LU TP 94-11
June 1994

HERWIG for $\gamma\gamma$ and γp Physics*

Michael H. Seymour
Department of Theoretical Physics, University of Lund,
Sölvegatan 14A, S-22362 Lund, Sweden

Overview

HERWIG is a general purpose event generator for simulating **H**adron **E**mission **R**eactions **W**ith **I**nterfering **G**luons, which includes any combination of lepton, hadron or photon scattering in a single package[1]. The emphasis is firmly on the perturbative description of jet evolution, and in particular the treatment of coherence effects in that evolution.

The philosophy of HERWIG is to provide as complete as possible an implementation of perturbative QCD[2,3], coupled with as simple as possible a model of non-perturbative effects[4,5], uniformly for as many processes as possible. It gives a good description of high-energy data with relatively few adjustable parameters, most of which are well constrained by those data. This means that when moving to new types of collision, such as $\gamma\gamma$ and γp, at correspondingly lower energy, it is highly predictive, with only a couple of parameters that can be freely tuned. This, of course, is a great advantage for understanding the physical processes, the extent to which they can be described by perturbative QCD, and how similar they are in different reactions. But on the other hand, if the aim is simply to find a good parametrisation of the experimental final state, then the lack of free parameters can certainly be a nuisance.

To illustrate the process by which HERWIG generates events, we describe the reaction with the most steps, doubly-resolved photon scattering in $e^{\pm}e^{-}$ events. The various other combinations can generally be obtained from this picture by replacing one or more steps by simpler equivalents.

When HERWIG is presented with an ee collision and asked to produce a hadronic event, the first step is to emit photons from the electrons. This is done using the Equivalent Photon Approximation, retaining the full dependence on the photon transverse momentum and virtuality, within user-defined limits. This is important both as a tool for calculating experimental corrections and in reducing the theoretical ambiguities inherent in the EPA[6]. Since the photons in general have a non-negligible transverse momentum, the entire event is then boosted to their centre-of-mass frame, where they collide head-on.

*Talk given at the Workshop on Two-Photon Physics at LEP and HERA, Lund, Sweden, May 26–28, 1994.

In this frame the hard process, usually QCD $2 \to 2$ scattering, is treated exactly like a hadron-hadron collision: kinematics and parton flavours are generated according to the lowest-order matrix elements, within phase-space limits supplied by the user, most importantly the p_t range of the scattering. A limited set of parton distribution functions is built into HERWIG, but it also is straightforward to link to the PDFLIB library[7].

Next, the perturbative evolution of the incoming and outgoing partons into jets is simulated, by generating initial- and final-state parton showers. Additional partons produced in the initial-state shower go on to produce their own final-state jets, with transverse momenta constrained to be smaller than that of the hard process. The coherent nature of QCD radiation is fully included in both cases by appropriate choice of ordering variable[2], as is interference between initial- and final-state radiation by initial conditions that are matched to the colour dynamics of the hard scattering[3]. In both cases the evolution is downwards in scale, moving 'outwards' from the hard process towards the outgoing hadrons in the final-state case, and 'backwards' towards the incoming photons in the initial-state case. Evolution is cut off when the generated transverse momenta fall below a global cutoff or, for initial-state evolution, when the evolution scale reaches a cutoff related to the structure function set being used (of which more later).

Once the cutoffs have been reached, non-perturbative confinement effects set in, and we must switch from using QCD calculations to a hadronisation model. The cluster model used in HERWIG[4] relies strongly on the *preconfinement* property of QCD, that in the large-N_c limit partons tend to end up close in phase-space to partons with which they form colour singlets. Terms that violate this picture are suppressed by colour factors of at least $1/N_c^2$, and tend to be dynamically suppressed as well. This is implemented by first splitting all remaining gluons to quark-antiquark pairs, then joining colour-singlet quark-antiquark pairs (which come from different gluons of course) into colourless hadronic clusters. These clusters decay into hadrons according to uniform phase-space, with the choice of hadron species controlled purely by their mass suppression in the cluster decay phase-space, without any additional parameters.

Secondary particle decays are then generated according to simplified PDG tables, except for heavy (b and higher) quark decays, which are treated in the spectator model, as new sources of coherent radiation[8].

An optional final step provides a 'soft underlying event'. This is modeled as an independent collision between the colourless clusters that contain the photon remnants, assumed to be identical to that of two hadrons at the same collision energy, with distributions taken as simple parametrisations of data[9]. Despite its apparent simplicity, this model is capable of reproducing seemingly collective phenomena such as the 'pedestal effect'[5] in which the amount of hadronic activity in events with hard jets is observed to be larger than in minimum bias events, even well away from the jets. Recently, attempts have been made to improve on this simple model by interfacing a multiple scattering model to HERWIG[10].

Lepton-Photon Vertex

The phase-space for the EPA is controlled by the user-specified parameters, `Q2WWMN` and `Q2WWMX`, which control the lower and upper limits respectively of the generated photon virtuality. If set ≤ 0 or $\geq s$ respectively, the kinematic limits are used. The defaults are 0 and 4 GeV^2 respectively. The energy cutoff of the photon is controlled by the requirement that there be enough energy to satisfy the user's phase-space limits in the hard process.

Hard Processes

Although the full range of HERWIG's hard processes is also available in $\gamma\gamma$ and γp collisions, only a small fraction of them give large enough cross-sections to be useful. One of them must be chosen by setting the parameter `IPROC` to one of the values listed below. More details can be found in [1].

For twice-resolved photon processes in ee, or resolved photon in ep we have:

Process	IPROC	Phase-space controls	Defaults
QCD $2 \to 2$	1500	PTMIN $< p_t <$ PTMAX	10 GeV, $\sqrt{s}/2$
Prompt photon	1800	YJMIN $< y^* <$ YJMAX	Kinematic limits
Heavy quarks	1700+ID	ID=4,5,6 for c,b,t	

For once-resolved $\gamma\gamma$ processes in ee, or pointlike photon in ep:

Process	IPROC	Phase-space controls	Defaults
Point-like photon	5000	as IPROC = 1500 etc	
Prompt J/ψ	9107	$p_t >$ PTMIN, $z <$ ZJMAX	10 GeV, 0.9

For once-resolved $\gamma\gamma^*$ processes in ee:

Process	IPROC	Phase-space controls	Defaults
DIS	9000	Q2MIN $< Q^2 <$ Q2MAX	0 GeV^2, s
		YBMIN $< y_{\rm Bj} <$ YBMAX	0, 1

For point-like processes in ee:

Process	IPROC	Phase-space controls	Defaults
$\gamma^{(*)}\gamma^{(*)} \to q\bar{q}$	500	PTMIN $< p_{t,\gamma\gamma} <$ PTMAX	10 GeV, $\sqrt{s}/2$
		EMMIN $< m_{\gamma\gamma} <$ EMMAX	0 GeV, $\sqrt{s}$
		$\|\cos\theta^*\| <$ CTMAX	0.9999

Final-State Parton Shower

The process described by parton showers, multiparton production in hard processes, is fundamentally non-classical, since it necessarily involves interference between many different diagrams, including both real and virtual partons. Nevertheless, it can be formulated as a semi-classical 'time-ordered' process, at least to logarithmic accuracy. This is because in the presence of a fixed infrared cutoff, the logarithmic terms from both real and virtual emission exponentiate to give a factor that can be interpreted as the probability to get no additional radiation within the same phase-space region. This can then be interpreted in a causal way, as a series of real $1 \to 2$ branches, one after another.

However, this only correctly includes the interference between diagrams if one carefully accounts for the coherence of emission from all the partons in the event. This coherence means that soft gluon emission from a quark-gluon pair, for example, at angles greater than their opening angle is identical to that from an on-shell quark with their total momentum, suggesting that decreasing opening angle is a convenient variable to measure increasing 'time' along the shower. In time-ordered language, this is described as the quark emitting the larger angle gluon 'before' the smaller angle one, the famous angular-ordering condition[2].

Angular-ordered final-state parton showers are implemented in HERWIG by defining an evolution scale $Q^2 \equiv E^2(1-\cos\theta)$, where E is the energy of the emitter and θ the opening angle of the emission. Since energies decrease away from the hard process, ordering in Q^2 is equivalent to ordering in θ. Although this may not seem a Lorentz-invariant prescription since it involves energies and angles, the combination Q^2 is approximately boost-invariant, and becomes exactly so in the soft and collinear limits in which the parton shower approach is valid. The initial conditions of the shower are controlled by the coherence of the hard process, as we discuss below. The infrared cutoff is on transverse momentum, and the shower stops when the evolution variable of all branches is too small to produce emission above the cutoff. This cutoff is $m + Q_0$, where m is the mass of the parton, and Q_0 is separately tunable for quarks and gluons. This form ensures the correct treatment of coherence for heavy quarks (the 'dead cone' effect). In the program, Q_0 is called `VGCUT` for gluons and `VQCUT` for quarks. The default values, 100 MeV and 480 MeV respectively, are tuned to fit OPAL data[11]. The other important input is the strength of the running coupling, which is controlled by the QCD scale $\Lambda =$ `QCDLAM`, which defaults to 180 MeV.

Initial-State Parton Shower

In hard processes in which the incoming partons are coloured, emission from them should also be considered. However, it is not as simple as the final-state case, and it is not yet clear whether all the relevant effects can be incorporated into a simple probabilistic picture.

Emission that is soft relative to the parton going into the hard process is rather similar to the final-state case, except that one finds that coherence results in ordering of angles from the fixed incoming hadron direction, rather than of opening angles. On the other hand, if the emission takes away most of the energy, corresponding to small-x processes, the factor that exponentiates is not simply the probability for non-emission, and a probabilistic interpretation is not possible. However, the main effects can be incorporated by generating emission within a p_t-ordered phase-space[3].

In HERWIG, this is again implemented by defining an evolution variable $Q^2 = E^2(1 - \cos\theta)$, where E is the energy of the emitter (which is this time the 'furthest' from the hard process), and θ is the angle from the hadron direction. It can be seen that when all Es are similar, Q^2-ordering gives angular-ordering, and when they increase rapidly away from the hard process, it gives p_t-ordering, interpolating the two regions. The 'backward evolution' algorithm is used[12], which ensures that at each stage of branching, the parton distributions produced agree with the input parton distribution functions. The evolution is stopped by the same cutoffs on transverse momentum as for final-state emission. The emitted partons go on to produce final-state showers of their own, with angular-ordering requiring their initial Q^2 to be set to the Q^2 of their emission.

Towards the end of the shower, non-perturbative effects begin to be important, coming from the fact that the parton being evolved originated in a hadron. These are modeled by first forcing all partons to be evolved back to valence quarks, and then giving these a Gaussian intrinsic p_t, of width `PTRMS`, which defaults to 0. At present, incoming photons are treated exactly like hadrons, with no account of the inhomogeneous term, which has been discussed in detail at this workshop. In the parton shower, this term corresponds to the probability for a quark to be evolved backwards into the photon before the infrared cutoff is reached.

When evaluating the pdfs, the factorisation scale is set to Q^2, since emission above this has been treated, and only that below it should be factorised. However, most available structure function sets, including all those built into the program, are only valid for Q^2 above some scale which is typically 4 or 5 GeV^2. Thus an additional cutoff is applied, $Q >$ `QSPAC`, which defaults to 2.5 GeV and should always be at least as large as the smallest allowed scale in the pdf.

Rapidity Gap Events!

A surprising feature of HERWIG is the fact that it produces too many rapidity gap events, despite containing none of the relevant physics. This is in fact a simple consequence of the `QSPAC` cut just described: when the evolution variable Q reaches `QSPAC`, all evolution stops. But the angle θ could still be reasonably large, if the valence quark happens to be reasonably soft, meaning that the evolution could stop well before the corresponding angles have left the detector in the remnant direction. Although there is still a colour connection spanning the rapidity gap, it is quite possible that in some fraction of events it does not get filled during the hadronisation process. The only way to correct this problem while staying within the spirit of HERWIG is to use pdf sets that are reliable at scales down to around 1 GeV, and lower `QSPAC` accordingly. However, it is probable that future versions will contain a mechanism for freezing the pdfs below `QSPAC` as a systematic cross-check, even if it is not the default.

Coherence in the Hard Process

In addition to its main role in determining the kinematics of the hardest partons in the event, the hard process is also important in determining the initial conditions for the evolution of those partons. To leading order in N_c, any hard process can be represented as the sum of contributions in which each parton has a unique colour connected partner. The coherence of radiation from the different partons can be incorporated by requiring emission from each to be confined to angular cones whose extent is set by the angle to the colour-connected partner[3]. Corrections to this picture are again dynamically suppressed, as well as being non-leading in colour. Only if this coherence is correctly treated is the interference between initial- and final-state radiation incorporated, so the separation of the two justified. Its importance in reproducing high energy collider data has recently been shown by CDF in studies of three-jet events in comparison with versions of HERWIG and PYTHIA[13] with and without this effect[14].

Cluster Model

The cluster hadronisation model, as it is described in the overview, has only one adjustable parameter, the effective gluon mass, which controls the energy release in the forced gluon splitting to quarks. In the program this is called `RMASS(13)`, and defaults to the OPAL value, 750 MeV. All other features of the cluster decays to hadrons are controlled by the perturbative distribution of cluster masses and momenta. The clusters decay uniformly to two hadrons in their cmf according to pure phase-space, except for those that contain quarks from the perturbative evolution (ie. not those from gluon splitting) in which the decay is more 'string-like', with the hadron that contains the quark following the quark's direction. Effects such as strangeness suppression arise from the smaller phase-space available, since strange mesons are heavier than non-strange mesons, rather than being an adjustable parameter.

However, the clusters left at the end of perturbative evolution are sometimes too massive for uniform two-body decay to be a reasonable assumption. All clusters above an adjustable cutoff, `CLMAX`, undergo 'string-like' splitting to two lighter clusters. The parameter `PSPLT` controls how quickly the mass spectrum is reduced, resulting in a few fast-moving clusters (small `PSPLT`) or more slower clusters (large `PSPLT`). The defaults are 3.35 GeV and 1 respectively. With these values, a significant fraction of final state hadrons is produced by the cluster splitting mechanism.

Soft Underlying Event

As already described, if required, the soft underlying event (sue) is modeled as a soft hadron-hadron collision between the pair of colourless clusters that contain the hadron remnants. Whether or not this takes place is controlled by a prefix to the process code, `IPROC`: the numbers given earlier switch on the sue, while adding 10000 switches it off. If it is switched on, a fraction `PRSOF` of events have a sue, which defaults to 1. In events without one, the remnant clusters decay like any other cluster.

In event types such as DIS or point-like photon processes, in which there is only one hadron remnant, the break-up of the remnant is modeled as a sue between the remnant and its nearest neighbour, contributing to the E_t flow in the remnant direction.

Conclusion

As with most talks on Monte Carlo event generators, my main conclusion is the advice to get to know your generator: different models stress different aspects, and none provide a thorough description of every aspect of hadroproduction reactions. In HERWIG, the emphasis is strongly on an accurate description of the perturbative phase of parton evolution, with the non-perturbative phase made as simple as possible, with as few adjustable parameters as feasible. It gives a good fit to high energy data, with most of the parameters being reasonably well constrained. This means the model will be well tested by comparisons with data from $\gamma\gamma$ and γp collisions, with little freedom to tune away any discrepancies.

References

1. G. Marchesini, B.R. Webber, G. Abbiendi, I.G. Knowles, M.H. Seymour and L. Stanco, Computer Phys. Comm. 67 (1992) 465.
2. G. Marchesini and B.R. Webber, Nucl. Phys. B238 (1984) 1.
3. G. Marchesini and B.R. Webber, Nucl. Phys. B310 (1988) 461.
4. B.R. Webber, Nucl. Phys. B238 (1984) 492.
5. G. Marchesini and B.R. Webber, Phys. Rev. D 38 (1988) 3419.
6. P. Kessler, these proceedings.
7. H. Plothow-Besch, Computer Phys. Comm. 75 (1993) 396.
8. G. Marchesini and B.R. Webber, Nucl. Phys. B330 (1990) 261.
9. UA5 Collaboration, G.J. Alner et al., Nucl. Phys. B291 (1987) 445.
10. J.M. Butterworth and J.R. Forshaw, J. Phys. G19 (1993) 1657.
    J.M. Butterworth, these proceedings.
11. OPAL Collaboration, M.Z. Akrawy et al., Zeit. Phys. C49 (1991) 375.
12. T. Sjöstrand, Phys. Lett. 157B (1985) 321.
13. H.-U. Bengtsson and T. Sjöstrand, Computer Phys. Comm. 46 (1987) 43;
    T. Sjöstrand, these proceedings.
14. CDF Collaboration, F. Abe et al., FERMILAB Pub-94-072-E.