#).
- [88] K. Hamilton, P. Nason, C. Oleari, and G. Zanderighi, “Merging H/W/Z + 0 and 1 jet at NLO with no merging scale: a path to parton shower + NNLO matching,” *JHEP* **1305** (2013) 082, [1212.4504](#).
- [89] B. Cooper, J. Katzy, M. Mangano, A. Messina, L. Mijovic, *et al.*, “Importance of a consistent choice of α_s in the matching of AlpGen and Pythia,” *Eur.Phys.J.* **C72** (2012) 2078, [1109.5295](#).
- [90] N. Lavesson and L. Lönnblad, “Extending CKKW-merging to one-loop matrix elements,” *JHEP* **12** (2008) 070, [0811.2912](#).
- [91] D. Amati and G. Veneziano, “Preconfinement as a Property of Perturbative QCD,” *Phys. Lett.* **B83** (1979) 87.
- [92] X. Artru and G. Mennessier, “String model and multiproduction,” *Nucl. Phys.* **B70** (1974) 93–115.
- [93] B. Andersson, *The Lund model*. Camb. Monogr. Part. Phys. Nucl. Phys. Cosm., 1997.
- [94] B. Andersson, G. Gustafson, and T. Sjöstrand, “A Model for Baryon Production in Quark and Gluon Jets,” *Nucl. Phys.* **B197** (1982) 45.
- [95] B. Andersson, G. Gustafson, and T. Sjöstrand, “Baryon production in jet fragmentation and Upsilon decay,” *Phys.Scripta* **32** (1985) 574.
- [96] P. Eden and G. Gustafson, “Baryon production in the string fragmentation picture,” *Z.Phys.* **C75** (1997) 41–49, [hep-ph/9606454](#).
- [97] T. Sjöstrand and P. Z. Skands, “Baryon number violation and string topologies,” *Nucl. Phys.* **B659** (2003) 243, [hep-ph/0212264](#).
- [98] T. Sjöstrand and P. Z. Skands, “Multiple interactions and the structure of beam remnants,” *JHEP* **03** (2004) 053, [hep-ph/0402078](#).
- [99] M. G. Bowler, “ e^+e^- production of heavy quarks in the string model,” *Z. Phys.* **C11** (1981) 169.
- [100] K. G. Wilson, “Confinement of Quarks,” *Phys.Rev.* **D10** (1974) 2445.
- [101] V. Khoze, F. Krauss, A. Martin, M. Ryskin, and K. Zapp, “Diffraction and correlations at the LHC: Definitions and observables,” *Eur.Phys.J.* **C69** (2010) 85–93, [1005.4839](#).
- [102] T. Sjöstrand and M. van Zijl, “A multiple-interaction model for the event structure in hadron collisions,” *Phys. Rev.* **D36** (1987) 2019.
- [103] M. Bähr, J. M. Butterworth, and M. H. Seymour, “The Underlying Event and the Total Cross Section from Tevatron to the LHC,” *JHEP* **01** (2009) 065, [0806.2949](#).
- [104] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour, “Multiparton Interactions in Photoproduction at HERA,” *Z. Phys.* **C72** (1996) 637–646, [hep-ph/9601371](#).

- [105] M. Bähr, J. M. Butterworth, S. Gieseke, and M. H. Seymour, “Soft interactions in Herwig++,” [0905.4671](#).
- [106] M. Ryskin, A. Martin, and V. Khoze, “High-energy strong interactions: from ‘hard’ to ‘soft’,” *Eur.Phys.J.* **C71** (2011) 1617, [1102.2844](#).
- [107] H. Schulz and P. Skands, “Energy Scaling of Minimum-Bias Tunes,” *Eur.Phys.J.* **C71** (2011) 1644, [1103.3649](#).
- [108] P. Z. Skands, “Tuning Monte Carlo Generators: The Perugia Tunes,” *Phys.Rev.* **D82** (2010) 074018, [1005.3457](#).
- [109] P. Z. Skands, “Soft-QCD and UE spectra in pp collisions at very high CM energies (a Snowmass white paper),” [1308.2813](#).
- [110] **ATLAS** Collaboration, G. Aad *et al.*, “Measurement of underlying event characteristics using charged particles in pp collisions at $\sqrt{s} = 900\text{GeV}$ and 7 TeV with the ATLAS detector,” *Phys.Rev.* **D83** (2011) 112001, [1012.0791](#).
- [111] A. Karneyeu, L. Mijovic, S. Prestel, and P. Skands, “MCPLLOTS: a particle physics resource based on volunteer computing,” [1306.3436](#).
- [112] D. Lombrana Gonzalez, L. Rivkin, M. Zacharov, I. an Giovannozzi, E. McIntosh, *et al.*, “LHC@home: a Volunteer Computing System for Massive Numerical Simulations of Beam Dynamics and High Energy Physics Events,” *Conf.Proc.* **C1205201** (2012) 505–507.
- [113] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and J. E. von Seggern, “Systematic event generator tuning for the LHC,” *Eur. Phys. J.* **C65** (2010) 331–357, [0907.2973](#).
- [114] P. Richardson and D. Winn, “Investigation of MC Uncertainties on Higgs Boson searches using Jet Substructure,” *Eur.Phys.J.* **C72** (2012) 2178, [1207.0380](#).
- [115] P. Skands, B. Webber, and J. Winter, “QCD Coherence and the Top Quark Asymmetry,” *JHEP* **1207** (2012) 151, [1205.1466](#).
- [116] R. Field, “Min-Bias and the Underlying Event at the LHC,” *Acta Phys.Polon.* **B42** (2011) 2631–2656, [1110.5530](#).
- [117] L. Lönnblad and T. Sjöstrand, “Modeling Bose-Einstein correlations at LEP-2,” *Eur.Phys.J.* **C2** (1998) 165, [hep-ph/9711460](#).
- [118] R. Corke and T. Sjöstrand, “Multiparton Interactions and Rescattering,” *JHEP* **1001** (2010) 035, [0911.1909](#).
- [119] J. Rathsman, “A generalised area law for hadronic string reinteractions,” *Phys. Lett.* **B452** (1999) 364–371, [hep-ph/9812423](#).
- [120] P. Z. Skands and D. Wicke, “Non-perturbative QCD effects and the top mass at the Tevatron,” *Eur.Phys.J.* **C52** (2007) 133–140, [hep-ph/0703081](#).
- [121] S. Gieseke, C. Rohr, and A. Siodmok, “Colour reconnections in Herwig++,” *Eur.Phys.J.* **C72** (2012) 2225, [1206.0041](#).
- [122] K. Werner, I. Karpenko, and T. Pierog, “Hydrodynamical evolution in heavy ion collisions and p p scatterings at the LHC: Ridges in A A and p p scattering,” *Acta Phys.Polon.Supp.* **4** (2011) 629–634.