CERN–TH/95–10
16 January 1995

Event Generators in Particle Physics*

Torbjörn Sjöstrand**

Theory Division, CERN
CH-1211 Geneva 23, Switzerland

Abstract

This presentation gives an introduction to the topic of event generators in particle physics. The emphasis is on the physics aspects that have to be considered in the construction of a generator, and what lessons we have learned from comparisons with data. A brief survey of existing generators is also included. As illustration, a few topics of current interest are covered in a bit more detail: QCD uncertainties in W mass determinations and γp/$\gamma\gamma$ physics.

*To appear in the Proceedings of the XV Brazilian National Meeting on Particles and Fields, Angra dos Reis, Brazil, October 1994.
**On leave of absence (until February 1) from the Department of Theoretical Physics 2, University of Lund, Lund, Sweden.

CERN–TH/95–10

1 Introduction

The ultimate goal of particle physics is to find and understand *the* underlying theory of the Universe. Experimental progress in this direction requires exploration at ever higher energies. This way one hopes to gain access to new particles and reactions not allowed at lower energies, to observe hints of symmetries that are spontaneously broken at some high energy scale, and to profit from a smaller $\alpha_{\mathrm{s}}(Q^2)$. The latter point also implies a reduced dependence of non-perturbative physics, such as confinement and hadronic wave functions, that are not all too well understood. Experimental data are consistent with the hypothesis that, above some grand unification scale, the three couplings of the standard model come together (at least if supersymmetry is included in the game) [1]. So it is no wonder that theorists' folklore says that physics is simpler at higher energies.

From an experimental point of view, this is all wrong. If the energy of a process is increased, the amount of cascading is also increased. With cascading I here mean, in a very broad sense, all the mechanisms that increase the number of quanta that are needed to describe the event as times goes by: decays of exotic new particles or the familiar W and Z resonances, initial- and final-state QCD parton showers, fragmentation mechanisms that turn partons into hadrons, and decays of ordinary unstable particles. The lower cut-off of the cascading is given essentially by the pion mass, independently of the full energy of the process. So higher energies means a larger energy range over which cascading can occur, i.e. higher multiplicities, and thus larger experimental challenges.

The evolution of experimental particle physics is therefore towards larger complexity. In the fifties and sixties, the emulsion and bubble chamber data of the time showed every single vertex of the processes studied, usually with two or three outgoing particles. Around 1980, the turn-on of higher-energy $\mathrm{e^+e^-}$ and $\overline{\mathrm{p}}\mathrm{p}$ colliders gave events with tens of charged particles, sometimes even above 100. When we today plan ahead for the LHC, the expected average charged multiplicity per event is above 100. At nominal luminosity, with around 20 events overlayed in each single beam crossing, any physics will have to be dug out among 4000 charged or neutral particles!

There are ways to get back to some kind of simplicity. One is to consider (semi-) inclusive quantities, such as jets, where a set of particles is characterized just by a summed energy and a direction vector. These jets approximate the partons of a simpler perturbative description, but with a non-negligible smearing that has to be understood for precision physics. Another is to search only for especially clean (parts of) final states. As an example, a 200 GeV Higgs particle can decay into two Z^0's, each of which can subsequently decay to a lepton pair. So the observation of events with four well isolated leptons, with invariant mass distributions peaked in the appropriate places, would provide clear evidence for new physics. However, there is a price to be paid, in that the clean states normally have small branching ratios. When we search for a process that is rare in the first place, this may not always be acceptable. Combine this with an unrealistic definition of 'isolation' of leptons among 4000 other particles, and a promising signal may be gone. Remove the isolation criterion, and leptons from the decays of top quarks and bottom mesons will completely overwhelm the new physics. It then becomes imperative to understand the finer details of detector acceptance and resolution, to devise an analysis strategy that cuts away these backgrounds without too much of a loss to the signal.

As we see, it is seldom possible to avoid all the complications of high-multiplicity events. Furthermore, even if we firmly believe in QCD today, it is somewhat embarrassing that we know so little about how confinement really works. After all, this is a unique chance that the standard model provides us with to study strong-coupling physics! One line of approach is to devise models that allow various aspects of confinement to be tested by a judicious analysis of data. The foremost place where the two aspects above (experimental demands and theoretical curiosity) come together is event generators.

2 Event Generator Survey

In real life, an accelerator provides events. These events are registered in a detector as electronics signals. A data acquisition system stores a digitized and compressed version of the information, often only for the 'promising' part of the full event rate. In an event reconstruction program, the digitized electronics signals are turned back into a list of particle momenta and charges. The reconstructed events can then be used for physics analysis.

In the 'virtual reality' world, event generators take the rôle of the accelerator in providing the events. The response of the detector is then modelled in a detector simulation program. The most frequently used such program is GEANT [2]. This program provides an output in a format identical to that of the experimental data acquisition system, except that the original physics event input is also kept for later reference. Therefore the same reconstruction programs can be used, and the same physics analysis strategies. Comparisons between real data and simulated data form an important ingredient for the final results that can be published.

While event generators are not always as fast as desirable, the real bottleneck is detector simulation, where (for some applications) it is necessary to trace hadronic and electromagnetic shower evolution in excruciating detail. For many studies it is therefore common to jump directly from the event generator to the physics analysis. Detector effects are then completely neglected, or simulated by simple geometrical cuts and rule-of-thumb smearing of momentum vectors.

The phenomenologist is not normally concerned with the detector-specific aspects, and therefore may use the event generator as it is to explore various potentially interesting aspects. He/she might, for instance, introduce 'crackpot' alternative models just to check whether one should expect any testable consequences.

The description above may help illustrate why generators are useful. Let us try to give a somewhat more formalized list, subdivided by interest group. For the event generator author

- it allows theoretical studies of very *complex* multiparticle physics, by a subdivision of the complete problem into more manageable subtasks;
- it gives a larger flexibility in the spectrum of physical quantities that can be studied (no need to worry whether an observable is infrared finite or not in perturbative QCD);
- it provides a vehicle for the dissemination of interesting theoretical ideas to the experimental community;
- it allows a larger feedback from the experimental community, and hence a faster path for improved understanding of the underlying physics; and
- it is a source of fun and satisfaction (for some of us) to attack the non-trivial challenges.

For the experimental physicists (and many phenomenological ones as well), an event generator can be used

- to predict event rates and topologies, and hence to estimate the feasibility of an intended physics study in the first place;

- to simulate possible backgrounds, and hence to devise analysis strategies that optimize signal-to-background ratios;
- to study detector requirements, and hence to optimize the detector design and trigger strategy; and
- to study detector imperfections, and hence to evaluate acceptance corrections.

To the lists above, a final point should be added: nature is random! We are all familiar with the quantum mechanical uncertainty principle, with the principle of superposition, with the collapse of the wave function at measurements, and so on. What this means is that each event is unique. Had we had a perfect understanding of QCD, and infinite computing power, it would still have been a formidable task to enumerate all possible hadronic final states allowed, e.g. at LEP, as a function of the complete setup of all quantum numbers of the event (flavours, momenta, decay vertices, spins, ...), to calculate the complete matrix element for each such state, and to sum it all up to arrive at something as straightforward as a charged multiplicity distribution. It is therefore natural to subdivide the complete process into a sequence of smaller steps: the Z^0 decaying to a specific $q\overline{q}$ flavour; these developing a shower by consecutive branchings $q \to qg$ and $g \to gg$; fragmentation of a complex partonic system as the iterative production of one particle at a time; a sequential chain of secondary hadronic decays; and so on. In each step, nature is assumed to make a random choice between the allowed possible outcomes, and the relative probabilities may be calculated or modelled. This sequence of random choices may be simulated in an event generator by the use of random numbers. After each new step the set of possible states that could be reached is larger and more varied, until the final output has the full complexity observable in nature. Therefore, just as each experimental event is unique, so is each generator event. It is the average over many events that should be compared, and the fluctuations around this average.

There is one catch: the basic description of quantum mechanics is in terms of amplitudes rather than probabilities. One therefore has to watch out that a probabilistic description does not lose some of the fundamental aspects associated with interference terms. There is no generic recipe to handle this problem, but often nature is kind to us, so that reasonable ways out can be found.

There exist a wide range of generators, and by now the zoology may be quite confusing. However, in view of their increasing importance, a number of workshops have been devoted in part to collect information and critically compare all main generators by topic. Surveys are available for LEP 1 [3], for HERA [4], and for hadron colliders [5], and another will appear within the framework of the current LEP 2 workshop.

Generators can be designed for different purposes, and therefore also be quite different in size. The two with the widest scope are HERWIG [6] and PYTHIA/JETSET [7], which can be used for e^+e^-, ep and $\overline{p}p$ collision alike, which contain a wide range of allowed subprocesses, and which are attempts to cover all the aspects of the way from a hard process to a complex multihadronic final state. Almost in the same class is ISAJET [8], which is primarily intended for $\overline{p}p$ physics. The development and support of programs like these can easily be full-time efforts, where many model aspects related to non-perturbative QCD have to be developed from scratch.

There is then a broad spectrum of other generators, with more specific scopes. Some are devoted to the study of QCD parton shower evolution, or to the non-perturbative fragmentation modelling, or to more precise descriptions of particle decays, or to the simulation of multiparton matrix elements, or to multiple QED radiation, or to higher-order (including loop graphs) corrections to electroweak processes, or to a multitude of other tasks. Most of them do not contain models for fragmentation. So long as nonhadronic final states are considered, or observables not so sensitive to fragmentation, these programs are often superior to the general-purpose ones above. Therefore, the choice is sometimes between describing a few things very precisely or 'everything' at a reduced level of precision. In practice, often both are needed.

In the following, the discussion will tend to be centred more around the former kind of approach, firstly because this is where my own interests lie, and secondly because it gives me the chance to address a wide range of topics.

3 e^+e^- Physics

Jet physics started in earnest when the experiments at PETRA observed clean two- and three-jet events. Experience has shown that a sensible approach is to divide the process into four consecutive steps:

1. A hard process $e^+e^- \to \gamma^*/Z^0 \to q\overline{q}$. This process is perturbatively calculable in the standard model. Often initial-state QED radiation is also included in the description. The 'final state' of this step is given by the q flavour and angular distribution, plus possibly the distribution of additional photons.
2. A stage where perturbative QCD is applicable. Full second-order matrix-element calculations have been performed for the jet rate, which means that the production of two-, three- and four-jet events can be described consistently to that order. The game gets increasingly more complicated for each new order, however, at the same time as higher-order effects are clearly visible in the data. The alternative is therefore to adopt the parton-shower approach, wherein the evolution towards higher parton multiplicities is described as a sequence of branchings at decreasing virtualities, of the kinds $q \to qg$, $g \to gg$ and $g \to q\overline{q}$. This is an approximation to the correct answer, which should be good in the collinear limit but less good for widely separated jets. A standard method is to match the first branching of the shower to the first-order matrix element, so that a reasonable description is thereby obtained over the full kinematical range.
3. The fragmentation stage. When the shower is evolved towards smaller virtualities, the running α_s becomes larger, and ultimately a limit is hit where perturbation theory breaks down. In models this scale typically comes out to be around 1 GeV. Below this scale, the coloured partons are somehow transformed into colourless hadrons. Currently only phenomenologically motivated models are available, today normally string or cluster fragmentation.
4. Secondary decays occur since many of the hadrons produced above are unstable. Normally also this step involves non-perturbative physics, but experimentally determined branching ratios [9] can here often be used as input.

3.1 Parton Showers

The parton-shower picture is derived within the framework of the leading-logarithm approximation, LLA. In this picture, only the leading terms in the perturbative expansion are kept in a systematic manner. Some subleading corrections are included, as we shall see, but most are neglected. The overall theoretical picture is rather encouraging: there is reason to believe that neglected effects are small, and the predictive power of this approach is increasing year by year.

Phenomenologically, the main reason for the LLA success is our ability to formulate it in terms of a probabilistic picture, suitable for event generation. The probability $\mathcal{P}$ that a branching $a \to bc$ will take place during a small change $\mathrm{d}t = \mathrm{d}Q^2_{\mathrm{evol}}/Q^2_{\mathrm{evol}}$ of the evolution parameter $t = \ln(Q^2_{\mathrm{evol}}/\Lambda^2)$ is given by the evolution equations [10]

$$\frac{\mathrm{d}\mathcal{P}_{a\to bc}}{\mathrm{d}t} = \int \mathrm{d}z \frac{\alpha_{\mathrm{s}}(Q^2)}{2\pi} P_{a\to bc}(z) \; . \tag{1}$$

For gluons it is necessary to sum over all allowed final-state flavour combinations b and c to obtain the total branching probability. The $P_{a\to bc}(z)$ are the Altarelli–Parisi splitting kernels

$$\begin{aligned}
P_{\mathrm{q}\to\mathrm{qg}}(z) &= C_F \frac{1+z^2}{1-z} \; , \\
P_{\mathrm{g}\to\mathrm{gg}}(z) &= N_C \frac{(1-z(1-z))^2}{z(1-z)} \; , \\
P_{g\to\mathrm{q\overline{q}}}(z) &= T_R(z^2+(1-z)^2) \; ,
\end{aligned} \tag{2}$$

with $C_F = 4/3$, $N_C = 3$, and $T_R = n_f/2$, i.e. T_R receives a contribution of 1/2 for each allowed $\mathrm{q\overline{q}}$ flavour. The z variable specifies the sharing of four-momentum between the daughters, with daughter b taking fraction z and c taking $1-z$.

Starting at the maximum allowed virtuality t_{max} for parton a, the t parameter may be successively degraded. This does not mean that an individual parton runs through a range of t values: each parton in the end is associated with a fixed t value, and the evolution procedure is just a way of picking that value. It is only the ensemble of partons in many events that evolve continuously with t, cf. the concept of structure functions. The probability that no branching occurs during a small range of t values, δt, is given by $(1 - \delta t\, \mathrm{d}\mathcal{P}/\mathrm{d}t)$. When summed over many small intervals, the no-emission probability exponentiates

$$\mathcal{P}_{\mathrm{no-emission}}(t_{\mathrm{max}}, t) = \exp\left(- \int_t^{t_{\mathrm{max}}} \mathrm{d}t' \frac{\mathrm{d}\mathcal{P}_{a\to bc}}{\mathrm{d}t'} \right) \; . \tag{3}$$

This is (almost) what is normally called the Sudakov form factor. Thus the actual probability for a branching of a given t is the naive probability, eq. (1), multiplied by the probability that a branching has not already taken place, eq. (3). Compare with the exponential decay law of radioactive decays, with a t-dependent decay probability.

Once the branching of parton a has been selected, the products b and c may be allowed to branch in their turn, and so on, giving a tree-like structure of branchings at successively smaller t values. The branching of a given parton is stopped whenever the evolution parameter is below t_{min}.

Very valuable input for model builders is provided by the theoretical studies of corrections beyond leading log, such as coherence effects [11, 12, 13]. The latter come in two kinds:

- The intrajet coherence phenomenon is responsible for a decrease of the amount of soft gluon emission inside jets. It has been shown that an ordering in terms of a decreasing emission angle takes into account the bulk of soft gluon interference effects. Algorithms which contain angular ordering are loosely said to produce coherent showers, while those without generate conventional ones.
- The interjet coherence phenomenon, responsible for the flow of particles in between jets, with constructive or destructive interference depending on colour configuration ('colour drag phenomena'), cf. [12]. This form of coherence is not just a direct

consequence of the ordering of (polar) emission angles mentioned above, but also requires that azimuthal angles of branchings be properly distributed.

The angular ordering may be understood as follows, for the example of a branching $q_0 \to qg$. In the branching, the original q_0 colour is inherited by the gluon, while the q and g share a new colour–anticolour pair. A soft gluon g' (emitted at large angles) corresponds to a large (transverse) wavelength, so the soft gluon is unable to resolve the separate colour charges of the q and the g, and only feels the net charge. This is nothing but the original charge carried by the q_0. Such a soft gluon (in the region $\theta_{q_0 g'} > \theta_{qg}$) could therefore be thought of as being emitted by the q_0 rather than by the q–g system. If one only considers the emission that should be associated with the q or the g, to a good approximation, there is a complete destructive interference in the regions of non-decreasing opening angles, while partons radiate independently of each other inside the regions of decreasing opening angles ($\theta_{qg'} < \theta_{qg}$ and $\theta_{gg'} < \theta_{qg}$), once azimuthal angles are averaged over. The details of the colour interference pattern are reflected in non-uniform azimuthal emission probabilities.

Parton-shower programs generally give a good account of LEP data [14]: thrust distributions, jet masses, the number of jets as a function of the resolution parameter, and so on. In some variables, deviations are visible, such as the four-jet relative angular distributions used to test the triple-gluon vertex [15]. This is not so surprising, since the parton shower does not contain any explicit information about the four-jet matrix elements. So, while the overall rate of four-jet emission is well described, some of the azimuthal angular distributions are not correctly simulated.

Many of the above distributions do not test coherence specifically. Better signals are a slower increase in multiplicity as a function of energy, or in a characteristic depletion of particle production at low momenta [13]. The data clearly speak in favour of coherence. Recently, two new tests have been performed. One is to consider the rate of particle pairs as a function of the relative angle, and specifically the difference in rate between the pairs almost back-to-back and those almost collinear [16]. Another is to classify events as either three- or two-jet ones at some given resolution scale, and then compare the average number of additional sub-jets that are found when the resolution parameter is reduced [17]. Again incoherent models fail to describe the data, while the coherent ones do very well. The only catch is that all the tests above probe not only the perturbative but also the non-perturbative aspects of the models, so some caution is necessary in not overinterpreting the results.

Recently, the b-quark rate in hadronic Z^0 decays has been a topic of some controversy. The experimentally observed $b\overline{b}$ rate is about two standard deviations above the theoretically predicted one in the standard model [18]. This may not seem enough of a deviation to worry anybody but, given the excellent agreement with the standard model in almost every other respect, it is natural to closely scrutinize every hint of even the slightest crack in the shiny wall. One possibility put forward is that there is a larger rate of secondary $b\overline{b}$ production, i.e. branchings $g \to b\overline{b}$, than assumed in current shower programs. Since the absolutely overwhelming contribution comes from primary $Z^0 \to b\overline{b}$ production, the small excess observed translates into a requirement to enhance the $g \to b\overline{b}$ rate by at least a factor of 4. The issue has been studied in some detail [19], by comparing parton-shower programs with resummed matrix-element calculations. The conclusion is that parton showers seem accurate in this respect to about the 20% level. This does not include general uncertainties such as what is the proper b quark mass (not necessarily the same for the high-virtuality $Z^0 \to b\overline{b}$ vertices as for the close-to-threshold $g \to b\overline{b}$ ones), but it is still difficult to imagine a total error of a factor of 2, let alone one of 4. The discrepancy, if there is one, is therefore likely to be found elsewhere. One example might be in uncertainties in the shape of the b fragmentation function, specifically how often a B hadron is so slow that no secondary vertex is registered in the detector.

3.2 Fragmentation

The fragmentation process has yet to be understood from first principles, starting from the QCD Lagrangian. This has left the way clear for the development of a number of different phenomenological models. Being models, none of them can lay claims to being 'correct'. The best one can aim for is a good representation of existing data, plus a predictive power for properties not yet studied or results at higher energies.

All existing models are of a probabilistic and iterative nature. This means that the fragmentation process as a whole is described in terms of one (or a few) simple underlying branchings, of the type jet $\to$ hadron + remainder-jet, string $\to$ hadron + remainder-string, cluster $\to$ hadron + hadron, or cluster $\to$ cluster + cluster. At each branching, probabilistic rules are given for the production of new flavours, and for the sharing of energy and momentum between the products.

Three main schools are usually distinguished, string fragmentation (SF), cluster fragmentation (CF) and independent fragmentation (IF) [20]. These need not be mutually exclusive; it is possible to have models which contain both cluster and string aspects, or models which interpolate between independent and string fragmentation. Local parton–hadron duality (LPHD) is a fourth approach, wherein the perturbatively calculable rate of partons is assumed directly translatable into corresponding rates of hadrons; it is not a complete scheme in the sense of the others, but it is useful for some observables.

While the evolution of fragmentation models was rapid in the early eighties, no really new algorithms have been introduced in the last ten years, and only a modest amount of refinement of the existing approaches has been performed.

For lack of time, and because of personal bias, in the following I concentrate on the string fragmentation approach [21].

While non-perturbative QCD is not solved, lattice QCD studies lend support to a linear confinement picture (in the absence of dynamical quarks), i.e. the energy stored in the colour dipole field between a charge and anticharge increases linearly with the separation between the charges, if the short-distance Coulomb term is neglected. This is quite different from the behaviour in QED, and is related to the presence of a three-gluon vertex in QCD. The details are not yet well understood, however.

The assumption of linear confinement provides the starting point for the string model, most easily illustrated for the production of a back-to-back $\mathrm{q\overline{q}}$ jet pair. As the partons move apart, the physical picture is that of a colour flux tube (or maybe colour vortex line) being stretched between the q and the $\overline{\mathrm{q}}$. The transverse dimensions of the tube are of typical hadronic sizes, roughly 1 fm. If the tube is assumed to be uniform along its length, this automatically leads to a confinement picture with a linearly rising potential. In order to obtain a Lorentz covariant and causal description of the energy flow due to this linear confinement, the most straightforward way is to use the dynamics of the massless relativistic string with no transverse degrees of freedom. The mathematical, one-dimensional string can be thought of as parametrizing the position of the axis of a cylindrically symmetric flux tube. From hadron spectroscopy, the string constant, i.e. the amount of energy per unit length, is deduced to be $\kappa \approx 1$ GeV/fm.

As the q and $\overline{\mathrm{q}}$ move apart, the potential energy stored in the string increases, and the string may break by the production of a new $\mathrm{q'\overline{q}'}$ pair, so that the system splits into two colour singlet systems $q\overline{q}'$ and $q'\overline{q}$. If the invariant mass of either of these string pieces is large enough, further breaks may occur. In the Lund string model, the string break-up process is assumed to proceed until only on-mass-shell hadrons remain, each hadron corresponding to a small piece of string.

In order to generate the quark–antiquark pairs $q'\overline{q}'$, which lead to string break-ups, the Lund model invokes the idea of quantum mechanical tunnelling. This leads to a flavour-independent Gaussian spectrum for the transverse momentum of $q'\overline{q}'$ pairs. Tunnelling also implies a suppression of heavy quark production, $u : d : s : c \approx 1 : 1 : 0.3 : 10^{-11}$. Charm and heavier quarks hence are not expected to be produced in the soft fragmentation.

A tunnelling mechanism can also be used to explain the production of baryons. This is still a poorly understood area. In the simplest possible approach, a diquark in a colour antitriplet state is just treated like an ordinary antiquark, such that a string can break either by quark–antiquark or antidiquark–diquark pair production. The production probabilities are then given by the effective diquark masses assumed, plus simple flavour Clebsch-Gordan coefficients of the baryon wave functions. In this approach, the baryon and antibaryon are produced next to each other, and share (at least) two quark flavours. A more complex scenario is the 'popcorn' one, where diquarks as such do not exist, but rather quark–antiquark pairs are produced one after the other. Part of the time, this scenario gives back an effective diquark picture, but in addition configurations are possible where one or more mesons are produced in between the baryon and antibaryon, and where therefore these two are no longer required to be as strongly correlated in flavour content.

In general, the different string breaks are causally disconnected. This means that it is possible to describe the breaks in any convenient order, e.g. from the quark end inwards. Results, at least not too close to the string endpoints, should be the same if the process is described from the q end or from the $\overline{q}$ one. This 'left–right' symmetry constrains the allowed shape of fragmentation functions $f(z)$, where z is the fraction of $E + p_{\mathrm{L}}$ that the next particle will take out of whatever remains. Here p_{L} is the longitudinal momentum along the direction of the respective endpoint, opposite for the q and the $\overline{q}$. Two free parameters remain, which have to be determined from data.

If several partons are moving apart from a common origin, the details of the string drawing become more complicated. For a $q\overline{q}g$ event, a string is stretched from the q end via the g to the $\overline{q}$ end, i.e. the gluon is a kink on the string, carrying energy and momentum. As a consequence, the gluon has two string pieces attached, and the ratio of gluon/quark string forces is 2, a number that can be compared with the ratio of colour charge Casimir operators, $N_C/C_F = 2/(1 - 1/N_C^2) = 9/4$. In this, as in other respects, the string model can be viewed as a variant of QCD, where the number of colours N_C is not 3 but infinite. Fragmentation along this kinked string proceeds along the same lines, as sketched for a single straight string piece. Therefore no new fragmentation parameters have to be introduced.

The more prominent features of LEP data are well described by the string model, when combined with the parton-shower approach mentioned before. One recent example is detailed comparisons of quark and gluon jet fragmentation, which have been made possible by the high statistics and good flavour tagging capabilities of LEP experiments [22]. It is now clearly established that gluon jets have a softer particle momentum spectrum and a broader angular distribution than quark jets of the same energy, and that the results are in excellent agreement with the model predictions. However, as a consequence of the increased statistics, also discrepancies start to show up. This is most notable in the flavour composition, i.e. in the rate of various mesons and baryons: even with a rather large number of free parameters available, the current string fragmentation approach is somewhat off in many places and fails miserably in some [23]. So the conclusion seems to be that the general space–time structure of fragmentation is under control, but that the detailed mechanism of flavour production and hadron formation is still not so well understood.

3.3 W^+W^- Events

Based on the above sections, the situation might seem rather satisfactory: if both perturbative and non-perturbative QCD aspects appear reasonably well under control, then, from now on, effectively we can forget about QCD whenever we go about the business of testing the standard model in the weak sector. However, this is not quite true, and as an example we can consider determinations of the W mass. The m_W will be *the* critical observable of LEP 2, which should allow new precision tests.

The problem is that QCD interference effects between the W^+ and W^- decays undermine the traditional meaning of a W mass in the process $e^+e^- \to W^+W^- \to q_1\overline{q}_2\, q_3\overline{q}_4$. Specifically, it is not even in principle possible to subdivide the hadronic final state into two groups of particles, one of which is produced by the $q_1\overline{q}_2$ system of the W^+ decay and the other by the $q_3\overline{q}_4$ system of the W^- decay: some particles originate from the joint action of the two systems.

In order to understand which QCD interference effects can occur in hadronic W^+W^- decays, it is useful to examine the space–time picture of the process. Consider a typical c.m. energy of 170 GeV, a W mass $m_W = 80$ GeV, and a width $\Gamma_W = 2.08$ GeV. The averaged (over the W-mass distribution) proper lifetime for a W is $\langle\tau\rangle \approx (2/3)\hbar/\Gamma_W \approx 0.06$ fm. This gives a mean separation of the two decay vertices of 0.04 fm in space and 0.07 fm in time. A gluon with an energy $\omega \gg \Gamma_W$ therefore has a wavelength much smaller than the separation between the W^+ and W^- decay vertices, and is emitted almost incoherently either by the $q_1\overline{q}_2$ system or by the $q_3\overline{q}_4$ one. Only fairly soft gluons, $\omega \lesssim \Gamma_W$, feel the joint action of all four quark colour charges. On the other hand, the typical distance scale of hadronization is about 1 fm, i.e. much larger than the decay vertex separation. Therefore the hadronization phase may contain significant interference effects.

A complete description of QCD interference effects is not possible since non-perturbative QCD is not well understood. The concept of colour reconnection/rearrangement [24] is therefore useful to quantify effects (at least in a first approximation). In a reconnection two original colour singlets (such as $q_1\overline{q}_2$ and $q_3\overline{q}_4$) are transmuted into two new ones (such as $q_1\overline{q}_4$ and $q_3\overline{q}_2$). Subsequently each singlet system is assumed to hadronize independently according to the standard algorithms, as outlined above. Depending on whether a reconnection has occurred or not, the hadronic final state is then going to be somewhat different.

In the following, we will first discuss perturbative effects and subsequently non-perturbative ones. Further details may be found in [25].

Until today, perturbative QCD has mainly been applied to systems of primary partons produced almost simultaneously. The radiation accompanying such a system can be represented as a superposition of gauge-invariant terms, in which each external quark line is uniquely connected to an external antiquark line of the same colour. The system is thus decomposed into a set of colourless $\widehat{q\overline{q}}$ antennae/dipoles. Neglecting interferences, the $e^+e^- \to W^+W^- \to q_1\overline{q}_2\, q_3\overline{q}_4$ final state can be subdivided into two separate dipoles, $\widehat{q_1\overline{q}_2}$ and $\widehat{q_3\overline{q}_4}$. Each dipole may radiate gluons from a maximum scale m_W downwards. Within the perturbative approach, colour transmutations can result only from the interferences between gluons (virtual as well as real) radiated in the W^+ and W^- decays. A colour reconnection then corresponds to radiation, e.g. from the dipoles $\widehat{q_1\overline{q}_4}$ and $\widehat{q_3\overline{q}_2}$. The emission of a single primary gluon cannot give interference effects, by colour conservation, so interference terms only enter in second order in α_s.

The general structure of the results is well illustrated by the interference between the graph where a gluon with momentum k_1 (k_2) is emitted off the $\widehat{q_1\overline{q}_2}$ ($\widehat{q_3\overline{q}_4}$) dipole and the same graph with k_1 and k_2 interchanged:

$$\frac{1}{\sigma_0}\,\mathrm{d}\sigma^{\mathrm{int}} \simeq \frac{\mathrm{d}^3\,\mathbf{k}_1}{\omega_1}\,\frac{\mathrm{d}^3\,\mathbf{k}_2}{\omega_2}\,\left(\frac{C_F\,\alpha_{\mathrm{s}}}{4\pi^2}\right)^2\,\frac{1}{N_C^2-1}\,\chi_{12}\,H(k_1)\,H(k_2)\;. \tag{4}$$

Note that the interference is suppressed by $1/(N_C^2-1) = 1/8$ as compared to the total rate of double primary gluon emissions. This is a result of the ratio of the corresponding colour traces.

The so-called profile function χ_{12} controls decay–decay interferences. It quantifies the overlap of the W propagators in the interfering Feynman diagrams. Near the W^+W^- pair threshold, χ_{12} simplifies to

$$\chi_{12} \approx \frac{\Gamma_{\mathrm{W}}^2}{\Gamma_{\mathrm{W}}^2 + (\omega_1-\omega_2)^2}\;. \tag{5}$$

Other interferences (real or virtual) are described by somewhat different expressions, but have the same general properties. The profile functions cut down the phase space available for gluon emissions with $\omega \gtrsim \Gamma_{\mathrm{W}}$ by the alternative quark pairs. The possibility for the reconnected systems to develop QCD cascades is thus reduced, i.e. the dipoles are almost sterile.

The radiation pattern $H(k)$ is given by

$$H(k) = \widehat{q_1\overline{q}_4} + \widehat{q_3\overline{q}_2} - \widehat{q_1q_3} - \widehat{\overline{q}_2\overline{q}_4}\;, \tag{6}$$

where the radiation antennae are

$$\widehat{ij} = \frac{(p_i\cdot p_j)}{(p_i\cdot k)(p_j\cdot k)}\;. \tag{7}$$

In addition to the two dipoles $\widehat{q_1\overline{q}_4}$ and $\widehat{q_3\overline{q}_2}$, which may be interpreted in terms of reconnected colour singlets, one finds two other terms, $\widehat{q_1q_3}$ and $\widehat{\overline{q}_2\overline{q}_4}$, which come in with a negative sign. The signs represent the attractive and repulsive forces between quarks and antiquarks. The effects of the reconnected almost sterile cascades should appear on top of a dominant background generated by the ordinary-looking no-reconnection dipoles $\widehat{q_1\overline{q}_2}$ and $\widehat{q_3\overline{q}_4}$. The negative-sign interference terms are therefore perfectly physical, and distort the overall radiation pattern in the same direction as the positive-sign ones.

We now turn to the possibility of reconnection occurring as a part of the non-perturbative hadronization phase. This requires model building, beyond what is already available in the standard string fragmentation approach. Specifically, the string model does not constrain the nature of the string fully. At one extreme, the string may be viewed as an elongated bag, i.e. as a flux tube without any pronounced internal structure. At the other extreme, the string contains a very thin core, a vortex line, which carries all the topological information, while the energy is distributed over a larger surrounding region. The latter alternative is the chromoelectric analogue to the magnetic flux lines in a type II superconductor, whereas the former one is more akin to the structure of a type

I superconductor. We use them as starting points for two contrasting approaches, with nomenclature inspired by the superconductor analogy.

In scenario I, the reconnection probability is proportional to the space–time volume over which the W^+ and W^- strings overlap, with saturation at unit probability. A consideration of distances in the W^+W^- system shows that each W can effectively be viewed as instantaneously decaying into a string spanned between the partons, from a quark end via a number of intermediate gluons to the antiquark end. The strings expand, both transversely and longitudinally, at a speed limited by that of light. They eventually fragment into hadrons and disappear. An overlap of the W^+ and W^- strings may be calculated by making an ansatz for each individual string field, uniform in the longitudinal direction and falling off as a Gaussian of approximately 0.5 fm width in the transverse direction, and an average proper time of decay of $\tau_{\mathrm{frag}} \approx 1.5$ fm. This gives a model with one free parameter, the constant of proportionality between the space–time integral of the overlap and the probability of a reconnection.

In scenario II it is assumed that reconnections can only take place when the core regions of two string pieces cross each other. This means that the transverse extent of strings can be neglected, which leads to considerable simplifications compared with the previous scenario. Such an approach does not introduce any new parameters. The reconnection probability comes out to be about 35% at 170 GeV; the free parameter of model I has been adjusted to give the same answer at this energy. This probability does not vary by more than a factor of 2 over the full LEP 2 energy range.

Comparing scenarios I and II above with the no-reconnection scenario, it turns out that reconnection effects are very small. The change in the average charged multiplicity is at the level of a per cent or less, and similar statements hold for rapidity distributions, thrust distributions, and so on. This is below the experimental precision we may expect, and so may well go unobserved. One would like to introduce more clever measures, which are especially sensitive to the interesting features, but so far we have had little success.

Ultimately, the hope would be to distinguish between scenarios I and II, and thereby to gain some insight into the nature of the confinement mechanism. In principle, there are such differences. For instance, the reconnection probability is much more sensitive to the event topology in scenario II, since the requirement of having two string cores cross is more selective than that of having two broad flux tubes overlap.

We now come to the W mass. Experimentally, m_{W} depends in a non-trivial fashion on all particle momenta of an event. Errors in the W mass determination come from a number of sources [26], which we do not intend to address here. Therefore we only study the extent to which the average reconstructed W mass is shifted when reconnection effects are added, but everything else is kept the same. Even so, results do depend on the reconstruction algorithm used. We have tried a few different ones, which however are all based on the same philosophy: a jet finder is used to define at least four jets, events with two very nearby jets or with more than four jets are rejected, the remaining jets are paired to define the two W's, and the average W mass of the event is calculated. Events where this number agrees to better than 10 GeV with the input average mass are used to calculate the systematic mass shift.

In scenario I this shift is consistent with being zero, within the 10 MeV uncertainty in our results from limited Monte Carlo statistics (160,000 events per scenario). Scenario II gives a negative mass shift, of about -30 MeV; this also holds for several variations of the basic scheme. A simpler model, where reconnections are always assumed to occur at the centre of the event, gives a positive mass shift instead: about $+30$ MeV if results are rescaled to a reconnection probability of 35%. We are therefore forced to conclude that not even the sign of the effect can be taken for granted, but that a real uncertainty of ± 30 MeV does exist from our ignorance of non-perturbative reconnection effects. Studies show that pure perturbative effects could add at most about 5 MeV to this, while the potential interplay between perturbative and non-perturbative effects (one W decaying inside the hadronic field of the other W) has been assumed no larger than that.

Since the three sources are not independent, the numbers are added linearly to get an estimated total uncertainty of 40 MeV. In view of the aimed-for precision, this is non-negligible, and should be a cause for worry. It is not impossible that one could tailor-make experimental algorithms that are less sensitive to these effects, however.

Potential reconnection effects may not be the only uncertainty. Currently we are studying the uncertainties that could come from Bose–Einstein effects [27]. The underlying reason is the same as for the W reconnections: the two W's decay so close to each other compared with typical hadronization distances and Bose–Einstein radii. For a pair of nearby π^0's, say, the production amplitude then should be symmetrized with respect to which π^0 comes from which W. Even neglecting reconnection phenomena. the concept of a W mass on the hadron level is then undermined. The Bose-Einstein phenomenon is very poorly known, so we cannot definitely claim that there have to be observable effects. However, attempts at modelling indicate that the uncertainty in the W mass can well turn out to be comparable with or even larger than the one quoted for reconnection phenomena.

4 Hadronic Physics

The need for event generators is excellently illustrated by the recent CDF top paper [28]. In order to reach a conclusion, generators are used at every step of the way: to study the signal, potential backgrounds, detector response and imperfections, and so on. This shows how critical the generator aspects are, for better or for worse.

While hadron colliders are well suited to reach energies higher than e^+e^- ones (cf. the Tevatron and LEP), there are disadvantages. Hadrons have a complicated internal structure of quarks and gluons. This means that hadronic collisions are more complex than leptonic ones. Leptoproduction ep events are intermediate to the e^+e^- and the $pp/\overline{p}p$ ones, with a simple probe on a complicated target. Therefore the experience from HERA will be invaluable in reaching a better understanding of the hadronic structure in its broadest sense, i.e. also including aspects such as initial-state QCD radiation, interference between initial- and final-state radiation, and beam-jet fragmentation. However, from an event-generator point of view, ep and e^+e^- processes may be viewed as special cases of the hadron–hadron description. I therefore now jump directly to the latter kind of processes.

A summary of the physics in hadronic event generators is the following:

- An event is normally classified by the 'hardest' (i.e. the one with largest momentum transfer) interaction that occurs. This can be a process such as $qg \to qg$, $q\overline{q}' \to W^+$, $q\overline{q} \to t\overline{t}$, or anything else. The corresponding matrix element is perturbatively calculable. Not all events need contain a hard, calculable subgraph, exceptions are found among elastic, diffractive and low-$p_\perp$ events.
- In order to calculate a cross section, the squared matrix element has to be multiplied by two parton-distribution functions, which describe the partonic content of the two incoming hadrons. The analogy with e^+e^- and ep physics is made more transparent if one introduces parton distributions also for leptons. The evolution is here not given by QCD processes but by QED branchings such as $e \to e\gamma$. The probability

that the electron retains a fraction x of the full momentum if it is probed at a scale Q^2 is fully perturbatively calculable, unlike the QCD case.

- The initial-state radiation that gave rise to the two incoming partons has to be reconstructed, i.e. the inclusive parton-distribution description has to be turned into an exclusive set of radiated partons.
- Also partons in the final state can radiate further, in the same spirit as described for e^+e^- events.
- Not all partons of an incoming hadron take part in the hard interaction. A remnant is left behind, 'attached' to the hard interaction by its colour charge. Nothing forbids several partons being kicked out, by independent (semi-)hard interactions. All this gives a 'beam jet' structure that still is not so well understood.
- Again outgoing coloured partons turn into colourless hadrons by fragmentation. Normally the fragmentation process is assumed universal, i.e. the same models can be used as in e^+e^-. This need not be correct — universality is known to break down if one tries to extrapolate from pp collisions to heavy-ion ones — but it is a reasonable starting point.
- Unstable particles decay, just as in e^+e^-.

Below we give some further details on the topics not already covered for e^+e^-.

4.1 Hard Processes and Parton Distributions

The wide range of physics processes that are of interest in hadronic physics leads to a need for generators to contain a bit of everything. For instance, PYTHIA contains the following major groups:

- Hard QCD processes, e.g. $qg \to qg$.
- Soft QCD processes, such as diffractive and elastic scattering, and minimum-bias events.
- Heavy-flavour production, e.g. $gg \to t\bar{t}$.
- Prompt-photon production, e.g. $qg \to q\gamma$.
- Photon-induced processes, e.g. $\gamma g \to q\bar{q}$.
- Deep inelastic scattering, e.g. $q\ell \to q\ell$.
- W/Z production, such as $e^+e^- \to \gamma^*/Z^0$ or $q\bar{q} \to W^+W^-$.
- Standard model Higgs production, where the Higgs is reasonably light and narrow, and can therefore still be considered as a resonance.
- Gauge boson scattering processes, such as $W_L W_L \to W_L W_L$ (L = longitudinal), when the standard model Higgs is so heavy and broad that resonant and non-resonant contributions have to be considered together.
- Non-standard Higgs particle production, within the framework of a two-Higgs-doublet scenario with three neutral (h^0, H^0 and A^0) and two charged ($H^\pm$) Higgs states.
- Production of new gauge bosons, such as a Z', W' and R (a horizontal boson, coupling between generations).
- Production of fourth-generation fermions.
- Leptoquark (L_Q) production.
- Technicolour, e.g. $gg \to \eta_{\mathrm{techni}}$.
- Compositeness, e.g. d^* and u^* production.

- Other deviations from standard model processes, e.g. due to contact interactions or a strongly interacting gauge boson sector.

The list is by no means a survey of all interesting physics. Most notable is the absence of supersymmetric particle production and decay, but many other examples could be found: axigluons, baryon number violating processes, and so on. Also, within the scenarios studied, not all contributing graphs have always been included, but only the more important and/or more interesting ones. In many cases, various approximations are involved in the matrix elements coded. Other generators contain also other processes, and sometimes in other approximations, so there is a lot of complementarity.

The cross-section for a process $ij \to k$ is given by

$$\sigma_{ij\to k} = \int \mathrm{d}x_1 \int \mathrm{d}x_2 \, f_i^1(x_1, Q^2) \, f_j^2(x_2, Q^2) \, \hat{\sigma}_{ij\to k}(\hat{s}) \; . \tag{8}$$

Here $\hat{\sigma}$ is the cross-section for the hard partonic process, as codified in the matrix elements for each specific process. For processes with many particles in the final state it would be replaced by an integral over the allowed final-state phase space.

The $f_i^a(x, Q^2)$ are the parton distribution functions, which describe the probability to find a parton i inside beam particle a, with parton i carrying a fraction x of the total a momentum, when the a is probed at some squared momentum scale Q^2 that characterizes the hard process. Since we do not understand QCD in the low-Q^2 region, a derivation from first principles of the parton distributions of hadrons does not yet exist. It is therefore necessary to rely on parametrizations, where experimental data are used in conjunction with the evolution equations for the Q^2 dependence, to pin down the parton distributions. The most complete selection of distributions is found in PDFLIB [29].

The input from HERA has provided further stimulus for studies in this field. What is the small-x behaviour? Do parton distributions saturate? What is the rôle of the pomeron and rapidity-gap events? A number of questions remain to be answered.

4.2 Initial- and Final-State Radiation

For parton showers, a separation of radiation into a hard scattering and initial- and final-state showers is arbitrary, but very convenient. There are also situations where it is appropriate: for instance, the process $\mathrm{e}^+\mathrm{e}^- \to \mathrm{Z}^0 \to \mathrm{q}\overline{\mathrm{q}}$ only contains final-state QCD radiation, while $\mathrm{q}\overline{\mathrm{q}} \to \mathrm{Z}^0 \to \mathrm{e}^+\mathrm{e}^-$ only contains initial-state QCD radiation. Similarly, the distinction of emission as coming either from the q or from the $\overline{\mathrm{q}}$ is arbitrary. In general, the assignment of radiation to a given mother parton is a good approximation for an emission close to the direction of motion of that parton, but not for the wide-angle emission in between two jets, where interference terms are expected to be important. For such configurations the matrix-element approach is better, if possible.

In both initial- and final-state showers, the structure is given in terms of branchings $a \to bc$, specifically $\mathrm{e} \to \mathrm{e}\gamma$, $\mathrm{q} \to \mathrm{qg}$, $\mathrm{q} \to \mathrm{q}\gamma$, $\mathrm{g} \to \mathrm{gg}$, and $\mathrm{g} \to \mathrm{q}\overline{\mathrm{q}}$. These processes are characterized by the splitting kernels and evolution equations given earlier.

Each parton is associated with some virtuality scale Q^2, which gives an approximate sense of time ordering to the cascade. In the initial-state shower, Q^2 values are space-like ($m^2 < 0$) and gradually increasing as the hard scattering is approached, while Q^2 is time-like ($m^2 > 0$) and decreasing in the final-state showers. Emission angles increase on the way in to the hard interaction and decrease again thereafter. Only the energy per parton is decreased at both stages, as more and more partons are created and share the original energy.

A closer look reveals further differences. In a final-state branching, the two daughters are on an equal footing, both being time-like (or on the mass shell). In the initial-state branching, the mother parton and one daughter parton are space-like, whereas the other daughter is time-like (or on the mass shell). It is the space-like daughter that goes on towards the hard interaction, while the other daughter may initiate a time-like cascade on a side branch, just as in final-state radiation. The initial-state cascade may be viewed as a virtual fluctuation, wherein an initial parton almost on the mass shell is resolved into a set of partons with the same net invariant mass. Such fluctuations are born and die continuously in the proton wave function. It is the hard interaction that provides the momentum transfer to turn the space-like virtualities of the two incoming partons into time-like virtualities of the outgoing partons. It thereby also allows all the side branches to be promoted from a status of virtual fluctuations into one of final-state partons. Space-like fluctuations in principle are allowed on the side branches, but then remain purely virtual and are not observable in the final state.

A sequential evolution of the shower in time is not very convenient for generator applications, since the momenta of the incoming partons are then not known beforehand, which makes a matching to a desired hard scattering very costly in terms of efficiency. A common solution is backwards evolution [30], wherein the evolution equations are rewritten to act in the opposite direction for inital-state showers, i.e. from a given daughter-parton, the mother that produced it (together with a sister) is reconstructed. The procedure can then be started at the hard scattering, with known kinematics, and traced back to the two shower initiators.

Shower evolution is cut off at some lower scale Q_0, typically around 1 GeV for QCD branchings. The same cut-off scale is also used to regularize the soft-gluon-emission divergences in the splitting kernels. From above, a maximum scale Q_{max} is introduced, where the showers are matched to the hard interaction itself. The relation between Q_{max} and the kinematics of the hard scattering is uncertain, and the choice made can strongly affect the amount of well-separated jets.

We already mentioned a few open questions in the description of parton distributions; clearly, also these are reflected in corresponding uncertainties in the structure of initial-state parton showers. On top of this, the coherence conditions that we encountered for final-state radiation have a much more complicated and poorly understood analogue for the initial state [31]. Although existing parton shower programs do rather well by comparison with experiments, it should therefore not be a surprise that the level of confidence is not as high as in $\mathrm{e}^+\mathrm{e}^-$ annihilation. However, progress is being made. For instance, CDF recently presented an interesting study that shows the importance of angular ordering in the initial-state radiation or, more specifically, how the angles in the shower are matched to the hard scattering angle [32].

4.3 Beam Remnants and Multiple Interactions

In a hadron–hadron collision, the initial-state-radiation algorithm reconstructs one shower initiator in each beam, by backwards evolution from the hard scattering. This initiator only takes some fraction of the total beam energy, leaving behind a beam remnant that takes the rest. For a proton beam, a u quark initiator would leave behind a ud diquark beam remnant, with an antitriplet colour charge. The remnant is therefore colour-connected to the hard interaction, and forms part of the same fragmenting system. Often the remnant is more complicated, e.g. a g initiator would leave behind a uud proton-remnant system in a colour octet state, which can conveniently be subdivided into a colour triplet quark and a colour antitriplet diquark, each of which are colour-connected to the hard interaction. The energy sharing between these two remnant objects, and their relative transverse momentum, introduces additional degrees of freedom.

So far we have assumed that each event only contains one hard interaction, i.e. that each incoming particle has only one parton that takes part in hard processes, and that all other constituents sail through unaffected. This is appropriate in e^+e^- or ep events, but not necessarily so in hadron–hadron collisions. Here each of the beam particles contains a multitude of partons, and so the probability for several interactions in one and the same event need not be negligible. The dominant mechanism is expected to be that disjoint pairs of partons, with one parton from each beam, undergo $2 \to 2$ scatterings.

The dominant $2 \to 2$ QCD cross sections are divergent for $p_\perp \to 0$, and drop rapidly for larger $p_\perp$. Probably the lowest-order perturbative cross-sections will be regularized at small $p_\perp$ by colour coherence effects: an exchanged gluon of small $p_\perp$ has a large transverse wave function and can therefore not resolve the individual colour charges of the two incoming hadrons; it will only couple to an average colour charge that vanishes in the limit $p_\perp \to 0$. Customarily, some effective $p_{\perp\mathrm{min}}$ scale is therefore introduced, below which the perturbative cross-section is either assumed completely vanishing or at least strongly damped. Phenomenologically, in some approaches, $p_{\perp\mathrm{min}}$ comes out to be a number of the order of 1.5–2.0 GeV [33].

In a typical 'minimum-bias' event one therefore expects to find one or a few scattering(s) at scales around or a bit above $p_{\perp\mathrm{min}}$, while a high-$p_\perp$ event also may have additional scatterings at the $p_{\perp\mathrm{min}}$ scale. The probability to have several high-$p_\perp$ scatterings in the same event is small, since the cross-section drops so rapidly with $p_\perp$.

The understanding of a multiple interaction is still very primitive, and even the experimental evidence that it exists at all is rather weak [34]. The approach taken to this problem therefore varies significantly from one generator to the next. This may not always be appreciated by the normal user, since any approach by necessity contains a number of free parameters, and these parameters have been tuned by comparisons with essentially the same experimental data. Only when a broad range of properties are studied, preferably at several different energies, can one hope to better understand this complex part of physics.

HERA may provide very interesting possibilities to test these issues: a deep inelastic scattering should only give one interaction, but in the $Q^2 \to 0$ limit the hadronic nature of the photon takes over (see below) and one comes back to a hadron–hadron physics scenario with the possibility of multiple interactions.

4.4 γp and $\gamma\gamma$ Events

There are many reasons for being interested in γp and $\gamma\gamma$ physics. The process ep $\to$ eγp $\to$ eX is a main one at HERA and $e^+e^- \to e^+e^-\gamma\gamma \to e^+e^-X$ will be a main one at LEP 2 and future linear e^+e^- colliders. Therefore, these events are always going to give a non-negligible background to whatever other physics one is interested in. However, more importantly, collisions between real photons provides the richest spectrum of (leading-order) processes that is available for any choice of incoming elementary particles. For instance, since the photon has a hadronic component, all of hadronic physics is contained as a subset of the possibilities. A correct description of the components of the total γp and $\gamma\gamma$ cross sections is therefore the ultimate challenge of 'minimum-bias' physics. (Leaving heavy-ion physics aside.) This also explains why γp and $\gamma\gamma$ events here appear under the heading of hadronic physics. In the following I will describe the approach developed in

[35].

To first approximation, the photon is a point-like particle. However, quantum mechanically, it may fluctuate into a (charged) fermion–antifermion pair. The fluctuations $\gamma \leftrightarrow \mathrm{q\overline{q}}$ are of special interest to us, since such fluctuations can interact strongly and therefore turn out to be responsible for the major part of the γp and $\gamma\gamma$ total cross sections, as we shall see. On the other hand, the fluctuations into a lepton pair are uninteresting, since such states do not undergo strong interactions to leading order, and therefore contribute negligibly to total hadronic cross sections. The leptonic fluctuations are perturbatively calculable, with an infrared cut-off provided by the lepton mass itself. Not so for quark pairs, where low-virtuality fluctuations enter a domain of non-perturbative QCD physics. It is therefore customary to split the spectrum of fluctuations into a low-virtuality and a high-virtuality part. The former part can be approximated by a sum over low-mass vector-meson states, usually (but not necessarily) restricted to the lowest-lying vector multiplet. Phenomenologically, this Vector Meson Dominance (VMD) ansatz turns out to be very successful in describing a host of data. The high-virtuality part, on the other hand, should be in a perturbatively calculable domain.

In total, the photon wave function can then be written as

$$|\gamma\rangle = c_{\mathrm{bare}}|\gamma_{\mathrm{bare}}\rangle + \sum_{V=\rho^0,\omega,\phi,\mathrm{J}/\psi,\Upsilon} c_V|V\rangle + \sum_{\mathrm{q=u,d,s,c,b}} c_{\mathrm{q}}|\mathrm{q\overline{q}}\rangle + \sum_{\ell=\mathrm{e},\mu,\tau} c_\ell|\ell^+\ell^-\rangle \; . \qquad (9)$$

In general, the coefficients c_i depend on the scale μ used to probe the photon. Introducing a cut-off parameter p_0 to separate the low- and high-virtuality parts of the $\mathrm{q\overline{q}}$ fluctuations, one obtains $c_{\mathrm{q}}^2 \approx (\alpha_{\mathrm{em}}/2\pi)2e_{\mathrm{q}}^2\ln(\mu^2/p_0^2)$. The VMD part corresponds to the range of $\mathrm{q\overline{q}}$ fluctuations below p_0 and is thus μ-independent (assuming $\mu > p_0$). The major contribution comes from the ρ^0, $c_\rho \approx 0.04$. Finally, c_{bare} is given by unitarity: $c_{\mathrm{bare}}^2 \equiv Z_3 = 1 - \sum c_V^2 - \sum c_{\mathrm{q}}^2 - \sum c_\ell^2$. In practice, c_{bare} is always close to unity. Usually the probing scale μ is taken to be the transverse momentum of a $2 \to 2$ parton-level process. Our fitted value $p_0 \approx 0.5$ GeV then sets the minimum transverse momentum of a perturbative branching $\gamma \to \mathrm{q\overline{q}}$.

The subdivision of the above photon wave function corresponds to the existence of three main event classes in γp physics:

1. The VMD processes, where the photon turns into a vector meson before the interaction, and therefore all processes allowed in hadronic physics may occur. This includes elastic and diffractive scattering as well as low-$p_\perp$ and high-$p_\perp$ non-diffractive events.
2. The direct processes, where a bare photon interacts with a parton from the proton.
3. The anomalous processes, where the photon perturbatively branches into a $\mathrm{q\overline{q}}$ pair, and one of these (or a daughter parton thereof) interacts with a parton from the proton.

All three processes are of $O(\alpha_{\mathrm{em}})$. However, in the direct contribution the photon structure function is of $O(1)$ and the hard-scattering matrix elements of $O(\alpha_{\mathrm{em}})$, while the opposite holds for the VMD and the anomalous processes. The VMD component contributes about 80% of the total γp cross section, but less in the jet cross section; at intermediate $p_\perp$ values the anomalous processes are contributing most and at large $p_\perp$ values the direct processes dominate.

The difference between the three classes is easily seen in terms of the beam jet structure. The incoming proton always gives a beam jet containing the partons of the proton that did not interact. On the photon side, the direct processes do not give a beam jet at all, since all the energy of the photon is involved in the hard interaction. The VMD ones give a beam remnant just like the proton, with a 'primordial $k_\perp$' smearing of typically up to half a GeV. The anomalous processes give a beam remnant produced by the $\gamma \to q\overline{q}$ branching, with a transverse momentum going from p_0 upwards. Thus the transition from VMD to anomalous should be rather smooth.

A generalization of the above picture to $\gamma\gamma$ events is obtained by noting that each of the two incoming photons is described by a wave function of the type given in eq. (9). In total, there are therefore three times three event classes. By symmetry, the 'off-diagonal' combinations appear pairwise, so the number of distinct classes is 'only' six: VMD$\times$VMD, VMD$\times$direct, VMD$\times$anomalous, direct$\times$direct, direct$\times$anomalous and anomalous$\times$anomalous. The pattern of their relative importance is the same as for the γp process: VMD$\times$VMD dominates the total cross section and direct$\times$direct the jet cross section at large $p_\perp$.

When pp (or $\overline{\mathrm{p}}$p), γp and $\gamma\gamma$ events are compared at a common energy, the above ansatz leads to characteristic differences. There are most jets in $\gamma\gamma$ events and least in pp ones, not surprisingly, and this is also reflected in the total transverse energy flow, in the multiplicity distribution, and so on. Indications along these lines now start to appear at HERA, e.g. in the inclusive $p_\perp$ spectrum of charged particles [36]. The excess of jets in $\gamma\gamma$ events is observed at TRISTAN [37] and LEP 2 should have much more to say.

5 Summary and Outlook

This talk has only scratched the surface of the aspects that have to be programmed in a modern, versatile generator. If one looks back, the evolution has been explosive. The first version of the Lund Monte Carlo, in 1978, was about 200 lines long (and coded on punched cards!). Today, PYTHIA/JETSET together is over 30,000 lines of code, supplemented by a 300 pages long physics description and manual [7]. The growth in physics potential of the programs has been fairly linear over these years (a roughly constant number of persons contributing a rather constant number of new aspects per year), whereas the increase in the code itself has been closer to an exponential. This in part reflects changes in programming style, in part the trend to address more subtle and difficult-to-program problems as the simpler are 'solved'.

Of course, Lund is only one family of generators. Historically one should presumably start with models based on pure phase space or longitudinal phase space, but the first event generator in a more modern sense (that I am aware of) is the Artru–Mennessier model of 1974 [38]. The Field–Feynman ansatz of 1978 [39] had an enormous impact (partly due to the magic of the name Feynman). Starting with PETRA, the use of event generators has taken off, so that today there is hardly any experimental analysis presented or planned without the help of generators.

Initially there was a lot of scepticism, and it is not so easy to say when that disappeared. In retrospect it is tempting to call the UA1 experiences of 1984 the watershed. The 'discoveries' of supersymmetry and of top [40] can be traced back at least in part to a poor understanding of the signal, of the backgrounds, and of the detector response, and the only way one has found to do better in cases like these is to have more elaborate event generation and detector simulation programs.

Today, the problem is rather the opposite: some people have too deep a faith in generators. For instance, in the LEP 1 workshop our main recommendation was that 'Due to the large uncertainties present in any realistic QCD Monte Carlo, physics studies must be based on the use of at least two complete and independent programs.' [3], but this rule is not always followed. It is therefore important to remember the size and complexity of current-day generators. Hopefully this talk has given you some insight into the different aspects and assumptions that enter, and the many question marks that still remain. Even if the 'pioneering days' may be past, there is every need for continuing studies, e.g. in the area of multiparticle production, to address the new problems that come along. This way, hopefully, event generators will remain in the heartland of phenomenological and experimental particle physics.

Acknowledgements

A warm thank to many friends and collaborators who have helped shape the field I have here tried to describe, and an apology to all of those who have not been explicitly mentioned in the references. Special thanks go to Valery Khoze and Gerhard Schuler, with whom I collaborate on the two physics topics described in somewhat more detail here. Finally, my sincerest gratitude to the organizers of the meeting, for having brought me to this marvellous part of the world.

References

[1] U. Amaldi, W. de Boer and H. Fürstenau, Phys. Lett. **B260** (1991) 447

[2] R. Brun et al., GEANT 3, CERN Report DD/EE/84-1 (1987, revised)

[3] Z Physics at LEP 1 — Vol. 3: Event Generators and Software, eds. G. Altarelli, R. Kleiss and C. Verzegnassi, CERN 89-08 (1989)

[4] Physics at HERA — Vol. 3: Monte Carlo Generators, eds. W. Buchmüller and G. Ingelman (DESY, 1992)

[5] T. Sjöstrand, Z. Phys. **C42** (1989) 301;
F. Anselmo et al., in 'Large Hadron Collider Workshop', eds. G. Jarlskog and D. Rein, CERN 90-10 (1990), Vol. II, p. 130;
I.G. Knowles and S.D. Protopopescu, in 'Workshop on Physics at Current Accelerators and Supercolliders', eds. J.L. Hewett, A.R. White and D. Zeppenfeld, ANL-HEP-CP-93-92 (1993)

[6] G. Marchesini, B.R. Webber, G. Abbiendi, I.G. Knowles, M.H. Seymour and L. Stanco, Comput. Phys. Commun. **67** (1992) 465

[7] T. Sjöstrand and M. Bengtsson, Comput. Phys. Commun. **43** (1987) 367;
H.-U. Bengtsson and T. Sjöstrand, Comput. Phys. Commun. **46** (1987) 43;
T. Sjöstrand, Comput. Phys. Commun. **82** (1994) 74, CERN-TH.7112/93

[8] F.E. Paige and S.D. Protopopescu, in 'Physics of the Superconducting Super Collider 1986', eds. R. Donaldson and J. Marx (1987), p. 320

[9] Particle Data Group, L. Montanet et al., Phys. Rev. **D50** (1994) 1173

[10] G. Altarelli and G. Parisi, Nucl. Phys. **B126** (1977) 298;
Yu.L. Dokshitzer, Sov. Phys. JETP **46** (1977) 641

[11] A.H. Mueller, Phys. Lett. **B104** (1981) 161;
B.I. Ermolaev and V.S. Fadin, JETP Lett. **33** (1981) 269

[12] Yu.L. Dokshitzer, V.A. Khoze and S.I. Troyan, Sov. J. Nucl. Phys. (Yad. Fiz.) **47** (1988) 1384

[13] Yu.L. Dokshitzer, V.A. Khoze, A.H. Mueller and S.I. Troyan, Basics of Perturbative QCD (Editions Frontières, Gif-sur-Yvette, 1991)

[14] T. Hebbeker, Phys. Rep. **217** (1992) 69;
S. Bethke and J.E. Pilcher, Annu. Rev. Nucl. Part. Sci. **42** (1992) 251

[15] L3 Collaboration, B. Adeva et al., Phys. Lett. **B248** (1990) 227;
OPAL Collaboration, M.Z. Akrawy et al., Z. Phys. **C49** (1991) 49;
VENUS Collaboration, K. Abe et al., Phys. Rev. Lett. **66** (1991) 280;
DELPHI Collaboration, P. Abreu et al., Phys. Lett. **B255** (1991) 466;
ALEPH Collaboration, D. Decamp et al., Phys. Lett. **B284** (1992) 151

[16] A.A. Syed, Particle Correlations in Hadronic Decays of the Z^0 Boson (Ph.D. Thesis, Nijmegen, 1994)

[17] S. Catani, B.R. Webber, Yu.L. Dokshitzer and F. Fiorani, Nucl. Phys. **B383** (1992) 419;
OPAL Collaboration, R. Akers et al., Z. Phys. **C63** (1994) 363;
ALEPH Collaboration, D. Busculic et al., preprint CERN-PPE/94-208

[18] D. Schaile, CERN-PPE/94-162, to appear in the proceedings of the XXVII International Conference on High Energy Physics, Glasgow, July 1994

[19] M.H. Seymour, Lund preprint LU TP 94-7 (1994), to appear in Nucl. Phys. **B**

[20] T. Sjöstrand, Int. J. Mod. Phys. **A3** (1988) 751

[21] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, Phys. Rep. **97** (1983) 31

[22] OPAL Collaboration, P.D. Acton et al., Z. Phys. **C58** (1993) 387, and in preparation

[23] R.J. Hemingway, OPAL-CR173, to appear in the proceedings of the XXVII International Conference on High Energy Physics, Glasgow, July 1994

[24] G. Gustafson, U. Pettersson and P. Zerwas, Phys. Lett. **B209** (1988) 90

[25] T. Sjöstrand and V.A. Khoze, Phys. Rev. Lett. **72** (1994) 28, Z. Phys. **C62** (1994) 281

[26] LEP 2 workshop presentation by L. Camilleri at the LEPC open meeting, CERN, November 1992

[27] L. Lönnblad and T. Sjöstrand, in preparation

[28] CDF Collaboration, F. Abe et al., Phys. Rev. **D50** (1994) 2966

[29] H. Plothow-Besch, Comput. Phys. Commun. **75** (1993) 396

[30] T. Sjöstrand, Phys. Lett. **157B** (1985) 321;
M. Bengtsson, T. Sjöstrand and M. van Zijl, Z. Phys. **C32** (1986) 67

[31] M. Ciafaloni, Nucl. Phys. **B296** (1987) 249;
S. Catani, F. Fiorani and G. Marchesini, Nucl. Phys. **B336** (1990) 18;
G. Marchesini and B.R. Webber, Nucl. Phys. **B349** (1991) 617

[32] CDF Collaboration, F. Abe et al., Phys. Rev. **D50** (1994) 5562

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

[34] AFS Collaboration, T. Åkesson et al., Z. Phys. **C34** (1987) 163;
UA2 Collaboration, J. Alitti et al., Phys. Lett. **B268** (1991) 145;
CDF Collaboration, L.J. Keeble et al., in 'The Fermilab Meeting DPF 92', eds. C.H. Albright, P.H. Kasper, R. Raja and J. Yoh (World Scientific, Singapore, 1993), Vol. 2, p. 1002

[35] G.A. Schuler and T. Sjöstrand, Phys. Lett. **B300** (1993) 169, Nucl. Phys. **B407** (1993) 539, CERN-TH.7193/94 (1994)

[36] H1 Collaboration, I. Abt et al., Phys. Lett. **B328** (1994) 176

[37] TOPAZ Collaboration, H. Hayashii et al., Phys. Lett. **B314** (1993) 149;
AMY Collaboration, B.J. Kim et al., Phys. Lett. **B325** (1994) 248

[38] X. Artru and G. Mennessier, Nucl. Phys. **B70** (1974) 93

[39] R.D. Field and R.P. Feynman, Nucl. Phys. **B136** (1978) 1

[40] UA1 Collaboration, G. Arnison et al., Phys. Lett. **139B** (1984) 115 and **147B** (1984) 493